

Library Declaration and Deposit Agreement

1. STUDENT DETAILS

Thomas Michael Loudon
12583863

2. THESIS DEPOSIT

- 2.1 I understand that under my registration at the University, I am required to deposit my thesis with the University in BOTH hard copy and in digital format. The digital version should normally be saved as a single pdf file.
- 2.2 The hard copy will be housed in the University Library. The digital version will be deposited in the University's Institutional Repository (WRAP). Unless otherwise indicated (see 2.3 below) this will be made openly accessible on the Internet and will be supplied to the British Library to be made available online via its Electronic Theses Online Service (EThOS) service. [At present, theses submitted for a Masters degree by Research (MA, MSc, LLM, MS or MMedSci) are not being deposited in WRAP and not being made available via EThOS. This may change in future.]
- 2.3 In exceptional circumstances, the Chair of the Board of Graduate Studies may grant permission for an embargo to be placed on public access to the hard copy thesis for a limited period. It is also possible to apply separately for an embargo on the digital version. (Further information is available in the Guide to Examinations for Higher Degrees by Research.)
- 2.4 (a) Hard Copy I hereby deposit a hard copy of my thesis in the University Library to be made publicly available to readers immediately.
I agree that my thesis may be photocopied.
- (b) Digital Copy I hereby deposit a digital copy of my thesis to be held in WRAP and made available via EThOS.
My thesis can be made publicly available online.

3. GRANTING OF NON-EXCLUSIVE RIGHTS

Whether I deposit my Work personally or through an assistant or other agent, I agree to the following: Rights granted to the University of Warwick and the British Library and the user of the thesis through this agreement are non-exclusive. I retain all rights in the thesis in its present version or future versions. I agree that the institutional repository administrators and the British Library or their agents may, without changing content, digitise and migrate the thesis to any medium or format for the purpose of future preservation and accessibility.

4. DECLARATIONS

(a) I DECLARE THAT:

- I am the author and owner of the copyright in the thesis and/or I have the authority of the authors and owners of the copyright in the thesis to make this agreement. Reproduction of any part of this thesis for teaching or in academic or other forms of publication is subject to the normal limitations on the use of copyrighted materials and to the proper and full acknowledgement of its source.
- The digital version of the thesis I am supplying is the same version as the final, hardbound copy submitted in completion of my degree, once any minor corrections have been completed.
- I have exercised reasonable care to ensure that the thesis is original, and does not to the best of my knowledge break any UK law or other Intellectual Property Right, or contain any confidential material.
- I understand that, through the medium of the Internet, files will be available to automated agents, and may be searched and copied by, for example, text mining and plagiarism detection software.

(b) IF I HAVE AGREED (in Section 2 above) TO MAKE MY THESIS PUBLICLY AVAILABLE DIGITALLY, I ALSO DECLARE THAT:


- I grant the University of Warwick and the British Library a licence to make available on the Internet the thesis in digitised format through the Institutional Repository and through the British Library via the EThOS service.
- If my thesis does include any substantial subsidiary material owned by third-party copyright holders, I have sought and obtained permission to include it in any version of my thesis available in digital format and that this permission encompasses the rights that I have granted to the University of Warwick and to the British Library.

5. LEGAL INFRINGEMENTS

I understand that neither the University of Warwick nor the British Library have any obligation to take legal action on behalf of myself, or other rights holders, in the event of infringement of intellectual property rights, breach of contract or of any other right, in the thesis.

Please sign this agreement and return it to the Graduate School Office when you submit your thesis.

Student's signature:Date:



**The composition and dynamics of exoplanet
atmospheres**

by

Thomas Michael Louden

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

June 2016

THE UNIVERSITY OF
WARWICK

Contents

List of Tables	v
List of Figures	vi
Acknowledgments	viii
Declarations	ix
Abstract	x
Abbreviations	xi
Chapter 1 Introduction	1
1.1 History of exoplanet detection	1
1.1.1 The First exoplanet	4
1.1.2 Planetary archeology	6
1.1.3 The planet detection gold-rush	7
1.2 Detection methods	10
1.2.1 Transits	10
1.2.2 Radial velocity	18
1.2.3 Microlensing	20
1.2.4 Direct imaging	21
1.2.5 Rossiter-McLaughlin effect	22
1.3 Observed population of exoplanets	23
1.3.1 Mass-density relations	25
1.3.2 The Radius anomaly	28
1.3.3 The occurrence of hot Jupiters	30
1.3.4 Transmission signal	31
1.3.5 Availability of comparison stars	35
1.4 Observable properties of Exoplanet Atmospheres	37

1.4.1	Transmission spectroscopy	37
1.4.2	Atmosphere erosion	42
1.4.3	Thermal emission and atmosphere dynamics	47
1.5	Conclusions	51
Chapter 2 Methods and Techniques		52
2.1	Introduction	52
2.2	CCD	52
2.2.1	Reduction	55
2.2.2	X-ray astronomy	57
2.3	Spectrograph	58
2.4	ACAM on the William Herschel Telescope	59
2.5	HARPS	68
2.6	COS on HST	69
2.7	XMM-Newton	72
2.8	Markov Chain Monte Carlo	72
2.8.1	Bayes Theorem	73
2.8.2	MCMC operation	76
2.8.3	Bayesian Information Criterion	77
2.9	Gaussian processes	77
2.10	Conclusions	82
Chapter 3 The hunt for hot Neptunes with NGTS		83
3.1	Introduction	83
3.2	Planet catch simulations	84
3.3	Zero level Pipeline design	89
3.3.1	Reduction	92
3.4	Astrometry	93
3.4.1	Aperture photometry	97
3.4.2	Frame rejection	98
3.4.3	Detrending	98
3.4.4	Light curves	100
3.4.5	Quality Control plots for the Current pipeline	103
3.5	Conclusions	105
Chapter 4 Reconstructing the high energy irradiation of the evaporating hot Jupiter HD 209458b		109
4.1	Introduction	110

4.2	Observations	112
4.3	Method	114
4.3.1	Differential Emission Measure	114
4.3.2	Fitting procedure	116
4.4	Results	116
4.5	Discussion	118
4.5.1	Comparison to previous work	118
4.5.2	Energy-limited mass-loss	121
4.6	Conclusions	123
Chapter 5	Rotation and Eastward winds on HD 189733b	124
5.1	Introduction	124
5.2	Observations	125
5.3	Transit model	127
5.4	Results	130
5.5	Discussion and conclusions	133
Chapter 6	Probing the atmosphere of WASP-52b	135
6.1	Introduction	136
6.2	Observations	137
6.3	Method	139
6.3.1	Extraction	139
6.3.2	Fitting procedure	140
6.3.3	Gaussian process and white noise model	142
6.3.4	Application of the model	144
6.3.5	Spot crossing event	145
6.4	Results	147
6.4.1	Transmission spectra	147
6.4.2	Updated system parameters	148
6.5	Discussion	148
6.5.1	Transit timing	148
6.5.2	Model atmospheres	152
6.5.3	An unknown absorber beyond 7750 Å?	157
6.6	Conclusions	160
Chapter 7	Conclusions and Future prospects	165
7.1	The hunt for hot Neptunes with NGTS	165

7.2	Reconstructing the high energy irradiation of the evaporating hot Jupiter HD 209458b	167
7.3	Rotation and Eastward winds on HD 189733b	168
7.4	Probing the atmosphere of WASP-52b	169
7.5	Concluding remarks	170

List of Tables

2.1	The order blocking filters for ACAM	66
4.1	UV line strengths for HD 209458	113
4.2	differential emission measure for HD 209458b	117
4.3	Recovered high energy spectrum fluxes	119
5.1	Stellar model used for HD 189733b	129
6.1	Derived system parameters for WASP-52b.	146
6.2	Model atmosphere goodness of fits for WASP-52b	155
6.3	Transmission spectra for WASP-52b	159

List of Figures

1.1	Planets Detected to date	3
1.2	NGTS fully installed	9
1.3	The observed population of exoplanets	13
1.4	A schematic representation of a transit	14
1.5	A high precision transit of HAT-P-7b from <i>Kepler</i>	16
1.6	Limb darkening profiles	17
1.7	The detection of 51 peg b	19
1.8	Diagram of the Rossiter-McLaughlin effect	22
1.9	Exoplanet mass-density relation	24
1.10	The radius inflation effect	32
1.11	Exoplanet atmosphere signal	34
1.12	Probability of finding a comparison star	38
1.13	Transmission spectrum of HD 189733b	43
1.14	Maxwell Boltzmann distribution	45
1.15	Signal to noise for eclipse observations	49
1.16	hotspot offset in HD 189733b	50
2.1	Uptake of CCDs in astronomy	54
2.2	A schematic representation of CCD	55
2.3	A raw HARPS spectrum	60
2.4	Weather at La Palma	61
2.5	The instrument design of ACAM	62
2.6	ACAM raw spectrum of WASP-52b and comparison star	65
2.7	ACAM instrument response	67
2.8	The HARPS instrument	69
2.9	Systematics in HST light-curve	70
2.10	A spectrum taken with HST COS	71
2.11	Effective area of XMM-Newton cameras	73

2.12	A uniform prior	75
2.13	A white noise kernel, squared exponential kernel and Matern kernel.	80
2.14	Three periodic kernels	81
3.1	The Follow up potential of NGTS simulated targets	90
3.2	The I-band follow up potential of NGTS simulated targets	91
3.3	An NGTS unit at Geneva	92
3.4	NGTS raw frame	94
3.5	A solved astrometry frame	95
3.6	An NGTS flux-RMS diagram from the Geneva run	99
3.7	Eclipsing binaries recorded with NGTS	101
3.8	A 1 mmag delta scuti recorded with NGTS	102
3.9	TrES-2b detected in WASP vs NGTS	104
3.10	RMS binning plot magnitudes	106
3.11	RMS binning plot magnitudes	107
4.1	Radiative loss function for coronal plasma	112
4.2	Best fitting DEM for HD 209458	115
4.3	Recovered high energy spectrum for HD 209458	119
5.1	Diagram of star-planet geometry for HD 189733b	128
5.2	Planetary absorption line profiles for HD 189733b	131
5.3	Bootstrap posteriors of atmospheric velocities	132
5.4	The effect of the Rossiter-McLaughlin effect on absorption profiles	132
6.1	Spectral trace for target and comparison star	138
6.2	WASP-52 periodogram of stellar activity	141
6.3	White light models for WASP-52b transit	149
6.4	Comparison of WASP-52b transmission spectra for two nights	150
6.5	Final combined transmission spectrum for WASP-52b	151
6.6	Model Temperature Pressure Profiles	151
6.7	Transit timing variations in WASP-52	156
6.8	Individual spectral lightcurves for night 1 on WASP-52b	161
6.9	Residuals for night 1 on WASP-52b	162
6.10	Individual spectral lightcurves for night 2 on WASP-52b	163
6.11	Residuals for night 2 on WASP-52b	164

Acknowledgments

First of all, this thesis would not have been possible without the guidance and encouragement of my supervisor, Peter Wheatley, so to him go my sincerest thanks.

To my family, I could not have done this without your love and support. My (perhaps too infrequent) visits home to you, and escapes to cottages and camping in France have been invaluable to reducing my levels of stress!

To all the friends that I have shared offices, pub trips, coffee and board games with throughout the years, thank you. Thank you for the endless laughter, for the companionship, for the stupid jokes, and above all for tolerating me and my antics!

The opportunity to study such an exciting subject has been worth all the sleepless nights and looming deadlines. Carl Sagan once described astronomy as a “humbling and character-building experience”, and I could not agree more. These 4 years have taught me a lot about myself, but above all that there is always more to learn and to discover.

This PhD has been one long adventure, I’m thankful for every minute of it.

Declarations

I declare that the work presented in this thesis is my own except where stated otherwise, and was carried out entirely at the University of Warwick, during the period October 2012 to June 2016, under the supervision of Dr. Peter Wheatley. The research reported here has not been submitted, either wholly or in part, in this or any other academic institution for admission to a higher degree.

The results of the planet catch simulations in Chapter 3 are published in *The conference Proceedings of “Hot Planets and Cool Stars”* entitled ‘The Next Generation Transit Survey (NGTS)’ (Wheatley et al., 2013).

The work in Chapter 4 is based on a paper submitted to the *Monthly Notices of the Royal Astronomical Society (MNRAS)* entitled ‘Reconstructing the high energy irradiation of the evaporating hot Jupiter HD 209458b’ (Louden et al., 2016).

The work in Chapter 5 is based on a paper published in *Astrophysical Journal Letters (ApJL)* entitled ‘Spatially resolved eastward winds and rotation of HD 189733b’ (Louden et al., 2016).

Other contributions based on this thesis are:

- Exoplanet UK community Meeting, Cambridge, UK, April 2014, Oral Presentation: *What drives the evaporation of HD209458b?*
- Exoplanet UK community Meeting, Warwick, UK, March 2015, Oral Presentation: *High resolution transmission spectroscopy of HD 189733b.*
- Institute of Astronomy, Cambridge, UK, May 2016, Invited talk: *A forecast for high winds and clouds: Exoplanet atmospheres at low and high resolution.*

Abstract

The study of exoplanets has rapidly developed in the last twenty years, and the detailed characterization of planetary atmospheres has become a key area of research.

For transiting planets around bright stars, atmospheric features can be detected with transmission spectroscopy. I will present a low resolution transmission spectrum of WASP-52b, and show that the most likely interpretation is that the planet is shrouded with an opaque cloud layer.

By using transmission spectroscopy at much higher resolution, in this thesis I will present the first *spatially resolved* measurements of a weather system in an exoplanet. By modeling the absorption profile of sodium on HD 189733b, I show that the planet atmosphere has an excess velocity not explained by planetary rotation.

HD 209458b is evaporating under intense irradiation from its star, and may lose as much as 10^{10} g s⁻¹. Mechanisms of mass loss are poorly understood, in particular the efficiency. To calculate this rate for HD 209458b a key component is missing - the high energy flux of the star. I will demonstrate that it is possible to recover this flux by building a coronal model for the star, using constraints for different temperatures of plasma from UV and X-ray sources.

I will present commissioning data gathered with NGTS that are the most precise ever gathered with a ground based wide field transit survey. Simulations of the performance of NGTS in this thesis show that the survey can be expected to detect ~200 low mass planets.

The simulations of NGTS also show that a sample of bright super-Earths and hot-Neptunes can be expected to be detected, which would be sensitive to the same techniques performed on hot Jupiters in this thesis. One day, these same techniques could be important tools for characterizing the atmospheres of Earth analogs.

Abbreviations

TTV Transit Timing Variation	ISM Interstellar Medium
RV Radial velocity	EUV Extreme Ultraviolet
WASP Wide Angle Search for Planets	XUV X-ray and UV
HST Hubble Space Telescope	HARPS High Accuracy Radial velocity Planet Searcher
NGTS Next Generation Transit Survey	
COS Cosmic Origins Spectrograph	STIS Space Telescope Imaging Spectrograph
WHT William Hershel Telescope	
ACAM Auxiliary-port camera	M_J Jupiter Mass, 1.898×10^{27} kg
CCD Charge Couple Device	R_J Jupiter Radius, 71492 km
PSF Point Spread Function	M_\odot Solar Mass, 1.989×10^{30} kg
RM RossiterMcLaughlin	R_\odot Solar Radius, 7.955×10^5 km
MCMC Markov Chain Monte Carlo	M_\oplus Earth Mass, 5.972×10^{24} kg
GP Gaussian Process	R_\oplus Earth Radius, 6371 km
DEM Differential emission measure	AU Astronomical Unit, 1.496×10^8 km

Chapter 1

Introduction

For then why may not every one of these Stars or Suns have as great a Retinue as our Sun, of planets, with their moons, to wait upon them? Nay, there's Manifest reason why they should.

Christiaan Huygens, 1698

1.1 History of exoplanet detection

We are currently living in what could be considered a revolutionary age of discovery. The last 20 years have been a veritable gold-rush, uncovering a dazzling variety of worlds that even writers of science fiction would have struggled to anticipate.

Humans have imagined the presence of other worlds for millenia, initially almost as an exercise in humility - what claim do we have to be unique, after all? In 300 BCE the Greek philosopher Epicurus wrote:

“For the atoms out of which a world might arise, or by which a world might be formed, have not all been expended on one world or a finite number of worlds, whether like or unlike this one. Hence there will be nothing to hinder an infinity of worlds.”

This touches on a question that would become a key point of contention, and we are only just now on the brink of being able to answer: How do planets form? Is it a common process, or a rare one? One could imagine being satisfied that our

Solar System formed through natural means, but finding that the process through which it formed was so staggeringly unlikely that only a vanishingly small fraction of stars harbored a planetary system.

During the 19th and early 20th century, various “catastrophic” theories of planetary formation were popular. The common theme was that the material of the planets in the Solar System must have come from the Sun, and so a significant, violent one-off event had occurred which stripped a fraction of the Sun’s material and allowed it to condense into the planets of the Solar System. At the time these theories seemed necessary, due to the angular velocity distribution of the Solar System, where the planets, with 0.13% of the total mass hold 99.5% of the angular momentum (Williams and Cremin, 1968). This appeared to invalidate any theory of Solar System formation that had the Sun and the Planets forming concurrently from a single primordial nebula, a theory originally proposed by Swedenborg (1734). As far back as the 18th century, Buffon (1745) suggested that the planets were formed by material ejected after a comet struck the Sun, the actual nature of comets not being known at the time. For a time, scientists spoke of collisions of stars (Bickerton, 1878; Arrhenius, 1913), until JEANS (1916) demonstrated that material could be drawn off merely by a close encounter. These “Tidal” theories of planetary formation were popular despite some criticism (Spitzer, Lyman, 1939) until a solution to the Angular Momentum Problem could be found.¹ What they have in common is that the universe suddenly becomes a much more lonely place, far from “An infinity of Worlds”, planetary systems would necessarily be immensely rare, with each star having only a 1 in 10 billion chance of hosting a planetary system (Winn and Fabrycky, 2014). Instead, the consensus of modern astronomy is that planet formation is a natural and common side effect of stellar formation (e.g. Williams and Cieza, 2011). The vast numbers of planets (approaching 2000 at the time of writing) detected over the last two decades gives conservative estimates of at least 50% of Sun-like stars hosting exoplanets, from transits and radial velocity measurements (Fressin et al., 2014; Mayor et al., 2011), with perhaps ~40% of M-dwarf stars having a planet in the habitable zone (Bonfils et al., 2013). Planets have been found to exist almost everywhere we look, across all the whole Hertzsprung-Russell diagram, and around both single and multiple stars. One only has to look at Figure 1.1 to see how rapidly things have progressed, and with no sign of slowing down.

