




Observing and Modelling the Atmospheres of Giant Exoplanets

by

Alastair Bryan Claringbold

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

May 2026

**UNIVERSITY
OF WARWICK**

Contents

List of Tables	v
List of Figures	vi
Acknowledgments	xi
Declarations	xiii
Abstract	xiv
Abbreviations and quantities	xv
Chapter 1 Introduction	1
1.1 The Exoplanet Population	2
1.1.1 Transiting planets	4
1.1.2 Radial velocity	8
1.1.3 Gas giant demographics	11
1.1.4 Formation and migration of giant planets	15
1.2 Planetary Atmospheres	18
1.2.1 Radiative transfer and opacities	18
1.2.2 Structure, transport and dynamics	24
1.2.3 Atmospheric chemistry	30
1.2.4 Aerosols and condensation	37
1.3 Exoplanet Atmosphere Characterization	39
1.3.1 Low resolution transmission spectroscopy	40
1.3.2 Low resolution eclipse spectroscopy	48
1.3.3 High resolution spectroscopy	51
1.3.4 Phase curves	53
1.3.5 Direct imaging	53
1.4 Overview and Thesis Structure	54
Chapter 2 Methods	55
2.1 Ground-based Observations	55
2.2 JWST Observations	56
2.2.1 Instrument setup	56

2.3	Data Reduction and Analysis	58
2.3.1	Raw image processing	59
2.3.2	Spectral extraction	60
2.3.3	Light-curve fitting	62
2.4	Atmospheric Modelling and Retrievals	65
2.4.1	Radiative transfer	65
2.4.2	Chemistry modelling	72
2.4.3	Atmospheric retrievals	74
Chapter 3 Transmission Spectroscopy of HAT-P-44b		78
3.1	LRG-BEASTS	78
3.2	The HAT-P-44 System	80
3.3	Observations	81
3.4	Data Reduction	85
3.5	Data Analysis	87
3.5.1	Transit model	87
3.5.2	Light-curve fitting	87
3.5.3	Transmission spectrum	92
3.5.4	Testing the sodium detection	96
3.6	Atmospheric Retrieval	98
3.6.1	Retrieval setup	98
3.6.2	Retrieval results	98
3.7	Discussion	101
3.7.1	The detection of sodium	101
3.7.2	Aerosols and evidence for methane and water	103
3.7.3	HAT-P-44b in context	105
3.7.4	Tests of background subtraction	107
3.8	Conclusions	110
Chapter 4 BOWIE-ALIGN: HAT-P-30b		112
4.1	BOWIE-ALIGN	112
4.1.1	Predictions	114
4.1.2	Sample	115
4.2	The HAT-P-30 system	117
4.3	Observations	118
4.4	Data reduction	118
4.4.1	Eureka!	119
4.4.2	Tiberius	120
4.4.3	The transmission spectrum of HAT-P-30b	122
4.5	Atmospheric modelling	125
4.5.1	petitRADTRANS forward models	125
4.5.2	petitRADTRANS retrievals	127
4.5.3	BeAR retrievals	133

4.5.4	Overview of retrieval results	137
4.6	Interior structure models	140
4.7	Discussion	141
4.7.1	The atmosphere of HAT-P-30 b	141
4.7.2	Implications for formation	142
4.8	Conclusions	144
Chapter 5 BOWIE-ALIGN: Complete Sample		146
5.1	BOWIE-ALIGN observations	146
5.1.1	The WASP-15 system	147
5.1.2	The TrES-4 system	147
5.1.3	The NGTS-2 system	148
5.2	Modelling Setup	148
5.2.1	Retrieval setups	148
5.3	WASP-15b	149
5.3.1	Forward models	149
5.3.2	Retrievals	151
5.3.3	Discussion	155
5.4	TrES-4b	159
5.4.1	Retrievals	159
5.4.2	Discussion	163
5.5	NGTS-2b	165
5.5.1	Retrievals	165
5.5.2	Discussion	167
5.6	Remaining Sample	171
5.7	Conclusions	174
Chapter 6 Atmospheric Modelling of Temperate Jupiters		176
6.1	Temperate Jupiters	176
6.2	Methods	178
6.2.1	Radiative-convective equilibrium	178
6.2.2	Chemical kinetics, photochemistry and condensation	179
6.3	Results	180
6.3.1	Energy transport and equilibrium processes	180
6.3.2	Disequilibrium chemistry	182
6.3.3	Condensation and aerosols	186
6.3.4	Transmission spectra	189
6.4	Mapping the Parameter Space	189
6.4.1	Further work	191
Chapter 7 Conclusions and Future Work		192
7.1	Summary	192
7.2	Ongoing and Future Work	194

7.2.1	The future of BOWIE-ALIGN	194
7.2.2	The temperate Jupiter TOI-1899b	195
7.2.3	Warm Jupiters with inner companions	195
7.3	Longer-term Outlook	197
7.3.1	Long period planet discovery	197
7.3.2	Ground-based atmospheric characterization	197
7.3.3	Space-based atmospheric characterization	198
7.3.4	Conclusions	199

List of Tables

3.1	Planetary parameters of HAT-P-44b obtained from white-light light-curve fits	88
3.2	Statistical comparison of different spectroscopic light-curve fitting models	95
3.3	Tabulated transmission spectrum of HAT-P-44b from both nights of observation	95
3.4	Statistical comparison between the retrieval models and simple parametric to the WHT/ACAM transmission spectrum of HAT-P-44b . .	102
3.5	Comparison of planetary and retrieved atmospheric properties of various hot Jupiters with confirmed sodium detections	106
4.1	System parameters of HAT-P-30 b used in this analysis	118
4.2	System parameters of HAT-P-30b from the JWST NIRSpec/G395H white-light light-curve fits	122
4.3	Statistical comparison between the retrieval models fit to the $R = 400$ Tiberius and Eureka! transmission spectra of HAT-P-30b	139
5.1	Comparison of the Bayesian evidence from different hybrid chemistry retrieval models to the JWST transmission spectrum of WASP-15b .	155

List of Figures

1.1	Histogram of cumulative exoplanet detections from the NASA Exoplanet Archive	3
1.2	Scatter plot of the exoplanet population in mass-period space from the NASA Exoplanet Archive	3
1.3	Transit light-curve of HD 209458b from Charbonneau et al. (2000) .	5
1.4	Schematic of a transit light-curve	7
1.5	Phase-folded radial velocity curve of the hot Jupiter 51 Pegasi b from Mayor & Queloz (1995)	10
1.6	Schematic of the Rossiter-McLaughlin effect at different obliquities from Gaudi & Winn (2007)	11
1.7	Histogram of estimate occurrence rates of giant planets at different orbital separations from (Fulton et al., 2021)	13
1.8	Plot of the observed and fitted mass-metallicity relationship of giant planets, from Chachan et al. (2025)	14
1.9	Protoplanetary disc C/O ratio as a function of orbital distance, from Öberg et al. (2011)	17
1.10	Absorption cross-section spectrum of CH ₄ based on the line-list of Yurchenko et al. (2024)	22
1.11	Pressure-temperature profiles of Solar System objects and hot Jupiter models from Madhusudhan & Seager (2009)	28
1.12	Vertical abundance ratios of CO in a brown dwarf, demonstrating the impact of vertical mixing	35
1.13	Condensation curves of various cloud species from Gao et al. (2021)	38
1.14	Schematic of atmospheric characterization techniques at different orbital phases from Gao et al. (2021)	40
1.15	JWST transmission spectra of hot Jupiters from Fu et al. (2025a) . .	44
1.16	Optical and infrared transmission spectra of hot Jupiters from Sing et al. (2016)	46
1.17	JWST emission spectrum of WASP-77 Ab from August et al. (2023)	50
1.18	Detection analysis of molecular features in the high resolution transmission spectrum of HD 189733b from Brogi et al. (2016)	52

2.1	Aperture extraction and background subtraction regions shown for ACAM observations of HAT-P-44b and JWST/NIRSpec observations of HAT-P-30b	61
2.2	Comparison of different limb-darkening laws performance from Espinoza & Jordán (2016)	63
2.3	Geometry of a grazing light ray in transmission spectroscopy from Mollière et al. (2019)	67
2.4	Opacity distribution of H ₂ O in frequency space and cumulative opacity distribution space from Grimm & Heng (2015a)	69
3.1	Transmission spectra of nine exoplanets published as part of the LRG-BEASTS survey	79
3.2	Example science frames from ACAM observations of HAT-P-44b . .	82
3.3	Ancillary data for the first night of observation of HAT-P-44b	83
3.4	Ancillary data for the second night of observation of HAT-P-44b . .	84
3.5	Plot of stellar spectra of target and comparison star, overplotted with the wavelength bins used to generate spectroscopic light-curves . . .	86
3.6	White-light light-curve fits of the first night of observation of HAT-P-44b	88
3.7	Spectroscopic light-curve fits of the first night of observation of HAT-P-44b	89
3.8	Red noise analysis of spectroscopic light-curve fit residuals from the first night of observation of HAT-P-44b	91
3.9	White-light light-curve fits of the second night of observation of HAT-P-44b	92
3.10	Spectroscopic light-curve fits of the second night of observation of HAT-P-44b	93
3.11	Transmission spectra of HAT-P-44b from each night of observation .	94
3.12	WHT/ACAM transmission spectrum of HAT-P-44b with simple parametric model fits	96
3.13	Transit depth of a binned light-curve centered on the sodium feature at 589 nm with incrementally increasing bin width	97
3.14	WHT/ACAM transmission spectrum of HAT-P-44b with the best-fitting retrieval model fits with and without clouds	99
3.15	Corner plot of posterior distributions of the retrieval analysis of the WHT/ACAM transmission spectrum of HAT-P-44b without clouds .	100
3.16	Corner plot of posterior distributions of the retrieval analysis of the WHT/ACAM transmission spectrum of HAT-P-44b including clouds	101
3.17	WHT/ACAM transmission spectrum of HAT-P-44b with the complete set of best-fitting retrieval and simple parametric models . . .	102
3.18	Comparison of the transmission spectra of HAT-P-44b and WASP-39b	105
3.19	Partial ACAM frames demonstrating the background subtraction performance from the first night of observation of HAT-P-44b	109

3.20	Full ACAM frames demonstrating the performance of an alternative background subtraction method for the first night of observation of HAT-P-44b	111
3.21	Comparison of transmission spectra obtained through different background subtraction methods	111
4.1	Schematic of the hot Jupiter origin scenarios as tested by the BOWIE-ALIGN programme from Kirk et al. (2024a)	113
4.2	Diagram of trends in atmospheric composition predicted by Penzlin et al. (2024) that will be tested by the BOWIE-ALIGN programme .	115
4.3	Parameter space plots of planets and their hosts of the BOWIE-ALIGN sample from Kirk et al. (2024a)	116
4.4	JWST NIRSpec/G395H white-light light-curves of HAT-P-30b for each detector from the Eureka! data reduction	120
4.5	Comparison of transmission spectra of HAT-P-30b obtained using Tiberius , the result of two analyses independently performed by me (ABC) and JK. Given the concordance between the results of these spectra, I favoured the homogeneous analysis produced by JK for further analysis, referred to simply as the Tiberius reduction for the rest of the chapter.	122
4.6	JWST NIRSpec/G395H transmission spectra of HAT-P-30b from each data reduction at $R = 100$ and $R = 400$	123
4.7	Differences between data reductions for the JWST transmission spectrum of HAT-P-30b	124
4.8	Equilibrium forward models of HAT-P-30b compared to the JWST transmission spectrum	126
4.9	Median retrieved spectrum from an equilibrium chemistry retrieval on the JWST transmission spectrum HAT-P-30b	128
4.10	Posterior corner plot from the equilibrium chemistry retrieval on the JWST transmission spectrum of HAT-P-30b	129
4.11	Posterior corner plot comparing the equilibrium chemistry retrievals on the JWST Eureka! spectrum of HAT-P-30b at $R = 400$ and $R = 100$	130
4.12	Spectral contributions to the median retrieved transmission spectrum of HAT-P-30b from the equilibrium chemistry retrieval	131
4.13	Posterior corner plot from the free chemistry retrieval of the JWST transmission spectrum of HAT-P-30b	132
4.14	Vertical mixing ratios of H_2O , CO_2 , CH_4 , CO , and H_2S from both the free and equilibrium chemistry retrievals of the JWST transmission spectrum of HAT-P-30b	133
4.15	Posterior corner plot from the hybrid chemistry retrieval of the JWST transmission spectrum of HAT-P-30b	134

4.16	Cornerplot showing the posteriors from the BeAR chemical equilibrium retrievals on the Tiberius (blue) and Eureka! (orange) reductions of HAT-P-30 b at $R = 400$. The top right insert shows the best-fit models for the Tiberius (blue) and Eureka! (orange) reductions. The Eureka! spectrum is offset by 1300 ppm for visualisation purposes. The legend in the bottom left indicates the reduced χ^2 values for each of the fits. Figure created by CEF.	136
4.17	Posterior corner plot of the favoured retrieval scenario to the $R = 400$ Eureka! reduction of the JWST transmission spectrum of HAT-P-30b, using equilibrium chemistry and no detector offset	138
5.1	Forward models of WASP-15b compared to the JWST transmission spectrum	150
5.2	Best fitting retrieval models to the JWST transmission spectrum of WASP-15b	152
5.3	Posterior corner plot from the free chemistry retrievals to the JWST transmission spectrum of WASP-15b	153
5.4	Posterior corner plot from the hybrid chemistry retrievals to the JWST transmission spectrum of WASP-15b	154
5.5	Opacity contributions to the best fitting hybrid and free chemistry retrievals models to the JWST transmission spectrum of WASP-15b	157
5.6	Posterior corner plot from free chemistry retrievals to the JWST transmission spectrum of TrES-4b	160
5.7	Posterior corner plot from equilibrium chemistry retrievals to the JWST transmission spectrum of TrES-4b	161
5.8	Posterior corner plots from hybrid chemistry retrievals to the JWST transmission spectrum of TrES-4b	162
5.9	Temperature-pressure profile posteriors from retrieval tests performed on the JWST transmission spectrum of TrES-4b using a Guillot temperature profile	164
5.10	Diagram of two potential formation scenarios that would explain the measured composition of TrES-4b, from Meech et al. (2025)	165
5.11	Posterior corner plot from the free chemistry retrievals of the JWST transmission spectrum of NGTS-2b, including SO as an opacity source	166
5.12	Overview of free chemistry retrievals of the JWST transmission spectrum of NGTS-2b using multiple models, reductions, and retrieval pipelines, from Fairman et al. (2026)	168
5.13	Posterior corner plots from equilibrium chemistry retrievals on the JWST transmission spectrum of NGTS-2b	169
5.14	Opacity contributions to the best fitting equilibrium chemistry retrieval of NGTS-2b	170
5.15	Comparison of posteriors from equilibrium retrievals of NGTS-2b	171

5.16	Transmission spectrum of KELT-7b with best fitting equilibrium retrieval model and spectral contributions from Ahrer et al. (2025b)	172
5.17	Transmission spectrum of WASP-94 Ab with best fitting equilibrium retrieval model and spectral contributions from Ahrer et al. (2025a)	173
6.1	Predicted pressure-temperature profiles of temperate Jupiter atmospheres with different instellations and stellar host spectral types	181
6.2	Vertical mixing ratios of key species in the atmospheres of temperate Jupiters based on photochemical modelling	183
6.3	Vertical mixing ratios of key atmospheric species in temperate Jupiters, highlighting key temperature and stellar host dependent transitions and the impact of vertical mixing strength	185
6.4	Column density of haze precursor species above 1 mbar as a function of planetary parameters including stellar host, temperature, vertical mixing strength, and metallicity	187
6.5	Synthetic infrared transmission spectra of two temperate Jupiters, showing the contribution spectra of different species	188
6.6	Schematic of the temperate Jupiter parameter space, highlighting key transitions that divide the parameter space into sextants	191
7.1	Diagram showing three gas giant formation scenarios I will test in JWST programme GO 7982 using a population of warm Jupiters with inner companions, adapted from Penzlin et al. (2024)	196

Acknowledgments

I would like to first and foremost thank my supervisor, Peter Wheatley, whose guidance, wisdom, and friendship have been invaluable, and express gratitude for the many hours he has dedicated to synthesizing new ideas, refining arguments and getting lost on tangents with me. I am also hugely grateful to Evie Ahrer and James Kirk, whose support and advice have been instrumental in getting me established in the field. I would also like to thank all the members of the BOWIE+ collaboration, for I could not have asked for a more supportive and positive environment in which to begin my academic career, with particular thanks to Hannah Wakeford, Mercedes López-Morales and James Owen for their mentorship. I would also like to dedicate my thanks to my master's supervisors Paul Rimmer and Sarah Rugheimer, and to Pete again, for giving me the opportunity to get involved in exoplanet atmosphere research, which I can truthfully say has been a near life-long goal. I also thank the supervisors and teachers that guided and inspired me through my academic journey, with a particular mention to Bob Dillon for his tutelage and friendship which guided me back towards physics when I was unsure of my direction.

I would also like to thank the many friends that have made the last four years of my life such a vibrant and formative time. This includes my friends at Warwick, especially Morgan, Ioannis, Jamie, James Blake, Marcelo and Anna, for the many fun times we have had, be it in the office, at the house, or on the road trip. It also includes the friends with whom I have spent many a late night playing games with – an essential component to me retaining my sanity I think – Josh, Leah, John and others. I of course also give thanks to the friends who have been by my side in Cambridge for the past eight years, particularly Seb, James, Emily, Arina and Nicky; I could not ask for more dedicated, compassionate company.

Lastly I thank my family, whose pride in me has been a constant source of motivation, and in particular my parents who have always supported and believed in me. Most of all I thank Laura, who has been my ardent partner in all things and

an endless font of encouragement and support; I simply could not imagine the last eight years without you.

“Trust in the Force. Do not return to the Temple... that time has past. And our future is uncertain. We will each be challenged. Our trust. Our faith. Our friendships. But we must persevere. And in time, a new hope will emerge.

May the Force be with you, always.”

OBI-WAN KENOBI, STAR WARS REBELS

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. I declare that this thesis is my own work except where explicitly stated otherwise, and has not been submitted in any previous application for any degree. The research was carried out during the period of October 2022 to April 2026 under the supervision of Prof. Peter Wheatley.

The content of Chapter 3 is mostly a reproduction of my first-author paper ‘LRG-BEASTS: detection of sodium and evidence for water absorption in the hot Saturn HAT-P-44b’ published in *Monthly Notices of the Royal Astronomical Society* (*MNRAS*), 547, stag384 (2026). The content of Chapter 4 is mostly a reproduction of my first-author paper ‘BOWIE-ALIGN: Sub-solar C/O ratio and metallicity atmosphere of the misaligned hot Jupiter HAT-P-30b’ published in *MNRAS*, 546, stag143 (2026).

The content of Chapter 5 is based on my research that was contributed to the following three papers on which I am a co-author: ‘BOWIE-ALIGN: JWST reveals hints of planetesimal accretion and complex sulphur chemistry in the atmosphere of the misaligned hot Jupiter WASP-15b’ published in *MNRAS*, 537, 3027-3052 (2025), on which I am the third of 26 authors; ‘BOWIE-ALIGN: Sub-stellar metallicity and carbon depletion in the aligned TrES-4b with JWST NIRSpec transmission spectroscopy’ published in *MNRAS*, 539, 1381-1403 (2025), on which I am the second of 22 authors; and ‘BOWIE-ALIGN: Modelling degeneracies in the NIRSpec/G395H transmission spectrum of the aligned hot Jupiter NGTS-2b’ accepted for publication in *MNRAS*, on which I am the third of 19 authors.

I provide a more detailed declaration as a preface to each of these results chapters describing individual contributions from collaborators, and work completed by collaborators is also stated explicitly within the text.

Abstract

Since the launch of JWST, we are in the midst of an ongoing revolution in our understanding of exoplanet atmospheres. This is particularly the case for giant planets, where we can characterize the physical and chemical processes present in their atmospheres in far greater detail than ever before, and compare across the population to gain insight into their formation. In this thesis, I present a combination of transmission spectroscopy observations and atmospheric models used to study a range of giant planets, from highly irradiated hot Jupiters, to a novel population of cooler gas giants, the temperate Jupiters.

I present the ground-based optical transmission spectrum of the hot Jupiter HAT-P-44b, detecting the presence of gaseous sodium, and evidence for Rayleigh scattering from hydrogen gas and hints of absorption from water vapour. This suggests HAT-P-44b possesses a relatively aerosol-free upper atmosphere.

I also present the JWST infrared transmission spectrum of the hot Jupiter HAT-P-30b, measuring its elemental abundances as part of the BOWIE-ALIGN programme to understand the connection between formation history and atmospheric composition. I also interpret the spectra of three additional hot Jupiters targeted by BOWIE-ALIGN: WASP-15b, TrES-4b, and NGTS-2b, revealing their diverse atmospheric compositions.

Finally, I include a modelling analysis simulating the atmospheric chemistry and aerosols present in temperate Jupiters. I predict strong transitions in sulfur condensation and nitrogen photochemistry, and trends in minor atmospheric species and haze precursors, that provide a cursory map of the parameter space to guide future observations.

Combined, these analyses contribute to our growing understanding of hot Jupiter aerosols and chemical composition, while laying the groundwork for the exploration of a whole new parameter space of exoplanets through temperate Jupiters. I will play a leading role in the latter exploration as the principal investigator of multiple JWST programmes targetting these planets.

Abbreviations and quantities

Abbreviations

ACAM Auxiliary CAMera

BIC Bayesian Information Content

BOTS Bright Object Time Series

BOWIE-ALIGN BOWIE – A Spectral Light Investigation into hot gas Giant
origiNs

CIA Collisionally-Induced Absorption

CRIRES CRyogenic InfraRed Echelle Spectrograph

EFOSC2 ESO Faint Object Spectrograph and Camera version 2

ELT Extremely Large Telescope

ESA European Space Agency

ESO European Southern Observatory

FWHM Full Width at Half Maximum

GCM Global Circulation Model

GP Gaussian Process

HATNet Hungarian-made Automated Telescope Network

HST Hubble Space Telescope

HWO Habitable Worlds Observatory

IGRINS Immersion GRating INfrared Spectrometer

JWST James Webb Space Telescope

LIFE Large Interferometer for Exoplanets

LRG-BEASTS Low-Resolution Ground-Based Exoplanet Atmosphere Survey using Transmission Spectroscopy

LTE Local Thermodynamic Equilibrium

MCMC Markov Chain Monte Carlo

MIRI Mid-InfraRed Instrument

MMW Mean Molecular Weight

NASA National Aeronautics and Space Administration

NIRCam Near-InfraRed Camera

NIRISS Near InfrarRed Imager and Slitless Spectrograph

NIRSpec Near InfraRed Spectrograph

NGTS Next Generation Transit Survey

NTT New Technology Telescope

PSF Point Spread Function

RMS Root Mean Square

RV Radial Velocity

SOSS Single Object Slitless Spectroscopy

TESS Transiting Exoplanet Survey Satellite

VLT Very Large Telescope

WASP Wide-Angle Search for Planets

WFC3 Wide Field Camera 3

WHT William Herschel Telescope

Quantities

R_J Jupiter radius: 71,492 km

M_J Jupiter mass: 1.898×10^{27} kg

R_☉ Solar radius: 695,000 km

M_☉ Solar mass: 1.988×10^{30} kg

au Astronomical unit: 1.496×10^8 km

pc Parsec: 3.086×10^{13} km

Chapter 1

Introduction

*“In the great commons at Gaia’s Landing we have
a tall and particularly beautiful stand of white
pine, planted at the time of the first colonies. It
represents our promise to the people, and to Planet
itself, never to repeat the tragedy of Earth.”*

LADY DEIRDRE SKYE, SID MEIER’S ALPHA CENTAURI

Life, at least as we know it, arose within the gravitational amalgamation of rock, metal, water, and gas that constitutes the planet Earth. Planets are therefore the most relatable of astronomical objects, and of fundamental importance to answering one of humanity’s most essential questions: *are we alone?* Locked in an orbital dance around their hosts, planets are complex environments of diverse physical and chemical processes forged in the dynamic setting of nascent stars.

Dating back to antiquity, humanity has known of six planets. Early introduced two more planets (Herschel, 1781; Galle, 1846), as well as some dwarf planets and many moons (e.g. Galilei, 1610; Huygens, 1656). For much of the 20th century, we were restricted to these Solar System objects to shape our understanding of planetary systems. Even the definition for a planet, as laid out by the International Astronomical Union, only recognizes these eight planets, defining planets as objects with sufficient mass to be spherical under their own gravity, having cleared their orbits of smaller objects, and in orbit around the Sun.

That planets be restricted to the Solar System, however, was not by and large expected, with speculation on worlds beyond the Solar System, *exoplanets*, dating back centuries (consider the philosophical treatise of Bruno, 1584). This was confirmed in 1992 with the detection of the first exoplanet, discovered orbiting the neutron star PSR B1257+12 due to variations in the timing of radio pulses (Wolszczan & Frail, 1992). This was swiftly followed by the first detection of an exoplanet orbiting a main sequence star, based on the Doppler shift induced in its stellar host (Mayor & Queloz, 1995). The field has since grown immensely, with the number of confirmed exoplanets, as listed by the NASA Exoplanet Archive (Akeson

et al., 2013), now numbering 6,128 as of 2026 March 3.

In this thesis, I focus on the observational and computational analysis of the atmospheres of giant exoplanets. As the study of exoplanets has grown into a enormous field of research, I here provide a focussed background for the study of exoplanet atmospheres. In Sec. 1.1, I give an overview of the exoplanet population, including the principal methods of observation, with a focus on giant exoplanets. In Sec. 1.2, I describe planetary atmospheres as astrophysical environments, including relevant physical and chemical processes. I then outline in Sec. 1.3 the observational methods of atmospheric characterization and an overview of past results. I provide a summary overview of the thesis focus, structure, and goals in Sec. 1.4.

1.1 The Exoplanet Population

The currently confirmed exoplanet population exceeds 6,000 planets. This population has been steadily built up from a combination of techniques (Fig. 1.1), and covers a diverse range of planetary parameters, including many varieties of planets not represented in the Solar System. These include hot Jupiters – gas giants at very short orbital distances – and super-Earths and sub-Neptunes – planets with masses intermediate between the Earth and Neptune. However, the exoplanet population remains mostly distinct from the Solar System regime, with no extrasolar analogues to most of the Solar System planets, owing to observational biases imposed by the various methods used to detect exoplanets (Fig. 1.2). Jupiter is the exception to this (e.g. Kipping et al., 2016), although our understanding of Jupiter-analogue exoplanets is still highly limited.

The majority of exoplanet detections are through the transit method, discussed in Sec. 1.1.1. This method is restricted to transiting planets, planets with fortuitously aligned orbits such that they eclipse their stellar host along the line of sight of an observer. These only represent a minor fraction of total exoplanets, but enable a level of characterization not typically possible for non-transiting planets with current observatories. The transit method has yielded a diverse population of planets with relatively short orbital periods (mostly less than 10 days), with radii ranging from smaller than Mercury (Barclay et al., 2013) to twice the size of Jupiter (Zhou et al., 2017).

The next most successful method, first proposed by Struve (1952) and used by Mayor & Queloz (1995) in their detection of the first exoplanet around a main sequence star, is the radial velocity method, described in Sec. 1.1.2. This method has contributed planets across the parameter space, including out to greater periods than the transit method, but has mainly been restricted to planets with masses considerably greater than Earth.

The main other methods for exoplanet detections to date are microlensing and direct imaging. The relativistic phenomenon of gravitational lensing (Einstein, 1936) enables a star to act as a lens to a background object through chance alignment




Figure 1.1: Cumulative histogram of detected exoplanets over time by detection method, generated from the NASA Exoplanet Archive (Akeson et al., 2013) on 2026 March 16.




Figure 1.2: Plot of the mass-period distribution of known exoplanets based on detection method, generated from the NASA Exoplanet Archive (Akeson et al., 2013) on 2026 March 16. Solar System planets are marked with their initial.

(Mao & Paczynski, 1991). If the lens star is orbited by a planet, it can impact the lensing in such a way that the planet can be detected (Bond et al., 2004). The properties of planets detected in this way are diverse – including planets with low masses and intermediate orbital periods which are difficult to detect using other methods – making them well suited for understanding demographics, but of limited use for deep investigations of individual systems, as observations cannot be repeated.

Direct imaging of exoplanets is currently restricted by the angular resolution and infrared precision of existing technology, due to the need to mask the far brighter stellar host, and the relative faintness of planetary thermal emission. Nonetheless, direct imaging has played a key role in the detection of exoplanets at extremely distant orbits (e.g., Chauvin et al., 2004; Marois et al., 2008; Lagrange et al., 2009; Macintosh et al., 2015), with periods greater than 1000 days (and much longer in most cases). These planets are all young, retaining residual heat from their formation to provide sufficient luminosity to permit detection. Direct imaging offers a number of advantages over other methods, as the planet is observed directly, enabling signal to noise to be built over continuous observation, removing contamination effects from the star, and permitting spectral characterization of the atmosphere.

Another method, which has only presented a handful of confirmed exoplanets to date but has the potential for a considerable volume of exoplanet detections (Perryman et al., 2014) in the future, is astrometry. Through long baseline astrometric observations, the component of the motion of the star due to orbiting the barycentre from an orbiting planet can be identified (Sahlmann et al., 2013; Curiel et al., 2022; Stefánsson et al., 2025). This method is particularly complementary to other methods, because it favours planets at longer orbits.

1.1.1 Transiting planets

As briefly mentioned, a fraction of exoplanets have orbits that eclipse the disc of their stellar host along the line of sight visible from Earth. The blocking of the stellar surface by the planet during the eclipse, referred to as a transit or primary eclipse, reduces the flux from the star. Through time-series photometric monitoring of stars, we can detect the characteristic dips of exoplanetary transits to deduce the presence of planets, known as the transit method. The first planetary transit was observed by Charbonneau et al. (2000), depicted in Fig. 1.3. The relative simplicity of this detection method, including the ability to observe multiple stars simultaneously, means that transiting exoplanets account for the largest component of the discovered exoplanet population. This has been primarily driven by the space-based surveys from the Kepler space telescope (Borucki et al., 2010) and more recently the Transiting Exoplanet Survey Satellite (TESS, Ricker et al., 2015), in addition to ground-based surveys, for example the Wide-Angle Search for Planets (WASP, Pollacco et al., 2006), the Hungarian-made Automated Telescope Network (HAT-Net, Bakos et al., 2004), and the Next Generation Transit Survey (NGTS, Wheatley et al., 2018).




Figure 1.3: Transit light-curve of HD 209458b, using two nights of phase-folded observations and the best fitting transit model (solid black), with the dashed lines representing model planets with 10% larger and smaller radii. Figure reproduced from Charbonneau et al. (2000).

Transiting planets have immense potential for furthering our understanding of exoplanets, as the shape of the transit light-curve (i.e. the photometric dip in time-series observations), as well as the periodicity of the transit, can be used to calculate a number of planetary properties like radius, inclination, and semi-major axis. Furthermore, spectroscopic observations during transit, and also secondary eclipse (when the star eclipses the planet), can be used to reveal information about the planetary atmospheres (see Sec. 1.3), which is the focus of the observational component of this thesis.

Transit geometry

I here outline a set of equations used to deduce planetary and orbital parameters of exoplanet systems, as well as the observational biases affecting transit observations, adapted from the description of Winn (2010).

The discovery of exoplanets via the transit method is strongly biased towards large planets at short orbital distances, with the majority of discovered transiting planets having orbital periods of less than ten days (see Fig. 1.2). The observational bias towards large planets can be simply identified by recognising the magnitude of the signal (i.e. the flux blocked by the transit), the transit depth δ , is simply given by the ratio of the areas of the planet and the star, and therefore

$$\delta = \left(\frac{R_p}{R_*} \right)^2, \quad (1.1)$$

where R_p and R_* are the planet and stellar radii respectively, and assuming a fully opaque, non-emitting circular planet and uniform stellar intensity. This technique is therefore able to directly measure the radii of exoplanets, given the stellar radius. As the signal is dependent on the ratio of the radii, it is also easier to detect planets around smaller stars.

The observational bias towards close-in planets is a combination of a geometric constraint and a temporal constraint. The former can be understood by considering that for a transit to be observable, the inclination i of an orbit with semi-major axis a must satisfy

$$\cos(i) \leq \frac{R_p + R_*}{a} \left(\frac{1 - e^2}{1 - e \sin(\omega)} \right), \quad (1.2)$$

where e is the eccentricity and ω is the argument of periastron, restated from (Winn, 2010). This means that by taking any given line of sight inclination from a random distribution of inclinations, the geometric transit probability is proportional to the right hand side of the equation, being inversely proportional to a , and also proportional to R_* . The temporal constraint is the result of the frequency of transit events being set by the orbital period P , which is in turn related to a simply via Kepler's third law:

$$P^2 = \left(\frac{4\pi^2}{G(M_* + M_p)} \right) a^3, \quad (1.3)$$

where M_* and M_p are the masses of the star and planet respectively. This is further compounded by the fact that multiple transits are generally required to detect a planet from the transit method alone.

In addition to the transit depth, which can be used to calculate the planetary radius, the duration of the transit is also set by orbital parameters. This requires the definition of four contact points during a transit, depicted in Fig. 1.4, 1 – the beginning of ingress, when the planetary and stellar discs first begin to overlap, 2 – the end of ingress, when the planetary and stellar discs first fully overlap, 3 – the beginning of egress, when the planetary and stellar discs first no longer fully overlap, and 4 – the end of egress, when the planetary and stellar discs last partially overlap. As described by Winn (2010), the time T between contact points 1 and 4 can be approximately expressed based on the impact parameter b , defined as the sky projected distance between the centres of the planetary and stellar discs at conjunction (mid-transit), for an eccentric orbit:

$$b = \frac{a \cos(i)}{R_*} \left(\frac{1 - e^2}{1 - e \sin(\omega)} \right), \quad (1.4)$$

and the transit depth δ , as

$$T_{14} = \frac{P}{\pi} \sin^{-1} \left[\frac{R_*}{a} \frac{\sqrt{(1 + \sqrt{\delta})^2 - b^2}}{\sin(i)} \right] \left(\frac{1 - e^2}{1 - e \sin(\omega)} \right), \quad (1.5)$$




Figure 1.4: Simple schematic defining the contact points T_1 , T_2 , T_3 , and T_4 in a light-curve including an exoplanetary transit. The ingress and egress are defined between T_1 and T_2 , and T_3 and T_4 , respectively, as the times where the planetary and stellar discs partially overlap.

and the time between contacts 2 and 3 written as

$$T_{23} = \frac{P}{\pi} \sin^{-1} \left[\frac{R_*}{a} \frac{\sqrt{(1 - \sqrt{\delta})^2 - b^2}}{\sin(i)} \right] \left(\frac{1 - e^2}{1 - e \sin(\omega)} \right). \quad (1.6)$$

These equations are exact in the case of circular orbits ($e = 0$), where the velocity of the planet is the same between ingress and egress, while the e -dependent terms provide a good approximation for eccentric orbits (Winn, 2010). Solving these two equations with $R_* \ll a$ gives the following expression for the scaled stellar radius a/R_*

$$\frac{a}{R_*} = \frac{2\delta^{1/4}}{\pi} \frac{P}{\sqrt{T_{14}^2 - T_{23}^2}}, \quad (1.7)$$

and in turn the inclination i . These calculations constitute part of the transit model that I use in Sec. 2.3.3, alongside a description of the depth, which depends on the spatial flux distribution of the stellar disc (limb darkening), as described below.

Limb darkening

If stars were uniform light sources, then the transit depth between contacts 3 and 4 would be constant, and the depth during ingress and egress would be determined solely by the geometry of overlapping circles. However, the brightness of the stellar disc varies as a function of distance from the centre – a phenomenon called limb darkening. This gives a characteristic curved shape to the transit light-curve, which is dependent on wavelength. This limb darkening needs to be modelled to obtain

accurate planetary parameters from light-curves.

Limb darkening is the result of viewing the stellar surface at different angles given the spherical nature of stars; as the line of sight impact parameter increases, the star is being viewed at a steeper angle, and therefore a shallower distance into the star. As the temperature of a star increases with depth, the light originating from the stellar limb corresponds to a lower temperature and is therefore fainter compared to the centre. The wavelength dependence in the limb darkening is due to the difference in thermal emission spectra of the stellar atmosphere at different temperatures.

Modelling limb darkening directly requires a complex treatment of the radiative transfer in the stellar atmosphere. It is therefore typical to parametrize the stellar intensity I as a function of $\mu = \cos(\theta)$, where θ is the angle between the line of sight and stellar surface normal, using a set of limb-darkening parameters u_n . A number of limb darkening laws of various complexity are used to describe the dependence of I on μ , including the following from Claret (2000):

$$\text{Linear : } \frac{I(\mu)}{I(1)} = 1 - u_1(1 - \mu) \quad (1.8)$$

$$\text{Quadratic (I) : } \frac{I(\mu)}{I(1)} = 1 - u_1(1 - \mu) - u_2(1 - \mu)^2 \quad (1.9)$$

$$\text{Four parameter : } \frac{I(\mu)}{I(1)} = 1 - u_1(1 - \mu)^{1/2} - u_2(1 - \mu) - u_3(1 - \mu)^{3/2} - u_4(1 - \mu)^2. \quad (1.10)$$

In this thesis I also use quadratic limb-darkening re-parametrized following Kipping (2013), to account for the strong degeneracy between u_1 and u_2 which makes fitting limb darkening coefficients difficult, where $q_{1,2}$ are new parameters calculable from $u_{1,2}$:

$$\text{Quadratic (II) : } \frac{I(\mu)}{I(1)} = 1 - 2q_2\sqrt{q_1}(1 - \mu) - \sqrt{q_1}(1 - 2q_2)(1 - \mu)^2. \quad (1.11)$$

The appropriate choice of limb darkening law depends on the precision and cadence of the light-curve, with more complex choices of limb darkening only outperforming simpler laws with higher precision or better sampled light-curves (e.g., Espinoza & Jordán, 2015, 2016). Departing from simple limb darkening laws means that transit models cannot be generated analytically (Mandel & Agol, 2002), and must be computed numerically (e.g. Kreidberg, 2015a).

1.1.2 Radial velocity

Another way that exoplanets can be detected from stellar observations is through the radial velocity method. In a planetary system, the star and planets orbit the centre of mass of the system, the barycentre. This will be offset from the centre

of the star by any orbiting planets, inducing a wobble in the star’s motion periodic with the planet’s orbit, the radial component of which can be observed through the Doppler shift of stellar spectral features. This requires extremely high precision observations of stellar spectra, repeated to sample the radial velocity curve of at least one planetary orbit. In addition to being a successful method of exoplanet detection, the radial velocity method is powerful as it permits constraints of planetary masses, based on the semi-amplitude of the induced radial velocity K_*

$$K_* = \left(\frac{2\pi G}{P} \right)^{\frac{1}{3}} \frac{M_p \sin i}{(M_p + M_*)^{\frac{2}{3}}} \frac{1}{(1 - e^2)^{\frac{1}{2}}}, \quad (1.12)$$

using the orbital period P , eccentricity e , and inclination i (Perryman, 2018).

The first radial velocity detections of exoplanets were of gas giants at short orbital distances, given the bias towards planets with high mass and short orbital periods present in Eq. 1.12 (e.g. Mayor & Queloz, 1995, as depicted in Fig. 1.5). However, while the RV signal depends directly on orbital period ($\propto P^{-1/3}$) unlike the transit method (where the transit depth is independent of orbital period), the frequency and geometric transit probability represents a much stronger selection bias against the detection of long period transiting planets. Radial velocity detections have therefore played the dominant role in building up the population of known giant planets at longer orbital periods, from 100 days, up to 50,000 days when combined with astrometric data (e.g., Feng et al., 2022; An et al., 2025).

In order to deduce planetary system parameters from radial velocity observations, the time-series radial velocity curve is fit using an orbital model (possibly including components from multiple planets), including the semi-amplitude K_* and period P , with the shape of the curve influenced by the eccentricity. However, because the amplitude is a measure of $M_p \sin i$, it is not possible to disentangle the mass of the planet from the inclination, meaning the masses obtained by radial velocity measurements constitute a lower limit.

As the inclination can be measured independently for transiting exoplanets, this allows precise measurements of the planetary mass. This means we can measure both the mass and the radius for transiting exoplanets, which can in turn be related to density, and therefore composition. For some planetary densities, this allows direct constraints on the bulk composition based on interior structure models, for instance identifying the envelope mass fraction for gas giants. In other parts of the parameter space, the planetary composition remains highly degenerate from density measurements, with variable fractions of a metal-rich core, silicates, water and other ices, and an atmosphere.

Another advantage of the radial velocity technique for transiting exoplanet systems is direct measurement of the surface gravity of the planet g_p , by combining knowledge of the period P , as well as K_* and e from radial velocity measurements, and R_p , a , and i from the transit light-curve, based on the following equation identified by Southworth et al. (2007)




Figure 1.5: Phase-folded radial velocity curve of the hot Jupiter 51 Pegasi b, demonstrating the fitted orbital motion of the planet induced on the star (black curve). Figure reproduced from Mayor & Queloz (1995).

$$g_p = \frac{2\pi}{P} \frac{(1 - e^2)^{\frac{1}{2}} K_* a^2}{R_p^2 \sin i}. \quad (1.13)$$

This method results in reduced uncertainty in comparison to calculating the gravity from the mass and radius, and can be understood by considering that the semi-amplitude of the radial velocity is directly caused by the planet's gravity.

Rossiter-McLaughlin Effect

Another advantage of measuring the radial velocity of transiting exoplanet systems is that it allows measurement of the Rossiter-McLaughlin effect (Rossiter, 1924; McLaughlin, 1924). While observing the stellar radial velocity during planetary transit, a portion of the stellar disc will be obscured by the planet, and therefore the radial velocity contributions from the obscured component will be subtracted




Figure 1.6: Diagram illustrating the Rossiter-McLaughlin effect of planets with different obliquities, λ , with the top panels demonstrating the orbital trajectories of the planet relative to the stellar rotation, and the bottom panel displaying the associated radial velocity curve. Figure reproduced from Gaudi & Winn (2007).

from the overall radial velocity measured. This is important because the star will be rotating, so part of the disc will be rotating towards the observer and hence blue-shifted and the other half will be rotating away and consequently red-shifted. The velocity of the missing flux induces a characteristic shape to the radial velocity signal during transit, which is dependent on the impact parameter of the planet's orbit, and the obliquity, the angle between the spin axis of the star and the sky-projected orbital plane of the planet. The Rossiter-McLaughlin effect was first measured for an exoplanet by Queloz et al. (2000).

Detection of the signal can in turn be used to deduce the spin-orbit alignment of the system, which can be related to the migration history of the planet (see Fig. 1.6), with multiple potential causes of orbital misalignment, notably including high-eccentricity migration of planets (e.g., Rasio & Ford, 1996; Wu & Murray, 2003; Muñoz et al., 2016), or primordial or induced misalignment of the protoplanetary disc (e.g., Bate et al., 2010; Batygin, 2012; Huber et al., 2013). The consideration of the impact of planet migration on system alignment constitutes an essential component of the hypothesis tested by the BOWIE-ALIGN programme described in Chapters 4 and 5.

1.1.3 Gas giant demographics

Gas giants are planets with masses dominated by envelopes of hydrogen and helium. Through our discovery surveys of exoplanets, we know that gas giants can be found at a wide range of orbital separations from their stars and therefore a wide range of instellations. Known gas giants vary from ultra-hot Jupiters with temperatures exceeding that of some stars (e.g. KELT-9b, Gaudi et al., 2017), to distant worlds with orbital periods exceeding a million years (e.g. COCONUTS-2b, Zhang et al.,

2021b). The mass of a giant planet is often referenced in its name (super-Jupiter, Jupiter, Saturn, sub-Saturn, or Neptune) depending on its mass relative to Solar System counterparts.

Of particular interest are hot Jupiters, often defined as gas giants with orbital periods of less than ten days (Wang et al., 2015), which are the best studied of the exoplanet sub-populations, given the strong observational biases favouring both their detection and characterization. The cold gas giants of the Solar System, at 5.2 and 9.6 au for Jupiter and Saturn respectively, are obviously our best understood individual giant planets, due to the wealth of information obtained by both spatially-resolved remote sensing and in-situ measurements.

Radial velocity surveys have yielded exoplanets across the range of orbital separations between hot Jupiters and the Solar System (e.g., Knutson et al., 2014; Rosenthal et al., 2021), but the number of transiting planets bridging this gap is relatively limited (e.g. the Jovian analogue Kepler-167e, Kipping et al., 2016). Recent advancements in planet detection, particularly photometric and spectroscopic follow-up of single transit events from TESS, have significantly increased the number of transiting giant planets at intermediate orbital separations orbiting sufficiently bright stars to be amenable to atmospheric follow-up and mass measurements (e.g., Cañas et al., 2020; Gill et al., 2020, 2024; Hobson et al., 2023; Mann et al., 2023; Mireles et al., 2023; Battley et al., 2024; Rescigno et al., 2024; Ulmer-Moll et al., 2025).

Giant planet occurrence rates

A combination of radial velocity, transiting planets, direct imaging, and gravitational microlensing have been used to build up a distribution describing the occurrence rates of giant planets at different masses and orbital separations, based on stellar host. Radial velocity surveys in particular have been critical in describing the occurrence rates of giant planets, providing constraints on the distribution out to orbital separations of about 10 au (e.g., Fernandes et al., 2019; Fulton et al., 2021). A finding of these analyses is that there is a correlation between the occurrence rate of giant planets and both stellar host metallicity and stellar host mass (e.g., Johnson et al., 2010; Wang & Fischer, 2015; Ghezzi et al., 2018). Giant planets are rare around M dwarf stars, with most gas giants orbiting M dwarfs found around the most massive M dwarfs (e.g., see Gan et al., 2023; Bryant et al., 2023), although a gas giant has recently been found transiting a $0.2 M_{\odot}$ M dwarf (Bryant et al., 2025).

As can be seen in Fig. 1.7, the occurrence rate of giant planets is significantly greater beyond $\sim 1\text{--}3$ au, meaning hot Jupiters are comparatively rare. This roughly maps onto the location of the water-ice line, the orbital separation in the protoplanetary disc where water condenses into solid form, which is proposed to enable more efficient formation of giant planets (Kennedy & Kenyon, 2008). It remains an ongoing question as to whether the gas giants interior to this, hot and warm Jupiters, therefore migrated inwards, or formed in-situ, which I discuss further in Sec. 1.1.4.




Figure 1.7: Estimated occurrence rates of giant planets at different semi-major axes (black), based on counts from the California Legacy Survey (Rosenthal et al., 2021) using high-precision radial velocity measurements (blue dotted) corrected for observational biases. The enhancement in occurrence between 1–10 au can be clearly seen, with the occurrence rate beyond that being relatively unconstrained. Figure reproduced from Fulton et al. (2021).

Mass-metallicity relationship

The giant planets of the Solar System demonstrate a relationship of decreasing metallicity (abundance of elements heavier than helium) with increasing planetary mass (e.g. Guillot, 2005). It is therefore possible to use exoplanets to probe whether this relationship holds on a larger population level. This can be done using the bulk composition of giant planets, based on their mass and radii to extract their density, which can in turn be related to the bulk metallicity of the planet through interior structure models. Thorngren et al. (2016) demonstrate the relationship between the mass of heavy elements, M_Z , relative to the planetary mass, M_p , roughly follows $M_Z \propto \sqrt{M_p}$, based on interior structure modelling using the densities and ages of known giant exoplanets. Chachan et al. (2025) further refine this relationship as

$$M_Z = M_{\text{core}} + f(M_p - M_{\text{core}}), \quad (1.14)$$

where M_{core} and f are fitted parameters, with the former representing a core composed purely of heavy metals, depicted in Fig. 1.8

This method has limited application to hot Jupiters however, which have inflated radii. The cause of this radius inflation is not completely understood, with a variety of methods put forward to explain this observed radius anomaly (see Burrows et al., 2007), and can be sorted into mechanisms involving an additional internal heat source (e.g., tidal heating, dynamical heating, or Ohmic heating, see Bodenheimer




Figure 1.8: Mass of heavy elements of giants planets, calculated using interior structure models based on their age and density, as a function of their mass. The model fitted by Chachan et al. (2025) is shown as a fit to the data, with the solid black line representing the mean models, blue representing samples of the models, and the dashed black lines representing the range of astrophysical scatter. Figure reproduced from Chachan et al. (2025).

et al., 2001; Showman & Guillot, 2002; Batygin & Stevenson, 2010, respectively), and mechanisms where the cooling of giant planets from formation is delayed (e.g., by increased opacity or suppressed convection by compositional gradients, see Burrows et al., 2007; Chabrier et al., 2007, respectively). This uncertainty in inflation mechanisms limits the ability of interior structure models to accurately determine the bulk metallicity of hot Jupiters.

Another measurable quantity is atmospheric metallicity, which in some cases can be related to the bulk elemental ratios within the hydrogen envelope. The metal enhancement of the Solar System giants has been measured for the volatile species carbon, nitrogen, and sulfur (e.g., Wong et al., 2004; Fletcher et al., 2009). Initial attempts have been made to characterize the mass-metallicity relationship using atmospheric observations of hot Jupiters (e.g., Barstow et al., 2017; Pinhas et al., 2019; Welbanks et al., 2019), although Sun et al. (2024) have demonstrated that a larger dataset is needed to properly discern a linear trend with strong evidence. The atmospheres of planets can be depleted in heavy elements relative to the interior (see Baraffe et al., 2008), depleted of certain volatile species by cloud formation (e.g. Helling et al., 2016), or enriched by later accretion by being chemically uncoupled from their interior (Müller & Helled, 2024).

1.1.4 Formation and migration of giant planets

There are two main theories for how giant planets form from the protoplanetary disc, the rotating disc of circumstellar material around a nascent star. In the core accretion scenario, a rocky or icy protoplanetary core will accrete a large quantity of gas from the disc. This requires that a large core can reach a critical mass via the accretion of smaller planetesimals, which will trigger run-away gas accretion (Perri & Cameron, 1974; Pollack et al., 1996). This mechanism is typically invoked to explain the mass distribution of planets in the Solar System, suggesting that planetary cores beyond the ice line grew larger due to the presence of solid ices, and were therefore able to undergo run-away gas accretion (Pollack et al., 1996). Core accretion also explains the mass-metallicity relationship of the Solar System giant planets (Lambrechts et al., 2014).

The alternative method of giant planet formation is gravitational instability, where regions of a protoplanetary disc reach sufficient temperatures or densities due to hydrodynamical interactions to trigger disc fragmentation and gravitational collapse (Boss, 1997; Durisen et al., 2007). This mechanism is required for the highly distant directly imaged exoplanets in stable, broadly circular orbits, where the timescales of core accretion are unfeasibly long (Dodson-Robinson et al., 2009), and scattered planets from the inner disc would not possess such orbits. Gravitational instability is not a potential mechanism for the formation of giant planets at close orbital distances, where the protoplanetary disc is too hot (Rafikov, 2005).

Origin of hot and warm Jupiters

While core accretion and gravitational instability can both explain the formation of giant planets beyond the water-ice line, the origin of hot Jupiter systems is less clear. There are three origin scenarios for hot Jupiters outlined by Dawson & Johnson (2018): in-situ formation, disc migration, and high-eccentricity migration.

Both methods of gas giant formation present issues for forming planets at the close orbital separations where hot Jupiters are found (Rafikov, 2005, 2006). High disc temperatures and rapid rotation at short orbital periods prevent gravitational instability within the inner protoplanetary disc, preventing the criterion of gravitational instability (Toomre, 1964) from being reached, and shearing the clump before it can cool even if the disc did fragment (Rafikov, 2005). Core accretion has also been traditionally disfavoured due to the impracticalities of forming a sufficiently massive core from the solids available at close orbital radii (e.g. Rafikov, 2006), however that has been revisited given the abundances of super-Earths of the necessary mass found at those orbital distances (Lee et al., 2014; Lee & Chiang, 2016; Batygin et al., 2016). Whether in-situ formation is a possible origin scenario of hot Jupiters therefore remains an outstanding question (see Dawson & Johnson, 2018, for more details).

The more likely origin of hot Jupiters is that they have migrated to their current location. There are two principle methods for this considered by Dawson &

Johnson (2018), disc migration and high-eccentricity migration. In disc migration, a proto-hot Jupiter will exchange angular momentum with the protoplanetary disc before dispersal, resulting in a reduction of orbital radius. Due to the constant braking applied by interaction with the disc, the orbit remains roughly circular as it migrates inwards, and remains in the orbital plane of the disc. Assuming that the disc is aligned with the stellar rotation, this results in a hot Jupiter with an orbit that is also aligned to the stellar rotation.

High-eccentricity migration on the other hand, transforms a cold Jupiter that formed beyond the water-ice line into a hot Jupiter some time after disc dispersal. This is a two step process: first some kind of perturbation from another planet or star excites the planet onto a highly eccentric orbit (e.g., Rasio & Ford, 1996; Wu & Lithwick, 2011; Petrovich, 2015). The planet’s orbit will then circularize through dissipative tidal interactions with the star at periastron, resulting in a close, circular orbit. Due to the nature of dynamical processes that promote planets into high eccentricity orbits, the final orbit of a planet which migrated in such a manner is expected to be misaligned from the stellar rotation.

Measurements of the obliquities of hot Jupiters demonstrate a lack of misaligned hot Jupiters around cooler stars (Winn et al., 2010; Albrecht et al., 2012). This transition in the obliquity distribution roughly lines up with an observed transition in stellar rotation rates with a similar stellar host temperature dependence ($T_{\text{eff}} = 6400\text{--}6600\text{ K}$ depending on metallicity), known as the Kraft break (Kraft, 1967). The Kraft break is caused by a transition in the depth of the stellar convective envelopes; stars above the Kraft break have shallower convective envelopes and therefore less efficient magnetic dynamos. With less efficient magnetic dynamos, stars above the Kraft break do not efficiently slow their rotation rate through magnetic braking, where the stellar magnetic field is coupled to the stellar wind causing loss of angular momentum to the interstellar medium. The lack of misaligned hot Jupiters around stars below the Kraft break is therefore also thought to be a consequence of the stellar convective envelope, with stars below the Kraft break able to realign the rotation of their convective envelope to the orbiting hot Jupiter through tidal interactions (Winn et al., 2010), causing an absence of misaligned planets, while stars above the Kraft break lack sufficient convective envelopes to allow tidal realignment. Recent work by Beyer & White (2024) has demonstrated that the transition in the hot Jupiter obliquity distribution occurs at cooler stellar effective temperatures than the Kraft break (by about 450 K), although the convective envelope size required for effective tidal realignment, and for an efficient magnetic dynamo, need not be the same, explaining the offset between these transitions.

Warm Jupiters, here considered as giant planets with orbital distances interior to the water-ice line, but exterior to the distances where strong tidal interactions with the stellar host dominate and override orbital properties ($\sim 0.1\text{--}2\text{ au}$), can also shed light on these formation processes (consider Su et al., 2024). Sub-populations of warm Jupiters with shared dynamical properties can be identified, including highly-




Figure 1.9: C/O ratio of the solids and gas of a protoplanetary disc as a function of orbital distance relative to solar, with the marked snowlines of H₂O, CO₂, and CO. Figure reproduced from Öberg et al. (2011).

eccentric warm Jupiters which are potentially undergoing high-eccentricity migration (Petrovich & Tremaine, 2016), warm Jupiters with inner companions that suggest a roughly in-situ formation scenario to retain such a planetary system (Huang et al., 2016), and the bulk of warm Jupiters, which are preferentially found in aligned orbits suggesting disc migration (Rice et al., 2022; Wang et al., 2024).

Ice lines and chemical composition

I have already mentioned the importance of the water-ice line in the formation location of giant planets; protoplanetary discs are expected to host a number of these ice lines at the sublimation temperatures of different species, notably including H₂O, CO₂, CO, NH₃ and N₂ (see Fig. 1.9). The location in the disc relative to these ice lines determines whether a species is in the form of gas or solids, meaning the eventual planetary composition of a planet is dependent on its formation location. This can in principle be traced using elemental ratios, which can be probed using atmospheric observations (e.g., Öberg et al., 2011; Madhusudhan et al., 2014; Schneider & Bitsch, 2021a). C/O is the main elemental ratio considered, given the abundance of both elements and strong spectral features of common carbon- and oxygen-bearing molecules, but additional information can be obtained by considering other species including nitrogen (Piso et al., 2016; Cridland et al., 2020), sulfur (Turrini et al., 2021), and silicon (Schneider & Bitsch, 2021b). Just as with relating atmospheric and bulk metallicity, an additional complication arises if elemental ratios are non-homogeneous throughout the planet, potentially making it difficult to relate the eventual atmospheric gas composition to the raw materials from which the planet

formed (e.g., Baraffe et al., 2008; Müller & Helled, 2024).

1.2 Planetary Atmospheres

Planetary atmospheres are complex environments, envelopes of gaseous material bound to a planetary interior, heated by radiation both incident from the star and latent heat from formation emerging from below, mixed by temperature contrasts and the rotation of the planet, and host to intricate chemical reaction networks. Studying these atmospheres gives us access to astrophysical laboratories of dizzying complexity. In the Solar System, we have a small number of atmospheres available to us, but the ability to study them in exquisite detail with both remote sensing and in-situ measurements. This includes the terrestrial atmospheres: the dense CO₂ atmosphere of Venus with thick clouds and torrid surface temperature (Krasnopolsky & Parshev, 1981), the tholin haze and hydrocarbon rains of the skies of Titan (Gupta et al., 1981; Brown et al., 2002), and the thin CO₂ envelope and dust storms of Mars (Barth, 1974); and the thick hydrogen-helium envelopes of the Solar System giants, with clouds of H₂O, NH₃ and CH₄ ices and the dynamical cells and storms induced by their rapid rotation (e.g., Lewis, 1969; Atreya & Romani, 1985; Romani & Atreya, 1988).

As discussed in Sec. 1.1, the growing number of exoplanets have dramatically increased the quantity of planetary atmospheres at our disposal, primarily covering parts of the planetary parameter space not sampled by the Solar System. Extrasolar gas giants are notable as being akin to their Solar System counterparts in bulk planetary parameters, but typically found much closer to their host stars, resulting in much greater instellation. In this thesis, I focus on these giant exoplanets, ranging in instellation from hot Jupiters, with equilibrium temperatures greater than 500 K, to temperate Jupiters, spanning the atmospheric temperatures of terrestrial Solar System objects, including the habitable temperatures of Earth (200–500 K).

While our knowledge of exoplanet atmospheres is limited by observational constraints, with limited access to spatial information and no opportunity for in-situ measurements, the power of spectroscopy can get us a long way in inferring the physical processes of exoplanet atmospheres. An understanding of the underlying physics and chemistry is necessary to make these inferences; in this section I discuss these core radiative transfer, transport, chemical, and aerosol processes.

1.2.1 Radiative transfer and opacities

Radiative transfer, the transport of photons through the atmosphere of a planet, is essential to both the key observational techniques used to study exoplanet atmospheres, and in transporting energy through the atmosphere. Exoplanet atmospheres contain sources of both absorption and scattering opacity from the atoms, molecules, and aerosol particles that make up the atmosphere, with each of these opacity sources being wavelength-dependent. A thorough analytical approach to

radiative transfer can be found in Chandrasekhar (1960), while specific theory and applications for Earth’s and exoplanet atmospheres can be found in Goody & Yung (1989) and Seager (2010) respectively.

Radiative transfer equation

I provide a brief derivation of the core radiative transfer equation, used in this thesis for both numerical simulations and synthetic observations, based on Seager (2010).

The static radiative transfer equation describes the time-independent change in the intensity, I , of a beam of light of frequency ν , as it travels along a light-path s in direction $\hat{\mathbf{n}}$, with emission $\varepsilon(\mathbf{x}, \hat{\mathbf{n}}, \nu)$:

$$\frac{dI(\mathbf{x}, \hat{\mathbf{n}}, \nu)}{ds} = -\kappa(\mathbf{x}, \nu)I(\mathbf{x}, \hat{\mathbf{n}}, \nu) + \varepsilon(\mathbf{x}, \hat{\mathbf{n}}, \nu), \quad (1.15)$$

where the spatial dimension is encoded in Cartesian coordinates $\mathbf{x}$, which is in turn a function of pressure and temperature (which typically vary spatially through an atmosphere). The opacity κ , with units of cm^{-1} , represents the total extinction of a volume of gas, which can in turn be computed from the mass attenuation coefficients k (with units of cm^2/g) and densities ρ of all the constituent species, i , as

$$\kappa = \sum_i \rho_i k_i, \quad (1.16)$$

and includes both absorption (α) and scattering (σ_s) terms:

$$\kappa(\mathbf{x}, \nu) = \alpha(\mathbf{x}, \nu) + \sigma_s(\mathbf{x}, \nu). \quad (1.17)$$

κ can typically be considered as a scalar (independent of $\hat{\mathbf{n}}$) in atmospheric radiative transfer, as most atmospheric extinction sources are isotropic, and even anisotropic materials (e.g. non-spherical aerosol particles) will frequently lack an overall alignment, due to a random distribution of orientations for individual aerosols (for counter examples, one can consider light pillars and sundogs visible on Earth caused by aligned ice crystals, as well as potentially mechanical alignment of aerosols by strong winds in hot Jupiters, see Mullens & Lewis, 2025).