¹Today this “Angular Momentum Problem” is considered to be solved by magnetic braking in the Solar wind (Schatzman, 1962; Weber and Davis, Leverett, 1967)




Figure 1.1: Exoplanets detected to date, colour coded by detection method. The “Timing” category is an amalgamation of several different related methods, but is mostly comprised of planets detected through Transit Timing Variations. The Pulsar planets detected in 1992 are also included in this category. Almost all of the “Space based Transit” category are from a single data release from the *Kepler* Program in 2015. These planets are in fact not formally confirmed, but are validated to greater than 3σ certainty through arguments of astrophysical blending probability (Source: NASA exoplanet Archive, <http://exoplanetarchive.ipac.caltech.edu/> accessed 17/06/16).

1.1.1 The First exoplanet

It is somewhat contentious to attempt to put a date on the discovery of the first exoplanet. The commonly accepted date is 1995, but there are several detections in the preceding years that could also have a fair claim to the title.

Campbell et al. (1988) produced the first tentative detection of an exoplanet around gamma Cephei A using the radial velocity technique, but the low significance prevented them from making a strong claim. This planet was subsequently retroactively confirmed as gamma Cep Ab (Hatzes et al., 2003), with a period of 906 days and a minimum mass of 1.7 M_J . However, since Campbell et al. (1988) found similar variation in 7/15 target stars, and 6 have not been confirmed with modern instruments, it is possible that they simply got lucky with gamma Cep A, so this probably can not be considered a true “first detection”.

Latham et al. (1989) provided the first compelling evidence for a sub-stellar companion to a star, again through radial velocity measurements. The companion they detected around HD 114762 had a minimum mass of 11 M_J and a period of 84 days. While not as extreme as the hot Jupiters that would later be discovered, it was early evidence for the migration of planet mass companions, as it is difficult for a planet that massive to form interior to the snow line in conventional planet formation models. HD 114762b is the object with the earliest detection date to be included in the NASA exoplanet archive². The high minimum mass raised an objection to christening this object the first exoplanet, as even moderate inclinations would push it into a mass range that is typically considered to be the realm of “brown dwarfs”, objects intermediate between planets and stars.

This raises the question of what exactly constitutes a planet, reminiscent of the debate over the planethood of Pluto, which has become a part of popular culture. Currently, a popular definition is that a planet must be below the “deuterium burning limit”. At a certain mass, a planet’s interior becomes hot and high pressure enough to fuse deuterium. The problem, however, is this mass is very model dependent, and can in fact vary between 11 and 16 M_J (Spiegel et al., 2011). Further, deuterium burning can start and stop at different periods during the lifetime of the object - does an object transition multiple times between planet and not-planet? The primary issue, given these uncertainties, is there is not a simple observational test one can make of a potential planet to tell if it is currently fusing, or has ever fused deuterium, making its use as a classification tool dubious.

Another suggestion was to distinguish between planets and brown dwarfs

²<http://exoplanetarchive.ipac.caltech.edu/index.html> accessed 10/06/16

based on their formation history. Whilst this is an appealing solution, it is complicated by the orbital migration of planets. In individual cases it may not be possible to determine exactly where the planet formed in the proto-stellar nebular, and is certainly possible that multiple formation mechanisms could give rise to otherwise indistinguishable giant planets. What then? Perhaps it is best to accept that there are no clear categories in nature, the $13 M_J$ limit is perhaps a useful guideline, but it should not be enforced too strictly. To this end, I include “brown dwarfs” in my discussions of planet detection and occurrence.

In 1992, a new strong detection of a planetary mass system came from an unexpected source - exquisitely small variations in the frequency of Pulsars (Wolszczan and Frail, 1992). This was a particularly fortuitous occurrence, as in subsequent years only one other pulsar system with a planet has been found (Richer et al., 2003). The history of these objects are particularly interesting - did they survive the supernova that led to the destruction of their parent star, or are they second generation planets formed from the debris of the previous solar system’s destruction? Regardless, these were the first unambiguously planetary mass objects detected outside the Solar System, so this landmark discovery is often considered the beginning of the exoplanet era. Unfortunately, the nature of this system means there is very little prospect for uncovering more information about these exotic objects than their mass.

Finally, we come to 1995, the year of the first detection of a sub-stellar companion to a main sequence star that was almost certainly below the deuterium burning limit, 51 Pegasi b (Mayor and Queloz, 1995). This planet is a prototype of the hot Jupiter class of planets that would come to dominate the detections of new systems for the next 20 years. With a minimum mass of 0.5 Jupiter masses and a period of 4.23 days, this exotic “Roaster planet” was unprecedented, and prompted theories of planet formation to be almost immediately turned on their head. The question of how giant planets, which must surely form at much greater orbital distances where the cooler gas allows much more efficient accretion, could move so close as to only be a few stellar radii away from the surface of their star demanded advances in theories of orbital migration, and raised questions of planet formation that remain unsolved to this day.

A rather lovely confirmation of the planetary nature of 51 Pegasi b came recently. Using the reflected daytime spectrum it was possible for Martins et al. (2015) to solve for the inclination of the system directly. They found that the system was in fact only slightly inclined with respect to the Earth, giving a true mass of the planet of $0.46^{+0.06}_{-0.01} M_J$. This also gave confirmation that water vapor

existed in the dayside spectrum.

1.1.2 Planetary archeology

In fact, one could claim that the earliest direct evidence of an exoplanetary system was discovered in 1917, though only very recently have we come to understand these observations (Zuckerman, 2015). Approximately 50% of white dwarfs show signs of metal pollution (Koester et al., 2014), which is about the same fraction of stars that show evidence of hosting a super Earth. Due to the extremely high surface gravity of white dwarfs the atmospheres are chemically stratified, with settling times for molecules heavier than helium being significantly shorter than the cooling time (Paquette et al., 1986). Therefore, if a white dwarf atmosphere appears anything but pristine, it must have accreted new material astronomically recently. As white dwarfs are the end state of almost all stellar systems, and planets appear to be very common byproducts of star formation, it stands to reason that this additional material may be the destroyed remnants of previous planetary systems. Other alternative explanations, such as accretion from roughly solar metallicity interstellar material, did not hold up to the observed lack of hydrogen in the atmospheres of DBZ white dwarfs (Farihi, 2016). Observations from *Spitzer* made it clear that debris disks were commonly found around polluted white dwarfs. The hypothesis that this material was planetary in nature was recently all but confirmed by the discovery of a disintegrating transiting planetesimal around a white dwarf (Vanderburg et al., 2015). This provides a particularly useful counterpoint to conventional exoplanet science, as it provides the only known way to measure the interiors of exoplanet material, allowing hypotheses such as carbon rich composition to be tested (Wilson et al., 2016).

Armed with this knowledge, one can with perfect hindsight go back to some of the earliest observations of white dwarfs, before their post-stellar nature was known. In 1917 Van Maanen published a short letter on *Two faint stars with high proper motion* (van Maanen, 1917), one of which he classified as an F type star, most likely due to the deep calcium H and K absorption features this object showed. As it turns out, this was the first detected metal-polluted white dwarf, and thus most likely the earliest detection of exo-planetary material, a full 80 years before the next earliest contender. Data from the photographic plate used to make these observations are still available, and the deep calcium feature is clearly visible and matches modern high resolution measurements (Farihi, 2016).

An interesting and useful fact about this population of white dwarf exoplanets is that due to stellar populations and lifetimes, most of these polluted

white dwarfs were originally A class stars, a population that is somewhat under-represented in the list of confirmed exoplanets due to various observational effects and biases. Not only are A class stars intrinsically rare, they are very difficult to measure accurate radial velocity signatures for, due to having low numbers of stellar line features. There are intriguing hints from Rossiter-McLaughlin studies that planet formation and migration processes may occur differently around hotter stars (see discussion in 1.2.5), so statistics from a large population of polluted white dwarfs that originally formed around hotter stars may in fact be able to help shed some light on the birth, as well as death, of planetary systems.

1.1.3 The planet detection gold-rush

After the first few planets were detected, a steady trickle of new detections quickly became a seemingly unstoppable flow. The advent of new, high resolution spectrographs such as CORALIE and SOPHIE made the detection of new planets significantly easier, and by the turn of the century a few dozen planets outside of the Solar System were known, all of them falling into the hot Jupiter category. Radial velocity surveys continue to be the workhorse of exoplanet discoveries, and continue to produce planets of smaller mass and greater orbital periods as sensitivity increases.

Charbonneau et al. (2000) provided the first detection of a planetary transit, and the potential of transit-surveys to rapidly discover large numbers of exoplanets became clear (Horne, 2003). Transiting exoplanets have an advantage over purely radial-velocity detected exoplanets as their true masses can be known, and not just a minimum mass. Transits also make atmospheric characterization significantly easier, as transmission spectroscopy is a powerful tool.

“Transit surveys” soon began in order to capitalize on this population of hot Jupiters with high transit probability. The first detection of transits from a dedicated survey- OGLE (Optical Gravitational Lensing Experiment) (Udalski et al., 2002), has detected 28 transiting planets towards galactic core. OGLE was a medium aperture 1.3m telescope with seeing limited resolution survey and a small field of view. The planets that were detected were all around very faint stars in the galactic bulge.

A different approach was taken by TrES (Trans-atlantic Exoplanet Survey), who realised that planets around brighter stars are valuable, and began a small aperture wide-field survey that could cover large regions of sky. This became the template for the other very successful surveys of the decade, WASP and HATNet (Hungarian-made Automated Telescope Network) (Alonso et al., 2004).

HATNet (Bakos et al., 2002) operates in both the northern and southern

hemisphere. There are observatories at several different latitudes, with the aim of increasing the length of a period of an exoplanet that can be detected from the ground. It has so far detected 68 planets.

WASP (Pollacco et al., 2006) has discovered 94 planets to date, it operates in north and south with an observatory at La Palma and in South Africa. WASP remains the single most successful ground based transit survey, and has detected almost as many transiting planets as all other surveys combined (94 out of 193 at the time of writing).

MEARTH (Irwin et al., 2014) is a specifically targeted search for low mass M-dwarf planets. A list of potential planet hosts are observed several times a night each, with an increasing cadence if a transit is seen to begin. MEARTH discovered the highly observed object GJ 1214b (Charbonneau et al., 2009) and more recently, the earth sized planet GJ 1132b (Berta-Thompson et al., 2015). The recently announced TRAPPIST system of three temperate earth sized planets was also detected with this directed technique (Gillon et al., 2016)

Numerous other surveys; XO, lupus, WTS, POTS and QES, KELT and KELT-south have all also detected at least one planet (McCullough et al., 2005; Weldrake et al., 2008; Birkby et al., 2011; Koppenhoefer et al., 2013; Alsubai et al., 2013; Pepper et al., 2007, 2012).

Mascara (Snellen et al., 2013) takes a slightly different approach, pushing the “small aperture wide field” concept to the extreme, covering virtually the entire sky from several fixed mounted 20mm F/2.8 wide-field Canon lens. The goal is to detect the handful of bright 4-8 magnitude planet hosts that have so far slipped through the grasp of the surveys discussed so far.

At the same time, the space based search for transits was beginning with CoRoT (Barge et al., 2008), which went on to discover 29 exoplanets as well as providing new insights into stellar oscillations. It also led to the discovery of the first rocky exoplanet, CoRoT-7b (Queloz et al., 2009).

Kepler (Koch et al., 2010) took this to the next level. Launched in 2009, it had the stated purpose of finding “Eta Earth”, or the fraction of Earth like planets. The survey strategy was to stare continuously at a 100 square degree field in the constellation of Cygnus. This would allow the detection of much longer period transiting objects than had been found before, as to be confident of a detection a minimum of three events would need to be seen. In fact, the majority of planets detected to date originated in this mission, and the data on populations of exoplanets has been immensely exciting. After the reaction wheels malfunctioned, the prime *Kepler* mission was declared over. But, due to an ingenious use of the pressure



Figure 1.2: The NGTS array of telescopes

of the solar radiation, a solution to the stability was found if *Kepler* was pointed towards the ecliptic plane, and so K2 was born from the ashes of *Kepler* (Howell et al., 2014). K2 has gone on to be a great success, despite the reduced sensitivity from poor guiding, and the relatively short amount of time spent on each field.

TESS is scheduled for launch in 2017, and will combine the large area approaches of ground based surveys with the precision achievable from space (Ricker et al., 2014). NGTS will likely play an important role in confirming TESS candidate planets which are too long period to be detected multiple times with the relatively short field observations of TESS.

PLATO, scheduled for launch in 2024 is set to be a revolution in exoplanet science, with the potential to detect thousands of planets, including many Earth like planets in the habitable zone (Rauer et al., 2014).

The Next Generation Transit Survey (NGTS) (Wheatley et al., 2013) is intended to be the spiritual successor to WASP, building on many of the lessons learned. The facility is comprised of 12 fast astrographs with 20cm apertures and a wide 3 degree FOV. Combined, the 12 telescopes cover a field similar to *Kepler*, but the survey is designed to regularly cycle fields. The additional precision gained from improvements to the CCD's, tracking stability and optics allows a much higher precision than other ground based surveys, allowing the detection of 1 mmag transits. *Kepler* has shown us that there is a large population of Neptune radius planets, that the precision of NGTS will allow us to detect. However, the wider sky area covered

by NGTS will provide these planets around brighter stars, which will allow their masses to be accurately measured and their atmospheres to be studied, which is rarely possible for the faint stars observed by *Kepler*. I have been involved with the development of the pipeline and data quality assurance products for this project, as described in Chapter 3.

The remainder of this Section is arranged as follows: In Section 1.2 I will describe in more detail a few important methods for detecting exoplanets, with particular focus on the *Transit method*. Section 1.3 will provide an overview of the currently known population of exoplanets, with a discussion on the relative signal size of the planets studied in this thesis. In Section 1.4 I will describe the observable properties of exoplanets, and how these can be used to characterise the planet’s atmospheric properties (Source: NASA exoplanet Archive, <http://exoplanetarchive.ipac.caltech.edu/> accessed 17/06/16).

1.2 Detection methods

1.2.1 Transits

Perhaps the oldest scientific use of a transit was the observation that the shadow of the Earth cast onto the moon during a lunar eclipse was circular. Aristotle presented this in 350 BC as evidence of the spherical nature of the Earth in his cosmological treatise *On the Heavens*. More recently, lunar eclipses have been used to observe the transmission spectrum of the Earth as though it were an extrasolar planet, which is particularly interesting from an astrobiological perspective, when considering what signatures may be accessible on nearby exoplanets (Pallé et al., 2009; Yan et al., 2015b).

Beyond the rather unobtrusive affair of lunar and solar eclipses, we have to wait until the invention of the telescope before the importance of transits becomes clear again. In 1639 Jeremiah Horrocks became the first person to observe the transit of another planet, when Venus passed in front of the Sun. In 1663 it was suggested by James Gregory that simultaneous observations of the transit of Mercury from widely separated points on the Earth could be used geometrically to make the first accurate determination of the astronomical unit, the average distance between the Sun and the Earth. Measuring the astronomical unit allows one to put physical units on the scale of the Solar System, and serves as the baseline for measurements of stellar parallax, thus forming the base of the cosmic distance ladder that allows the entire observable universe to be put into perspective. It could be argued that the eventual determination of the astronomical unit by an international collaboration to

observe the transit of Venus in 1761 and 1769 is one of the most important results in the history of astronomy. Remarkably, once the results of the various observers had been collated by de Lalande (1771), the published result was within 2% of the modern value from space probe telemetry.

Finally we come to the transits of extrasolar planets, the first of which was HD 209458b (Charbonneau et al., 2000), the radial velocity detection of which is described in Mazeh et al. (2000). At the time, radial velocity surveys had already detected 11 presumed extrasolar planets. While their planetary nature was strongly suspected, it was not certain due to the $\sin i$ dependence of the minimum mass. They were not likely to be stellar in nature, as tidal synchronization is observed in all short period spectroscopic binaries (Mayor and Queloz, 1995), but massive brown dwarf companions on orbits inclined by greater than $\sim 10\%$ were possible. The companions were known to have orbits with semi-major axes, a , of ~ 0.05 AU around solar type stars. Viewing the problem purely geometrically, assuming that planet orbits are oriented with random inclinations (there being no a priori reason why they wouldn't be) then the probability that the orbit of an extrasolar planet crosses the face of its star from our perspective on Earth can be derived quite simply from geometry,

$$P_{transit} = \left(\frac{R_* \pm R_p}{a} \right) \left(\frac{1 + e \sin \omega}{1 - e^2} \right) \quad (1.1)$$

where e is the eccentricity and ω is the argument of periapsis, which is defined to be the position of the orbit with respect to a reference plane when the planet is at its closest approach. Here, the plus or minus signifies whether grazing transits, where the planet's disc does not fully cover the star, are to be included. The term is positive if grazing transits are allowed and negative if they are not. A full derivation of the transit geometry can be seen in Winn (2010). Some simplifying assumptions can be made to this equation, for the case where the planet radius is significantly smaller than the stellar radius, and for where eccentricity is negligible. From the observational side, eccentricity is a measurable parameter for planets detected through radial velocity, and is indeed very close to 0 for the many hot Jupiters. In the absence of additional system components the tidal forces should quickly circularise the orbits of hot Jupiters, though there are some dramatic examples of warm Jupiters with very high eccentricities, such as HD 80606b (Barge et al., 2008). Setting both R_p and e to 0,

$$P_{transit} = \frac{R_*}{a} = 0.103 \left(\frac{R_*}{R_\odot} \right) \left(\frac{a}{0.045 AU} \right)^{-1} \quad (1.2)$$

so for a typical hot Jupiter the probability of a transit is about 10% (or slightly

higher if we assume a non-negligible planet radius and allowed for grazing transits). By contrast, a planet in the habitable zone of an M dwarf has a transit probability of about 1%, and the Earth has a transit probability of 0.5%. This is one component that makes transit searches heavily biased towards short period planets. Another factor that preferences short period planets is the “window function”, which is a description of the likelihood of observing a planet multiple times (which is typically considered a requirement for becoming a viable candidate, with some exceptions) given a set of observations. Ground based transit surveys typically have window functions that make detections of periods of much more than 10 days extremely inefficient, due to day night cycles, seasonal changes in observability, and weather. Indeed, the longest period transiting planet detected directly from the ground is HAT-17b, with a period of 16.3 days (Brahm et al., 2016). HATNEet have attempted to improve their window function by observing from multiple longitude sites. MASCARA has a similar planned strategy.

In Figure 1.3 I have plotted the currently known population of exoplanets. The bias of the transit method to close in planets is very clear from the extreme clumping of transit planets towards the top left of the diagram - this is not a true feature of the underlying planet population, as can be seen from the relatively much more even distribution of radial velocity planets.