The emission term, ε , can in turn be written in terms of a thermal emission $\varepsilon_{\text{therm}}$ and a scattering term $\varepsilon_{\text{scat}}$. Both of these terms can have very complex behaviours. For the thermal emission term, we can assume local thermodynamic equilibrium (LTE) in order to write the term in terms of the Planck function describing black body radiation B through Kirchoff’s Law:

$$\varepsilon_{\text{therm}} = \alpha(\mathbf{x}, \nu)B(\mathbf{x}, \nu). \quad (1.18)$$

This assumption demands that within a small region of the atmosphere, where differences in temperature, pressure, and chemical composition are small compared to the mean free path of a photon, the atmosphere is in complete thermodynamic equilibrium with the exception that the local radiation field can depart from that

of a black body. LTE is valid in most parts of a planetary atmosphere, as densities are sufficiently high for collisions to thermalize the atmosphere (such that energy level distributions follow that of the local kinetic temperature). Departure from LTE can occur in the upper atmosphere, shown to cause non-negligible differences from LTE in the temperature profile of the solar system gas giants atmospheres at μ bar pressures (Appleby, 1990).

Scattering is even more complex, as the scattered intensity can depend on the scattering angle (anisotropic scattering). In the simple case of isotropic scattering, where every scattering angle is equally likely, the scattering function can be simply written in terms of the solid angle-averaged intensity $J \equiv \langle I \rangle$

$$\varepsilon_{\text{scat}} = \sigma_s(\mathbf{x}, \nu) J(\mathbf{x}, \nu). \quad (1.19)$$

In the case the length scale of the atmosphere is small compared to the radius of the planet, which is typically true for the purposes of exoplanet characterization, we can usually make the assumption of a plane-parallel atmosphere. This effectively collapses the problem to a single dimension, by assuming axial and azimuthal symmetry such that the temperature T , pressure P , and density ρ only vary with altitude z . This assumption requires that the atmospheric height be small compared to the planetary radius (which is true for any conventional planet). This allows us to rewrite Eq. 1.15 as follows:

$$\cos \theta \frac{dI(z, \theta, \nu)}{dz} = -\kappa(z, \nu) I(z, \theta, \nu) + \varepsilon(z, \theta, \nu), \quad (1.20)$$

where θ is the angle between z and the beam, such that

$$\frac{dz}{ds} = \cos \theta. \quad (1.21)$$

This equation can be further simplified by introducing the quantity optical depth, τ , defined thus:

$$d\tau(s, \nu) = -\kappa(s, \nu) ds, \quad (1.22)$$

which can be shown to equal the mean free path of a photon, as well as by defining the source function S :

$$S(\mathbf{x}, \hat{\mathbf{n}}, \nu) = \frac{\varepsilon(\mathbf{x}, \hat{\mathbf{n}}, \nu)}{\kappa(s, \nu)}. \quad (1.23)$$

This allows us to arrive at this canonical form of the static radiative transfer equation, with the further definition of $\mu \equiv \cos(\theta)$:

$$\mu \frac{dI(\tau_\nu, \mu, \nu)}{d\tau_\nu} = I(\tau_\nu, \mu, \nu) - S(\tau_\nu, \mu, \nu). \quad (1.24)$$

This is the equation solved for when considering both energy transport through a planetary atmosphere, and the transmitted and emitted flux measured in exoplanet

observations. The complexity of the equation is entirely captured within the source function S , and describes flux from both below ($\mu < 0$) and above ($\mu > 0$).

Molecular and atomic opacities

One of the main opacity sources in exoplanet atmospheres is the absorption of photons by atoms and molecules to enter a higher quantum energy level. This transition occurs at a discrete wavelength λ relating to the change in energy ΔE of the absorber via the Planck relation $\lambda = hc/\Delta E$. This discrete wavelength-dependent absorption is responsible for the line core, which in some cases can be directly identified even in low resolution spectroscopy. In high resolution observations, the unique wavelength-dependence of forests of lines can be used as fingerprints to identify molecular and atomic species. Most commonly in low resolution spectroscopy, increased broadband opacity in a wavelength containing many different spectral lines is used to identify absorption due to molecular species. As an example, the infrared absorption cross-section spectrum of CH_4 is depicted in Fig. 1.10. I here give an outline of the line opacity sources and profiles relevant to atmospheric observations; a more detailed account of these opacity sources can be found in Bernath (1995) and Goody & Yung (1989, chapter 3).

Electronic transitions from atomic species, where bound electrons move between energy states, are particularly notable sources of opacity at optical wavelengths. In particular, the alkali metals sodium and potassium have strong absorption features identifiable even at low resolution (Burrows et al., 2000; Seager & Sasselov, 2000), corresponding to promotion of their single valence electron into an excited state. These occur at 589 nm and 768 nm respectively, and are actually doublets due to spin-orbit coupling. Other neutral metallic species, like iron and titanium, also have observable optical spectra from electronic transitions, but these consist of much more complex distributions of lines arising from their wealth of valence electrons. Certain molecules like TiO and VO also have electronic transitions expected to give rise to considerable optical opacity in exoplanet atmospheres (e.g., Burrows & Sharp, 1999; Hubeny et al., 2003).

Most near-infrared and mid-infrared opacity sources relevant to exoplanet atmospheres are the result of rotational-vibrational (ro-vib) transitions of molecules. At its simplest, these transitions originate as quantum modes of centre of mass motion of atoms in the presence of a permanent or induced dipole moment (i.e. where the centre of mass is different from the centre of charge). Homonuclear molecules (for example diatomic N_2) by and large do not undergo these transitions as a result of lacking a permanent dipole moment, but heteronuclear molecules without permanent dipole moments are able to undergo transitions through asymmetric vibrational modes, which can induce a dipole moment. Rotational transitions arise from the quantization of the rotational energy of the whole molecule, with energies at microwave wavelengths. Vibrational states arise from the motion of atoms relative to one another within a molecule, and cause transitions with energies corresponding




Figure 1.10: Absorption cross-section spectrum of CH_4 at temperatures of $T = 2000$ and $T = 300$ K at infrared wavelengths, using the CH_4 line-list of Yurchenko et al. (2024). Broad peaks correspond to vibrational transitions, with the sub-structure within these peaks is due to different rotational energy levels. Pressure and thermal broadening, and the temperature-dependent distribution of rotational energy states, serve to blend the forest of narrow lines into these broad absorption features. Figure reproduced from the ExoMol website (<https://www.exomol.com/data/molecules/CH4/12C-1H4/MM/>), accessed 2026 April 29.

to infrared wavelengths. As vibrational transitions will always induce a rotational transition as well, each vibrational transition actually corresponds to a series of closely-packed lines, referred to as ro-vib structure. The broadband spectral features observed in low resolution spectroscopy typically correspond to either a single vibrational transition (broadened by its ro-vib structure), or sometimes to branches with the ro-vib structure (distinct substructure resultant from an increase, decrease, or in some molecules no change in, the total angular momentum quantum number), depending on spectral resolution. High resolution observations can be sensitive to the entire forest of lines from ro-vib spectra, contingent on the wavelength range and spectral resolution of the instrument. Molecules predicted (e.g. Burrows et al., 2000; Seager & Sasselov, 2000) and/or observed (see Sec. 1.3) to have key roles in determining atmospheric opacities include H_2O , CH_4 , CO_2 , NH_3 , CO , HCN , H_2S , and SO_2 , with the first three playing a large role in the greenhouse heating of Earth’s atmosphere.

Spectral lines are not infinitely narrow. In addition to a natural quantum broadness arising from the finite lifetime of excited states combined with Heisenberg’s uncertainty principle (minimal compared to thermal broadening above sub-Kelvin temperatures and therefore not relevant to spectral lines of planetary origin), external interactions can cause both thermal and pressure broadening of quantum lines. Thermal, or Doppler, broadening, is a result of Doppler shifts induced by the thermal velocities of gas particles, which are described by the Maxwell-Boltzmann distribution. The frequency distribution around a spectral line with frequency ν_0 resulting from the thermal broadening of spectra lines can be found by substituting

the Maxwell-Boltzmann distribution into the equation for Doppler shift:

$$P_{\text{therm}}(\nu)d\nu = \sqrt{\frac{mc^2}{2\pi kT\nu_0^2}} \exp\left(-\frac{mc^2(\nu - \nu_0)^2}{2kT\nu_0^2}\right) d\nu \quad (1.25)$$

which is a Gaussian profile with standard deviation:

$$\sigma_\nu = \sqrt{\frac{kT}{mc^2}} \nu_0. \quad (1.26)$$

Pressure broadening, or collisional broadening, is the other main source of spectral line broadening in planetary atmospheres, and is caused by collisions interrupting the absorption of photons, reducing the lifetime of excited states and therefore increasing the uncertainty in energy (via Heisenberg’s uncertainty principle) and therefore frequencies of absorption. Pressure broadening is thus named because it depends on the rate of collisions, which is highly pressure dependent. The distribution associated with pressure broadening can be written

$$P_{\text{press}}(\nu)d\nu = \frac{1}{\pi} \frac{\gamma}{(\nu - \nu_0)^2 + \gamma^2} d\nu, \quad (1.27)$$

where γ is a function of pressure, temperature, and chemical composition, and typically needs to be determined from laboratory measurements or numerical calculations. This distribution is a Lorentz profile, resulting in broad wings. Detection of broad Lorentzian wings in a spectral feature can therefore be used to infer that it is being observed deep in an atmosphere (Nikolov et al., 2018).

Continuum opacities

Homonuclear diatomic molecules like H_2 , generally lack infrared absorption features due to their lack of permanent dipole or asymmetric vibration modes. However, dipoles can be induced through close-range intermolecular interactions present during collisions, in turn making even diatomic molecules temporarily able to absorb or emit in infrared vibrational modes. These collisions are very short-lived, resulting in very strongly broadened features which are closer in practice to continuum sources than spectral lines, although they do peak at wavelengths associated with the vibrational transition. These opacities are called collisionally induced absorption (CIA). CIA absorption from $\text{H}_2\text{--H}_2$ and $\text{H}_2\text{--He}$ interactions are a key source of continuum absorption in the infrared for giant planets (e.g. Martin et al., 1976).

One of the main opacity sources at short wavelengths is Rayleigh scattering, with a strong, broadly λ^{-4} , wavelength dependence that gives rise to blue skies on Earth (Rayleigh 1871). Rayleigh scattering describes the interaction of particles with radii much smaller than the wavelength of scattered light, including both gas molecules, and very small aerosol particles. The Rayleigh scattering cross-section due to spherical particles of radius r and refractive indices n , averaged across all scattering angles, is given by:

$$\sigma_s = \frac{8\pi}{3} \left(\frac{2\pi}{\lambda} \right)^4 \left(\frac{n^2 - 1}{n^2 + 2} \right)^2 r^6. \quad (1.28)$$

Small aerosols particles, often referred to as haze, can be treated through Rayleigh scattering. However, larger condensates are often of a similar length scale to the wavelength of incoming light. Mie theory (Mie, 1908) can be used in this case, to determine the scattering and extinction cross-sections, but only in the case of spherical particles. Mie theory can be used to calculate the various parameters to describe physical condensate opacity sources in radiative transfer models (e.g. Kitzmann & Heng, 2018). The absorption spectra of Mie opacity is often relatively spectrally flat at optical and near-infrared wavelengths and is significant at short optical wavelengths, for example being responsible for the grey water clouds on Earth.

1.2.2 Structure, transport and dynamics

In this section I describe a number of fundamental physical processes that dictate the overall pressure-temperature structure of a planetary atmosphere. This includes the static balance of forces that establish the vertical pressure and density variation, a simplistic treatment of energy balance as well as a more complex approach to energy transport, and a brief overview of dynamical mixing processes (for more detail on these topics, see Seager, 2010).

Hydrostatic equilibrium

A planetary atmosphere is a relatively thin envelope of gas bound to a planet by gravity (or in the case of planets with large gaseous envelopes and therefore no clear surface, the atmosphere is typically considered as the outermost part of that envelope, often defined above a select pressure level). Gravity plays an important role in dictating the structure of the atmosphere, with the atmospheric pressure (and hence density) decaying rapidly with altitude. The force of gravity on the atmospheric gas is opposed by the pressure gradient of the gas. In a static atmosphere, this condition is called hydrostatic equilibrium.

To derive an equation for hydrostatic equilibrium, we first consider a plane-parallel atmosphere (see Sec. 1.2.1). We can consider the balance of forces on an element of a small volume of gas with area δA and height δz , and density ρ :

$$\rho g \delta A \delta z = - \frac{dP}{dz} \delta z \delta A, \quad (1.29)$$

where the left hand side is the downwards force of gravity from the gravitational acceleration g , and the right hand side is the net upwards force of the pressures acting on the element, due to the lower pressure from the top surface of the volume. We can then cancel the δA and δz terms, and rearrange for the pressure gradient:

$$\frac{dP}{dz} = -\rho g. \quad (1.30)$$

Density is itself a function of pressure however, so we can use the ideal gas equation of state to substitute in for density:

$$P = \frac{\rho k T}{\mu_m m_H}, \quad (1.31)$$

where T is the temperature, k is Boltzmann's constant, m_H is the atomic mass unit, and μ_m is the mean molecular weight. Combining these equations gives:

$$\frac{dP}{P} = -\frac{\mu_m m_H g}{k T} dz. \quad (1.32)$$

We then solve this to find an expression for the pressure structure of the atmosphere. Usually we would expect the temperature to vary as a function of altitude, $T(z)$. Taking the simple isothermal case where T is constant, and also assuming constant g , we can write:

$$P = P_0 \exp\left(-\frac{\mu_m m_H g}{k T} z\right). \quad (1.33)$$

From this equation we can therefore identify the characteristic exponential length scale of the atmosphere, known as the atmospheric scale height H

$$H = \frac{k T}{\mu_m m_H g}, \quad (1.34)$$

which is proportional to temperature and inversely proportional to gravity and mean molecular weight.

Additional modifications to hydrostatic equilibrium need to be made, particularly in the extremely irradiated and large scale-height atmospheres of hot Jupiters. For example, the assumption that gravity is constant is frequently inaccurate; this is simply addressed by implementing a Newtonian z^{-2} dependence above a reference pressure layer (above which the atmospheric contributions to the gravity cancel following Gauss' law), which is done in all exoplanet atmosphere modelling codes features in this thesis. Centrifugal force from the rotation of the planet can also be important in planets with very weak gravity (Banerjee et al., 2023).

Energy balance

Temperature is an essential parameter in determining atmospheric properties, determining both local behaviour (e.g. chemistry and condensation), and global structure (as seen in the scale height) and circulation. A first approach to consider a planet's atmospheric temperature is to calculate the equilibrium temperature, T_{eq} , considering a simple energy balance approach. This approach can be done using only orbital parameters, plus an assumption of the bond albedo, A_B , the fraction of irradiation scattered back into space, and by neglecting the internal heat of the planet. The

equation of energy balance can be written:

$$\frac{\pi}{f} R_p^2 F_p = (1 - A_B) F_* \left(\frac{R_*}{a} \right)^2 \pi R_p^2, \quad (1.35)$$

with the left hand side representing the power radiated by the planet's flux F_p , and the right hand side representing the power absorbed by the planet from the incident stellar flux F_* (Seager, 2010). f is a correctional factor used to model how heat is redistributed around the planet. In the case of complete heat redistribution, $f = 1/4$, which is relevant to all Solar System atmospheres. We may instead consider the dayside temperature of a planet with no heat redistribution, in which case the appropriate value is $f = 2/3$. This equation simply states the star and planet are in equilibrium (their energies are time-independent).

Assuming the temperature of both the star and the planet radiate as black bodies, we can rewrite the fluxes of each planet using the Stefan-Boltzmann law in terms of an effective temperature:

$$F = \sigma T_{\text{eff}}^4. \quad (1.36)$$

Using T_* to represent the effective temperature of the star and letting the equilibrium temperature equal the effective temperature of the planet $F_p = \sigma T_{\text{eq}}^4$, we can substitute these in to Eq. 1.35 to obtain:

$$T_{\text{eq}} = T_* \left(\frac{R_*}{a} \right)^{\frac{1}{2}} [f(1 - A_B)]^{\frac{1}{4}}. \quad (1.37)$$

The equilibrium temperature, calculated from appropriate albedo and redistribution values, will depart from the effective temperature of the planet, which is determined by directly measuring the flux, if internal heating is significant compared to irradiation. In the Solar System, both Jupiter and Saturn have $T_{\text{eff}} > T_{\text{eq}}$ by about 15 K due to internal heating (e.g. Ingersoll & Porco, 1978).

T_{eq} can be a good approximation of the temperature at the optical photosphere of the planet, but it fails to account for differences in opacity to short-wave incoming radiation from the star, and long-wave outgoing thermal radiation from the planet. This can cause the true temperatures in a planetary atmosphere to differ, failing to account for the greenhouse effect and thermal inversions. For example, while the equilibrium temperature of Venus describes the cloud-top temperature of 230 K, it fails to account for the much hotter surface temperature of 740 K (Basilevsky & Head, 2003). A proper description of energy transport through an atmosphere (either via radiation or convection) is required to properly calculate the vertical temperature distribution.

Radiative equilibrium

For a more sophisticated approach to determine the temperature structure of an atmosphere, we require that every layer be in thermal equilibrium with its adja-

cent layers. This is effectively the requirement that each layer is in static energy balance; the amount of energy entering a given volume is the amount of energy leaving. Energy can be transported through an atmosphere through either radiation or convection. I describe here a brief treatment of how radiative equilibrium can be calculated.

Let us first consider a one dimensional atmosphere in radiative equilibrium, such that the total radiative flux F is constant at every altitude z , and therefore every optical depth τ (see Eq. 1.22)

$$\frac{dF(\tau)}{d\tau} = 0. \quad (1.38)$$

By definition, the total radiative flux is related to the intensity, I , by taking the flux in the upwards direction, multiplying by $\cos(\theta)$, and integrating over all solid angles $d\Omega = \sin\theta d\theta d\phi$:

$$F = \int_0^\infty \int_{-\pi}^\pi \int_0^{2\pi} I \cos(\theta) \sin(\theta) d\theta d\phi = 2\pi \int_0^\infty \int_{-1}^1 I \mu d\mu d\nu. \quad (1.39)$$

We therefore examine the radiative equilibrium condition by simply integrating the radiative transfer equation (Eq. 1.24) over all solid angles and over all wavelengths

$$\int_0^\infty \int_{-1}^1 \mu \frac{dI}{d\tau} d\mu d\nu = \int_0^\infty \int_{-1}^1 (I - S) d\mu d\nu = 0 \quad (1.40)$$

and identifying that the left hand term is proportional to $\frac{dF}{d\tau}$ to set the equation to 0. In the case of pure absorption (no scattering), this differential equation can be integrated analytically for the intensity I (e.g. Chandrasekhar, 1960; Heng et al., 2014). In the case of scattering, due to the appearance of the intensity within the source term, approximations and numerical methods are required. Appropriate boundary conditions are used for the flux at the top of the atmosphere, set with an incident stellar flux spectrum, and the bottom of the atmosphere, set with an internal flux $F_{\text{int}} = \sigma T_{\text{int}}^4$.

One such approach is the two-stream approximation (e.g., Meador & Weaver, 1980; Heng et al., 2014), wherein the flux is split into two different constant values for the upwards and downwards hemispheres. There are also different approaches to incorporating opacity, ranging from the simple “double-grey” approach, where an average opacity in two separate bands of radiation from stellar (optical) and thermal (infrared) radiation are considered, to opacity sampling based on real gas mixtures. With some approximations, the pressure-temperature profile can be solved analytically, for example the commonly used double grey model temperature profile of Guillot (2010). A selection of measured Solar System and modelled hot Jupiter pressure temperature profiles are depicted in Fig. 1.11.




Figure 1.11: Published pressure-temperature profiles of various planets (solid lines), each fitted with the parametric pressure-temperature profile of Madhusudhan & Seager (2009). Data for Solar System objects are taken from Lindal et al. (1981, 1985, 1987, 1990) respectively for Jupiter, Saturn, Uranus, and Neptune. Exoplanet data is from radiative-convective equilibrium modelling of hot Jupiters from Fortney et al. (2008); Tinetti et al. (2007), plus F06, B08, and K09 referring to Fortney et al. (2006), Burrows et al. (2008), and Knutson et al. (2009) respectively. Figure reproduced from Madhusudhan & Seager (2009).

Convective equilibrium

In regions of high opacity and steep temperature gradients, convection can become a more efficient method of energy transport than radiation. To determine which transport mechanism will dominate, the Schwarzschild criterion (Schwarzschild, 1958) can be considered in each layer of the atmosphere to determine whether the atmosphere there is stable against convection. If a small volume of gas within a temperature gradient dT/dz is displaced by a distance δz , it must become less buoyant than the surrounding gas to not continue to rise. This is the condition for convective stability, as otherwise any small perturbation would grow unstably, and the temperature gradient dT/dz would become set by the change in temperature of moving material

(i.e. by convection), rather than whatever originally set the temperature gradient (i.e. the radiation field).

To determine the criterion, assume that the perturbation occurs adiabatically, such that the displaced volume of gas does not exchange energy with its surroundings. Therefore, the temperature change of the gas volume is set by the adiabatic lapse rate $(dT/dz)_{\text{adi}}$. For an adiabatic change, the change in internal energy dU equals the work done dW by the system:

$$m(C_P - R)dT = -PdV, \quad (1.41)$$

where the left hand side of the equation is the internal energy $dU = mC_V dT$, in terms of the mass m , heat capacity at constant pressure $C_p = C_v + R$, and change in temperature dT ; and the right hand side is the work done $dW = PdV$. R is the specific gas constant. Using the equation of state of an ideal gas (Eq. 1.31) and $V = m\rho^{-1}$, the right hand side can be rewritten $mRdT - m\rho^{-1}dP$, and therefore the conservation of energy equation becomes:

$$\frac{dT}{dP} = \frac{1}{\rho C_P}. \quad (1.42)$$

Finally substituting in dP using the expression for hydrostatic equilibrium (Eq. 1.32):

$$\left(\frac{dT}{dz}\right)_{\text{adi}} = -\frac{g}{C_P}. \quad (1.43)$$

An atmosphere will be convectively unstable then if the atmospheric (i.e. radiative) temperature gradient is steeper (more negative) than this adiabatic profile, in which case the energy transport will be convective and the temperature profile will equal this adiabatic lapse rate.

Therefore to determine the overall pressure-temperature structure of an exoplanet atmosphere, an initial calculation of radiative equilibrium can be made to determine an initial temperature profile $T(P)$. This solution must then be investigated and adjusted for convective stability, and the radiative equilibrium re-calculated between convective regions of the atmosphere. This combined approach is typically referred to as radiative-convective equilibrium.

Dynamics and mixing

Atmospheric dynamics play a key role in both energy and chemical transport, and is modelled in three dimensions as part of global circulation models (GCMs, consider Showman et al., 2010). A full consideration of atmospheric dynamics is beyond the scope of this thesis however, as I restrict my analyses to a single dimension in both modelling and observation. However, it is important to introduce the concept of vertical mixing, as it factors into both the distribution of molecules through the atmosphere, and the size and vertical distribution of aerosol species.

Vertical mixing is typically parametrized using the eddy diffusion coefficient,

K_{zz} , collapsing all mixing processes into a single dimension and including various unresolvable processes including turbulence, convection, and the vertical component of 3D circulation. The vertical flux of a species Φ_i can be written in terms of the K_{zz} as

$$\Phi_i = -K_{zz}n\frac{dX_i}{dz}, \quad (1.44)$$

where X_i is the mixing ratio of species i , and n is the total number density of molecules (for a derivation, see Hu et al., 2012). K_{zz} is a very poorly constrained parameter, able to vary by orders of magnitude, particularly in radiative regions, where no clear dynamical process can be identified for setting K_{zz} . In convective regions, the K_{zz} can be loosely estimated through mixing length theory (e.g. Lee et al., 2024), such that

$$K_{zz} \sim wL, \quad (1.45)$$

where w is a characteristic vertical velocity, which can be estimated based on the difference between the local temperature profile and the adiabatic temperature profile, and L is some characteristic length, usually taken as some fraction of the scale height.

Vertical mixing plays a key role in setting the particle size distribution of condensate aerosols (Ackerman & Marley, 2001), and in inducing disequilibrium chemistry, both through transport processes (e.g. Tsai et al., 2017) and photochemistry (e.g. Ohno & Fortney, 2023b).

1.2.3 Atmospheric chemistry

As discussed in 1.2.1, molecular absorption plays a key role in determining atmospheric opacity. This means that atmospheric composition has a strong influence on the temperature structure of the atmosphere, as seen on Earth and Venus, which both have considerable heating of the lower atmospheres due to infrared absorption from CO_2 (the greenhouse effect). Furthermore, the opacity from atmospheric species is imprinted on both the transmitted and emitted atmospheric flux, allowing the chemistry to be directly constrained from observations. It is therefore useful to consider how atmospheric chemistry is determined. Our knowledge of Solar System atmospheres demonstrates that the atmospheres of rocky objects are highly diverse: Earth, Venus, and Titan all contain N_2 , but found alongside O_2 on Earth, CO_2 on Venus, and CH_4 on Titan. Meanwhile the giant planets all possess somewhat similar atmospheric compositions dominated by H_2 , He, and CH_4 . The gas giant atmospheric composition can be understood broadly in the context of chemical equilibrium, while for the former atmospheres a consideration of chemical kinetics and the interior-atmosphere connection is required. In this section I first outline chemical equilibrium, and then discuss departures from chemical equilibrium relevant to gas giant atmospheres.

Chemical equilibrium

An isolated system at fixed temperature and pressure should eventually reach chemical equilibrium. This is a chemical composition that for a given temperature, pressure, and elemental abundances, maximizes the entropy of a system. Chemical reactions still occur under equilibrium conditions, but the overall ratio of different chemical species remain constant. Here I outline a thermodynamic derivation of the central calculation of equilibrium chemistry (for more detail on the central thermodynamics, consider e.g. Blundell & Blundell, 2006), and the relevant equilibrium transitions for moderately irradiated giant exoplanet atmospheres. Our understanding of the equilibrium chemistry of giant exoplanets, and the relevant pressure-temperature dependent transitions, is particularly built on that of brown dwarfs (e.g. see Lodders & Fegley Jr, 2006). For more strongly irradiated planets, thermal dissociation and ionization must also be considered (e.g. Lothringer et al., 2018).

The second law of thermodynamics states that the entropy of an isolated system can only increase, therefore the equilibrium state of any isolated system is the one with maximum entropy. When considering chemistry however, an individual reaction does not constitute an isolated system, as while a reaction will change the entropy of its constituent atoms, it will also release or absorb energy into its surroundings based on the change in internal energy, changing the entropy of the surroundings as well. To capture both of these changes in entropy, the chemical system must be allowed to exchange energy with its surroundings. Instead of through entropy, the second law can be more easily analysed through the change in Gibbs free energy, which is effectively the overall entropy change of both the system and its coupled surroundings, which are at fixed temperature and pressure. This Gibbs free energy is a thermodynamic potential related to the internal energy U , temperature T , entropy S , pressure P , and volume V of a system as follows:

$$G = U - TS + PV = \sum_i \mu_i dn_i, \quad (1.46)$$

where μ_i is the Gibbs free energy per particle of species i with number densities n_i . The Gibbs free energy can be understood as the maximum amount of energy that can be extracted to do work in a closed system of fixed temperature and pressure.

By invoking the first law of thermodynamics that energy is conserved $dU = TdS - PdV$, and the second law of thermodynamics $dS \geq dQ/T$, such that the entropy S can only increase, we can state the Gibbs free energy in differential form:

$$dG \leq VdP - SdT \quad (1.47)$$

Therefore we can see that for a change at constant temperature ($dT = 0$) and pressure ($dP = 0$), the Gibbs free energy can only decrease. Therefore, any sequence of chemical reactions $dG < 0$ will result in the Gibbs free energy being minimized. The dimensionless Gibbs free energy of a system with N total particles

can be written:

$$\frac{G}{RT} = \sum_i n_i \left[\frac{\mu_i}{RT} + \ln P + \ln \frac{n_i}{N} \right]. \quad (1.48)$$

To solve for an atmosphere in thermodynamic equilibrium, you therefore minimize Eq. 1.48 with respect to the additional constraint that the number of atoms of each element are conserved. This is done by treating each atmospheric layer as an isolated system at fixed temperature and pressure.

The dependence of chemical equilibrium on temperature and pressure can be understood by considering reversible atmospheric reactions, with a forwards and backwards reaction. The reactant and product species have different potential energies stored in their bonds, and by undergoing the forward reaction, this energy can be absorbed from (endothermic) or released into the environment (exothermic), with the opposite occurring in the backwards reaction. At low temperatures, the more entropic outcome is to release energy from the bonds, favouring the exothermic reaction, with the reverse occurring at high temperatures. This is because the overall entropy increase of the surroundings for a given release of energy is greater at low temperatures than at high temperatures, understood using the $\propto T^{-1}$ dependency of the following:

$$\Delta S_{\text{surr}} = -\frac{\Delta H}{T}, \quad (1.49)$$

where ΔH is the enthalpy change of the chemical reaction (effectively the internal energy change of the reaction, with an additional term describing the work done against constant pressure from any change in volume, $\Delta H = \Delta U + P\Delta V$).

There is a similar dependence on pressure (within a fixed volume): if the number of molecules on either side of the equation is uneven, the reaction forming more gaseous molecules in total will be increasingly favoured at low pressures, while the reverse reaction will be favoured at high pressures. These changes can be understood through changes in the value of the equilibrium constant, K_{eq} , which is equal to the rate constant of the forward reaction k_f over the rate constant of the backwards reaction k_b , using the equilibrium condition that the forward and backwards reaction are identical. From the definition of the rate constants, the equilibrium constant can also be related to the partial pressures of the products $[B_j]$ and reactants $[A_i]$

$$K_{\text{eq}} = \frac{k_f}{k_b} = \frac{\prod_j [B_j]}{\prod_i [A_i]} = e^{-\frac{\Delta G}{RT}}. \quad (1.50)$$


The equilibrium constant itself is a function of the change in Gibbs free energy of the reaction ΔG (e.g. see Blundell & Blundell, 2006, chapter 22). The dependence on pressure and temperature is all built into this expression.

The tendency of equilibrium to resist changes in conditions is Le Chatelier's principle, and is used to optimize industrial reactions on Earth. These industrial

gas-phase reactions are often identical to the key reactions of exoplanet atmospheres; for example the Haber process for ammonia production is identical to the following expression of dominant nitrogen species in equilibrium:




Le Chatelier’s principle can be used to understand the above reaction, as well as the reaction determining the dominant carrier of carbon:




which are both exothermic in the forward reaction. At high temperatures, CO and N₂ are the dominant carriers of the carbon and nitrogen respectively, while at low temperatures (like in the Solar System gas giants), CH₄ and NH₃ dominate (e.g. Zahnle & Marley, 2014; Fortney et al., 2020). The position of equilibrium is also pressure sensitive however, so there can be different dominant species in equilibrium in different parts of the atmosphere (also considering temperature is not constant), which is relevant when discussing the impact of transport processes.

The equilibrium also depends on the elemental ratios present in the atmosphere. The elemental ratios of species apart from hydrogen and helium are often parametrized as scaled solar values (e.g. Asplund et al., 2009), referred to using metallicity. Throughout this thesis I use [M/H] to define the ratios of M to hydrogen ratios on a log₁₀ scale normalized to solar, where M is all non-hydrogen or helium atoms, to parametrize the metallicity. At higher metallicities, the equilibrium is pushed towards N₂ in Eq. 1.51, and also towards the production of CO₂ (Lodders & Fegley Jr, 2002; Zahnle et al., 2009) through the following reaction:



At high carbon to oxygen ratios (>1), the equilibrium can also be pushed towards the production of HCN and other carbon-rich species (Madhusudhan, 2012; Mollière et al., 2015), for instance through the following reaction:



Condensation processes (the phase transition from gas to solid or liquid phase) are also governed by chemical equilibrium, with sublimation and vaporization being reversible endothermic reactions which increase the number of gas molecules. The composition of the Solar System gas giant atmospheres is primarily driven by equilibrium chemistry at low temperature, with CH₄ as the dominant carrier of car-

bon, and oxygen and nitrogen sequestered from the upper atmosphere due to the condensation of NH_3 and H_2O into ices (e.g., Lewis, 1969; Atreya & Romani, 1985).

Disequilibrium chemistry and quenching

A key assumption in simulating an atmosphere in chemical equilibrium is that each layer can be treated as an isolated system. However, layers of an atmosphere are not isolated systems, for two key reasons. Firstly, material can be transported between atmospheric layers (in a 1D system, this is vertical mixing). Secondly, planetary atmospheres are irradiated, meaning there is a constant flux of high-energy photons which can chemically interact with gases in the atmospheres in photochemical reactions.

These processes are both time-dependent (transport) or irreversible (photochemistry), and therefore cannot be solved using thermodynamics. Instead, let us consider where the assumption of equilibrium chemistry can hold. If chemical reactions are far more rapid than advection processes responsible for moving material, chemical equilibrium will be re-established before the material can travel any notable distance (a characteristic dynamical timescale, t_{dyn}). A condition for the assumption of chemical equilibrium to be valid can be stated thus:

$$t_{\text{chem}} < t_{\text{dyn}}. \quad (1.55)$$

The rate of reaction in a gas (and hence the characteristic chemical timescale, t_{chem}) is strongly dependent on temperature and pressure. Therefore, in the hot, high-pressure deep atmosphere of a planet, t_{chem} is very small, meaning the assumption of equilibrium chemistry typically holds. Equilibrium chemistry is consequently also referred to as thermochemistry.

In the upper atmosphere, where temperatures and pressures are lower, this inequality inevitably reverses. In this case, material can travel a long distance before being transformed by chemical reactions. Therefore, the chemical composition can be held as approximately constant in the spatial direction (altitude, in a model 1D atmosphere).

To model this behaviour, we can consider the concept of quenching (e.g. Visscher & Moses, 2011). The quench pressure, P_{quench} , is the point at which Eq. 1.55 becomes an equality. Below the quench pressure, the chemical composition is set by chemical equilibrium. At altitudes above the quench pressure, the composition is “frozen in” (hence the term quenching) and held vertically constant. This reduces the impact of dynamics in 1D on chemistry, to a single free parameter P_{quench} , which is primarily a function of the dynamical timescale, t_{dyn} – see Fig. 1.12. As each net reaction that can be considered (e.g. $\text{CO}-\text{CH}_4$) has its own chemical timescale, t_{chem} , there should be theoretically a different quench point for each chemical transition. In this case, the ratio between molecules included in that quenched reaction should remain constant above P_{quench} .




Figure 1.12: Vertical abundance ratios of CO in the brown dwarf Gliese 229B assuming equilibrium chemistry (dashed line) and vertical transport with different K_{zz} values (solid lines), demonstrating the abundance of CO being held constant above the quench pressure, which varies with K_{zz} . The drop in abundance at low pressures arises as a result of molecular diffusion dominating over eddy diffusion. Figure reproduced from Visscher & Moses (2011).

The quench pressure can be crudely estimated analytically (e.g. Madhusudhan & Seager, 2011) by directly finding where the chemical and dynamical timescales are equal, using the following expression for the chemical timescale:

$$t_{\text{chem}} = n \left(\frac{dn}{dt} \right)^{-1}, \quad (1.56)$$

where $\frac{dn}{dt}$ is the rate-limiting step (the rate of reaction of species n); and the following expression for the dynamical timescale:

$$t_{\text{dyn}} = \frac{L^2}{K_{zz}}, \quad (1.57)$$

where L is some characteristic length scale, generally taken as the scale height of the atmosphere, or some fraction thereof (varying by up to an order of magnitude). Given this uncertainty in L , the fact K_{zz} is also poorly constrained by multiple orders of magnitude, and difficulty in identifying the rate-limiting step, attempting to use this approximation *a priori* is often not effective (e.g. Tsai et al., 2017). Zahnle & Marley (2014) use chemical kinetic models to determine chemical timescales as an emergent property of the chemical network.

Proper consideration of vertical mixing as well as the internal heat profile is important to understand where key transitions (e.g. CO–CH₄ and N₂–NH₃) occur as a function of planet mass and vertical mixing. Thorough considerations of the dependence of these transitions on equilibrium temperature and planet mass for the strongly- and moderately-irradiated giant planets that are the focus of this thesis are presented in Thorngren et al. (2019) and Fortney et al. (2020) respectively. For example, carbon monoxide is a disequilibrium species found in the atmosphere of Jupiter due to vertical mixing from the high-temperature deep atmosphere (Prinn & Barshay, 1977).

Chemical kinetics

To truly simulate the chemistry of an atmosphere while incorporating time-dependent and irreversible processes, we must evolve the composition of an atmosphere in time considering the creation and destruction of species by chemical reactions, as well as the transport of species through the atmosphere. This is chemical kinetics, and effectively involves solving the mass continuity equation in each layer of an atmosphere, presented here in 1D

$$\frac{\partial n_i}{\partial t} = P_i - L_i - \frac{\partial \Phi_i}{\partial z}, \quad (1.58)$$

where P_i and L_i are the production and loss terms of species n_i , and Φ_i is the transport flux defined in Eq. 1.44. This is the equation solved by the photochemistry simulation codes of Rimmer & Helling (2016) and Tsai et al. (2017), amongst others.

The production and loss terms are defined in terms of some chemical network, which defines all the possible chemical reactions that can take place, with rate

constants k such that the production and loss terms can be written as the following (here simplified) sum of reactions:

$$P_i = \sum_j k_j n_j + \sum_{jk} k_{jk} n_j n_k + \sum_{jkl} k_{jkl} n_j n_k n_l \quad (1.59)$$

$$L_i = \sum_i k_i n_i + \sum_{ij} k_{ij} n_i n_j + \sum_{ijk} k_{ijk} n_i n_j n_k \quad (1.60)$$

including the following terms: left – interactions between particles and photons (thereby including photochemistry) – central – collisions between two particles – and right – three-body interactions involving two reactant particles and a generic third particle. Kinetics codes also permit the addition of source and sink terms into the production and loss equations respectively to model interactions with a non-simulated chemical inventory. Constructing an accurate chemical network requires including all relevant reactions in the network, and using accurate rate constants based on experimental values or theoretical calculations. Typically a chemical kinetics simulation is run until a steady-state is reached, but it can also be used to deduce time-dependent effects like varying stellar flux from a flare or eccentric orbit.

Chemical kinetics gives the potential to simulate photochemistry, which can play a considerable role in planetary atmospheres. The terrestrial planet atmospheres of the Solar System are all dramatically shaped by chemical kinetics and photochemistry, without the constant exchange with a large reservoir of high temperature and pressure gas in equilibrium that maintains the atmospheric composition of gas giants – for instance the SO_2 of Venus is volcanic in origin (Esposito, 1984), while photolysis is likely responsible for converting primordial NH_3 to atmospheric N_2 on Titan (Atreya et al., 1978). Even in gas giants, photochemistry can have a notable impact, both in exoplanets (e.g. forming SO_2 in high metallicity hot Jupiters, Tsai et al., 2023), and the Solar System (e.g. producing complex hydrocarbons in the upper atmosphere of Jupiter, Gladstone et al., 1996).

1.2.4 Aerosols and condensation

Aerosols, solid and liquid particles suspended in the atmosphere, also contribute significantly to atmospheric opacities, and therefore energy budgets and observed spectra. Aerosols appear in every atmosphere of the Solar System, including the water clouds of Earth, the sulfuric acid clouds of Venus (Hansen & Hovenier, 1974), the ammonia and ammonium hydrosulfide (NH_4HS) clouds of Jupiter and Saturn (Lewis, 1969), and the tholin hazes of Titan (Brown et al., 2002). Condensation can also induce gradients in mean molecular weight. On Earth, where the mean molecular weight of the atmosphere is greater than that of H_2O , this serves to reduce stability with respect to convection, which is responsible for the turbulent dynamics of the troposphere. On giant planets, where the mean molecular weight is lower than that of condensible species, this instead can inhibit convection (Leconte et al.,




Figure 1.13: Condensation curves of an array of cloud forming species, presented alongside example pressure-temperature profiles for the exoplanets HD 209458b and HR 8799b (Saumon & Marley, 2008), and the Solar System planets Jupiter and Neptune (Moses & Poppe, 2017), as well as the transition curves of CH_4 –CO and NH_3 – N_2 (Lodders & Fegley Jr, 2002), all presented for solar metallicity gas. Condensation curves are taken from Helling et al. (2001); Lodders & Fegley Jr (2002); Visscher et al. (2006, 2010); Morley et al. (2012); Wakeford et al. (2016). Figure reproduced from Gao et al. (2021).

2017; Innes et al., 2023).

Atmospheric aerosols are frequently split into clouds and hazes, although the terminology for this is not always consistent. In this thesis, I follow the definitions of cloud and haze presented by Gao et al. (2021), which are distinguished based on their formation mechanism. In brief, clouds are particles formed through thermochemical reactions and phase changes, while hazes are the product of irreversible energetic bombardment processes, principally photochemistry. Clouds tend to be composed of larger particles than hazes, however this is not diagnostic in this definition. Nonetheless, when parametrizing aerosol opacities for interpreting observations, it is common to refer to small particle opacity (i.e. Rayleigh scattering, see Sec. 1.2.1) with the term “haze” (e.g. Pont et al., 2008), see Sec. 2.4.1 for more details. Gao et al. (2021) provide a thorough overview of the aerosols relevant to exoplanet atmospheres.

As touched on in the definition, clouds are formed through equilibrium reactions and phase changes, often collectively referred to as condensation (although strictly this should only refer to the latter). In hot Jupiters, this is of species typically associated with minerals on Earth, principally silicates like MgSiO_3 , Mg_2SiO_4 , and SiO_2 (Gao et al., 2020; Grant et al., 2023b), but also iron (Visscher et al., 2010), salt species including KCl and ZnS (Morley et al., 2012), and corundum (Al_2O_3 Wakeford et al., 2016). Lower temperature cloud species of potential importance for giant exoplanets include NH_3 and NH_4HS (Lodders & Fegley Jr, 2002), H_2O (Mang et al., 2022), and S_8 (Gao et al., 2017).

Cloud formation can be simply understood in terms of saturation curves.

The saturation curve is determined through Gibbs energy minimization, as with any equilibrium process, and defines a partial pressure at each point in the atmosphere above which a gaseous species will condense into a solid or liquid. Example condensation/saturation curves of cloud species are displayed against planetary $T(P)$ profiles in Fig. 1.13. Expressions for the saturation pressure P_V are defined in terms of temperature, for instance the following formulation based on the Clausius-Clapeyron equation from Sánchez-Lavega et al. (2004):

$$\ln P_V = \ln C + \frac{1}{R} \left[-\frac{L_0}{T} + \Delta\alpha \ln T + \frac{\Delta\beta}{2} T \right], \quad (1.61)$$

where C and L_0 are integration constants, $\Delta\alpha$ and $\Delta\beta$ describe the change in the heat-capacity at constant pressure C_p between the vapour and solid phase using the approximation $C_p = \alpha + \beta T$, and R is the specific gas constant. Condensation can also be incorporated into chemical kinetic models by including condensation reactions.

Determining the size, shape, and location of cloud particles once they have formed however is a more complicated endeavour. The analytical prescription of Ackerman & Marley (2001) uses the balance of sedimentation and vertical mixing, parametrized using f_{sed} as the ratio of mass-weighted sedimentation velocity to convective velocity, to derive an exponential decline in cloud species mixing ratio above a cloud base. The geometric mean radius of the particles can also be found from this approach, by further prescribing a log-normal distribution, with parametrized geometric standard deviation σ_g , for the particle radii. A more rigorous treatment of clouds requires a complex microphysical simulation of clouds (e.g. Turco et al., 1979; Toon et al., 1979; Ohno & Okuzumi, 2017). Some of this complexity can be captured by using a variable f_{sed} , tuned to match the results of microphysical models.

Photochemical hazes are predicted to be important sources of opacity for cooler hot Jupiters below 950 K, particularly where CH_4 is abundant to form hydrocarbon hazes (Gao et al., 2020). Hazes are typically simulated using microphysical models, with the production of haze being an irreversible process that must be either parametrized using some haze production flux, or coupled to the abundance of some haze precursor abundance from a chemical kinetics model (e.g. Kawashima & Ikoma, 2018, 2019). Hazes can also possess distinct spectral absorption features (e.g. Khare et al., 1984b; Corrales et al., 2023; He et al., 2024), and serve as nucleation points for condensation processes (heterogeneous nucleation, e.g. Mang et al., 2022).

1.3 Exoplanet Atmosphere Characterization

In comparison to its host star, an exoplanet is intrinsically both small and faint. As briefly discussed in Sec. 1.1, direct imaging of exoplanets requires very distant orbits and high intrinsic luminosity. While this does permit atmospheric characterization, the suitable targets represent only a very minor fraction of the exoplanet population. Fortunately, there are high-precision techniques that take advantage of




Figure 1.14: Simple schematic demonstrating atmospheric characterization methods for transiting exoplanets at different orbital phases, with the transmission spectrum observable during primary transit, and the dayside emission and reflection spectrum identifiable based on the in- and out-of-secondary elclipse contrast. Measuring the phase curve requires continuous observation of an entire orbit, while nightside emission can be disentangled through phase curves or combined transmission and emission measurements. Figure reproduced from Gao et al. (2021).

transiting exoplanets to observe exoplanet atmospheres. This includes measuring the transmission of starlight through the atmosphere during transit – transmission spectroscopy – as well as detection of planetary emission during secondary eclipse – emission spectroscopy. My thesis focusses on using the former technique at low spectral resolution, with resolving powers of $R = 100\text{--}400$, where R is given by the smallest difference in wavelength $\Delta\lambda$ that can be distinguished as follows:

$$R = \frac{\lambda}{\Delta\lambda}. \quad (1.62)$$

I provide an overview of the method and key results from transmission spectroscopy observations at different wavelengths in Sec. 1.3.1. I also provide a brief discussion of low resolution emission spectroscopy in Sec. 1.3.2 and high resolution spectroscopy techniques in Sec. 1.3.3, and finally touch on some of the findings of phase curves in Sec. 1.3.4 and direct imaging in Sec. 1.3.5.

1.3.1 Low resolution transmission spectroscopy

Transmission spectroscopy is effectively the transit method of exoplanet detection to measure the transit depth of the planet (see Sec. 1.1.1), but performed simultaneously at multiple wavelength using spectroscopy. As exoplanet atmospheres have strongly wavelength dependent opacity (see Sec. 1.2.1), the apparent size of the atmosphere (and hence planet) should also be wavelength dependent.

The impact of a planetary atmosphere on the transit depth of an exoplanet can be illustrated by considering the starlight blocked by a non-emitting opaque circular planet (the ratio of stellar flux during transit divided by the flux out of transit) with radius R_p plus an atmosphere of wavelength-dependent height $z(\lambda)$ transiting in front of a star of radius R_* :

$$\frac{F_{\text{in}}}{F_{\text{out}}}(\lambda) = 1 - \left(\frac{R_p + z(\lambda)}{R_*} \right)^2 \simeq 1 - \frac{R_p^2}{R_*^2} - \frac{2R_p}{R_*^2} z(\lambda), \quad (1.63)$$

where using for a typical atmosphere $z \ll R_p$, allowing us to ignore terms quadratic in z . Finding the wavelength dependence of the in- and out-of-transit flux ratio can therefore be used to determine the wavelength dependence of the atmospheric height $z(\lambda)$. This formulation of the flux ratio in transmission spectroscopy assumes no emission from the planet's nightside (which can be non-negligible in many cases, e.g., Lustig-Yaeger et al., 2025; MacDonald et al., 2026), and a homogeneous stellar surface, which is frequently false, resulting in contamination of the transmission spectrum by the star (e.g. Rackham et al., 2018) These effects are discussed in Sec. 1.3.2 and later this section respectively.

Transmission spectroscopy can be performed by directly evaluating the transmission of flux as described in Eq. 1.63. However, rather than simply sampling the flux values in- and out-of-transit, the conventional method is to fit the whole light-curve with a transit model in each wavelength bin to find the transit depth as a function of wavelength (see Sec. 1.1.1), due to the dependence of the transit shape on limb-darkening (which is also wavelength dependent).

The height of the atmosphere, and therefore the atmospheric signal, is in turn proportional to the scale height of the planet (Eq. 1.34), allowing us to identify the dependence of the wavelength-dependent part of the transit depth $\Delta_{\text{transit}}(\lambda)$ as

$$\Delta_{\text{transit}}(\lambda) \propto \frac{R_p T}{\mu_m g R_*^2}. \quad (1.64)$$

This proportionality explains why the best targets for transmission spectroscopy are high temperature, weak gravity (low density), large planetary radius, and small stellar radius. Most notably this includes hot Jupiters, which have high temperature and low density, and planets orbiting small stars, particularly M dwarfs. Another key factor is mean molecular weight, which is ~ 2.3 for H_2 and He dominated atmospheres, but can be much higher for terrestrial planet atmospheres (for example, a pure CO_2 atmosphere has $\mu_m = 44$). The ongoing challenge of transmission spectroscopy is generally to push towards smaller and cooler planets. Other relevant factors include stellar magnitude (brighter stars increase the signal-to-noise of the time-series stellar spectra and therefore the transmission spectrum), as well as scheduling concerns like transit frequency, and for ground-based observations, transit duration, which can exceed the duration of a night for some longer period planets.

Transmitted flux

A proper description of the transmitted flux in Eq. 1.63 can be solved by considering the radiative transfer equation Eq. 1.24 in the case of no source function $S = 0$, as the emission of the planet is considered to be negligible compared to the stellar flux:

$$\mu \frac{dI(z, \mu, \nu)}{d\tau_\nu} = I(\tau_\nu, \mu, \nu). \quad (1.65)$$

This can be simply integrated from an optical depth τ_ν , where the intensity is $I(0, \mu, \nu)$ (see Eq. 1.22 for the definition of τ_ν) to a transparent edge of $\tau_\nu = 0$, where $I(\tau_\nu, \mu, \nu)$ is measured:

$$[\ln I]_{I(0)}^{I(\tau_\nu)} = [\tau_\nu / \mu]_{\tau_\nu}^0 \quad (1.66)$$

which gives:

$$I(\tau_\nu, \mu, \nu) = I(0, \mu, \nu) e^{-\tau_\nu / \mu}. \quad (1.67)$$

This is the Beer-Lambert law. In the case of transmission spectroscopy (i.e. observing at an effectively infinite distance), only the direct ray $\mu = 1$ is relevant. We can improve the model by replacing an opaque atmosphere of height $z(\lambda)$ with a succession of atmospheric layers of height δz with fixed τ_λ ,

$$z(\lambda) = \sum_i \delta z_i (1 - e^{-\tau_{\lambda,i}}), \quad (1.68)$$

with the full expression for the transmission from summing all of these layers. In a given layer, at wavelengths where the atmosphere is fully opaque $e^{-\tau_{\lambda,i}} = 0$, and at wavelengths where the atmosphere is fully transparent $e^{-\tau_{\lambda,i}} = 1$.

As discussed in Sec. 1.2.1, different opacity sources dominate at different wavelengths, while different spectral resolutions are sensitive to some forms of opacity over others. I discuss the key findings from low resolution transmission spectroscopy observations for the remainder of the section.

Measuring chemistry

Low resolution observations are performed at a resolution of $R < 5000$, giving limited sensitivity to individual spectral lines. The actual resolution used is frequently much lower, only giving sensitivity to the broad molecular bands produced by vibrational transitions and continuum opacity sources. The exception to this is the alkali metals sodium and potassium, where the line core and the shape of the wings can be probed.

The first result from low resolution transmission spectroscopy, and exoplanet atmospheric characterization in general, was the detection of sodium by Charbonneau et al. (2002) in the hot Jupiter HD 209458b. Sodium has since been detected in numerous hot Jupiters (e.g., Sing et al., 2012; Nikolov et al., 2016; Fischer et al., 2016; Nikolov et al., 2018; Alderson et al., 2020; Alam et al., 2021; Ahrer et al., 2022), as has potassium (e.g., Sing et al., 2011b, 2015; Feinstein et al., 2023). These observations were all performed at relatively short wavelengths ($0.3 < \lambda < 1 \mu\text{m}$, hereafter “optical”). Also visible in low resolution at these wavelengths are titanium oxide (TiO) and vanadium oxide (VO), with evidence found in ultra-hot Jupiter atmospheres (Sedaghati et al., 2017; Chen et al., 2021; Kirk et al., 2021).

Near-infrared transmission spectroscopy at low spectral resolution is primarily performed using space-based telescopes, given the considerable infrared opacity of the Earth’s atmosphere. The Hubble Space Telescope (HST) is capable of providing transmission spectroscopy from 1.1–1.7 μm using the Wide Field Camera 3 (WFC3) instrument. This wavelength range is sensitive to absorption from water vapour, with detections in hot Jupiters (e.g., Deming et al., 2013; Kreidberg et al., 2015; Evans et al., 2016; Wakeford et al., 2017a) and Neptunes (Fraine et al., 2014; Wakeford et al., 2017b). This narrow region is degenerate between H_2O and CH_4 as the molecules have an overlapping absorption feature. A feature originally attributed to H_2O in the temperate sub-Neptune K2-18b (Benneke et al., 2019; Tsiaras et al., 2019) was later identified as fully or partly arising from CH_4 (Madhusudhan et al., 2023).

The increased wavelength coverage, resolution, and precision of the James Webb Space Telescope (JWST) has dramatically increased the sensitivity of low resolution transmission spectroscopy to molecular absorption (e.g. Fisher et al., 2024). Carbon dioxide was detected in Early Release Science (ERS) observations of WASP-39b (JWST Transiting Exoplanet Community Early Release Science Team et al., 2023; Alderson et al., 2023; Rustamkulov et al., 2023; Carter et al., 2024), as well as in other hot Jupiters (e.g., Fu et al., 2024; Xue et al., 2024) as well as sub-Neptunes (Benneke et al., 2024). Sulfur dioxide (SO_2) was also identified in WASP-39b (Alderson et al., 2023; Rustamkulov et al., 2023; Powell et al., 2024), as well as the Neptune-mass planets WASP-107b (Dyrek et al., 2024) and HAT-P-26b (Gressier et al., 2025). Methane has now also been discovered in transiting giant exoplanets (Bell et al., 2023), as has hydrogen sulfide (Fu et al., 2024) and silicon monoxide (Gapp et al., 2025; Evans-Soma et al., 2025). A selection of hot Jupiter transmission spectra are presented in Fig. 1.15.

Cross-correlation techniques (see Sec.1.3.3) on medium resolution JWST data have also been used to detect carbon monoxide (Grant et al., 2023a; Esparza-Borges et al., 2023), which is hard to strongly constrain from low resolution data alone. Observations of sub-Neptune exoplanets have also exhibited some evidence of carbon disulfide (CS_2 , Benneke et al., 2024; Constantinou et al., 2026) and more complex aliphatic compounds (Madhusudhan et al., 2023, 2025; Stevenson et al., 2025; Constantinou et al., 2026), hinting at the rich chemistry present in these cooler, higher metallicity atmospheres (with mean molecular weights of up to ~ 5).

Medium resolution transmission spectroscopy using the Lyman- α line in the ultraviolet and metastable helium lines in the near infrared are used to detect very deep transits attributed to giant hydrogen-helium exospheres attributed to evaporation of the atmosphere (e.g., Bourrier et al., 2016; Spake et al., 2018; Nortmann et al., 2018; Kirk et al., 2020; Vissapragada et al., 2020; Orell-Miquel et al., 2024; Krishnamurthy et al., 2025). This has demonstrated that these highly irradiated planets with gaseous envelopes are frequently subject to mass loss, and has been used to suggest that the lower bound radius of the hot Jupiter population is shaped




Figure 1.15: Near-infrared transmission spectra of hot Jupiters observed with JWST, with highlighted bands covering spectral features of commonly observed molecules (the L band primarily overlaps with H_2O absorption), as well as the absorption spectra of common molecules. Figure reproduced from Fu et al. (2025a).

by planet formation and not photoevaporation (Vissapragada et al., 2022).

The spectral features of atoms and molecules can be used to deduce atmospheric abundances and temperatures, and thereby place constraints on elemental ratios, often through the assumption of chemical equilibrium. By assuming the elemental ratios of the atmosphere are reflective of the bulk planetary composition, the chemical abundances from atmospheric observations can in principle be used to understand formation and migration processes of the planets (e.g., Öberg et al., 2011; Helling et al., 2014; Madhusudhan et al., 2014; Booth et al., 2017; Espinoza et al., 2017), as well as constrain the mass-metallicity relationship of planets (e.g., Barstow et al., 2017; Welbanks et al., 2019). CO_2 and SO_2 both serve as excellent indicators of metallicity. JWST transmission spectra of giant exoplanets have exhibited a diverse range of metallicities without a clear correlation with mass (e.g., Bean et al., 2023; Kanodia et al., 2026), and also show great diversity in C/O ratio, from $0.11^{+0.12}_{-0.06}$ (Xue et al., 2024) to $0.92^{+0.02}_{-0.03}$ (Evans-Soma et al., 2025).

The detection of some atmospheric species out of equilibrium can provide direct evidence of chemical processes including photochemistry, like the detection of SO_2 in WASP-39b (Tsai et al., 2023). The apparent absence of CH_4 that would be expected in cooler hot Jupiters could be the result of quenching from hotter-than-anticipated interiors (Yu et al., 2026).

Measuring aerosols

The size and location of aerosol species can be constrained by low resolution observations at optical wavelengths through the identification of scattering slopes and the muting of the wings or cores of sodium and potassium absorption. Strong Rayleigh scattering slopes in optical spectra indicate the presence of small aerosol particles (e.g., Lecavelier des Etangs et al., 2008a; Sing et al., 2015; Kirk et al., 2017). The broadly λ^{-4} dependence of Rayleigh scattering can be used to measure the atmospheric scale height H , and therefore the temperature of the atmosphere, based on the variation in transit depth R_p at short wavelength (Lecavelier des Etangs et al., 2008a,b)

$$H = -4 \frac{d \ln R_p}{d \ln \lambda}. \quad (1.69)$$

Rayleigh scattering can deviate from λ^{-4} at short wavelengths, especially into the UV. Transmission spectroscopy observations have also yielded scattering slopes with an even steeper wavelength dependence, known as “super-Rayleigh” scattering slopes (e.g., Lecavelier des Etangs et al., 2008a; Pont et al., 2013; Palle et al., 2017; Alderson et al., 2020; Luque et al., 2020; Ahrer et al., 2022). Potential causes for these scattering slopes include photochemical hazes (Ohno & Kawashima, 2020) or combinations of cloud and haze layers (Pont et al., 2013; Pinhas et al., 2019). It could be instead a result of stellar contamination, discussed later in this section (McCullough et al., 2014; Espinoza et al., 2019).




Figure 1.16: Optical and infrared transmission spectra of hot Jupiters observed with HST and Spitzer with the best fitting models, demonstrating the variance in scattering slope and feature strength interpreted as differences in aerosol properties. Figure reproduced from Sing et al. (2016).

Flat, featureless optical spectra in giant planets are attributed to larger aerosol particles (e.g., Kreidberg et al., 2014a; Kirk et al., 2016; Lendl et al., 2016; Louden et al., 2017; Spyros et al., 2021; Ahrer et al., 2023a). In small planets, a flat spectrum may instead be interpreted as a suppression of the scale height through a high mean molecular weight atmosphere, or the complete absence of an atmosphere. A relative lack of aerosols in the atmosphere can be inferred from the detection of a pressure-broadened sodium feature (Seager & Sasselov, 2000), as seen in WASP-39b, WASP-96b, and WASP-62b (Fischer et al., 2016; Nikolov et al., 2018; Alam et al., 2021). Optical transmission spectroscopy is a highly complementary tool to infrared spectroscopy for reducing degeneracies in chemical composition and distinguishing between clear, hazy, and cloudy atmospheres (Fairman et al., 2024). The diversity of optical transmission spectra attributed to aerosols is demonstrated in Fig. 1.16.

The mid infrared spectral coverage of JWST has also permitted characterization of cloud composition, including the detection of SiO_2 (Grant et al., 2023b) and silicate clouds (Dyrek et al., 2024), based on the vibrational absorption modes of cloud species (Wakeford & Sing, 2015). At near infrared wavelengths, the main role of aerosols is simply to mute spectral features (e.g. Gao et al., 2021). Observations with both HST and JWST have demonstrated that clouds are extremely common in hot Jupiter atmospheres.