To go back to Charbonneau in 2000, since 11 planets were known from radial velocity surveys, and the transit probability was 10%, it was a pretty reasonable chance that at least one would transit, and the expected signal of such a transit event was well within reach of contemporary equipment. Similar to the transit probability, the amount of flux blocked by the transit can be estimated by straightforward geometry,

$$F_{transit} = \left(\frac{R_p}{R_*} \right)^2 = 0.01 \quad (1.3)$$

so assuming Jupiter mass planets are approximately Jupiter radius, the drop in flux from a transit event is about 1% (I will show in Section 1.3 that this is a fair assumption, as planets have an approximately Jupiter sized radius over 2 order of magnitude in mass). The Earth on the other hand, only blocks about 0.01% of the Sun’s light, but for Earth mass planets around low mass stars, the depth is also about 1%. This formula is of course only an approximation, as it assumes that the star is a disk of uniform intensity. In reality, limb darkening can be important, but it is accurate as a first approximation. Indeed, the transit depth of a hot Jupiter is seen to typically be of order 1%.

Considering the geometry of an exoplanet transit as seen in Figure 1.6, a




Figure 1.3: The currently known exoplanets, colour coded by their discovery method. The planets of the Solar System are also marked for comparison. In particular, it is important to be aware of the selection effects of the different methods. The large concentration of hot Jupiters at the top right is not a real feature of the underlying population, this can be seen in the radial velocity surveys, which are now largely complete for giant planets within 1 AU. The traditional $13 M_J$ mass limit for planets is given as a dotted line.




Figure 1.4: Schematic diagram of a transit. Figure from Winn (2010)

useful parameter to define is the *impact parameter*, b . This is the distance of the chord of the transit from the centre of the star from the perspective of an observer on Earth, and is obviously related to the inclination of the orbit. It is useful to describe this parameter as a fraction of the stellar radius,

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

this factor is important in calculating the duration of a transit, which for a circular orbit the total time, i.e., between first and final contact is

$$T_{tot} = \frac{P}{\pi} \sin^{-1} \left(\frac{R_*}{a} \frac{\sqrt{\left(1 + \frac{R_p}{R_*}\right)^2 - b^2}}{\sin i} \right) \quad (1.5)$$

which once more with some simplifying assumptions, assuming $R_p \ll R_* \ll a$, and excluding grazing events, reduces to

$$\begin{aligned} T_{tot} &= \frac{R_*}{a} \frac{P}{\pi} \sqrt{1 - b^2} \\ T_{tot} &= 3.2hr \sqrt{1 - b^2} \left(\frac{P}{4days} \right) \left(\frac{R_*}{R_\odot} \right) \left(\frac{a}{0.045AU} \right)^{-1}. \end{aligned} \quad (1.6)$$

The characteristic timescale for a typical hot Jupiter transit is thus 2-3 hours, which is fortunately easy to do in a single night of observations. Through the same process,

one can derive a similar approximation for the time spent in ingress or egress, τ ,

$$\begin{aligned}\tau &= \frac{R_p}{a} \frac{P}{\pi} \sqrt{1 - b^2} \\ \tau &= 0.3hr \frac{1.0}{\sqrt{1 - b^2}} \left(\frac{P}{4days} \right) \left(\frac{R_p}{R_\odot} \right) \left(\frac{a}{0.045AU} \right)^{-1}\end{aligned}\tag{1.7}$$

We have now what can be thought of as the primary observables of a transit. If we neglect for now the limb darkening and imagine the idealized case of a circular orbit and a uniform stellar surface, in addition to the period we have three parameters; the depth, transit duration and ingress/egress duration, that can fully describe the shape of the event. As can be shown from Equation 1.3 and by rearranging Equations 1.6 and 1.7, the same shape can also be described by four model parameters; P , R_p/R_* , a/R_* and b . In some cases, it can be more convenient to use i instead of b through Equation 1.4.

From the information of the transit alone, it is not possible to make direct measurements of the planetary radius, the knowledge of the radius of the star will often be a limiting factor. This can be an issue for large surveys such as *Kepler*, where accurate information was not available for all stars in the initial catalogue. Huber et al. (2013) show that when more accurate stellar parameters were calculated through astroseismology, many *Kepler* planet candidates were in fact around evolved stars and had radii underestimated by a factor of about 1.5. However, if the effective temperature of the star is known, which can be estimated simply from multi-band photometry of the star, then the equilibrium temperature of the planet can be estimated directly from a/R_* , as,

$$T_{eq} = T_*(1 - A)^{1/4} \frac{1}{\sqrt{2}} \left(\frac{a}{R_*} \right)^{-1/2}\tag{1.8}$$

where A is the planetary albedo. This is not the only parameter that can be derived directly from a transit lightcurve. A particularly useful result derived by Seager and Mallen-Ornelas (2003) using Kepler's third law, is that the stellar density, given some assumptions can be determined uniquely,

$$\rho_* = \frac{3\pi}{GP^2} \left(\frac{a}{R_*} \right)^3 - \rho_p \left(\frac{R_p}{R_*} \right)^3\tag{1.9}$$

since, as shown in Equation 1.3, a typical hot Jupiter transit depth is of order 1%, the second term on the right hand of the equation is small, and can be neglected. This leaves the stellar density to be determined only from the period and a/R_* ,




Figure 1.5: A high precision transit of HAT-P-7b from *Kepler*, with secondary transit and phase curve visible. Figure from Borucki et al. (2009).

which is itself directly measurable from the transit shape.

In some exceptional cases, the stellar parameters might be known to very high precision, through astroseismology, for example. In this situation the stellar density could be used as a prior in a Markov Chain Monte Carlo (MCMC) fit to constrain the other transit parameters. This result must be used cautiously however, as it implicitly assumes a circular orbit. This assumption, as well as a systematic bias when fitting b , led to incorrect stellar density, and hence stellar type estimates, for *Kepler* objects of interest in Batalha et al. (2013), as discussed in Huber et al. (2013).

Surprisingly, with the addition only of K_* , the semi-velocity amplitude from radial velocity measurements, the surface gravity of the planet, g , can be calculated independently of any model stellar parameters,

$$g = \frac{2\pi}{P} \left(\frac{a/R_*}{R_p/R_*} \right)^2 \frac{K_* \sqrt{1-e^2}}{\sin i} \quad (1.10)$$

as shown by Southworth et al. (2007). This can serve as a “sanity check” of the planetary nature of a companion.

Interestingly, since both equilibrium temperature and surface gravity can be measured independently of stellar parameters, the quantity $H\mu$, the expected scale height multiplied by the mean molecular weight, is actually a directly measured




Figure 1.6: Example limb darkening profiles for a 3000 K and 6000 K star (blue and red, respectively), in B and I wavebands (dashed and dotted lines). Limb darkening is more severe for cooler stars, and at bluer wavelengths

quantity. In some sense, we can speak more confidently about the expected properties of a detected planet’s atmosphere than we can about the bulk properties of the planet itself.

In this discussion so far we have ignored the impact of limb darkening. Limb darkening is an effect where the intensity at the center of the stellar disk appears brighter than at the edge. Limb darkening is a consequence of optical depth and the Temperature Pressure profile of the stellar atmosphere. At optical wavelengths, an observation at a high angle of incidence (i.e, towards the center of the star) will look deeper into the atmosphere than a low incident angle (towards the limb). At the region of the star’s atmosphere relevant to optical wavelengths, the temperature increases with depth, so the high angle of incidence observes a hotter, and therefore brighter region of the star.

Tables of limb darkening coefficients can be calculated from stellar atmosphere models, such as ATLAS (Kurucz, 1993), as a function of both stellar type and instrument bandpass (e.g. Sing, 2010). In Chapter 6, since the transit must be fit to arbitrary wavelength channels to construct the transmission spectrum, we use the LDTK code of Parviainen and Aigrain (2015) to calculate priors on the limb darkening distribution using the stellar values and their errors as inputs.

Mandel and Agol (2002) derived a useful set of equations for the modelling of transit lightcurves, which are widely used, as evidenced by the over 1000 citations

for the paper. I use the Mandel and Agol equations to fit the spectroscopic transit lightcurves in Chapter 6. In Chapter 5 it was not practicable to model the transit analytically in this way, since the interactions between the rotation of the star, stellar spectrum and rotationally broadened planetary spectrum are non-trivial. Instead, the transit was modelled with a pixel grid.

1.2.2 Radial velocity

The radial velocity method for detecting exoplanets, where the reflex motion of the star due to the orbit of the planet is detected spectroscopically, was first proposed by Struve (1952). The basic principle can be derived simply from Keplers third law,

$$GM_{tot} = \left(\frac{2\pi}{P_{orb}} \right)^2 a^3 \quad (1.11)$$

where G is the gravitational constant, P_{orb} is the orbital period, a the semi-major axis and M_{tot} is the total mass of the system. For the case of an exoplanet around a main sequence star this can be approximated to a good level of accuracy as just the mass of the star, since $M_J/M_\odot \sim 1 \times 10^{-3}$. Periods can be measured precisely, while the mass of the star can be estimated spectroscopically and through evolutionary models, and in some exceptional cases through astroseismology (Silva Aguirre et al., 2015). For the case of exoplanets in binary systems the mass can be determined much more precisely and directly, since the stellar binary orbit can be fully solved (Doyle et al., 2011). Mass estimates do not greatly affect the semi-major axis, as it depends on them only to the 1/3rd power.

Now, if a_1 and a_2 are the distances of the two objects to their mutual center of mass (barycenter), such that $a = a_1 + a_2$, then, with M_1 and M_2 as the masses of the two objects,

$$M_1 a_1 = M_2 a_2 \quad (1.12)$$

follows trivially. Since a_1 and a_2 can not typically be measured directly, they can be recast in terms of the radial velocity semi amplitude, K_n , which is defined as

$$K_n \equiv \frac{2\pi}{P} \frac{a_n \sin i}{(1 - e^2)^{1/2}} \quad (1.13)$$

where i is the inclination of the orbit and e is the eccentricity. The radial velocity of the *star* is something that can be directly observed. If we let object 1 be the star and object 2 be the planet, then by combining 1.11, 1.12 and 1.13 we can construct




Figure 1.7: The radial velocity curve that led to the detection of the first exoplanet around a Sunlike star, 51 peg b (Figure from Mayor and Queloz (1995))

the *binary mass function*, f ,

$$f = \frac{M_p^3 \sin^3 i}{(M_* + M_p)^2} = \frac{P_{orb} K_*^3}{2\pi G} (1 - e^2)^{3/2} \quad (1.14)$$

Now, since the mass of the star, M_1 , can be estimated from other means, and if we assume that it is significantly greater than the planet's, then it is trivial from this definition to construct an equation for the minimum mass of the exoplanet, the caveat being that the inclination cannot be determined directly.

$$M_p \sin i = K_* M_*^{2/3} \left(\frac{P_{orb}}{2\pi G} \right)^{1/3} (1 - e^2)^{1/2} \quad (1.15)$$

In some cases it has been possible to break this degeneracy by measuring the radial velocity of the planet itself through cross correlation of the molecular lines in the atmosphere (Snellen et al., 2010). 51 Peg b itself has recently had its dayside reflection spectrum measured with HARPS, allowing its velocity to be measured and hence removing the $\sin i$ dependence, effectively making the system a double lined eclipsing binary (Martins et al., 2015). The inclination was found to be 80^{+10}_{-19} degrees, which corresponds to a mass of $0.46^{+0.06}_{-0.01} M_J$ for 51 peg b, which is in fact

within 2 sigma of the originally published minimum mass of 0.47 ± 0.02 , finally confirming beyond reasonable doubt the planetary nature of the object.

An advantage of radial velocity surveys is their relatively weak dependence on period for the signal to noise, making them less biased towards the hot Jupiter population that dominates the transit surveys. Additionally, since they are not bound by the relatively low transit probability, they can be detected from a smaller sample of brighter stars. Indeed, the two targets for atmospheric follow up, HD 209458b and HD 189733b, were both detected around bright stars with radial velocity surveys.

1.2.3 Microlensing

Microlensing is a particularly fascinating consequence of General Relativity, and was directly predicted by Einstein (1936). Due to the distortion of spacetime, a massive object is capable of briefly magnifying a background object to enormous degrees³, though the alignment must be precise, and is an ephemeral occurrence. The lensing lightcurve of a star can last several weeks, and has a distinct shape that can be fit by calculations from GR. If the lensing system is a star with a planetary mass companion, there will be an additional smaller lensing event superimposed on the primary event, which typically has a timescale of a few hours. So as soon as a stellar lensing event is observed, multiple observatories across the world will observe as frequently as possible in case a planetary event is still to come. There are often severe degeneracies in the combined system solutions, but these can potentially be solved by microlensing parallax (Alcock et al., 1995), when the event is observed by two observatories with a sufficiently large separation (i.e, one on Earth, and one at L2).

Micro-lensing surveys typically observe areas of high stellar density, such as the galactic bulge, as this maximizes the probability of detecting a lensing event. A significant disadvantage is the one-time nature of the events, whilst typically highly significant, there is no way to follow up on a system and gain additional information or constraints on the system. However, it provides very valuable input to our understanding of planet populations, as compared to other planet finding techniques it covers a very different and complementary parameter space, having a maximal sensitivity beyond the snow-line in most systems (Beaulieu et al., 2008). See Figure 1.3, where the microlensing planets are typically the closest to the planets in our own Solar System.

Microlensing has provided a tantalising hint of a large population of free-

³There are suggestions that the gravitational lens of the Sun could one day be used for astronomy. The major issue is in getting to the focal point at 1000 AU (Eshleman, 1979)

floating planets (Sumi et al., 2011) which may be twice as common as stars in the Milky Way (creating another challenge to the traditional definition of a “planet”). This enormous population of free planets presents a challenge for theoretical models of planet formation and ejection (e.g. Veras and Raymond, 2012). Remarkably, microlensing also has the potential to detect planets in other galaxies (Ingrosso et al., 2011).

1.2.4 Direct imaging

Direct imaging, as the name implies, is the only planet detection technique where the light of the planet is detected directly. Due to the extreme contrasts and narrow separations from their host stars, there are enormous instrumental difficulties in directly detecting exoplanets, and the absolute number is a relatively small fraction of the total detected planet population. The science advantages of having direct access to the photons from the planet are becoming apparent.

Currently, the technique is only possible for young, giant planets which emit significant IR fluxes due to heat from their gravitational contraction. Referring to Figure 1.3, the planets discovered by direct imaging occupy a distinctive region towards the top right of the diagram, straddling the 13 M_J line typically used to differentiate between planets and brown dwarfs. Mass measurements of directly imaged planets are not direct measurements, but are inferred from models of young planet cooling, and are therefore very model dependent.

Without the contamination of the photons from the host star, spectra of directly imaged planets can reach excellent signal to noise ratios, and are rapidly approaching the quality of results obtained for field brown dwarfs (e.g. Ingraham et al., 2014).

β -Pictoris b is an example of a planet detected through direct imaging, on an orbit between 8 and 15 AU (Lagrange et al., 2010). The planet later had its rotation speed directly measured, the first time this was accomplished for an exoplanet (Snellen et al., 2014).

The technique is also heavily biased towards planets on extremely wide orbits, which may have different formation channels and histories than the close-in planets detected by other techniques. A full understanding of planet formation will require explanations of both of these populations.

The next generation of high contrast imaging surveys including SPHERE, GPI, ScEXAO and Project 1640 (Beuzit et al., 2008; Macintosh et al., 2014; Jovanovic et al., 2015; Hinkley et al., 2011) are all currently operating, and promise exciting results in the near future.

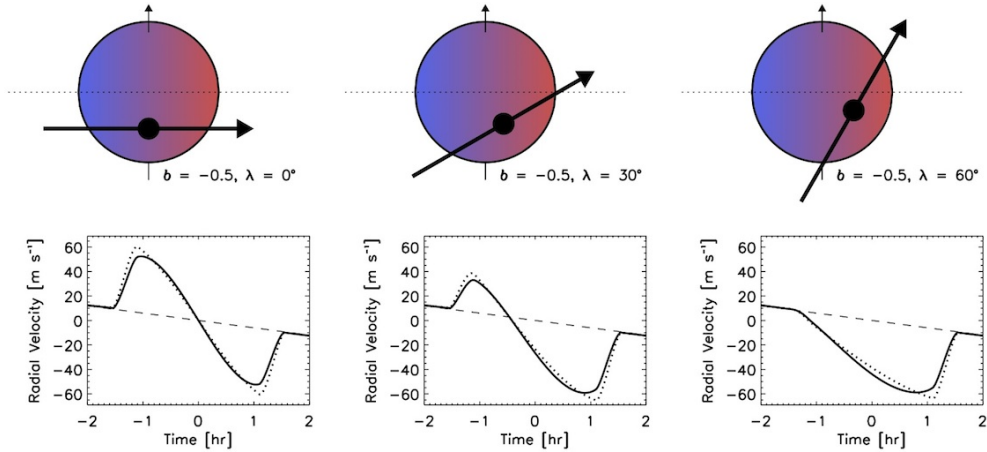


Figure 1.8: A Diagram of the Rossiter-McLaughlin effect for three different trajectories (Figure adapted from Gaudi and Winn (2007))

1.2.5 Rossiter-McLaughlin effect

Whilst not directly an exoplanet detection technique, the Rossiter-McLaughlin effect is an important tool that can be used to confirm the planetary nature of an object (Gaudi and Winn, 2007), and it has important consequences for other observational techniques.

The Rossiter-McLaughlin effect has a long history. Holt (1893) realized that in an eclipsing binary, if the stellar rotation axis was aligned with the orbital plane, then an eclipse would first cover the half of the star rotating towards the observer, and then the half rotating away. This would yield a characteristic signal if the radial velocity of the system was being measured through absorption lines, with a characteristic red shift followed by a blue shift. This effect is measurable because the natural width of the stellar lines is typically smaller than the rotational broadening (a diagram of the effect can be seen in Figure 1.8). A tentative detection of the effect was claimed by Schlesinger (1910), but it was convincingly demonstrated by Rossiter (1924) and McLaughlin (1924), from where the effect gets its name. The first theoretical framework for the Rossiter-McLaughlin effect was developed by Petrie (1938), and a numerical model of the effect caused by an exoplanet transit was first made by Queloz et al. (2000). The Rossiter-McLaughlin effect has been seen both inside and outside the Solar System, with a detection in the transit of Venus in 2012 (Molaro et al., 2013).

The Rossiter-McLaughlin effect is a vital input to theories of planetary formation and migration. The prevailing opinion is the planets form in a disk, that

should at least initially be aligned with the angular momentum vector of the star (e.g. Pollack et al., 1996). Hot Jupiters are believed to form beyond the snow-line and only subsequently migrate inwards. The mechanism by which they move in, either through disk interactions (e.g. Lin and Papaloizou, 1986), planet-planet scattering (e.g. Chatterjee et al., 2007) or the Kozai-Lidov effect (e.g. Naoz et al., 2011), make different predictions about whether they should be expected to still be aligned with the stellar spin vector after migration.