Stellar inhomogeneity

A common assumption in transmission spectroscopy is that the star is a steady and radially homogeneous source of flux. Sadly (for exoplanet astronomers at least), this is not universally true. Hotter and cooler regions on the stellar disc, including star spots and faculae, are common sources of stellar heterogeneity, particularly in active stars. These features can cause measurable features in transmission spectra, particularly apparent at optical wavelengths where the stellar flux is most variable, and can be caused by both occulted (which are also apparent in the light-curve) and unocculted regions (which do not appear in the light-curve). This is the transit light source effect (Rackham et al., 2018, 2019). Attempts can be made to account for this using spectral models, by letting a fraction of the star possess a spectrum from a stellar model of a different temperature (e.g., Pinhas et al., 2018; Welbanks & Madhusudhan, 2021; MacDonald & Lewis, 2022). This is limited however by the accuracy of stellar models, particularly for M dwarfs (e.g. Witzke et al., 2022; Moran et al., 2023), and the strong degeneracies introduced by this parametrization (e.g., Wakeford et al., 2019; Barclay et al., 2021). This degeneracy is weaker in the optical, making short wavelength observations useful for constraining stellar activity. An additional concern that can be introduced when considering the coolest stars is the presence of molecular opacity like H_2O in star spots, making it very difficult to discern the source of the absorption (Moran et al., 2023), particularly at purely infrared wavelengths.

Stellar contamination is the greatest challenge to transmission spectroscopy

in the age of JWST. Fortunately none of the observations analysed in this thesis have been impacted by stellar contamination, so I will stop at this brief summary. For a more thorough overview of the challenges presented by stellar inhomogeneity, consider Rackham & de Wit (2024).

1.3.2 Low resolution eclipse spectroscopy

By observing an exoplanet host star during the secondary eclipse of its planet, we obtain the stellar spectrum without the planetary emission. Therefore by comparing the in- and out-of-eclipse measurements, we can identify the component of the combined stellar plus planetary spectrum arising from the planet’s dayside. This eclipse spectrum is in turn a combination of reflected stellar spectrum and the thermal emission of the planet. The former naturally peaks at short wavelengths, depending on the temperature of the star, while the latter peaks in the infrared, depending on the dayside temperature T_{day} of the planet according to Wien’s displacement law:

$$\lambda_{\text{peak}} = \frac{b}{T_{\text{day}}} \quad (1.70)$$

where b is Wien’s displacement constant.

In eclipse spectroscopy the wavelength-dependent quantity of interest that can be measured, equivalent to the transit depth for transmission spectroscopy, is the secondary eclipse depth, which is equal to the planet-star flux ratio. This can be measured at each wavelength by fitting an eclipse model similar to spectroscopic light-curves, as with transmission spectroscopy.

Because the eclipse spectrum measures the planet to star flux ratio, rather than the spatial scale of the planetary atmosphere as in transmission spectroscopy, it is affected by different observational biases. The thermal secondary eclipse depth δ_{eclipse} is simply proportional to

$$\delta_{\text{eclipse}}(\lambda) \propto \frac{B_{\lambda}(T_{\text{day}})}{B_{\lambda}(T_{*})} \left(\frac{R_P}{R_{*}} \right)^2, \quad (1.71)$$

where B is the Planck function describing black body radiation. This means that emission spectroscopy is instead strongly dependent on temperature (via the Planck function), but does not depend on gravity or mean molecular weight. This makes emission spectroscopy much better suited for the characterization of dense (high gravity) planets compared to transmission spectroscopy, as well as small planets that are expected to possess a high mean molecular weight atmosphere. On the other hand, emission spectroscopy is less suitable for cooler planets, even when observing at very long wavelengths to capture the peak of the planet’s thermal emission. The reflected light spectrum instead depends on the planet-star radius ratio and the albedo. Like for transmission spectroscopy, other observational constraints include the brightness of the star (which affects the noise of the stellar spectra), and the eclipse frequency and duration.

As the atmosphere is observed head-on in emission spectroscopy, rather than

at a highly-grazing angle, emission observations are sensitive to deeper parts of the atmosphere than in transmission spectroscopy. A different longitude of the atmosphere is probed as well, exploring the dayside of the planet, while the transmission spectroscopy technique is only sensitive to the terminator regions, plus an opening angle around the terminator (e.g. Caldas et al., 2019). Combining strategically-observed transmission and emission spectroscopy observations allows subtraction of the pure stellar spectrum (measured during the secondary eclipse) from the combined planet-star spectrum measured by the out-of-transit transit observation, allowing evaluation of the infrared excess created by the nightside emission spectrum of the planet (Stevenson, 2020; Lustig-Yaeger et al., 2025).

Emitted flux

The thermal emitted flux of a planet, or emergent flux, can be calculated from the radiative transfer equation (Eq. 1.24). By assuming LTE and neglecting the impact of scattering, we can therefore write the source function $S = B$, and the radiative transfer equation becomes:

$$\mu \frac{dI(z, \mu, \nu)}{d\tau_\nu} = -I(\tau_\nu, \mu, \nu) + B(\tau_\nu, \nu) \quad (1.72)$$

which can be solved as a first-order ordinary differential equation and integrated over all angles μ to give (see Seager, 2010, chapters 5–6) the following expression for the emergent flux F_p

$$F_p(\lambda) = 2\pi \int_0^\infty B(\tau_\lambda, \lambda) e^{-\tau_\lambda} d\tau_\lambda. \quad (1.73)$$

From this expression, we can identify that the flux at a given wavelength is effectively an exponential-weighted sum of the blackbody radiation at each layer of the atmosphere. This can effectively be considered by treating the flux at a given wavelength as emerging from the altitude at which the atmosphere becomes optically thick $\tau_\lambda = 1$. The amount of flux that emerges therefore depends on the temperature at that optical depth. Where the opacity is weaker, the flux emerges from deeper in the atmosphere. If the deep atmosphere is warmer than the upper atmosphere, the flux in line cores will originate from a cooler part of the atmosphere and we therefore observe absorption features. If the upper atmosphere is warmer than the lower layers (a temperature inversion), we will instead observe emission features. An emission spectrum is therefore very sensitive to the temperature structure of the atmosphere. An isothermal atmosphere will simply yield a blackbody spectrum. This simplified picture is complicated by scattering, which can alter the observed brightness temperature and give spectral signatures even in an isothermal atmosphere (Taylor et al., 2021).




Figure 1.17: NIRSpec G395H emission spectrum of the hot Jupiter WASP-77 Ab observed as part of JWST programme GTO 1274. Figure reproduced from August et al. (2023).

Overview of existing results

Low resolution emission spectroscopy observations have detected a variety of molecules in exoplanet atmospheres from space-based data, including CO, H₂O and CO₂ (e.g., Kreidberg et al., 2014b; Fu et al., 2022; Bean et al., 2023; August et al., 2023), as well as potential evidence of opacity from TiO, VO, H⁻, and NH₃ (e.g., Haynes et al., 2015; Evans et al., 2017; Coulombe et al., 2023; Yang et al., 2024). Like with low resolution observations, these have been used to determine the C/O ratios and metallicities of various planets and assess the presence of clouds, but have also been used to interpret vertical mixing based on quenched abundances and the measured pressure-temperature profile. An example JWST emission spectrum is depicted in Fig. 1.17.

Emission spectroscopy can also be used to spatially resolve a flux map of the planet using high-cadence observations which probe different slices of the planet’s atmosphere over the course of ingress and egress, a technique known as eclipse mapping (Rauscher et al., 2007, 2018). Limited eclipse mapping was possible with Spitzer (Majeau et al., 2012; de Wit et al., 2012), however much more in depth eclipse mapping to reveal the spatial temperature structure of the dayside of hot Jupiters has been made possible with JWST, revealing non-uniform emission (Coulombe et al., 2023; Hammond et al., 2024; Valentine et al., 2024), typically interpreted as evidence of an equatorial jet predicted by GCMs (Showman & Guillot, 2002).

1.3.3 High resolution spectroscopy

Transmission and emission spectroscopy can both also be performed at very high spectral resolution. High resolution transmission spectroscopy operates quite differently to low resolution transmission spectroscopy. Rather than evaluating the flux ratio directly by fitting light-curves, high resolution transmission spectroscopy takes advantage of the unique spectral distribution of abundant line cores of an atmospheric absorber, which can probe the planetary transmission even when the signal to noise is incredibly poor at each individual wavelength. Cross-correlation techniques between the stellar spectrum and a template containing the molecular line lists are used to detect each molecule and identify the Doppler shift from the planet’s orbit. As this method probes the narrow cores of spectral lines, it is primarily sensitive to higher in the atmosphere compared to low resolution transmission spectroscopy, and it is unable to deduct continuum opacity sources. High resolution emission spectroscopy uses more-or-less the same technique, and is also unlike its low resolution counterpart, observing the dayside spectrum directly near eclipse, rather than observing the subtracted spectrum caused by eclipse.

Due to the immense bulk of high resolution spectrographs, these observations are restricted to ground-based telescopes. The spectrographs used are often the same used for radial velocity measurements of exoplanets (Sec. 1.1.2), as well as notable contributions from the CRyogenic InfaRed Echelle Spectrograph (CRIRES, Ulrich Käufl, 2008) on the Very Large Telescope (VLT) and the Immersion GRating INfrared Spectrometer (IGRINS, Park et al., 2014) on Gemini South. Sophisticated methods are required to remove telluric and stellar lines (e.g. Dash et al., 2024).

High resolution spectroscopy observations have been used to detect numerous species, including CO, H₂O, NH₃, HCN, CH₄, OH and CO₂ in the near infrared (e.g., Snellen et al., 2010; de Kok et al., 2013; Birkby et al., 2013; Brogi et al., 2016; Jacobbe et al., 2021), and TiO, Fe, Fe⁺⁺, Na, Ti, Cr, Mg, V, and Mn in the optical (e.g., Redfield et al., 2008; Snellen et al., 2008; Wyttenbach et al., 2015; Hoeijmakers et al., 2019; Ehrenreich et al., 2020; Prinoth et al., 2022; Gibson et al., 2022; Gandhi et al., 2023). The optical species except sodium are predominantly a feature of ultra-hot Jupiters, where temperatures are sufficient to vaporize these refractory species. Like with low resolution observations, abundances and atmospheric temperatures can be inferred based on these observations (Brogi & Line, 2019). High resolution observations can be used to infer dynamical properties of the atmosphere, including day-to-nightside winds and rotation (e.g., Snellen et al., 2010; Wyttenbach et al., 2015; Brogi et al., 2016). The sodium feature is particularly powerful for providing dynamical constraints, either using the shape of the sodium line profile, which is impacted by Doppler broadening (Seidel et al., 2020, 2023), or spatially-resolving the atmosphere and identifying a difference in redshift between the limbs (Louden & Wheatley, 2015). These dynamical constraints are consistent with the prediction that hot Jupiters are tidally locked (e.g. Brogi et al., 2016), and have shown evidence for a super-rotating equatorial jet and/or vertical winds (e.g., Louden &




Figure 1.18: Analysis of molecular absorption from the high resolution transmission spectrum of HD 189733b obtained using VLT/CRIRES, showing the model spectra (left column), and sensitivity maps for two detection analyses (middle and rights columns). The dark regions of the sensitivity maps correspond to regions of significance, in this case H_2O and CO_2 are detected at the expected radial velocity amplitude (K_p), while CO_2 is non-detected and the evidence of CH_4 is located at the incorrect K_p for the planet and therefore treated as a non-detection. Figure reproduced from Brogi et al. (2016).

Wheatley, 2015; Seidel et al., 2020, 2023). High resolution observations can be performed on non-transiting planets, by identifying atmospheric features with redshifts corresponding to the planet’s orbital velocity (e.g., Brogi et al., 2014; Snellen et al., 2014; Parker et al., 2024).

As high resolution spectroscopy is sensitive to both different opacity sources and different altitudes in a planetary atmosphere, it is a highly complementary technique to low resolution transmission spectroscopy. Combining low resolution and high resolution observations can grant stronger constraints than any individual technique (Brogi et al., 2017; Pino et al., 2018; Khalafinejad et al., 2021). However, there have also been strong discrepancies between low and high resolution observa-

tions that remain to be completely resolved; for example in the case of HD 209458b, JWST observations (Xue et al., 2024) ruled out the presence of multiple species detected in high resolution (Giacobbe et al., 2021), at least at the altitudes probed by low resolution transmission spectroscopy.

1.3.4 Phase curves

Any time-series observation of an exoplanet, even a non-transiting exoplanet, will contain an unresolved combination of the star and planet’s emission. If the planet has a different emission spectrum between the dayside and the nightside, either due to reflected starlight or a difference in thermal emission arising from a day-night temperature contrast, the planetary contribution to the combined flux will vary with the phase of the orbit. Therefore by observing the complete orbit of an exoplanet system, the flux will vary based on the proportion of the brighter dayside exposed.

This technique, known as a phase curve, can be used to longitudinally map the brightness of the planet, and hence the temperature when using longer wavelengths corresponding to thermal emission, or albedo when using shorter wavelengths corresponding to reflected starlight. The technique is even more powerful for transiting planets, where a phase curve observation also yields a secondary eclipse (or typically two, by extending the observation slightly) and a transit, the former of which can be used to completely subtract out the stellar spectrum from the whole observation. Meaningful constraints from phase curves require strong planet-star flux ratios, making this technique most effective for strongly illuminated planets, although phase curves are also powerful for providing constraints on rocky planet atmospheres (Gillon et al., 2026). As the planetary orbits of even the most close-in exoplanets can last for days, phase curves are observationally expensive and only realistic with space-based telescopes.

Phase curve observations with Spitzer and HST demonstrate that the brightest (and hence hottest) part of the planet is offset from the substellar point (e.g., Knutson et al., 2007; Mikal-Evans et al., 2023), fitting the prediction from dynamical modelling of strong equatorial jets (Showman & Guillot, 2002). Phase curve observations also constrain the day-night temperature redistribution (e.g., Cowan et al., 2012; Kreidberg et al., 2018), demonstrating strong day-night temperature contrasts in hot Jupiters as well as evidence for dayside temperature inversions. JWST observations greatly improve our ability to characterize energy balance and horizontal energy transport for a variety of different classes of exoplanet atmospheres, including hot Jupiters (Mikal-Evans et al., 2023) and warm Neptunes (Kempton et al., 2023), and are able to provide more reliable constraints than for emission alone in terrestrial planets (Hammond et al., 2025; Gillon et al., 2026).

1.3.5 Direct imaging

For directly imaged exoplanets, the planet’s emission spectrum can be directly measured without contamination from the stellar spectrum. This enables higher resolu-

tion and signal to noise observations than is typically possible for unresolved planets. Although directly imaged planets occupy a very unique part of the exoplanet parameter space, highly distinct from both Solar System and transiting planets, one of the main efforts of the field is to bridge this gap. This gap will continue to close as technology improves, with upcoming observatories like the Extremely Large Telescope (ELT, Gilmozzi & Spyromilio, 2007), Habitable Worlds Observatory (HWO, Feinberg et al., 2024), and Large Interferometer for Exoplanets (LIFE, Quanz et al., 2022).

Directly imaged exoplanet atmospheres have exhibited many of the same opacity sources seen in transiting planets, including H_2O , CH_4 and CO (Konopacky et al., 2013; Janson et al., 2013), as well as unambiguous detections of NH_3 (Mâlin et al., 2025; Voyer et al., 2025). An example of the benefit offered by the combination of high resolution and high signal to noise possible in direct imaging observations is the detection of isotopologues in directly imaged planets and brown dwarfs, including ^{13}CO (Zhang et al., 2021a) and $^{15}\text{NH}_3$ (Barrado et al., 2023; Kühnle et al., 2025).

1.4 Overview and Thesis Structure

In this thesis, I investigate the atmospheres of giant transiting exoplanets, ranging in temperature from hot Jupiters, to a more moderately irradiated population I refer to as temperate Jupiters. My overall goal is twofold, to investigate how the atmospheric processes described in Sec. 1.2 shape trends and transitions in gas giant atmospheres, and to understand the origin and evolution of giant planets (as introduced in Sec. 1.1.4). To do this, I use the analysis and interpretation of low resolution transmission spectroscopy observations (Sec. 1.3.1), combined with predictions made using numerical models targeted towards a new regime of lower temperature giant planets.

In Chapter 2, I outline the methods I use to do this, describing the use of both optical ground-based and infrared JWST observations and the associated data reduction and analysis techniques, as well as the computational modelling used for both making predictions and interpreting observations. In Chapter 3 I present the optical transmission spectrum of the hot Saturn HAT-P-44b, based on ground-based observations with the William Herschel Telescope (WHT). In Chapters 4 and 5 I present JWST transmission spectra obtained as part of the BOWIE-ALIGN programme investigating the origin of hot Jupiters, with the full analysis of the hot Jupiter HAT-P-30b in Chapter 4, and atmospheric modelling to aid in the interpretation of additional planets within the sample in Chapter 5. In a move towards bridging the gap between the previously discussed hot gas giants and the Solar System regime, I present a modelling analysis of temperate Jupiters in Chapter 6, aimed at making testable predictions and to support the interpretation of future JWST observations. Finally I discuss my conclusions, the opportunities for future work, and the long-term outlook for the field in Chapter 7.

Chapter 2

Methods

2.1 Ground-based Observations

The exoplanet atmosphere characterization techniques discussed in Sec. 1.3 leverage both ground-based and space-based telescopes. Ground-based instruments contribute in two major ways: they are the only way to obtain high-resolution spectroscopy of exoplanets, due to the bulky nature of high resolution spectrographs prohibiting their usage on space telescopes; and they can also be used to measure low resolution transmission spectra at optical and very near-infrared wavelengths (typically 300–950 nm), where telluric absorption is less of a concern. The latter approach is a highly effective supplement to space-based observations, given the scarcity of HST observing time (the only space telescope suitable for these observations), and is the method employed in the ground-based observations analysed in this thesis (Chapter 3). The low-resolution optical transmission spectroscopy measurements that can be obtained from ground-based telescopes is a powerful tool in reducing degeneracies in chemical composition and distinguishing between clear, hazy, and cloudy atmospheres (see Sec. 1.3.1).

Due to the strong systematics introduced by measuring through the Earth’s atmosphere, which are often spectroscopic in nature, it is typical to use a comparison star to obtain differential spectroscopic light-curves which are then fit to determine the transmission spectrum. This differential spectroscopy requires a comparison star, which is ideally a good match in both colour ($|\Delta(B - V)| < 0.5$) and magnitude ($\Delta V < 1$) to the target star. In theory, a good nearby comparison star should be sensitive to identical changes in the Earth’s atmosphere, removing telluric effects from the differential light-curve. This requires an instrument with a sufficiently large field of view, such that a suitable comparison star can be found. It is not unusual for no suitable comparison star to be present, but sophisticated detrending methods including Gaussian processes (GPs) have been used to effectively mitigate for the lack of a comparison star (Panwar et al., 2022; Spyros et al., 2023). My analysis in Chapter 3 makes use of a nearby, well-matched comparison star.

Like in all transmission spectroscopy, ground-based observations are long time-series observations (often an entire night), with a cadence set to minimize

readout times while also avoiding saturation and maintaining a sufficiently rapid cadence to sample the shape of the transit (300s exposures generally being the longest considered). For my differential observations involving a comparison star, a wide slit is employed to avoid differential slit losses, rotated to include the target and comparison. This does not impact the practical wavelength resolution of the observation, which is typically set by the atmospheric seeing. Broad wavelength bins are generally used to create sufficiently precise spectroscopic light-curves. Narrow bins (commonly 3 nm, e.g. Alderson et al., 2023) are sometimes used to investigate the line absorption core of alkali metals sodium and potassium, which can be further probed using bins of incrementally-increasing width.

2.2 JWST Observations

Launched in 2021 December, JWST has become the premier tool for exoplanet atmosphere characterization. A 6.5m space telescope at the Earth-Sun Lagrange point 2 (Gardner et al., 2006; Rigby et al., 2023), JWST has unmatched precision and stability for observing the time-series stellar spectra required for low-resolution transmission and emission spectroscopy, while unlocking near- and mid-infrared spectroscopic wavelength coverage only previously accessible at lower spectral resolution and precision with the Spitzer Space Telescope (Werner et al., 2004). This enables deep atmospheric characterization, with the detection of new molecular species including CO₂ (JWST Transiting Exoplanet Community Early Release Science Team et al., 2023), CH₄ (Bell et al., 2023), and SO₂ (Alderson et al., 2023; Rustamkulov et al., 2023), as well as the identification of cloud species (Grant et al., 2023b). I analyse JWST observations of hot Jupiters in Chapters 4 and 5.

2.2.1 Instrument setup

All four of JWST’s instruments are used for time-series spectroscopy, each with different resolutions, wavelength coverage, and brightness limits: the Near InfraRed Spectrograph (NIRSpec), the Near InfraRed Imager and Slitless Spectrograph (NIRISS), the Near-InfraRed Camera (NIRCam), and the Mid-InfraRed Instrument (MIRI). Each instrument has multiple filters and/or modes suited for transmission spectroscopy.

NIRSpec

Transmission and emission spectroscopy observations with the NIRSpec instrument (Jakobsen et al., 2022) are conducted in the Bright Object Time Series (BOTS) mode. NIRSpec offers a choice of six paired medium (M, $R \sim 1000$) and high (H, $R \sim 2700$) resolution diffraction gratings; in the latter the light is dispersed across two detectors, NRS1 and NRS2, with a detector gap resulting in lost wavelength coverage, while in the former the light is solely dispersed across NRS1 with therefore no gap. The combinations of the following dispersing elements and filters can be

used to achieve different wavelength coverages: G395H or G395M with the F290LP filter (2.87–5.10 μm), G235H or G235M with the F170LP filter (1.66–3.07 μm), and G140H or G140M with either the F100LP (0.97–1.84 μm) or F070LP (0.70–1.27 μm) filters. There is also the PRISM dispersal element, a low resolution ($R \sim 100$) double-pass prism with no (CLEAR) filter, which can provide the entire wavelength range of 0.60–5.30 μm in a single observation as there is no need for wavelength range limiting transmission filters required to prevent order contamination (as is needed for diffraction gratings). Each choice of dispersal element has its own associated brightness limit to avoid saturation, with the high-resolution gratings having the brightest saturation limit, while the PRISM setting is particularly sensitive to full or partial saturation, limiting the targets for which it can be used. The sensitivity of the NIRSpec instrument to saturation is particularly notable for target acquisition, as for many exoplanet targets the acquisition integration will saturate ($J < 11.4$ mag). This is usually remedied using fixed offset acquisition from a nearby background star. The JWST data analysed and interpreted in Chapter 4 and Chapter 5 are all acquired using the NIRSpec instrument with the G395H/F290LP grating/filter combination.

Other instruments

NIRISS (Doyon et al., 2012) is used for time-series spectroscopy of exoplanet targets in Single Object Slitless Spectroscopy (SOSS) mode. The instrument uses a grism with a resolving power of $R \sim 700$ to disperse three spectral orders in unique traces onto the detector, with the first order providing coverage of 0.9–2.8 μm , the second order covering 0.6–1.4 μm , and the third order generally being too weak to be useful. The NIRISS detector has a choice of subarray modes, with SUBSTRIP256 containing all three orders, while SUBSTRIP96 only contains the first order but can be used for brighter targets. There are also new multistripe modes to observe even brighter targets. As a slitless spectrograph with multiple cross-dispersed traces per source, SOSS observations potentially suffer from contamination from nearby objects, which can restrict its usage for observing exoplanet hosts in crowded fields. This can be somewhat mitigated using constraints on position angle to schedule observations.

NIRCam (Horner & Rieke, 2004) can also be used for transmission spectroscopy, using the grism time series mode. In this mode, NIRCam can provide spectroscopy at a resolving power of $R \sim 1600$. Two main filters are used for transiting exoplanet observations: F444W (3.9–5.0 μm) and F322W2 (2.4–4.0 μm). Different readout modes and subarrays can be used to observe particularly bright exoplanet host stars with NIRCam. NIRCam also provides simultaneous short wavelength photometry.

MIRI (Rieke et al., 2015) is the only JWST instrument capable of observing at mid-infrared wavelengths ($> 5.3 \mu\text{m}$). The Low Resolution Spectrograph (LRS) slitless mode is the main mode used in spectroscopic transiting exoplanet observations, providing coverage from 5–14 μm at a low ($R \sim 100$) resolution. Like NIRISS

SOSS it is a slitless spectrograph so contamination from nearby sources is a potential issue. There is also the Medium Resolution Spectrograph (MRS), which can provide higher resolution and longer wavelength coverage (4.9–27.9 μm). However, its usage for transiting exoplanet observations is limited, as only a single wavelength channel can be provided at once, with four total transits/eclipses required to cover the entire wavelength range. MRS has been successful for the characterization of directly imaged planetary-mass companions (Miles et al., 2023). The imaging mode of MIRI is also used for photometric eclipse observations to characterize the dayside temperature of small planets (Greene et al., 2023; Zieba et al., 2023).

Detectors and readout

JWST is fitted with mercury-cadmium-telluride detectors for the near infrared instruments, and arsenic doped silicon detectors for MIRI. These infrared detectors operate in such a manner that the count of each pixel can be sampled without removing the accumulated charge, a process known as a non-destructive readout. Each measurement of the charge is referred to as a frame, with multiple frames (plus skipped frames in some readout patterns) being averaged into groups onboard to reduce data volume. Most readout patterns used for transmission spectroscopy (including the NRSRAPID pattern of NIRSpec and NISRAPID pattern of NIRISS) use a single frame per group. Multiple groups are taken consecutively to form an integration, which takes the form of a flux ramp consisting of the flux at each group measurement, after which the charge in the detector is reset. To obtain an overall flux value from each integration, the slope of this ramp is fit as part of the data reduction. Each integration is a flux measurement of the time-series spectrum. Multiple integrations are combined into a single exposure, usually covering the entire time series of interest (in the case of transmission spectroscopy, the transit plus out-of-transit baseline). Due to a hard limit of 10,000 s, some exoplanet transit observations require multiple exposures, causing brief interruptions in data acquisition (40 s in Kempton et al., 2023).

2.3 Data Reduction and Analysis

There are a number of open-source data reduction pipelines for low-resolution transmission spectroscopy observations, including *Eureka!* (Bell et al., 2022), *ExoTiC-JEDI* (Alderson et al., 2022), and *exoTEDRF* (Radica, 2024). In this thesis I focus on the *Tiberius* pipeline (Kirk et al., 2017, 2021), which I use for the analysis of ground-based observations of HAT-P-44b and JWST observations of HAT-P-30b in Chapters 3 and 4 respectively. In this section I give a concise overview of the workflow used to analyse low-resolution transmission spectroscopy data, including processing of the raw images, spectral extraction to obtain light-curves, post-extraction processing, and finally the fitting of the light-curves to obtain transmission spectra. I also describe where the approach for ground-based and JWST data diverges.

2.3.1 Raw image processing

For ground-based observations, the raw image processing is relatively straightforward. A master bias, created from median-stacking multiple bias frames, is subtracted from the raw image to remove the bias voltage. A common step in astronomical image processing is to divide the science frames by a flat field, obtained through observation of either a lamp-illuminated dome or the sky, in order to remove the effect of pixel-to-pixel variations in sensitivity. While this can be done for ground-based transmission spectroscopy observations, it is not standard practice. This is because of repeated findings from multiple independent analyses and instruments that flat-fielding either results in no improvement in noise properties, or actively introduces red noise into observations (Jordán et al., 2013; Kirk et al., 2017; Rackham et al., 2017; Huitson et al., 2017; Gibson et al., 2017; Bixel et al., 2019; Alderson et al., 2020; Kirk et al., 2021; Spyratos et al., 2021; Ahrer et al., 2022). I therefore do not use flat-field corrections in my analyses. The lack of improvement from flat-fielding follows from the fact that thousands of pixels contribute to each wavelength bin in conventional ground-based transmission spectroscopy, minimizing the impact of pixel-to-pixel variations. Meanwhile, the potential to introduce red noise is a result of the difficulty of removing sky lines from sky flats, which could realistically introduce notable systematics into the data (e.g. Bixel et al., 2019). Cosmic rays can be removed from the raw images at this time, or they can be removed from the stellar spectra; both approaches have been used in published transmission spectra (e.g., Alderson et al., 2020; Ahrer et al., 2022).

For JWST observations, much of the initial calibration and processing of raw images is performed using Stage 1 of the open-source `jwst` pipeline. All of the Stage 1 steps are performed in the following order: flagging data quality, detection of saturation, bias subtraction, linearity correction, reference pixel correction, dark current subtraction, ramp-fitting, and gain scale correction. The `jump` step is skipped following the finding of Rustamkulov et al. (2023) that it can increase noise in exoplanet transit light curves. A custom correction is implemented within `Tiberius` to mitigate for $1/f$ noise, an instrumental correlated noise source caused by the detector readout and current biases in the electronics (Birkmann et al., 2022). This is applied before ramp-fitting, using a trace obtained by running a cursory spectral extraction on the uncalibrated data to mask the spectrum to permit proper background estimation. A wavelength solution is obtained from running the `assign_wcs` and `extract_2d` steps from Stage 2 and Stage 3 of the `jwst` pipeline respectively. Custom cosmic ray identification and removal using `Tiberius` is then applied by subtracting a running median, calculated from a moving box of three pixels for the time series of each pixel, to identify outliers and replace them with the running median. A bad pixel mask is also typically constructed at this time from a median science image from the final group of each integration. There are two methods for this within `Tiberius`, using either a row-by-row running median method or the median absolute deviation of the entire frame to identify outliers, which are then added

to a mask of flagged bad pixels from the `jwst` pipeline.

2.3.2 Spectral extraction

Tiberius uses normal aperture extraction to extract the photometry in the cross-dispersion direction – each row for ACAM or column for NIRSpec data (although rotated prior to this step in the code and hereafter referred to accordingly as rows for simplicity). This is done by summing the flux in the spatial direction within a defined aperture. This contrasts with the optimal extraction method (Horne, 1986) used in **Eureka!** and other long-slit spectroscopy pipelines, where pixels within the point spread function (PSF) are weighted accordingly to their signal to noise. These approaches have been shown to have similar performance in the high signal to noise stellar spectra used in transmission spectroscopy, verified by consistent results between long-slit spectroscopy pipelines across a wealth of datasets (e.g. Rustamkulov et al., 2023, as well as in Chapter 4). The trace centre is fitted row-by-row using a Gaussian, with the trace centres subsequently fit with a fourth-order polynomial to ensure a smooth trace.

The background is subtracted from the flux for each row by estimating the background level within the aperture. This is done using a defined background region, using either the median of the background level or a polynomial fit to the background level on either side of the trace. The former approach is well-suited for NIRSpec observations, but it performs poorly in ground-based data, where the background level has a complex functional form in the cross-dispersion direction. For ground-based data, it is typical to use a narrow region with simple (linear or quadratic) polynomial for background subtraction. The width of the aperture region and the location of the background region can all be customized in **Tiberius** for each observation. These regions are chosen to balance a number of factors, with the aperture width a balance of the inclusion of stellar flux versus background, while the background region needs to be broad enough to well-characterize the level of the background across the aperture, narrow enough that this doesn't capture large-scale variations that require a more complicated polynomial to fit, and sufficiently offset from the trace to avoid including stellar flux. It is typical to iterate upon these decisions, choosing the final set of parameters based on the statistics of intermediate data products (e.g. the residuals of white-light light-curve fits). See Sec. 3.7.4 for a practical test of different background subtraction methods from ACAM data. For ground-based data, both the target and comparison stellar spectra are extracted to enable creation of differential light-curves. Example aperture and background regions for ACAM and JWST/NIRSpec data are shown in Fig. 2.1.

There are a few processing steps used to improve data quality and create spectroscopic and white-light light-curves from the time-series stellar spectra. If cosmic ray removal has not already been performed, it can be done at this stage, identifying outliers in the stellar spectra attributed to cosmic rays using a running median or comparison to a median reference spectrum. **Tiberius** can then resample spectra




Figure 2.1: Aperture extraction and background subtraction regions overlaid upon raw images from ACAM observations of HAT-P-44b and its comparison star (top panel) and JWST/NIRSpec observations of HAT-P-30b (bottom panel). The aperture region is shown with solid blue (HAT-P-44b) and red (HAT-P-30b) lines, while the background region is marked with dotted blue (HAT-P-44b) and red (HAT-P-30b) lines.

onto a consistent wavelength grid using either cross-correlation or Moffat profiles fit to stellar absorption features to identify pixel shifts, which are then typically smoothed using a polynomial function or spline. Both arc calibration frames and stellar and telluric features can be used for wavelength calibration in ground-based data, with the latter strongly favoured given the stretch and shifts that occur in wavelength solution during the night. Finally, white-light and spectroscopic light-curves can be created by summing the time-series stellar spectra into wavelength bins. For ground-based data, a differential light-curve, with the target star spectra divided by the comparison star spectra, is used instead. Different wavelength resolutions are used for different purposes. $R = 100$ is generally suitable for the detection of broad molecular features and continuum opacities, with generally only minor improvements in posteriors from increasing the data resolution (at the expense of increased computation time in both light-curve fitting and the generation of models used for interpretation). Higher resolution data, for example the $R \sim 2700$ pixel-level spectra from NIRSpec G395H observations, can be useful in cross-correlation analyses to detect CO (Grant et al., 2023a; Esparza-Borges et al., 2023) and other species (Esparza-Borges et al., 2025) and even isotopologues, as well as distinguish between species sharing broad-band absorption features.

2.3.3 Light-curve fitting

The final step to obtain a transmission spectrum is to calculate the transit depth in each wavelength bin, done by fitting a transit model to the spectroscopic light-curves. As only relative differences in transit depth are relevant in transmission spectroscopy, it is usual to fix the system parameters other than transit depth. These system parameters can be obtained from the joint fit of light-curves from multiple observations, but it is more common to just use the system parameters obtained from a fit to the white-light light-curve.

Tiberius uses **batman** (Kreidberg, 2015b) for the transit light curve model, which is either an implementation of the analytic model of Mandel & Agol (2002), or a numerical model, depending on the limb darkening law used. All system parameters are typically fitted for in the white-light light-curve except period and eccentricity, which are frequently fixed given a single transit time series is not informative on these quantities. A variety of approaches to limb darkening are possible, including different laws (e.g., linear, quadratic, square-root, logarithmic, three- and four-parameter) and parametrizations (e.g. Kipping, 2013) with different complexities required for different observing cadences, light-curve precisions, and stellar hosts (e.g. Espinoza & Jordán, 2016, Fig 2.2). Limb-darkening parameters can be free parameters, or fixed to calculations from stellar models (e.g., Husser et al., 2013; Magic et al., 2015) implemented using **LDTK** (Parviainen & Aigrain, 2015) or **ExoTiC-LD** (Grant & Wakeford, 2024).

It is common, especially for ground-based data, for the light-curves to contain correlated noise, whether that be instrumental, telluric, or astrophysical (e.g. stellar




Figure 2.2: Comparison of the performance, measured using the mean-squared error, of different limb-darkening laws for fitting hot Jupiter transit light-curves as a function of light-curve precision. Higher cadence (more in-transit points) data is shown with dotted lines. Figure reproduced from Espinoza & Jordán (2016).

activity) in origin. This can potentially be detrended from the light-curves by multiplying the transit model with a systematics model. A variety of basis functions can be summed to construct a systematics model. In this thesis I focus on polynomial detrending using observational parameters, which include time, full-width at half-maximum (FWHM) of the PSF, airmass, sky background, x-pixel position, y-pixel position, and derotator angle. An alternative approach is to use Gaussian Process (GP) regression (Ramussen & Williams, 2006). As I do not use GPs in any of the analyses included in this thesis, I do not cover the method here. However, GP detrending has been used successfully in transmission spectroscopy (e.g., Gibson et al., 2012a,b; Kirk et al., 2017; Loudén et al., 2017; Ahrer et al., 2022); more detail on the use of GPs in ground-based transmission spectroscopy can be found in Kirk (2018) and Ahrer (2023). Using simple polynomials of observational basis functions has the advantage of leveraging astronomically tractable causes of red noise, for example full width at half maximum (FWHM) can deal with the bleeding of stellar features between bins on nights with varying seeing, while sky background detrending can be related to non-linearity of the detector. When spectroscopic light-curves are fitted with GPs independently, it tends to give rise to considerably reduced precision compared with polynomial detrending. In contrast, the simultaneous fitting of all spectroscopic bins with shared GP hyperparameters for length scale has been shown to give rise to more closely comparable precision, but at the expense of significant computational resources (Ahrer et al., 2022).

Fitting methods and model comparison

A few different numerical methods can be used for light-curve fitting. For simple applications where short computation time is desirable, including comparing large numbers of different detrending methods or when using very simple models on high quality data (for example spectroscopic light-curve fitting with JWST), I use the Levenberg-Maquardt algorithm (Levenberg, 1944; Marquardt, 1963) for light-curve fitting. This is an iterative, numerical algorithm for least-squares minimization, which can rapidly find a local minimum. However, where sampling of the posterior distribution is desired, as well as to better avoid getting trapped in a local minimum, Bayesian methods should be used. This is necessary for white-light light-curve fits, where understanding the posterior distribution of planetary parameters is important, and is also preferable when using more complicated systematic models.

For light-curve fitting, I use the Markov Chain Monte Carlo (MCMC) method to sample the posterior probability distribution. This is a stochastic (Monte Carlo) method, using n walkers sampling different parts of the multi-dimensional model parameter space with the Metropolis-Hastings algorithm (Hasting, 1970). At each step of the chain and at each walker, a new location in model parameter space is considered, where the posterior distribution ratio of the new location relative to the old ratio is calculated and compared to a random number between 0 and 1. When the probability ratio is greater than the random number, the move is accepted, updating the location of the walker in the parameter space, otherwise the move is rejected and the walker doesn't move. This moves the walkers on average into a higher probability (better fitting) part of the parameter space without directly computing the Bayesian evidence (refer ahead to Sec. 2.4.3 for the definition of Bayesian evidence), which slows computation. Over time, the sample distribution of the Markov chains better approximates the true posterior probability distribution function. MCMC is implemented in *Tiberius* via the *EMCEE* package (Foreman-Mackey et al., 2013), using an initial burn-in chain, after which the photometric uncertainties are scaled to give a reduced χ^2 of 1, followed by a production chain to estimate the posterior distribution.

To choose between different systematic models and limb-darkening methods, I use the Bayesian Information Criterion (BIC, Schwarz, 1978) to make comparisons. This is an estimate of the Bayesian evidence of the model, which is not directly computed in MCMC methods. BIC is calculated from the goodness of fit with a penalization term based on the number of free parameters, rewarding models that can achieve comparable goodness of fit with fewer free parameters. This ensures the complexity of the model fits the detail provided by the data. The difference of BIC is an estimate of the Bayesian evidence ratio between two models, with a ΔBIC of 6 amounting to strong evidence against the higher BIC model (Kass & Raftery, 1995). The BIC is calculated as follows:

$$\text{BIC} = \chi^2 + k \ln n, \quad (2.1)$$

where k is the number of model free parameters, n is the size of the sample (in this case number of points in the light-curve), and χ^2 is calculated by:

$$\chi^2 = \sum_i^n \frac{(y_{\text{data},i} - y_{\text{model},i})^2}{\sigma_i^2}, \quad (2.2)$$

where y is the variable (in this case flux), i in this case is image number, and σ is the error (in the flux).

2.4 Atmospheric Modelling and Retrievals

Many aspects of exoplanet atmospheres are most readily modelled numerically. This includes the simulation of energy transport through the atmosphere, to determine the temperature-profile. Chemical composition can be determined either by assuming equilibrium conditions and solving the time-independent thermodynamic state of highest entropy, or time-dependently by solving for chemical kinetics until a steady state is achieved. These techniques can be solved using a 1D approximation of a planetary atmosphere, or in more complex general circulation models (GCMs) that simulate atmospheric dynamics and transport processes in three dimensions. The outputs of these simulations, or simple parametric approximations, can be used as inputs to simulate observations. Of particular relevance to this thesis is the ability to create synthetic transmission spectra by simulating the radiative transfer of light through a model atmosphere of some collection of opacity sources. This allows interpretation of observations through one of two methods: forward modelling, where synthetic transmission spectra are directly generated, which can then be compared to the data; or atmospheric retrievals, the inverse problem, wherein a sampling method and Bayesian inference with the observation is used to determine the posteriors of some parametric transmission spectrum model.

In Sec. 2.4.1 I detail the radiative transfer models used to create synthetic transmission spectra, and in Sec. 2.4.2 I describe how chemistry can be determined and parametrized with these models, both relevant to all four research chapters. In Sec. 2.4.3 I give a theoretical overview of atmospheric retrievals, which are the primary interpretation technique used in Chapters 3, 4 and 5. The detailed physical and chemical simulations used in Chapter 6 do not appear elsewhere in the thesis, and are therefore described there.

2.4.1 Radiative transfer

To simulate a transmission spectrum, the goal is to calculate the effective cross-sectional area of a planet at a given wavelength, given the opacity sources in the atmosphere. This is identical to saying what radius of a fully opaque disc would give rise to an identical transmission of flux – this is relevant because when fitting light-curves, a fully opaque disc is the exact assumption used in the transit model. The effective area at a given wavelength, A_λ , of a planet of radius R_p can be written

in terms of the impact parameter r and the optical depth τ_λ as follows:

$$A_\lambda = 2\pi \int_0^{R_p} r(1 - e^{-\tau_\lambda})dr, \quad (2.3)$$

where the transmission is evaluated from a grazing light ray. This effective area of the planet can be in turn related to the ratio of in- and out-of-transit flux ratio ($\frac{F_{\text{in}}}{F_{\text{out}}}$) as follows:

$$\frac{A(\lambda)}{\pi R_*^2} = 1 - \frac{F_{\text{in}}(\lambda)}{F_{\text{out}}} \simeq \left(\frac{R_p(\lambda)}{R_*} \right)_{\text{obs}}^2 \quad (2.4)$$

which is in turn closely related to the square of the transit depth ratio $\left(\frac{R_p(\lambda)}{R_*} \right)_{\text{obs}}$ measured in transmission spectroscopy. This is not an equality, because the transit depth ratio is not computed simply from the in- and out-of-transit fluxes, but rather from fitting an opaque disc model to a whole light-curve. While for a radially symmetric planet with “sharp edges” (rapid spatial transition from optically thick to transparent) this should be identical, it is worth highlighting that there is a discrepancy between the physical quantities we compare when fitting conventional model transmission spectra to transmission spectra from transit light curve fitting.

In this thesis I use the open-source radiative transfer modelling software `petitRADTRANS`¹ (Mollière et al., 2019), which solves for Eq.2.3. To set up the model, it is necessary to parametrize the atmosphere in terms of a sequence of spherically-symmetric pressure layers, in each of which the temperature and chemical composition is defined. A reference pressure, an arbitrary pressure level corresponding to an inputted radius of the planet, which along with the gravity evaluated at this reference radius (which varies according to Newton’s law as r^{-2}) and the assumption of hydrostatic equilibrium and the ideal gas law, is used to geometrically convert from a 1D pressure coordinate system to a 1D impact parameter coordinate system (r) – see Fig.2.3.

The remaining step is therefore to calculate the transmission $e^{-\tau_\lambda}$ at each wavelength. This transmission can include multiple opacity sources, including line absorption, collision-induced absorption, Rayleigh scattering, and additional opacity from aerosols. This requires the mass fraction of any species contributing opacity to be defined in each pressure layer. Line opacity sources are obtained from both computational and experimental sources, mainly including the spectroscopic databases HITRAN² and ExoMol³, and converted into pressure-temperature-opacity tables.

¹<https://petitradtrans.readthedocs.io/en/latest/>

²<https://hitran.org/>

³<https://exomol.com/>




Figure 2.3: Geometry of transmission spectroscopy, demonstrating the grazing light ray (red) over which the transmission is evaluated for a given impact parameter r_i (defined using pressure layers P_i). Figure reproduced from Mollière et al. (2019).

Correlated- k

There are two methods to calculate the transmission in each wavelength bin in `petitRADTRANS`: correlated- k and line-by-line. Correlated- k is an assumption that enables accurate band-integrated transmission calculations at reduced spectral resolution, and is implemented at a maximum spectral resolution of $R = 1000$ in `petitRADTRANS`. Line-by-line on the other hand can be used at a much higher spectral resolution of $R = 10^6$, but at this resolution is much more computationally expensive. I use correlated- k opacity sampling in this thesis, mainly at the native $R = 1000$ resolution, which is ideal for the low-resolution data I use, which has a maximum spectral resolution of $R = 400$ (to ensure Nyquist-sampling, the model resolution should be at least twice the data resolution).

The key idea of correlated- k (e.g., Lacis & Oinas, 1991; Fu & Liou, 1992) is to consider an opacity source in a single layer with a wavelength bin, $\Delta\nu$, and construct a cumulative opacity distribution function, $\gamma(k)$:

$$\gamma(k) = \int_0^k f(k') dk' \quad (2.5)$$

from the opacity distribution function $f(k)$, which describes how “often” (i.e. the probability at a random infinitesimal wavelength) the layer has a given opacity per unit mass k . $k(\gamma)$ is effectively a reordering of the opacities to a monotonic and

typically smooth function of ascending opacity strength (see Fig. 2.4).

This cumulative opacity distribution function can in turn be used to convert the opacity integral from frequency-space, where it is an extremely rapidly changing function due to the many lines and troughs within a wavelength bin $k(\nu)$, to the smooth function $k(\gamma)$:

$$\int_{\nu}^{\nu+\Delta\nu} e^{-k(\nu)u} d\nu \Rightarrow \int_0^1 e^{-k(\gamma)u} d\gamma, \quad (2.6)$$

where u is the absorber column mass. This integral needs to be performed in every pressure layer, at every wavelength bin, for every opacity source. By performing this conversion, the integral can instead be computed using the smooth function $k(\gamma)$, which can be estimated using a numerical approximation far more accurately. This is done by recording the opacity at a discrete number of Gaussian quadrature coordinates, N_γ , in γ -space (16 are used per wavelength bin in `petitRADTRANS`). An $R = 1000$ correlated- k opacity table for `petitRADTRANS` therefore constitutes the opacity of a species evaluated over 16 γ coordinates in each wavelength bin (set by the $R = 1000$ resolution), in turn evaluated across a grid of temperatures and pressures.

The correlated- k assumption is to now assume that the distribution of opacities $k(\gamma)$ is preserved between different layers of the atmosphere, allowing the same γ coordinates to be used across all layers. This means the transmission from the top of the atmosphere ($P=0$) to a given layer i at pressure P_i , from a given species j with opacity κ_j and mass fraction X_j can be written (adapted from Mollière et al., 2019):

$$e^{-\tau_{ij}} = \sum_{l=1}^{N_\gamma} \exp\left(-\int_0^{P_i} \frac{X_j \kappa_{jl}}{g \cos \theta} dP\right) \Delta\gamma_l, \quad (2.7)$$

where g is the gravitational acceleration and θ is the angle between the atmospheric normal and the ray (which is handled when the geometry is converted to a grazing geometry with impact parameter r). The calculation is further simplified by assuming that the opacity distribution functions between molecular species are uncorrelated (random overlap), allowing the total transmission to be simply calculated as the product of the transmission of each species:

$$e^{-\tau_i} = \prod_j e^{-\tau_{ij}}. \quad (2.8)$$

Correlated- k is significantly more accurate than line-by-line resampled to the same resolution, as the cumulative opacity distribution can be downsampled while preserving information better than the extremely high frequency information of the opacity distribution in wavelength-space. This means line-by-line opacity tables need to be many times the resolution of the data to ensure sufficient accuracy. This can be understood by comparing both methods at the same spectral resolution, lower





Figure 2.4: Demonstration of the conversion of an opacity distribution from frequency space ($\kappa(\nu)$, top panel, equivalent to $k(\nu)$ in the main text) to cumulative opacity distribution space ($\kappa(y)$, bottom panel, equivalent to the $k(\gamma)$ in the main text), using data for H_2O at 1500 K and 1 atm pressure (Rothman et al., 2010). The impact of cutting off the Lorentzian wings of the spectral lines is demonstrated by contrasting the green (no clipping of the wings) and red (wings clipped at $500 \Gamma_L$) distributions. Figure adapted from Grimm & Heng (2015a).

than the wavelength resolution of the spectral lines. Correlated- k is able to simulate the contribution of much higher resolution spectral information (i.e. the individual spectral lines), by evaluating the opacity at a set of γ -coordinates which accurately sample the true cumulative opacity distribution calculated using a complete line-list. This comes at some computational (the transmission is calculated for each γ -coordinate in each wavelength bin) and storage cost (the opacity at each γ -coordinate needs to be recorded). In contrast, in line-by-line, each wavelength bin is calculated using a single opacity per wavelength bin per species, which is downsampled from a much higher resolution line-list through interpolation, which results in much less accurate downsampling and therefore needs much higher resolution opacity tables to model the same data. Correlated- k is unnecessary when the data resolution, and hence desired model resolution for comparison, is high enough to be sensitive to individual lines.

Temperature structure

For the radiative transfer model, the temperature has to be set in each pressure layer of the atmosphere. There are numerous parametrizations for temperature that can be used, including simple assumptions like an isothermal atmosphere, where the temperature is held constant throughout; analytic formulations like the Guillot profile (which considers both intrinsic and irradiation flux with a double-grey opacity atmosphere, Guillot, 2010); and parametric profiles designed to capture physical features including temperature inversions (e.g. Madhusudhan & Seager, 2009). While the temperature structure is very important in determining the emission properties of an atmosphere, in transmission the temperature has a more minor role, scaling the feature strength via the scale height through hydrostatic equilibrium, and to a lesser extent modifying the feature width due to thermal broadening. The other key role of temperature is setting the equilibrium chemistry in each layer, which I discuss in Sec. 2.4.2. Using an isothermal temperature profile is usually sufficient in transmission modelling to capture the main effects of temperature on the spectrum without resulting in too much degeneracy, however it has been suggested that this can bias the interpretations from atmospheric retrievals (Rocchetto et al., 2016; Welbanks & Madhusudhan, 2022). Furthermore, differences in atmospheric temperature between the morning and evening limbs of tidally-locked exoplanets (which includes hot Jupiters) are both expected and observed (Murphy et al., 2024), undermining the accurate 1D modelling of exoplanet atmospheres (Caldas et al., 2019; MacDonald et al., 2020).

Aerosol parametrizations

`petitRADTRANS` has numerous options for incorporating opacity due to aerosols (see Sec. 1.2.4 for a theoretical discussion of aerosols). A common method used to match the observation of muted spectral features frequently observed in hot Jupiters (Gao et al., 2021) is to simply use a completely opaque grey cloud deck. This is used

to represent the relatively flat optical to near-infrared Mie scattering spectrum of large particle-size cloud species. Another simple parametrization is to inflate the gas Rayleigh scattering opacity by some fixed scaling factor (often referred to as a haze factor), representing Rayleigh scattering from small particle-size aerosols. This gives a λ^{-4} dependence and is typically only significant at optical wavelengths.

Optical transmission spectroscopy observations have also revealed scattering slopes with an even steeper wavelength-dependence than the λ^{-4} of Rayleigh scattering, known as “super-Rayleigh” scattering slopes (e.g., Sing et al., 2011a; Luque et al., 2020; Ahrer et al., 2022). These have been attributed to photochemical hazes (Kawashima & Ikoma, 2019; Ohno & Kawashima, 2020). To model this, **petitRADTRANS** includes the following option for a more freely parametrized scattering slope, equivalent to that from MacDonald & Madhusudhan (2017):

$$\kappa = \kappa_0 \left(\frac{\lambda}{\lambda_0} \right)^\alpha, \quad (2.9)$$

with free parameters κ_0 , the opacity at $0.35 \mu\text{m}$ in units of cm^2g^{-1} , and α the slope gradient. For super-Rayleigh slopes $\alpha < -4$. **petitRADTRANS** also allows for a cloud opacity with an arbitrarily parametrized functional dependence on pressure and wavelength.

There is also the option to include opacity from real cloud species, including a variety of mineral cloud species expected to condense in the high temperatures of hot Jupiter atmospheres, as well as liquid and solid H_2O and liquid H_2SO_4 . These are implemented using opacity calculated from Mie theory, either directly for spherical particles, or using an extended formulation with a distribution of hollow spheres to model irregular particles (Min et al., 2005). For this parametrization it is necessary to input the mass fraction of the relevant cloud species, as well as the mean and standard deviation of the cloud particle radius distribution. The cloud particle radius distribution can instead be set using the analytical prescription of Ackerman & Marley (2001), using the eddy diffusion parameter (K_{zz}) to describe the vertical mixing strength, and the settling parameter f_{sed} , to express the settling velocity. These more sophisticated aerosol models are best suited to transmission observations with broad wavelength coverage, as usually both optical scattering slopes and spectral features in the mid-infrared are required to place sufficient constraints on cloud species (e.g., Grant et al., 2023b; Fairman et al., 2024).

Any of these aerosol parametrizations (except the Rayleigh scaling factor) can also be implemented in such a way to allow only partial coverage (typically referred to as patchiness) of aerosol opacity. The patchy cloud transit radius (r_λ , simply calculated from the effective area A_λ) is calculated by simply considering linear contributions from cloudy and cloud-free models, weighted by a cloud fraction parameter (c_f) as follows:

$$r_\lambda = c_f r_{\text{cloudy},\lambda} + (1 - c_f) r_{\text{clear},\lambda}. \quad (2.10)$$

In transmission spectroscopy, the effect of cloud-top pressure versus patchiness can be considerably degenerate (although consider the counter-example of MacDonald & Madhusudhan, 2017), especially in the absence of broad wavelength coverage to constrain the scattering properties.

2.4.2 Chemistry modelling

The key appeal of atmospheric characterization using transmission spectroscopy is the sensitivity to chemical composition. Often the main feature desired from a radiative transfer model is to demonstrate the absorption due to molecular features. Like temperature, the mass fraction of each potential opacity species needs to be set in each pressure layer of the atmosphere. This can be done parametrically, for example by setting the mass fraction of a species to be vertically constant through the atmosphere, or using the output of chemical modelling. For more detail on the theoretical background of the relevant atmospheric chemistry, refer to Sec. 1.2.3.

Equilibrium chemistry

A common approach to atmospheric composition used in both forward models and retrievals is to assume equilibrium chemistry, where the mass fractions in each layer are such that they minimize the Gibb’s free energy (see Sec. 1.2.3). For ease of calculation, this is usually done via interpolation from a pre-computed grid, based on the temperature, pressure, and elemental ratios within the layer. As I primarily deal with hydrogen- and helium-dominated gas giant planets in this thesis, the elemental ratios are parametrized using two quantities: metallicity, and C/O ratio. In this parametrization, each ratio of every element to hydrogen, besides helium and oxygen, is scaled from the solar elemental ratios by a constant factor, the metallicity. The oxygen abundance is instead calculated from the C/H (which is simply set by the metallicity) and the C/O ratios as follows:

$$\frac{\text{O}}{\text{H}} = \frac{1}{\text{C/O}} \left(\frac{\text{C}}{\text{H}} \right). \quad (2.11)$$

I note that this is not the only implementation method of this parametrization, with other radiative transfer codes instead fixing the O/H ratio to the metallicity, and further calculating the C/H ratio from the C/O ratio, or even using the metallicity to set (O+C)/H. This can result in an apparent discrepancy between different retrieval codes. The ratios of other elements, for example sulfur, nitrogen, and silicon, are also scaled by the metallicity. The remaining gas is assumed to be hydrogen and helium in a solar ratio.

For equilibrium chemistry calculations, I use a pre-computed grid included in `petitRADTRANS`, calculated using `easyCHEM` (Mollière et al., 2017), in turn based on (Gordon & McBride, 1994). The grid uses 100 equidistant temperature points $T \in [60, 4000]$ K, 100 log-equidistant pressure points $P \in [10^{-8}, 1000]$ bar, 20 equidistant points in C/O ratios $\in [0.1, 1.6]$, and 40 equidistant points in $[\text{M}/\text{H}] \in [-1, 3]$.

The solar values of Asplund et al. (2009) are used to define $[M/H]$. I note that there is no method to vary the nitrogen or sulfur ratios relative to carbon in this parametrization; while in many cases this is not expected to be important, with equilibrium carriers of nitrogen and sulfur rarely having a large impact on transmission spectra, it is still worth noting this is a non-physical condition, and potentially an issue for temperate planets where there is expected to be significant contribution from NH_3 .

Disequilibrium chemistry

The assumption of equilibrium chemistry (i.e. the minimization of Gibbs free energy) is dependent on an isolated system. In the above method, equilibrium chemistry is solved independently for each layer of the atmosphere, and therefore it is assumed that each layer is an isolated system. Transport processes, as well as photochemical processes from incoming stellar radiation, are therefore neglected.

A proper consideration of these processes would require modelling involving chemical kinetics (i.e. modelling the time-dependence of chemical reactions) and 3D transport processes (e.g., Baeyens et al., 2021; Zamyatina et al., 2023, 2024). However, this is highly computationally expensive, and while powerful for performing individual simulations or coarse model grids, it is not typically feasible for a fine grid of forward models or the rapid model calculation required for atmospheric retrievals (see Al-Refaie et al., 2024; Bardet et al., 2025, for an early implementation of sufficiently rapid chemical kinetics calculations to enable coupling atmospheric retrievals).

To model transport processes we can therefore make use of the quenching approximation discussed in Sec. 1.2.3. This allows us to use a single free parameter P_{quench} to model the impact of vertical mixing on the chemical composition, using equilibrium chemistry below the quench point, and simply holding it constant above. `petitRADTRANS` includes a simple implementation of CH_4 -CO quenching, thereby accounting for parametrized vertical transport (but not photochemistry).

Chemical kinetics

The most complete method of chemical modelling considered in this thesis is 1D chemical kinetic modelling. This can be used to properly solve for the impact of vertical transport (and therefore determine P_{quench} based on K_{zz}) and photochemistry. A chemical kinetics model works by numerically solving the mass continuity equations outlined in Sec. 1.2.3, using a network of chemical reactions, until the model converges and reaches a steady state. The outputs of a chemical kinetics code are the vertical mixing ratios of each considered species, which can then be directly read into a radiative transfer code to create a transmission spectrum. However, as chemical kinetics simulations are computationally intensive, it is challenging to produce a sufficiently fine grid of models for direct data interpretation (like atmospheric retrievals). Instead, kinetic models are generally used to identify trends or transitions

in chemistry, for example the occurrence of photochemically-produced SO_2 in hot Jupiters (Tsai et al., 2023), although early chemical kinetics retrievals have been employed in the analysis of hot Jupiters (Bardet et al., 2025).

In Chapter 6 I use the chemical kinetics code `VULCAN` (Tsai et al., 2017; Tsai et al., 2021) to simulate the atmospheric chemistry of temperate Jupiters; I describe the model setup and operation there.

2.4.3 Atmospheric retrievals

1D radiative transfer models can rapidly simulate exoplanet transmission spectra. This allows the use of inverse methods to interpret an observed transmission spectrum with a high dimensionality parametric model, and thereby extract posterior information on the model parameters. This technique is called an atmospheric retrieval, and while originally applied to remote sensing of Solar System atmospheres (e.g. Irwin et al., 2008), it has now become the premier method for data interpretation of exoplanet observations (e.g., Madhusudhan & Seager, 2009; Line et al., 2013; MacDonald & Madhusudhan, 2017; Mollière et al., 2019).

An atmospheric retrieval is fundamentally the inverse problem of the forward modelling techniques presented in 2.4.1. In forward modelling, a state vector of model parameters $\mathbf{x}$ is applied to a forward map (the radiative transfer model) F , to give rise to a synthetic observation $\mathbf{y}$:

$$\mathbf{y} = F\mathbf{x}. \quad (2.12)$$

The goal of a retrieval is to invert this problem, to find the model parameters $\mathbf{x}$:

$$\mathbf{x} = F^{-1}\mathbf{y}. \quad (2.13)$$

This problem is typically solved using Bayesian inference to determine the posterior probability distribution of the model parameters. This distribution, $P(\mathbf{x}|\mathbf{y})$, can be calculated via Bayes' theorem:

$$P(\mathbf{x}|\mathbf{y}) = \frac{P(\mathbf{y}|\mathbf{x})P(\mathbf{x})}{P(\mathbf{y})}, \quad (2.14)$$

where $P(\mathbf{x})$ is the prior probability distribution of model parameters $\mathbf{x}$, $P(\mathbf{y})$ is the model evidence $P(\mathbf{y}) = \int P(\mathbf{y}|\mathbf{x})P(\mathbf{x})d\mathbf{x}$, and $P(\mathbf{y}|\mathbf{x})$ is the likelihood function of $\mathbf{x}$. This is done by coupling a posterior sampling method, which explores the multidimensional parameter space of $\mathbf{x}$ based on a prior distribution, to a radiative transfer model, which computes a forward model $\mathbf{y}' = F\mathbf{x}$ which can be compared to the observation $\mathbf{y}$ to calculate the likelihood, $\mathcal{L}(\mathbf{x}) \equiv P(\mathbf{y}|\mathbf{x})$.