A particularly interesting result of measuring the spin-orbit alignment for many systems is that there is a significant dichotomy between stars cooler than 6250 K and those hotter, with planets around hot stars being significantly more likely to be misaligned (Winn et al., 2010). The authors speculate that this may indicate a fundamental difference between formation and migration channels for cool and hot stars, or that all planets have their obliquities randomised during migration, but cool stars are able to subsequently re-align the planet through efficient tidal interactions with the convective zone, which hotter stars do not possess.

In Chapter 5, I show that the Rossiter-McLaughlin effect has an effect on the transmission spectra, and may have led to spurious detections in past observations. The Rossiter-McLaughlin effect is implicitly accounted for by the two-dimensional transit model I describe in Chapter 5.

1.3 Observed population of exoplanets

Figure 1.9 shows the currently known planets for which a determination of the density exists. In practice, these are all transiting planets which have had their masses measured by an additional method. The combination of transit and radial velocity (RV) measurement is particularly powerful, as it removes the $\sin i$ dependence from the mass determination, transforming the minimum mass into a true mass measurement. It is immediately clear from Figure 1.9 that for the majority of planets below about a Saturn mass the accuracy of the density determination is very poor. This is because as well as having an intrinsically weaker signal due to their size, the majority of these planets are from the *Kepler* mission, and on average have visual magnitudes of greater than 12. Both of these factors conspire to make following up on these systems particularly difficult. For most currently known sub-Saturn mass planets it is not possible, or prohibitively expensive to determine their mass through conventional radial velocity measurements with current generation instruments. Instead, the majority of measurements in this range come from the fraction of systems in the *Kepler* sample that exhibit Transit Timing Variations (TTVs). There is a




Figure 1.9: The mass-density relation for the currently known exoplanets. Planets with masses measured from radial velocity are in pink, and from TTV's with green. In this plot the points with fractional density errors of greater than 30% have been faded to place more emphasis on the well characterized sample. The dotted grey line signifies the 13 M_J cutoff commonly used to define a planet. Blue dashed lines mark broad trends in the data. The thick black line are the model results of Fortney et al. (2007), for a coreless H/He planet on a 0.045 AU orbit around a Sun-like star.

small but significant difference in the sample of planets with RV and TTV measurements. It can be seen that on average the RV sample of intermediate mass planets have a systematically higher density than the TTV sample in this range. Steffen (2016) explains this difference in terms of the sensitivity biases of the two techniques: the signal to noise of RV measurements improves as the period shortens, while the TTVs become weaker. This implies that for a given mass, the RV sample will be more sensitive to planets on shorter orbits, which are susceptible to higher radiation fluxes and hence more likely to have higher levels of atmosphere erosion, which leads to a statistically denser sample (I explore the efficiency of atmosphere erosion in Chapter 4). The difference between the two samples is most likely physical in origin, instead of implying a systematic bias in either of the well understood measurement techniques. Additionally, it is possible that the dynamically-packed, near-resonance systems that lead to TTVs have different formation histories, which could yield a statistically different sample. Both methods are useful, and as more measurements are made, untangling the differences in the populations they probe will provide rich fodder for dynamicists and formation theorists.

One of the key science goals of the NGTS project is to provide a statistical sample of planets in this size range that are bright enough to have their mass measured to a high level of precision.

1.3.1 Mass-density relations

Examining the population of exoplanets in Figure 1.9, we can see that they are broadly separated into 3 characteristic regimes.

Firstly, for planets less massive than about 5 Earth masses the density is approximately constant. Admittedly, for the time being this regime is dominated by the terrestrial planets in the Solar System, as at this time no exoplanet has had its density measured to better than 30% below 3 Earth masses (10^{-2} Jupiter masses). At this size, planets do not have significant surface gravity to hold onto significant hydrogen/helium atmospheres, so are comprised of rocks, ices and metals. To first order, these materials are incompressible at the pressures typically encountered within the interiors of these planets, so for a given composition the density will remain similar as the planet increases in mass. Of course this is only approximate, there is gravitational compression of materials in the Earths interior (this is why the Earth is slightly more dense than Mercury despite having a smaller iron core fraction), but the effect is small compared to the other broad trends in the mass-density diagram.

At some point, a planet becomes large enough to accrete a significant primary

atmosphere of hydrogen and helium. The exact point of the transition is unclear, and is not likely to be a sharp discontinuity, but it is believed to be around 1.6 Earth radii or ~ 5 Earth masses (Rogers, 2015; Buchhave et al., 2014). The nature of core accretion is that the planets will quickly gain a thick atmosphere of hydrogen and helium after this threshold, and decrease in average density.

Most planets that we have well characterised so far are particularly low density, if one were to throw this sample of planets into a sufficiently large body of water, a little over half would float.

Planets in the second mass range, roughly between a few Earth masses and a Saturn mass, are observed empirically to roughly follow an *inverse* mass-density relationship $\rho \propto M^{-1/2}$. The index was originally found to be -0.485 by fitting a relation between Earth and Saturn, and appears to be at least approximately correct for describing the broad trend of known exoplanets in this range, though the scatter around the relation is significant. The trend describes the tendency of more massive planets to have a larger fraction of their mass in hydrogen/helium atmospheres. This makes sense, as there is simply not enough material to build arbitrarily massive rocky planets in a stellar accretion disk, and a larger planet is more likely to begin the runaway process of accreting a giant atmosphere. This feature is an expected outcome of the core-accretion theory of planet formation. In all, getting to Saturn mass and above does not appear to be possible without accreting a large gaseous envelope.

However, for a given mass there can be a large range of core-mass fractions, leading to a scatter in density of up to two orders of magnitude. Lopez and Fortney (2014) show that when the atmosphere fraction exceeds $\sim 1\%$ the radius is virtually independent of the mass. Additionally, planets with high gas fractions have low surface gravities, and as a result can sustain very large scale heights - this makes their observed radii particularly sensitive to the amount of radiation they receive. The very low density planets are whimsically referred to as “Super-Puffs” by Lee and Chiang (2016). Forming them is a difficult task theoretically, as at the orbital positions where they are detected their low core masses would not have had sufficient gravity to accrete their thick atmospheres. It is believed they formed further out and later migrated in. Beyond 1 AU the gas is cooler and less opaque, which decreases the radiative timescales to levels where the planet is able to accrete its thick hydrogen atmosphere, however, they must then quickly migrate inwards before a catastrophic runaway accretion takes over and they end up as gas giants. This knife edge formation history probably explains why they are comparatively rare, but their existence further reinforces the rich complexity and diversity of planet

formation.

A recent paper by Lammer et al. (2016) suggests that some of the low density planets observed may have ‘true’ radii significantly smaller than their transit depths suggest. For low mass planets with high scale heights, the difference between the effective transit radius and the 1 bar level typically used as a reference point for gas dominated planets can be significant. Optically thick clouds can potentially form at low pressures (< 0.1 mbar), which is about 7 scale heights above where a clear atmosphere would become optically thick, since broad-band transit observations are sensitive to pressures of ~ 100 mbar (Fortney et al., 2003).

Take GJ 1214b as an example (Berta-Thompson et al., 2015). The planet has a radius of 0.254 Jupiter radii, or 17570 km, and a scale height of about 272 km, assuming a hydrogen-helium dominated atmosphere with a mean molecular weight of 2.3. 7 scale heights is 1904 km, or 11% of the reported radius, if we imagine a planet otherwise the same as GJ 1214b without clouds it would have an estimated density 36% higher. Therefore caution must be used when using a mass-density diagram to infer composition trends, as without a common reference pressure for the radii measurements there can be a significant additional source of scatter.

Another factor that effects the envelope fraction of these planets, and therefore their position on the mass-density diagram is atmosphere erosion. There are hints within the *Kepler* planets that evaporation may be a dominant force sculpting the observed population of low mass planets, for example by explaining the seeming bi-modality in occurrence rates, whereby 2 Earth radii planets are comparatively rare. (Lopez et al., 2012; Lopez and Fortney, 2013; Owen and Wu, 2013). Atmosphere erosion may also help explain the presence of a small number of anomalously dense objects in this mass range that cannot be explained by conventional formation theories. The most bizarre cases by far being Kepler-128 b and c, which both have radii of 1.13 ± 0.03 Earth Radii, but TTV measured masses of 30.7 ± 6.0 and 33.3 ± 3.3 Earth masses - this implies densities of over 100 g/cm^3 , which is approaching the density of a stellar core. If these detections are genuine they may be the leftover cores from much more massive objects whose atmospheres have been completely stripped away, and have not yet relaxed to a lower pressure state (Mocquet et al., 2014). Atmosphere erosion will be explored in more detail in Section 1.4.2 and in Chapter 4.

We come to the third regime, where most of the well characterized planets we know of currently reside - the Jupiters. Referring back to Figure 1.9, the final trend is for the density to increase with mass.

The trend of increasing density can also be seen in the mass-radius relation-

ship, shown in Figure 1.10. The plot includes theoretical mass-radius relationships for coreless H/He planets at several different irradiation levels from Fortney et al. (2007). Jupiter sits quite neatly on the prediction for a 4.5 Gyr old planet, but Saturn is a little denser than expected. This is because Saturn is believed to have an approximately $20 M_{\oplus}$ core (Fortney, 2004), though the exact mass is uncertain due to our lack of knowledge of the equation of state of hydrogen and helium mixtures (Saumon and Guillot, 2004).

The decrease in planet radius with mass is due to *electron degeneracy pressure*. Because the electron density is proportional to M/R^3 , the Fermi energy is proportional to $(M/R^3)^{2/3}$. The gravitational potential energy of a proton, on the other hand, is proportional to M/R . Equating these two relationships it can be shown that the radius is proportional to $M^{-1/3}$. Electron-degenerate objects span a factor of $\sim 47,000$ in mass, so the least massive electron-degenerate object has 17 times the radius of the most massive object under this simple relationship.

However, this relationship is only exactly true for *cold spheres* (Zapolsky and Salpeter, 1969). Hot Jupiters have a large amount of internal heat from formation, which is larger for more giant planets, and also the intense irradiation from the star. In fact, these competing effects all conspire such that with the typical irradiation of a hot Jupiter ($\sim 500 F_{\oplus}$) the radius is approximately constant at $1.1 R_J$ and the mass radius relationship is flat over 2 orders of magnitude, this can be seen in Figure 1.10. However, there is a population of planets that are significantly inflated with respect to this theoretical curve, for which there is still no well known solution.

1.3.2 The Radius anomaly

Baraffe et al. (2003) showed that the hot Jupiter planets can have their radii increased by the high levels of irradiation they receive (or more correctly, that their contraction from their formation radii can be retarded.) This explains why the majority of Jupiter mass planets have densities somewhat lower than Jupiter. However, a sub-category of hot Jupiters exist that are significantly larger than can be explained by the incident radiation alone, among them the prototypical inflated hot Jupiter, HD 209458b. The mechanism behind this inflation is not known, but the theories behind it can broadly be categorized into two types: delayed contraction and direct inflation.

An example of a delayed contraction theory is fingering, or double diffusive, convection. The basic idea is that compositional gradients in the planet's interior can inhibit large scale convection. Indeed, double diffusive convection is thought to play a role in the interior of the giant planets in our Solar System, being invoked

to explain some anomalies in the heat transport on Uranus (Gierasch and Conrath, 1987), and proposed as being active on Jupiter and Saturn (Leconte and Chabrier, 2012) where it would have observable consequences.

The other subset of theories are direct inflation, where an additional source provides the energy for inflation. The actual energy required is relatively small, the radius of HD 209458b can be replicated by depositing only 1% of the incident irradiation into the convective region of the planet (See Figure 1.10). Tidal interactions have been suggested (Bodenheimer et al., 2001), but this cannot be the cause for the majority of cases as an additional large planet would need to be present in the system to prevent the orbit from circularising on very short timescales. This certainly is not the case for HD 209458b, as an additional massive companion would have already been detected in this well studied system, and if large hot Jupiter companions were common enough to explain the large inflated population they would be detected more regularly. One observational factor in favour of direct inflation theories is that there is a correlation between the irradiation level of the planet and its radius anomaly, as noted by Laughlin et al. (2011). This can be seen in Figure 1.10, where there is a clear tendency for the most inflated planets to be those with the highest incident irradiation. This is known to be an effect of temperature, and not orbital distance. Indeed, when effective temperature has been removed the radius actually increases as the semi-major axis increases. This may be an artifact of evaporation on closer in planets (Enoch et al., 2012).

A key piece of the radius anomaly puzzle will be detecting a statistical sample of planets around evolved stars, as this will allow the degeneracy between delayed-contraction and re-inflation models to be broken (Lopez and Fortney, 2015). If delayed contraction models such as Double diffusive convection are correct, we should expect that in old systems the originally cold planets beyond the snow line will have reached their final radius, and will not inflate appreciable when heated up by the evolving star. On the other hand, if re-inflation models are correct, then these formerly cold planets will begin to display the radius inflation effect after they are heated.

A popular re-inflation model is ohmic dissipation (Batygin and Stevenson, 2010), though questions remain as to whether the ionisation levels in the atmosphere can be sufficiently high to sustain the necessary currents. Alternatively, it is possible that strong winds (Guillot and Showman, 2002) could transport energy from the planet’s radiative atmosphere to the deep adiabat, where it can effect the bulk structure of the planet. However, at present it is not clear that winds that span such a large range of pressures can form in an atmosphere. Another suggestion, which

resurrects the idea of tides, and which would fit well with the observed irradiation dependence, is that *thermal* tides could be raised, where the local scale height is increased by the high temperatures, changing the vertical mass distribution, and these would be offset from the sub-stellar point by strong winds. This would provide the torque for a gravitational interaction. There is some potential support for this idea in that the hot-spots of hot Jupiters in secondary eclipse are indeed commonly observed to be offset (e.g. Knutson et al., 2007a). The method of directly detecting planetary winds I demonstrate in Chapter 5 could provide constraints for these theories, and help move us a step closer to solving the radius inflation problem.

1.3.3 The occurrence of hot Jupiters

It's important again to emphasize that the dominance of hot Jupiters in the picture so far is due to the relative ease of detecting and characterizing them due to their large dimensions, and the short orbits of the hot Jupiters. In fact, these planets are comparatively rare, occurring around only about 1% of stars. For comparison the occurrence rate of Super Earth sized planets per star appears to be at least 50% (Mayor et al., 2011), and the rates of Super Earth and Earth sized planets around cool stars may be greater than 1, Dressing and Charbonneau (2015) find a cumulative planet occurrence of 2.5 1-4 R_{\oplus} planets per M star in the *Kepler* sample.

Intriguingly there is a significant discrepancy between the occurrence rates reported from different survey techniques. Radial velocity surveys, headed by The Elodie Planet Search found a hot Jupiter occurrence rate of $0.7 \pm 0.5\%$, the California Planet Survey finds an occurrence rate of $1.2\% \pm 0.38\%$ with the Keck and Lick observatories (Wright et al., 2012), which is consistent with the value of $0.89\% \pm 0.36\%$ found by the HARPS survey (Mayor et al., 2011). Transit surveys, on the other hand, typically find a significantly lower occurrence rates. Results from *Kepler* gave an occurrence rate of $0.4\% \pm 0.1\%$ (Howard et al., 2012) and with a slightly improved sample $0.43\% \pm 0.05\%$ (Fressin et al., 2013).

The results of ground based transit surveys are consistent with the values of *Kepler*. SWEEPS finds a value of $\sim 0.42\%$ (Sahu et al., 2006). The OGLE-III survey finds $0.14^{+0.15}_{-0.08}\%$ occurrence for planets on orbits between 1 and 3 days, and $0.31^{+0.42}_{-0.18}\%$ for orbits between 3 and 5 days, so about 0.5% in total (Gould et al., 2006). Finally, SuperLupus find a rate of $0.10^{+0.27}_{-0.08}\%$ (Bayliss and Sackett, 2011).

We see that on average, radial velocity surveys find hot Jupiter occurrence rates that are over twice as high as those produced by transit surveys. It is not obvious why this should be the case, since the observational biases of both techniques are well understood, and the statistics have been built up over a decade by multiple

independent surveys. The implication is therefore that there is a subtle difference in the populations observed by the two techniques, and therefore that the population of hot Jupiters varies across the Galaxy. A systemic difference in metallicity has been suggested as a cause, as stars with higher metallicity are well known to be much more likely to host a hot Jupiter (e.g. Gonzalez, 1997; Fischer and Valenti, 2005; Wang and Fischer, 2014). This is in line with the expectations of core accretion theory, where large cores must be built before triggering runaway accretion to create a giant planet, this is much less likely in less metal rich discs. Though note, this metallicity correlation does not extend to smaller planets (Sousa et al., 2008), although it has been suggested that their occurrence saturates beyond a certain metallicity threshold, and that a similar power law could exist below that value (Zhu et al., 2016). There is also known to be a weak anticorrelation between metallicity and close stellar companion rates (Grether and Lineweaver, 2007; Raghavan et al., 2010), which reinforces the idea that giant planet formation follows a fundamentally different pathway to star formation. A metallicity difference of just 0.15 dex between the transit and radial velocity population may be enough to explain a factor of 2 difference in the hot Jupiter occurrence rate, however, a comparison of the metallicities of the samples of the *Kepler* and California Planet Survey samples by Wang et al. (2015) found an average 0.05 dex difference in metallicity between the samples, with a potential 0.05 dex systematic offset, which may lead to a difference in hot Jupiter occurrence of only 35%.

Another potential solution for the discrepancy is that the *Kepler* hot Jupiter rate is underestimated due to dilution from blending scenarios, where giant planets become misidentified as small planets. This could potentially boost the occurrence rate by a further 10%, however, recent improvements to the *Kepler* planet validating system make this solution unlikely (Morton et al., 2016).

1.3.4 Transmission signal

Transmission spectroscopy is a method where the planet transit is observed in multiple wavelengths. As atmospheric species are expected to have a wavelength dependent absorption profile, the planet will exhibit transit depth variations as a function of wavelength. It is an important tool for characterizing the atmospheres of exoplanets, and is the subject of two chapters in this thesis (5, 6), and will also be described in more detail in 1.4.1. It is useful to determine which out of the known population of exoplanets are expected to be the best targets for atmospheric characterisation, in order to choose the best targets to apply for telescope time on. We require an objective classification of the ‘observability’ of a planet’s atmosphere.