Older retrieval methods use a simple grid search (e.g. Madhusudhan & Seager, 2009) to sample the posterior distribution, which can be inefficient and restrictive as much of the computation is performed far from the solution. Another historical method is optimal estimation (e.g. Irwin et al., 2008), which is a direct inversion

of the forward map matrix. Optimal estimation uses iterative fitting to simultaneously minimize the distance between the model output and measured spectrum, and the state vector (model parameters) and a Gaussian prior vector. Optimal estimation is limited by the need to use Gaussian priors, which often is not appropriate for exoplanet observations where uninformative priors on reference radius, aerosols and chemical composition are desired. MCMC (see Sec. 2.3.3) and nested sampling (Skilling, 2004) are currently the dominant approaches for exploring the posterior parameter space using flexible priors.

There are many open source codes that use Bayesian inference with radiative transfer models to interpret exoplanet observations, for example **Poseidon** (MacDonald & Madhusudhan, 2017; MacDonald, 2023), **BeAR** (Kitzmann et al., 2020), and **PLATON** (Zhang et al., 2019). A catalogue of these codes is maintained by MacDonald & Batalha (2023). In this thesis, I use the retrieval capability of **petitRADTRANS** (Mollière et al., 2019; Nasedkin et al., 2024), which uses a nested sampling algorithm implemented using **MultiNest** (Feroz & Hobson, 2008) with **PyMultiNest** (Buchner et al., 2014).

Nested sampling

Nested sampling (Feroz & Hobson, 2008) is a Bayesian algorithm to sample the posterior probability distribution through complete exploration of the multidimensional parameter space and direct calculation of the Bayesian evidence $Z \equiv P(\mathbf{y}) = \int \mathcal{L}(\mathbf{x})P(\mathbf{x})d\mathbf{x}$.

Nested sampling operates by generating a hypercube of model parameters, consisting of N draws from the prior distributions (where N is known as the number of live points). For each step i , the live point with lowest likelihood, $\mathcal{L}_i$, is replaced by a new draw from the prior with higher likelihood. This results in a reduction in prior volume, with each rejected point assigned a weight w_i which is used to approximate the evidence $Z \approx \sum_i \mathcal{L}_i w_i$. This step is iterated until the local and global maxima are found, with the termination criterion set in **MultiNest** by reaching a certain precision in the Bayesian evidence. Running a nested sampling algorithm outputs a distribution of weighted posterior samples with their individual likelihoods, and the Bayesian evidence of the model, which can be used in model comparison. The best fitting (likelihood maximizing) $\mathbf{x}_{\mathcal{L}_{\max}}$ mode parameters, and a vector composed of the median of each parameters $\mathbf{x}_{\text{med}}$ can both be identified. The results of nested sampling are typically presented as a corner plot of posterior probability distribution functions, examples of which are used to present atmospheric retrieval results in Chapters 3, 4, and 5.

Model comparison

As nested sampling directly outputs the Bayesian evidence of any given model (that is, a given set of priors and forward model setup parameters), it is straightforward to statistically compare different models. This can be used to derive the evidence

supporting the inclusion or omission of certain free parameters. This is commonly employed to determine the statistical evidence supporting the presence of an individual molecule (or other opacity source), or determine the level of detail required in an aerosol or temperature-profile prescription. Retrieval model comparisons in this thesis use the following method, outlined below and adapted from Thorngren et al. (2025).

To compare between two models, A and B , we should consider the Bayes factor, $\mathcal{B}$, defined as the ratio of the evidences ($Z = P(y)$) of the two models:

$$\mathcal{B}_{AB} = 1/\mathcal{B}_{BA} = \frac{Z_A}{Z_B} = \frac{P(\mathbf{y}|A)}{P(\mathbf{y}|B)}. \quad (2.15)$$

The quantities that we care about is the probability of models A and B given the observation $\mathbf{y}$. In model comparison, we assume that these two models are the only possibilities to represent the data, and therefore both the priors and the posteriors must sum to one:

$$P(A) + P(B) = 1 \quad \text{and} \quad P(A|\mathbf{y}) + P(B|\mathbf{y}) = 1. \quad (2.16)$$

We also assign these two models equal prior probabilities, such that $P(A) = P(B)$, allowing us to rewrite the expression for the Bayes factor as follows:

$$\mathcal{B}_{AB} = \frac{P(\mathbf{y}|A)}{P(\mathbf{y}|B)} = \frac{P(A|\mathbf{y})P(B)}{P(B|\mathbf{y})P(A)} = \frac{P(A|\mathbf{y})}{1 - P(A|\mathbf{y})}, \quad (2.17)$$

by substituting in Eq. 2.14 and Eq. 2.16.

Now we can rearrange for $P(A|\mathbf{y})$, the quantity of interest effectively expresses the statistical probability of the data supporting model A over model B :

$$P(A|\mathbf{y}) = \frac{1}{1 + \mathcal{B}_{AB}}. \quad (2.18)$$

This probability can be used intuitively to determine which model should be favoured, and to what extent. Compared to simple approximations like the BIC (see Sec. 2.3.3), it properly penalizes increased model complexity and rewards improvements in goodness of fit, thanks to the Bayesian evidence calculation. However, this value should not be understood as a frequentist p -value. Nonetheless, it can be used to determine a significance (σ), frequently used to report the significance of detecting a molecule or other opacity source, by using the Bayes factor of models with or without that opacity source included. The inverse cumulative distribution function of the normal distribution (Φ^{-1}) can be applied to $P(A|\mathbf{y})$, just as one would with a p -value, to create an estimate of significance:

$$\sigma(A) = \Phi^{-1}(P(A|\mathbf{y})/2). \quad (2.19)$$

However, it is important to note that the Bayesian evidence is hugely sensitive to the choice of priors, which must be carefully considered when comparing models. Furthermore, this simple two-sided model comparison cannot alone be used

to determine the presence of a feature. Many reported significances from exoplanet observations use an approximation of Trotta (2008) to convert between Bayes factor and frequentist statistics. This conversion is miscalibrated, and can inflate the significance of detection by $\sim 1\sigma$ (Thorngren et al., 2025).

Chapter 3

Transmission Spectroscopy of HAT-P-44b

Declaration

This chapter is mainly a reproduction of the paper “LRG-BEASTS: Detection of sodium absorption and Rayleigh scattering in the hot Saturn HAT-P-44b” Claringbold et al. (2026b), with modifications to make my contributions clear and additional context on the LRG-BEASTS survey. I performed all the analyses and writing within this chapter. The observations for this analysis were taken by Ian Skillen and Peter J. Wheatley prior to the start of my PhD.

3.1 LRG-BEASTS

Optical transmission spectroscopy is a powerful technique that can measure the clouds and hazes present in exoplanet atmospheres (see Sec. 1.3.1). In order to investigate the diversity of aerosol properties across the hot gaseous exoplanet parameter space, a large sample size is required. This has the added benefit of providing optical transmission spectra that can both synergize with space-based infrared observations, as well as inform target selection for science cases which require the high precision detections facilitated by clear atmospheres.

The Low Resolution Ground-Based Exoplanet Atmosphere Survey using Transmission Spectroscopy (LRG-BEASTS; “large beasts”) was launched in 2016 to provide a large sample of hot Jupiters with homogeneously analysed optical transmission spectra. LRG-BEASTS has demonstrated the ability to achieve comparable precision on 4m-class telescopes to that obtained by 10-m class telescopes and HST. The LRG-BEASTS programme primarily makes use of data from the ACAM instrument of the 4.2 m William Herschel Telescope (WHT) at Roque de los Muchachos in La Palma, Spain, and the EFOSC2 (ESO Faint Object Spectrograph and Camera 2) instrument on the 3.6 m New Technology Telescope (NTT) at La Silla in Chile. Both of these instruments have since been removed from their respective telescopes, prior




Figure 3.1: Transmission spectra of nine exoplanets published as part of the LRG-BEASTS survey, demonstrating the diversity of features present in hot Jupiter optical spectra, and the potentially sub-scale height precision offered by 4-m class telescopes. The red lines indicate atmospheric models fit to the observed spectra, using a mixture of simple linear models (WASP-80b, WASP-52b, HAT-P-18b) and the best-fitting models from atmospheric retrievals (WASP-25b, HATS-46b, WASP-94Ab, WASP-130b, WASP-21b, WASP-39b). Figure assembled by James Kirk based on the works of Kirk et al. (2016), Kirk et al. (2017), Loudén et al. (2017), Kirk et al. (2018), Kirk et al. (2019), Alderson et al. (2020), Kirk et al. (2021), Ahrer et al. (2022), Ahrer et al. (2023a), and McGruder et al. (2023), used with permission.

to this a large number of exoplanet transit observations (of at least 40 exoplanets) have been acquired since 2016. Due to the great precision required for transmission spectroscopy, and the long observing duration required for a satisfactory quality light-curve, many of these observations are impacted by weather or technical difficulties.

Previous characterizations with LRG-BEASTS have detected hazes in HAT-P-18b (Kirk et al., 2017) and WASP-80b (Kirk et al., 2018), grey clouds in WASP-52b (Louden et al., 2017) and HATS-46b (Ahrer et al., 2023a), and sodium absorption and super-Rayleigh slopes in WASP-21b (Alderson et al., 2020) and WASP-94 Ab (Ahrer et al., 2023a). Combined analyses involving LRG-BEASTS data have aided in the interpretation of the atmospheres of WASP-39b (Kirk et al., 2019), WASP-103b (Kirk et al., 2021), WASP-25b (McGruder et al., 2023) and WASP-94Ab (Ahrer et al., 2024). These join results from other ground-based surveys using the Gran Telescopio Canarias (e.g. Parviainen et al., 2016; Murgas et al., 2017), Magellan/IMACS (ACCESS, e.g. Rackham et al., 2017; Espinoza et al., 2019), Gemini/GMOS (e.g. Huitson et al., 2017; Panwar et al., 2022), and VLT/FORS2 (e.g. Nikolov et al., 2016; Sedaghati et al., 2016; Spyrtatos et al., 2021).

The data reduction and analysis for LRG-BEASTS data is implemented using the now open-access **Tiberius** pipeline (Kirk et al., 2017, 2021)¹, which is now also used for JWST data. The full pipeline, as well as specific details for this kind of observation with the WHT telescope, is fully described in Kirk (2018). A concise overview of this workflow relevant to this thesis is outlined in Sec. 2.3.

3.2 The HAT-P-44 System

In this chapter, I use the ACAM instrument on the WHT to present the transmission spectrum of the inflated hot Saturn HAT-P-44b (Hartman et al., 2014), the tenth transmission spectrum published as part of the LRG-BEASTS programme. With a radius of $1.24^{+0.11}_{-0.05} R_{\text{Jup}}$, mass of $0.35 \pm 0.03 M_{\text{Jup}}$, and equilibrium temperature of 1110^{+50}_{-30} K (Hartman et al., 2014), HAT-P-44b is a close sibling to the well-characterized WASP-39b (the target of the JWST Early Release Science programme as well as a previous LRG-BEASTS target), which has a comparable radius of $1.27 \pm 0.04 R_{\text{Jup}}$, mass of $0.281 \pm 0.032 M_{\text{Jup}}$, and equilibrium temperature of 1166 ± 14 K (Mancini et al., 2018). The deep transit depth, low gravity, and high equilibrium temperature of HAT-P-44b gives it a strong expected signal of 288 ppm for a single scale height. HAT-P-44b has also been observed using the Exoplanet Transmission Spectroscopy Imager on the McDonald Observatory Otto Struve 2.1m telescope, with a reconnaissance spectrum obtained in eight wavelength bins (Oelkers et al., 2025). The variance of transit depth across the eight bins is used to suggest that HAT-P-44b may have a clear atmosphere.

¹<https://github.com/JamesKirk11/Tiberius>

3.3 Observations

Two transits of HAT-P-44b were observed on the nights of 2018 March 26 and 2019 March 18 using the ACAM instrument (Benn et al., 2008) on the WHT, providing low resolution spectroscopy over a wide wavelength range ($\sim 350 - 940$ nm) with a wide field of view (8 arcmin). The fast readout mode was used, with no order-blocking filter employed for either night. This is the same instrument setup for long-slit spectroscopy as Kirk et al. (2017, 2018, 2019) and Alderson et al. (2020). The volume phase holographic grism dispersion is 0.33 nm/pixel.

A 7.6 arcmin long, 40 arcsec wide slit was rotated to a position angle of 80.9 degrees to observe the target and a comparison star while avoiding differential slit losses. The same comparison, 2MASS J14124195+4701058, located 1.28 arcmin from HAT-P-44, was used for both nights. The comparison is both a good brightness ($\Delta V = 0.21$) and colour match ($\Delta(B - V) = 0.11$), and therefore enables effective differential spectroscopy.

For the first night, a long exposure time of 300 s was used due to the relative faintness of HAT-P-44 ($V_{\text{mag}} = 13.2$). The moon was at 77% illumination and 72 degrees from the target, setting during the post-transit observation. The first hour of the observation was affected by clouds and poor seeing (4"), and seeing remained high throughout the night (1.5 – 3"). There is a small gap after the fifth exposure, as guiding briefly failed due to clouds. The effective spectral resolution of this night is 12 nm pre-transit, and around 5 nm for the rest of the night, once the seeing had improved.

The second night of observation took place during nearly full moon (94% illumination at 60 degrees from the target) and had better seeing (1.3"), so a reduced exposure time of 180 seconds was chosen. Clouds affected 5 of the pre-transit frames, causing a dip in the flux and spike in the background. During the transit a sudden enhancement in the background can be seen localised to the blue end of the target spectrum that lasted for ~ 40 minutes, which can be seen in Fig. 3.2. I tentatively attribute this to indirect scattered moonlight, given its abrupt appearance and disappearance and the presence of the full moon. The effective resolution for the second night is 4 nm.

Both observations used fast readout mode and a smaller-than-standard window to minimise the overhead time to 6.4 s. For both nights, lamp and sky flats, biases, and HeAr arc frames were also taken. I did not use any flat frames in our final data reduction as I found that it increased the noise in our data, in line with previous analyses from LRG-BEASTS (e.g., Kirk et al., 2017; Alderson et al., 2020; Ahrer et al., 2023a), ACCESS (e.g. Rackham et al., 2017), VLT/FORS2 (e.g. Gibson et al., 2017), and Gemini (e.g. Huitson et al., 2017). The use of flat-fielding in ACAM transmission spectroscopy was explored by Kirk (2018), who found it was not possible to construct a well-illuminated flat field, and the use of flat fielding introduced red noise across multiple datasets. In addition, as 1620-5400 pixels are included in each wavelength bin of our final transmission spectrum, the impact of




Figure 3.2: Example ACAM frames demonstrating the used extraction (solid blue lines) and background (dashed blue lines) regions from the first night (first panel) and second night (second panel). I also observed a strong 2D enhancement on the detector in the second night for some of the transit (third panel), and highlighted with a difference frame between without-scattering and with-scattering frames with similar full width at half maximum (FWHM) for the second night in the final panel. The approximate wavelength of each Y pixel is indicated. Each X pixel spans 0.25 arcseconds.




Figure 3.3: Ancillary data for the first night of observation of HAT-P-44b. The blue crosses correspond to HAT-P-44b and the red crosses correspond to the comparison star. Each panel represents the variation of various quantities throughout the night, from top to bottom: target airmass, the shift in trace x and y centroid across the slit, the FWHM of the stellar lines, the normalized flux, the total sky background counts, and the white-light light-curve differential flux. The impact of clouds is evident as dips in the normalized flux in the fifth and final twelve exposures. The poor seeing at the beginning of the night is apparent in the FWHM plot, and the kink in the sky background curve near the end of the night is due to moonset.




Figure 3.4: Ancillary data for the second night of observation of HAT-P-44b. The blue crosses correspond to HAT-P-44b and the red crosses correspond to the comparison star. From top to bottom: target airmass, the shift in trace centroid across the slit, the FWHM of the stellar lines, the total sky background counts, and the white-light light-curve differential flux. The minimum in normalized flux immediately before transit is due to clouds, while I believe the differential enhancement in the background to be indirect scattered moonlight due to its abrupt appearance and disappearance.

pixel-to-pixel variations is negligible. I constructed a median-combined master bias frame and subtracted it from each science frame. The airmass, x pixel shift, y pixel shift, trace full width at half maximum (FWHM), normalized flux, sky background flux, and differential flux for each night are presented in Fig. 3.3 and Fig. 3.4 respectively. Manual guiding corrections were applied in the second night based on real-time spectral analysis, which is apparent as discontinuities in x and y positions in Fig. 3.4.

3.4 Data Reduction

I use the *Tiberius* pipeline for spectral extraction and light-curve fitting introduced in Kirk et al. (2017, 2021), as is standard for LRG-BEASTS data (e.g. Kirk et al., 2018, 2019; Ahner et al., 2022), and has also been used for JWST observations (e.g. Alderson et al., 2023; Rustamkulov et al., 2023; Moran et al., 2023; Kirk et al., 2024b). My analysis closely follows that of Alderson et al. (2020), with the key steps summarized below.

To extract the traces of HAT-P-44 and the comparison star, I fit the position of the trace with a quartic function in the dispersion direction and a Gaussian for the trace width, and place an aperture over it. To remove the sky background I subtract a polynomial fitted to background regions on either side of the trace on a row-by-row basis. I clip outliers in the sky background of three or more standard deviations and then refit the polynomial. I iterate the aperture width and background region and polynomial order, to minimise the rms of the residuals in the white-light light-curves. For the first night I use a wide aperture of 60 pixels (~ 15 arcsecs) full-width, with a large 20 pixel offset to a 55 pixel background region on either side of the target spectrum. For the second night I use a narrower 24 pixel (~ 6 arcsecs) full-width aperture and a 50 pixel background region offset by 12 pixels. The aperture is much wider in the first night because the seeing was much greater, and is a compromise between fitting the 40 pixel FWHM pre-transit, and the much lower 15 pixel FWHM for the rest of the night. In both cases I find that a quadratic fit for background subtraction minimises the white-light light-curve rms. As the seeing was highly variable on the first night, the chosen width is a compromise value to avoid systematics associated with a time-varying aperture width.

After spectral extraction, I remove cosmic rays from the spectra using a running median to identify outliers, replacing identified pixels with the mean of their neighbours. I also remove outliers directly from the original frame in the background region when performing background subtraction. I then resample the spectra onto consistent pixels in wavelength by fitting Moffat profiles to ten stellar absorption features spanning from 486 to 866 nm to calculate the pixel shifts of the spectrum (including stretching and compression), which I then smooth in time using a cubic function before resampling. I then use a single fit to 13 stellar and telluric lines from a reference spectrum in the middle of the night to assign a wavelength scale to




Figure 3.5: Flux ratio between the target and comparison (top panel), and normalized flux of the target (bottom panel, blue) and comparison (bottom panel, red) as a function of wavelength, for an example exposure. The wavelength bins used in the creation of the transmission spectra as displayed as black dotted lines, and the wavelengths of Na I and K I absorption are highlighted as green lines at 589 nm and 770 nm respectively. Additional bins excluded from the final analysis are indicated with grey dotted lines.

the aligned spectra. I permit differences in radial velocities between the target and comparison star.

A single bin spanning the entire wavelength range is used to define the white-light light-curve. I use 24 wavelength bins of the stellar spectra to define spectroscopic light-curves. I attempt to use constant 20 nm width bins, with minor modifications in width to avoid bin edges being located near absorption features in the stellar spectrum. I use three narrower bins of 10 nm (9 nm for the central bin) width to explore the core and wings of the Na I doublet feature at 589 nm. I choose to use a single 30 nm bin to cover the entire telluric O₂ A band feature at 762 nm, to avoid any bin edges falling in the wings of the feature. This bin happens to also include the K I doublet at 770 nm, which is found in the wing of the O₂ absorption. I present example stellar spectra from the first night annotated with the adopted wavelength bins in Fig. 3.5.

I create normalized differential white-light and spectroscopic light-curves by summing the flux within the relevant wavelength bins for the target, dividing by the corresponding flux of the comparison star, and normalizing by the median out-of-

transit flux.

3.5 Data Analysis

3.5.1 Transit model

I fit the differential light-curves using analytic light curves (Mandel & Agol, 2002) from the `Batman` package (Kreidberg, 2015a), with detrending based on physical basis functions to remove red noise (see Sec. 3.5.2). I use the Limb-Darkening Toolkit (LDTK) package (Parviainen & Aigrain, 2015) implemented through `Tiberius` to compute the limb-darkening coefficients using `PHOENIX` models (Husser et al., 2013), based on the following stellar parameters of HAT-P-44 from Hartman et al. (2014): $T_{\text{eff}} = 5295 \pm 100$ K, $\log g = 4.46 \pm 0.06$, and $\text{Fe}/\text{H} = 0.33 \pm 0.10$.

First I fit the white-light light curves for the planet-to-star radius ratio R_p/R_* , the system inclination i , the mid-transit time T_c , the ratio of semi-major axis to stellar radius a/R_* , and the first quadratic limb-darkening coefficient u_1 , using only wide uniform priors. I fix the period as ($P = 4.30119043$ d) based on the most precise literature value from Kokori et al. (2023). I perform fits using different limb-darkening parametrizations, including linear and quadratic limb-darkening laws with each parameter fixed or fitted (using the Kipping parametrization (Kipping, 2013) to avoid degeneracy when both u_1 and u_2 are fitted). I let the data inform us on the most appropriate level of detail by choosing the fit with the best BIC, which slightly favoured quadratic limb-darkening with u_2 fixed and u_1 fitted for the white-light light-curves in both nights, and significantly favoured a quadratic law with both u_2 and u_1 fixed for the spectroscopic light-curves over any other method.

I present the results of these white-light light curve fits in Table 3.1. All the parameters are consistent between the nights, except for u_1 , which is strongly discrepant between the nights. The value from the first night is consistent with the calculated value of $u_1 = 0.584 \pm 0.003$. I attribute the discrepancy to the scattered light feature seen during the second night in Fig. 3.2 and Fig. 3.4. For the spectroscopic light-curves I fix i , T_c , and a/R_* to the best-fitting values from the white-light light-curve of the first night data.

3.5.2 Light-curve fitting

For both nights, I fit the white-light and spectroscopic light-curves with the transit model multiplied by a systematics model, composed of the sum of polynomial functions of a range of physical parameters (basis functions), with the coefficients of these polynomial c_n becoming fitted parameters. I investigate all combinations of linear and quadratic basis functions including sky background, FWHM, and x and y pixel shifts, plus linear in time and airmass. I compare the resulting BIC to determine the best combination of detrending parameters. I choose the detrending parameters independently for the white-light and spectroscopic light-curves, letting the BIC guide my choice in both cases.

Parameter	Prior	Night 1 Fitted Value	Night 2 Fitted Value	Hartman et al. (2014)
R_p/R_*	$\mathcal{U}(k_0 - 3\sigma_k, k_0 + 3\sigma_k)$	$0.1367^{+0.0010}_{-0.0008}$	$0.1363^{+0.0011}_{-0.0010}$	0.1343 ± 0.0010
a/R_*	$\mathcal{U}(s_0 - 3\sigma_s, s_0 + 3\sigma_s)$	$11.95^{+0.13}_{-0.22}$	$11.91^{+0.20}_{-0.21}$	$11.50^{+0.50}_{-0.80}$
i ($^\circ$)	$\mathcal{U}(i_0 - 3\sigma_i, i_0 + 3\sigma_i)$	$89.2^{+0.5}_{-0.4}$	$88.9^{+0.5}_{-0.4}$	89.1 ± 0.4
u_1	$\mathcal{U}(u_{1,0} + 4\sigma_{u_1})$	0.58 ± 0.02	0.45 ± 0.02	
u_2	Fixed	0.0042	0.0042	

Table 3.1: Parameter values from white-light light-curve fits for both nights. Values for the planet-to-star radius ratio R_p/R_* , the ratio of semi-major axis to stellar radius a/R_* , and the system inclination i , are consistent between nights, while the values for the first quadratic limb-darkening coefficient u_1 are inconsistent, an effect I attribute to the scattering event during the transit in the second night. The calculated value for u_1 is 0.584 ± 0.003 , consistent with the first night. k_0 , s_0 and i_0 represent initial guessed values for R_p/R_* , a/R_* , and i respectively, with σ_k , σ_s , and σ_i representing the error bars used to define the prior ranges, taken from Hartman et al. (2014). The u_1 priors are determined with $u_{1,0}$ and σ_{u_1} which are calculated from the stellar parameters of Hartman et al. (2014) using LDTK (Parviainen & Aigrain, 2015).




Figure 3.6: Top panel: White-light light-curve of the first night of observation of HAT-P-44b (black data points), with an analytic transit model (grey dotted), polynomial systematics model (green), and the combined fit to the data (red). Middle panel: Detrending components contributing to the systematics model: quadratic in sky background (poly 1, blue) and quadratic in FWHM (poly 2, orange). Bottom panel: Residuals from the combined systematics and transit model.

In the first night I find that the detrending parameters resulting in the best BIC for the white-light light curve are quadratic sky and quadratic FWHM. These parameters are able to account for the variable seeing and remove the kink in the out-of-transit baseline corresponding to moonset, which I correct for using quadratic




Figure 3.7: Spectroscopic light-curve fits from the first night of observation of HAT-P-44b (black data points) with the combined transit model and polynomial de-trending fit (red) with different y offsets. Residuals to these fits are displayed in the right-hand panel, with the central wavelength and residual rms displayed above each fit.

sky background detrending. I clip the fifth and final four data points from the light-curves, as they are strongly affected by clouds, visible as dips in the normalized flux in Fig. 3.3. For the spectroscopic light-curves I find that a simpler detrending with linear sky and linear FWHM gives the best BIC, by accounting for the impact of variable seeing. I present the white-light light-curve fit in Fig. 3.6 and the spectroscopic fits in Fig. 3.7. In order to assess the impact of red noise, I include plots in Fig. 3.8 demonstrating how the root mean square (RMS) residuals of the spectroscopic light-curve fits vary as a function of binning. For all but one bin I find that the residuals are consistent with $1/\sqrt{N}$ behaviour of white noise. The only bin to deviate from this behaviour is the bin centred at 759 nm, which contains the telluric O₂ feature. This bin does not impact my interpretation of the transmission spectrum.

In the second night, there is an interval of spatially-localised background enhancement on the detector during the transit, which is clearly visible in Fig. 3.2 and Fig. 3.4. With careful background subtraction I can eliminate most of its impact on the differential white-light light-curve, but there is still a small impact on the residuals. However, there is a much stronger impact on the spectroscopic light-curves, given the complicated 2D structure of the scattering feature (see Fig. 3.2), differentially affecting the blue wavelengths. Since the scattered light feature only effects one of the stars, power of differential spectro-photometry is lost at these times. Using best BIC detrending, I find that linear basis functions of time, airmass, and sky background are favoured for the white-light light-curve, while the combination of just linear airmass and linear sky is favoured for the spectroscopic light-curves. I also clip the final hour of observation, which is affected by moonset and poor seeing. I also attempted fits where the drop in flux due to clouds for five exposures just before transit are both included and cropped, finding the inclusion has a negligible effect on the resulting transmission spectrum. The results of our white-light and spectroscopic light-curve fits are illustrated in Fig. 3.9 and Fig. 3.10 respectively. I find generally poorer fits in the lower signal-to-noise bins at the edge of the detector (not shown), so I restrict our transmission spectrum to a wavelength range of 495–874 nm.

The fit for each light-curve was performed using MCMC from the `emcee` package (Foreman-Mackey et al., 2013), with an initial run with 5000 steps and 12 walkers per parameter, followed by a production run with a further 5000 steps. 10000 steps were used for both the initial run and production run for fitting the white-light light-curves. A Nelder-Mead algorithm (Nelder & Mead, 1965) was employed for initial parameter optimization, after which photometric uncertainties were rescaled.




Figure 3.8: Root mean square of the spectroscopic light-curve fit residuals from the first night as a function of binning in time, compared to the $1/\sqrt{N}$ expected from white noise. For each wavelength bin the measured noise is consistent with the uncertainties scaled to the binning, which is ~ 200 ppm at 10 points per bin. The single exception to this is the wavelength bin centered at 7590 Å, which contains the telluric O₂ absorption feature. This single bin does not drive our interpretation of the transmission spectrum.




Figure 3.9: Top panel: White-light light-curve of the second night of observation of HAT-P-44b, with an analytic transit model (grey dotted), polynomial systematics model (green), and the combined fit to the data (red). Middle panel: Detrending components contributing to the systematics model: linear in time (poly 1, blue), linear in sky background (poly 2, orange), and linear in airmass (poly 3, green). Bottom panel: Residuals from the combined systematics and transit model.

3.5.3 Transmission spectrum

I construct my highlighted transmission spectra for both nights using the transit depth of the spectroscopic light-curves, presented in Table 3.3 and Fig. 3.11. The recovered spectra cover a wavelength range of 495–874 nm using ~ 20 nm bins, with evidence for absorption at the wavelength of the sodium doublet in night 1. When comparing the spectra, a strong disagreement can be seen, with night 2 displaying broad absorption at blue wavelengths which cannot readily be attributed to any atmospheric feature. Given the strong scattering feature that impacted a large amount of the in-transit spectrum (Fig. 3.2), I believe that despite my best efforts in background subtraction and detrending, I was not able to remove its impact on the transmission spectrum. I therefore repeated the fits, cropping all exposures impacted by the scattering event. This results in cropping out almost half the transit including the entire ingress. Using the best BIC detrending for this cropped light-curve does not result in an improved transmission spectrum (not shown). I therefore conclude that the need to either detrend out or crop the scattering event results in a loss of information in the spectroscopic light-curves making it impossible to recover an accurate transmission spectrum for this night of observation. I therefore proceed using only the transmission spectrum from the first night.




Figure 3.10: Spectroscopic light-curve fits from the second night of observation of HAT-P-44b (black data points) with the combined transit model and polynomial detrending fit (red) with different y offsets. Residuals to these fits are displayed in the right-hand panel, with the central wavelength and residual rms displayed above each fit.




Figure 3.11: Transmission spectrum of HAT-P-44b from spectroscopic light-curve fits from the first night (black) and the second night (blue), using the refitted limb darkening coefficients. The wavelength of Na I absorption at 589 nm is highlighted with a yellow dotted line. The bottom panel shows both nights on the same axes for comparison.

Night 1 Spectroscopic Fit	k	δ BIC
Linear Sky, Linear FWHM	4	+0
Linear Sky, Linear FWHM, Linear Airmass	5	+70
Linear Sky, Linear Airmass	4	+76
Quadratic Sky, Linear FWHM	5	+87
Linear Sky	3	+166
Linear Time	3	+189
Quadratic Sky, Quadratic FWHM	6	+205
Quadratic LD, Both Fixed	4	+0
Quadratic LD, u_2 Fitted	5	+158
Quadratic LD, Both Fitted	6	+295
Linear LD, Fixed	4	+237
Linear LD, Fitted	5	+163

Table 3.2: The combined BICs of different systematic models and limb-darkening parametrizations on the spectroscopic light-curve fits to the first night, relative to the best BIC model (δ BIC = 0), with the number of free parameters (k).

Wavelength (nm)	Night 1 $(R_p/R_*)^2$ (%)	Night 2 $(R_p/R_*)^2$ (%)
495–515	$1.90702^{+0.02719}_{-0.02735}$	$1.71500^{+0.03586}_{-0.03606}$
515–535	$1.85703^{+0.02050}_{-0.02079}$	$1.72197^{+0.03021}_{-0.03010}$
535–555	$1.90239^{+0.02205}_{-0.02189}$	$1.80706^{+0.02550}_{-0.02566}$
555–575	$1.88760^{+0.02254}_{-0.02238}$	$1.86445^{+0.02342}_{-0.02344}$
575–585	$1.87296^{+0.02731}_{-0.02746}$	$1.90324^{+0.02997}_{-0.02993}$
585–594	$1.97394^{+0.0282}_{-0.02797}$	$1.84642^{+0.03402}_{-0.03346}$
594–604	$1.83302^{+0.02914}_{-0.02899}$	$1.72752^{+0.02967}_{-0.02963}$
604–624	$1.83031^{+0.01627}_{-0.0162}$	$1.76061^{+0.02372}_{-0.02401}$
624–644	$1.85276^{+0.0193}_{-0.01925}$	$1.77854^{+0.0221}_{-0.02202}$
644–664	$1.87838^{+0.01728}_{-0.01737}$	$1.74364^{+0.02106}_{-0.02114}$
664–680	$1.84468^{+0.02384}_{-0.02377}$	$1.79627^{+0.01976}_{-0.01973}$
680–704	$1.84604^{+0.01549}_{-0.01551}$	$1.79091^{+0.01825}_{-0.0185}$
704–724	$1.86088^{+0.01756}_{-0.01739}$	$1.76789^{+0.02185}_{-0.0219}$
724–744	$1.86729^{+0.02033}_{-0.02020}$	$1.83242^{+0.02159}_{-0.02165}$
744–774	$1.84829^{+0.02302}_{-0.02309}$	$1.78420^{+0.02375}_{-0.02380}$
774–794	$1.87709^{+0.02640}_{-0.02632}$	$1.76778^{+0.02412}_{-0.02433}$
794–814	$1.86943^{+0.02896}_{-0.02912}$	$1.82434^{+0.02651}_{-0.02646}$
814–834	$1.90112^{+0.03290}_{-0.03265}$	$1.76635^{+0.02449}_{-0.02458}$
834–854	$1.92327^{+0.03572}_{-0.03592}$	$1.85979^{+0.02652}_{-0.02674}$
854–874	$1.89349^{+0.03795}_{-0.03812}$	$1.83917^{+0.03012}_{-0.03047}$

Table 3.3: Tabulated transmission spectra from both nights of observation, plotted in Fig. 3.11.

I achieve a mean transit depth precision of 246 ppm in the first night, which is less than the 288 ppm expected transit signal of one scale height of the planetary atmosphere, permitting detailed atmospheric characterization.




Figure 3.12: Transmission spectrum of HAT-P-44b from the first night with simple parametric model fits, including a linear model (blue) with the sodium bin masked, and a linear model plus a Gaussian feature with fixed mean centred on the sodium feature and fitted width (red) fitted to the whole spectrum.

3.5.4 Testing the sodium detection

With a strong feature at the wavelength of Na I absorption in our spectrum from the first night, I conduct simple tests to estimate the significance and reliability of the detection. Following Nikolov et al. (2016) and Alderson et al. (2020), I fit a linear polynomial to the transmission spectrum with the sodium bin masked and obtain a 3.9σ significance based on the difference between the fit and the sodium bin (see Fig. 3.12). I also fit the entire spectrum with three simple models: a linear polynomial; a linear polynomial with a Gaussian feature fixed at 589.3 nm with fixed 6nm width and fitted amplitude; and a Voigt feature fixed at 589.3 nm with free width, amplitude, and Lorentzian half-width at half-maximum. The Voigt feature represents giving the sodium freedom to be pressure and temperature broadened, while the Gaussian feature represents a narrow feature with the width set by the spectral resolution. The linear plus fixed Gaussian sodium feature gives the best reduced χ^2_ν of 1.33, representing a 3.5σ improvement on the linear fit based on the change in χ^2_ν . I believe the lack of improvement of the Voigt fit to be representative of an inability to directly detect broadening as a result of the broad bins set by the high FWHM of our observations. Switching from a linear polynomial to a simple Rayleigh scattering slope (Pont et al., 2008) does not improve the fit because it cannot reproduce the overall redward shape of the spectrum (with a positive temperature). I plot a selection of our parametric fits in Fig. 3.12.

I conducted a further investigation into the presence of sodium by fitting light-curves with various bin widths centred on the Na I feature at 589 nm. This




Figure 3.13: Transit depth of a binned light-curve centred on the sodium absorption feature at 589 nm as a function of bin width in the first night. These results are consistent with a feature with the width of the spectral resolution (which varies from 5 to 11 nm throughout the night), with decreasing transit depth at increasing bin width above this resolution. In the second night, this trend is much less clear, likely caused by the equally high transit depth in an adjacent bin.

can be used to identify the width of the feature and test that it is centred at the correct wavelength. For a narrow feature, the expected behaviour is a sharp decrease in transit depth with increasing bin width, approaching the surrounding continuum level (see Alderson et al., 2020). This is the behaviour I observe in the first night, with a more gradual decrease in width than in Alderson et al., remaining above the continuum level at 50 nm bin width, likely due to the strength of the feature I observe. The spectral resolution deteriorates to ~ 12 nm initially due to the poor seeing, just above the 9 nm bin width corresponding to the highest transit depth. As can be seen in Figure 3.13, the transit depth decreases with bin width above this, as expected. The effect of the stellar sodium feature being smoothed out into adjacent bins due to the strongly varying spectral resolution throughout the night is particularly apparent in the narrowest bins. The FWHM detrending does a good job of eliminating this effect in wider bins, but struggles in the smallest bin.

I repeated the analysis on the K I feature at 768 nm, but see no evidence for increased absorption. The close proximity of the strong telluric O₂ feature can result in large systematic errors when using narrow bins, particularly with variable FWHM. When using a 20 nm bin, as is used for the rest of the spectrum, the transit depth is consistent with the adjacent bins (see Fig. 3.11).

3.6 Atmospheric Retrieval

3.6.1 Retrieval setup

I carried out an atmosphere retrieval analysis using the `petitRADTRANS` 2.7.7 package (Mollière et al., 2019; Nasedkin et al., 2024) to interpret the transmission spectrum from the first night. `petitRADTRANS` uses Bayesian nested sampling with `MultiNest` (Feroz et al., 2009) and radiative transfer with correlated- k opacity tables to compute and fit synthetic transmission spectra to the observed spectrum.

I performed a collection of retrievals with different aerosol parameterizations and including various opacity sources. In all of our models I use an isothermal atmosphere and equilibrium chemistry, with a fixed solar C/O ratio of 0.55 (Asplund et al., 2009), a Gaussian prior on the gravity with mean and standard deviation from Hartman et al. (2014), and uniform priors on the planetary radius at a reference radius of 0.01 bar (1.0–1.6 R_J), isotherm temperature (700–2600 K), and metallicity ($[M/H] = -1$ –2.5). I include combinations of opacity from H_2O (Polyansky et al., 2018), CH_4 (Yurchenko et al., 2017), and Na (Piskunov et al., 1995), as well as collisionally induced absorption from H_2 – H_2 and H_2 –He and Rayleigh scattering from H_2 and He. I parameterize the chemical abundances using metallicity, which is scaled based on the solar abundances of Asplund et al. (2009). I fix the C/O ratio, since the wavelength range covered by our spectrum is insufficiently sensitive to relative abundances of carbon- and oxygen-bearing species (with CH_4 opacity in particular only present in the reddest two bins and overlapping with H_2O absorption).

For the aerosol parametrizations, I ran retrievals including opacity from haze, modelled by the expected Rayleigh scattering for H_2 and He multiplied by a fitted haze factor, and grey cloud, modelled using an opaque grey cloud deck at a fitted cloud-top pressure, as well as retrievals including neither source of opacity as well as both. By including both sources, I am effectively assuming that aerosol opacity in this case is driven by some combination of condensate clouds with large particle sizes fixed below some altitude, as is observed in spectral muting of absorption features in transmission spectra, plus some photochemically produced aerosols which are sufficiently small to be well modelled by Rayleigh scattering.

In previous optical transmission spectra, scattering slopes even steeper than Rayleigh have been observed (Sing et al., 2011b; Alderson et al., 2020; Luque et al., 2020; Ahrer et al., 2022), often attributed to photochemical hazes (Kawashima & Ikoma, 2019; Ohno & Kawashima, 2020). Therefore I also run a retrieval in which I permit an arbitrary scattering slope of the form $\lambda = \kappa^{-\alpha}$, with the gradient permitted to vary between 4 and 20 (it is fixed to 4 for Rayleigh scattering). Both the gradient α and the multiplicative factor S are fitted parameters in this retrieval.

3.6.2 Retrieval results

The retrieval with the highest Bayesian evidence includes absorption from Na, H_2O , and CH_4 , and no aerosol opacity from cloud or haze. This very simple setup gives us a




Figure 3.14: Transmission spectrum from the first night with (black data points) and the best-fitting `petitRADTRANS` model from the highest evidence atmospheric retrieval (blue). The spectrum demonstrates absorption due to sodium (the sharp winged feature at 589 nm), Rayleigh scattering without haze-enhancement (the upwards slope at the blue end of the spectrum), and H₂O and CH₄ absorption (the upwards slope at the red end of the spectrum, and the small feature at 730 nm). I also plot the best-fitting model from the second highest evidence atmospheric retrieval (orange), which also includes a grey cloud layer as a free parameter, with the 2σ minimum cloud-top pressure at ~ 1 mbar. Wavelengths corresponding to absorption features of Na, H₂O, and CH₄ are shaded and labelled.

concordant spectrum using opacity from Na, H₂O, and CH₄ with a single atmospheric metallicity of $7^{+16}_{-5} \times$ solar, and gives tight constraints of the temperature of 790 ± 30 K. This is cooler than the 1100 K equilibrium temperature, as is commonly observed in 1D isothermal retrievals (MacDonald et al., 2020; Welbanks & Madhusudhan, 2022). I plot the best fitting model from this retrieval in Fig. 3.14 and the posteriors in Fig. 3.15.

The Bayesian evidence is nearly identical when I include opacity from a grey cloud deck, which yields an identical best-fitting model with essentially no muting from clouds (see Fig. 3.14). This places a 2σ upper limit on the cloud-top pressure of ~ 1 mbar, ruling out high altitude clouds such as those responsible for muting the optical transmission spectrum of HATS-46b (Ahner et al., 2023a). The near equal evidence of this retrieval demonstrates that it is agnostic to the inclusion of mid-altitude (> 1 mbar) cloud. However, the posterior distributions in Fig. 3.16 highlight that a degeneracy is present in the temperature, metallicity, and cloud-top pressure posteriors when clouds are considered. There are two modes to this retrieval, with the best-fitting mode being the same as the retrieval without clouds, and an additional mode including muted water features due to clouds at ~ 10 mbar and unconstrained metallicity and a higher temperature of ~ 1600 K.

Comparing the evidence with the other retrieval setups in Table 3.4 shows a slight preference for not including a Rayleigh haze enhancement, and notably dis-




Figure 3.15: Corner plot of posterior distributions of retrieval parameter values obtained by fitting the transmission spectrum from the first night with the best evidence `petitRADTRANS` atmospheric model using a Rayleigh scattering slope, equilibrium chemistry, and an isothermal atmosphere, as depicted in blue in Fig. 3.14.

favouring a free (super-Rayleigh) haze slope (with a Bayes factor of 25). Using Bayesian model comparison between the best-fitting retrieval without sodium absorption (including both clouds and haze), and the corresponding model with sodium absorption included, reveals weak evidence for sodium absorption using a retrieval analysis, corresponding to a significance of 1.4σ .

I summarise our Bayesian evidences, best-fit reduced χ^2_ν values, degrees of freedom, and detection significances of all our retrievals and simple model fits in Table 3.4. Repeating the atmospheric retrieval while omitting the data point centred at 759 nm containing the telluric O_2 feature resulted in no change to the posteriors or best fit spectrum, highlighting that any red noise associated with this feature does not impact my interpretation of the transmission spectrum. I also include the full set of fits to the transmission spectrum in Fig. 3.17.




Figure 3.16: Corner plot of posterior distributions of retrieval parameter values obtained by fitting the transmission spectrum from the first night with the second best evidence **petitRADTRANS** atmospheric model, including an isothermal atmosphere, equilibrium chemistry, a Rayleigh scattering slope, and a grey cloud deck, as depicted in orange in Fig. 3.14.

3.7 Discussion

3.7.1 The detection of sodium

My analysis using simple parametric model fits results in a 3.9σ detection of Na I absorption in HAT-P-44b, using our sub-scale height precision transmission spectrum from the first night of observation (see Sec. 3.5.4 for details). The observed sodium feature spans 5 scale heights of transit depth even with a broad 9 nm bin.

Previous ground-based optical transmission spectra have given diverse results for sodium absorption, including non-detection of sodium (e.g., Kirk et al., 2017; Ahrer et al., 2023a), detection of the narrow core of a sodium feature (e.g., Redfield et al., 2008; Chen et al., 2018; Alderson et al., 2020), and detection of the pressure-broadened wings of the sodium feature (Nikolov et al., 2016, 2018). This is typically

Retrieval / Simple Model	$\ln Z$	B_m	χ_ν^2	DOF	Rejection σ
No clouds	-20.1	-	1.20	16	
Clouds	-20.1	1.0	1.32	15	-
Clouds + Haze	-20.4	1.6	1.45	14	-
Clouds + Haze, without Na	-21.7	5.0	1.89	14	1.38
Clouds + super-Rayleigh	-23.2	25.1	1.68	13	2.07
Linear Fit	-	-	2.01	18	3.53
Linear + Gaussian Fit	-	-	1.33	17	1.52
Linear + Voigt Fit	-	-	1.50	15	-

Table 3.4: The Bayesian evidence ($\ln Z$), Bayes factor (B_m), reduced χ_ν^2 , degrees of freedom (DOF), and significance of being rejected relative to the no clouds retrieval reference model (σ reject), of the different retrievals and simple model fits applied to the transmission spectrum. For retrievals, the significance of rejection is converted from the Bayes factor via the p-value using $p = \frac{1}{1+B_m}$, while a likelihood ratio test using the χ_ν^2 values is used for the simple model fits. The rejection σ is not given for retrievals where the Bayes factor is too small for it to be meaningful.




Figure 3.17: Transmission spectrum from the first night with (black data points) and the full set of best-fitting `petitRADTRANS` model from our atmospheric retrieval analysis, plus the best-fitting simple parametric fit (red dotted line). The plotted retrieval fits include a model with no aerosols (dark blue), a model with free grey cloud-top pressure (orange), a model with free cloud-top pressure and a free enhancement of the Rayleigh scattering slope representing haze particles (light blue), and a model with free cloud-top pressure and a free (permitting super-Rayleigh) scattering slope (magenta).

attributed to varying cloud altitude suppressing the feature strength and obscuring the feature wings or core with a scattering slope. Detection of a strong sodium feature is usually taken as evidence for a clear atmosphere, especially when obvious pressure-broadened wings are also detected.

The low and variable resolution of the first night (5–11 nm) hinders direct detection of pressure-broadening in HAT-P-44b, with no improvement in χ_ν^2 from using a linear plus Voigt fit over a linear plus fixed-width Gaussian fit for the whole

spectrum. However, the very strong transit depth even with a relatively broad 9 nm bin requires a very strong and hence probably broad sodium feature. I emphasize that the 9 nm bin used in this analysis is much broader than the typical 3 nm bin sodium detections (e.g. Alderson et al., 2020). As highlighted in my retrieval analysis in Sec. 3.6, the observed features imply a haze-free atmosphere, with additional absorption from some combination of H_2O and CH_4 , and optionally mid-altitude clouds (2σ upper limit of 1 mbar on the cloud altitude), to accurately fit the shape of the spectrum. The presence of a strong sodium feature and the detectability of the weak molecular features at these wavelengths strongly fit the narrative of a cloud-free atmosphere.

With the detection of a strong sodium feature, HAT-P-44b joins a small group of exoplanets with potentially cloud-free upper atmospheres: WASP-39b (Fischer et al., 2016; Nikolov et al., 2016; Sing et al., 2016), WASP-96b (Nikolov et al., 2018; McGruder et al., 2022), WASP-62b (Alam et al., 2021), and WASP-94 Ab (Ahrer et al., 2022). These planets have proven ideal for high SNR infra-red follow-up with JWST (Ahrer et al., 2023b; Alderson et al., 2023; Feinstein et al., 2023; Radica et al., 2023; Rustamkulov et al., 2023; Taylor et al., 2023), which can provide exceptional chemical and structural detail.

Despite my detection of sodium, I do not detect potassium in HAT-P-44b, although I stress that due to the proximity of potassium to the telluric O_2 feature, the observation is poorly sensitive to potassium absorption. Evidence for potassium depletion can be seen in the transmission spectra of WASP-21b (Alderson et al., 2020), WASP-96b (Nikolov et al., 2018; Taylor et al., 2023) and WASP-62b (Alam et al., 2021). In contrast there are strong detections of potassium in WASP-39b (Fischer et al., 2016; Feinstein et al., 2023). The cause of this discrepancy is unknown; the depletion of potassium could be due to some primordial depletion, or sequestering of potassium by condensates like KCl (sufficiently deep in the atmosphere to not give rise to cloud impacting the transmission spectrum).

3.7.2 Aerosols and evidence for methane and water

Outside of the sodium feature, there is some structure in the observed spectrum, including both increasing transit depth towards both the blue and red ends of the spectrum (see Fig. 3.14). Blueward slopes in the optical spectra of hot Jupiter atmospheres are common, and can be evidence for scattering slopes, or stellar inhomogeneity arising from activity. The gradient of the slope combined with the strength of the sodium feature can explain the aerosol properties of the atmosphere. A slope gradient consistent with Rayleigh scattering is typical of small particles, whether that is the H_2 gas, haze, or mineral condensate particles (e.g. Lecavelier des Etangs et al., 2008a). A flat spectrum and suppressed sodium absorption indicates large particles in the form of mineral clouds (e.g. Ackerman & Marley, 2001; Fortney, 2005). A steeper slope can arise from a photochemical haze (Kawashima & Ikoma, 2019; Ohno & Kawashima, 2020), a combination of cloud and haze layers

(Pont et al., 2013), sulfide clouds (Pinhas & Madhusudhan, 2017), or stellar activity such as unocculted spots or occulted faculae (McCullough et al., 2014; Oshagh et al., 2014; Espinoza et al., 2019). Redward slopes are typically the result of molecular absorption, like H_2O , CH_4 , TiO , and VO .

I use atmospheric retrievals to investigate these features in HAT-P-44b, allowing me to investigate opacity from Rayleigh scattering (with or without haze enhancement), a grey-cloud deck, and molecules. The best-fit model from my highest evidence retrieval gives an improvement in goodness of fit equivalent to being favoured at 1.5σ compared to my parametric fit (χ^2 of 19.2 with 16 degrees of freedom for the retrieval versus χ^2 of 22.6 with 17 degrees of freedom for the linear plus Gaussian fit).

In this retrieval, the bluewards slope is attributed to Rayleigh scattering from H_2 gas, while the increased absorption at the red end is attributed to molecular absorption from H_2O and CH_4 , with relative abundances in chemical equilibrium consistent with the scaled sodium abundance relative to solar. Combined, these features point towards a cool (800 K), aerosol-free, super-Solar metallicity ($7^{+16}_{-5} \times \text{Solar}$) upper atmosphere. This remains the favoured explanation even when a grey cloud deck is permitted, although adding a grey cloud deck into the model adds degeneracies between metallicity, cloud-top pressure, temperature, and radius, complicating the picture and preventing a direct measurement of metallicity or temperature. The addition of a Rayleigh opacity enhancement term, representative of haze particles, is very slightly disfavoured in the Bayesian evidence, while inclusion of a free (i.e. potentially super-Rayleigh) scattering slope parameterization is moderately disfavoured.

Stellar activity can also cause bluewards slopes in optical transmission spectra that mimic Rayleigh scattering slopes (McCullough et al., 2014; Oshagh et al., 2014). This is notable in planets around more active stars like HD 189733b, where the super-Rayleigh slope can also be explained with occulted faculae or unocculted spots. Stellar activity is a potential alternative explanation for the slope in HAT-P-18b (Fournier-Tondreau et al., 2024), but HAT-P-44 is not an active star, with a median $\log R_{\text{HK}} = -5.247 \pm 0.058 \pm 0.10$ (Hartman et al., 2014) and a low $v \sin i = 0.2 \pm 0.5 \text{ km s}^{-1}$, indicating low levels of chromospheric activity. I note however that it is possible for calcium emission to be masked by a calcium outflow from some planets (Haswell et al., 2020), and a low $v \sin i$ can potentially arise from low stellar inclination.

My observation of features in the transmission spectrum of HAT-P-44b appears consistent with the results of the reconnaissance spectrum obtained by Oelkers et al. (2025), which suggests HAT-P-44b may be clear based on the variance in its transit depth in different wavelength bins.




Figure 3.18: Transmission spectrum of HAT-P-44b compared against the transmission spectrum of WASP-39b, a similar inflated hot Saturn. The WASP-39b spectrum includes VLT, HST, and WHT data from Fischer et al. (2016); Nikolov et al. (2016); Sing et al. (2016); Wakeford et al. (2017a) as combined by Kirk et al. (2019). The spectra have similar sodium absorption and continuum levels.

3.7.3 HAT-P-44b in context

With the combination of a strong sodium feature and hints of H_2 gas Rayleigh scattering and H_2O and CH_4 absorption, I infer a relatively haze- and cloud-free upper atmosphere, placing a 2σ upper limit on a grey cloud-top pressure at 1 mbar, and find the scattering slope strength to be consistent with 800 K H_2 gas scattering. I therefore conclude that the parts of the atmosphere probed using transmission spectroscopy are likely lacking in small particle aerosols, as well as large particle aerosols above 1 mbar.

In this regard, the optical transmission spectrum of HAT-P-44b has similarities to those of WASP-21b (Alderson et al., 2020; Chen et al., 2020), WASP-39b (Fischer et al., 2016; Nikolov et al., 2016; Sing et al., 2016; Kirk et al., 2019), WASP-94 Ab (Ahrer et al., 2022, 2024), WASP-96b (Nikolov et al., 2018; McGruder et al., 2022), and WASP-62b (Alam et al., 2021). A comparison of these planets’ planetary parameters and the features of their optical transmission spectra are presented in Tab. 3.5.

WASP-39b and WASP-21b are notable points of comparison for HAT-P-44b, with similar equilibrium temperatures and surface gravities, with WASP-39b also possessing a similar host star to HAT-P-44. In terms of its optical spectrum, HAT-P-44b sits closer to WASP-39b in aerosol properties, with a stronger sodium feature and weaker slope than WASP-21b indicative of less haze. I depict the spectra of both HAT-P-44b and WASP-39b in Fig. 3.18 for comparison. Like WASP-21b, HAT-P-44b shows no evidence of potassium in its transmission spectrum, although both data sets are ground-based and are potentially affected by the same issue with the telluric O_2 feature. WASP-39b on the other hand does have potassium absorption detected

	HAT-P-44b	WASP-39b	WASP-21b	WASP-96b	WASP-62b	WASP-94Ab
T_{eq} (K)	1110^{+50}_{-30}	1116 ± 14	1340 ± 30	1285 ± 40	1475^{+25}_{-20}	1604^{+25}_{-22}
Gravity (ms^{-2})	5.7 ± 1.1	4.3 ± 0.6	5.6 ± 0.3	8.3 ± 0.1	7.4 ± 1.4	3.8 ± 0.4
Stellar T_{eff} (K)	5300	5400	5800	5500	6200	6200
Scattering slope	Rayleigh	Rayleigh	Super-Rayleigh	Rayleigh	Not observed	Super-Rayleigh
Haziness	Weak	Weak	Moderate	Weak	Weak	Moderate
Sodium	Strong	Strong, broad	Strong, narrow	Strong, broad	Strong, broad	Strong, broad
Potassium	None*	Present	None*	Tentative	None	Unknown

* From only ground-based data, so may be impacted by the telluric O₂ feature.

Table 3.5: Comparison of hot Jupiters with evidence for relatively cloud-free atmospheres from their optical transmission spectra due to strong sodium absorption, scattering slopes, or both. I use moderate haziness to refer to planets with strong (relative to the sodium core, or obscuring the sodium wings) or super-Rayleigh scattering slopes, and weak haziness to planets with clear sodium wings and weak or non-detected scattering slopes. Strong haziness would refer to planets with strong Rayleigh or super-Rayleigh scattering slopes that lead to no or weak sodium features like HAT-P-18b (Kirk et al., 2017). I summarise the results of the following atmospheric characterizations: HAT-P-44b (this paper), WASP-39b (Fischer et al., 2016; Nikolov et al., 2016; Sing et al., 2016; Kirk et al., 2019; Feinstein et al., 2023), WASP-21b (Alderson et al., 2020; Chen et al., 2020), WASP-96b (Nikolov et al., 2018; McGruder et al., 2022; Radica et al., 2023; Taylor et al., 2023), WASP-62b (Alam et al., 2021), and WASP-94Ab (Ahrer et al., 2022, 2024). The discovery papers for the included exoplanets are as follows: HAT-P-44b (Hartman et al., 2014), WASP-39b (Faedi et al., 2011), WASP-21b (Bouchy et al., 2010), WASP-96b (Hellier et al., 2014), WASP-62b (Hellier et al., 2012), WASP-94Ab (Neveu-VanMalle et al., 2014).

using space-based observations (Fischer et al., 2016; Sing et al., 2016; Feinstein et al., 2023).

Gao et al. (2020) predict that silicates are the dominant aerosol species responsible for muting spectral features at these temperatures, while photochemical hazes are implicated in super-Rayleigh slopes (Kawashima & Ikoma, 2019; Ohno & Kawashima, 2020). The similarities between the optical spectra of the sibling planets HAT-P-44b and WASP-39b are therefore suggestive that these processes are driven mainly by these planetary properties: gravity, temperature, and stellar host. On the other hand, HATS-46b ($T_{\text{eq}} = 1100$ K, $g = 5.1 \text{ ms}^{-2}$, stellar $T_{\text{eff}} = 5400$ K) is another close sibling to both WASP-39b and HAT-P-44b yet has a cloudy optical transmission spectrum (Ahrer et al., 2023a). Looking at the hotter, higher gravity side of hot Saturns, a featureless spectrum indicative of clouds is present for WASP-49b (Lendl et al., 2016), which is otherwise a close analogue ($T_{\text{eq}} = 1400$ K, $g = 7.8 \text{ ms}^{-2}$, stellar $T_{\text{eff}} = 5600$ K) to WASP-96b in these parameters.

The impact of both temperature and gravity on cloudiness has been explored with both computer modelling and statistical analysis of observations, with generally higher temperature and gravity associated with less cloud (Heng, 2016; Stevenson, 2016; Fu et al., 2017; Dymont et al., 2022; Estrela et al., 2022). The growing roster of optical transmission spectra make it apparent that while there is some island of relatively cloud-free hot Saturns, with temperatures 1000–1500 K and surface gravities

$4\text{--}8\text{ ms}^{-2}$, whether a given exoplanet exists in this island cannot be predicted from these properties alone. This could be explained by the chemical inventories of those planets, or complicated dependencies on various planetary and stellar parameters involving photochemistry, 3D dynamics, and aerosol microphysics (see Gao et al., 2021, and references within).

3.7.4 Tests of background subtraction

At the suggestion of the reviewer, I implemented the following tests to demonstrate that the transmission spectrum of HAT-P-44b is not overly sensitive to choices in the background subtraction. I use a well-established method of background subtraction following the methodology of previous LRG-BEASTS studies (e.g., Kirk et al., 2017; Alderson et al., 2020; Ahrer et al., 2022), using a row-by-row polynomial fit within a relatively narrow background region on either side of the aperture. I iterated values for the width of the aperture, background region, and the background offset, and the polynomial order used to fit the background, to minimize the RMS residuals of an initial white-light light-curve fit. I also visually examine the post-subtraction frames and background model to ensure the subtraction method was well-behaved. My final background subtraction for the first night, as described in Section 3.4, uses an aperture width of 60 pixels, an offset of 20 pixels, a background region of 55 pixels, and a linear polynomial to fit the background (the background regions are indicated in Fig. 3.2). The pre- and post-subtraction frames, and the fitted background model, for a single frame are depicted in Fig. 3.19. This background subtraction can be seen to perform well, with no residual sky lines visible. The curvature of the sky lines is negligible (less than one pixel across the regions of interest); if this were not the case there would be near-horizontal residuals of the sky lines in the post-subtraction image. The diagonal feature in the comparison star frame is a bad column which doesn't extend into the background or source region. Background subtraction reveals the presence of very faint background stars (too faint to be seen in the raw images) in the post-subtraction background region of the target star, as well as the very tail of the PSF of another background star in the edge of the comparison star background region. Collapsing the flux vertically reveals these stars amount to $\sim 0.1\%$ of the total background, resulting in a negligible offset in the transit depth that is constant in time. I also note the poor performance of the background subtraction for y-pixel values above 1600, but highlight that I do not use the flux from these wavelengths in either the white-light or spectroscopic light-curves. The initial white-light light-curve fit for this subtraction (using the same white-light light-curve fitting and detrending described in Section 3.5) yields a residual RMS of 324 ppm.

To demonstrate the robustness of our spectrum to other choices in background subtraction methodology, I repeated the full host of reduction and analysis steps with alternative choices in subtraction methods, using both the white-light light-curves and resultant transmission spectra to compare between methods. I highlight that three alternative choices: using a wider background region (up to and

including the entire detector), using a wider mask to mitigate for any stellar PSF within the background region, and masking the background stars in the primary reduction, all give rise to consistent transmission spectra, and perform equally or worse in terms of residual RMS.

The level of the background far from the source is uninformative on the level of the background within the aperture because the background varies in the cross-dispersion direction in a relatively complicated manner. This motivates the usage of a relatively narrow background region, where the background level has a simpler functional form. This typically means that a higher-order polynomial is required when a wider background region is used. This is confirmed when I test wider background regions, for example with a wider region of 100 pixels, the subtraction fares poorly with a linear fit (white-light light-curve RMS of 442 ppm, 188 ppm worse than my primary reduction), but improves notably with a quadratic fit (white-light light-curve RMS of 348 ppm, 24 ppm worse). When using a background fit to the entire detector (excluding 10 pixels at each edge and masking both stars), this behaviour holds, with linear, quadratic, and cubic fits all performing very badly. With a quartic fit to the entire detector, the performance approaches that of the primary subtraction, with a white-light light-curve RMS of 370 ppm, 46 ppm worse than the primary reduction. I present the pre- and post-subtraction frames and fitted background model with this method in Fig. 3.20. I highlight again the lack of any diagonal features from curvature of the sky lines, and the presence of the same background stars in the post-subtraction frame that appeared in the primary subtraction.