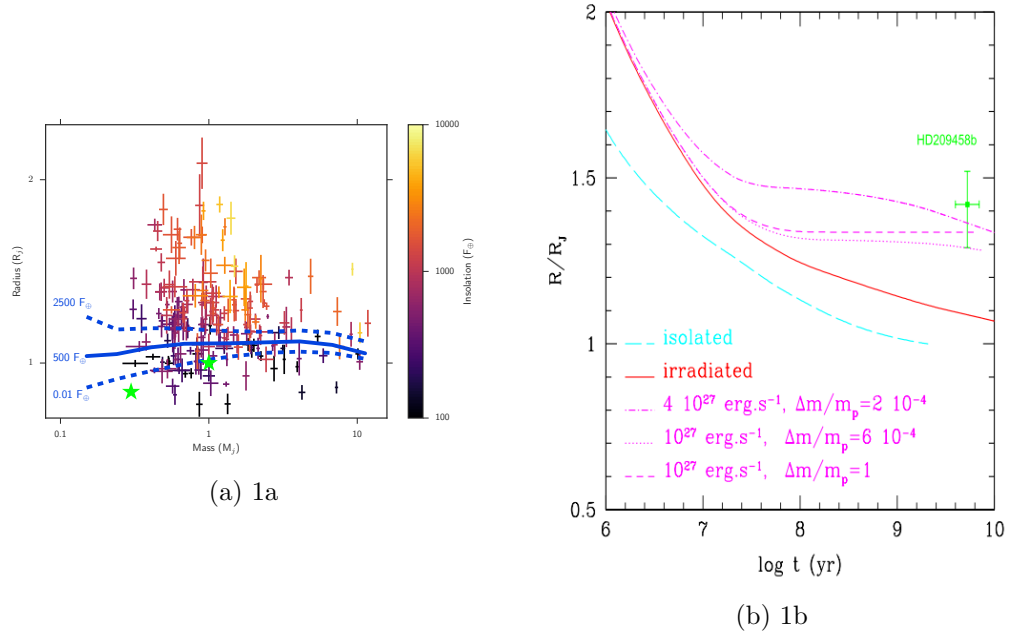


Figure 1.10: Left: mass-radius diagram of planets greater than $100 R_{\oplus}$, planets are shaded based on their incident irradiation. The model mass-radius relationships of Fortney et al. (2007) for a pure H/He planet at 4.5 Gyr are overplotted for several different flux levels. A large population of planets exist that are significantly larger than their expected radii, and the discrepancy is more severe for planets with higher irradiation. For comparison, Jupiter and Saturn are plotted as green stars. Right: Figure from Baraffe et al. (2003). The radius of HD 209458b cannot be explained by conventional planet evolution models, and requires additional energy input into the convective region of the planet. The additional required energy is only of order 1% of the irradiation when deposited in the planet's deep interior.

Fundamentally, the signal we are interested in is a difference in occulted area as a function of wavelength. If we imagine comparing the observed disc of the planet at two wavelengths, the difference in occulting area will be the annulus of the optically thick atmosphere. The difference in heights where the atmosphere becomes optically thick will be the width of the annulus. Features in an atmosphere may have different strengths, but will typically only cover a small number of “scale heights”, H , which is the characteristic pressure decay length scale of an atmosphere. For the purposes of comparing the expected strength of similar features across different planets, only the proportionality matters, so a single scale height is assumed. The scale height is defined as

$$H = \frac{k_B T_{eq}}{\mu g} \quad (1.16)$$

where k_B is the Boltzmann constant, T_{eq} is the equilibrium temperature of the planet’s atmosphere, μ is the average molecular weight of the atmosphere and g is the surface gravity of the planet. Brown (2001), neglecting the squared term in H gives the fractional increase in area from the occulting annulus as

$$\frac{\delta A}{A} = \frac{2\pi R_p H}{\pi R_*^2} \quad (1.17)$$

so

$$\frac{\delta A}{A} = \frac{2R_p^2 k_B T_{eq}}{\mu G M_p R_*^2} \quad (1.18)$$

Brown (2001) go on to show that the flux ratio as a function of wavelength goes as

$$\frac{\delta F_2}{F_2} \frac{\delta F_1}{F_1} = \frac{\delta A}{A} \log \frac{\sigma_1}{\sigma_2} \quad (1.19)$$

where sigma is the opacity at the two wavelengths This is the underlying principle behind the construction of a *transmission spectrum*, which is the flux ratio as a function of wavelength. For a molecular line the change in opacity from the wings to the core is of order 10^4 , which will make a typical spectral feature a signal of about 1 part in 10^3 in each wavelength bin, with sufficient resolution to resolve the feature.

This does make the assumption that the scale height for the absorbing species matches that of the bulk of the atmosphere. For well mixed molecular components such as H_2O this should be close to correct, but it is not necessarily the case for heavy cloud particles, which could settle and have a small fractional scale height. It is also not the case for an atmosphere with a strong thermal inversion, where the temperature and as such the scale height can become much greater at larger alti-




Figure 1.11: A plot to demonstrate the detectability of exoplanet atmospheres. Same color scheme as Figure 1.3. The dashed lines are equal S/N contours, spaced by factors of 10 with respect to the signal of HD 189733b’s atmosphere. Several important planets are highlighted. The use of Visual magnitudes in this plot de-emphasises the molecular signals achievable for very red stars such as GJ 1214b

tudes. This can significantly increase the slope of broad band features like Rayleigh scattering (Pont et al., 2013).

Using Equation 1.18, I plot the expected transit signal of every known transiting planet vs the V band magnitude of the host star in Figure 1.11. Note that in all of these cases, we have assumed a pure hydrogen helium atmosphere with reduced molecular mass 2.3. We do not know the typical composition of the atmospheres of super-Earths and mini-Neptunes, so it is possible that their signal will be significantly smaller than shown here due to a reduced scale height.

We see that despite the large number of candidates, the relative faintness of the *Kepler* planets limits their usefulness for transmission spectroscopy. A limitation to the lines of constant S/N in this plot is that at some point the systematic noise sources in the instrument, or the fluctuations in the sky-conditions will become significantly larger than the photon noise. So in reality, there is a point where

increasing the brightness of the star no longer improves the signal, but since this is difficult to quantify and depends on the exact observing instrument in question we do not attempt to include it in this plot.

This metric is useful for comparing the expected strength of signals across planets, but there is an important subtlety to be discussed; the atmospheric temperature will not be constant across the whole planet, especially if heat transport is inefficient. This could lead to atmospheric signals being stronger on one side of the planet than the other. Indeed, taking this a step further, Kataria et al. (2016) suggest that the chemical composition could be noticeably different between the two sides of the planet, which could lead to observable differences between the two hemispheres. Approaching this from a different angle, I show in Chapter 5 that it is possible to distinguish between the leading and trailing edges of a transiting planet not by the strength of the signal, but by its velocity signature, induced by rotation and winds.

1.3.5 Availability of comparison stars

In practice, when it comes to target selection for transmission spectroscopy, additional concerns may come into play. For example, the availability of a nearby comparison star. For broad band spectroscopy from the ground a comparison star is required to remove the variable opacity from the Earth's atmosphere. This becomes less important for narrow features, which can be distinguished from telluric signals by their unique velocity signatures. (e.g. Snellen et al., 2010; Loudon and Wheatley, 2015; Brogi et al., 2016). In order that both comparison star and target star are both well exposed within the same image, they should ideally be close in magnitude. If one is observing with broad filters they should also preferentially be of a similar spectral type, otherwise the band-weighted response of the stars to changes in the atmosphere will not match and differential photometry will be inefficient. Perversely, this requirement for a close by comparison star biases target selection towards fainter objects. A 1 magnitude increase in brightness would be expected to decrease the density of stars in an average field by a factor of $\sim 60\%$, if stars were uniformly distributed. Of course, this is not the case, and the actual distribution of stars in the sky is much more complicated. We can improve this estimate by using galactic population models. Bahcall and Soneira (1984) construct a model of the Milky Way galaxy from a simple two component model, with a disk and a spheroidal component. This provided a good fit to star counts available at the time (e.g. McLaughlin, 1983), and is still used to model stellar number densities at visual wavelengths to estimate the performance of scientific instruments, such

as ULTRACAM (Dhillon et al., 2007). More complex models which contain both thick and thin disk populations (e.g. Gilmore and Reid, 1983) do not significantly differ from the predictions of the Bahcall-Soneira model for visual and submicron wavelengths for magnitudes < 20 (Simons, 1995).

The Bahcall-Soneira model gives the expected frequency of stars as a function of magnitude and galactic latitude, from this it is straightforward to calculate the probability that there will be at least two stars within the same magnitude range within a given area. As the distribution should follow Poisson statistics, it is trivial to invert this problem by using the complement probability of there being 0 comparison stars,

$$\begin{aligned} P(c > 0) &= 1 - P(c = 0) \\ P(c > 0) &= 1 - e^{-AN} \end{aligned} \tag{1.20}$$

where N is the number density of stars and A is the area. For this example I used the area of a circle with radius of 5 arcminutes, as this is the maximum usable length of the slit for ACAM on WHT. The results for three different galactic latitudes are plotted in 1.12. For a galactic latitude of 30 degrees, which is the whole sky average, and V magnitude of 10, only about 25% of stars have a suitable comparison within 5', but this fraction rises steeply to 80% by 12th magnitude. This prediction can be tested practically by using a catalog search to check for comparison stars for all currently known transiting planet hosts⁴. Using the UCAC4 catalogue (Zacharias et al., 2013) I searched for comparison stars within 5' and 0.5 mag of all known exoplanets, the results are shown in Figure 1.12.

When applying for telescope access, potential targets are selected from a ranked list based on their expected atmospheric signal, observability during an observing season, and presence of a comparison star. WASP-52b scored highly on all of these metrics, so was a prime candidate for transmission spectroscopy, I present a transmission spectrum of this planet in Chapter 6.

It is in fact possible in some circumstances to construct broad-band transmission spectra from the ground using the chromatic Rossiter-McLaughlin effect (Gloria et al., 2015) and chromatic line profile tomography (Borsa et al., 2016).

Transmission spectroscopy with *HST* does not typically use a comparison star. The systematics on *HST* are due largely to thermal breathing effects during the orbit and are quite repeatable, so can typically be removed by parametric models (Wakeford et al., 2016). Another consideration is that the field of view is relatively small for Hubble, and additionally, the scanning spectroscopy mode developed to

⁴accessed from NASA exoplanet Archive, <http://exoplanetarchive.ipac.caltech.edu/>, accessed 26/05/16

increase the Signal-to-noise ratio usually precludes the use of a comparison star. In some cases where planet hosts have close companion stars it is possible to perform differential spectroscopy, giving similar results to single target modes (Wakeford et al., 2013).

1.4 Observable properties of Exoplanet Atmospheres

1.4.1 Transmission spectroscopy

As soon as transiting planets were detected, efforts began to characterise their atmospheres. Predictions from multiple independent atmosphere modelling efforts agreed that the atmospheric transmission signal in the optical would be dominated by large, pressure broadened alkali metal absorption lines, with the sodium feature at 589 nm being the strongest (Seager and Sasselov, 2000; Brown, 2001; Hubbard et al., 2001). Indeed, measurements of the spectra of brown dwarfs with similar equilibrium temperatures had been known for some time to contain appreciable alkali metal signatures. GJ 229B was the first unambiguous detection of a brown dwarf (Nakajima et al., 1995), (coincidentally detected in the same year as the first unambiguous detection of an exoplanet, 51 Peg b), with a mass of between 20–50 M_J and an equilibrium temperature of 900 K. It is not terribly far removed from the parameters of a hot Jupiter, though of course, brown dwarfs are internally heated by their contraction, and hot Jupiters externally by extreme insolation. Oppenheimer (1998) report the spectrum of GJ 229B which has clearly visible sodium, potassium, lithium, rubidium and cesium lines, these features were also known to be present in other very low mass stars (e.g. Kirkpatrick et al., 1997). Interestingly, the K star companion to GJ 229B appears to also host a Neptune mass planet just outside the habitable zone (Tuomi et al., 2014).

There was both good theoretical and observational evidence to suspect that hot Jupiters might host atmospheres which would yield significant alkali metal absorption signatures during transit. Early attempts to detect the exospheres of 51 Peg b⁵ (Coustenis et al., 1997, 1998; Rauer et al., 2000) and the then recently detected HD 209458b (Bundy and Marcy, 2000; Moutou et al., 2001) were unsuccessful due to insufficient Signal to Noise.

The first successful detection of an exoplanet atmosphere was made by Charbonneau et al. (2002), who used the STIS spectrograph on *HST* to measure sodium absorption in HD 209458b. The signal however, was significantly weaker than model

⁵While the planet itself does not transit, the hope was that an inflated atmosphere may have a grazing incident transit




Figure 1.12: How likely is a star to have a comparison star within 1 magnitude of brightness, and within 5 arcmin. The dashed lines show the theoretical expectations from galactic population models for galactic latitudes of 90, 30 and 20. 30 is close to the average stellar density for the whole sky.

predictions would suggest.

One hypothesis was that the sodium abundance in the atmosphere was depleted on HD 209458b, at the time of formation or at some point afterwards. The giant planets in the Solar System display enhanced atmosphere metallicities with respect to Solar for most elements measured (e.g. Young, 2003; Guillot, 2009), which is thought to be a consequence of formation by core accretion (Matsuo et al., 2007) but sodium is conspicuously absent in Jupiter’s atmosphere, it is thought that it differentiated into the planet’s interior (Atreya et al., 2003). A similar process may operate on hot Jupiters, though at a lower efficiency as at those temperatures the alkali metals will typically be free and not bound up in heavier molecules. Alternatively, the sodium that is visible may have been delivered significantly after formation by asteroids and comets ablating in the upper atmosphere. Indeed, this is the source of the well known sodium layer in the Earth’s own atmosphere (e.g. McNeil et al., 1995; Plane, 1991, and references within).

Alternatively, the same reduction in signal could be due to a high altitude layer of cloud or haze. To reproduce the muted sodium signal measured on HD 209458b, assuming Solar metallicity, would require a layer of opaque cloud at 0.4 mbar (Charbonneau et al., 2002). As we shall see, clouds and hazes appear to be ubiquitous on exoplanets, as they are in our Solar System (e.g. Karkoschka, 1994), though their formation is poorly understood.

The detection of sodium on HD 209458b remains one of the most robust results in the study of exoplanet atmospheres, it has been confirmed by an independent re-analysis of the same dataset (Sing et al., 2008a), and excellent agreement was found with measurements from the ground with the High Dispersion Spectrograph (SDS) on the Subaru telescope (Snellen et al., 2008), and also the High Resolution Echelle Spectrometer (HIRES) on Keck (Langland-Shula et al., 2009) and the High Resolution Spectrograph (HRS) on the 9m Hobby-Eberly Telescope (Jensen et al., 2011). The HIRES measurements found a significantly stronger absorption than the other observations, by a factor of 4.7, this could be a signature of variability in the sodium layer of HD 209458b, or an indication of the difficulty in removing some systematic errors from the data.

The *HST* data reanalysed by Sing et al. (2008a) were of sufficient quality that the absorption profile of sodium could be modeled to extract atmospheric parameters (Sing et al., 2008b; Vidal-Madjar et al., 2011b,a), these authors find evidence of high altitude thermal inversion layer above 10^{-7} bar, as well as evidence that the sodium abundance varies as a function of altitude, between 2 and 0.2 times Solar levels. The observations also favour a clear atmosphere, with a hint of enhanced opacity

blueward of 4000 Å due to Rayleigh scattering. Désert et al. (2008) argue that the enhanced opacity at ~6250 Å may be a signature of TiO and VO absorption, though cold-trapping may remove both of these molecules from the atmosphere. These observations demonstrate the significant amount of information on exoplanet atmospheres that can be gained by detailed study of individual lines and high quality transmission spectra.

A thermal inversion layer was also initially claimed to be present in thermal emission data from *Spitzer* for HD 209458b (Knutson et al., 2008), but later analysis with new and archival data showed that this probably was not the case (Line and Yung, 2013; Diamond-Lowe et al., 2014; Zellem et al., 2014).

Clouds and hazes in exoplanet atmospheres are becoming an increasingly important topic of research. Clouds and hazes are distinguished by their formation mechanism ‘clouds’ are formed by material condensing directly out of the gas phase into droplets, whilst a ‘haze’ implies that some form of chemistry, typically photochemical, is responsible for the formation of small droplets of the produced material. Typically, haze droplets are much smaller than cloud droplets, which droplet sizes smaller than a micrometre, so in the optical their distinguishing observational feature is a Rayleigh scattering slope, while clouds present a flat, gray opacity source, though perhaps with some molecular features in the far infrared.

Clouds and hazes are detected, and play important roles in every body with an atmosphere in the Solar System, from the cloud belts on Jupiter to the thick hydrocarbon hazes on Titan - so it should come as no surprise that they appear to be common features of exoplanet atmospheres.

Several high profile exoplanets have been shown to have atmospheres dominated by hazes or clouds, notably, the mini-neptune GJ 1214b, which has an extraordinarily flat transmission spectrum, believed to be due to clouds (Kreidberg et al., 2014). The super earth HD 97658b also displays a flat transmission spectrum, indicative of clouds (Knutson et al., 2014). In Chapter 6 I show that the hot Jupiter WASP-52b has a mostly flat spectrum, that is probably best explained by a thick cloud layer.

A survey of 8 hot Jupiters observed with *HST* (Sing et al., 2015), showed that exoplanets exist on a continuum between clear and cloudy. The understanding of cloud formation and composition, and their impact on other observable features, will therefore be an extremely important topic to the development of the field (Wakeford and Sing, 2015).

HD 189733b, discovered in 2005 (Bouchy et al., 2005) remains the planet with the highest theoretical atmosphere signal (see Figure 1.11), due to its deep

transit and bright K-class host star. Observations of this planet have been fruitful, and yielded some particularly interesting results. It was soon found that similar to HD 209458b, this planet had a strong sodium absorption signature by Redfield et al. (2008) and Jensen et al. (2011) with HRS on the Hobby-Eberly Telescope. However, the signal was found to be at least 3 times as strong as on HD 209458b, hinting that the atmospheres of these two iconic planets could be quite different. It soon emerged that the broad-band transmission spectrum of HD 189733b was dramatically different to that of HD 209458b, being dominated by a steep Rayleigh scattering slope over the entire optical and near-infrared spectrum, as seen with multiple instruments on *HST* (Pont et al., 2008; Sing et al., 2009; Gibson et al., 2011, 2012). The spectrum eventually flattens out to a cloud deck with no strong molecular features in the infrared, as measured by the *Spitzer* 3.6, 4.5, 5.8 and 8 μm channels (Désert et al., 2009, 2011; Knutson et al., 2012; Agol et al., 2010). The state of the broad band photometry is summarized by Pont et al. (2013) (see Figure 1.13), and is a clear demonstration that wide wavelength coverage with multiple instruments is required to begin to understand even the best exoplanetary target.

Although the planet is covered by haze and clouds, narrow line cores formed at low pressures and high altitudes are still accessible, but require high resolution measurements. Birkby et al. (2013) detect water in the spectrum of HD 189733b by cross-correlating a molecular absorption template with secondary eclipse measurements made with CRIRES on the VLT, and using the same technique Brogi et al. (2016) detect both CO and H₂O. This technique was first used by Snellen et al. (2010) on CO in HD 209458b and led to the detection of high altitude winds on the planet, as well as an absolute solution to the orbital motion of the system, by effectively making it a double lined binary system. This technique can even be used on the atmospheres of non-transiting planets, and has been demonstrated on τ Bootis b (Rodler et al., 2012; Brogi et al., 2012) and the first detected exoplanet, 51 peg b (Martins et al., 2015).

Huitson et al. (2012) detected the sodium signal in medium resolution data of HD 189733b with data from *HST* STIS, but was unable to find the broad components due to the aforementioned haze layer. The sodium signature of the planet was resolved however, and from this a basic Temperature Pressure profile could be constructed. Similarly to HD 209458b, a high altitude thermal inversion layer is implied, with temperatures reaching over 1000 K higher than the equilibrium temperature of the planet (~ 1400 K). This is in fact an expected result for a planet with high degrees of optical scattering in the upper atmosphere. Heng et al. (2012) developed parameterized models of temperature-pressure profiles which include the

effect of clouds and hazes, they found that an “anti-greenhouse” effect exists for atmospheres with high levels of shortwave scattering, such as that produced by a high altitude haze, whereby the lower atmosphere can be cooled significantly below the equilibrium temperature, and the high atmosphere can be heated. I use analytical T-P profile formalism of Heng et al. (2012) in Chapter 6 to produce model transmission spectra for WASP-52b.

Wyttenbach et al. (2015) showed that it was possible to detect the line core of sodium in archival data from HARPS, and in Chapter 5, (Louden and Wheatley, 2015) I will show that these same observations can be used to spatially resolve the atmosphere of HD 189733b, providing evidence of strong planetary winds and rotation. Crucially, I also implicitly account for the Rossiter-McLaughlin effect implicitly in my model, which was likely the cause of some of the systematic effects seen in earlier observations, (see Redfield et al., 2008; Jensen et al., 2011; Wyttenbach et al., 2015). Brogi et al. (2016) independently came to the same conclusion on the necessity of modeling the Rossiter-McLaughlin effect. The combination of high spectral resolution with a time resolved model provides an exciting new step in characterizing the atmospheres of exoplanets.

1.4.2 Atmosphere erosion

The atmospheres of planets are not fixed and static things, but are in constant flux. To take the Earth as an example, the atmosphere we enjoy today bears virtually no resemblance to that of the young Earth. If we were to visit the Earth shortly after its formation, we would see it shrouded in hydrogen and helium from the protostellar nebular. This was quickly lost to space, and replaced with an atmosphere outgassed from the planet’s interior and water vapor delivered by asteroid impacts. The Earth, of course, has ended up with a “tertiary” atmosphere rich in oxygen as the result of biological processes.

From the Earth, we can directly observe the process of atmosphere evolution through the evaporation of planets. It was quickly realised that the high levels of irradiation of the earliest detected hot Jupiters could lead to them having greatly extended exospheres. Guillot et al. (1996) extended their theory of giant planet atmosphere to the case of the then recently discovered 51 Peg b and found, assuming that the stellar characteristics were similar to that of the Sun, that the planet was stable against catastrophic photo-disintegration.

For the hot Jupiters, the rates of evaporation over the lifetime of the planet are small, and typically add up to only a few percent of their total mass (e.g. Sanz-Forcada et al., 2011), but for the intermediate mass planets on short orbits,

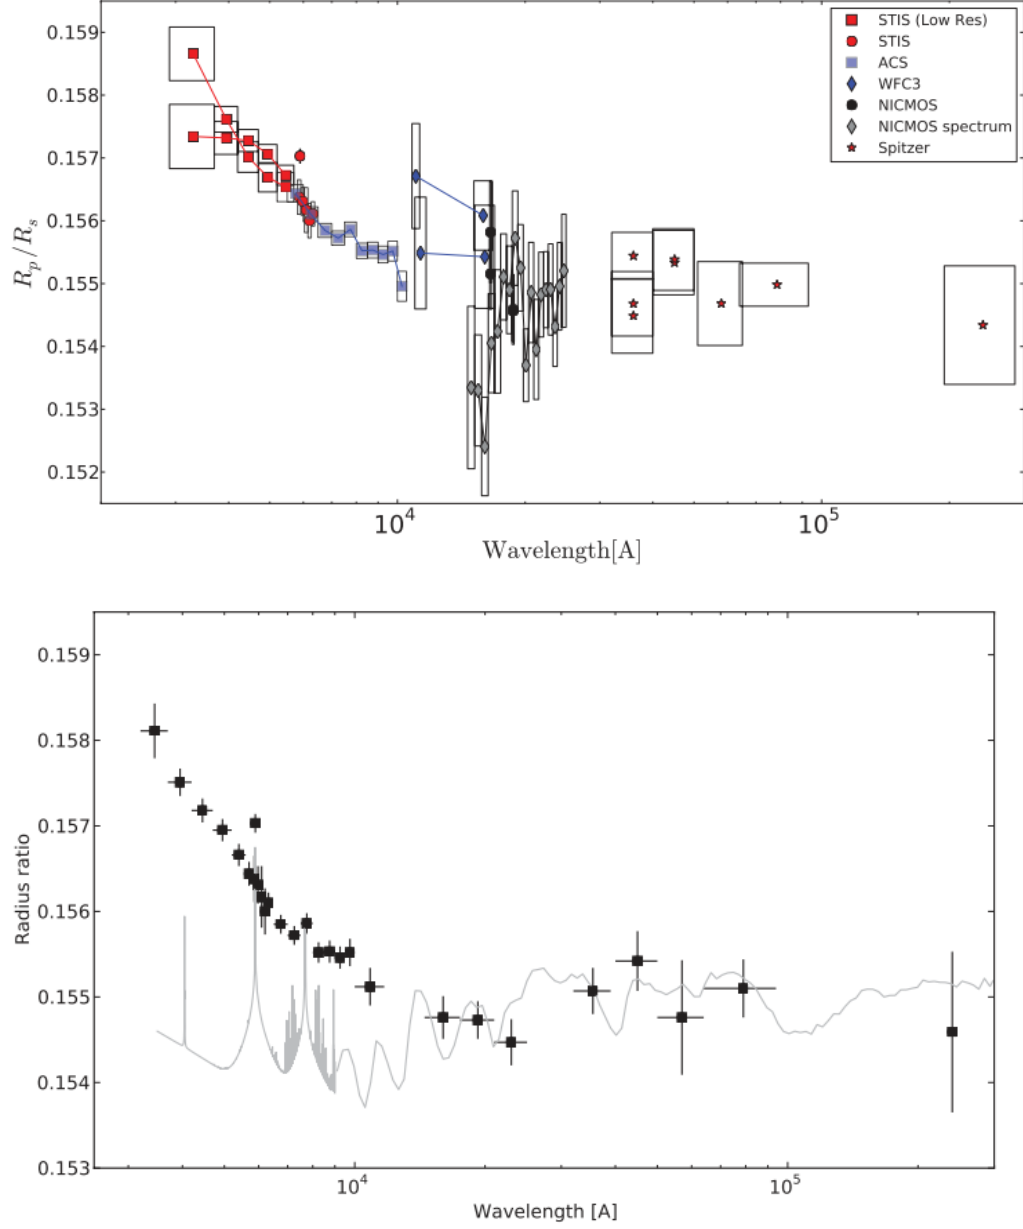


Figure 1.13: Dominated by haze top: The instruments required to build the full picture of HD 189733b's atmosphere (Pont et al., 2013) Bottom: Spectrum of HD 189733b compared to a clear model atmosphere.

that seem to be particularly common, evaporation may be a key sculptor of the observed population. Looking ahead to the characterisation of smaller planets, particularly habitable ones, evaporation may literally be a case of life or death. Too little evaporation may lead to surfaces crushed by the weight of a significant surviving hydrogen envelope, making liquid water impossible (Owen and Wu, 2016), and we have seen the effect of too much evaporation with the fate of Mars, and the atmosphere of Venus is currently evaporating (Brace et al., 1987).

The most straightforward way for a particle to leave the atmosphere is simply for it to have a velocity higher than the local escape velocity, at this point the particle can vanish off into space with no impediment. The velocity distribution of species in an atmosphere are described by the Maxwell-Boltzmann distribution

$$f(v) = \sqrt{\left(\frac{m}{2\pi k_B T}\right)^3} 4\pi v^2 e^{-\frac{mv^2}{2k_B T}}. \quad (1.21)$$

The Maxwell-Boltzmann distribution for oxygen and hydrogen is plotted in Figure 1.14. For oxygen at surface temperatures, there is effectively no chance of a molecule reaching escape velocity. Hydrogen on the other hand, being a much lighter molecule has a significantly higher mean velocity, and the high end of the distribution reaches much further. In fact, the exosphere is quite a bit hotter than the surface (~ 1500 K), so the distribution is wider still. When this is accounted for a non-trivial fraction of hydrogen atoms at the exosphere will have velocities higher than the escape velocity.

The exobase is an important transition point in the atmosphere, moving into the “exosphere”, which is the region where molecules are in principle still gravitationally bound to the planet, but the density is low enough that the medium is essentially collisionless. As a result, normal gas laws that dominate the rest of the atmosphere no longer apply. The exobase can be defined as the point where the scale height is equal to the mean free path of a molecule, so that a particle traveling through this region would be expected to only have one collision on average. Recalling the definition of a scale height in Equation 1.16,

$$H = \frac{k_B T_{eq}}{\mu g} \quad (1.22)$$

then the exobase is the region where the following condition is met,

$$H n \sigma_c = 1 \quad (1.23)$$

where n is the total number density, and σ_c is the collision cross-section. Note that




Figure 1.14: The maxwell Boltzmann distribution for an oxygen molecule in the Earths atmosphere at 287 K (blue) and for hydrogen at 287 K (green) and for hydrogen at exosphere temperatures, 1500 K (orange). The escape velocity of the Earth (11.2 km/s)is marked as a dotted line.

both H and n are dependent on the atmospheric temperature. At these altitudes the temperature is controlled largely by the Solar activity level, so the exobase can vary between about 500 and 1000 km. Due to heating from the Solar atmosphere, exosphere temperatures are of order 1000 K.

Starting from the Maxwell-Boltzmann distribution, one can integrate to find the fraction of molecules in the high tail with velocities greater than the escape velocity, and from there the total escaping flux of particles can be calculated (JEANS, 1916),

$$F_J = n_c \frac{V_{th}}{\sqrt{4\pi}} (1 + \lambda_c) \exp(-\lambda_c) \quad (1.24)$$

where V_{th} is the most probable velocity from the Maxwell-Boltzmann distribution, which is

$$V_{th} = \left(\frac{2k_B T}{m} \right)^{\frac{1}{2}} \quad (1.25)$$

and λ_c is the “escape parameter”, which is the square of the ratio of the escape velocity to the thermal velocity, where the escape velocity is

$$V_{esc} = \left(\frac{2GM}{R} \right)^{\frac{1}{2}} \quad (1.26)$$

so Equation 1.28 can be written as

$$F_J = n_c \frac{V_{th}}{\sqrt{4\pi}} \left(1 + \frac{k_B T R}{G M m} \right) \exp \left(-\frac{k_B T R}{G M m} \right). \quad (1.27)$$

Note in particular, that the escape rate depends on mass and temperature with an exponential term, so is both extremely sensitive to temperature and also comparatively inefficient for larger molecules. Recall also that for hot Jupiters, R is approximately constant with respect to mass, so the Jeans flux will scale exponentially with the flux. All other things being equal, a $0.5 M_J$ hot Jupiter can have a Jeans escape flux 100 million times higher than a $10 M_J$ planet.

In Jeans' time, the high temperature of the exosphere was not known, so the calculated loss rate implied that the Earth could not lose a significant fraction of gas through these means. Jeans hypothesized that since hydrogen was rare in the atmosphere, the Earth must have been hotter in the past. Analogously, when the mass loss rates of hot Jupiters are compared to the calculations from Jeans escape, they are much too low, unless the exosphere of these planets is correspondingly heated to very high temperatures.

The mass loss rate of the hot Jupiter HD 209458b was first measured by Vidal-Madjar et al. (2003), and later by (Linsky et al., 2010) to be a minimum of 10^{10} g/s, by observing UV absorption of metal lines with depths that implied material from the planet extended over the Roche lobe. (See discussion in Chapter 4)

In fact, the mass loss rates on hot Jupiters can become so high that they can reach a state of *hydrodynamic outflow* (e.g. Owen and Jackson, 2012). In this situation, heavier elements have been observed to be carried off of the planet, including C, Si and Mg (Linsky et al., 2010; Vidal-Madjar et al., 2013). This is strong evidence that a transition to hydrodynamic escape has occurred, due to the exponential dependence on mass in the Jeans escape rate.

Optical light from the photosphere cannot heat the atmosphere to the required temperatures of $\sim 10,000$ K (it would violate the second law of thermodynamics) so the heating is presumably caused by high levels of X-ray and EUV flux from the star. In Chapter 4, I calculate the amount of high energy radiation available to heat the atmosphere of the evaporating hot Jupiter HD 209458b.

Non-thermal effects are also important for planets in the Solar System, for example, interactions with the Solar wind are believed to have stripped the atmosphere of Mars. For hot Jupiters, interactions with the stellar wind could shape the outflowing gas (Bourrier and Lecavelier des Etangs, 2013) and act to remove

material that has been lifted above the magnetosphere of the planet.

The study of mass loss has significant implications for life outside the Solar System, considering that many of the most habitable planets detected to date orbit M dwarfs, which are particularly active (Armstrong et al., 2015).

1.4.3 Thermal emission and atmosphere dynamics

Transmission spectroscopy relies on the light of the star being blocked by the thin annulus of the planet’s atmosphere. A separate and complementary set of techniques relies on the emitted thermal flux of the planet itself; these are the secondary eclipse and thermal phase curve measurements. Because these photons are directly emitted by the planet itself, these methods give an insight into the thermal structure and heat transport within the planetary atmosphere.

Analogous to the primary transit, the thermal flux of the planet can be blocked out as it passes behind the star. Note that the presence of a primary transit does not necessarily mean that a secondary transit will occur, and vice versa, due to non-zero eccentricities. Approximating both the star and the planet as black body radiators, the expected occultation signal is

$$\frac{\delta F}{F} = \left(\frac{R_p}{R_*} \right)^2 \frac{B_\lambda(T_p)}{B_\lambda(T_*)} \quad (1.28)$$

where B_λ is the Planck function, and T_p and T_* are the equilibrium temperatures of the planet and the star, respectively.

A diagram of this expected contrast ratio against J band magnitude is shown in Figure 1.15. In contrast to the optical transmission spectroscopy diagram shown in Figure 1.11, HD 209458b and particularly HD 189733b have significantly stronger signals than any other exoplanet. However, the absolute signal strengths are typically higher than found in transmission spectroscopy for all planets.

The thermal emission of an exoplanet was detected for the first time using *Spitzer* for TrES-1b by Charbonneau et al. (2005) and for HD 209458 by Deming et al. (2005).

A carbon rich chemistry was inferred for WASP-12b by Madhusudhan et al. (2011), by combining J H and K data from ground based observatories with *Spitzer* data of the secondary eclipse. Their analysis suggested a deep carbon monoxide signature in the dayside spectrum, but very low, or no absorption from water. Their atmospheric recovery suggested a scenario where the carbon to oxygen ratio was greater than 1, whereas C/O is only 0.54 in the Sun (Asplund, 2005). Studies of planetary debris on white dwarfs suggest that such C/O ratios are uncommon

(Wilson et al., 2016). More recent results from Kreidberg et al. (2015) do detect a very significant water signal in the atmosphere of WASP-12b, and with a new method of atmospheric fitting, they showed that C/O ratios > 1 were ruled out at greater than 3σ significance by these new data. For the time being, carbon planets, while in principle very interesting objects, have no observational evidence of their existence.

As well as the secondary eclipse, if the thermal emission is not emitted evenly from the planet’s surface (and due to the tidal locking, the extreme flux contrast between the two hemispheres makes this seem unlikely), there will be a thermal phase curve, as the orbit of the planet changes the portion of the emitting surface that is observed from the Earth, transitioning from being almost entirely the permanent night side close to the primary transit, to almost entirely the day side close to the secondary transit.

The first detection of a planet’s thermal phase curve was made by Harrington et al. (2006) using the *Spitzer* space telescope. The 24 micrometer light curve for ν Andromedae b showed that there was indeed a variation in flux between the day and the night side, indicating that the timescale for heat transport was comparable to the radiative timescale in the atmosphere. Other hot Jupiters show less significant flux differences between the day and night sides of the planet (e.g. Cowan et al., 2007; Zellem et al., 2014), with indications that the heat transport efficiency may be proportional to the planet’s irradiation level.

Knutson et al. (2007a, 2009) found that the maxima of the thermal phase curve of HD 189733b was significantly offset from the secondary eclipse time, and were able to construct a longitudinal map of the brightness temperature of the planet (see Figure 1.16). They found that the brightest point of the planet’s atmosphere was offset from the substellar point by approximately 30 degrees. This was interpreted as evidence of strong equatorial winds blowing around the equator of the planet, with speeds of up to 10 km s^{-1} which was an early prediction of General Circulation Models of hot Jupiters (Showman and Guillot, 2002). The equatorial jet is a natural consequence in highly irradiated, tidally locked atmospheres, and is caused by planetary scale Rossby waves interacting with the day to night flow caused by the extreme temperature contrast between the two hemispheres (Showman and Polvani, 2011).

It was predicted that these strong winds would also imprint signatures on the transmission spectra, through Doppler shifting the transmission lines (Showman et al., 2013; Miller-Ricci Kempton and Rauscher, 2012). In addition to an equatorial jet, these models predicted an average blue shift to the planet’s atmosphere from




Figure 1.15: A plot of J magnitude against the expected thermal eclipse signal for transiting planets. Planets detected through radial velocity are in pink, through ground based transit in blue, and space based transit in yellow. Planets with radii smaller than $0.1 R_J$ are plotted as triangles. Several important exoplanets are labeled, exquisite phase curves of a super Earth recently (Demory et al., 2016a), (Demory et al., 2016b). The signal sizes are significantly larger than transmission spectroscopy (though the systematics in *Spitzer* more than make up for it), also not that unlike in transmission spectroscopy, HD 189733b is significantly better than any other target.




Figure 1.16: The *Spitzer* lightcurve and inferred temperature distribution of HD189733b from Knutson et al. (2007a). The hottest point of the atmosphere is clearly offset from the substellar point, indicating planetary winds.

a polar wind. Measuring these wind speeds directly through transmission spectroscopy, instead of indirectly with thermal phase curves, in principle allows direct constraints on theories of heat transport in planets to be made.