To verify that starlight from the wings of the PSF bleeding into the background region is not affecting the results, I repeated the background subtraction with a broader background offset of 30 pixels and reduce the background width to 45 pixels (so no new pixels are considered). This has virtually no impact on the post-subtraction data, with a white-light light-curve fit residual RMS of 323 ppm (a 1 ppm improvement). I note that the white-light R_p/R_* is effectively identical as well (72 ppm deeper, compared to an error of 1300 ppm), demonstrating that the choice of a wider mask has a negligible impact on the transit depth. I also repeated my primary reduction but masking the pixels including the background stars centered at 67 pixels to the left of, and 85 pixels to the right of, the target star. These stars are very faint (about 35 and 20 counts at their peak pixels respectively), and masking them has negligible impact again, with an unchanged white-light light-curve residual of 324 ppm, and a change in R_p/R_* of only 138 ppm shallower.

All of these different background subtraction methods produce highly concordant transmission spectra, presented in Fig. 3.21, demonstrating that my results are robust to all of these decisions. Therefore, the interpretations of this chapter, including sodium absorption and the presence of blue-ward and red-ward slopes, are not sensitive to the specific choices made in the background subtraction. I note that the use of the whole detector gives rise to more noise in both white-light and spec-




Figure 3.19: Image of the detector region selected for background subtraction in the primary reduction for a single frame of the first night of observation, centred on the fitted trace of the target star (top row) and comparison star (bottom row). The first column shows the region before background subtraction, the second column shows the region after background subtraction, and the fitted linear background model is shown in the third column. The aperture region used for spectral extraction is between the solid red lines, with the background region used for fitting a linear polynomial is without the dotted lines, between which is the masked offset region. In the post-subtraction image for the target star you can see very faint background stars not visible in the raw image, and in the post-subtraction image of the comparison star you can see a bad pixel column, which is fully contained within the mask for the entire night.

troscopic light-curves and requires a higher order polynomial to fit the background, confirming that the background flux far from the trace is not helpful in determining the shape of the background within the aperture.

3.8 Conclusions

I present the low-resolution optical transmission spectrum of the inflated hot Saturn HAT-P-44b, the tenth planet in the LRG-BEASTS survey. Making use of the ACAM instrument on the WHT and using data from one of two nights of observation (the other being compromised by systematics from scattered light), I obtain a transmission spectrum in the range 495–874 nm with an average transit depth precision of 246 ppm using ~ 20 nm bins, which is $0.85\times$ the signal strength of one scale height. Using a comparison star and detrending based on physical parameters including sky background, airmass, and FWHM, I am able to effectively remove systematic effects from variable seeing, scattered light, moonset, and atmospheric extinction, with the median RMS-noise equal to the expected photon noise.

I detect strong sodium absorption spanning ~ 5 scale heights in the spectrum of HAT-P-44b with 3.9σ confidence using a broad 9 nm bin. There is also evidence for Rayleigh scattering from H_2 gas and absorption due to H_2O and CH_4 in our highest evidence retrieval, consistent with a relatively cloud-free, 800 K atmosphere with a super-solar metallicity ($7^{+16}_{-5}\times$ solar). A super-Rayleigh slope indicative of strong photochemical haze like those observed in WASP-21b and WASP-94Ab is disfavoured at 2.1σ significance, while the 2σ upper limit cloud-top pressure is placed at ~ 1 mbar.

In this regard, the optical transmission spectrum of HAT-P-44b resembles that of WASP-39b, a close sibling in surface gravity, equilibrium temperature, and stellar host. While it is compelling to suggest that this similarity in planetary properties is the cause of these planets’ similar aerosol properties, it is worth noting that another close sibling in temperature, gravity, and stellar host, HATS-46b, shows evidence for clouds in the transmission spectrum. The physics and chemistry determining aerosol properties is therefore evidently more complex than a simple function of these three planetary parameters, despite HAT-P-44b apparently joining a group of similar 1000–1500 K hot Saturns with relatively cloud-free atmospheres.

The presence of identifiable features in the atmosphere of HAT-P-44b makes it an ideal candidate for further atmospheric characterization, particularly high SNR infra-red follow-up with JWST, as has been successfully done for other relatively cloud-free hot Saturns WASP-39b and WASP-96b. Extending the transmission spectrum further into the blue would also permit more detailed analysis of the scattering slope and any potential impact of stellar activity.




Figure 3.20: Image of the entire detector selected for background subtraction in an alternative reduction for a single frame of the first night of observation. The first column shows the region before background subtraction, the second column shows the region after background subtraction, and the fitted quartic background model is shown in the third column.




Figure 3.21: Transmission spectrum from the primary data reduction from the first night (black data points), compared to alternative methods using different background subtraction techniques. The same spectral features are present in each of the reductions, including sodium absorption, and the red-ward and blue-ward slopes associated with CH_4 and H_2O absorption, and rayleigh-scattering respectively.

Chapter 4

BOWIE-ALIGN: HAT-P-30b

Declaration

This chapter is mainly a reproduction the paper “BOWIE-ALIGN: Sub-solar metallicity and C/O ratio atmosphere of the misaligned hot Jupiter HAT-P-30b” Clar- ingbold et al. (2026a), with modifications to make my contributions clear, as well as additional detail on a **Tiberius** data reduction that I performed for HAT-P-30b but did not include in the paper, as it was consistent with the other **Tiberius** reduc- tion. I overall led the analysis and interpretation of these data, including writing the majority of the text (not including listed sections specific to individual reductions and models). Contributions from my collaborators are from Chloe E. Fisher (CEF), who contributed the **BeAR** atmospheric retrieval analysis (Section 4.4.3), Eva-Maria Ahrer (EA), who provided the **Eureka!** data reduction (Section 4.3.1), James Kirk (JK), who included the **Tiberius** data reduction featured in the paper and used for atmospheric modelling and interpretation in this chapter (Section 4.3.2), and Daniel P. Thorngren (DPT), who ran interior structure models (Section 4.5). These col- laborators are referred to by their initials in the following sections. I contributed a **Tiberius** data reduction as the lead investigator for this planet, while JK provided a **Tiberius** data reduction as part of a homogeneous analysis for all of the planets. As the spectrum are consistent and use very similar methodology, I utilize his analysis as the dominant **Tiberius** spectrum for all modelling and interpretations purposes.

4.1 BOWIE-ALIGN

In principle, we can use the elemental abundances of a planet’s composition to infer where a planet formed in the protoplanetary disc (Öberg et al., 2011) (see Sec. 1.1.4). In particular, the C/O ratio can be used to suggest where a planet formed relative to the ice lines of H_2O , CO_2 , and CO .

In practice, relating the atmospheric composition for a single planet to its for- mation history challenges this simple picture, with numerous physical processes ob- scuring the relationship between formation models and observed atmospheres. This is because, fundamentally, modelling planet formation includes a high dimensional-




Figure 4.1: Schematic demonstrating the mechanisms of hot Jupiter migration tested by the BOWIE-ALIGN programme, reproduced from Kirk et al. (2024a). The left diagram shows the initial migration scenario for a hot Jupiter, which then proceeds through one of the two mechanisms depicted in the right diagram: mechanism i) where migration is completed through the disc prior to dispersal, resulting in an aligned hot Jupiter with a composition reflecting the entire disc; mechanism ii) where migration occurs after disc dispersal due to a perturbation onto an eccentric orbit before undergoing tidal circularization resulting in a misaligned hot Jupiter with a composition reflected only the outer disc.

ity of poorly constrained parameters, including disc temperature profile, dust-to-gas ratio, and composition of the protoplanetary disc (e.g., Mollière et al., 2022). This can be additionally complicated by including physical processes important during formation, such as relative accretion of gas and solid (Espinoza et al., 2017), the evolution of ice lines (Morbidelli et al., 2016; Owen, 2020), the drift of volatile-carrying solids (Booth et al., 2017; Schneider & Bitsch, 2021a), and the trapping of volatiles within ice (Ligterink et al., 2024), as well as the complexity arising due to migration, as planets can accrete in multiple environments through their evolution. Observations of protoplanetary discs have also revealed them to be diverse objects, with their radial composition varying between host stars (Law et al., 2021a; Law et al., 2021b). Furthermore, directly relating a planet’s atmospheric composition to its bulk elemental ratios is non-trivial, with potential impacts from cloud formation (Helling et al., 2016) or enrichment of an atmosphere uncoupled from the interior (Müller & Helled, 2024).

As demonstrated by Penzlin & Booth et al. (2024), while any single planet is a poor tracer of planet formation, we can use populations of exoplanets with similar dynamical properties implying shared formation histories to test whether atmospheric composition traces planet formation. In particular, Penzlin & Booth et al. (2024) show how planets which underwent disc-migration should be chemically distinct from planets that underwent disc-free migration based on their elemental ratios, with the precise nature of this distinction depending upon the behaviour of accreted silicates and the dominant carbon-bearing species within the disc. We can leverage the two distinct populations of aligned hot Jupiters, that likely migrated through the disc, and misaligned hot Jupiters, that likely underwent high-eccentricity migration

after disc dispersal (Rasio & Ford, 1996; Wu & Murray, 2003; Muñoz et al., 2016). Comparing these two dynamically distinct populations we can test the predictions of Penzlin & Booth et al. (2024), depicted in Fig. 4.1.

This is precisely the goal of the BOWIE-ALIGN (BOWIE¹ – A spectral Light Investigation into gas Giant origiNs) survey. The BOWIE-ALIGN collaboration originates from a JWST Cycle 2 programme (GO 3838; PIs: Kirk & Ahrer) designed to spectroscopically characterizing eight planets: four aligned hot Jupiters believed to have migrated through the disc, and four misaligned hot Jupiters believed to have migrated via high-eccentricity migration (Kirk et al., 2024a). All eight planets orbit stars above the Kraft Break ($T_{\text{eff}} \gtrsim 6100$ K), an observed shift in rotation rates attributed to much thinner surface convective zones than found in cooler stars, in turn causing less efficient magnetic breaking (Kraft, 1967; Beyer & White, 2024). As a result of this less efficient breaking, realignment of hot Jupiters due to tidal interactions with the star is unlikely to occur above the Kraft break, fitting the observation that hot Jupiters around cooler stars tend to have low stellar obliquities (Winn et al., 2010; Albrecht et al., 2012). This ensures the aligned planets in the BOWIE-ALIGN sample are not aligned due to re-alignment of the star, and are instead likely aligned due to migration through the disc.

4.1.1 Predictions

The BOWIE-ALIGN hypothesis test leverages the planet formation models of Penzlin & Booth et al. (2024), where the calculations are tailored for the hotter discs of the F stars included in the BOWIE-ALIGN sample. The model used simulates the final assembly of gas giant planets through pebble, planetesimal, and gas accretion from an initial seed of super-Earths, tracking the elemental composition of accreted solid and gaseous material. A total of 1.8×10^5 planets are simulated to span the parameter range of disc accretion rate, disc temperature, the ratio of dust to gas accretion, initial formation locations of planets, and the amount of planetesimals in the disc. The high-eccentricity (i.e. misaligned) hot Jupiters are identified as those with final separations > 1 au (Muñoz et al., 2016), to be transported into the inner system by some later, un-modelled mechanism, while the disc migration (i.e. aligned) hot Jupiters are those with final orbits < 0.1 au. Planets between these separations are excluded, including warm Jupiters which underwent only weak migration. The primary result of these tests are that in the fiducial models, where the accreted dust is assumed to be destroyed and accreted in the atmosphere, the disc migrated planets have a lower C/O ratio than the high-eccentricity migrated planets when the metallicity is super-solar (see also Madhusudhan et al., 2014), as well as a general anti-correlation between C/O ratio and metallicity (recreating the findings

¹BOWIE originally stood for Bristol Oxford Warwick Imperial Exeter, which were the original institutions of the BOWIE+ collaboration; BOWIE+ is no longer considered an acronym, having since incorporated collaborators from many more institutions




Figure 4.2: Diagram of trends in planet formation aimed to be probed in the BOWIE-ALIGN programme, reproduced from Penzlin et al. (2024). Blue is used to represent parts of the parameter space occupied by misaligned planets, while orange is used for misaligned planets. The specific values vary based on various modelling assumptions, while maintaining trends.

of Cridland et al., 2019)

Penzlin & Booth et al. (2024) also highlight that an inverted scenario can occur when silicate grains from dust accretion are permitted to rain-out from the atmosphere, as predicted below an equilibrium temperature of ~ 1700 K by 2D cloud modelling (Powell & Zhang, 2024). Observations of spectral muting attributed to silicate clouds, as well as direct detection of SiO_2 clouds by Grant et al. (2023b) suggest that silicate rain-out from the atmosphere is either incomplete, or these kind of hot Jupiters atmospheres are fully mixed. In this silicate rain-out scenario the trend in the fiducial model reverses, with the C/O ratio of disc migration planets becoming higher than that of high-eccentricity migration planets. These two scenarios are used as the two end-case scenarios for C/O ratio differences between the populations to be probed by the BOWIE-ALIGN programme. A rough overview of these regions of compositional parameter space that could lead to an understanding of planet formation are presented in Fig. 4.2.

4.1.2 Sample

As described in Kirk et al. (2024a), the sample of BOWIE-ALIGN planets was drawn from a list of 16 hot Jupiters ($M_P > 0.3 M_J$, $R_P > 0.5 R_J$, $1000 < T_{\text{eq}} < 2050$ K) with precisely measured obliquities ($\sigma(|\lambda|) < 15^\circ$) such that they could definitely be




Figure 4.3: The BOWIE ALIGN sample of hot Jupiters, reproduced from Kirk et al. (2024a). The top-panel shows the sky-projected obliquity of as a function of equilibrium temperature for each planet. The bottom-panel shows the effective temperature as a function of metallicity for each host star, highlighting that they are all above the Kraft Break defined by Spalding & Winn (2022). The open circles correspond to planets from the GO 3838 programme, while the crosses are observed as part of other JWST programmes.

identified as either misaligned ($|\lambda| > 45^\circ$) or aligned ($< 30^\circ$). The upper temperature bound was set to avoid thermal dissociation of H_2O , observed in WASP-19b ($T_{\text{eq}} = 2077 \pm 34 \text{ K}$, Bernabò et al. 2024) which would complicate inferences of the oxygen abundance and therefore C/O ratio. To estimate the appropriate sample size to test this prediction, Kirk et al. performed atmospheric retrievals on simulated JWST transmission spectroscopy observations, using synthetic transmission spectra assuming equilibrium chemistry with elemental ratios obtained from the models of Penzlin & Booth et al. (2024), to estimate the posterior precision on C/O and metallicity for each planet. Nested sampling was then used to determine the significance of detecting a difference in C/O ratio between the aligned and misaligned planets, using samples of the best four, eight, and 16 planets. This was tested for different formation modelling assumptions (with and without silicate rainout), fixed metallicities ($3\times\text{solar}$ and $30\times\text{solar}$), and with and without clouds at 1 mbar. A sample of eight planets (four misaligned and four aligned) was found to be sensitive to differences in C/O ratio at 2σ except in scenarios without silicate rainout at $3\times\text{solar}$ metallicity. Kirk et al. found diminishing returns from using a larger sample, due to the worse signal-to-noise of the additional planets included. The final sample of 8 planets is depicted in Fig. 4.3, leveraging observations of WASP-94 Ab (GO 3154), WASP-17b (GTO 1353), and HD 149026 (GTO 1274) from other programmes.

WASP-15b, HAT-P-30b, TrES-4b, KELT-7b, and NGTS-2b were selected as the five targets for observation in GO 3838, using the NIRSpec instrument of JWST with the G395H filter to obtain a broadband transmission spectrum from $2.8\text{--}5.2 \mu\text{m}$, sensitive to absorption arising for key carbon- and oxygen-bearing species CO_2 , CO , H_2O , and CH_4 , required for determination of the metallicity and C/O ratio. All of these observations use the BOTS mode with the 2048 subarray and F290LP filter. The number of groups per integration was adjusted for each observation to fill 80% of the full well. Each visit was the length of a full transit, plus four hours to provide sufficient out-of-transit baseline and additional flexibility for scheduling and relaxation of systematic effects.

For each planet, at least two independent reductions using different open-source pipelines are used to analyse the data, which are then interpreted with at least two independent retrieval analyses. An end-to-end analysis with unchanged setup parameters using the *Tiberius* pipeline (see Sec. 2.3) and atmospheric retrievals with *petitRADTRANS* (see Sec. 2.4) is used to ensure a homogeneous comparison can be made between members of the sample.

4.2 The HAT-P-30 system

In this chapter, I present the JWST NIRSpec/G395H transmission spectrum of the misaligned hot Jupiter HAT-P-30b from $2.8\text{--}5.2 \mu\text{m}$ as part of the BOWIE-ALIGN survey.

HAT-P-30 b is a misaligned hot Jupiter with a measured obliquity of $70.5^{+2.9}_{-2.8}^\circ$

(Cegla et al., 2023). HAT-P-30 b, also known as WASP-51 b, has a mass of $0.746 \pm 0.021 M_{\text{Jup}}$ (Bonomo et al., 2017), a radius of $1.42 \pm 0.03 R_{\text{Jup}}$, equilibrium temperature of $1630 \pm 42 \text{ K}$ (Blažek et al., 2022), and an orbital period of 2.8106013 ± 0.0000006 days (Ivshina & Winn, 2022). I include a full summary of our adopted system parameters in Table 4.1. Like all the other planets in the BOWIE-ALIGN sample, HAT-P-30 b orbits a star above the Kraft break (Kraft, 1967).

Table 4.1: System parameters for HAT-P-30 b. References are for: [1] Johnson et al. (2011), [2] Bonomo et al. (2017), [3] Blažek et al. (2022), [4] Ivshina & Winn (2022), and [5] Cegla et al. (2023).

Parameter	Value	Reference
<i>Stellar parameters</i>		
Mass, M_* ($M_{\odot}$)	1.242 ± 0.041	[2]
Radius, R_* ($R_{\odot}$)	1.215 ± 0.051	[2]
Effective Temperature, T_{eff} (K)	6304 ± 88	[1]
Age (Gyr)	$1.0^{+0.8}_{-0.5}$	[2]
Surface gravity, $\log g$ (cgs)	4.36 ± 0.04	[2]
Metallicity [Fe/H] (dex)	0.13 ± 0.08	[2]
<i>Planetary parameters</i>		
Mass, M_p (M_{Jup})	$0.746^{+0.020}_{-0.021}$	[2]
Radius, R_p (R_{Jup})	1.417 ± 0.033	[3]
Equ. Temperature, T_{eq} (K)	1630 ± 42	[3]
Gravity, (ms^{-2})	9.2 ± 0.5	
Orbital period, P (days)	2.8106013 ± 0.0000006	[4]
Obliquity, λ (degrees)	$70.5^{+2.9}_{-2.8}$	[5]

4.3 Observations

One transit of HAT-P-30b was observed with JWST on 10 March 2024, 22:20:47 UT as part of the GO 3838 programme, using the NIRSpec instrument (Jakobsen et al., 2022) in Bright Object Time Series (BOTS) mode, with the G395H grating, F290LP filter, and SUB2048 subarray. This setup provides spectroscopy from $2.8 - 5.2 \mu\text{m}$ at a resolution of $R \sim 2700$, with a gap in coverage from $3.72 - 3.82 \mu\text{m}$ due to a physical gap between the NRS1 and NRS2 detectors. The number of groups per integration was set to 26, with a total of 925 integrations over 6.26 hours, including 2.25 hours during transit. A nearby faint star (2MASSJ08154583+0550218) was used for target acquisition in WATA mode with the SUB32 subarray and CLEAR filter.

4.4 Data reduction

My collaborators and I performed three independent data reductions, as prescribed by the BOWIE-ALIGN data analysis strategy found in Kirk et al. (2024a). We used two pipelines, **Eureka!** (Bell et al., 2022) and **Tiberius** (Kirk et al., 2017, 2021), for the data reduction, broadly following the methodology for other BOWIE-ALIGN

targets WASP-15 b (Kirk et al., 2025) TrES-4 b (Meech et al., 2025), and KELT-7 b (Ahrer et al., 2025b), with exact details of the spectral extraction and light curve fitting detailed below.

4.4.1 Eureka!

EA conducted a data reduction of HAT-P-30 b using the open-source `python` package **Eureka!** (Bell et al., 2022), which has been used extensively for analysing JWST exoplanet transmission spectra (e.g., Ahrer et al., 2023b; Moran et al., 2023; Alderson et al., 2024; Wallack et al., 2024; Xue et al., 2024; Teske et al., 2025).

Light curve extraction

EA followed previous BOWIE-ALIGN analyses when conducting her **Eureka!** analysis. She started with the uncalibrated files and run Stage 1 and 2 of **Eureka!** which is wrapped around the `jwst` pipeline (version 1.12.2, CRDS context pmap: 1253), with the following common modifications: a correction factor was applied to the `jwst` superbias (group 1, smooth, window length of 30 pixels), the `jump_rejection_threshold` was increased to 10.0σ (e.g., see Alderson et al., 2023), a group-level column-by-column background subtraction was run (with an outlier rejection threshold of $3\times$ median), and the `photom_step` was skipped. Note that in order to compute the bias scale factor and run the $1/f$ background subtraction the trace was masked using 10 pixels around the central trace.

In Stage 3, the stellar spectrum was extracted for each integration. A constant column-by-column background subtraction (excluding the 8/10 pixel area on both sides from the central trace pixel for NRS1/NRS2, respectively) was applied using a 5-sigma threshold for outlier rejection along both the time and spatial axes. The stellar spectrum was extracted using optimal spectral extraction (Horne, 1986) with an aperture half-width of 4 pixels. Columns were further masked manually where the count is $> 15\sigma$ from a rolling mean of 20 pixels across the frame.

The extracted stellar spectra were binned in Stage 4 following the common BOWIE-ALIGN grid, at $R=100$ and $R=400$. Outliers $> 5\sigma$ were masked on the rolling median (25 pixels) of the binned light curves.

Light curve fitting

The broadband white light curves for NRS1 and NRS2 were fitted first following common practice, using the transit depth R_p/R_* , the scaled stellar radius a/R_* , the mid-transit time T_0 , and inclination i as free parameters. EA further fitted for a baseline offset and a linear slope. The orbital period and eccentricity of HAT-P-30 b were fixed, to 2.8106 days (Ivshina & Winn, 2022) and 0, respectively, and the limb-darkening was fixed using the quadratic limb-darkening law and parameters u_1, u_2 . The limb-darkening values were computed for each bin in **Eureka!**'s Stage 4 using




Figure 4.4: White-light light-curves for HAT-P-30b for each detector, NRS1 (green) and NRS2 (magenta) from the *Eureka!* data reduction (left) and the residuals from the light curve fits (right), offset for clarity.

ExoTiC-LD (Grant & Wakeford, 2024) and the 3D Stagger grid (Magic et al., 2015), based on the stellar parameters from Table 4.1.

The best-fit parameters were derived using the *Python* MCMC sampler *emcee* (Foreman-Mackey et al., 2013), using 50 walkers and 3000 steps (of which 1000 were discarded as burn-in steps). The retrieved parameters for HAT-P-30 b’s orbit based on the white light curve from NRS1 and NRS2 are in Table 4.2, with the light curve fits presented in Fig. 4.4. For the spectroscopic NRS1 and NRS2 fits, EA then fixed the orbital parameters (a/R_* , i , T_0) to the best-fit values by the NRS1 or NRS2 white-light light-curves, respectively. The limb-darkening parameters were fixed to those calculated for each spectroscopic bin. Therefore, each spectroscopic light-curve fit included three free parameters: the transit depth R_p/R_* and the parameters describing the baseline offset and linear slope term of the light curve. The resulting transmission spectra at $R = 100$ and $R = 400$ are displayed in Fig. 4.6.

EA also performed a simultaneous fit to the white-light light-curves from both detectors from which is derived the combined fit system parameters presented in Table 4.2. Using these system parameters to generate a transmission spectrum in the same manner causes negligible differences in the spectrum when compared to that generated with the individual system parameters for each detector, except for a minor vertical offset (with an average of 32 ppm for NRS1, 26 ppm for NRS2) with no slope or change of shape.

4.4.2 Tiberius

We also conducted data reductions of HAT-P-30 b using *Tiberius* (Kirk et al., 2017, 2021), an open-source package for data reduction and light-curve fitting originally developed for ground-based exoplanet atmosphere transmission spectroscopy, now used extensively for JWST observations (e.g., Rustamkulov et al., 2023; Moran et al., 2023; Kirk et al., 2024b). JK and I both provided *Tiberius* data reductions, with JK using identical methodology to previous LRG-BEASTS analyses for homogeneous

comparison, while I used reduction parameters optimised for the observation of HAT-P-30 b, described below.

Light curve extraction

For JK’s **Tiberius** reduction, he followed an identical process as used in our BOWIE-ALIGN analyses of WASP-15b (Kirk et al., 2025) and TrES-4b (Meech et al., 2025), and used the same pipeline version (**Tiberius**: v1.0.4, **jwst**: v1.8.2) and calibration reference files. A detailed explanation of this process is detailed in Kirk et al. (2025). In short, he performed stage 1 using a modified version of the **jwst** pipeline, skipping the **jump_step** and adding a custom 1/f correction. Stage 2 is performed by **Tiberius** after the custom bad pixel flagging and correction procedure (detailed in Kirk et al., 2025). Standard aperture photometry was performed with an aperture full width of 8 pixels and an additional background subtraction at this stage, using a linear polynomial fitted to each pixel column after masking the 22 pixels centered on the stellar trace.

My **Tiberius** reduction followed the same steps as detailed above, with the following modifications. I used a wider aperture full width of 10 pixels, and a narrower mask of 18 pixels for background subtraction. I also used the most up-to-date reference files for the observation of HAT-P-30 b. I also used a different bad pixel mask, using a row-by-row outlier subtraction rather than the median absolute deviation of the entire frame. More detail is provided in Sec. 2.3.

To construct our light curves, the same two binning schemes are used as in the other BOWIE-ALIGN papers (Kirk et al., 2025; Meech et al., 2025; Ahrer et al., 2025b), namely a lower resolution $R = 100$ wavelength grid and a higher resolution $R = 400$ wavelength grid. Our NRS1 light curves span a wavelength range of 2.75–3.72 μm with NRS2 spanning 3.82–5.18 μm .

Light-curve fitting

For our light-curve fits, both JK and I followed an identical procedure to that detailed in Kirk et al. (2025). We fitted our light-curves with a quadratically limb-darkened analytic transit light-curve (implemented through **batman**, Kreidberg 2015b) multiplied by a linear-in-time polynomial. For our white-light light-curve fits, the free parameters are the planet-to-star radius ratio (R_p/R_*), the planet’s inclination (i), the semi-major axis to stellar radius ratio (a/R_*), the time of mid-transit (T_0), and the two parameters of a linear polynomial systematics model (c_1 , c_2). We both held the period fixed to 2.81060126 d (Ivshina & Winn, 2022) and the eccentricity to 0². We fixed the limb darkening coefficients to the values found using the Stagger grid of 3D stellar atmosphere models (Magic et al., 2015) as computed by **ExoTiC-LD** (Grant & Wakeford, 2024) using the stellar parameters given in Table 4.2.

²JK also tested fixing eccentricity to 0.035 according to the results of Johnson et al. (2011) but found that this led to a negligible difference in the resulting transmission spectra with a median difference of 1.5 ppm.




Figure 4.5: Comparison of transmission spectra of HAT-P-30b obtained using **Tiberius**, the result of two analyses independently performed by me (ABC) and JK. Given the concordance between the results of these spectra, I favoured the homogeneous analysis produced by JK for further analysis, referred to simply as the **Tiberius** reduction for the rest of the chapter.

Table 4.2: The resulting system parameters from our fits to the individual JWST NIRSpec/G395H NRS1 and NRS2 white-light light-curves, the system parameters from the combined fit to both the NRS1 and NRS2 data (**Eureka!** only), and the weighted mean of both individual NRS1 and NRS2 fits (**Tiberius** only).

Pipeline	Instrument	T_0 (BJD)	R_p/R_*	a/R_*	i ($^\circ$)
Eureka!	NRS1	$2460380.639378 \pm 0.000030$	0.110156 ± 0.000090	6.819 ± 0.021	82.746 ± 0.031
Eureka!	NRS2	$2460380.639344 \pm 0.000034$	0.10991 ± 0.00011	6.799 ± 0.018	82.717 ± 0.026
Eureka!	Combined fit	$2460380.639363 \pm 0.000018$	0.109920 ± 0.000067	6.791 ± 0.009	82.699 ± 0.013
Tiberius	NRS1	$2460380.639369 \pm 0.000023$	0.110265 ± 0.000070	6.798 ± 0.012	82.716 ± 0.017
Tiberius	NRS2	$2460380.639335 \pm 0.000029$	0.109843 ± 0.000088	6.782 ± 0.017	82.690 ± 0.021
Tiberius	Weighted mean	$2460380.639356 \pm 0.000018$	0.110101 ± 0.000055	6.792 ± 0.009	82.706 ± 0.013

The best-fitting system parameters are derived from the white-light light-curves from JK’s reduction and are given in Table 4.2. These are the result of optimization using a Levenberg-Marquardt algorithm. Following the fitting of the white light curves, we fit our spectroscopic light curves with the same model setup but with the system parameters fixed to the weighted mean values as given in Table 4.2. The final two **Tiberius** spectra were highly concordant, depicted in Fig. 4.5. I therefore proceeded in my analysis with just JK’s spectrum, given its utility as a point of comparison in a homogeneous analysis between BOWIE-ALIGN targets. The result of this was the transmission spectra shown in Fig. 4.6, hereafter referred to simply as the **Tiberius** reduction.

4.4.3 The transmission spectrum of HAT-P-30b

I present the transmission spectra obtained from each of our reduction pipelines at spectral resolutions of $R = 100$ and $R = 400$ in Fig. 4.6. While these reductions agree well in overall feature shape and depths (with a 7 ppm mean offset), I do find that the difference between the reductions depends linearly on wavelength, as depicted




Figure 4.6: Transmission spectra of HAT-P-30 b from Tiberius (top) and Eureka! (bottom) data reductions, binned to spectral resolutions of $R=400$ and $R=100$.




Figure 4.7: Comparison between transmission spectra obtained by our **Tiberius** and **Eureka!** reductions at $R = 100$, including the difference between the $R = 100$ and $R = 400$ spectra in green in the bottom panel. The shaded grey regions are the $1\times$ and $2\times$ the median transit depth uncertainty of the **Eureka!** $R = 100$ reduction. I find that the difference between reductions depends linearly on wavelength, with a mean offset of $+35$ ppm and -21 ppm in NRS1 and NRS2 respectively.

in Fig. 4.7. I performed a linear fit on the differences, demonstrating a difference in slope of 100 ppm across the whole spectrum. The light curve fitting, while performed independently for the two reductions, arrived at the same systematic models and used the same limb-darkening coefficients, only differing in the use of MCMC versus Levenberg-Marquardt, which has been demonstrated to not impact the results. While our derived **Eureka!** transmission spectrum uses a different approach to system parameters (using the system parameters from each detector individually) than the **Tiberius** spectrum (using the weighted mean), this was verified to not be the cause of the difference by testing a **Eureka!** transmission spectrum using system parameters from a combined white light curve fit. Comparing this to the original **Eureka!** spectrum resulted in a near-identical transmission spectrum with no slope and minimal detector offset (~ 5 ppm). Due to the high impact parameter of HAT-P-30 b ($0.854^{+0.008}_{-0.010}$ Johnson et al., 2011), the limb-darkening is fixed in both of our

reductions.

A detailed investigation into differences in reduction method was performed for the previous BOWIE-ALIGN target WASP-15b (Kirk et al., 2025), finding that the offset and slopes between the *Tiberius* and *Eureka!* reductions occur at the spectral extraction stage, as opposed to stage 1 extraction, system parameters, or limb-darkening coefficients. Understanding the difference between reductions is an ongoing effort for the field (e.g. Carter et al., 2024).

4.5 Atmospheric modelling

To interpret the *Eureka!* and *Tiberius* transmission spectra, I employed a grid of simple 1D equilibrium forward models for an initial comparison to the data. I follow this with an atmospheric retrieval analysis to provide a detailed interpretation, implemented using the publicly available package *petitRADTRANS*. CEF provided an independent retrieval analysis using the alternative software *BeAR*. The modelling setup and results of each package are detailed below.

4.5.1 *petitRADTRANS* forward models

I simulated the atmosphere of HAT-P-30 b with the *petitRADTRANS* (v3.1.3, Mollière et al., 2019) package, assuming a 1D isothermal atmosphere and chemical equilibrium. In all cases I assumed a H₂/He-dominated atmosphere (assuming a solar H/He ratio), and included $R = 1000$ correlated- k line opacities from CO (Rothman et al., 2010), H₂O (Polyansky et al., 2018), CO₂ (Yurchenko et al., 2020), CH₄ (Yurchenko et al., 2017), H₂S (Azzam et al., 2016), SO₂ (Underwood et al., 2016), SO (Brady et al., 2024), HCN (Barber et al., 2014), and NH₃ (Coles et al., 2019), as well as collisionally-induced absorption from H₂–H₂ and H₂–He (Borysow et al., 1988, 2001; Borysow, 2002), and Rayleigh scattering from H₂ and He (Dalgarno & Williams, 1962; Chan & Dalgarno, 1965). I modelled the atmosphere using 100 log-spaced pressure layers from 10^{-6} to 10^2 bar.

I initially constructed a simple grid of cloud-free simulations with the temperature set to the equilibrium temperature, varying the C/O ratio from 0.2–1.0, and the log10 metallicity relative to solar ($[M/H]$) from -1 to 2. These simple models demonstrate visibly stronger spectral features than those present in the observed spectra, highlighting that the spectral features are muted to some degree, as shown in the first panel of Fig. 4.8.

I then expanded this initial grid by including an opaque grey cloud deck at varying pressures (up to 0.1 mbar), and by reducing the isotherm temperature (down to 1000 K). By tuning the C/O ratio, isotherm temperature, and cloud-top pressure, I was able to create comparable, adequate fits to the data at a range of metallicities from $0.3 - 30\times$ solar, presented in the second panel of Fig. 4.8. This exemplifies the strong degeneracies in atmospheric models, particularly given the limited wavelength coverage of G395H data, with only subtle differences between these fits mainly visible




Figure 4.8: 1D isothermal chemical equilibrium forward models of HAT-P-30 b compared to the Eureka! $R = 100$ transmission spectrum. The top panel depicts cloud-free models with the isotherm temperature set to 1630 K, the equilibrium temperature of HAT-P-30 b. The bottom panel depicts a selection of models at different metallicities with comparable fits to the data using different C/O ratios, isotherm temperatures, and opaque grey cloud-top pressures, all noted in the legend. These models demonstrate that the spectral features of HAT-P-30 b are muted compared to those predicted by cloud-free models. The degeneracies between composition, temperature, and clouds highlighted in the bottom panel demonstrate the need for a full exploration of the parameter space using atmospheric retrievals to interpret the spectrum of HAT-P-30 b.

at $4.4\ \mu\text{m}$ and $3.3\ \mu\text{m}$. Our ability to constrain the atmosphere of HAT-P-30 b by eye with a coarse forward model grid alone is therefore limited, and I turn to the statistical power offered by atmospheric retrievals and nested sampling to fully explore the model parameter space.

4.5.2 **petitRADTRANS retrievals**

I performed a variety of atmospheric retrievals on the HAT-P-30 b transmission spectrum with **petitRADTRANS** (Nasedkin et al., 2024), using a similar setup to that used to analyse the spectrum of the BOWIE-ALIGN targets WASP-15 b (Kirk et al., 2025) and TrES-4 b (Meech et al., 2025). **petitRADTRANS** explores the multi-dimensional model transmission spectrum parameter space using Bayesian nested sampling (Skilling, 2004) implemented through **MultiNest** (Feroz & Hobson, 2008) with **PyMultiNest** (Buchner et al., 2014). As with other BOWIE-ALIGN targets, I use three principle setups: equilibrium chemistry, free chemistry, and hybrid chemistry.

For all of our retrievals, I use the model setup properties and opacity sources described in Section 4.5.1. The stellar radius is fixed to the value of $1.215 R_{\odot}$ from Bonomo et al. (2017). For our standard retrieval setups, I assume an isothermal pressure-temperature profile with a wide uniform temperature prior of 500–3000 K. I also adopt a Gaussian prior for the gravity based on the mass and radius (Bonomo et al., 2017), and a wide uniform prior for the planetary radius of $0.8\text{--}2.2 R_{\text{J}}$, defined at a reference pressure of 1 mbar. I parametrize the impact of aerosols by including a grey cloud deck, with a log-uniform prior on the cloud-top pressure from 10^{-6} bar to 10^2 bar.

I ran the retrievals on the $R = 100$ and $R = 400$ transmission spectra from both the **Eureka!** and **Tiberius** reductions. I permit for an offset between NRS1 and NRS2, with a uniform prior of ± 200 ppm. I determined the detection significances of different species by computing the difference in Bayesian evidence $\ln Z$ between the retrieval with all species included, and the retrieval with the considered species omitted, converted into a frequentist significance value from the Bayes factor B_m via the p -value using $p = \frac{1}{1+B_m}$ (Jeffreys, 1939).

Equilibrium chemistry

In our equilibrium retrievals, I assume all the atmospheric species are in chemical equilibrium, with the abundances of CH_4 , CO_2 , CO , H_2O , H_2S , HCN , and NH_3 interpolated from a pre-computed grid parametrized with temperature, pressure, C/O ratio, and [M/H]. I use a wide uniform prior in both C/O ratio (0.1–1.5) and [M/H] (−2–3). In this parametrization, the C/H ratio is fixed by the [M/H] value (using solar values from Asplund et al., 2009), while the O/H ratio is computed by the ratio of the C/H to the C/O.

I obtain consistent results across both reductions and resolutions, inferring a sub-Solar to Solar metallicity and a sub-Solar C/O ratio, with values of [M/H]




Figure 4.9: Median retrieved spectrum (red) of the **Eureka!** $R = 400$ data (orange) from the **petitRADTRANS** equilibrium chemistry retrieval with no offset, with the posterior probability distributions of C/O ratio and [M/H].

$= -0.42^{+0.48}_{-0.38}$ and $\text{C/O} \leq 0.55$ (2σ upper limit) from the **Eureka!** $R = 400$ reduction, depicted in Fig. 4.9. I also infer cool limb temperatures of 1010^{+110}_{-70} K, as is commonly observed in transmission spectroscopy retrievals (MacDonald et al., 2020; Welbanks & Madhusudhan, 2022), and poor constraints on the cloud-top pressure (with a peak in the posterior at 10 mbar - see Fig. 4.10). There is evidence for H_2O and CO_2 at 3.3 and 2.1 σ significance respectively. The choice of spectral resolution has negligible impact on the posteriors, depicted in Fig. 4.11.

The **Eureka!** reduction is consistent with no detector offset, while the **Tiberius** reduction has a median detector offset of -56 ± 35 ppm. Considering the metallicity posterior in Fig. 4.10 highlights a degeneracy present in the **Tiberius** reduction, with a lower evidence secondary mode with high metallicity ($\sim 60\times$ solar) and a large ~ 110 ppm detector offset. This degeneracy is far weaker in the **Eureka!** reduction, only representing a slight tail in the posterior metallicity distribution rather than a second peak.

Repeating the retrievals on the $R = 400$ reductions with no offset, I confirm that no offset is necessary for the **Eureka!** reduction, slightly increasing the Bayesian evidence with a Bayes factor of 4.5, and removing the high metallicity tail, resulting in a somewhat tighter posterior on the metallicity of $-0.39^{+0.30}_{-0.27}$. Not including an offset for the **Tiberius** reduction, on the other hand, is greatly disfavoured with a Bayes factor ($\frac{Z_1}{Z_2}$) of 1339, and gives discrepant results, with a lower C/O ratio, metallicity, and deeper clouds. I present the median model fit to **Eureka!** in the no-offset case in Fig. 4.9, and the opacity contributions to this model in Fig. 4.12.

Free chemistry

For our free chemistry retrievals, the abundances of CH_4 , CO_2 , CO , H_2O , H_2S , SO_2 , and SO are free parameters, with a wide log-uniform prior in mass fraction from




Figure 4.10: Corner plot showing the posterior probability distributions from the `petitRADTRANS` equilibrium chemistry retrievals on the $R = 400$ transmission spectra from *Tiberius* (blue) and *Eureka!* (orange, with black text). The top right panel displays the best-fit models from each retrieval, offset by 500 ppm, with the reduced χ^2 of the fit indicated in the legend.

10^{-12} to $10^{-0.5}$. I also tested retrievals including opacity from other species including NH_3 , HCN , OCS , CS_2 , and N_2O , but chose not to include them from our primary retrievals when the posteriors provided minimal constraints on their abundances. I obtain consistent results across all reductions and resolutions, retrieving a $\log(X_{\text{H}_2\text{O}})$ volume mixing ratio of -3.3 ± 0.8 , and a $\log(X_{\text{CO}_2})$ volume mixing ratio of $-8.1^{+0.9}_{-1.8}$ from the *Eureka!* $R = 400$ reduction. I also retrieve a very cold limb temperature of 650^{+120}_{-80} K, far cooler than the planet's equilibrium temperature of 1630 K. I place 2σ upper limits on the $\log(X_{\text{CH}_4})$ volume mixing ratio of -6.1 , $\log(X_{\text{CO}})$ volume mixing ratio of -3.6 , and $\log(X_{\text{SO}_2})$ volume mixing ratio of -5.3 , respectively (see the posterior corner plot in Fig. 4.13). There is also a strong peak in the posterior of the SO abundance at ~ 10 ppm, with a feature at $4.5 \mu\text{m}$, although repeating the




Figure 4.11: Corner plot showing the posterior probability distributions comparing the `petitRADTRANS` equilibrium chemistry retrieval on the $R = 400$ *Eureka!* transmission spectrum (orange, with black text), and the $R = 100$ *Eureka!* transmission spectrum (red).

retrievals without SO finds its inclusion is slightly disfavoured by the Bayesian evidence ($\Delta \ln Z = -0.4$), and it does not affect the other posteriors. The free retrieval abundance posteriors are consistent with the chemistry implied by the equilibrium retrievals above 100 mbar, as demonstrated in Fig. 4.14. The derived H_2O and CO_2 abundances correspond to an atmospheric O/H ratio of $0.51^{+2.72}_{-0.43} \times \text{Solar}$, which is consistent with, but far less precise than, the atmospheric O/H ratio constraint from equilibrium chemistry of $\text{O/H} = 0.77^{+0.84}_{-0.56} \times \text{Solar}$. Without meaningful constraints on the CO abundance, the atmospheric C/H ratio cannot be derived from the free chemistry retrievals.

Hybrid chemistry

In the hybrid chemistry retrievals, I combine the equilibrium chemistry approach to set the abundances of the carbon- and nitrogen-bearing species CH_4 , CO_2 , CO , H_2O ,




Figure 4.12: Median retrieved transmission spectrum of the **Eureka!** $R = 400$ data from the **petitRADTRANS** equilibrium chemistry retrieval convolved to a spectral resolution of $R = 400$, with the spectral contributions from H_2O , CO_2 , CH_4 , CO , and a grey cloud deck, and the combined absorption of all opacities (red).

HCN , and NH_3 , parameterized with C/O ratio and $[\text{M}/\text{H}]$, and the free chemistry approach to set the abundances of sulfur species H_2S , SO_2 , and SO to be free parameters. This allows the model atmosphere to be broadly in chemical equilibrium, but with the abundance of sulfur species modified by photochemistry or a variable S/O ratio.

The hybrid chemistry results are broadly consistent with the equilibrium chemistry results, however, there is stronger evidence for a high metallicity, high detector offset mode in both **Eureka!** and **Tiberius** reductions now, becoming the higher evidence mode in **Tiberius** $R = 400$. This appears to be caused by SO_2 , which has a notable peak in the posterior at ~ 1 ppm abundance. The abundances of H_2S and SO are unconstrained in the hybrid retrievals. All the hybrid retrievals have poorer evidence than their equilibrium chemistry counterparts. Each sulfur species added further decreases the evidence (disfavoured with a Bayes factor of 3 with all three species for the **Eureka!** reductions), except in the **Tiberius** $R = 400$ reduction, where the inclusion of SO_2 minimally improves the Bayesian evidence (Bayes factor of 1.3). The $R = 400$ hybrid chemistry retrieval posteriors are depicted in Fig. 4.15.




Figure 4.13: Corner plot showing the posterior probability distributions from the *petitRADTRANS* free chemistry retrievals on the $R = 400$ transmission spectra from *Tiberius* (blue) and *Eureka!* (orange, with black text). Abundances are given in units of logmass fractions. The top right panel displays the best-fit models from each retrieval, offset by 500 ppm, with the reduced χ^2 of the fit indicated in the legend.




Figure 4.14: Vertical mixing ratios of H₂O, CO₂, CH₄, CO, and H₂S from the Eureka! $R = 400$ petitRADTRANS equilibrium chemistry retrieval (lines), plotted with the median retrieved abundances with 1σ error bars for the detected species (H₂O and CO₂), and upper limits from the other species (CH₄, CO, and H₂S) from the free chemistry retrieval (data points).

4.5.3 BeAR retrievals

CEF also performed retrievals using the open-source GPU-accelerated Bern Atmospheric Retrieval code (BeAR³) (Kitzmann et al., 2020). BeAR uses the MultiNest library (Feroz & Hobson, 2008) to perform the retrieval using Bayesian nested-sampling (Skilling, 2004), and line-by-line opacity sampling. For the BeAR retrievals, CEF used a similar setup as in Kirk et al. (2025). The opacities were sampled at a resolution of 0.1 cm^{-1} in wavenumber (equivalent to $R \sim 20,000\text{--}30,000$), and include the following molecules and their associated ExoMol and HITRAN line-lists: H₂O (Polyansky et al., 2018), CH₄ (Yurchenko & Tennyson, 2014), CO (Li et al.,

³Formerly known at Helios-r2. BeAR can be found at <https://github.com/newstrangeworlds/bear>




Figure 4.15: Corner plot showing the posterior probability distributions comparing the *petitRADTRANS* hybrid chemistry retrievals on the $R = 400$ Eureka! transmission spectrum (orange, with black text) and the $R = 400$ Tiberius transmission spectrum (blue). Abundances are in units of mass mixing ratio.

2015), CO₂ (Yurchenko et al., 2020), H₂S (Azzam et al., 2016), SO₂ (Underwood et al., 2016), NH₃ (Coles et al., 2019), and HCN (Barber et al., 2014). All these opacities are computed using **Helios-k** (Grimm & Heng, 2015b) and are taken from the **DACE** database (Grimm et al., 2021). Also included was opacity due to H₂ Rayleigh scattering (Cox, 2000), and collision-induced absorption from H₂-H₂ (Abel et al., 2011) and H₂-He (Abel et al., 2012). The model atmosphere is divided into 200 levels, equal in log-pressure space, ranging from 10 bar to 10⁻⁸ bar.

For the retrievals, the stellar radius is again fixed to 1.215 $R_{\odot}$ (Bonomo et al., 2017). The planet’s gravity and radius at the 10 bar pressure level are free parameters in the retrieval. The gravity has a gaussian prior on $\log g$ with a mean of 2.96 and a standard deviation of 0.03 in cgs units, and the radius has a uniform prior of 1.25–1.65 R_J . The atmosphere is assumed to be isothermal, with a uniform prior on the temperature of 500–2500 K. A grey cloud deck is included, and the cloud-top pressure is a free parameter in the retrieval, with a log-uniform prior of 10⁻⁷–1.0 bar. CEF also included an offset between the NRS1 and NRS2 detectors, with a uniform prior from -100 to +100 ppm.

CEF ran two types of chemistry for the **BeAR** retrievals. Firstly, free chemistry retrievals are considered. In these retrievals, each molecule’s volume mixing ratio is a free parameter, with a log-uniform prior from 10⁻¹² to 0.3. The mixing ratios are constant with altitude. The rest of the atmosphere is then filled with H₂ and He, in a solar ratio of 0.17 (Asplund et al., 2009). Secondly, equilibrium chemistry is assumed. For this, the free parameters are the [M/H], which has a log-uniform prior of 0.1–1000 $\times$ solar, and the C/O ratio, which has a uniform prior of 0.1–2.0. [M/H] and C/O are related to the carbon and oxygen abundance in the same manner as in the **petitRADTRANS** models. The chemistry is calculated using **FastChem** (Stock et al., 2018; Stock et al., 2022), which is already integrated into **BeAR**. Both types of retrievals were applied to the **Eureka!** and **Tiberius** spectra, at $R = 100$ and $R = 400$.

In the equilibrium chemistry case, there is a good agreement across the two data reductions. The only difference occurs in the detector offset, for which the retrieval on the **Eureka!** reduction favours no offset, while the **Tiberius** reduction retrieval finds an offset of -60^{+23}_{-21} ppm. For the **Eureka** $R = 400$ retrieval, the retrieved [M/H] is $-0.52^{+0.31}_{-0.25}$, and the C/O ratio is $0.26^{+0.17}_{-0.11}$. The results are consistent with the **Tiberius** and $R = 100$ retrievals. The $R = 400$ retrievals are presented in Fig. 4.16

In the free chemistry case, there is also good agreement between **Eureka!** and **Tiberius** reductions, with the exception of a bimodality in the **Eureka!** $R = 100$ case. The alternative solution in the **Eureka!** case corresponds to an exceptionally high CO₂ abundance, which is compensated by a higher temperature, larger planet radius, lower water abundance, and a substantial negative offset of NRS2. However, at $R = 400$, the second solution in the **Eureka!** case is not present (as the corresponding offset suggested by the **petitRADTRANS** retrieval is outside of the offset




Figure 4.16: Cornerplot showing the posteriors from the BeAR chemical equilibrium retrievals on the **Tiberius** (blue) and **Eureka!** (orange) reductions of HAT-P-30 b at $R = 400$. The top right insert shows the best-fit models for the **Tiberius** (blue) and **Eureka!** (orange) reductions. The **Eureka!** spectrum is offset by 1300 ppm for visualisation purposes. The legend in the bottom left indicates the reduced χ^2 values for each of the fits. Figure created by CEF.

prior range), and the retrievals agree with both the $R = 100$ and $R = 400$ **Tiberius** solution.

Finally, CEF also tested the effects of adding stellar contamination to our free chemistry retrievals. For this, the PHOENIX grid of stellar models (Husser et al., 2013) was used, with the stellar values set to those in Table 4.1. The stellar effective temperature was retrieved, using a Gaussian prior with a mean of 6304 K and a standard deviation of 88 K (Johnson et al., 2011). The effects of hot and cold spots on the star were then included, retrieving their temperatures and covering fractions. For the hot spots, the temperature is retrieved as a positive ΔT with respect to the stellar effective temperature, using a uniform prior of 0–1000 K. For the cold spots, the temperature is retrieved as a negative ΔT with respect to the stellar effective temperature, using a uniform prior of 0–1500 K. For both, the covering fractions use uniform priors of 0–0.5. For both reductions and resolutions, adding stellar contamination to the retrievals has little effect on the results, and a Bayesian evidence comparison favours the retrieval without contamination in all cases except using the **Eureka!** reduction at $R = 100$. In the latter case, however, the Bayes factor is negligible at only 1.65. Overall CEF and I conclude that this constitutes no evidence for stellar contamination in the spectrum of HAT-P-30 b.

4.5.4 Overview of retrieval results

Using two independent retrieval analyses with **petitRADTRANS** and **BeAR** on two independent data reductions at two spectral resolutions, with a variety of retrieval setups, I am able to put together a consistent picture for the atmosphere of HAT-P-30 b. I find that retrievals on the **Tiberius** reductions require an offset between the detectors, while the **Eureka!** reduction retrievals do not. I also generally find that there are slightly tighter constraints using the $R = 400$ retrievals, given the better ability to sample narrow spectral features. Therefore I favour the interpretations from the **Eureka!** $R = 400$ retrieval with no offset. I present the median retrieved **petitRADTRANS** transmission spectrum to the **Eureka!** $R = 400$ spectrum in Fig. 4.9 with equilibrium chemistry and the posteriors for metallicity and C/O ratio, as well as the spectral contributions of each opacity source in Fig. 4.12, and the posterior corner plot in Fig. 4.17. I present a statistical comparison of the different $R = 400$ **petitRADTRANS** retrievals in Tab. 4.3.

All the equilibrium retrievals have a consistent main result: sub-solar C/O ratio (≤ 0.55) and metallicity ($-0.39^{+0.30}_{-0.27} \times \text{solar}$), with exact numbers taken from the **petitRADTRANS Eureka!** $R = 400$ equilibrium retrieval with no offset. The free chemistry retrievals are similarly concordant, and yield consistent abundances with those implied by the equilibrium chemistry, with the vertical abundance profiles depicted in Fig. 4.14. Free and equilibrium chemistry give rise to consistent posteriors for the O/H ratio, although the latter approach results in much more precise constraints.

In a sub-sample of the retrievals, there is a degree of bimodality, with an




Figure 4.17: Corner plot showing the posterior probability distributions from the favoured retrieval setup: the `petitRADTRANS` equilibrium chemistry retrieval on the $R = 400$ *Eureka!* transmission spectrum with no detector offset.

alternative explanation of a high-metallicity, CO_2 -dominated spectrum. This secondary mode appears in the *BeAR* free retrieval of the *Eureka!* $R = 100$ spectrum, the `petitRADTRANS` equilibrium retrievals of both *Tiberius* spectra, and all the `petitRADTRANS` hybrid retrievals. The mode is highly degenerate with detector offset, requiring an offset of ~ 60 ppm in the *Eureka!* data and ~ 110 ppm in the *Tiberius* data. This mode represents a $50\text{--}100\times$ solar metallicity, the lower end of which does fit within the constraints of a fully-mixed atmosphere from the interior structure modelling (see ahead to Section 4.6). This mode is only favoured over the low-metallicity mode in a single retrieval, the *Tiberius* $R = 400$ `petitRADTRANS` hybrid retrieval. However, the evidence generally does not support the extra free parameters necessary for utilising a hybrid retrieval over equilibrium chemistry, and the ~ 110 ppm offset required is much larger than the typical offset observed in other data sets (e.g., Alderson et al., 2023; Moran et al., 2023). I therefore highly emphasize the lower metallicity mode as our favoured interpretation.

I also note that all of our free chemistry retrievals have very cold limb temper-

Retrieval	$\ln Z$	χ^2_ν
<i>Equilibrium Chemistry:</i>		
Eureka! $R = 400$	1797.8 ± 0.1	0.97
Tiberius $R = 400$	1803.9 ± 0.1	1.00
<i>Eq. Chemistry, no offset:</i>		
Eureka! $R = 400$	1799.3 ± 0.1	0.96
Tiberius $R = 400$	1796.7 ± 0.1	0.96
<i>Free Chemistry:</i>		
Eureka! $R = 400$	1798.0 ± 0.1	0.96
Tiberius $R = 400$	1805.6 ± 0.1	1.04
<i>Free Chemistry, no SO:</i>		
Eureka! $R = 400$	1797.6 ± 0.1	0.97
<i>Hybrid Chemistry:</i>		
Eureka! $R = 400$	1796.9 ± 0.1	0.97
Tiberius $R = 400$	1803.5 ± 0.1	0.94

Table 4.3: The Bayesian evidence ($\ln Z$) of different `petitRADTRANS` atmospheric retrievals performed on the $R = 400$ **Tiberius** and **Eureka!** transmission spectra of HAT-P-30b, and the corresponding reduced χ^2_ν of the best-fitting model from each retrieval.

atures of ~ 650 K. Cool limb temperatures are commonly observed in transmission spectroscopy of hot Jupiters, and can be attributed to the 1D modelling of the atmosphere (MacDonald et al., 2020) (i.e. limb asymmetries) or the isothermal parametrization used (Welbanks & Madhusudhan, 2022). This case, however, seems particularly extreme, with a 1000 K discrepancy from equilibrium temperature to observed isothermal temperature, although within the range of biases explored by MacDonald et al. (2020). By comparing to the equilibrium chemistry retrievals, which are effectively forced into a higher temperature by the absence of strong CH_4 features in the spectrum (the dominant carbon-carrier at low temperatures), that the retrievals may in fact be probing the cloud-top pressure-scale height degeneracy. The particular noise profile of the observation may be affecting the feature width and pushing us towards a low temperature explanation for the muted scale height, while in reality this is more likely an impact of clouds. This is demonstrated when restricting the temperature prior to higher temperatures, as the cloud top pressure is pushed higher to compensate. This does not impact our interpretation of the composition, however, as the abundances are consistent with the higher temperature equilibrium chemistry retrieval counterparts. A limb asymmetry analysis of the complete sample of BOWIE-ALIGN planets, including HAT-P-30 b, will be presented in future work.

I do note that the contribution from CH_4 appears to be non-negligible (see Fig. 4.12), despite CH_4 not being expected at the high temperatures of HAT-P-30 b. If the retrieved low temperature of HAT-P-30 b is a result of the temperature profile parametrization or limb asymmetry, the retrieval may then be trying to minimise the contribution from CH_4 by suppressing the carbon content of the planet, reflected as a low C/O ratio and metallicity. To account for this, I ran an equilibrium retrieval

discarding CH_4 as an opacity source. This results in a temperature posterior of ~ 650 K, matching the free retrievals, as the absence of strong CH_4 features no longer sets a lower limit on the temperature. I account for this by repeating the retrieval without CH_4 with a temperature lower bound of 1000 K. This results in a completely consistent best-fit spectrum to the equilibrium retrieval including CH_4 , although the constraints on temperature, cloud-top pressure, and metallicity are all weaker. I therefore demonstrate that rather than the low temperature driving the inclusion of CH_4 in the spectrum, it is in fact the other way around, with the enhanced opacity at $\sim 3.3 - 3.5 \mu\text{m}$ driving the retrieval to a lower temperature. I therefore do not believe that the low observed temperatures are suppressing the C/O ratio, but that our low temperature conclusion may instead be the result of enhanced transit depths at $\sim 3.3 - 3.5 \mu\text{m}$ from the particular random noise draw of this observation, the presence of disequilibrium CH_4 (e.g. an enhancement of CH_4 due to vertical or horizontal transport) or species with similar opacity, or a non-grey continuum opacity putting in a bluewards-increasing slope in that wavelength range. I conclude that the observed sub-solar C/O ratio is robust, while the metallicity, while consistent with sub-solar, may be less well constrained ($-0.38^{+0.83}_{-0.68}$) than our equilibrium models suggest, should the real temperature be higher and the opacity discovered to not be CH_4 . Fixing the temperature at the equilibrium temperature of 1630 K biases the metallicity posterior towards the higher end of that distribution ($0.43^{+0.22}_{-0.46}$) but is still consistent with the nominal sub-solar metallicities within 2σ , while ruling out metallicities greater than $10\times\text{Solar}$. The C/O ratio is unchanged by the fixed temperature, while the need for clouds around 0.1-1 mbar to sufficiently mute the spectral features becomes evident. The Bayesian evidence does not support fixing the temperature (decrease in $\ln Z$ of 6.7), and I note the true temperature profile could very well diverge from an isotherm fixed to the equilibrium temperature.

For the purpose of studying the implications of our retrieved parameters for planet formation and relation to other planets in the BOWIE-ALIGN sample, I recommend the use of the posteriors from the `Eureka!` $R = 400$ `petitRADTRANS` retrieval with no offset between the detectors (Fig. 4.17), as it is not impacted by the degeneracy with the higher-metallicity mode. I therefore proceed with a C/O ratio of ≤ 0.55 and metallicity of $-0.39^{+0.30}_{-0.27}$.

4.6 Interior structure models

Using the Bayesian framework of Thorngren & Fortney (2019), DPT estimated the bulk composition of HAT-P-30 b, accounting for uncertainties in the mass, radius, and age of the planet. The evolution models underlying the statistical model are again from Thorngren & Fortney (2018), and solve the equations of hydrostatic equilibrium, conservation of mass, and the equation of state in one dimension. The planet’s thermal state is evolved forward from a hot initial state using the atmosphere models of Fortney et al. (2007) to regulate heat flow out of the planet’s interior. The equilibrium temperature is well into the hot Jupiter range, so the anomalous heating

was accounted for using the flux-heating relation of Thorngren & Fortney (2018). This implies a quite high intrinsic temperature of 653 ± 45 K.

The interior retrieval finds a bulk metallicity of $Z = 0.28 \pm 0.03$ (this is the statistical error from parameter uncertainties, not modeling uncertainty). This is within the expected range for a planet of this mass (Thorngren et al., 2016), being only moderately above average. If the planet were fully mixed, it would imply an atmospheric metallicity 2-sigma upper limit of $71.1 \times \text{Solar}$. The observed atmospheric metallicity could be less than this if the bulk metal is contained within a core or beneath a compositional gradient. It could not be greater, however, as this would be unstable to convection.