Redfield et al. (2008) found a very high average blueshift in the signal of sodium in HD 189733b, though they suspected that this was a systematic effect due to the unphysically high velocity, which was significantly higher than the speed of sound. In Chapter 5 I will show that this was likely caused by not accounting for the Rossiter-McLaughlin effect on the stellar absorption lines. Using high dispersion spectroscopy, Snellen et al. (2010) was able to measure the average windspeed of HD 209458b, and found that there was an average blueshift of $2 \pm 1 \text{ km s}^{-1}$. In Chapter 5 I find a very similar average wind velocity in the atmosphere of HD 189733b, $1.9^{+0.7}_{-0.6} \text{ km s}^{-1}$, but I will also show that with high dispersion spectroscopy it is possible to *spatially resolve the atmosphere*. This allows direct measurements of wind speeds on the two limbs of the planet, and is no longer limited to being a single velocity for the whole planet. Refinement of this technique, and application to other exoplanets, will enable tests of the predictions of heat transport in exoplanet atmospheres.

1.5 Conclusions

In this Chapter I have given an introduction to the detection and characterization of extrasolar planets, and put recent results in the context of the known population of exoplanets. I have pointed out key areas where there is room for development, and indicated my own contributions to improvements of knowledge in these areas. I have described the parameters that lead to a planet being a good candidate for atmospheric follow up, and in the following chapters I shall describe the methods by which I have done this.

Chapter 2

Methods and Techniques

There is this TREMENDOUS
MESS of waves, all over in space,
which is the the light bouncing
around the room, and going from
one thing to the other

Richard Feynman – Fun to imagine

2.1 Introduction

In this Chapter I will introduce the main instruments, and provide additional background on the methods and theory used in Chapters 3-6.

This Chapter is organized as followed:

In Section 2.2 I will give a brief introduction to the functioning of the CCD, followed by a general overview of spectroscopy in 2.3. From there, I will give an overview of the instruments from which data in this thesis was taken: ACAM on the WHT (2.4), HARPS (2.5), HST COS (2.6) and XMM-Newton (2.7). I will then discuss Bayesian statistics, and it's application through MCMC and Gaussian Processes in Sections 2.8 and 2.9.

2.2 CCD

It is fair to say that the Charge-Coupled Device (CCD) has revolutionized astronomy. First developed in 1969 at Bell Laboratories, the potential for CCD use in astronomy was immediately obvious, but it did not fully supplant traditional photographic plates until the late 80's (see Figure 2.1). The advantages of modern CCD's

are numerous, the most important being their high quantum efficiency, providing very high low light sensitivity, and their linearity over a wide range of intensities, making it easy to compare the brightness of many different objects. There are practical advantages also, removing the need to prepare and develop photographic plates, which must be stored carefully to prevent degradation.

CCD's are composed of a thin wafer of p doped silicon on top of a substrate material. The basics of a CCD's function can be understood by an analogy, presented in Figure 2.2. In the bucket analogy, each pixel is represented by a bucket arranged in a regular grid, which are "filled" by photons (here the raindrops). A photon interacts with the silicon, generating a electron-hole pair. The doped silicon is designed such that the band-gap energies correspond to visible light.

Traditional CCD's are not directly sensitive to UV light, one solution is to cover the CCD with a phosphor which fluoresces under UV illumination, generating optical photons which can then interact with the silicon. However, this phosphor can be lost over time, compromising the detector wavelength sensitivity. The UV instruments on the HST; COS and STIS, also have MAMA detectors, which are serial photon counting devices. Unfortunately, these detectors also have been seen to lose gain over time, causing the sensitivity of the instrument to seriously degrade. Therefore the lifetime of UV instruments are often strongly limited.

"Saturation" occurs when the charge on a pixel becomes greater than the well depth, which can cause charge to overflow into surrounding pixels and electronics. This is not always necessarily the end of the world, depending on the CCD design, it can still be possible to perform photometry on saturated sources, so long as the extraction region is expanded to capture all the charge. This does not work if the charge in the sensor is not conserved, i.e., by charge escaping into the readout electronics, or into regions of the CCD which are not read out. In fact, *Kepler* has made use of photometry on saturated sources, with good success (Gilliland et al., 2010).

The "buckets" in the analogy are of course a simplification. The photo-electrons are held in place by a positively charged "gate", which is separated from the silicon by a thin insulating layer (typically glass so it is transparent to the incoming photons). Charge can then be transferred from one well to the next by "clocking" the gate voltages - increasing the charge on the neighboring well causes the electrons to be transferred. Doing this in a finely timed series moves each column in turn to the readout column (which has not been exposed to light), where the same process takes place in the vertical direction, feeding charge to the Analogue to Digital Converter (ADC). Before entering the ADC, the small charges from the pixels must




Figure 2.1: Mentions of “CCD” verses “photographic” on the NASA ADS abstract service. CCD mentions do not overtake photographic mentions until the late 80’s.

be converted to a voltage and amplified to a level that the ADC can register using a capacitor. The amplification process is not perfect, and so contributes to the “readout noise” of the system. The amplifier also adds a constant “bias level” to the signal which must be removed in reduction. The purpose of the bias level is to keep all voltages positive, so that the data can be stored as unsigned integers to save storage space. Finally the ADC, as it’s name suggests, converts the analogue voltage from the amplifier to a quantized output value (counts) that can be stored as digital data. The number of electrons it takes to generate one count is called the gain. In this way the charge of each pixel on the chip is measured sequentially, and since the order they were read out in is easy to recreate, all the data required to reconstruct the image digitally is obtained.

CCD’s can be either front illuminated or back illuminated. Front illuminated CCD’s are cheaper to produce, but suffer from decreased sensitivity and intra-pixel sensitivity issues, as the control electronics and gates are in front of the charge generating silicon.

Interestingly, a modern development of the CCD, the “Binary Pixel Device” more closely resembles the operation of photographic plate than a conventional CCD. Driven by the mobile technology market’s ever increasing miniaturization, the pixels are much smaller than the diffraction limit, and posses a quantisation limit of one or two photons. This translates into an exceptional dynamical range, but a strongly non-linear response function. Whether this technology has any application




Figure 2.2: The bucket analogy for the working of a CCD, from Janesick and Blouke (1987)

for astronomy remains to be seen (Yang et al., 2009).

2.2.1 Reduction

To obtain useful science data from a CCD image, several additional steps must first be taken:

De-biasing

This step refers to the removal of both the bias level offset and the readout noise, which usually has a fixed pattern component in addition to the stochastic element. A bias frame can be generated simply by taking an extremely short exposure with the shutter closed. Typically several 10s of bias frames are taken and combined in order to remove the Poisson noise.

Dark noise removal

In addition to electron-hole pairs generated by photons, the electrons can also be thermally excited, which provides an additional background noise source. The rate at which thermal electrons are generated decreases exponentially with decreasing temperature (Birney et al., 2006). In most professional astronomical observatories, the CCD's are cooled by liquid nitrogen to cryogenic temperatures, drastically reducing the generation of thermal electrons. For these situations, dark currents are

negligible. For example, the AUXCAM detector in ACAM on the WHT has a dark current of < 4 electrons an hour¹, which is far below the sky brightness level.

In situations where the dark noise is an important contribution, a “dark frame” can be taken. Because different regions of the CCD will be at slightly different temperatures due to the precise layout of the electronics and airflow, the dark frame will contain a fixed pattern component, as well as a stochastic one. Similar to the bias correction, the shutter is closed, but the exposure length is set to be the same as the science data. Many dark frames are combined to average over the noise, and then the dark frame is subtracted from the science data. Since the dark frame will already contain the effect of the bias, one should only use one or the other correction, never both.

Flat fielding

CCD production is not perfect, and slight differences in the sensitivities and well depths of pixels can produce percent level variations in count rates from pixel to pixel. External factors such as vignetting in the instrument or dust settling on the chip or optics, also add to this effect. This obviously would wreak havoc high precision measurements if the star of interest travels across pixels.

Flat fielding is a method of attempting to remove these effects. The CCD is well exposed on a uniform intensity field, and the resulting image is bias/dark corrected, and then the science frame is *divided* by this flat frame to remove non-uniform responses.

An additional concern is the wavelength dependence of flat fields. This is a particular problem in thinned, back-illuminated CCD’s (such as the ones used by NGTS) due to slight variations in the etching of the surface layers. The flat-fielding strategy of NGTS is based around taking regular dawn and twilight flats (the sky is a convenient approximately flat illumination source). However, the sky is very blue, in contrast to the majority of NGTS targets which are quite red. Blue light is more sensitive to surface properties of CCD’s, leading to what is known as the “Blue Diamond effect” (Grange and Goad)².

In some cases the importance of flat-fielding can be diminished, if the guiding is precise enough to keep the stars centered on the same pixels, or if the star is heavily defocussed, so that many hundreds of pixels are sampled (Southworth, 2010).

¹http://www.ing.iac.es/PR/wht_info/opticalCCDs.html, accessed 14/06/16

²A conference poster on this topic can be found at <http://www.mrao.cam.ac.uk/wp-content/uploads/2014/04/AGrange.pdf>

Sky removal

The sky emits and scatters light, which is typically of no interest astronomically. This is mainly a problem for ground based observing, although geocoronal contamination can be a major limiting factor for space-based UV observations. Therefore it is necessary to estimate the sky contribution to the pixels containing the target and subtract this value. A larger number of non-target pixels used to estimate this background provides a better estimate, however if the sky background is spatially varying the pixels used must be kept close to the target.

When conducting high resolution spectroscopy and attempting to measure specific absorption lines, it is critical to separate out the absorption of the Earth's atmosphere from that of the target, particularly if there are shared spectral lines (as is the case for sodium). Fortunately, for high resolution observations, the velocity signal of the Earth's orbital motion makes it easy to separate out and remove this contribution. (This is what makes my measurement of the line profile of sodium in HD 189733b in Chapter 5 possible).

CCD equation

The total error on a pixel count value depends on several components, related to the different reduction steps mentioned so far, this is known as the CCD equation (Howell, 2006)

$$\frac{S}{N} = \frac{N_*}{\sqrt{N_* + n_{pix} \left(1 + \frac{n_{pix}}{n_B}\right) (N_S + N_D + N_R^2 + G^2 \sigma_f^2)}} \quad (2.1)$$

Where N_* are the source counts, n_{pix} are the number of source pixels, n_B are the number of pixels used to estimate the background level, N_S are the number of sky counts *per pixel*, N_D is the dark current, N_R is the readout noise (this term is squared since it is a Gaussian component, not poissonian) G^2 is the gain and σ_f^2 is factor that depends on details of the ADC, typically about 0.29.

For the applications in this thesis, where the targets are relatively bright and well exposed, the N_* term is dominant, and the signal to noise ratio therefore scales roughly as $\sqrt{N_*}$, but, it is worth being aware of other potential sources of error.

2.2.2 X-ray astronomy

CCD's can be used just as effectively for X-ray astronomy as well as for optical. However, the high energy photons are sufficient to liberate many hundreds of elec-

trons. A major advantage of this is that the number of electrons liberated by an X-ray is proportional to the energy of the photon. Since X-ray sources are typically very faint, the CCD can be read out between each impact, so every incoming photon has a time, position, and energy stored in an “Event list”. This effectively means that an X-ray source imaged with a CCD also contains a spectrum. This can also be used to distinguish between background and target sources, as the characteristic energy ranges may be different. Since the CCDS can also respond to visible light, it is important that an appropriate blocking filter is used on bright sources to prevent large numbers of optically generated electrons being mistaken for X-ray events, for instance the early false detection of HD 209458b by Kashyap et al. (2008) was likely due to an inappropriately thin blocking filter was likely responsible (Wheatley, P. J., personal communication).

2.3 Spectrograph

Spectrographs are one of the most important tools in all of astronomy. The interpretations of exoplanet atmospheres in this thesis rely on the composition and dynamics insights that can only be gained by analysing exoplanet spectra. Most modern spectrographs disperse light using diffraction gratings, the basic principle of which is summed up by the grating equation

$$m\lambda = d(\sin(\alpha) + \sin(\beta)) \quad (2.2)$$

where α is the angle of incidence and β is the angle of diffraction, and d is the separation between the lines in the grating. This implies that the dispersion, $d\beta/d\lambda$ is:

$$\frac{d\beta}{d\lambda} = \frac{m}{d \cos \beta} \quad (2.3)$$

This does not depend upon the angle of incidence. Calculating the spectral resolution, R , from this relationship is straightforward if we consider the Rayleigh criterion, whereby a process becomes diffraction limited when the diffraction maxima of one image overlaps with the minima of another. In the case of a spectrum, the condition is to find the wavelength separation, $\Delta\lambda$, where the width of the diffraction spike of wavelength λ and order n is equal to the separation from the diffraction spike at $\lambda + \Delta\lambda$.

$$m(\lambda + \Delta\lambda) = d(\sin(\alpha) + \sin(\beta + \Delta\beta)) \quad (2.4)$$

It is known that an additional pathlength of λ/N , where N is the number of slits, induces a phase lag of $\pi/2$, giving the minima (Jha, 2009). Therefore we can also solve for the minima at angle $\beta + \Delta\beta$,

$$m\lambda + \frac{\lambda}{N} = d(\sin(\alpha) + \sin(\beta + \Delta\beta)) \quad (2.5)$$

Now Equations 2.4 and 2.6 can be combined to find

$$R \equiv \frac{\lambda}{\Delta\lambda} = mN \quad (2.6)$$

There is no resolution at all at the zeroth order, and this increases linearly with the order, m . Crucially, there is no dependence at all on the angle of incidence.

The intensity distribution can be described in the same way as the intensity from a single slit; Fraunhofer diffraction.

Clearly this spectrograph would be inefficient - most of the flux remains in the zeroth order where there is no resolving power. In practice, most spectrographs are “Blazed”, whereby the grating is designed so that each of the grating elements is at an angle. This causes the majority of the flux to end up in a higher order.

$$\lambda_b = \frac{2d}{Nm} \sin(\theta_b) \cos(\alpha - \theta_b) \quad (2.7)$$

where θ_b is the blaze angle and λ_b is the blaze wavelength, or the wavelength where intensity is maximal for each order, m . Harps has a high blaze angle of 75%

The extreme case of this is in an echelle spectrum, where the blaze angle can be set so high that the peak of the grating efficiency function might be over a hundred orders away from the zeroth order. This has the advantage of being able to reach extremely high resolutions, without having to have extremely finely ruled diffraction gratings that are impractical to produce. The disadvantage is that the spectral orders begin to overlap, so an additional step is needed to separate them.

An additional dispersive element is placed to “cross disperse”, dispersing the orders perpendicular to their dispersion direction from the primary disperser. An example of an echelle spectrum is shown in Figure 2.3.

2.4 ACAM on the William Herschel Telescope

In Chapter 6 I will present a ground based optical transmission spectrum of WASP-52b, which was obtained with the auxillary-port camera (ACAM) instrument on the William Herschel Telescope (WHT). The WHT is a 4.2m alt-azimuth telescope




Figure 2.3: A raw spectrum from HARPS, artificially colored. The laser frequency comb appears as a set of dots adjacent to each order for accurate wavelength calibration. The dispersion direction is vertical, and the cross dispersion is horizontal. The many separate orders are clearly visible. The Blue (Jasmine) and red (Fire) chips are seen to be separated in the middle of the image.




Figure 2.4: Percentage of WHT nights lost to bad weather from 1989-2016, individual years in grey, monthly mean in red.³

at the Roque de los Muchachos Observatory on La Palma. At the time of its construction in 1987 it was one of the largest telescopes in the world, and it remains the largest single-mirror telescope in Europe ⁴

Roque de los Muchachos is one of the worlds premier astronomical observing sites, with a median seeing of 0.8 arcsec in the summer months, with around 75% of nights being clear (see Figure 2.4 for a monthly breakdown). Whilst not as dry as sites in Hawaii or the Atacama desert, the atmosphere above La Palma has a particularly small water column, which makes it the best site in Europe for Near Infra-red observations that could be effected by high atmospheric transmission (Selby and Mampaso, 1991). However, during the summer months, dust from the Sahara desert can become suspended above the observatory for periods of several days. The Saharan dust has a grey opacity profile, which can cause an optical extinction of up to a magnitude.

The sky above La Palma is exceptionally dark, a median moonless night brightness of 21.9 and rms 0.1 mag/arcsec² in V band. The main contributions to sky brightness are airglow and zodiacal light, with light pollution (e.g. sodium D line emission from street lamps) being negligible except for a very narrow region

⁴Limitations to casting the glass for mirrors of sizes much greater than this, as well as practical advantages related to mirror maintenance such as Active Optics systems means that larger telescopes are frequently segmented mirrors on actuators, such as the 10m GTC, which is also at La Palma




Figure 2.5: The instrument design of ACAM

above the horizon at azimuths corresponding to nearby population centers. The observations I present in Chapter 6 benefit from this, as the 7000–9000 Å region can become contaminated by OH band airglow.

The WHT has a Cassegrain configuration, with the option to mount visitor instruments at dual Cassegrain-Naysmith focii. The focal ratio at these mount points is $f/11$. ACAM is mounted permanently at one of these points, which makes requesting for time to be scheduled on this instrument particularly easy. Due to this setup, a *de-rotator* is required to keep a consistent sky angle on the CCD. One concern is that the transmission of the de-rotator is not constant, which could inject systematic errors into the lightcurves.

For the purpose of my science goals, ACAM in spectroscopy has a number of advantages over the other common user spectrograph that is often mounted at WHT, ISIS:

- ACAM has a higher average throughput to ISIS, by a factor of 1.5
- The throughput of ACAM is a factor of 2.1 higher in the V band, which covers the important sodium doublet feature. However, ISIS is 1.7 times more sensitive

in U band, which may have been useful for the detection of Rayleigh slope like features.

- ACAM is optically simpler, which limits the potential for unexplained systematics to enter the data
- The slit on ISIS is motor controlled, and has a maximum separation of $22''$, which we consider to be slightly too narrow, as it increases the risk of differential slit losses (the optimum being about $27''$)
- Although ISIS has higher resolution, it is dual-armed, with two different optical paths, for red and blue light, These paths are known to possess different systematics. This would preclude the use of the *common systematic model* I use in Chapter 6 to construct a relative transmission spectrum

ACAM is a dual purpose imager and spectrograph, it has access to a wide range of filters, so is capable of broad or narrow band photometry, in addition to long slit spectroscopy at low resolution. An instrument diagram is shown in Figure 2.5. In imaging mode the plate scale is 0.25 arcseconds/pixel, which gives a field of view of 8 arcminutes. In spectroscopic mode the maximum useful field of view is set by internal vignetting at about 6 arcminutes, but due to the curved edges at the edge of the slits it is usually recommended that not more than the central 5 arcminutes be used for slit spectroscopy. Since the technique of differential spectroscopy requires two similar brightness stars to be observed simultaneously, this 5 arcminute limit is a significant constraint on the target selection process (see Section 1.3.5 for details)

Originally the largest slit available was $10''$, which is too narrow for precision differential spectroscopy. In December 2013 a $27''$ slit was kindly made available by the ING due to requests from members of our group, expressly for the purpose of precision transit spectroscopy. The ability of slits to be cut on request and inserted in the filter wheel is, as mentioned, an advantage over the ISIS instrument. More recently a $40''$ slit was added, which has better machined edges, so it is safer to use the full $6'$ range. This will allow more flexibility in future observing applications.