The substantial enhancement of the bulk metallicity relative to that observed in its atmosphere is not an inconsistency – it instead implies that the planet has a massive core. To match the radius and atmospheric metallicity, the core must be at least $50M_{\oplus}$. This could be either with a sharp core boundary or as a fuzzy core featuring an extensive compositional gradient (see e.g. Helled & Stevenson, 2017) – structure models such as these don’t distinguish between those cases (Bloot et al., 2023). A planet like this could have formed if most of the metal accretion occurred early on, with the final stages being the accretion of metal-poor gas. Alternatively, if the planet formed in the ice line, there might be a substantial quantity of rock and iron mixed into its interior that has either settled out or cannot be seen in the atmosphere due to condensation.

4.7 Discussion

4.7.1 The atmosphere of HAT-P-30 b

The transmission spectrum, obtained from two consistent independent reductions, describes the wavelength-dependent opacity in the atmosphere of HAT-P-30 b from $2.8 \mu\text{m}$ and $5.2 \mu\text{m}$ at a resolution of $R = 400$. We have used parametric radiative transfer models, implemented using atmospheric retrievals from the **petitRADTRANS** and **BeAR** retrieval packages, with a variety of retrieval setups including equilibrium and free chemistry, to analyze this spectrum and make inferences about the atmospheric properties of HAT-P-30 b. In 22 of the 23 retrievals performed, the favoured interpretation is one of a low-metallicity and C/O ratio atmosphere, representing sub-solar abundances of carbon and potentially also oxygen, derived from the equilibrium chemistry retrievals ($\text{C}/\text{H} = 0.41^{+0.40}_{-0.19} \times \text{Solar}$ and $\text{O}/\text{H} = 0.77^{+0.84}_{-0.56} \times \text{Solar}$, computed from the C/H, which is fixed by the metallicity, and the C/O ratio). In this interpretation, some muting of spectral features from cooler-than-expected temperatures is required (~ 1000 K in the equilibrium retrievals, ~ 650 K in the free retrievals due to its freedom to supply a CH_4 -free low temperature model) and clouds. No offset between detectors is favoured by the Bayesian evidence for the **Eureka!** spectra, while a ~ 55 ppm offset is required for the **Tiberius** spectra, matching the observed linear offset between the reductions. There is evidence for

both H₂O and CO₂ at 3.3 and 2.1 σ respectively from the equilibrium retrievals.

Favoured in the Tiberius $R = 400$ **petitRADTRANS** hybrid chemistry retrieval, and present as a notable secondary mode in 4 other retrievals, is an alternative interpretation, with a highly super-solar metallicity of $\sim 60\times$ solar. This is consistent with the upper limits of metallicity of a fully-mixed atmosphere suggested by our interior structure modelling. This mode requires freedom in both detector offset, needing a ~ 60 ppm offset in the **Eureka!** retrievals and ~ 110 ppm in the **Tiberius** retrievals, and SO₂ abundance, to appropriately fit the spectra with CO₂ features.

I favour the low metallicity interpretation for the atmosphere of HAT-P-30 b for a number of reasons. Firstly, it is favoured in all of the standard equilibrium and free chemistry retrievals, and it is only not favoured with a single combination of reduction, resolution and retrieval pipeline, using the hybrid retrieval method - where the increase in complexity from the equilibrium method is not justified by the Bayesian evidence. Secondly, the large detector offset of 110 ppm that gives rise to the higher metallicity solution is greater than the offsets that have previously been present in NIRSpec/G395H observations of hot Jupiters (e.g., Alderson et al., 2023). Thirdly, the high metallicity modes only notably appears when the SO₂ abundance is a free parameter, to appropriately scale the transit depth at 4 μm . While SO₂ has been observed in exoplanet atmospheres (e.g., Alderson et al., 2023) and is predicted by photochemical models (e.g., Tsai et al., 2023; Kirk et al., 2025) in high metallicity hot Jupiters, including SO₂ does not improve the Bayesian evidence in all free retrievals, and 3 of 4 hybrid retrievals.

I therefore highly emphasize our low metallicity ($-0.39^{+0.30}_{-0.27}$) and C/O ratio (≤ 0.55) interpretation as that most favoured by our data and physical models. I do however note the existence of this physically-plausible alternative interpretation, with a high metallicity ($1.77^{+0.31}_{-2.07}$) and potential SO₂, representing a fully-mixed atmosphere, although I highlight it is by and large not favoured by our data. Further observations could definitively resolve between these scenarios, including NIRISS SOSS observations taken in March 2025 as part of JWST programme 5924 (PI: Sing). I also note that the suggestion of enhancements absorption at the wavelengths of the CH₄ feature at 3.3 μm is somewhat impacting our interpretation of the chemistry and atmospheric temperature, and while this does not affect our inferred C/O ratio, disregarding CH₄ as a potential equilibrium opacity source does weaken the constraint on the metallicity to $-0.38^{+0.83}_{-0.68}$.

4.7.2 Implications for formation

The preferred atmospheric retrieval solutions of HAT-P-30 b show a sub-solar C/O ratio and sub-solar, sub-stellar metallicity with $\text{C/O} \leq 0.45$ and $[\text{M}/\text{H}] = -0.83$ to -0.05 (1σ ranges), derived from **petitRADTRANS** equilibrium chemistry retrievals of the $R = 400$ **Eureka!** reduction. The C/O ratio deviates from the solar value at the $\sim 1.5\sigma$ level. In formation models, supersolar metallicities and subsolar C/O

are the expected outcome of high solid accretion of oxygen-rich dust and ice (e.g. Madhusudhan et al., 2014). However, I disfavour the high metallicity solutions with large detector offsets that appear in a subset of our retrievals (see Section 4.7.1).

Constraints from the C/O ratio and metallicity alone lead to large degeneracies in the possible formation history of the planet. However, a sub-solar C/O ratio and a sub-solar metallicity together are harder to reproduce because the gas in protoplanetary discs typically has sub-solar metallicity but super-solar C/O (see, e.g. Bergin et al., 2024). For this reason, the retrieved atmosphere is inconsistent with the fiducial models of Penzlin et al. (2024) by at least 1σ . The closest match is for a planet that accreted most of its gas far from the star, where the gas is extremely metal-poor and all carbon and oxygen carriers are frozen out on the dust. Since the solid-phase abundance at this location has a solar C/O ratio, adding solids at this location produces an atmosphere with a solar C/O. If the planet accretes solids after migrating further in, this is no longer the case – the composition of solids can reach a sub-solar C/O ratio due to, e.g., the sublimation of CO and CH₄. Thus, by accreting gas far out but being enriched by a modest amount of solids as it migrates, a planet can reach sub-solar C/O ratios and metallicities in the outer disc in our models (Penzlin et al., 2024). Despite this, the fiducial simulations do not reach C/O ratios below 0.35 for sub-solar metallicity and, hence, exceed the best fit values of HAT-P-30 b. This mismatch is, however, sensitive to the composition of the disc – the depletion of CO into CH₄ or C₂H₆ instead (e.g. Bosman et al., 2019) can lower the C/O ratio in the solids further and lead to a better agreement between models and observation, and is less than 1σ .

As noted by Meech et al. (2025) when discussing TrES-4b (which has similar metallicity and C/O ratio to HAT-P-30 b: $\log(Z/Z_{\odot}) = -0.41$ to -0.04 and $C/O = 0.30$ – 0.42), this is not the only possible channel of formation of sub-solar C/O and metallicities. Two further mechanisms allow a planet to reach sub-solar C/O and metallicity while migrating through the disc to a close-in orbit. In the first case, the abundances are dominated by the gas accreted. This requires the gas runaway phase of planet growth to happen just inside the water ice line (where the gas-phase C/O ratio can be sub-solar if the dust grains contain a significant amount of refractory carbon) and, therefore, accretes its bulk mass through water vapour-enriched gas and would predict that HAT-P-30 b’s metallicity is at most only slightly sub-stellar. The second case relies on accretion in an environment where the gas is depleted in metallicity by trapping volatile species in a dust trap further out in the disc.

TrES-4b differs from HAT-P-30 b due to its orbital alignment: TrES-4b has an aligned orbit while HAT-P-30 b is misaligned. If this difference in orbital alignment reflects an essential role of high-eccentricity migration in HAT-P-30 b but not in TrES-4b, then TrES-4b will likely have migrated further through the disc to a close-in orbit and have probably accreted more gas and planetesimals in regions close to the star. This puts an additional caveat on the gas accretion-dominated scenario, in that the water snow line would have to be far enough out that HAT-P-30 b formed

inside the water snow line while remaining far enough out that high-eccentricity migration remains possible ($\gtrsim 0.6$ AU, Muñoz et al. 2016). As it stands, the similarity in composition between TrES-4b and HAT-P-30 b may, or may not, point to similar formation histories despite their different migration history, making it difficult to draw strong conclusions by comparing them. However, with the full sample of eight planets in the BOWIE-ALIGN survey, we can make statistically more meaningful comparisons between the populations, hopefully gaining new insights into the formation of gas giants.

4.8 Conclusions

In this chapter, I have presented the JWST NIRSpec/G395H transmission of HAT-P-30 b, the fourth planet observed as part of the BOWIE-ALIGN programme, and one of four misaligned planets in the eight planet BOWIE-ALIGN sample. My collaborators and I have used two independent data reductions and a host of atmospheric retrievals from two different pipelines to demonstrate a consistent interpretation for the planet’s atmosphere. I find that HAT-P-30 b has a sub-solar metallicity of $0.41^{+0.40}_{-0.19} \times \text{solar}$ and a sub-solar C/O ratio of $0.29^{+0.16}_{-0.13}$. I detect H_2O (3.1σ) and see evidence for CO_2 (2.1σ), with consistent abundances between free and equilibrium chemistry retrievals, and evidence for muting of spectral features due to clouds and/or cool limb temperatures. I do not find evidence for stellar contamination or sulfur chemistry. I see no evidence for a detector offset in the **Eureka!** reduction, while I observe a ~ 55 ppm detector offset between the detectors in the **Tiberius** reduction, becoming consistent with **Eureka!** with the free offset and therefore matching the observed linear offset between the reductions. I see evidence for a second possible atmospheric composition in a sub-set of the retrievals corresponding to a higher metallicity of $\sim 60 \times \text{solar}$, consistent with the upper limit of a fully-mixed atmosphere from our internal structure modelling. This alternative mode is favoured in only one of 23 performed retrievals, and is absent in most of the performed retrievals, requiring a larger offset (~ 110 ppm for **Tiberius**, ~ 65 ppm for **Eureka!**) and SO_2 -abundance as a free parameter to adequately fit the observed transmission spectrum. I therefore strongly favour the low metallicity interpretation (discussed in more detail in Section 4.7.1), although I note that the metallicity constraints may be weaker depending on the cause of the low observed temperature.

The retrieved abundances suggest a composition similar to another BOWIE-ALIGN target, TrES-4 b (Meech et al., 2025), which has a metallicity of $0.58^{+0.59}_{-0.32} \times \text{solar}$ and a C/O of $0.32^{+0.11}_{-0.08}$, comparing **petitRADTRANS** equilibrium chemistry retrievals with identical setups. Obtaining a low metallicity and low C/O ratio deviates from the expected outcome of formation models, being inconsistent with Penzlin et al. (2024) by $> 1\sigma$. Several scenarios could explain this discrepancy, including accretion in the outer disc together with some CO depletion relative to the fiducial models, a composition dominated by gas accreted inside the water ice line, or accretion in

a metallicity-depleted environment due to the trapping of volatile species further out in the disc. While HAT-P-30 b is misaligned in its orbit, TrES-4 b is aligned, suggesting that TrES-4 b may have migrated through disc migration rather than high-eccentricity migration. Given the constraints these migration mechanisms may place on the planet’s evolution relative to the ice lines, we could very well be probing different explanations that give rise to similar atmospheric compositions. I therefore await the full BOWIE-ALIGN sample of eight planets to fully explore the formation of these two populations of hot Jupiters and draw statistically meaningful inferences. I do note that HAT-P-30 b has a significant difference in metallicity and C/O ratio compared to its fellow misaligned hot Jupiter WASP-15 b (Kirk et al., 2025), demonstrating significant diversity can exist in planets with presumed similar formation histories.

Chapter 5

BOWIE-ALIGN: Complete Sample

Declaration

This chapter is based on results and interpretation that I contributed to the following papers: “BOWIE-ALIGN: JWST reveals hints of planetesimal accretion and complex sulphur chemistry in the atmosphere of the misaligned hot Jupiter WASP-15b” Kirk et al. (2025), “BOWIE-ALIGN: Sub-stellar metallicity and carbon depletion in the aligned TrES-4b with JWST NIRSpec transmission spectroscopy” Meech et al. (2025), and “BOWIE-ALIGN: Modelling degeneracies in the NIRSpec/G395H transmission spectrum of the aligned hot Jupiter NGTS-2b” Fairman et al. (2026) as part of the BOWIE-ALIGN collaboration. All atmospheric models and retrievals presented here are exclusively my own work, interpreting JWST data obtained as part of the BOWIE-ALIGN programme and reduced by Eva-Maria Ahrer (EA), Charlotte Fairman (CF), James Kirk (JK), and Annabella Meech (AM). I include additional context and analyses beyond what is found in the contributed sections to the highlighted papers. For completeness, I also report results for the remaining sample, including WASP-94Ab and KELT-7b from Ahrer et al. (2025b) and Ahrer et al. (2025a), for which I did contribute to the analysis but did not directly provide models used in the final interpretation, and preliminary results for HD 149026 and WASP-17b, in which I was not involved.

5.1 BOWIE-ALIGN observations

BOWIE-ALIGN is a comparative study of hot Jupiter systems to determine whether differences in planet formation mechanisms are reflected in atmospheric composition, using transmission spectroscopy observations with the NIRSpec instrument of JWST (Kirk et al., 2024a). Eight hot Jupiters, four misaligned and four aligned, are included in the BOWIE-ALIGN sample, with five of these planets obtained as part of JWST programme GO 3838. The analysis of one of these five planets, HAT-P-30b,

constitutes Chapter 4. More detail on the BOWIE-ALIGN programme is provided in Sec. 4.1.

In this chapter I describe the results of the other planets of the BOWIE-ALIGN sample. In particular, I present the atmospheric models and retrievals that I used in the interpretation of the transmission spectra of three planets included in the BOWIE-ALIGN programme: WASP-15b (Kirk et al., 2025), TrES-4b (Meech et al., 2025), and NGTS-2b (Fairman et al., 2026). Details on these systems are presented in the following sub-sections. I provide a description of my default atmospheric modelling and retrieval setup in Sec. 5.2, and present my results for each planet in Sec. 5.3 (WASP-15b), Sec. 5.4 (TrES-4b), and Sec. 5.5 (NGTS-2b). I report results from other analyses in Sec. 5.6 to provide additional context.

5.1.1 The WASP-15 system

WASP-15b is a misaligned hot Jupiter ($1.428 \pm 0.077 R_J$, $0.542^{+0.054}_{-0.053} M_J$, $T_{\text{eq}} = 1676 \text{ K}$; Bonomo et al., 2017) with a precisely measured sky-projected obliquity of $-139.6^{+4.3}_{-5.2}^\circ$ (Triaud et al., 2010). It was the first planet observed and published by the BOWIE-ALIGN programme (Kirk et al., 2025), and orbits an F7 star above the Kraft break ($1.477 \pm 0.072 R_\odot$, $1.18 \pm 0.12 M_\odot$, Bonomo et al. 2017; $T_{\text{eff}} = 6372 \pm 13 \text{ K}$, Gaia Collaboration et al. 2023). WASP-15 was discovered by West et al. (2009).

For my modelling analysis of WASP-15b, I use the 2.8–5.2 μm NIRSpec/G395H transmission spectra produced by Kirk et al. (2025), the first transmission spectra obtained for this planet. I interpret data reductions from two independent pipelines, **Tiberius** (reduced by JK) and **Eureka!** (reduced by EA), at two spectral resolutions, $R = 100$ and $R = 400$.

5.1.2 The TrES-4 system

TrES-4b is an aligned hot Jupiter ($R_p = 1.838^{+0.081}_{-0.090} R_J$, $M_p = 0.494 \pm 0.035 M_J$, $T_{\text{eq}} = 1795 \text{ K}$; Sozzetti et al., 2015), with a stellar obliquity of $\lambda = 6.3 \pm 4.7^\circ$ (Narita et al., 2010). Like other planets in the BOWIE-ALIGN programme it orbits an F type host star above the Kraft break ($1.81 \pm 0.08 R_\odot$, $1.45 \pm 0.05 M_\odot$, $T_{\text{eff}} = 6295 \pm 65 \text{ K}$, Sozzetti et al. 2015). This, combined with the orbital distance of 0.05 AU, implies the alignment of TrES-4b did not occur post-migration through tidal re-alignment (Spalding & Winn, 2022). TrES-4b was discovered by Mandushev et al. (2007).

I use the 2.8–5.2 μm NIRSpec/G395H transmission spectra of TrES-4b from Meech et al. (2025) in my analysis. I analyse two of the independent data reductions, the **Tiberius** reduction produced by AM, and the **Eureka!** reduction produced by EA, at both $R = 100$ and $R = 400$.

TrES-4b has previously been observed using the WFC3 instrument of HST, yielding a featureless transmission spectrum between 1.1–1.6 μm (Ranjan et al., 2014), but I do not consider these data in my analysis, due to the more marginal precision and increased potential for systematics in HST data.

5.1.3 The NGTS-2 system

NGTS-2b is an aligned hot Jupiter ($R_p = 1.595^{+0.047}_{-0.045} R_J$, $M_p = 0.74^{+0.13}_{-0.12} M_J$, $T_{\text{eq}} = 1468 \text{ K}$; Raynard et al., 2018) with an obliquity of $\lambda = -11 \pm 5^\circ$ (Anderson et al., 2018). NGTS-2b orbits an F5 host above the Kraft break ($1.64^{+0.19}_{-0.22} M_\odot$; $1.70^{+0.05}_{-0.04} R_\odot$, $T_{\text{eff}} = 6480 \pm 90 \text{ K}$, Raynard et al. 2018). NGTS-2b was independently discovered by Raynard et al. (2018) and Anderson et al. (2018), and is also known as WASP-179b.

I interpret the 2.8–5.2 μm NIRSpec/G395H transmission spectra of NGTS-2b from Fairman et al. (2026), using the independent **ExoTiC-JEDI** (reduced by CF) and **Tiberius** (reduced by JK) reductions at both $R = 100$ and $R = 400$. This is the only published atmospheric characterization of NGTS-2b.

5.2 Modelling Setup

I used the open-source atmospheric radiative transfer code **petitRADTRANS** (Mollière et al., 2019) for both atmospheric forward modelling and atmospheric retrievals (Nasedkin et al., 2024) in my analysis of the BOWIE-ALIGN data sets. I used the v2.7.7 version of **petitRADTRANS** for the analysis of WASP-15b and TrES-4b, and the newer v3.1.3 version (Blain et al., 2024) for the analysis of NGTS-2b. The modelling setup is otherwise kept highly consistent between planets, to ensure that fair comparisons can be made.

For the forward modelling and retrieval analyses, I use correlated- k radiative transfer at $R = 1000$, using the opacity sources described in Sec. 4.5.1, with the exception that I use the HITEMP line-list for H_2O (Rothman et al., 2010) for my analyses of WASP-15b and TrES-4b. For equilibrium chemistry models, the composition is interpolated from the same **easyCHEM** (Mollière et al., 2017) grid described in Sec. 4.5.2. As before, the abundance of all atoms, is scaled by metallicity relative to solar (defined by Asplund et al., 2009), except oxygen, where the O/H ratio is scaled by the C/O ratio relative to the C/H ratio.

5.2.1 Retrieval setups

For retrievals, **petitRADTRANS** uses Bayesian nested sampling (Skilling, 2004) implemented through **MultiNest** (Feroz & Hobson, 2008) with **PyMultiNest** (Buchner et al., 2014). For my default retrieval setup, the stellar radius is fixed, and the reference pressure set to 1 mbar, at which the radius and gravity are defined. I therefore use a wide uniform prior on the planet’s radius; as the true reference pressure is unknown, prior knowledge of the radius cannot be used in the retrieval without potentially biasing it. This is the most common approach for atmospheric retrievals (e.g., Alderson et al., 2023; Taylor et al., 2023; Banerjee et al., 2024). My default retrieval setup uses an isothermal temperature profile with a broad prior on limb temperature (typically in the range of 500–2500 K), and a grey cloud deck with a

wide log-uniform prior on cloud-top pressure (encompassing the full simulated pressure range). I use a Gaussian prior on the planet’s gravity, calculated from the radius and mass of the planet.

For each of the three planets considered, I use three different chemical parametrizations: free chemistry, equilibrium chemistry, and hybrid chemistry, as was done for HAT-P-30b in Chapter 4. For the free chemistry retrievals, I permit the log mass mixing ratio to vary from -12 to -0.5 for each molecule. For our chemical equilibrium retrievals, the elemental ratios were parametrized by C/O ratio and metallicity, as previously described, with a uniform prior of 0.1 to 1.5 for C/O, and a log uniform prior of -2 to 3 on metallicity.

In the hybrid chemistry retrieval approach, the abundances of all species are interpolated from an equilibrium grid, as in equilibrium chemistry retrievals, except for certain species, the abundances of which are a free parameter. For the hot Jupiter retrievals of the BOWIE-ALIGN programme, I allow sulfur-bearing species including H_2S , SO_2 , OCS and SO to be free parameters, to allow the model to account for the abundances of these highly photochemically-sensitive species to be pushed out of equilibrium. The exact sulfur species included in the production retrievals was varied based on the constraints provided by preliminary retrievals and forward models, with H_2S and SO_2 always included based on previous detections of these molecules in hot Jupiters (Alderson et al., 2023; Fu et al., 2024). This hybrid approach grants the models sensitivity to a non-solar S/C ratio, such as primordial enrichment or depletion of sulfur. This approach is similar to the retrieval setup of Xue et al. (2024), where key species detected in the high-resolution measurements of Giacobbe et al. (2021) are set as free parameters (CH_4 , C_2H_2 , HCN , and NH_3).

5.3 WASP-15b

5.3.1 Forward models

For an initial interpretation of the WASP-15b dataset, I implemented a grid of equilibrium chemistry simulations, with an isothermal temperature profile and the system parameters of WASP-15b reported in Section 5.1.1, using a metallicity of $0.1\text{--}100\times$ solar and a C/O ratio from 0.1–1.0.

A by-eye comparison of these forward models to the **Tiberius** and **Eureka!** $R = 100$ spectra demonstrated that a super-solar metallicity of $10 - 100\times$ solar and a C/O ratio from 0.2 – 0.8 gives rise to a similar spectrum, shown in the top panel of Figure 5.1. I examined the contributions from each molecule to identify H_2O and CO_2 as the primary opacity sources, with potential minor contributions from CO . However, it was not possible to match the strength of the dominant H_2O and CO_2 features using these clear atmosphere models, as at $< 80\times$ solar metallicity the strength of the H_2O feature is over-predicted, while at $> 50\times$ solar metallicity the CO_2 feature is under-predicted. By extending the grid with variable grey cloud decks and isotherm temperatures, I found that a cooler temperature, clouds (0.1 mbar), or




Figure 5.1: Top panel: Selected synthetic transmission spectra from my `petitRADTRANS` 1D equilibrium chemistry forward model grid, plotted against the Tiberius $R = 100$ spectrum of WASP-15b. The temperature profile is fixed to an isotherm at $T_{\text{eq}} = 1676$ K. I find that a roughly solar C/O ratio, and a metallicity of $> 10\times$ solar is required to give the appropriate feature shapes, identifiable as H_2O and CO_2 , but I was not able to match the feature amplitudes with this simple model. Bottom panel: Selected synthetic transmission spectra from my expanded `petitRADTRANS` 1D equilibrium chemistry forward model grid. This demonstrates that the inclusion of higher altitude clouds at 0.1 mbar (the cloudy model, orange) or fixing the isothermal temperature to a cooler value of 1000 K (the cool limbs model, green), combined with a 10 – 50 $\times$ solar metallicity and C/O = 0.55 allows me to more successfully fit the feature amplitude compared to the original grid with no clouds and equilibrium temperature limbs (top panel). I additionally include a model with 5 ppm of SO_2 (purple dashed line, bottom panel).

some combination thereof was necessary to give a good fit to both features, shown in the bottom panel of Figure 5.1.

No equilibrium chemistry model was able to fit the apparent features at $4.1\ \mu\text{m}$ and $4.9\ \mu\text{m}$. Given the presence of SO_2 in the hot Jupiter WASP-39b (Alderson et al., 2023), visible as a feature at $4.0\ \mu\text{m}$, I tested including a similar abundance (5 ppm) of SO_2 in the model, which fit the feature well (the purple dashed line in the bottom panel of Figure 5.1). To explain the feature at $4.9\ \mu\text{m}$, I tested a wide variety of molecules, finding that both OCS and O_3 had absorption at the correct wavelengths, but gave rise to broader features than the observed peak. I used the **TriArc** package (Claringbold et al., 2023) to demonstrate that given an equilibrium chemistry forward model and the precision of the observation, abundances of ≥ 10 and 1 ppm for SO_2 and OCS respectively would be needed to result in a 3σ detections.

5.3.2 Retrievals

As standard for my BOWIE-ALIGN analyses, I performed equilibrium, free, and hybrid chemistry retrievals on the $R = 100$ spectra from both **Tiberius** and **Eureka!**, with additional hybrid retrievals on the $R = 400$ spectra to probe for potential higher resolution information. I followed the default retrieval setup described in Sec. 5.2, with wide, uniform priors for the planet’s radius ($0.8 - 1.8\ R_J$) and limb temperature ($500 - 3000\ \text{K}$), a stellar radius fixed to $1.477\ R_\odot$, a free grey cloud deck, and Gaussian prior for gravity calculated from the planetary parameters of Bonomo et al. (2017). To help identify a sensible approach to setting the reference pressure, I used an experimental retrieval where the reference pressure was fitted, while the planetary radius was set by a Gaussian using the white-light light-curve transit depth. The resultant reference pressure posterior peaked sharply at 1 mbar, for which I defined the reference pressure in the following retrievals.

For the free and hybrid chemistry retrievals, I included opacity from CH_4 , H_2O , CO , CO_2 , H_2S , SO_2 , and OCS (with the final three species as free parameters for the hybrid retrievals). For the equilibrium chemistry retrievals, I used the same species excluding OCS and SO_2 .

In each retrieval setup, I found that the retrieved parameters from the **Tiberius** and **Eureka!** spectra, at both $R = 100$ and $R = 400$, were completely consistent. I show the best-fit models from each retrieval in Figure 5.2.

From the results of the equilibrium chemistry retrievals I found a C/O ratio for WASP-15b of $0.48^{+0.11}_{-0.16}$ for the **Tiberius** spectrum and $0.53^{+0.09}_{-0.16}$ for the **Eureka!** spectrum, which is consistent with the solar value of 0.55 (using the solar metallicity of Asplund et al., 2009). The metallicity posterior indicated a super-solar composition for WASP-15b, at $18^{+22}_{-8} \times \text{solar}$ and $22^{+27}_{-9} \times \text{solar}$, respectively. These retrievals also yielded a limb temperature of $\sim 900\ \text{K}$, notably colder than the equilibrium temperature of $\sim 1700\ \text{K}$, and weak constraints on the cloud-top pressure, setting a 3σ lower limit of 0.1 mbar.




Figure 5.2: Best fitting retrievals models of WASP-15b using free chemistry (blue), equilibrium chemistry (orange), and hybrid chemistry (equilibrium chemistry plus free S-bearing species, green), with the top panel showing retrievals on the **Tiberius** $R = 100$ reduction and the bottom panel showing retrievals on the **Eureka!** $R = 100$ reduction.

The results of the free chemistry retrievals were mostly consistent, confirming the presence of H_2O and CO_2 by ruling out abundances of less than 1 ppm. The log vertical volume mixing ratios were found to be $-2.83^{+0.59}_{-0.95}$ for H_2O and $-4.40^{+0.73}_{-1.02}$ for CO_2 in the **Tiberius** spectrum, with similar results for the **Eureka!** spectrum. The posteriors showed no firm constraints for other species in the **Tiberius** spectrum, but the posterior corner plot from the **Eureka!** spectrum (see Figure 5.3) highlights that one of CO or OCS is likely present in the spectrum (the posteriors are inconsistent with neither being present to 3σ), but they cannot be distinguished. A suggestion of OCS is also present in the **Tiberius** spectrum in the form of a peaked posterior. The free chemistry retrievals also demonstrated stronger constraints on the cloud-top




Figure 5.3: Posterior distributions from the *petitRADTRANS* free chemistry retrieval on the $R = 100$ transmission spectra of WASP-15b from **Tiberius** (black) and **Eureka!** (red). Species abundances are quoted in mass fractions. The best fitting models are shown in the top right.

pressure, requiring a cloudy atmosphere with a log cloud-top pressure of $-3.19^{+0.81}_{-0.74}$ for the **Tiberius** spectrum.

The best fit to the data came from the hybrid retrievals, based on a reduced $\chi^2_{\nu} = 1.04$ for the **Tiberius** $R = 100$ data and $\chi^2_{\nu} = 1.05$ for **Eureka!** $R = 100$ data, depicted in Figure 5.2. By allowing the sulfur species to depart from equilibrium as free parameters, the features at $4.0 \mu\text{m}$ and $4.9 \mu\text{m}$ were fit with SO_2 and OCS , respectively. This approach also allowed H_2S to be depleted, which can be the result of photochemical destruction, or a low atmospheric S/H ratio. This hybrid approach gave more precise constraints on the C/O ratio, with a retrieved C/O ratio of $0.53^{+0.09}_{-0.15}$ for the **Tiberius** spectrum and $0.56^{+0.07}_{-0.13}$ for the **Eureka!** spectrum. The hybrid retrieval corner plot is shown in Figure 5.4.

In order to investigate the significance of including different opacity sources,




Figure 5.4: Posterior distributions from the *petitRADTRANS* hybrid chemistry retrieval on the $R = 100$ transmission spectra of WASP-15b from Tiberius (black) and Eureka! (red). Species abundances are quoted in mass fractions. The best fitting models are shown in the top right.

I ran additional retrievals and used the Bayesian evidence for model comparison. I ran free chemistry retrievals omitting a single species each time on the **Tiberius** $R = 100$ spectrum and compared the Bayesian evidence as described in Sec. 2.4.3¹, placing the detection significance of CO_2 at 8.9σ , and H_2O at 4.2σ . To compare between different equilibrium and hybrid setups, I repeated the hybrid chemistry retrievals, omitting in turn SO_2 , OCS , H_2S and then all three. A comparison of the Bayesian evidences of these retrievals is found in Table 5.1. In the **Eureka!** spectrum, $\Delta \ln Z = +2.2$ between the equilibrium and hybrid retrievals, equivalent to a 2.6σ preference. This is the sum of small improvements of evidence, each individually insignificant, from including H_2S , SO_2 , and OCS in turn as free parameters. The **Tiberius** spectrum is agnostic between the various hybrid and equilibrium models, with a $\Delta \ln Z$ of only $+0.5$ between the best (hybrid with H_2S and OCS) and worst (hybrid with H_2S , SO_2 , and OCS) models. The $R = 400$ retrievals were completely consistent, with a marginal improvement in evidence for each S-species added with the **Eureka!** spectrum, but changing by less than 0.4 between different **Tiberius** setups.

Table 5.1: The Bayesian evidence ($\ln Z$) and degrees of freedom (DOF) from each of our **petitRADTRANS** retrievals of WASP-15b, including equilibrium chemistry, free chemistry, and hybrid retrievals with equilibrium CH_4 , H_2O , CO , and CO_2 abundances, but the abundance of specified S-bearing species as free parameters. The evidence is very agnostic about the inclusion of various S-species in the **Tiberius** spectrum, but there is a moderate (2.6σ) preference for the hybrid chemistry model in the **Eureka!** spectrum, implying some combination of H_2S depletion, OCS enrichment, and the presence of SO_2 .

Retrieval	Tiberius $\ln Z$	Eureka! $\ln Z$	DOF
Equilibrium	-40.0	-47.5	6
Hybrid (H_2S only)	-39.9	-46.3	7
Hybrid ($\text{H}_2\text{S}+\text{SO}_2$)	-40.0	-46.0	8
Hybrid ($\text{H}_2\text{S}+\text{OCS}$)	-39.8	-45.7	8
Hybrid ($\text{H}_2\text{S}+\text{OCS}+\text{SO}_2$)	-40.3	-45.3	9
Free	-43.0	-48.2	11

5.3.3 Discussion

Atmosphere of WASP-15b

My atmospheric retrievals yield a consistent atmospheric composition between reductions and resolutions, demonstrating a broadly equilibrium composition of $\sim 20\times$ solar metallicity and an approximately solar C/O ratio. The main opacity contributions to the best-fit spectrum from both the hybrid and free retrievals to the **Tiberius** $R = 100$ reduction are portrayed in Fig. 5.5. The most notable difference between retrievals is in the interpretation of minor opacity sources, including grey

¹I note that in Kirk et al. (2025), it is reported that the approximation of Trotta (2008) is used to calculate these detection significances; this is a mistake, and the accurate method described in Chapter 2 was in fact correctly applied.

clouds, CH_4 , OCS , and SO_2 , which have minimal impact on the overall C/O ratio and metallicity results used in diagnosing planet formation. Kirk et al. (2025) favour the interpretation presented by my hybrid retrieval of the Tiberius $R = 100$ data, with a $[\text{M}/\text{H}]$ of $1.34^{+0.28}_{-0.26}$ and C/O ratio of $0.53^{+0.09}_{-0.15}$.

Both the forward models and the free retrievals point towards shallower features than predicted by a clear atmosphere. It is notable that the retrieved isothermal temperature is significantly lower than the equilibrium temperature of 1676 ± 29 K, which is potentially related to these muted features. A lower temperature interpretation can also be caused by the presence of asymmetry in the limb temperature (MacDonald et al., 2020) or the temperature-pressure profile parametrization (Welbanks & Madhusudhan, 2022), and is a common feature of hot Jupiter retrievals, including other BOWIE-ALIGN targets (Meech et al., 2025; Claringbold et al., 2026a). A limb asymmetry analysis of WASP-15b will be the subject of a future BOWIE-ALIGN paper alongside the other targets.

The shallow NRS1 slope further requires an additional opacity source, which is CH_4 in the equilibrium retrieval, and a grey cloud deck in the free retrieval. Between these two opacity sources, I find a grey cloud deck more convincing, due to the improbability of finding CH_4 in hot Jupiters at this instellation (e.g. Fortney et al., 2020). This is also correlated with the retrieval of SO_2 , with the potential feature at $4.0 \mu\text{m}$ being fit with SO_2 or clouds depending on both the retrieval and the reduction used (see Fig. 5.2). As discussed, the Tiberius reduction is agnostic to the inclusion of all three sulfur species, while the Eureka! reduction experiences incremental improvements in Bayesian evidence from the inclusion of each of H_2S (through depletion relative to equilibrium rather than presence), SO_2 , and OCS .

Although not unambiguously favoured by the retrieval analysis, the presence of SO_2 in the atmosphere of WASP-15b, particularly in the presence of a super-solar metallicity atmosphere, fits with the photochemical models of Tsai et al. (2023). The 5 ppm abundance injected in the forward models to fit the magnitude of the apparent feature at $4.0 \mu\text{m}$ is consistent with the SO_2 abundances detected in the super-solar metallicity atmospheres of WASP-39b (0.5–25 ppm in Powell et al., 2024) and WASP-107b (6–9 ppm in Dyrek et al., 2024; Sing et al., 2024). The temperature and retrieved metallicity of WASP-15b place it close to the SO_2 shoreline posited by Crossfield et al. (2025), making a deeper investigation into the presence of SO_2 in WASP-15b an important test in tracing this shoreline. Confirming the presence of or a tighter upper limit on the SO_2 abundance in WASP-15b would require follow-up JWST observations, either using the MIRI instrument to detect the SO_2 feature at $7.4 \mu\text{m}$, a repeated observation at the same wavelength to improve the precision, or some combination thereof.

A notable feature of the WASP-15b spectrum is a peak at $4.9 \mu\text{m}$, spanning multiple wavelength bins in the $R = 400$ datasets. Kirk et al. (2025) performed a number of tests on the data to suggest that the feature is not an instrumental artefact and therefore astrophysical in origin. I injected a number of species into




Figure 5.5: Figure showing the contributions of opacities from the best fitting `petitRADTRANS` hybrid chemistry (equilibrium plus free S) retrieval (top panel) and free chemistry retrieval (bottom panel) to the Tiberius $R = 100$ spectrum of WASP-15b. The difference in Bayesian evidences, $\Delta \ln Z$, is 2.7 in favour of the hybrid model (top panel).

the forward models in an attempt to explain this feature, testing CH_4 , CO , H_2S , HCN , O_3 , PH_3 , NH_3 , C_2H_2 , SO_2 , NO , HC_3N , H_2CO , CS_2 , SO , SH , OCS , OH , AlH , AlO , CaH , CrH , FeH , MgH , MgO , NaH , SiO , and SiO_2 . I found that only OCS (Owens et al., 2024, linelist from) and ozone (O_3) were able to fit the wavelength of the observed feature. The width of O_3 absorption at $4.9\ \mu\text{m}$ is too broad to fit the feature, furthermore the molecule would be extremely unstable in a H_2 -rich atmosphere, making accumulation of a sufficient abundance unphysical.

Including OCS in the atmospheric retrievals does improve the fit in most combinations of retrievals and reductions (by a maximum Bayes factor of 1.8), including with the additional resolving power of $R = 400$. However, I do note that the width of the model OCS feature is also broader than the observed feature, as can be seen in Fig. 5.2. To investigate whether this feature could therefore be arising from an OCS feature at a lower temperature or pressure (and therefore with less thermal or pressure broadening), I injected OCS localized between 1 and $100\ \mu\text{bar}$ into a forward model atmosphere and varied the temperature and OCS abundance in an attempt to fit the feature. I was unable to fit the width of the observed feature in the $R = 400$ spectrum this way, even using an unphysically low temperature of 300 K with an unphysically high 2% mixing ratio of OCS . Furthermore, while OCS is anticipated to be present in moderate quantities in hot Jupiter atmospheres, it is expected to be destroyed photochemically by UV radiation and therefore not detectable at the altitudes probed by transmission spectroscopy (Tsai et al., 2021; Tsai et al., 2023). The interpretation of OCS is therefore uncertain, and the feature may instead arise from a molecule not yet tested with commonly available line lists.

Implications for formation

The atmospheric retrieval analysis, regardless of the precise details on the minor opacity sources, favours a super-solar metallicity and solar C/O ratio for the atmosphere of WASP-15b, with $\text{C/O} = 0.38\text{--}0.62$ and $[\text{M}/\text{H}] = 1.08\text{--}1.62$ in the preferred retrieval. WASP-15b is a misaligned hot Jupiter, suggesting that it formed in the outer disc beyond the water-ice line and migrated inwards after disc dispersal via high eccentricity migration (Muñoz et al., 2016). Kirk et al. (2025) discuss the implication of this retrieved composition for interpreting the formation history of WASP-15b, summarised below, but I emphasise the importance of combining the entire sample when attempting to make planet formation arguments given the degeneracies present (Penzlin et al., 2024).

As with the other individual planets from BOWIE-ALIGN, constraining the formation of an individual planet from C/O ratio is highly degenerate. In the disc-free migration scenario, a super-stellar metallicity composition is suggestive of the late accretion of solid material enriching the oxygen content of the planet (Booth et al., 2017; Schneider & Bitsch, 2021a; Schneider & Bitsch, 2021b). The tentative presence of sulfur-bearing molecules in the atmosphere of WASP-15b also supports this interpretation, with sulfur being bound in the solid material of the disc at

the formation location suggested by high-eccentricity migration (Chachan et al., 2023). The large mass of WASP-15b, greatly exceeding the pebble isolation mass of $\sim 20M_{\oplus}$, implies that planetesimal accretion would be the dominant source of late solid accretion.

An alternative explanation for the observed atmospheric composition would be through more-or-less in situ formation, where the high metallicity and sulfur content could be explained by enrichment by smaller pebbles in the inner disc (Morbiddelli et al., 2023) and the presence of sulfur in the accreted gases at high temperatures. However, in situ formation is disfavoured generally for hot Jupiters (e.g. Dawson & Johnson, 2018), and does not explain the misaligned orbit of the planet.

5.4 TrES-4b

5.4.1 Retrievals

I performed atmospheric retrievals to interpret the TrES-4b dataset, following the default setup described in Sec. 5.2. I used free chemistry, equilibrium chemistry, and hybrid chemistry retrieval setups. I assume an isothermal pressure-temperature profile, with a wide uniform prior on the isotherm temperature from 500–3000 K. I also used a wide uniform prior for the planetary reference radius ($0.8\text{--}1.8 R_J$) and a Gaussian prior for gravity based on the planetary parameters of (Sozzetti et al., 2015), defined at a fixed reference pressure of 10^{-3} bar. This reference pressure was determined using a retrieval with free reference pressure and a Gaussian prior on the radius, based on the mean white-light transit depth from both detectors. I also included an opaque grey cloud deck, with a log-uniform prior on the cloud-top pressure, from 10^{-6} bar to 10^2 bar, and line opacity from H_2O , CH_4 , CO , CO_2 , HCN , H_2S , as well as SO_2 for the free and hybrid chemistry retrievals.

I ran the retrievals on the *Tiberius* $R = 100$ and $R = 400$ spectra, as well as the *Eureka!* spectrum at $R = 100$. Across all three retrieval setups, I retrieve consistent radius, gravity, isotherm temperatures, and cloud-top pressure posteriors. I retrieved limb temperatures of $T = 1200 \pm 100$ K, which are notably much lower than equilibrium temperature of TrES-4b ($T_{\text{eq}} = 1800$ K), is a commonly observed feature in 1D isothermal retrievals (see e.g., Welbanks et al., 2019; Kirk et al., 2025). Potential causes for this include unresolved limb asymmetries (MacDonald et al., 2020) or the choice of PT profile parametrization (Welbanks & Madhusudhan, 2022). All the retrieval setups suggest a relatively cloud-free upper atmosphere, by placing a 3σ upper limit on the grey cloud-top pressure of 10^{-3} bar.

All reductions and resolutions were consistent in the free chemistry retrievals, revealing contributions from H_2O , CO , and CO_2 , with retrieved volume mixing ratios of $\log X_{\text{H}_2\text{O}} = -3.29^{+0.57}_{-0.59}$, $\log X_{\text{CO}} = -3.71^{+0.85}_{-0.84}$ and $\log X_{\text{CO}_2} = -7.21^{+0.56}_{-0.51}$ respectively. There are also hints of H_2S abundances of 10 ppm, but the posterior is not confidently defined. The 3σ upper limits on SO_2 and HCN abundances are set at 4 ppm and 2 ppm respectively. The posterior distributions are presented in




Figure 5.6: Posterior distributions from the `petitRADTRANS` free chemistry retrieval on the $R = 100$ transmission spectra of TrES-4b from *Tiberius* (pink) and *Eureka!* (blue). Species abundances are quoted in mass fraction. The best-fitting models are shown in the top right.

Fig. 5.6.

In the equilibrium chemistry retrievals, I retrieve a sub-solar C/O ratio posterior of $0.33^{+0.13}_{-0.11}$ and a sub-solar metallicity of $-0.33^{+0.24}_{-0.19}$ for the *Tiberius* $R = 400$ spectrum. The $R = 100$ values are consistent, with slightly reduced precision. The *Eureka!* reduction demonstrated a slightly higher, but still consistent, C/O ratio of $0.43^{+0.17}_{-0.14}$. The resulting posterior distributions from these retrievals are depicted in Fig. 5.7.

For the hybrid chemistry retrievals, the sulfur-bearing species H_2S and SO_2 are free parameters with the same broad, uniform priors as for free chemistry, while the abundances of the other species computed with equilibrium chemistry. As in other BOWIE-ALIGN analyses, this allows the sulfur-bearing species abundances to be in disequilibrium and permit measurement of elemental sulfur enrichment or




Figure 5.7: Posterior distributions from the **petitRADTRANS** equilibrium chemistry retrieval on the $R = 100$ transmission spectra of TrES-4b from Tiberius (pink) and Eureka! (blue). The best-fitting models are shown in the top right.




Figure 5.8: Posterior distributions from the `petitRADTRANS` hybrid chemistry retrieval on the $R = 100$ transmission spectra of TrES-4b from Tiberius (pink) and Eureka! (blue). Species abundances are quoted in mass fraction. The best-fitting models are shown in the top right.

depletion. As with the free chemistry retrieval, I see no evidence for SO₂, and hints of H₂S. As seen with WASP-15 b (5.3.2), the hybrid retrieval results in more precise constraints on the C/O ratio, yielding a C/O posterior of $0.30^{+0.11}_{-0.09}$, as well as a somewhat lower metallicity of -0.41 ± 0.18 . These hybrid retrieval posteriors are presented in Fig. 5.8. In the $R = 400$ retrievals, the Bayesian evidence for the equilibrium and hybrid retrievals are identical; I therefore have no cause to favour one parametrization over the other. I attribute this to the absence of SO₂ in the spectrum, as well as the poor constraints on H₂S. In the $R = 100$ retrievals, the hybrid chemistry is slightly disfavoured; this difference could be due the increased sensitivity to H₂S at higher resolutions, which could be realistically depleted by photochemistry.

5.4.2 Discussion

Atmosphere of TrES-4b

I obtain a consistent chemical composition for TrES-4b based on my retrieval analysis, with spectral features from H_2O , CO , and CO_2 consistent with an equilibrium chemistry atmosphere with a metallicity less than the stellar host metallicity of 0.28 ± 0.09 , and a sub-solar C/O ratio. The retrieved abundances are consistent between retrieval types, and the Bayesian evidence supports the usage of equilibrium or hybrid chemistry over free chemistry, supporting the inference that the atmosphere is in chemical equilibrium. Meech et al. (2025) do not highlight a particular retrieval analysis for interpretation in their conclusion, due to this consistency between retrievals.

Spectral muting due to aerosols is limited, with no evidence for strong clouds from the retrievals, placing an approximation 2σ upper limit on the cloud-top pressure of ~ 30 mbar. I do retrieve a lower isotherm temperature than expected from the equilibrium temperature, which is evidence of some degree of muted features. As discussed, the temperature profile parametrization is one potential cause for this (Welbanks & Madhusudhan, 2022). I therefore also ran retrievals using a Guillot temperature profile (Guillot, 2010), an analytic pressure-temperature profile where the equilibrium temperature T_{eq} , internal temperature T_{int} , ratio of optical and infrared opacities γ , and infrared opacity $\log \kappa$ are all free parameters. Using this retrieval parametrization gave consistent results for the chemical and aerosol opacities, while demonstrating somewhat broader constraints for the temperature profile, as depicted in Fig. 5.9. I determine from this that while an isothermal temperature profile may be biasing the inferences on the temperature profile, being overly optimistic on the certainty, it ultimately does not impact the chemical abundances used for atmospheric retrieval. Another potential cause mentioned is unresolved limb asymmetries, an analysis of which will be conducted in a future work examining limb asymmetries in the complete sample of BOWIE-ALIGN planets.

AM also provided a full set of atmospheric retrievals using the *Poseidon* retrieval code, presented in Meech et al. (2025). The results of these retrievals are fully consistent with my results from *petitRADTRANS*, accounting for the difference in definition in $[\text{M}/\text{H}]$ between the two codes. The *Poseidon* retrievals also include a stellar contamination analysis, and find no additional evidence to support models including stellar contamination. Constraining stellar contamination at these wavelengths is challenging, due to the strong degeneracies in contamination parameters (Pinhas et al., 2018; Rackham et al., 2018). AM also tested for an offset between the detectors in her retrievals, finding no evidence for an offset. I consequently did not test for an offset in the *petitRADTRANS* retrievals.

Both the free and hybrid retrieval analyses find minimal evidence of opacity from sulfur-bearing species including H_2S or SO_2 , with a clear non-detection of SO_2 as expected from a low metallicity atmosphere (Tsai et al., 2023; Crossfield et al.,




Figure 5.9: Confidence intervals on temperature profile posteriors for TrES-4b, based on `petitRADTRANS` equilibrium atmospheric retrieval using a Guillot temperate profile parametrization, where T_{int} , T_{eq} , γ , and $\log \kappa$ are all free parameters. The temperature profile is consistent with isothermal in the transmission region (~ 1 – 100 mbar), but the constraints are broader. This demonstrates the poor ability of transmission spectroscopy in constraining the temperature profile, but does not impact the chemical inferences for TrES-4b.

2025), and weak evidence for H_2S consistent with equilibrium abundances. Given this lack of disequilibrium chemistry, of the `petitRADTRANS` retrievals presented here, I favour the usage of the *Tiberius* $R = 400$ equilibrium retrievals. This also follows the preference for $R = 400$ established as the BOWIE-ALIGN programme evolved (Ahrer et al., 2025b; Claringbold et al., 2026a; Fairman et al., 2026).

Interpretation for formation

The consistent interpretation for the atmosphere of TrES-4b yields a sub-stellar metallicity and sub-solar C/O ratio, with $\text{C/O} = 0.22$ – 0.46 and $[\text{M/H}] = -0.52$ – -0.09 in my favoured equilibrium retrieval. As an aligned hot Jupiter, it is thought that TrES-4b evolved to its current location through disc migration. This composition is significantly different from the super-solar metallicity composition of WASP-15b, but is similar to the results for HAT-P-30b (see Chapter 4) – both of which are misaligned planets, further demonstrating the need to compare across samples of planets, rather than individual planets. Like for HAT-P-30b, a sub-stellar metallicity and sub-solar C/O composition challenges traditional formation models, where metallicity is anti-correlated with C/O ratio, with a higher metallicity composition generally representing solid accretion of oxygen-rich ices (e.g., Öberg et al., 2011; Bergin et al., 2024). Meech et al. (2025) consider two methods for this, which I




Figure 5.10: Diagram depicting two potential formation scenarios that can explain a low C/O ratio, low metallicity composition for TrES-4b. On the left is a gas-dominated accretion scenario, where the planet accreted its observable atmosphere interior to the water-ice line but exterior to the soot line, resulting in a low C/O ratio gas. On the right is a volatile trapping scenario, where the trapping of volatile solids in the outer disc leaves the inner disc depleted in solids, requiring less solid accretion to achieve a low C/O ratio. Figure reproduced from Meech et al. (2025).

summarize below.


One potential scenario is the accretion of oxygen-rich gases, requiring that the observable atmosphere be accreted at a location interior to the water-ice line (where H₂O is therefore in the gas) but exterior to the soot line (where carbon grains remain in the solid). This implies accretion of the atmosphere in the inner disc, or potentially later accretion further out in the 1–10 au region following destruction of CO and CO₂ (Eistrup et al., 2018). In the alternative scenario presented in Meech et al. (2025), the inner disc could become depleted in metals through the trapping of volatiles species in the outer system (Mah et al., 2024), allowing the planet to obtain a low C/O ratio with reduced accretion of solids. These two scenarios are depicted in Fig. 5.10.

5.5 NGTS-2b

5.5.1 Retrievals

I used the default `petitRADTRANS` atmospheric retrieval setup outlined in Sec. 5.2 to aid in the interpretation of the transmission spectra of NGTS-2b. As with other data sets, I used equilibrium chemistry and hybrid chemistry setups, and I also performed free chemistry setups with (labelled as “Model A”) and without (labelled as “Model B”) opacity from SO. I performed retrievals on all four reductions synthesised by Fairman et al. (2026), from both `Tiberius` and `ExoTiC-JEDI` at $R = 400$, except for the free chemistry Model II retrievals which were only performed on the $R = 400$ spectra.

Following other analyses in `BOWIE-ALIGN`, I use a wide uniform prior on the planetary radius ($0.8\text{--}2.2 R_J$) and a Gaussian prior for gravity calculated using the planet’s mass and radius (Raynard et al., 2018), defined at a reference pressure of 1 mbar. I use a grey cloud deck with a log-uniform prior on cloud-top pressure from 10^{-6} to 10^2 bar to parametrize the impact of aerosols. I permit for a detector offset in the retrievals, using a uniform prior of up to 200 ppm in either direction



retrievals (Bayes factor of 1.8), so I therefore ran an equivalent set of free retrievals without including opacity from SO (Model B) in an attempt to find a more physical solution.

Without SO as an opacity source, the free retrievals instead tend to a very high MMW atmosphere, bunching up around the upper bound of the H_2O prior (equivalent to a mixing ratio of $\sim 4\%$), and with a total mean molecular weight of 4.4. The CO_2 abundance is well constrained with $\log X_{\text{CO}_2} = -4.90^{+0.71}_{-0.87}$, and there is slight evidence for a high abundance of CO. The retrieved temperature is also low with this model (914^{+174}_{-145} K), although the discrepancy is less extreme. There is strong evidence for muted spectral features due to clouds, with tight constraints on the cloud-top pressure at $\log P_{\text{cloud}} = -2.43^{+1.11}_{-0.72}$.

Comparing my Model B retrievals to the Model II free chemistry *Poseidon* retrievals (see Fairman et al., 2026, for more details) shows that these results actually constitute the tail end of an extremely high MMW mode, demonstrating that the high abundances arise as a result of the retrieval attempting to suppress the atmospheric scale height to match the feature shapes. However, these high metallicities are also unphysical as demonstrated by the upper limit on the planetary metal content placed by interior structure models included in Fairman et al. (2026). CF ran an additional set of *Poseidon* retrievals restricting the maximum MMW to 3 based on results of the interior models, which yield similar results to the Model B *petitRADTRANS* retrievals.

My *petitRADTRANS* equilibrium chemistry results are consistent between reductions, demonstrating a cooler-than-equilibrium isotherm temperature of ~ 1000 K, and the same 56 ± 20 ppm offset as found in the free retrievals. This approach yields poor constraints on the cloud-top pressure. The C/O ratio is found to be consistently sub-solar, at $0.36^{+0.18}_{-0.16}$ in the *ExoTiC-JEDI* $R=400$ reduction and $0.29^{+0.17}_{-0.13}$ in the *Tiberius* $R=400$ reduction. On the other hand, the metallicity does vary between reductions, with the lowest value of $[\text{M}/\text{H}] = 0.23^{+0.55}_{-0.39}$ in the *Tiberius* $R=400$ reduction and the highest value of $[\text{M}/\text{H}] = 0.69^{+0.59}_{-0.46}$ in the *ExoTiC-JEDI* $R=400$ reduction. This amounts to a difference of about triple the median carbon and twice the median oxygen content between the *ExoTiC-JEDI* and the *Tiberius* reductions, but still consistent within 1σ . The results from the hybrid retrievals follow their equilibrium counterparts closely, demonstrating no evidence that would favour considering any of the three sulfur species: H_2S , SO, or SO_2 , as free parameters (and therefore out-of-equilibrium, in the case of H_2S , or present, in the case of SO and SO_2).

5.5.2 Discussion

Interpreting the retrieval analysis of NGTS-2b is not straightforward, with the naive first approach to free retrievals (Model A of *petitRADTRANS*, Model I of *Poseidon*) yielding a completely unphysical solution, requiring $100\times$ solar enrichment of sulfur. This is verified by a photochemical analysis presented in Fairman et al. (2026), where




Figure 5.12: Free chemistry retrievals of NGTS-2b on the ExoTiC-JEDI $R=400$ spectrum for the three models tested with **Poseidon** (Models I and II) and two models tested with **petitRADTRANS** (Models A and B): Model I and Model A including SO opacity (purple and pink); Model II and Model B excluding SO opacity (blue and orange); Model II excluding SO opacity and limiting the MMW to $\mu_m < 3$ (light green). Due to **Poseidon** and **petitRADTRANS** retrieving offsets for NRS1 and NRS2 respectively, the **petitRADTRANS** offset has been inverted to directly compare to the **Poseidon** retrieved value. Figure reproduced from Fairman et al. (2026).




Figure 5.13: Posterior distributions from the `petitRADTRANS` equilibrium chemistry retrievals of NGTS-2b.




Figure 5.14: Figure showing the contributions of opacities to the best fitting `petitRADTRANS` equilibrium chemistry retrieval of NGTS-2b.

it is demonstrated that an observable abundance of SO_2 should be present in any scenario with sufficient SO , making the non-detection of SO_2 evidence against this solution.

Furthermore, even when excluding SO (Model B of `petitRADTRANS`, Model II of *Poseidon*), the interpretation is complicated by degeneracies present between MMW, temperature, and cloud-top pressure in muting spectral features. This is similar to that observed in HAT-P-30b (see Chapter 4) and in the non-`petitRADTRANS` retrievals of WASP-15b (Kirk et al., 2025). In NGTS-2b this is more severe however, with a physical scenario only obtainable by placing physical upper limits on the MMW. Furthermore, a phase space analysis performed by CF (outlined in Fairman et al., 2026) show that the G395H data are effectively uninformative on the H_2O abundance and therefore unable to constrain the atmosphere. It is clear that the spectral muting present in NGTS-2b, which could be the result of clouds or potentially unresolved limb asymmetry (MacDonald et al., 2020), puts it in a highly degenerate part of the exoplanet parameter space that requires additional observation to obtain a definitive atmospheric composition. Expansion of the wavelength coverage of observation of NGTS-2b will be provided by both NIRISS/SOSS as part of the JWST programme GO 5924 (PI: Sing) and ground-based observations as part of LRG-BEASTS, will be able to provide improved constraints on both the H_2O abundance and the cloud-top pressure.

The reduced degeneracy of equilibrium retrievals does however present a stronger possibility for meaningful interpretation, but I do emphasize that without consistency between free and equilibrium retrievals, the atmosphere of NGTS-2b should not be treated as a solved quantity. Leveraging the consistent C/O and




Figure 5.15: Comparison of posterior distributions of metallicity, C/O ratio, isotherm temperature, and cloud-top pressure obtained from `petitRADTRANS` equilibrium retrievals of NGTS-2b.

metallicity obtained by all eight `petitRADTRANS` equilibrium and hybrid retrievals (Fig. 5.15), it can be seen that NGTS-2b has an approximately solar to super-solar metallicity, and sub-solar to solar C/O ratio. Due to these relatively weak constraints, the composition of NGTS-2b could be the product of a variety of formation scenarios, which therefore limits our ability to place constraints on the formation history of NGTS-2b as an individual planet. However, planet formation arguments can be strengthened when combined into the larger BOWIE-ALIGN sample.

5.6 Remaining Sample

KELT-7b

KELT-7b (Bieryla et al., 2015) is the last of the five planets obtained as part of JWST programme 3838 present here. An aligned ($\lambda = -10.6 \pm 0.3^\circ$, Tabernero et al., 2022) hot Jupiter with a slightly higher temperature ($T_{\text{eq}} = 2047 \pm 27$), KELT-7b was




Figure 5.16: Transmission spectrum of KELT-7b with the spectral contributions from the best fitting model obtained from an equilibrium chemistry retrieval, as well as the best fitting models from both the free and equilibrium chemistry retrievals. Figure reproduced from Ahrer et al. (2025b).

found to have very weak spectral features due to either a high cloud deck or low atmospheric metallicity, limiting the ability to place constraints on the C/O ratio and metallicity (Ahrer et al., 2025b, depicted in Fig. 5.16). The higher temperature and irradiation of KELT-7b has been implicated in causing water dissociation (Lewis et al., 2020) through photochemical and collisional processes, with the presence of H^- absorption detected in HST observations (Pluriel et al., 2020; Gascón et al., 2025), which could impact the inferred elemental ratios used to constrain planet formation.

WASP-94 Ab

A NIRSpec G395H transmission spectrum for the misaligned ($\lambda = -123.0 \pm 3.0^\circ$, Ahrer et al., 2024) hot Jupiter WASP-94 Ab (Neveu-VanMalle et al., 2014) was obtained through the separate JWST programme Go 3154 (PI: Ahrer). This built upon previous optical ground-based data, including low resolution data from the LRG-BEASTS programme (Ahrer et al., 2022) and high-resolution data from HARPS (Ahrer et al., 2024), which demonstrated a relatively cloud free atmosphere and a super-Rayleigh scattering slope. The G395H transmission spectrum demonstrated robust detections of H_2O and CO_2 , as well as evidence for H_2S and CO , favouring a sub-stellar C/O ratio of $0.45^{+0.089}_{-0.13}$ and a metallicity consistent with the host star (Ahrer et al., 2025a, shown in Fig. 5.17), which the authors find is best explained through either pebble or planetesimal accretion combined with the strong migration anticipated for a misaligned planet.

WASP-94 Ab has also been observed with NIRISS/SOSS as part of JWST programme GO 5924 (PI: Sing), demonstrating absorption due to H_2O and strong asymmetries between limbs attributed to aerosols (Fu et al., 2025b). Due to the lack




Figure 5.17: Transmission spectrum of WASP-94 Ab with the spectral contributions from the best fitting model obtained from an equilibrium chemistry retrieval, as well as the best fitting models from both the free and equilibrium chemistry retrievals. Figure reproduced from Ahrer et al. (2025a).

of carbon-bearing species in the NIRISS wavelength range, Fu et al. (2025b) do not report a C/O ratio, and they find the metallicity is highly sensitive to the treatment of the limb asymmetries.

WASP-17b

WASP-17b (Anderson et al., 2010) is a hot Jupiter in a retrograde orbit ($\lambda = -148.5^{+4.2}_{-5.4}$ ° Triaud et al., 2010), making it a misaligned planet in the BOWIE-ALIGN sample. WASP-17b has been extensively observed using JWST, using multiple instruments to achieve full wavelength coverage in both transmission and emission, as part of the JWST-TST DREAMS programme GTO 1353 (PI: Lewis). These observations have demonstrated quartz clouds (Grant et al., 2023b), nonuniform dayside emission (Valentine et al., 2024), nightside emission and SO₂ (Lustig-Yaeger et al., 2025), and a precise H₂O abundance (Louie et al., 2025). However, the G395H transmission spectrum analogous to the other BOWIE-ALIGN data sets, which is best suited to retrieving the elemental ratios necessary to probe planet formation, is yet to be published.

HD 149026b

The emission spectrum of the aligned ($\lambda = 12.0 \pm 7.0$, Albrecht et al., 2012) planet HD 149026b (Sato et al., 2005) was observed using JWST NIRSpec G395H through the programme GTO 1274 (PI: Lunine). HD 149026b was found to have a strongly super-solar metallicity ($59\text{--}276\times$ solar) and C/O ratio (0.84 ± 0.03) (Bean et al., 2023), with robust detections of H₂O and CO₂. As emission spectroscopy probes a

different part of the atmosphere than transmission spectroscopy, these inferences are not necessarily direct comparable to the constraints obtained for the other planets, which are in transmission. There is currently no scheduled transmission spectroscopy observation for HD 149026b.

5.7 Conclusions

I have provided atmospheric modelling analyses, mainly using atmospheric retrievals, to interpret the transmission spectra of WASP-15b, TrES-4b, and NGTS-2b. Combined with HAT-P-30b (Chapter 4) this represents half of the total BOWIE-ALIGN sample. For each of these planets I have obtained values for the metallicity and C/O ratio, which are the quantities that will be used to draw inferences for planet formation.

WASP-15b shows a clear super-solar metallicity, solar C/O composition, with uncertain minor opacity contributions from OCS and SO₂ to explain spectral features at 4.9 μm and 4.0 μm respectively, as well as either CH₄ or clouds to explain some level of spectral muting. The misaligned nature of the planet, high metallicity, and sulfur chemistry present evidence for planetesimal accretion.

TrES-4b demonstrates a highly consistent and well-explained composition with sub-stellar metallicity and carbon depletion, and no evidence for clouds. This composition challenges traditional formation scenarios, but can be explained with oxygen-rich gas accretion or metal depletion of the inner disc by volatile trapping. HAT-P-30b exhibits a similar composition, but is misaligned, whereas TrES-4b is aligned, which suggests that contrasting formation mechanisms can give rise to the same atmospheric composition.

NGTS-2b on the other hand defies concrete characterization due to the strong degeneracies present in the retrieval analysis. As the spectral features are highly muted, it is likely that this degree of spectral muting is responsible for the degeneracy. While it is difficult to present a concrete composition, equilibrium retrievals do suggest a solar to sub-solar C/O ratio and solar to super-stellar metallicity.

Some notable traits that are present in all the BOWIE-ALIGN analyses presented here are a tendency to retrieve a notably lower than equilibrium temperature, and degenerate solutions based on metallicity. The two main explanations for the former feature are unresolved limb asymmetries (MacDonald et al., 2020), and the choice of temperature profile parametrization (Welbanks & Madhusudhan, 2022). However, multiple instances of non-isothermal temperature profiles have been used in the examination of these planets (e.g., Meech et al., 2025; Fairman et al., 2026), and in no cases have they improved the Bayesian evidence or altered the chemical constraints obtained. The question of limb asymmetries will be examined for the homogeneous BOWIE-ALIGN sample in a future publication, with evidence already found in WASP-94 Ab (Ahrer et al., 2025a; Fu et al., 2025b).

I do not attempt an inter-comparison of BOWIE-ALIGN results beyond what

I have covered, as this will be the subject of two upcoming publications, performing a colour-band analysis and complete retrieval analysis respectively to determine differences between the misaligned and aligned populations, led by JK and EA respectively. These results will be further enriched by the inclusion of additional datasets, including NIRISS/SOSS observations of NGTS-2b, HAT-P-30b, and KELT-7b obtained in the JWST programme GO 5924 (PI: Sing), as has already been done for WASP-94 Ab (Fu et al., 2025b).

Chapter 6

Atmospheric Modelling of Temperate Jupiters

Declaration

This chapter forms the basis of the paper “Mapping the Parameter Space of Temperate Jupiter Atmospheres”, which is currently in preparation. I performed all the analysis and writing for this chapter. The adaptation of *Virga* to handle S_8 condensation was performed by Matt G. Lodge.

6.1 Temperate Jupiters

As discussed in Chapter 1, the atmospheric characterization of giant planets to date has mainly focussed on highly irradiated hot Jupiters (as covered in Chapters 3, 4, and 5), as well as directly imaged self-luminous planets on distant orbits (e.g. Konopacky et al., 2013), and the weakly irradiated gas giants of our Solar System (e.g. Atreya et al., 1999). The exoplanet population amenable to atmospheric remote sensing has previously been limited by the precision and wavelength coverage of our observatories and a sparsity of suitable candidates. Yet very recently, the increase in precision granted by JWST combined with an abundance of recent detections from the follow-up of TESS planet candidates (e.g., Cañas et al., 2020; Hobson et al., 2023; Gill et al., 2024; Ulmer-Moll et al., 2025) allows us to explore a new population of cooler transiting planets, temperate Jupiters.

Temperate Jupiters bridge the gap between hot Jupiters, directly imaged exoplanets, and the Solar System gas giants. With intermediate temperatures around that of Earth (~ 300 K), they are also laboratories for the atmospheric physics present in habitable atmospheres, including processes such as photochemistry, water cloud formation, and haze microphysics, and could even be host to habitable satellites. Five gas giants within this temperate range have been observed with JWST: Kepler-86b (Wang et al., 2013), TOI-1899b (Cañas et al., 2020), TOI-2525c (Trifonov et al., 2023), TOI-199b (Hobson et al., 2023), and TOI-6894b (Bryant

et al., 2025) all using the NIRSpec instrument (GO 3235, 4227, 4938, 5177 and 8696 respectively), with none of these observations yet published as of 2026 May (observations of TOI-199b by Bello-Arufe et al., 2025, are under review, but their results are compromised by a target acquisition issue that resulted in reduced precision that only permits the detection of spectral features from CH₄).

Furthermore, the formation and migration history of gas giant planets can in principle be understood through atmospheric characterization via the retrieval of elemental abundance ratios like C/O (Öberg et al., 2011; Helling et al., 2014; Madhusudhan et al., 2014; Booth et al., 2017; Mollière et al., 2022), as well as relative nitrogen and sulfur content (Piso et al., 2016; Cridland et al., 2020; Turrini et al., 2021). The equilibrium temperatures of temperate gas giants include a regime warm enough to avoid water condensation and hence depletion from the photosphere like in Jupiter, but cool enough that the dominant bearer of nitrogen is ammonia. NH₃ has strong absorption features that can be detected with transmission spectroscopy, in contrast to diatomic N₂ gas, which is the dominant carrier of nitrogen in warmer planets but isn't detectable using spectroscopic observations. This makes temperate Jupiters ideal candidates for simultaneously measuring the C/N/O ratio (Encrenaz et al., 2018).

JWST transmission spectroscopy observations should be sensitive to various physical and chemical processes, which can be simulated using numerical models. Chemical kinetics modelling is necessary to accurately obtain elemental abundance ratios based on observed abundances; Ohno & Fortney (2023a,b) demonstrate that this is particularly true in the case of the nitrogen abundance, which cannot be simply equated to the ammonia abundance due to thermochemical production of undetectable N₂ gas in the deep atmosphere, as well as photochemical destruction of NH₃ and production of HCN in the upper atmosphere. Cloud condensation and haze modelling can also be used to interpret transmission spectra and connect observed continuum features and scattering slopes to the composition and location of aerosols (e.g., Kreidberg et al., 2014a; Dyrek et al., 2024).

While hot Jupiters have been the subject of extensive physical modelling, work on temperate gas giants is much rarer, and generally focussed on young, directly imaged planets (e.g., Moses et al., 2016; Zahnle et al., 2016). Fortney et al. (2020) calculated 1D pressure-temperature profiles with chemical disequilibrium from vertical mixing to explore the onset of chemical transitions in both carbon (CH₄-CO) and nitrogen (NH₃-N₂) extending down to equilibrium temperatures of 200 K. Hu (2021) consider the effect of photochemistry on the atmospheres of both cold and temperate gas giants, including using Kepler-86b as a case study. In particular, they find that the stellar host spectral type has a noticeable impact on the transmission spectrum, with depletion of NH₃ in G stars that is not present in the case of an M dwarf host. Mukherjee et al. (2025) explore the chemistry of the entire parameter space of irradiated planets, including the sensitivity to various parameters including vertical mixing, internal temperature, and carbon-to-oxygen ratio, while Crossfield

et al. (2025) focus on transitions in sulfur chemistry. Both of these analyses extend down to temperatures of 400 K. Gao et al. (2017) propose sulfur hazes as a dominant source of aerosol opacity below 325 K. H_2O condensation, and hence the formation of water clouds, occurs at cooler temperatures around 250 K (MacDonald et al., 2018; Mang et al., 2022). The photochemical haze and mineral clouds responsible for attenuation of spectral features in warm exoplanets are expected to have less of an impact below 400 K (Brande et al., 2024).

In this chapter, I use a grid of 1D physical models to explore trends and transitions across the temperate Jupiter parameter space, creating the outline of a map to aid in the planning and interpretation of future observations. My modelling setup includes a hierarchical approach to radiative-convective equilibrium, chemical kinetics, and aerosol formation, described in Section 6.2. I present the results of my modelling in Section 6.3, including the chemical profiles, aerosols, and simulated transmission spectra. I put this together to arrive at a physically-motivated and observationally-testable definition of the temperate Jupiter parameter space in Section 6.4. This constitutes a summary of the transitions that define the external boundaries, as well as internal transitions that sub-divide the parameter space.

6.2 Methods

I use 1D physical models of a canonical planet with a mass of $0.5 M_J$ and a radius of $0.9 R_J$ around two stars, an M0V star representing later spectral type stars currently known to host temperate Jupiters (Cañas et al., 2020), and a G2V star representing the earlier type stars which host most of the known temperate Jupiters (e.g. Gill et al., 2024). The planetary mass and radius were chosen to be roughly representative of the radii and masses of known temperate Jupiter JWST targets. For each star I include a grid of simulations at different orbital distances intended to span the temperate Jupiter parameter space, with zero albedo, full heat redistribution equilibrium temperatures of 220 K (M0 – 0.449 au, G2 – 1.588 au), 300 K (M0 – 0.242 au, G2 – 0.854 au), 380 K (M0 – 0.151 au, G2 – 0.537 au), and 500 K (M0 – 0.087 au, G2 – 0.308 au). I use a hierarchical modelling approach, first computing the pressure-temperature profile using a radiative-convective equilibrium model, which I use with a chemical kinetics model including photochemistry and condensation to obtain the chemical vertical abundance profiles. I also use an analytic cloud condensation model to predict the vertical and radial distributions of aerosols. I process the temperature profiles and chemical abundances into synthetic transmission spectra.

6.2.1 Radiative-convective equilibrium

I use HELIOS (Malik et al., 2017, 2019) to compute the 1D pressure-temperature structure of the model atmospheres, solving for two-stream radiative transfer with the correlated-k approximation (see Sec. 1.2.2) and including line opacities from H_2O , CO_2 , CO , CH_4 , NH_3 , HCN , PH_3 , C_2H_2 , H_2S , Na, and K, as well as Rayleigh scat-

tering from H_2 , He, and H_2O , and collisionally-induced absorption scattering from $\text{H}_2\text{--H}_2$ and $\text{H}_2\text{--He}$ (see Grimm & Heng, 2015b). The code checks for convective stability and performs convective adjustment to ensure the planet is in radiative-convective equilibrium. I use the planetary parameters of the canonical planet ($M_{\text{P}} = 0.5 \text{ M}_{\text{J}}$, $R_{\text{P}} = 0.9 \text{ R}_{\text{J}}$, $g = 15.3 \text{ ms}^{-2}$), and stellar parameters for TOI-1899 ($T_{\text{eff}} = 3926 \text{ K}$, $R_{\star} = 0.607 \text{ R}_{\odot}$, $\log g = 4.672 \text{ (cgs)}$, $Z = 0.28$, from Lin et al., 2023) as a representative M0 host and the Sun ($T_{\text{eff}} = 5772 \text{ K}$, $R_{\star} = 1.0 \text{ R}_{\odot}$) as a representative G2 host. For the stellar spectrum of the M0 host I interpolate between PHOENIX stellar models (Husser et al., 2013) using the stellar parameters for TOI-1899b normalised for the effective temperature, while for the G2 host I use the solar spectrum of Gueymard (2004). Elemental abundances, cloud structure, and internal temperature are all free parameters in my HELIOS model. For the radiative-convective equilibrium calculation I assume chemical equilibrium throughout the atmosphere computed with a FastChem grid (Stock et al., 2018). This reduces the chemical free parameters to simply a question of elemental abundances, which I fix by setting the N, C, S, P, Na, K, and O abundances as 3 times solar (using the solar abundances of Asplund et al., 2009). I do not consider the radiative impact of aerosols in this model, which could have large impacts on the temperature profile in cases where optically thick aerosols form (e.g. Mak et al., 2023). I fix the internal temperature to 80 K as a representative value for a several-gigayear-old planet of this mass (compare the use of 75 K for a Saturn-mass planet by Fortney et al., 2020).

6.2.2 Chemical kinetics, photochemistry and condensation

I use the 1D chemical kinetics and photochemistry code VULCAN (Tsai et al., 2017; Tsai et al., 2021) to determine the chemical structure for each simulation in the grid, by solving the Eulerian continuity equations (see Sec.1.2.3). I include diffusion, advection transport parametrized using the eddy-diffusion coefficient K_{zz} , as well as condensation of S_8 and H_2O where relevant (determined using comparison of the pressure-temperature profiles to condensation curves). I compute a main grid considering sulfur chemistry using the C-H-N-O-S photochemical network (the 2024 version based on Tsai et al., 2023) for each of the pressure-temperature profiles computed, fixing the elemental abundances to $3\times\text{Solar}$ and the value of K_{zz} to $10^8 \text{ cm}^2\text{s}^{-1}$. I also calculate an extended grid with different values of metallicity ($1\times, 3\times, 10\times, 30\times$), K_{zz} ($10^8 \text{ cm}^2\text{s}^{-1}$ representing strong-moderate vertical transport, $10^6 \text{ cm}^2\text{s}^{-1}$ representing weaker vertical transport), and C/O ratio (0.5, 1), using the simpler C-H-N-O photochemical network. I do not run additional radiative-convective equilibrium models for the extended grid, instead selecting the relevant pressure-temperature profile computed from the $3\times\text{Solar}$ metallicity models, as the desire is to investigate the elemental ratio dependence, rather than the temperature differences that would be caused by changing the metallicity and C/O ratio. For the UV stellar spectra (50–300 nm), I use the solar spectrum of Gueymard (2004) to represent the G2 host, and the v2.2 MUSCLES spectrum of GJ 832 (France et al., 2016;

Youngblood et al., 2016; Loyd et al., 2016) to approximate a relatively old (8.1 ± 4.2) and not overly active M0 host. I set the zenith angle to 58° to compute the dayside average composition (Tsai et al., 2021), with the diurnal flux set to 0.5, treating the planets as non tidally-locked. I set the initial distribution to equilibrium chemistry, and use the temperature profiles from the **HELIOS** models. To include condensation (only done for the main grid), I include an initial run without condensation to set the initial abundances for a subsequent run with condensation, in which the abundances of condensable species are held constant once condensation-evaporation equilibrium has been reached.

I also use the **Virga** cloud condensation code (Batalha et al., 2026) to investigate the aerosol properties of condensed species, including saturation vapour pressure curves (Gao et al., 2017) and refractive indices (Fuller et al., 1997) for the sulfur allotrope S_8 added by Matt Lodge. I use the obtained temperature-pressure profiles from the **HELIOS** model and estimate the gas mixing ratio of S_8 based on the abundances in the condensable parts of the atmosphere from **VULCAN**. I vary the sedimentation efficiency, f_{sed} , as a free parameter (0.1, 0.5, 1, and 5).

6.3 Results

6.3.1 Energy transport and equilibrium processes

The pressure-temperature profiles obtained from the radiative-convective equilibrium models are depicted in Fig. 6.1. Examination of these profiles, and comparison to the pressure-temperature space of chemical transitions, reveal several interesting features of temperate Jupiters. The general shape of the temperature profile is a semi-isothermal radiative stratosphere (pressures below 10–100 mbar), above a region of steep temperature gradient, followed by a “knee” structure at ~ 1 bar where the gradient flattens, and finally steepens into a universal adiabat at deep pressures (~ 100 bar). The steep gradient above the “knee” is such that the atmospheres becomes unstable with respect to convection at these altitudes, creating a detached convective zone separated from the deep adiabat by a radiative layer (the “knee”). The 500 K simulations do not result in the formation of a detached convective zone, which is also the case for hot Jupiters. A mild temperature inversion is also seen in all the simulations except the 500 K. I do not account for moist convection, where the condensability of H_2O can inhibit convection in H_2 dominated atmospheres, affecting the temperature of convective regions (Leconte et al., 2017).

The difference between M and G dwarf hosts can be seen as a general flattening of the overall temperature differences with an M host, resulting in a warmer stratosphere but cooler deep radiative region, but minimal changes in the intermediate or deep convective zones. This can be understood as the incoming stellar radiation with a M dwarf host is more weighted towards the infrared, where the atmospheric opacity is higher due to CH_4 and H_2O , resulting in more absorption in the stratosphere, heating it and correspondingly cooling the deeper radiative region.




Figure 6.1: Predicted pressure-temperature profiles of temperate Jupiter atmospheres with different instellations (labelled with their zero albedo equilibrium temperatures) and stellar host spectral types (M0 and G2) calculated using *HELIOS*. Also plotted is the equilibrium isocompositional curve of NH_3 (which dominates at temperatures below this) and N_2 (which dominates at temperatures above this), as well as the condensation curves of KCl , H_2O and S_8 , calculated using *Virga*. As the isocompositional curve is nearly parallel to the universal deep adiabat shared by all the models, above which the chemistry is quenched in all considered vertical mixing regimes, the NH_3/N_2 ratio is constant. Condensation approximately occurs where the pressure-temperature profile first (starting from the deepest pressure) crosses the condensation curve.

The deep adiabatic temperature profile is constant between the models, which is consistent with a key finding of Ohno & Fortney (2023a) that at these temperatures the deep temperature structure is insensitive to the upper boundary condition of the atmosphere. It does however depend upon the internal temperature of the planet.

Also plotted in Fig. 6.1 is the equilibrium isocompositional curve of $\text{NH}_3\text{--N}_2$, which traces the pressure-temperature space where the abundances of NH_3 and N_2 are equal in equilibrium. The ratio of the abundance of one species over the other is dictated by the distance from the isocompositional curve at a given temperature and pressure (one could construct a family of similar curves defining different $\text{NH}_3:\text{N}_2$ ratios offset from the isocompositional curve). Many of the pressure-temperature profiles of temperate Jupiters cross the isocompositional curve in multiple places. With equilibrium chemistry, this would mean that the two different species would dominate at different altitudes. However, the isocompositional curve is near-parallel

to the shared deep adiabat. As vertical transport typically quenches the atmospheric composition to that of the deep atmosphere (abundances are approximately vertically constant above the quench level due to vertical mixing), and therefore set by the same universal deep adiabat, I expect the NH_3/N_2 ratio to be insensitive to the equilibrium temperature within the range of temperatures considered (also seen by Ohno & Fortney, 2023a). In the case of extremely sluggish vertical mixing, the abundances may be quenched from the radiative layer, causing strong depletion (at higher instellations) or enhancement (at lower instellations) of NH_3 .

I note that the isocompositional curve does not indicate a sharp transition but instead the pressure-temperature where the NH_3 and N_2 abundances are equal. The ratio of the abundance of one species over the other is dictated by the distance from the isocompositional curve at a given temperature and pressure (one could construct a family of similar curves defining different $\text{NH}_3:\text{N}_2$ ratios offset from the isocompositional curve).

Plotting the condensation curves of S_8 and H_2O reveals that these species can condense in the middle to upper atmospheres of temperate Jupiters, dependent on both instellation and stellar host spectral type. For H_2O , this corresponds with an equilibrium temperature between 220 and 300 K, while for S_8 this is at a higher temperature between 300 and 380 K depending on the host spectral type. I also consider KCl , as an example of a mineral condensate with a low condensation temperature. It can be seen that KCl condenses much higher in the atmosphere in the 500 K simulations, potentially causing a sharp transition in aerosol properties between warmer planets where mineral species can condense above 1 bar, and cooler planets where mineral species can only condense in the deep adiabat. I note that I do not consider radiative feedback from any clouds that may form as a result of this condensation, which could inhibit (or strengthen) cloud formation by altering the temperature profile (e.g. Mak et al., 2023).

6.3.2 Disequilibrium chemistry

The vertical mixing ratios of key species from the chemical kinetics model of the main grid, run without condensation, are depicted in Fig. 6.2. At the temperatures present in temperate Jupiters, H_2O and CH_4 dominate as the principle carriers of atmospheric oxygen and carbon respectively.

Nitrogen chemistry

The nitrogen is mainly split between NH_3 and N_2 , with near constant vertical abundances set by quenching from the deep atmosphere and a NH_3/N_2 ratio insensitive to temperature due to the universal deep adiabat, as expected. However, the first hints of photochemistry can be seen by looking at the 500 K planet with G2 host, by noticing that NH_3 is depleted in the upper atmosphere, enhancing the N_2 abundance as well as HCN , which is clearly identifiable as an important photochemical product in the 380 and 500 K planets due to its increasing abundance with altitude. The




Figure 6.2: Vertical mixing ratios of the key species in temperate Jupiters, computed using the C-H-N-O-S photochemistry network of *VULCAN*, not including condensation. The M0 220 K model is omitted as it did not converge.

increased abundance of HCN and decreased abundance of NH_3 can be attributed to the stronger near-UV flux of the G dwarf host (arising from its higher blackbody temperature), which is primarily responsible for the photochemical destruction of NH_3 (Hu, 2021). This feature becomes much more prominent in the extended grid when the vertical mixing is weaker, demonstrated in Fig. 6.3. Weaker mixing allows NH_3 to be depleted to a much greater depth in the G2 host models: down to 100 mbar at 300 K and 20 mbar at 380 K. The HCN abundance is enhanced by up to two orders of magnitude as a result of weaker vertical mixing.

Carbon chemistry

The presence of CH_4 gives rise to diverse carbon chemistry. CO is mainly present quenched from the deep atmosphere, while CO_2 at these temperatures is primarily a photochemical product. With a relatively weak dependence on temperature and metallicity, the main factor impacting the abundance of photochemical CO_2 is vertical mixing, where the abundance can be enhanced by up to two orders of magnitude for planets with a G2 host in the weaker mixing ($K_{zz} = 10^6 \text{ cm}^2 \text{ s}^{-1}$) model. This makes CO_2 an independent tracer of vertical mixing. This is provided that quenched CO_2 can be ruled out, as quenched CO_2 can become dominant over photochemical CO_2 with intense vertical mixing and high metallicity ($\geq 30 \times \text{solar}$), which is apparent in the extended model grid.

Ethane (C_2H_6) becomes a relevant species in the mid to upper atmosphere for the 500 K and 380 K planets. The extended grid reveals that the C_2H_6 abundance is strongly correlated with metallicity: a change of ± 0.5 in the $[\text{C}/\text{H}]$ results in a change of $\sim \pm 1$ in the log vertical mixing ratio of C_2H_6 . Acetylene (C_2H_2), propyne (C_3H_4), and the nitrogen-bearing cyanoacetylene (HC_3N) are also somewhat abundant species in the upper atmosphere with complex behaviors: C_2H_2 and HC_3N are both photochemical species favoured with stronger vertical mixing, reversing the usual trend, and without strong temperature or metallicity dependence. C_3H_4 is more prominent at higher temperatures and with weaker vertical mixing. All of these hydrocarbon species, but particularly C_3H_4 and HC_3N , are strongly enhanced in the $\text{C}/\text{O}=1$ atmospheres (1.6 and 1.2 orders of magnitude respectively). HC_3N is of particular interest, as a prebiomarker (molecule associated with the origin of life) with very strong spectral features (Claringbold et al., 2023).

Sulfur chemistry

H_2S is the dominant carrier of sulfur in the deep atmosphere, however it is very vulnerable to photo-dissociation, being strongly depleted above a certain altitude set by both the temperature and the stellar host. The photochemical H_2S depletion pressure peaks for the 380 K planets, and is deeper for planets with the G2 host. Most of the sulfur in the upper atmosphere is in the form of sulfur allotropes, primarily S_8 , with bands of S and S_2 . Another important carrier of sulfur in some of these planets is NS (a result also found by Crossfield et al., 2025). I do not find CS and




Figure 6.3: Vertical mixing ratios of various species in planets representative of four observationally distinct regions within the temperate Jupiter parameter space. Mixing ratios from two values of vertical mixing are displayed, weak-moderate mixing ($K_{zz} = 10^6 \text{ cm}^2 \text{ s}^{-1}$) in solid lines, and strong mixing ($K_{zz} = 10^8 \text{ cm}^2 \text{ s}^{-1}$) in dashed lines. The pressures at which solid S_8 condenses in the kinetics model are highlighted in translucent gold, the regions of optically thick S_8 haze from the cloud model are shown in more opaque, dark gold. Regions that contribute strongly to the transmission spectrum but don't form S_8 are highlighted in blue. I highlight the following key transitions in chemistry between these planets in the case of weaker vertical mixing case (solid lines): I first draw attention to the behaviour of NH_3 (light blue line), which is photochemically depleted (turns off to the left) at a much lower altitude with a G2 host than an M0 host; secondly, I note that HCN (red line) is notably photochemically enriched (shifted to the right at mid-altitudes) in the G2 host at 380 K; finally, I highlight that the impact both of these photochemical effects (NH_3 depletion and HCN formation), as well as the photochemical formation of CO_2 (navy line), are dramatically reduced by stronger vertical mixing (dotted lines). The formation of optically thick S_8 hazes in the stratosphere occurs in the 300 K planets. These four regions correspond to the upper four sextants of Fig. 6.6.

CS₂ carry notable amounts of sulfur contrary to Mukherjee et al. (2025), which can be attributed to the lack of certain CS-forming reactions in the version of the VULCAN chemical network I have used (Veillet et al., 2025). One interesting and unexpected feature is the formation of a considerable, highly localized layer of OCS at ~ 1 bar. This can be clearly identified as forming from CO, with a corresponding drop in the CO abundance, and is associated with the pressure of sulfur dissociation. Comparing otherwise identical runs performed with the C-H-N-O-S and C-H-N-O chemical networks reveals that the inclusion of sulfur in the network has minimal impact on the abundances of non-sulfur bearing species. These small differences include the aforementioned drop in CO abundance associated with the formation of an OCS layer, and a slight decrease in the N₂ and NH₃ abundances associated with the production of NS.

6.3.3 Condensation and aerosols

Organic hazes

Hydrocarbon-rich environments in the Solar System give rise to complex organic hazes, such as the tholins of Titan (e.g. Khare et al., 1984a), and have been implicated in forming super-Rayleigh scattering slopes and muting the spectral features of cooler hot Jupiters (Ohno & Kawashima, 2020; Gao et al., 2021). To investigate the presence of photochemical hydrocarbon hazes in temperate Jupiters, I follow the approach of Mukherjee et al. (2025) and measure the column density of certain haze precursor species above 1 mbar. In my first approach, I use the following species: C₂H₂, C₂H₄, C₂H₆, HCN, C₃H₄, HC₃N, C₄H₂, and C₆H₆, all of which are implicated in the formation of the polyaromatic hydrocarbons and nitriles which polymerize to form hazes (similar to the precursors used by Morley et al., 2015; Tsai et al., 2021; Mukherjee et al., 2025). In a second approach, I omit C₂H₆ and HCN, which dominate the precursor abundances in the first approach, but are stable molecules that are less likely to readily polymerize. The column densities of haze precursors above 1 mbar are depicted in Fig. 6.4.

As seen in Fig. 6.4, the abundance of haze precursors monotonically increases with metallicity in all cases, being typically directly proportional in the case of weak vertical mixing, but 2–6 $\times$ less sensitive to metallicity when the mixing is stronger. The first approach to measuring haze precursors yields universally higher precursor abundances, mainly as a result of HCN, which drives the peak abundances to occur in the 380 K planet with a G2 host and weak mixing. The second approach demonstrates an overall weaker sensitivity to temperature and stellar host, with an at most two orders of magnitude difference in column density between models with this approach (with both extremes interestingly in the 300 K case with weak mixing). I note that the proposed decrease in haziness with temperature below 400 K (Brande et al., 2024) is not universal in my models, and depends on the stellar host, vertical mixing, and precursor species considered.




Figure 6.4: Column density of haze precursor species above 1 mbar for different equilibrium temperatures (300 and 380 K), vertical mixing, stellar host (M0 and G2), and metallicity (1 $\times$, 3 $\times$, and 10 $\times$ solar). The left two grids consider the abundances of C₂H₂, C₂H₄, C₂H₆, HCN, C₃H₄, HC₃N, C₄H₂, and C₆H₆ as haze precursors, while the right two grids exclude the less reactive molecules HCN and C₂H₆ that otherwise dominate the precursor abundances. In most cases the precursor abundances do increase with temperature (Brande et al., 2024) and metallicity, but the relationship is overall complicated and also depends on the vertical mixing strength, stellar host spectral type, and the precursor species considered.

Sulfur Condensation

Post-processing the initial condensation-free **VULCAN** models with S₈ condensation reveals that condensation of sulfur does in fact take place in all the 220 K and 300 K models, as well as to a lesser extent in the 380 K model with G2 host. The vertical extent where condensed S₈ in these models is highlighted in Fig.6.3. This highlights a clear transition in aerosol properties within the temperate Jupiter regime depending on whether the stratosphere is cool enough for S₈ to form. Using **Virga** to investigate the optical properties of these S₈ hazes, I find that the S₈ condenses into a haze of highly scattering (single scattering albedo of 0.93 at optical wavelengths 0.3–0.7 μ m, and $\simeq 1$ at near–mid infrared wavelengths 0.7–15 μ m), small radius particles (with the radius distributions peaking between 10–100 nm). The most opaque region of the haze is highly localized in pressure space regardless of the f_{sed} used, achieving an infrared optical thickness of unity between 20 and 100 mbar for the 300 K models, and 30 and 500 mbar for the 220 K models, at infrared wavelengths (0.7–15 μ m). The haze is not optically thick in the 380 K model with G2 host. Using a lower value for the vertical mixing ($K_{zz} = 10^6 \text{ cm}^2 \text{ s}^{-1}$) results in smaller particle sizes and reduced asymmetry factor, but otherwise has minimal impact on the haze distribution, while using very strong vertical mixing ($K_{zz} = 10^{10} \text{ cm}^2 \text{ s}^{-1}$) increases the peak particle radius to around 1 μ m, but greatly reduces the overall opacity of the sulfur haze, only becoming optically thick in the 220 K model. Although I do not explore radiative




Figure 6.5: Synthetic infrared transmission spectra of two of my model planets, demonstrating the contribution of different species to the overall spectrum (plotted as transmission spectra containing a single line-absorbing species, but the same H_2 and He background, resulting in a shared continuum of collisionally induced absorption). The spectra share the same CH_4 and H_2O major contributions, but have notably different contributions from NH_3 , HCN , CO_2 , and HC_3N , due to differences in photochemistry.

feedback, the high level of infrared scattering from S_8 would likely significantly cool the lower atmosphere, while absorption of short wavelengths ($< 0.4 \mu\text{m}$) could cause a thermal inversion. Sulfur hazes would have a strong impact on the geometric albedo spectrum of a temperate Jupiter, as discussed by Gao et al. (2017).

Water Condensation

Condensation of H_2O only occurs in the 220 K models. Condensation results in the depletion of H_2O , and therefore oxygen, from the parts of the atmosphere visible to transmission spectroscopy. I do not focus on the aerosol properties of H_2O here, which are highly sensitive to aerosol microphysics (including likely requiring heterogeneous nucleation), as discussed in detail in Mang et al. (2022). The strong sensitivity of H_2O clouds to microphysics is apparent in my *Virga* models based on a strong dependence of the cloud properties and physical extent to both f_{sed} and K_{zz} . Nonetheless H_2O condensation represents one of the strongest temperature dependent transitions, visible through both the depletion of H_2O , and potentially through the pressure at which the cloud continuum is found.

6.3.4 Transmission spectra

I use `petitRADTRANS` (Mollière et al., 2019) to generate transmission spectra based on the calculated vertical mixing ratios and temperature profiles, in order to determine the potential detectability of different species. I include opacity from CH_4 (Yurchenko et al., 2017), H_2O (Polyansky et al., 2018), NH_3 (Coles et al., 2019), HCN (Barber et al., 2014), CO_2 (Yurchenko et al., 2020), CO (Rothman et al., 2010), HC_3N (Rimmer et al., 2021), H_2S (Azzam et al., 2016), NS (Yurchenko et al., 2018), OCS (Owens et al., 2024), SO_2 (Underwood et al., 2016), and HC_3N (Rimmer et al., 2021) at a spectral resolution of $R = 250$, between 0.3 and $10\ \mu\text{m}$. I also include opacity from H_2 and He Rayleigh scattering and $\text{H}_2\text{--H}_2$ and $\text{H}_2\text{--He}$ collisionally induced absorption.

I find that CH_4 is the dominant component of all of the spectra, as expected (see Fig. 6.5). NH_3 is also a major component, with features at 1.5 , 2 , 3 , and $6\ \mu\text{m}$, although it is noticeably depleted in the G2 host, weaker vertical mixing models. H_2O is the third major component, except in the $220\ \text{K}$ models, where it is completely absent from the transmission spectrum due to condensation. CO_2 has a prominent feature at $4.2\ \mu\text{m}$ in the weaker vertical mixing models, especially strong with a G2 host. HCN has a strong feature at $3\ \mu\text{m}$, that only features in the spectra of warmer ($380\ \text{K}$ and $500\ \text{K}$), weaker vertical mixing ratio models with a G2 host (such as the top panel of Fig. 6.5). HC_3N has very strong, narrow features, allowing these features to appear above the total spectrum at 3 and $4.4\ \mu\text{m}$ even when the abundance is very low, which occurs in the stronger vertical mixing $300\ \text{K}$ model with M0 host (bottom panel of Fig. 6.5). HC_3N is also visible in all of the $\text{C}/\text{O}=1$ models. I do not see contributions from any sulfur species (H_2S , SO_2 , OCS , or NS) in the spectrum of any of the models.

6.4 Mapping the Parameter Space

Based on the shared features of my models, I propose temperate Jupiters as an observationally distinct class of exoplanet. Features shared by all the temperate Jupiters considered here include the dominance of CH_4 opacity, the presence of nitrogen-bearing species visible with transmission spectroscopy, the absence of mineral condensates above the deep adiabatic part of the atmosphere, and the presence of a detached convective zone in the middle atmosphere. I use the former two features to suggest the following definition of a temperate Jupiter – a gas giant with instellation such to sculpt an atmosphere with detectable nitrogen-bearing species, but an absence of mineral condensates. This definition of temperate Jupiters ties it together with one of the unique insights offered by the regime, the ability to measure the $\text{C}/\text{N}/\text{O}$ ratio.

The condition of lacking mineral condensates is primarily a function of the temperature profile. Due to the characteristic shape of the temperate Jupiter temperature profile, the condensation pressures of mineral species like KCl experience

a sharp transition based on equilibrium temperature (see Fig. 6.1). The presence of an intermediate radiative layer between the detached convective zone (found only in more temperate planets) and the deep adiabat may serve to further trap these mineral species in the lower atmosphere. The detached convective zone may also strongly alter the vertical mixing, the effect of which could be detectable through tracers like CO_2 and HCN . The condition of detectable nitrogen sets the upper temperature bound by the equilibrium transition from N_2 to NH_3 , and the lower bound by the condensation of NH_3 . The transition from N_2 to NH_3 is not a simple function of temperature however (also shown by Fortney et al., 2020; Ohno & Fortney, 2023a,b; Mukherjee et al., 2025), and more dependent on internal temperature and photochemistry. While detectable nitrogen is a compelling boundary criterion, the strong impact of internal temperature must be considered, as the NH_3 abundance probes the deep adiabat temperature more closely than the irradiation temperature.

Based on this definition, and the resulting observational tracers, I propose an approximate parameter space for temperate Jupiters as giant planets with equilibrium temperatures between ~ 200 and ~ 500 K (assuming zero albedo and full heat redistribution). The precise boundary temperatures are dependent on metallicity, vertical mixing, planet mass, and other properties. Upcoming JWST observations of two planets near this upper boundary, NGTS-11b and TOI-216c (observed as part of GO 7982), will probe and refine this transition.

Through this modelling approach, I have also identified a number of strong transitions in atmospheric properties within the temperate Jupiter parameter space that depend on both the quantity and wavelength of instellation received (equilibrium temperature and stellar host, respectively), dividing the parameter space into speculative sextants presented in Fig. 6.6. This first of these transitions divides the parameter space based on the observed nitrogen chemistry depending on the stellar host spectral type (vertical red dash-dotted line in Fig. 6.6), as NH_3 can be strongly depleted by Sun-like and earlier stellar hosts with their relatively strong NUV flux, becoming HCN or N_2 depending on temperature. The other two main transitions divide the parameter space using the condensation of sulfur and water based on equilibrium temperature (horizontal dotted lines in Fig. 6.6). S_8 can condense to form an optically thick, highly scattering haze layer if the stratosphere is sufficiently cool, while at yet colder temperatures H_2O condensation creates a strong transition between warmer temperate Jupiters with atmospheric H_2O vapour, and cooler temperate Jupiters with dry stratospheres. These transitions, in particular in nitrogen photochemistry, are all sensitive to vertical mixing, allowing them to probe atmospheric dynamics. I have also identified photochemically produced CO_2 as an important tracer of vertical mixing which is relatively independent of temperature. I have found that C_2H_6 is a sensitive indicator of metallicity, and explored trends in the abundances of potentially-haze forming hydrocarbons including C_3H_4 and HC_3N on various atmospheric properties, identifying HC_3N as a readily visible species in planets with high atmospheric C/O ratios (>1). These six sextants hopefully serve




Figure 6.6: Schematic of key transitions responsible for shaping the temperate Jupiter parameter space, over-plotted with JWST gas giants targets from Cycles 1–5 in blue. The upper bound of the temperate Jupiter parameter space is set between 380 and 500 K based on a) the atmosphere is too cool for the condensation of mineral species above the deep convective region, and b) the temperature gradient is sufficient for the formation of a detached convective zone. The lower bound is set by the condensation of NH_3 , below which the atmospheric nitrogen content cannot be measured using transmission spectroscopy. Internal transitions that subdivide the parameter space are the condensation of H_2O and S_8 , and a key transition in photochemistry between early and late spectral type stars driven by the photodissociation of NH_3 , depicted in Fig. 6.3, which correspond to the four higher temperature quadrants. These transitions are sensitive to vertical mixing, metallicity, and radiative feedback from aerosols; these sextants represent a guide to be refined by observations.

as a useful guide to the planning and interpretation of observations of this fascinating class of exoplanet atmosphere.

6.4.1 Further work

The remaining tasks for this work include finalizing the 220 K photochemistry models, investigating CS_2 with an updated chemical network, and generating transmission spectra with additional opacity sources (C_2H_6 , C_3H_4 , and S_8 haze).

Chapter 7

Conclusions and Future Work

7.1 Summary

In this thesis I have analysed and interpreted a series of hot Jupiter transmission spectra, with complete observational analyses of HAT-P-44b at optical wavelengths (Chapter 3) and HAT-P-30b at infrared wavelengths (Chapter 4), as well as modelling analyses to interpret the spectra of both planets, and infrared observations of WASP-15b, TrES-4b, and NGTS-2b (Chapter 5). I also calculate a grid of simulations of cooler gas giants, to map key transitions and trends within this temperate Jupiter parameter space, to make predictions for and aid in the interpretation of upcoming JWST observations (Chapter 6).

The ground-based transmission spectrum of HAT-P-44b was taken as part of the LRG-BEASTS survey using the WHT, and is a valuable additional data point in understanding the parameter space of aerosol properties of hot Jupiters. The precise light-curves and robust transmission spectrum demonstrate the effectiveness of ground-based observations in obtaining optical transmission spectra, a key wavelength region for complete atmospheric characterization, but one not covered by JWST, the current flagship observatory. The strong detection of sodium and inference of a relatively cloud free atmosphere for HAT-P-44b is similar to the observed properties of WASP-39b (Fischer et al., 2016; Nikolov et al., 2016; Sing et al., 2016; Wakeford et al., 2017a; Kirk et al., 2019), which is a close sibling in planetary properties. This points to the existence of a potential population of inflated hot Saturns with equilibrium temperatures around 1100 K that demonstrate clear atmospheres. In contrast, another close sibling HATS-46b demonstrates a featureless spectrum indicative of cloud (Ahrer et al., 2023a), suggesting that simple planetary properties (temperature, gravity, and stellar host) are not completely predictive of aerosol parameters. The precision of my HAT-P-44b spectrum is enough to find evidence for H₂O absorption, measure the temperature of the H₂ gas Rayleigh from the scattering slope, and place constraints on the metallicity as super-solar, all of which is rarely achieved for ground-based observations, although these interpretation are sensitive to the aerosol parametrization used.

All four infrared datasets analysed in this thesis are JWST observations obtained as part of the BOWIE-ALIGN programme, a survey investigating the origin of hot Jupiters by contrasting the atmospheric composition of its misaligned and aligned targets. For TrES-4b, I found a robust atmospheric composition with minimal cloud, sub-stellar metallicity, and a sub-solar C/O ratio. For WASP-15b, I was also able to obtain a relatively robust result of super-stellar metallicity and approximately solar C/O ratio with some uncertainty on the presence of minor opacity contributions from SO_2 , OCS, CH_4 , and clouds. For the atmosphere of HAT-P-30b the interpretation is less conclusive, due to the presence of a high metallicity mode and a difference in temperature between free and equilibrium retrievals driven by CH_4 opacity. Nonetheless, I was able to determine robust results for the C/O ratio as strongly sub-solar, while also favouring a sub-stellar metallicity. NGTS-2b was the most resistant to definitive characterization due to its highly muted spectral features, which introduce degeneracies based on metallicity (via MMW), temperature, and cloud-top pressure, preventing me from obtaining a single preferred chemical solution, although equilibrium retrievals do favour a solar to super-solar metallicity and a sub-solar to solar C/O ratio.

For the BOWIE-ALIGN targets, formation arguments are made for each individual planet based on the atmospheric inferences, but always with the caveat that drawing conclusions from any individual planet can only be tentative due to the degeneracies and uncertainties in formation models. It is therefore the aim of the BOWIE-ALIGN programme to combine the four planets analysed in this thesis with an additional four targets discussed in Sec. 4.1 for a more definitive comparison. The intermediate goal is to infer the impact of formation mechanism on composition, by identifying significant differences between the two groups. I do not attempt this inter-comparison in this thesis, as it will be the subject of future publications led by the programme principal investigators, however just from the four planets analysed here it can be seen that the sample covers a diverse range of atmospheric compositions.

In my study of temperate Jupiters, I have used numerical simulations of radiative-convective equilibrium, chemical kinetics, and cloud condensation to draw up a preliminary map of the parameter space, divided into six observationally distinct regions based on key transitions in chemistry and aerosol properties. This includes the condensation of S_8 and H_2O , and the photodissociation of NH_3 . I also explore the dependence of various potentially detectable species, including HCN, CO_2 , HC_3N , and an ensemble of haze precursor molecules, on temperature, metallicity, vertical mixing, stellar host, and C/O ratio. This preliminary map will aid in the interpretation of JWST observations and guide future observing programmes for both planet detection and characterization.

Overall my work for this thesis has contributed complete transmission spectroscopy analyses for two hot Jupiters with limited previous atmospheric characterization, and aided in the interpretation of an additional three planets, contributing

to the advancement of both the LRG-BEASTS and BOWIE-ALIGN programmes. I have thereby contributed to our understanding of hot Jupiter aerosols and origins respectively. I have also laid theoretical groundwork for the observation of temperate Jupiters, a key upcoming area of exoplanet research that will help bridge the gap between hot Jupiters and the Solar System.

7.2 Ongoing and Future Work

7.2.1 The future of BOWIE-ALIGN

Analysis of the eight NIRSpec/G395H observations that constituted the original BOWIE-ALIGN survey is now complete. This enables the sub-population comparison of C/O ratios and metallicities originally envisioned by the survey. A homogeneous set of data reductions, using **Tiberius** and **Eureka!**, have been used for this purpose, with comparison proceeding along two lines. In one upcoming analysis, the comparison of the misaligned and aligned samples will take place using a colour-band analysis (Fu et al., 2025a) and stacking of transmission spectra (Kirk & Owen, 2025). In the other main summary analysis, a homogeneous retrieval analysis will be applied to all seven planets observed in transmission. Having provided many retrievals for the programme, I have advised and contributed particularly to the latter analysis, especially with regards to hybrid retrievals. I will also continue to play a role in potential future analyses of the programme, which are outlined below. The inclusion of already observed NIRISS and LRG-BEASTS data in the interpretation should also play a powerful role in reducing degeneracies in, and therefore greatly enhancing the precision of, our compositional constraints for multiple BOWIE-ALIGN planets.

Beyond its primary goal of investigating planet formation and migration, the BOWIE-ALIGN sample represents a set of homogeneously analysed hot Jupiters with very similar planetary parameters (particularly in temperature and stellar host). This represents an opportunity to perform comparative planetology with other atmospheric properties. One particularly strong contrast already seen in the BOWIE-ALIGN programme is a division between planets with strongly muted atmospheres (KELT-7b, NGTS-2b, and to a lesser extent HAT-P-30b) and those with relatively unmuted features (WASP-15b, WASP-94 Ab, WASP-17b, and TrES-4b). While this does not fall on alignment lines, understanding the source of this muting, be that aerosols or unresolved limb asymmetries, can help us understand atmospheric processes and refine our modelling approach to better interpret elemental ratios.

One particular aspect that we will explore as part of the BOWIE-ALIGN programme is a limb asymmetry analysis, using analysis techniques like fitting the entire light-curves with two hemispheres of different radii (Catwoman, Jones & Espinoza, 2020), fitting the light-curves just at ingress and egress, and investigating the wavelength dependence of the mid-transit time (all described in Espinoza et al., 2024). An unresolved limb asymmetry can be responsible for the muting of spectral features

(MacDonald et al., 2020), and evidence for a cloudy morning limb has already been observed in the BOWIE-ALIGN target WASP-94 Ab (Fu et al., 2025b). Another avenue for investigation of the BOWIE-ALIGN sample is a deeper investigation into choices of parametrization of temperature and aerosols. While initial investigations into patchy clouds and non-isothermal temperature profiles (Meech et al., 2025; Fairman et al., 2026) did not result in improvement of the evidence, Lueber et al. (2024) and Fairman et al. (2024) demonstrate the importance of broad wavelength coverage in accessing aerosol and temperature information content. Upcoming NIRISS/SOSS observations will provide this for NGTS-2b, KELT-7b, and HAT-P-30b, which also happen to be the three least well constrained of the BOWIE-ALIGN targets, and therefore grant the potential for more definitive chemical constraints from all three planets.

7.2.2 The temperate Jupiter TOI-1899b

I am leading the analysis of the transmission spectrum of TOI-1899b (Cañas et al., 2020; Lin et al., 2023), a temperate Jupiter ($T_{\text{eq}} \sim 360$ K) observed with the NIR-Spec PRISM instrument as part of JWST programme GO 4227 (PIs: Claringbold & Wheatley). Observed on 2025 May 10, the transmission spectrum covers a wavelength range of 0.60–5.30 μm , providing sensitivity to a wide range of opacity sources, including molecular opacity from H_2O , CH_4 , NH_3 , HCN , CO_2 , and HC_3N , as well as aerosol opacities. This should in turn permit measurement of the C/N/O ratio, which will provide powerful constraints on planet formation, as well as probing atmospheric processes like photochemistry and haze formation relevant to the understanding of smaller temperate planets. TOI-1899b is one of five gas giants with equilibrium temperatures below 500 K that have been observed with JWST as of 2026 May, and it is unique in orbiting an early M-type stellar host. This cool stellar host makes it an excellent test of the impact of stellar host temperature, particularly the impact of NUV flux discussed in Chapter 6, although it does also present issues due to the potential impact of stellar inhomogeneity, which appear to be ubiquitous for exoplanets orbiting M dwarfs (Cañas et al., 2025; Kanodia et al., 2026). The analysis of TOI-1899b is ongoing, with one preliminary data reduction with *Eureka!* already completed by Eva-Maria Ahrer, and a *Tiberius* reduction in progress by me. I plan to use atmospheric retrievals and forward models to interpret the resulting spectra.

7.2.3 Warm Jupiters with inner companions

In Sec. 1.1.4, I alluded to the existence of warm Jupiter sub-populations with shared dynamical properties suggestive of shared formation and migration histories. One such sub-population is warm Jupiters with smaller inner companions. Simulations of the evolution of such planetary systems suggest that they must have undergone minimal migration, to justify why the smaller planet did not migrate more rapidly than its outer companion (Huang et al., 2016; Su et al., 2024). The chemical compo-




Figure 7.1: Diagram showing the three gas giant formation scenarios I will test by combining the sample of warm Jupiters with companions from my programme (red) with the aligned (orange) and misaligned (blue) hot Jupiters from the BOWIE-ALIGN programme. Figure prepared by Anna Penzlin for the accepted proposal for the JWST programme GO 7982 (PIs: Claringbold & Owen), adapted from Penzlin et al. (2024), used with permission.

sition of these warm Jupiters should therefore be reflective solely of the inner disc. This makes this sub-population a point of contrast to the BOWIE-ALIGN samples: both the misaligned hot Jupiters (thought to form via high-eccentricity migration and therefore representative of the composition of the outer disc); and the aligned hot Jupiters (thought to form via disc migration and therefore reflective of the whole disc). This is depicted in the Fig.7.1.

I am leading the JWST programme GO 7982 (PIs: Claringbold & Owen) to extend the BOWIE-ALIGN analysis to this new sub-population, combining observations of three new warm Jupiters with inner companions (NGTS-11b, TOI-216c, and TOI-1670c) with two sets of archival observations (TOI-2525c and TOI-1130c). This will allow the measurement of the C/O ratios and metallicities of all five planets, which can be compared to the BOWIE-ALIGN sample. GO 7982 consists of six new observations to obtain as broad as possible spectral coverage of each planet, which should help combat the discussed degeneracies that have complicated the BOWIE-ALIGN analyses. I will use the G395M grating of NIRSpec for coverage from 2.9–5.2 μm for all three planets, combined with the SOSS mode of NIRISS for TOI-216c and NGTS-11b, and the G235M grating of TOI-1670c due to contamination from nearby stars preventing the use of NIRISS, which is more sensitive to contamination owing to its slitless design. The use of medium resolution NIRSpec gratings also eliminates the gap in wavelength coverage and potential offset between detectors that have hampered observation with the high resolution gratings. These observations will take place between 2026 July and 2027 March.

NGTS-11b and TOI-216c are of particular interest to me, as they both have equilibrium temperatures less than 500 K, fitting my definition of temperate Jupiters (Chapter 6). With similar temperatures and stellar hosts, comparison of these planets will probe the temperate-hot Jupiter boundary as well as the onset of NH_3 dissociation due to photochemistry.

7.3 Longer-term Outlook

7.3.1 Long period planet discovery

The characterization of temperate and cold Jupiters is primarily limited by the lack of known transiting planets around sufficiently bright stars. This is a direct consequence of the strong observational biases against the detection of long period transiting planets. This has been partly addressed in recent years through the follow-up of single transit events from TESS (e.g., Cañas et al., 2020; Ulmer-Moll et al., 2025), as well as planets in the TESS continuous viewing zone (e.g. Mireles et al., 2023). The continued operation of TESS will increase the number of both mono- and duotransit events (e.g. Hawthorn et al., 2024) which can result in the discovery of long period planet targets suited for atmospheric characterization, and thereby populate the temperate and cold Jupiter parameter space. The ESA PLATO (PLANetary Transits and Oscillations of stars, Rauer et al., 2014, planned to launch in 2027) mission will greatly improve the sensitivity to transiting long period planets within its field of view (Eschen et al., 2024).

In addition to transiting planets, astrometry from future data releases of the GAIA mission (Collaboration et al., 2016) has the potential to detect a significant number of exoplanets, which will primarily be long-period gas giants (Perryman et al., 2014). This could result in as many as $120,000 \pm 22,000$ total planets from the full DR5 data release (Lammers & Winn, 2026), but this number is sensitive to the noise floor of GAIA observations. Although the vast majority of these systems will be non-transiting, a significant number (25–42 predicted by Perryman et al., 2014) may be amenable to transmission spectroscopy. Advances in direct imaging techniques offered by the Nancy Grace Roman Space Telescope and the ELT will enable detection and characterization of nearby non-transiting giant planets around the Solar System regime (e.g., Carrión-González et al., 2021; Houllé et al., 2021). The next decade of exoplanetary science therefore has the potential to bridge the gap between the currently studied giant exoplanets and the giants of the Solar System using a variety of techniques, a truly exciting prospect for the study of giant planet atmospheres.

7.3.2 Ground-based atmospheric characterization

While JWST is currently in the spotlight for the revolutionary understanding of exoplanet atmospheres it offers, ground-based telescopes will continue to offer unique insights not available to JWST. Right now this is mainly the potential to use a much higher spectral resolution, in turn granting access to dynamical information from Doppler shifts and line profiles, isotopologues, certain species with weak broad-band opacities but forests of strong individual lines (CO being a notable example), and the ability to detect atmospheric lines from non-transiting planets (Snellen, 2025). In the future this will also include improvement in angular resolution granted by 30m-class telescopes, enabling direct imaging at closer orbital separations.

Given the finite lifetime of HST, currently the only source of space-based optical spectroscopy, ground-based telescopes may be necessary to provide spectroscopic observations blueward of 600 nm in the future. LRG-BEASTS and other programmes have demonstrated the ability of ground-based telescopes to provide precise optical transmission spectra. Although both the EFOSC2 and ACAM instruments used by LRG-BEASTS are now decommissioned, there are still spectrographs able to perform similar analyses (e.g. OSIRIS on the GTC). Furthermore, advances in analysis techniques have eliminated the absolute need for comparison stars (e.g., Panwar et al., 2022), expanding the pool of potential targets and instruments for ground-based transmission spectroscopy. Optical wavelengths are also of interest for emission spectroscopy, due to the ability to detect reflected starlight at these wavelengths and therefore constrain the scattering properties of atmospheres (e.g. Evans et al., 2013). This in turn allows measurement of the net energy budget of exoplanets. The combination of low resolution and high resolution measurements is also highly synergistic and an active avenue for research (Brogi et al., 2017; Pino et al., 2018), in particular by using the low resolution continuum to break degeneracies in the high resolution model fitting and thereby obtain precise abundances of atomic and molecular species (Khalafinejad et al., 2021).

7.3.3 Space-based atmospheric characterization

In its near four years of operation, JWST has already revolutionized many aspects of exoplanet atmosphere characterization. This trajectory is likely to continue, as the focus of the telescope shifts from reconnaissance to detailed exploration. A particular dedicated effort is being put towards finally pushing into the terrestrial planet regime and understanding the cosmic shoreline (Zahnle & Catling, 2017) and the presence (or lack thereof) of atmospheres around smaller temperate planets. This includes 500 hours of Director’s Discretionary Time dedicated to the Rocky Worlds programme to detect the presence of atmospheres around M dwarfs, as well as multiple other large programmes like GO 3730 (PIs: Diamond-Lowe & Mendonca) and GO 7073 (PIs: Lustig-Yaeger & Stevenson). These efforts are proceeding through both emission and transmission surveys, with the main limitations of each being degeneracies in heat redistribution (Hammond et al., 2025) and stellar contamination, respectively. Another major area of advancement possible with JWST includes measuring spatial information on exoplanets, primarily proceeding via limb asymmetries, eclipse mapping, and phase curves. Given the wealth of data obtained, we will be able to combine large surveys of planets to compare chemical composition and cloud properties across the parameter space. Furthermore, I anticipate a deeper exploration of temperate sub-Neptunes, which have already revealed a fascinating array of chemical diversity but are also subject to a wide range of interpretations (e.g., Benneke et al., 2024; Shorttle et al., 2024; Wogan et al., 2024; Madhusudhan et al., 2025; Constantinou et al., 2026).

Another space-based mission, scheduled for the near future, is Ariel (Tinetti

et al., 2018), currently planned to launch in 2031. Designed for the purpose of exoplanet characterization, Ariel will be able to obtain a large survey of exoplanet spectra at infrared wavelengths. With a much smaller telescope aperture but able to fully dedicate its time to exoplanets, Ariel will be able to complement the deep analyses of JWST with time-intensive broad analyses. This could enable insights such as studying variability and constraining the impact of stellar activity using multiple observations, and dedicating significant time to phase curves and eclipse mapping (Valentine et al., 2025). Further into the future, the HWO (Feinberg et al., 2024) and LIFE (Quanz et al., 2022) missions represent the future of space-based characterization, through direct imaging of habitable exoplanets at the wavelengths of reflected starlight and thermal emission respectively.

7.3.4 Conclusions

Through a combination of the discussed observational advances, we will finally be able to explore individual planets in immense detail, and compare atmospheric properties across large samples of exoplanets. This truly opens a new frontier in exoplanetology, accessible through two main philosophies: a planetary science philosophy, where individual planets are attempted to be understood as holistic systems; and an astronomy philosophy, where astrophysical measurements from individual planets are combined and synthesized in such a manner as to gain a systematic understanding of planetary systems (consider the Hertzsprung-Russell diagram for stars). I am thrilled to be a part of this frontier as both an explorer, experiencing the colours and smells of new worlds (in a manner of speaking), and an archaeologist, piecing together the histories of entire planets from the evidence that remains.

Bibliography

- Abel M., Frommhold L., Li X., Hunt K. L. C., 2011, The Journal of Physical Chemistry A, 115, 6805
- Abel M., Frommhold L., Li X., Hunt K. L. C., 2012, J. Chem. Phys., 136, 044319
- Ackerman A. S., Marley M. S., 2001, The Astrophysical Journal, 556, 872
- Ahrer E.-M., 2023, PhD thesis, University of Warwick
- Ahrer E., Wheatley P., Kirk J., Gandhi S., King G., Loudén T., 2022, Monthly Notices of the Royal Astronomical Society, 510, 4857
- Ahrer E., Wheatley P., Gandhi S., Kirk J., King G. W., Loudén T., Welbanks L., 2023a, Monthly Notices of the Royal Astronomical Society, 521, 5636
- Ahrer E.-M., et al., 2023b, Nature, 614, 653
- Ahrer E., et al., 2024, Monthly Notices of the Royal Astronomical Society, 530, 2749
- Ahrer E.-M., et al., 2025a, Monthly Notices of the Royal Astronomical Society, 540, 2535
- Ahrer E.-M., et al., 2025b, Monthly Notices of the Royal Astronomical Society, 543, 2442
- Akeson R., et al., 2013, Publications of the Astronomical Society of the Pacific, 125, 989
- Al-Refaie A. F., Venot O., Changeat Q., Edwards B., 2024, The Astrophysical Journal, 967, 132
- Alam M. K., et al., 2021, The Astrophysical Journal Letters, 906, L10
- Albrecht S., et al., 2012, The Astrophysical Journal, 757, 18
- Alderson L., et al., 2020, Monthly Notices of the Royal Astronomical Society, 497, 5182
- Alderson L., Grant D., Wakeford H., 2022, Zenodo
- Alderson L., et al., 2023, Nature, 614, 664

- Alderson L., et al., 2024, *The Astronomical Journal*, 167, 216
- An Q., Brandt T. D., Brandt G. M., Venner A., 2025, *The Astrophysical Journal Supplement Series*, 280, 61
- Anderson D. R., et al., 2010, *ApJ*, 709, 159
- Anderson D. R., et al., 2018, arXiv e-prints, p. arXiv:1809.07709
- Appleby J. F., 1990, *Icarus*, 85, 355
- Asplund M., Grevesse N., Sauval A. J., Scott P., 2009, *Annual review of astronomy and astrophysics*, 47, 481
- Atreya S., Romani P., 1985, *Recent advances in planetary meteorology*, 17, 68
- Atreya S. K., Donahue T. M., Kuhn W. R., 1978, *Science*, 201, 611
- Atreya S., Wong M., Owen T., Mahaffy P., Niemann H., De Pater I., Drossart P., Encrenaz T., 1999, *Planetary and Space Science*, 47, 1243
- August P. C., Bean J. L., Zhang M., Lunine J., Xue Q., Line M., Smith P. C., 2023, *The Astrophysical Journal Letters*, 953, L24
- Azzam A. A., Tennyson J., Yurchenko S. N., Naumenko O. V., 2016, *Monthly Notices of the Royal Astronomical Society*, 460, 4063
- Baeyens R., Decin L., Carone L., Venot O., Agúndez M., Mollière P., 2021, *Monthly Notices of the Royal Astronomical Society*, 505, 5603
- Bakos G., Noyes R., Kovács G., Stanek K., Sasselov D. D., Domsa I., 2004, *Publications of the Astronomical Society of the Pacific*, 116, 266
- Banerjee A., Barstow J. K., Haswell C. A., Lewis S. R., 2023, *Monthly Notices of the Royal Astronomical Society: Letters*, 523, L64
- Banerjee A., et al., 2024, *The Astrophysical Journal Letters*, 975, L11
- Baraffe I., Chabrier G., Barman T., 2008, *Astronomy & Astrophysics*, 482, 315
- Barber R., Strange J., Hill C., Polyansky O., Mellau G. C., Yurchenko S., Tennyson J., 2014, *Monthly Notices of the Royal Astronomical Society*, 437, 1828
- Barclay T., et al., 2013, *Nature*, 494, 452
- Barclay T., Kostov V. B., Colón K. D., Quintana E. V., Schlieder J. E., Louie D. R., Gilbert E. A., Mullally S. E., 2021, *The Astronomical Journal*, 162, 300
- Bardet D., Changeat Q., Venot O., Panek E., 2025, *Astronomy & Astrophysics*, 699, A342
- Barrado D., et al., 2023, *Nature*, 624, 263

- Barstow J. K., Aigrain S., Irwin P. G., Sing D. K., 2017, *The Astrophysical Journal*, 834, 50
- Barth C. A., 1974, *Annual Review of Earth and Planetary Sciences*, Vol. 2, p. 333, 2, 333
- Basilevsky A. T., Head J. W., 2003, *Reports on Progress in Physics*, 66, 1699
- Batalha N. E., et al., 2026, *The Astronomical Journal*, 171, 98
- Bate M., Lodato G., Pringle J., 2010, *Monthly Notices of the Royal Astronomical Society*, 401, 1505
- Battley M., et al., 2024, *Astronomy & Astrophysics*, 686, A230
- Batygin K., 2012, *Nature*, 491, 418
- Batygin K., Stevenson D. J., 2010, *The Astrophysical Journal Letters*, 714, L238
- Batygin K., Bodenheimer P. H., Laughlin G. P., 2016, *The Astrophysical Journal*, 829, 114
- Bean J. L., et al., 2023, *Nature*, 618, 43
- Bell T. J., et al., 2022, *JOSS*, 7, 4503
- Bell T. J., et al., 2023, *Nature*, 623, 709
- Bello-Arufe A., et al., 2025, arXiv preprint arXiv:2511.15835
- Benn C., Dee K., Agócs T., 2008, in *Ground-based and Airborne Instrumentation for Astronomy II*. pp 2384–2395
- Benneke B., et al., 2019, *The Astrophysical Journal Letters*, 887, L14
- Benneke B., et al., 2024, arXiv preprint arXiv:2403.03325
- Bergin E. A., et al., 2024, *The Astrophysical Journal*, 965, 147
- Bernabò L. M., Csizmadia S., Smith A. M. S., Rauer H., Hatzes A., Esposito M., Gandolfi D., Cabrera J., 2024, *A&A*, 684, A78
- Bernarth P. F., 1995, *Spectra of Atoms and Molecules*. Oxford University Press
- Beyer A., White R., 2024, *American Astronomical Society Meeting Abstracts*, 243, 205.08
- Bieryla A., et al., 2015, *AJ*, 150, 12
- Birkby J., De Kok R., Brogi M., De Mooij E., Schwarz H., Albrecht S., Snellen I., 2013, *Monthly Notices of the Royal Astronomical Society: Letters*, 436, L35
- Birkmann S., et al., 2022, *Astronomy & Astrophysics*, 661, A83

- Bixel A., et al., 2019, *The Astronomical Journal*, 157, 68
- Blain D., Mollière P., Nasedkin E., 2024, *Journal of Open Source Software*, 9, 7028
- Blažek M., et al., 2022, *MNRAS*, 513, 3444
- Bloot S., Miguel Y., Bazot M., Howard S., 2023, *Monthly Notices of the Royal Astronomical Society*, 523, 6282
- Blundell S. J., Blundell Katherine M., 2006, *Concepts in Thermal Physics*. Oxford University Press
- Bodenheimer P., Lin D., Mardling R., 2001, *The Astrophysical Journal*, 548, 466
- Bond I. A., et al., 2004, *The Astrophysical Journal Letters*, 606, L155
- Bonomo A. S., et al., 2017, *A&A*, 602, A107
- Booth R. A., Clarke C. J., Madhusudhan N., Ilee J. D., 2017, *MNRAS*, 469, 3994
- Borucki W. J., et al., 2010, *Science*, 327, 977
- Borysow A., 2002, *Astronomy & Astrophysics*, 390, 779
- Borysow J., Frommhold L., Birnbaum G., 1988, *Astrophysical Journal*, Part 1 (ISSN 0004-637X), vol. 326, March 1, 1988, p. 509-515. NASA-supported research., 326, 509
- Borysow A., Jørgensen U. G., Fu Y., 2001, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 68, 235
- Bosman A. D., Banzatti A., Bruderer S., Tielens A. G., Blake G. A., van Dishoeck E. F., 2019, *Astronomy & Astrophysics*, 631, A133
- Boss A. P., 1997, *Science*, 276, 1836
- Bouchy F., et al., 2010, *Astronomy & Astrophysics*, 519, A98
- Bourrier V., Des Etangs A. L., Ehrenreich D., Tanaka Y., Vidotto A. A., 2016, *Astronomy & Astrophysics*, 591, A121
- Brady R. P., Yurchenko S. N., Tennyson J., Kim G.-S., 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 6675
- Brande J., et al., 2024, *The Astrophysical Journal Letters*, 961, L23
- Brogi M., Line M. R., 2019, *The Astronomical Journal*, 157, 114
- Brogi M., De Kok R., Birkby J., Schwarz H., Snellen I., 2014, *Astronomy & Astrophysics*, 565, A124
- Brogi M., Kok R. d., Albrecht S., Snellen I., Birkby J., Schwarz H., 2016, *The Astrophysical Journal*, 817, 106

- Brogi M., Line M., Bean J., Désert J.-M., Schwarz H., 2017, *The Astrophysical Journal Letters*, 839, L2
- Brown M. E., Bouchez A. H., Griffith C. A., 2002, *Nature*, 420, 795
- Bruno G., 1584, *De l'infinito universo et mundi*
- Bryant E. M., Bayliss D., Van Eylen V., 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 3663
- Bryant E. M., et al., 2025, *Nature Astronomy*, 9, 1031
- Buchner J., et al., 2014, *Astronomy & Astrophysics*, 564, A125
- Burrows A., Sharp C., 1999, *The Astrophysical Journal*, 512, 843
- Burrows A., Marley M., Sharp C., 2000, *The Astrophysical Journal*, 531, 438
- Burrows A., Hubeny I., Budaj J., Hubbard W., 2007, *The Astrophysical Journal*, 661, 502
- Burrows A., Budaj J., Hubeny I., 2008, *The Astrophysical Journal*, 678, 1436
- Caldas A., Leconte J., Selsis F., Waldmann I., Bordé P., Rocchetto M., Charnay B., 2019, *Astronomy & Astrophysics*, 623, A161
- Cañas C. I., et al., 2020, *The Astronomical Journal*, 160, 147
- Cañas C. I., et al., 2025, arXiv e-prints, pp arXiv-2502
- Carrión-González Ó., Muñoz A. G., Santos N., Cabrera J., Csizmadia S., Rauer H., 2021, *Astronomy & Astrophysics*, 651, A7
- Carter A., et al., 2024, *Nature Astronomy*, 8, 1008
- Cegla H., et al., 2023, *Astronomy & Astrophysics*, 674, A174
- Chabrier G., Gallardo J., Baraffe I., 2007, *Astronomy & Astrophysics*, 472, L17
- Chachan Y., Knutson H. A., Lothringer J., Blake G. A., 2023, *The Astrophysical Journal*, 943, 112
- Chachan Y., Fortney J. J., Ohno K., Thorngren D., Murray-Clay R., 2025, *The Astrophysical Journal*, 994, 43
- Chan Y., Dalgarno A., 1965, *Proceedings of the Physical Society*, 85, 227
- Chandrasekhar S., 1960, *Radiative transfer*. Dover
- Charbonneau D., Brown T. M., Latham D. W., Mayor M., 2000, *The Astrophysical Journal Letters*, 529, L45

- Charbonneau D., Brown T. M., Noyes R. W., Gilliland R. L., 2002, *The Astrophysical Journal*, 568, 377
- Chauvin G., Lagrange A.-M., Dumas C., Zuckerman B., Mouillet D., Song I., Beuzit J.-L., Lowrance P., 2004, *Astronomy & Astrophysics*, 425, L29
- Chen G., et al., 2018, *Astronomy & Astrophysics*, 616, A145
- Chen G., Casasayas-Barris N., Pallé E., Welbanks L., Madhusudhan N., Luque R., Murgas F., 2020, *Astronomy & Astrophysics*, 642, A54
- Chen G., Pallé E., Parviainen H., Murgas F., Yan F., 2021, *The Astrophysical Journal Letters*, 913, L16
- Claret A., 2000, *Astronomy and Astrophysics*, v. 363, p. 1081-1190 (2000), 363, 1081
- Claringbold A. B., Rimmer P., Rugheimer S., Shorttle O., 2023, *The Astronomical Journal*, 166, 39
- Claringbold A. B., Fisher C. E., Kirk J., Ahrer E.-M., Penzlin A. B. T., Thorngren D. P., et al., 2026a, *Monthly Notices of the Royal Astronomical Society*, 546, stag143
- Claringbold A. B., Wheatley P. J., Kirk J., Ahrer E.-M., Skillen I., Brogi M., et al., 2026b, *Monthly Notices of the Royal Astronomical Society*, 547, stag384
- Coles P. A., Yurchenko S. N., Tennyson J., 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 4638
- Collaboration G., et al., 2016, arXiv preprint arXiv:1609.04153
- Constantinou S., Madhusudhan N., Holmberg M., 2026, *Astronomy & Astrophysics*, 705, A25
- Corrales L., Gavilan L., Teal D., Kempton E. M.-R., 2023, *The Astrophysical Journal Letters*, 943, L26
- Coulombe L.-P., et al., 2023, *Nature*, 620, 292
- Cowan N. B., Machalek P., Croll B., Shekhtman L. M., Burrows A., Deming D., Greene T., Hora J. L., 2012, *The Astrophysical Journal*, 747, 82
- Cox A. N., 2000, *Allen's astrophysical quantities*
- Cridland A. J., van Dishoeck E. F., Alessi M., Pudritz R. E., 2019, *Astronomy & Astrophysics*, 632, A63
- Cridland A. J., van Dishoeck E. F., Alessi M., Pudritz R. E., 2020, *Astronomy & Astrophysics*, 642, A229
- Crossfield I. J., et al., 2025, *The Astrophysical Journal*, 994, 184

- Curiel S., Ortiz-León G. N., Mioduszewski A. J., Sanchez-Bermudez J., 2022, *The Astronomical Journal*, 164, 93
- Dalgarno A., Williams D., 1962, *Astrophysical Journal*, vol. 136, p. 690-692, 136, 690
- Dash S., Brogi M., Gandhi S., Lafarga M., Meech A., Bello-Arufe A., Wheatley P. J., 2024, *Monthly Notices of the Royal Astronomical Society*, 530, 3100
- Dawson R. I., Johnson J. A., 2018, *ARA&A*, 56, 175
- Deming D., et al., 2013, *The Astrophysical Journal*, 774, 95
- Dodson-Robinson S. E., Veras D., Ford E. B., Beichman C. A., 2009, *The Astrophysical Journal*, 707, 79
- Doyon R., et al., 2012, in *Space telescopes and instrumentation 2012: Optical, infrared, and millimeter wave*. pp 1005–1017
- Durisen R. H., Boss A. P., Mayer L., Nelson A. F., Quinn T., Rice W., 2007, *Protostars and planets V*, p. 607
- Dymont A. H., Yu X., Ohno K., Zhang X., Fortney J. J., Thorngren D., Dickinson C., 2022, *The Astrophysical Journal*, 937, 90
- Dyrek A., et al., 2024, *Nature*, 625, 51
- Ehrenreich D., et al., 2020, *Nature*, 580, 597
- Einstein A., 1936, *Science*, 84, 506
- Eistrup C., Walsh C., van Dishoeck E. F., 2018, *A&A*, 613, A14
- Encrenaz T., Tinetti G., Coustenis A., 2018, *Experimental Astronomy*, 46, 31
- Eschen Y. N. E., Bayliss D., Wilson T. G., Kunimoto M., Pelisoli I., Rodel T., 2024, *Monthly Notices of the Royal Astronomical Society*, 535, 1778
- Esparza-Borges E., et al., 2023, *The Astrophysical Journal Letters*, 955, L19
- Esparza-Borges E., et al., 2025, *Monthly Notices of the Royal Astronomical Society*, 543, 3456
- Espinoza N., Jordán A., 2015, *Monthly Notices of the Royal Astronomical Society*, 450, 1879
- Espinoza N., Jordán A., 2016, *Monthly Notices of the Royal Astronomical Society*, 457, 3573
- Espinoza N., Fortney J. J., Miguel Y., Thorngren D., Murray-Clay R., 2017, *The Astrophysical Journal Letters*, 838, L9

- Espinoza N., et al., 2019, *Monthly Notices of the Royal Astronomical Society*, 482, 2065
- Espinoza N., et al., 2024, *Nature*, 632, 1017
- Esposito L. W., 1984, *Science*, 223, 1072
- Estrela R., Swain M. R., Roudier G. M., 2022, *The Astrophysical Journal Letters*, 941, L5
- Evans-Soma T. M., et al., 2025, *Nature Astronomy*, 9, 845
- Evans T. M., et al., 2013, *The Astrophysical Journal Letters*, 772, L16
- Evans T. M., et al., 2016, *The Astrophysical Journal Letters*, 822, L4
- Evans T. M., et al., 2017, *Nature*, 548, 58
- Faedi F., et al., 2011, *Astronomy & Astrophysics*, 531, A40
- Fairman C., Wakeford H. R., MacDonald R. J., 2024, *The Astronomical Journal*, 167, 240
- Fairman C., Wakeford H. R., Claringbold A. B., Kirk J., Ahrer E.-M., Thorngren D. P., et al., 2026, *Monthly Notices of the Royal Astronomical Society*, p. stag528
- Feinberg L., et al., 2024, in *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*. pp 548–559
- Feinstein A. D., et al., 2023, *Nature*, 614, 670
- Feng F., et al., 2022, *The Astrophysical Journal Supplement Series*, 262, 21
- Fernandes R. B., Mulders G. D., Pascucci I., Mordasini C., Emsenhuber A., 2019, *The Astrophysical Journal*, 874, 81
- Feroz F., Hobson M. P., 2008, *MNRAS*, 384, 449
- Feroz F., Hobson M., Bridges M., 2009, *Monthly Notices of the Royal Astronomical Society*, 398, 1601
- Fischer P. D., et al., 2016, *The Astrophysical Journal*, 827, 19
- Fisher C., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 535, 27
- Fletcher L., Orton G., Teanby N., Irwin P., Bjoraker G., 2009, *Icarus*, 199, 351
- Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
- Fortney J. J., 2005, *Monthly Notices of the Royal Astronomical Society*, 364, 649

- Fortney J. J., Saumon D., Marley M., Lodders K., Freedman R., 2006, *The Astrophysical Journal*, 642, 495
- Fortney J. J., Marley M. S., Barnes J. W., 2007, *The Astrophysical Journal*, 659, 1661
- Fortney J. J., Lodders K., Marley M. S., Freedman R. S., 2008, *The Astrophysical Journal*, 678, 1419
- Fortney J. J., Visscher C., Marley M. S., Hood C. E., Line M. R., Thorngren D. P., Freedman R. S., Lupu R., 2020, *The Astronomical Journal*, 160, 288
- Fournier-Tondreau M., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 528, 3354
- Fraine J., et al., 2014, *Nature*, 513, 526
- France K., et al., 2016, *The Astrophysical Journal*, 820, 89
- Fu Q., Liou K., 1992, *Journal of Atmospheric Sciences*, 49, 2139
- Fu G., Deming D., Knutson H., Madhusudhan N., Mandell A., Fraine J., 2017, *The Astrophysical Journal Letters*, 847, L22
- Fu G., et al., 2022, *The Astrophysical Journal Letters*, 925, L3
- Fu G., et al., 2024, *Nature*, 632, 752
- Fu G., et al., 2025a, *The Astrophysical Journal*, 986, 1
- Fu G., et al., 2025b, *The Astrophysical Journal Letters*, 989, L17
- Fuller K. A., Downing H. D., Querry M. R., 1997, in Palik E. D., ed., , *Handbook of Optical Constants of Solids*. Academic Press, Burlington, pp 899–922, doi:<https://doi.org/10.1016/B978-012544415-6.50139-4>, <https://www.sciencedirect.com/science/article/pii/B9780125444156501394>
- Fulton B. J., et al., 2021, *The Astrophysical Journal Supplement Series*, 255, 14
- Gaia Collaboration et al., 2023, *A&A*, 674, A1
- Galilei G., 1610, *Sidereus Nuncius*. Thomas Baglioni
- Galle J., 1846, *Monthly Notices of the Royal Astronomical Society*, Vol. 7, p. 153, 7, 153
- Gan T., et al., 2023, *The Astronomical Journal*, 165, 17
- Gandhi S., et al., 2023, *The Astronomical Journal*, 165, 242
- Gao P., Marley M. S., Zahnle K., Robinson T. D., Lewis N. K., 2017, *The Astronomical Journal*, 153, 139

- Gao P., et al., 2020, *Nature Astronomy*, 4, 951
- Gao P., Wakeford H. R., Moran S. E., Parmentier V., 2021, *Aerosols in exoplanet atmospheres*
- Gapp C., et al., 2025, *The Astronomical Journal*, 169, 341
- Gardner J. P., et al., 2006, *Space Science Reviews*, 123, 485
- Gascón C., et al., 2025, *The Astronomical Journal*, 170, 178
- Gaudi B. S., Winn J. N., 2007, *The Astrophysical Journal*, 655, 550
- Gaudi B. S., et al., 2017, *Nature*, 546, 514
- Ghezzi L., Montet B. T., Johnson J. A., 2018, *The Astrophysical Journal*, 860, 109
- Giacobbe P., et al., 2021, *Nature*, 592, 205
- Gibson N., Aigrain S., Roberts S., Evans T., Osborne M., Pont F., 2012a, *Monthly notices of the royal astronomical society*, 419, 2683
- Gibson N. P., et al., 2012b, *Monthly Notices of the Royal Astronomical Society*, 422, 753
- Gibson N. P., Nikolov N., Sing D. K., Barstow J. K., Evans T. M., Kataria T., Wilson P. A., 2017, *Monthly Notices of the Royal Astronomical Society*, 467, 4591
- Gibson N. P., Nugroho S. K., Lothringer J., Maguire C., Sing D. K., 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 4618
- Gill S., et al., 2020, *The Astrophysical Journal Letters*, 898, L11
- Gill S., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 532, 1444
- Gillon M., et al., 2026, *Nature Astronomy*, pp 1–21
- Gilmozzi R., Spyromilio J., 2007, *The Messenger*, 127, 3
- Gladstone G. R., Allen M., Yung Y., 1996, *Icarus*, 119, 1
- Goody R. M., Yung Y. L., 1989, *Atmospheric radiation : theoretical basis*. Oxford University Press
- Gordon S., McBride B. J., 1994, Technical report, Computer program for calculation of complex chemical equilibrium compositions and applications. Part 1: Analysis. NASA
- Grant D., Wakeford H. R., 2024, arXiv preprint arXiv:2408.10341
- Grant D., et al., 2023a, *The Astrophysical Journal Letters*, 949, L15
- Grant D., et al., 2023b, *The Astrophysical Journal Letters*, 956, L29

- Greene T. P., Bell T. J., Ducrot E., Dyrek A., Lagage P.-O., Fortney J. J., 2023, *Nature*, 618, 39
- Gressier A., et al., 2025, *The Astronomical Journal*, 170, 292
- Grimm S. L., Heng K., 2015a, *ApJ*, 808, 182
- Grimm S. L., Heng K., 2015b, *The Astrophysical Journal*, 808, 182
- Grimm S. L., et al., 2021, *ApJS*, 253, 30
- Gueymard C. A., 2004, *Solar energy*, 76, 423
- Guillot T., 2005, *Annu. Rev. Earth Planet. Sci.*, 33, 493
- Guillot T., 2010, *Astronomy & Astrophysics*, 520, A27
- Gupta S., Ochiai E., Ponnampereuma C., 1981, *Nature*, 293, 725
- Hammond M., et al., 2024, *The Astronomical Journal*, 168, 4
- Hammond M., et al., 2025, *The Astrophysical Journal Letters*, 978, L40
- Hansen J. E., Hovenier J., 1974, *J. atmos. Sci*, 31, 1137
- Hartman J. D., et al., 2014, *The Astronomical Journal*, 147, 128
- Hasting W. K., 1970, *Biometrika*, 57, 97
- Haswell C. A., et al., 2020, *Nature Astronomy*, 4, 408
- Hawthorn F., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 528, 1841
- Haynes K., Mandell A. M., Madhusudhan N., Deming D., Knutson H., 2015, *The Astrophysical Journal*, 806, 146
- He C., et al., 2024, *Nature Astronomy*, 8, 182
- Helled R., Stevenson D., 2017, *The Astrophysical Journal Letters*, 840, L4
- Hellier C., et al., 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 739
- Hellier C., et al., 2014, *Monthly Notices of the Royal Astronomical Society*, 440, 1982
- Helling C., Oevermann M., Lüttke M., Klein R., Sedlmayr E., 2001, *Astronomy & Astrophysics*, 376, 194
- Helling C., Woitke P., Rimmer P. B., Kamp I., Thi W.-F., Meijerink R., 2014, *Life*, 4, 142
- Helling C., et al., 2016, *MNRAS*, 460, 855

- Heng K., 2016, *The Astrophysical Journal Letters*, 826, L16
- Heng K., Mendonça J. M., Lee J.-M., 2014, *The Astrophysical Journal Supplement Series*, 215, 4
- Herschel W., 1781, *Philosophical Transactions of the Royal Society of London*, pp 492–501
- Hobson M. J., et al., 2023, *The astronomical journal*, 166, 201
- Hoeijmakers H. J., et al., 2019, *Astronomy & Astrophysics*, 627, A165
- Horne K., 1986, *Publications of the Astronomical Society of the Pacific*, 98, 609
- Horner S. D., Rieke M. J., 2004, in *Optical, infrared, and millimeter space telescopes*. pp 628–634
- Houllé M., et al., 2021, *Astronomy & Astrophysics*, 652, A67
- Hu R., 2021, *The Astrophysical Journal*, 921, 27
- Hu R., Seager S., Bains W., 2012, *The Astrophysical Journal*, 761, 166
- Huang C., Wu Y., Triaud A. H., 2016, *The Astrophysical Journal*, 825, 98
- Hubeny I., Burrows A., Sudarsky D., 2003, *The Astrophysical Journal*, 594, 1011
- Huber D., et al., 2013, *Science*, 342, 331
- Huitson C., Désert J.-M., Bean J., Fortney J., Stevenson K., Bergmann M., 2017, *The Astronomical Journal*, 154, 95
- Husser T.-O., Wende-von Berg S., Dreizler S., Homeier D., Reiners A., Barman T., Hauschildt P. H., 2013, *Astronomy & Astrophysics*, 553, A6
- Huygens C., 1656, *De Saturni Luna Observatio Nova*. Adrian Vlacq
- Ingersoll A. P., Porco C. C., 1978, *Icarus*, 35, 27
- Innes H., Tsai S.-M., Pierrehumbert R. T., 2023, *The Astrophysical Journal*, 953, 168
- Irwin P., et al., 2008, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 109, 1136
- Ivshina E. S., Winn J. N., 2022, *ApJS*, 259, 62
- JWST Transiting Exoplanet Community Early Release Science Team et al., 2023, *Nature*, 614, 649
- Jakobsen P., et al., 2022, *A&A*, 661, A80
- Janson M., et al., 2013, *The Astrophysical Journal Letters*, 778, L4

- Jeffreys H., 1939, *Theory of Probability*
- Johnson J. A., Aller K. M., Howard A. W., Crepp J. R., 2010, *Publications of the Astronomical Society of the Pacific*, 122, 905
- Johnson J. A., et al., 2011, *ApJ*, 735, 24
- Jones K., Espinoza N., 2020, *The Journal of Open Source Software*, 5, 2382
- Jordán A., et al., 2013, *The Astrophysical Journal*, 778, 184
- Kanodia S., et al., 2026, arXiv preprint arXiv:2603.16464
- Kass R. E., Raftery A. E., 1995, *Journal of the american statistical association*, 90, 773
- Kawashima Y., Ikoma M., 2018, *The Astrophysical Journal*, 853, 7
- Kawashima Y., Ikoma M., 2019, *The Astrophysical Journal*, 877, 109
- Kempton E. M.-R., et al., 2023, *Nature*, 620, 67
- Kennedy G. M., Kenyon S. J., 2008, *The Astrophysical Journal*, 673, 502
- Khalafinejad S., et al., 2021, *Astronomy & Astrophysics*, 656, A142
- Khare B., et al., 1984a, *Advances in Space Research*, 4, 59
- Khare B., Sagan C., Arakawa E., Suits F., Callcott T., Williams M., 1984b, *Icarus*, 60, 127
- Kipping D. M., 2013, *Monthly Notices of the Royal Astronomical Society*, 435, 2152
- Kipping D. M., et al., 2016, *The Astrophysical Journal*, 820, 112
- Kirk J., 2018, PhD thesis, University of Warwick
- Kirk J., Owen J. E., 2025, arXiv preprint arXiv:2510.27386
- Kirk J., Wheatley P. J., Loudén T., Littlefair S., Copperwheat C., Armstrong D. J., Marsh T., Dhillon V., 2016, *Monthly Notices of the Royal Astronomical Society*, 463, 2922
- Kirk J., Wheatley P. J., Loudén T., Doyle A., Skillen I., McCormac J., Irwin P., Karjalainen R., 2017, *Monthly Notices of the Royal Astronomical Society*, 468, 3907
- Kirk J., Wheatley P. J., Loudén T., Skillen I., King G. W., McCormac J., Irwin P., 2018, *Monthly Notices of the Royal Astronomical Society*, 474, 876
- Kirk J., López-Morales M., Wheatley P. J., Weaver I. C., Skillen I., Loudén T., McCormac J., Espinoza N., 2019, *The Astronomical Journal*, 158, 144

Kirk J., Alam M. K., López-Morales M., Zeng L., 2020, *The Astronomical Journal*, 159, 115

Kirk J., et al., 2021, *The Astronomical Journal*, 162, 34

Kirk J., et al., 2024a, *RAS Techniques and Instruments*, 3, 691

Kirk J., et al., 2024b, *AJ*, 167, 90

Kirk J., et al., 2025, *Monthly Notices of the Royal Astronomical Society*, p. staf208

Kitzmann D., Heng K., 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 94

Kitzmann D., Heng K., Oreshenko M., Grimm S. L., Apai D., Bowler B. P., Burgasser A. J., Marley M. S., 2020, *ApJ*, 890, 174

Knutson H. A., et al., 2007, *Nature*, 447, 183

Knutson H. A., Charbonneau D., Burrows A., O'Donovan F. T., Mandushev G., 2009, *The Astrophysical Journal*, 691, 866

Knutson H. A., et al., 2014, *The Astrophysical Journal*, 785, 126

Kokori A., et al., 2023, *The Astrophysical Journal Supplement Series*, 265, 4

Konopacky Q. M., Barman T. S., Macintosh B. A., Marois C., 2013, *Science*, 339, 1398

Kraft R. P., 1967, *ApJ*, 150, 551

Krasnopolsky V., Parshev V., 1981, *Nature*, 292, 610

Kreidberg L., 2015a, *Publications of the Astronomical Society of the Pacific*, 127, 1161

Kreidberg L., 2015b, *PASP*, 127, 1161

Kreidberg L., et al., 2014a, *Nature*, 505, 69

Kreidberg L., et al., 2014b, *The Astrophysical Journal Letters*, 793, L27

Kreidberg L., et al., 2015, *The Astrophysical Journal*, 814, 66

Kreidberg L., et al., 2018, *The Astronomical Journal*, 156, 17

Krishnamurthy V., et al., 2025, *Nature Astronomy*, pp 1–13

Kühnle H., et al., 2025, *Astronomy & Astrophysics*, 695, A224

Lacis A. A., Oinas V., 1991, *Journal of Geophysical Research: Atmospheres*, 96, 9027

- Lagrange A.-M., et al., 2009, *Astronomy & Astrophysics*, 506, 927
- Lambrechts M., Johansen A., Morbidelli A., 2014, *Astronomy & Astrophysics*, 572, A35
- Lammers C., Winn J. N., 2026, *The Astronomical Journal*, 171, 18
- Law C. J., et al., 2021a, *The Astrophysical Journal Supplement Series*, 257, 3
- Law C. J., et al., 2021b, *ApJS*, 257, 3
- Lecavelier des Etangs A., Pont F., Vidal-Madjar A., Sing D., 2008a, *Astronomy & Astrophysics*, 481, L83
- Lecavelier des Etangs A., Vidal-Madjar A., Désert J.-M., Sing D., 2008b, *Astronomy & Astrophysics*, 485, 865
- Leconte J., Selsis F., Hersant F., Guillot T., 2017, *Astronomy & Astrophysics*, 598, A98
- Lee E. J., Chiang E., 2016, *The Astrophysical Journal*, 817, 90
- Lee E. J., Chiang E., Ormel C. W., 2014, *The Astrophysical Journal*, 797, 95
- Lee E. K., Tan X., Tsai S.-M., 2024, *Monthly Notices of the Royal Astronomical Society*, 529, 2686
- Lendl M., et al., 2016, *Astronomy & Astrophysics*, 587, A67
- Levenberg K., 1944, *Quarterly of applied mathematics*, 2, 164
- Lewis J. S., 1969, *Icarus*, 10, 365
- Lewis N. K., et al., 2020, *The Astrophysical Journal Letters*, 902, L19
- Li G., Gordon I. E., Rothman L. S., Tan Y., Hu S.-M., Kassi S., Campargue A., Medvedev E. S., 2015, *ApJS*, 216, 15
- Ligterink N. F., Kipfer K., Gavino S., 2024, *Astronomy & Astrophysics*, 687, A224
- Lin A. S., et al., 2023, *The Astronomical Journal*, 166, 90
- Lindal G. F., et al., 1981, *Journal of Geophysical Research: Space Physics*, 86, 8721
- Lindal G. F., Sweetnam D., Eshleman V., 1985, *Astronomical Journal* (ISSN 0004-6256), vol. 90, June 1985, p. 1136-1146. NASA-supported research., 90, 1136
- Lindal G. F., Lyons J., Sweetnam D., Eshleman V., Hinson D. P., Tyler G. L., 1987, *Journal of Geophysical Research: Space Physics*, 92, 14987
- Lindal G., Lyons J., Sweetnam D., Eshleman V., Hinson D., Tyler G., 1990, *Geophysical Research Letters*, 17, 1733

Line M. R., et al., 2013, *The Astrophysical Journal*, 775, 137

Lodders K., Fegley Jr B., 2002, *Icarus*, 155, 393

Lodders K., Fegley Jr B., 2006, in , *Astrophysics Update 2*. Springer, pp 1–28

Lothringer J. D., Barman T., Koskinen T., 2018, *The Astrophysical Journal*, 866, 27

Louden T., Wheatley P. J., 2015, *The Astrophysical Journal Letters*, 814, L24

Louden T., Wheatley P. J., Irwin P. G., Kirk J., Skillen I., 2017, *Monthly Notices of the Royal Astronomical Society*, 470, 742

Louie D. R., et al., 2025, *The Astronomical Journal*, 169, 86

Loyd R. P., et al., 2016, *The Astrophysical Journal*, 824, 102

Lueber A., Novais A., Fisher C., Heng K., 2024, *Astronomy & Astrophysics*, 687, A110

Luque R., et al., 2020, *Astronomy & Astrophysics*, 642, A50

Lustig-Yaeger J., et al., 2025, *The Astrophysical Journal Letters*, 994, L4

MacDonald R. J., 2023, *Journal of Open Source Software*, 8, 4873

MacDonald R. J., Batalha N. E., 2023, *Research Notes of the AAS*, 7, 54

MacDonald R. J., Lewis N. K., 2022, *The Astrophysical Journal*, 929, 20

MacDonald R. J., Madhusudhan N., 2017, *Monthly Notices of the Royal Astronomical Society*, 469, 1979

MacDonald R. J., Marley M. S., Fortney J. J., Lewis N. K., 2018, *The Astrophysical Journal*, 858, 69

MacDonald R. J., Goyal J. M., Lewis N. K., 2020, *ApJ*, 893, L43

MacDonald R. J., et al., 2026, *Nature*, 655, 76

Macintosh B., et al., 2015, *Science*, 350, 64

Madhusudhan N., 2012, *The Astrophysical Journal*, 758, 36

Madhusudhan N., Seager S., 2009, *The Astrophysical Journal*, 707, 24

Madhusudhan N., Seager S., 2011, *The Astrophysical Journal*, 729, 41

Madhusudhan N., Amin M. A., Kennedy G. M., 2014, *ApJ*, 794, L12

Madhusudhan N., Sarkar S., Constantinou S., Holmberg M., Piette A. A., Moses J. I., 2023, *The Astrophysical Journal Letters*, 956, L13

- Madhusudhan N., Constantinou S., Holmberg M., Sarkar S., Piette A. A., Moses J. I., 2025, *The Astrophysical Journal Letters*, 983, L40
- Magic Z., Chiavassa A., Collet R., Asplund M., 2015, *A&A*, 573, A90
- Mah J., Savvidou S., Bitsch B., 2024, *Astronomy & Astrophysics*, 686, L17
- Majeau C., Agol E., Cowan N. B., 2012, *The Astrophysical Journal Letters*, 747, L20
- Mak M., Mayne N., Sergeev D., Manners J., Eager-Nash J., Arney G., Hébrard E., Kohary K., 2023, *Journal of Geophysical Research: Atmospheres*, 128, e2023JD039343
- Malik M., et al., 2017, *The astronomical journal*, 153, 56
- Malik M., Kitzmann D., Mendonça J. M., Grimm S. L., Marleau G.-D., Linder E. F., Tsai S.-M., Heng K., 2019, *The astronomical journal*, 157, 170
- Mâlin M., et al., 2025, *Astronomy & Astrophysics*, 693, A315
- Mancini L., et al., 2018, *Astronomy & Astrophysics*, 613, A41
- Mandel K., Agol E., 2002, *The Astrophysical Journal*, 580, L171
- Mandushev G., et al., 2007, *The Astrophysical Journal*, 667, L195
- Mang J., Gao P., Hood C. E., Fortney J. J., Batalha N., Yu X., de Pater I., 2022, *The Astrophysical Journal*, 927, 184
- Mann C. R., et al., 2023, *The Astronomical Journal*, 166, 239
- Mao S., Paczynski B., 1991, *Astrophysical Journal, Part 2-Letters (ISSN 0004-637X)*, vol. 374, June 20, 1991, p. L37-L40., 374, L37
- Marois C., Macintosh B., Barman T., Zuckerman B., Song I., Patience J., Lafrenière D., Doyon R., 2008, *science*, 322, 1348
- Marquardt D. W., 1963, *Journal of the society for Industrial and Applied Mathematics*, 11, 431
- Martin T. Z., Cruikshank D. P., Pilcher C. B., Sinton W. M., 1976, *Icarus*, 27, 391
- Mayor M., Queloz D., 1995, *nature*, 378, 355
- McCullough P., Crouzet N., Deming D., Madhusudhan N., 2014, *The Astrophysical Journal*, 791, 55
- McGruder C. D., et al., 2022, *The Astronomical Journal*, 164, 134
- McGruder C. D., et al., 2023, *The Astronomical Journal*, 166, 120

- McLaughlin D. B., 1924, *ApJ*, 60, 22
- Meador W., Weaver W., 1980, *Journal of Atmospheric Sciences*, 37, 630
- Meech A., et al., 2025, *Monthly Notices of the Royal Astronomical Society*, p. 530
- Mie G., 1908, *Annalen der physik*, 330, 377
- Mikal-Evans T., et al., 2023, *The Astrophysical Journal Letters*, 943, L17
- Miles B. E., et al., 2023, *The Astrophysical journal letters*, 946, L6
- Min M., Hovenier J., de Koter A., 2005, *Astronomy & Astrophysics*, 432, 909
- Mireles I., et al., 2023, *The Astrophysical Journal Letters*, 954, L15
- Mollière P., van Boekel R., Dullemond C., Henning T., Mordasini C., 2015, *The Astrophysical Journal*, 813, 47
- Mollière P., van Boekel R., Bouwman J., Henning T., Lagage P.-O., Min M., 2017, *Astronomy & Astrophysics*, 600, A10
- Mollière P., Wardenier J., Van Boekel R., Henning T., Molaverdikhani K., Snellen I., 2019, *Astronomy & Astrophysics*, 627, A67
- Mollière P., et al., 2022, *The Astrophysical Journal*, 934, 74
- Moran S. E., et al., 2023, *The Astrophysical Journal Letters*, 948, L11
- Morbidelli A., et al., 2016, *Icarus*, 267, 368
- Morbidelli A., Batygin K., Lega E., 2023, *Astronomy & Astrophysics*, 675, A75
- Morley C. V., Fortney J. J., Marley M. S., Visscher C., Saumon D., Leggett S., 2012, *The Astrophysical Journal*, 756, 172
- Morley C. V., Fortney J. J., Marley M. S., Zahnle K., Line M., Kempton E., Lewis N., Cahoy K., 2015, *The Astrophysical Journal*, 815, 110
- Moses J. I., Poppe A. R., 2017, *Icarus*, 297, 33
- Moses J. I., et al., 2016, *The Astrophysical Journal*, 829, 66
- Muñoz D. J., Lai D., Liu B., 2016, *MNRAS*, 460, 1086
- Mukherjee S., Fortney J. J., Wogan N. F., Sing D. K., Ohno K., 2025, *The Astrophysical Journal*, 985, 209
- Mullens E., Lewis N. K., 2025, *The Astrophysical Journal Letters*, 988, L43
- Müller S., Helled R., 2024, *The Astrophysical Journal*, 967, 7
- Murgas F., Pallé E., Parviainen H., Chen G., Nortmann L., Nowak G., Cabrera-Lavers A., Iro N., 2017, *Astronomy & Astrophysics*, 605, A114

- Murphy M. M., et al., 2024, *Nature Astronomy*, 8, 1562
- Narita N., Sato B., Hirano T., Winn J. N., Aoki W., Tamura M., 2010, *Publications of the Astronomical Society of Japan*, 62, 653
- Nasedkin E., Mollière P., Blain D., 2024, *Journal of Open Source Software*, 9, 5875
- Nelder J. A., Mead R., 1965, *The Computer Journal*, 7, 308
- Neveu-VanMalle M., et al., 2014, *Astronomy & Astrophysics*, 572, A49
- Nikolov N., Sing D. K., Gibson N. P., Fortney J. J., Evans T. M., Barstow J. K., Kataria T., Wilson P. A., 2016, *The Astrophysical Journal*, 832, 191
- Nikolov N., et al., 2018, *Nature*, 557, 526
- Nortmann L., et al., 2018, *Science*, 362, 1388
- Öberg K. I., Murray-Clay R., Bergin E. A., 2011, *ApJ*, 743, L16
- Oelkers R. J., et al., 2025, *The Astronomical Journal*, 169, 134
- Ohno K., Fortney J. J., 2023a, *The Astrophysical Journal*, 946, 18
- Ohno K., Fortney J. J., 2023b, *The Astrophysical Journal*, 956, 125
- Ohno K., Kawashima Y., 2020, *The Astrophysical Journal Letters*, 895, L47
- Ohno K., Okuzumi S., 2017, *The Astrophysical Journal*, 835, 261
- Orell-Miquel J., et al., 2024, *Astronomy & Astrophysics*, 689, A179
- Oshagh M., Santos N., Ehrenreich D., Haghighipour N., Figueira P., Santerne A., Montalto M., 2014, *Astronomy & Astrophysics*, 568, A99
- Owen J. E., 2020, *MNRAS*, 495, 3160
- Owens A., Yurchenko S. N., Tennyson J., 2024, *Monthly Notices of the Royal Astronomical Society*, 530, 4004
- Palle E., et al., 2017, *Astronomy & Astrophysics*, 602, L15
- Panwar V., Désert J.-M., Todorov K. O., Bean J. L., Stevenson K. B., Huitson C., Fortney J. J., Bergmann M., 2022, *Monthly Notices of the Royal Astronomical Society*, 510, 3236
- Park C., et al., 2014, in *Ground-based and Airborne Instrumentation for Astronomy V*. pp 510–521
- Parker L. T., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 531, 2356

- Parviainen H., Aigrain S., 2015, *Monthly Notices of the Royal Astronomical Society*, 453, 3821
- Parviainen H., Pallé E., Nortmann L., Nowak G., Iro N., Murgas F., Aigrain S., 2016, *Astronomy & Astrophysics*, 585, A114
- Penzlin A. B. T., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 535, 171
- Perri F., Cameron A. G., 1974, *Icarus*, 22, 416
- Perryman M., 2018, *The exoplanet handbook*. Cambridge university press
- Perryman M., Hartman J., Bakos G. Á., Lindegren L., 2014, *The Astrophysical Journal*, 797, 14
- Petrovich C., 2015, *The Astrophysical Journal*, 799, 27
- Petrovich C., Tremaine S., 2016, *The Astrophysical Journal*, 829, 132
- Pinhas A., Madhusudhan N., 2017, *Monthly Notices of the Royal Astronomical Society*, 471, 4355
- Pinhas A., Rackham B. V., Madhusudhan N., Apai D., 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 5314
- Pinhas A., Madhusudhan N., Gandhi S., MacDonald R., 2019, *Monthly Notices of the Royal Astronomical Society*, 482, 1485
- Pino L., et al., 2018, *Astronomy & Astrophysics*, 612, A53
- Piskunov N., Kupka F., Ryabchikova T., Weiss W., Jeffery C., 1995, *Astronomy and Astrophysics Supplement*, v. 112, p. 525, 112, 525
- Piso A.-M. A., Pegues J., Öberg K. I., 2016, *The Astrophysical Journal*, 833, 203
- Pluriel W., et al., 2020, *The Astronomical Journal*, 160, 112
- Pollacco D. L., et al., 2006, *Publications of the Astronomical Society of the Pacific*, 118, 1407
- Pollack J. B., Hubickyj O., Bodenheimer P., Lissauer J. J., Podolak M., Greenzweig Y., 1996, *icarus*, 124, 62
- Polyansky O. L., Kyuberis A. A., Zobov N. F., Tennyson J., Yurchenko S. N., Lodi L., 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 2597
- Pont F., Knutson H., Gilliland R., Moutou C., Charbonneau D., 2008, *Monthly Notices of the Royal Astronomical Society*, 385, 109
- Pont F., Sing D. K., Gibson N. P., Aigrain S., Henry G., Husnoo N., 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 2917

- Powell D., Zhang X., 2024, *The Astrophysical Journal*, 969, 5
- Powell D., et al., 2024, *Nature*, 626, 979
- Prinn R. G., Barshay S. S., 1977, *Science*, 198, 1031
- Prinoth B., et al., 2022, *Nature Astronomy*, 6, 449
- Quanz S. P., et al., 2022, *Astronomy & Astrophysics*, 664, A21
- Queloz D., Eggenberger A., Mayor M., Perrier C., Beuzit J., Naef D., Sivan J., Udry S., 2000, *Astronomy and Astrophysics*, 359, L13
- Rackham B. V., de Wit J., 2024, *The Astronomical Journal*, 168, 82
- Rackham B., et al., 2017, *The Astrophysical Journal*, 834, 151
- Rackham B. V., Apai D., Giampapa M. S., 2018, *The Astrophysical Journal*, 853, 122
- Rackham B. V., Apai D., Giampapa M. S., 2019, *The Astronomical Journal*, 157, 96
- Radica M., 2024, *Astrophysics Source Code Library*, pp ascl-2412
- Radica M., et al., 2023, *Monthly Notices of the Royal Astronomical Society*, 524, 835
- Rafikov R. R., 2005, *ApJ*, 621, L69
- Rafikov R. R., 2006, *ApJ*, 648, 666
- Ramussen C., Williams C., 2006, *The MIT Press*, Boston
- Ranjan S., Charbonneau D., Désert J.-M., Madhusudhan N., Deming D., Wilkins A., Mandell A. M., 2014, *The Astrophysical Journal*, 785, 148
- Rasio F. A., Ford E. B., 1996, *Science*, 274, 954
- Rauer H., et al., 2014, *Experimental Astronomy*, 38, 249
- Rauscher E., Menou K., Seager S., Deming D., Cho J. Y.-K., Hansen B. M., 2007, *The Astrophysical Journal*, 664, 1199
- Rauscher E., Suri V., Cowan N. B., 2018, *The Astronomical Journal*, 156, 235
- Rayleigh, 3rd Lord (Strutt, J. W.) 1871, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 41, 107
- Raynard L., et al., 2018, *MNRAS*, 481, 4960
- Redfield S., Endl M., Cochran W. D., Koesterke L., 2008, *The Astrophysical Journal*, 673, L87

- Rescigno F., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 5385
- Rice M., et al., 2022, *The Astronomical Journal*, 164, 104
- Ricker G. R., et al., 2015, *Journal of Astronomical Telescopes, Instruments, and Systems*, 1, 014003
- Rieke G. H., et al., 2015, *Publications of the Astronomical Society of the Pacific*, 127, 584
- Rigby J., et al., 2023, *Publications of the Astronomical Society of the Pacific*, 135, 048001
- Rimmer P. B., Helling C., 2016, *The Astrophysical Journal Supplement Series*, 224, 9
- Rimmer P. B., Majumdar L., Priyadarshi A., Wright S., Yurchenko S., 2021, *The Astrophysical Journal Letters*, 921, L28
- Rocchetto M., Waldmann I., Venot O., Lagage P.-O., Tinetti G., 2016, *The Astrophysical Journal*, 833, 120
- Romani P. N., Atreya S. K., 1988, *Icarus*, 74, 424
- Rosenthal L. J., et al., 2021, *The Astrophysical Journal Supplement Series*, 255, 8
- Rossiter R. A., 1924, *ApJ*, 60, 15
- Rothman L. S., et al., 2010, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 111, 2139
- Rustamkulov Z., et al., 2023, *Nature*, 614, 659
- Sahlmann J., Lazorenko P., Ségransan D., Martín E., Queloz D., Mayor M., Udry S., 2013, *Astronomy & Astrophysics*, 556, A133
- Sánchez-Lavega A., Pérez-Hoyos S., Hueso R., 2004, *American Journal of Physics*, 72, 767
- Sato B., et al., 2005, *The Astrophysical Journal*, 633, 465
- Saumon D., Marley M. S., 2008, *The Astrophysical Journal*, 689, 1327
- Schneider A. D., Bitsch B., 2021a, *A&A*, 654, A71
- Schneider A. D., Bitsch B., 2021b, *Astronomy & Astrophysics*, 654, A72
- Schwarz G., 1978, *The annals of statistics*, pp 461–464
- Schwarzschild M., 1958, *Structure and Evolution of Stars*. Princeton University Press

- Seager S., 2010, *Exoplanet Atmospheres: Physical Processes*. Princeton University Press
- Seager S., Sasselov D. D., 2000, *The Astrophysical Journal*, 537, 916
- Sedaghati E., et al., 2016, *Astronomy & Astrophysics*, 596, A47
- Sedaghati E., et al., 2017, *Nature*, 549, 238
- Seidel J., Ehrenreich D., Pino L., Bourrier V., Lavie B., Allart R., Wyttenbach A., Lovis C., 2020, *Astronomy & Astrophysics*, 633, A86
- Seidel J., et al., 2023, *Astronomy & Astrophysics*, 673, A125
- Shorttle O., Jordan S., Nicholls H., Lichtenberg T., Bower D. J., 2024, *The Astrophysical Journal Letters*, 962, L8
- Showman A. P., Guillot T., 2002, *Astronomy & Astrophysics*, 385, 166
- Showman A. P., Cho J. Y., Menou K., 2010, *Exoplanets*, 526, 471
- Sing D. K., et al., 2011a, *Monthly Notices of the Royal Astronomical Society*, 416, 1443
- Sing D. K., et al., 2011b, *Astronomy & Astrophysics*, 527, A73
- Sing D. K., et al., 2012, *Monthly Notices of the Royal Astronomical Society*, 426, 1663
- Sing D. K., et al., 2015, *Monthly Notices of the Royal Astronomical Society*, 446, 2428
- Sing D. K., et al., 2016, *Nature*, 529, 59
- Sing D. K., et al., 2024, *Nature*, 630, 831
- Skilling J., 2004, *Bayesian inference and maximum entropy methods in science and engineering*, 735, 395
- Snellen I. A., 2025, *Annual Review of Astronomy and Astrophysics*, 63
- Snellen I., Albrecht S., De Mooij E., Le Poole R., 2008, *Astronomy & Astrophysics*, 487, 357
- Snellen I. A., De Kok R. J., De Mooij E. J., Albrecht S., 2010, *Nature*, 465, 1049
- Snellen I. A., Brandl B. R., De Kok R. J., Brogi M., Birkby J., Schwarz H., 2014, *Nature*, 509, 63
- Southworth J., Wheatley P. J., Sams G., 2007, *Monthly Notices of the Royal Astronomical Society: Letters*, 379, L11

Sozzetti A., et al., 2015, *Astronomy and Astrophysics*, 575, L15

Spake J. J., et al., 2018, *Nature*, 557, 68

Spalding C., Winn J. N., 2022, *The Astrophysical Journal*, 927, 22

Spyratos P., et al., 2021, *Monthly Notices of the Royal Astronomical Society*, 506, 2853

Spyratos P., Nikolov N. K., Constantinou S., Southworth J., Madhusudhan N., Sedaghati E., Ehrenreich D., Mancini L., 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 2163

Stefánsson G., et al., 2025, *The Astronomical Journal*, 169, 107

Stevenson K. B., 2016, *The Astrophysical Journal Letters*, 817, L16

Stevenson K. B., 2020, *The Astrophysical Journal Letters*, 898, L35

Stevenson K. B., et al., 2025, *The Astronomical Journal*, 170, 257

Stock J. W., Kitzmann D., Patzer A. B. C., Sedlmayr E., 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 865

Stock J. W., Kitzmann D., Patzer A. B. C., 2022, *MNRAS*, 517, 4070

Struve O., 1952, *The Observatory*, Vol. 72, p. 199-200 (1952), 72, 199

Su X.-N., Zhang H., Zhou J.-L., 2024, *Monthly Notices of the Royal Astronomical Society*, 529, 3958

Sun Q., Wang S. X., Welbanks L., Teske J., Buchner J., 2024, *The Astronomical Journal*, 167, 167

Tabernero H., et al., 2022, *Monthly Notices of the Royal Astronomical Society*, 515, 1247

Taylor J., Parmentier V., Line M. R., Lee E. K., Irwin P. G., Aigrain S., 2021, *Monthly Notices of the Royal Astronomical Society*, 506, 1309

Taylor J., et al., 2023, *Monthly Notices of the Royal Astronomical Society*, 524, 817

Teske J., et al., 2025, *The Astronomical Journal*, 169, 249

Thorngren D. P., Fortney J. J., 2018, *AJ*, 155, 214

Thorngren D., Fortney J. J., 2019, *ApJ*, 874, L31

Thorngren D. P., Fortney J. J., Murray-Clay R. A., Lopez E. D., 2016, *ApJ*, 831, 64

Thorngren D., Gao P., Fortney J. J., 2019, *The Astrophysical Journal Letters*, 884, L6

- Thorngren D. P., Sing D. K., Mukherjee S., 2025, arXiv preprint arXiv:2510.00169
- Tinetti G., et al., 2007, *Nature*, 448, 169
- Tinetti G., et al., 2018, *Experimental astronomy*, 46, 135
- Toomre A., 1964, *Astrophysical Journal*, vol. 139, p. 1217-1238 (1964)., 139, 1217
- Toon O. B., Turco R., Hamill P., Kiang C., Whitten R., 1979, *J. atmos. Sci*, 36, 718
- Triaud A. H. M. J., et al., 2010, *A&A*, 524, A25
- Trifonov T., et al., 2023, *The Astronomical Journal*, 165, 179
- Trotta R., 2008, *Contemporary Physics*, 49, 71
- Tsai S.-M., Lyons J. R., Grosheintz L., Rimmer P. B., Kitzmann D., Heng K., 2017, *The Astrophysical Journal Supplement Series*, 228, 20
- Tsai S.-M., Malik M., Kitzmann D., Lyons J. R., Fateev A., Lee E., Heng K., 2021, *ApJ*, 923, 264
- Tsai S.-M., et al., 2023, *Nature*, 617, 483
- Tsiaras A., Waldmann I. P., Tinetti G., Tennyson J., Yurchenko S. N., 2019, *Nature Astronomy*, 3, 1086
- Turco R., Hamill P., Toon O., Whitten R., Kiang C., 1979, *Journal of Atmospheric Sciences*, 36, 699
- Turrini D., et al., 2021, *The Astrophysical Journal*, 909, 40
- Ulmer-Moll S., et al., 2025, *Astronomy & Astrophysics*, 703, A258
- Ulrich Käufel H., 2008, in *Precision Spectroscopy in Astrophysics: Proceedings of the ESO/Lisbon/Aveiro Conference held in Aveiro, Portugal, 11–15 September 2006*. pp 227–230
- Underwood D. S., Tennyson J., Yurchenko S. N., Huang X., Schwenke D. W., Lee T. J., Clausen S., Fateev A., 2016, *Monthly Notices of the Royal Astronomical Society*, 459, 3890
- Valentine D., et al., 2024, *The Astronomical Journal*, 168, 123
- Valentine D., Wakeford H. R., Hammond M., Challener R. C., Edwards B., Lüftinger T., Günther M. N., 2025, *Monthly Notices of the Royal Astronomical Society*, 544, 3647
- Veillet R., et al., 2025, arXiv preprint arXiv:2505.12152
- Vissapragada S., et al., 2020, *The Astronomical Journal*, 159, 278

- Vissapragada S., et al., 2022, *The Astronomical Journal*, 164, 234
- Visscher C., Moses J. I., 2011, *The Astrophysical Journal*, 738, 72
- Visscher C., Lodders K., Fegley Jr B., 2006, *The Astrophysical Journal*, 648, 1181
- Visscher C., Lodders K., Fegley Jr B., 2010, *The Astrophysical Journal*, 716, 1060
- Voyer M., et al., 2025, *The Astrophysical Journal Letters*, 982, L38
- Wakeford H. R., Sing D. K., 2015, *Astronomy & Astrophysics*, 573, A122
- Wakeford H. R., Visscher C., Lewis N. K., Kataria T., Marley M. S., Fortney J. J., Mandell A. M., 2016, *Monthly Notices of the Royal Astronomical Society*, p. stw2639
- Wakeford H. R., et al., 2017a, *The Astronomical Journal*, 155, 29
- Wakeford H. R., et al., 2017b, *Science*, 356, 628
- Wakeford H. R., et al., 2019, *The Astronomical Journal*, 157, 11
- Wallack N. L., et al., 2024, *The Astronomical Journal*, 168, 77
- Wang J., Fischer D. A., 2015, *The Astronomical Journal*, 149, 14
- Wang J., et al., 2013, *The Astrophysical Journal*, 776, 10
- Wang J., Fischer D. A., Horch E. P., Huang X., 2015, *The Astrophysical Journal*, 799, 229
- Wang X.-Y., et al., 2024, *The Astrophysical Journal Letters*, 973, L21
- Welbanks L., Madhusudhan N., 2021, *The Astrophysical Journal*, 913, 114
- Welbanks L., Madhusudhan N., 2022, *ApJ*, 933, 79
- Welbanks L., Madhusudhan N., Allard N. F., Hubeny I., Spiegelman F., Leininger T., 2019, *The Astrophysical Journal Letters*, 887, L20
- Werner M. W., et al., 2004, *The Astrophysical Journal Supplement Series*, 154, 1
- West R. G., et al., 2009, *AJ*, 137, 4834
- Wheatley P. J., et al., 2018, *Monthly Notices of the Royal Astronomical Society*, 475, 4476
- Winn J. N., 2010, *Exoplanet transits and occultations*. University of Arizona Press Tucson, AZ
- Winn J. N., Fabrycky D., Albrecht S., Johnson J. A., 2010, *The Astrophysical Journal*, 718, L145

- Witzke V., Shapiro A. I., Kostogryz N. M., Cameron R., Rackham B. V., Seager S., Solanki S. K., Unruh Y. C., 2022, *ApJ*, 941, L35
- Wogan N. F., Batalha N. E., Zahnle K. J., Krissansen-Totton J., Tsai S.-M., Hu R., 2024, *The Astrophysical Journal Letters*, 963, L7
- Wolszczan A., Frail D. A., 1992, *Nature*, 355, 145
- Wong M. H., Mahaffy P. R., Atreya S. K., Niemann H. B., Owen T. C., 2004, *Icarus*, 171, 153
- Wu Y., Lithwick Y., 2011, *The Astrophysical Journal*, 735, 109
- Wu Y., Murray N., 2003, *ApJ*, 589, 605
- Wytttenbach A., Ehrenreich D., Lovis C., Udry S., Pepe F., 2015, *Astronomy & Astrophysics*, 577, A62
- Xue Q., Bean J. L., Zhang M., Welbanks L., Lunine J., August P., 2024, *The Astrophysical Journal Letters*, 963, L5
- Yang J., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 532, 460
- Youngblood A., et al., 2016, *The Astrophysical Journal*, 824, 101
- Yu X., Glein C. R., Thorngren D. P., Murray D. F., 2026, *arXiv preprint arXiv:2604.01672*
- Yurchenko S. N., Tennyson J., 2014, *MNRAS*, 440, 1649
- Yurchenko S. N., Amundsen D. S., Tennyson J., Waldmann I. P., 2017, *Astronomy & Astrophysics*, 605, A95
- Yurchenko S. N., Bond W., Gorman M. N., Lodi L., McKemmish L. K., Nunn W., Shah R., Tennyson J., 2018, *Monthly Notices of the Royal Astronomical Society*, 478, 270
- Yurchenko S., Mellor T. M., Freedman R. S., Tennyson J., 2020, *Monthly Notices of the Royal Astronomical Society*, 496, 5282
- Yurchenko S. N., Owens A., Kefala K., Tennyson J., 2024, *Monthly Notices of the Royal Astronomical Society*, 528, 3719
- Zahnle K. J., Catling D. C., 2017, *The Astrophysical Journal*, 843, 122
- Zahnle K. J., Marley M. S., 2014, *The Astrophysical Journal*, 797, 41
- Zahnle K., Marley M. S., Freedman R. S., Lodders K., Fortney J., 2009, *The Astrophysical Journal*, 701, L20
- Zahnle K., Marley M. S., Morley C. V., Moses J. I., 2016, *The Astrophysical Journal*, 824, 137

- Zamyatina M., et al., 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 3129
- Zamyatina M., et al., 2024, *Monthly Notices of the Royal Astronomical Society*, 529, 1776
- Zhang M., Chachan Y., Kempton E. M. R., Knutson H. A., 2019, *PASP*, 131, 034501
- Zhang Y., et al., 2021a, *Nature*, 595, 370
- Zhang Z., Liu M. C., Claytor Z. R., Best W. M., Dupuy T. J., Siverd R. J., 2021b, *The Astrophysical Journal Letters*, 916, L11
- Zhou G., et al., 2017, *The Astronomical Journal*, 153, 211
- Zieba S., et al., 2023, *Nature*, 620, 746
- de Kok R. J., Brogi M., Snellen I. A., Birkby J., Albrecht S., de Mooij E. J., 2013, *Astronomy & Astrophysics*, 554, A82
- de Wit J., Gillon M., Demory B.-O., Seager S., 2012, *Astronomy & Astrophysics*, 548, A128