The WEAVE instrument will be mounted permanently on the WHT in 2017, and from 2018 approximately 70% of telescope time will be dedicated to survey science with WEAVE. ACAM will still be available, but it will no longer be permanently mounted, so competing for and scheduling observations with this instrument will become significantly more difficult. The capabilities and advantages of EFOSC2 mounted on the NTT at the La-silla observatory in Chile are similar to ACAM, and

I have also successfully applied for time on this instrument during the course of my PhD.

Only one grating is available for spectroscopy with ACAM, a VPH (Volume Phase Holographic) grating with 400 lines/mm and a blaze wavelength of 5650 Å. The VPH is placed in the filter wheel (see Figure 2.5). The VPH delivers a spectral resolution (R) of 430 at the blaze wavelength with a 1 arc-second slit. Since the slit used in our observations is significantly larger than the PSF of the stars, the PSF sets the spectral resolution. To attempt to improve the stability of the resolution, we defocussed to 2", somewhat above the seeing.

For this thesis I used ACAM on the WHT for precise differential spectroscopy of WASP-52b in Chapter 6.

Transit observations can take many hours; a typical transit duration is 2-3 hours, and at least as long is needed outside of the transit as in it in order to characterize the baseline trends and noise level. The guiding on the WHT, whilst very good, with an RMS of $< 0.05''$ over a timescale of several minutes, is not adequate to keep the spectrum centered on the same pixels for an entire transit. The 1 hour rms of the guider is $0.3''$, which is primarily caused by differential flexure between the instruments and the autoguider, and can be much worse at some sky positions than others. If an inappropriate filter is chosen for the autoguider, or at relatively low elevations, differential refraction can also decrease the accuracy of the autoguider.

An example of a raw spectral frame is given in Figure 2.6. The dispersion direction is a few degrees off of the vertical, and the spectral traces can be seen to curve slightly. Due to the low resolution of the measurements, the sky features appear as broad bands, which must be removed by background subtraction. The exposure time of this image was 100 seconds, resulting in the accumulation of several cosmic ray hits. A description of the reduction and extraction of high precision spectra from these raw data is given in Chapter 6.

The amount of drift over the expected approximately 6 hour observing period is great enough that differential slit losses could become an issue, and the spectrum will move over a large number of pixels, which increases the level of noise, particularly with the issues with flat fielding these data (described in Chapter 6). Therefore, to maximise the precision of the transit measurements, manual guiding corrections were calculated and supplied at regular intervals to the telescope operator in both X and Y directions. The X direction offsets are simple to calculate, as the horizontal position of the stars can simply be found by fitting for the center of the PSF in a CCD row. Fitting for the pixel drift in the Y (dispersion) direction requires the




Figure 2.6: ACAM raw spectrum of WASP-52b and comparison star

Table 2.1: The order blocking filters for ACAM

Filter	Transmitted wavelengths (Å)	Uncontaminated (Å)	Contaminated by second-order (Å)
None	3300-9500	3300-6600	6600-9500
GG395A	3950-9500	3950-7900	7900-9500
GG495A	4950-9500	4950-9500	None

spectra to be extracted and cross correlated. A script (provided by John Rostron) was used during observations to calculate the pixel shifts with these techniques, and small corrections were supplied to the telescope operator approximately every 15 minutes for all our observations.

The use of the order blocking filter is suggested in the instrument manual, the parameters of the options can be seen in table 2.1. While GG495A would completely remove any order overlap, it would also block out a significant fraction of the blue end of the spectrum, which is particularly important due to its sensitivity to Rayleigh scattering, so GG495A is not an appropriate option. The combination of low intrinsic brightness, (particularly for red stars like WASP-52), and lower detector response, makes regions blueward of 4000 Å difficult to measure with this instrument.

The contribution of the second order can be estimated as follows. Consulting the table (and recalling that the spectral dispersion scales proportionally to the spectral order), we see that the range 7900-9500 Å is contaminated with flux from the second order in the wavelength range ~3950-4750 Å. We can estimate the fractional contamination at 7900 Å from 3950 Å portion of the spectrum in this way. The specifications for the VPH grating show that the second order efficiency is ~0.1 at 400 nm, and extrapolating, ~0.4 at 7900 Å. The *specific intensity* will be an extra factor of two lower.

Using Plancks formula,

$$B_{\lambda} = \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda k_B T}} - 1} \quad (2.8)$$

the ratio of intensities at the two wavelengths, for a 5000 K blackbody is 0.8, the ratio of quantum efficiencies is 50/80 (The quantum efficiency curve for the AUXCAM instrument and the VPH response is shown in 2.7). Multiplying all of these factors

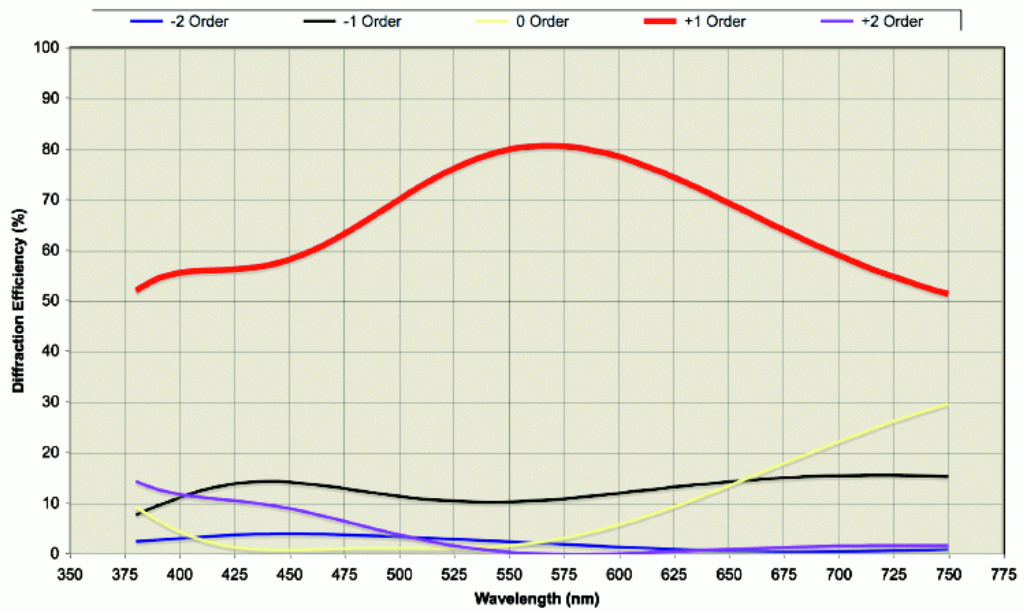



Figure 2.7: Top: the QE curve for the Redplus CCD is the same as the AUXCAM on ACAM. Bottom: The wavelength and order efficiency of the VPH grating for spectroscopy with ACAM.

together, the contamination at 7900 from 3950 Å is

$$0.25 * 0.8 * 0.5 * 0.625 = 0.06 \quad (2.9)$$

As the Limb darkening functions are very different between these wavelengths, and it is known that there is a correlation between limb darkening coefficients and transit depth, it is possible that this could inject a systematic. Assuming liberally that this limb darkening error could cause an error in the depth estimate of a 2.7% transit of order 10%, the total expected signal in planet radius ratio is of order 0.0005, (for a radius ratio of 0.164, which is appropriate for WASP-52b). For a pathological case where the contaminating light somehow does not contribute at all to the transit depth, the maximum fractional radius variation is still only 0.005.

Additionally, it is possible that an extra component in the optical path could add an extra systematic. However, there is no evidence for this in the data I present in Chapter 6.

2.5 HARPS

The High Accuracy Radial Velocity Planet Searcher (HARPS) is a fiber fed high resolution echelle spectrograph mounted permanently on the ESO 3.6m telescope at La Silla. The instrument has a spectral resolution of 115,000 (Mayor et al., 2003). The system’s primary goal and the majority of the allocated observing time, is dedicated to the detection of exoplanets through the radial velocity technique. At the time of writing, there are 40 papers in “The HARPS search for southern extra-solar planets” series (e.g. Pepe et al., 2004; Faria et al., 2016). Well over a hundred planets have been detected by HARPS, and many interesting inferences can be made from this population, including finding a direct constraint on the fraction of M-dwarfs with habitable planets to be $0.41^{+0.54}_{-0.16}$ (Bonfils et al., 2013).

One of the key strengths of HARPS is the high mechanical and thermal stability of the instrument, and of the wavelength solution.

HARPS has two fibers, for simultaneous observations of two targets (e.g. object + sky) or simultaneous wavelength calibration with a calibration source. Originally, the instrument was calibrated with a ThAr gas cell. The spectrum of ThAr contains over 8400 lines in the visible range, and combining the measurement of these lines on HARPS gives a wavelength solution with a single frame precision of 20 cm s^{-1} (Lovis and Pepe, 2007). Using this technique, HARPS has been shown to have a long term stability of $\sim 1 \text{ m s}^{-1}$ (Lovis et al., 2006). Recently, a laser frequency comb was installed on HARPS, similar to the instrument mounted on




Figure 2.8: Inside the HARPS vacuum chamber before sealing and evacuation. The echelle grating is clearly visible.

HARPS-N, which has been shown to have a single frame calibration accuracy of 6 cm s^{-1} , capable of binning to less than 2 cm s^{-1} on combined exposures (Cosentino et al., 2014). Laser frequency combs are in principle capable of reaching precisions long term stability of 1 cm s^{-1} (Li et al., 2008). The Earth induces an RV signature of 9 cm s^{-1} on the Sun. With this technique it is in principle possible to detect Earth analogs (In reality it is more difficult due to activity on the star, such as starspot and activity cycles.). A raw spectrum from HARPS can be seen in Figure 2.3, with the laser frequency comb visible.

The data I use in Chapter 5 to measure the windspeed on HD 189733b were obtained from the HARPS reduced data products archive, and had already been partly analysed by the Data Reduction System (DRS).

2.6 COS on HST

The Hubble Space Telescope is a 2.4m diameter Ritchey-Chrétien reflector telescope in a low Earth orbit. The telescope was launched in 1990, but due to a fault in the polishing of the primary mirror it suffered significant spherical aberration until a servicing mission was able to install a corrector in 1993.

Hubble has become a mainstay of exoplanet research, due to its precision and wavelength coverage with multiple instruments. However, there are complications




Figure 2.9: Three of the main systematic effects in HST transit lightcurves, from Wakeford et al. (2016)

involved in using the observatory for exoplanet transit observations. Due to the low orbit, most targets can only be observed for a maximum of ~ 45 minutes continuously before becoming unobservable for the remainder of the 95 minute orbit. This means that exoplanet transits, which frequently take upwards of two hours, can be left with poor phase coverage. Additionally, the observations suffer from several well known systematic effects, but the origin of these trends is relatively well understood and reproducible (see Figure 2.9). Wakeford et al. (2016) describe the three dominant systematic trends present in WFC3 transit light curves: “breathing” effects caused by thermal contraction and expansion as the telescope passes in and out of the Earth’s shadow, a “ramping” or “hook” effect believed to be caused by charge trapping in the CCD, and a roughly linear visit long trend for which the cause is unknown. Shifts in the wavelength solution can also induce systematics (Deming et al., 2013). Despite these difficulties, transit observations with HST have revealed a wealth of information about exoplanet atmospheres (e.g. Sing et al., 2015).

One area in which HST is currently unparalleled is in observations in the FUV and NUV, which cannot be done from the ground due to atmospheric transmission. Observing exoplanet transits in these wavelengths led to the discovery of evaporating planetary atmospheres (Vidal-Madjar et al., 2003). In Chapter 4 of this thesis I use UV data from HST COS to recover the unobserved EUV spectrum of HD 209458,




Figure 2.10: The COS spectrum of HD 209458 from which the linestrengths in Chapter 4 are taken. Figure from France et al. (2010).

an important component to understanding the evaporation efficiency on exoplanets.

The Cosmic Origins Spectrograph (COS) is a medium resolution spectrograph designed for observations between 115 and 320 nm. The instrument is slitless, and has access to 7 gratings, with spectral resolving powers ranging from 2100 to 21,000. The FUV sensor is a Cross Delay Line detector (XDL) with two segments that can be read out as accumulated images, or in a “TIME-TAG” mode, which provides an event list with 32 ms time resolution. Events detected by the photon-counter are stored on two 16384x1024 pixel arrays. COS is particularly susceptible to geo-coronal contamination, and the sustained impact of this illumination decreased the detector sensitivity to Lyman- α by 30% in the first year operation, though this can be offset somewhat by increasing the detector voltage (Osterman et al., 2011).

All of the information on the evaporation of extrasolar planets has been gathered with HST. No other currently available instruments can cover the same wavelength range, and there are currently no firm plans to produce one. Early proposals for a next generation space telescope, the Advanced Technology Large-Aperture Space Telescope (ATLAST), envisioned to be launched in the 2020s, has a strong emphasis on UV wavelengths (Postman, 2009). As it stands, the end of Hubble could signal the closing of our scientific eyes on a huge region of the electromagnetic spectrum for several decades.

2.7 XMM-Newton

The X-ray Multi-Mirror Mission (XMM-Newton) was launched into a elliptical high altitude orbit in 1999.

XMM-Newton is in fact a collection of three telescopes. Each telescope consists of a layered cylindrical array of 58 Wolter type I grazing incidence mirrors (Wolter, 1952a,b). The Wolter mirror design consists of two surfaces, a parabolic plane followed by a hyperbolic one. The two surfaces have a common focal point. Each mirror is vapor plated with gold and is ~ 1 mm thick, and separated from the next layer by between 1.5 and 4 mm. Each mirror cluster is 250 cm long and weighs 425 kg. The effective area of the mirrors depends on the wavelength, but is ~ 1500 cm² for photons at 1 keV - equivalent to a 45 cm diameter circular aperture.

The three telescopes feed two MOS cameras in the focal plane, and the PN and several spectrographs off axis, where roughly 50% of the photons are directed. The relevant detectors to this thesis are the EPIC MOS and PN cameras, which are cooled CCDs operating in pixel counting mode. The amount of charge generated by an X-ray photon in the CCD is proportional to the energy of the photon, so they function simultaneously as imagers and spectrographs. Due to the instrument setup, all three cameras can observe simultaneously, and their output combined to boost sensitivity. Their relative effective areas can be seen in Figure 2.11

In Chapter 4 I use data from XMM-Newton to place constraints on the X-ray flux of HD 209458. Combining this information with constraints from UV lines from HST COS I recover the high energy spectrum of the stellar corona. The X-rays probe much hotter gas than that probed by UV emission lines. The fluxes in Chapter 4 were taken from the 3XMM-DR5 (Rosen et al., 2016).

2.8 Markov Chain Monte Carlo

John von Neumann once said *“Anyone using Monte Carlo is in a state of sin”*. One can assume he was being flippant, considering he was involved in developing the technique in the 1940’s. Today, judging by its popularity, it might be more accurate to say that anyone *not* using Monte Carlo is in a state of sin. The original Metropolis algorithm is considered one of the ten most influential algorithms of the 20th century (Beichl and Sullivan, 2000). The Metropolis algorithm is a popular example of a class of sampling algorithms, collectively known as Markov Chain Monte Carlo (MCMC) techniques.

MCMC algorithms are used to sample from a probability distribution in




Figure 2.11: The effective area curves for the instruments in XMM-Newton. The step like features are due to absorption by the inner electron shell energies of elements in the optical blocking filters, such as carbon.

such a way that the distribution of stored samples in the “chain” asymptotically approaches, or converges towards, the underlying probability distribution. A key property of these algorithms is that they are stochastic, and do not have any memory of previous steps.

This is a computationally very simple way to approximate complex multi-dimensional integrals. Indeed, there are classes of high dimensional problems for which an MCMC approach is the only known way to recover an acceptable solution within a reasonable length of time (Dyer et al., 1991). The more steps in the chain, the more accurate the approximation of the underlying distribution becomes.

One particularly powerful use of MCMC, and the sense in which it is used in this Thesis, is to solve the complex integrals involved in generating the joint posterior distribution to perform Bayesian inference.

2.8.1 Bayes Theorem

Bayesian statistics is a branch of probability theory that holds that probabilities should be interpreted as ‘degrees of belief’ about the state of the physical world. This is in contrast to the frequentist interpretation that the likelihood of a given event is fixed, and will converge to a given value after an infinite amount of trials.

Bayes theorem is simple to derive from statements of conditional probability. Consider $M(\theta)$, a model with parameters θ , and a dataset, D . We want the

probability of the model being accurate given the data, from the rules of conditional probability, this can be written as:

$$P(M(\theta)|D) = \frac{P(M(\theta) \cap D)}{P(D)} \quad (2.10)$$

or, the probability of the model given the data is the combined probability of the model and the data divided by the probability of the model. It is not immediately obvious how this equation is helpful, but we can also write

$$P(D|M(\theta)) = \frac{P(M(\theta) \cap D)}{P(M(\theta))} \quad (2.11)$$

And by combining these Equations we recover Bayes theorem, which is written:

$$P(M(\theta)|D) = \frac{P(M(\theta))P(D|M(\theta))}{P(D)} \quad (2.12)$$

The terms on the right side are generally referred to as the prior, $P(M(\theta))$, the evidence, $P(D)$, and the likelihood, $P(D|M(\theta))$, and each has its own interpretation.

The prior is a feature that distinguishes Bayesian statistics from more traditional statistical techniques. It represents prior belief in the likelihood of the proposed model *irrespective* of the current dataset. This can represent knowledge from previous experiments being applied as a constraint on the current dataset. An example of a prior in practice in this thesis is in Chapter 6, where information from simulated limb darkening profiles is used to constrain the limb darkening coefficients fit during the transit. If parameters are correlated, constraining one parameter with prior knowledge can improve the precision of the estimate of the parameters of interest. There is an element of danger though, as an incorrect or inappropriate prior can lead to a false conclusion. In practice, unless well defined constraints exist, it is common to pick an “uninformative prior” chosen to add no bias to the fit. Even this can be fraught with danger, however, as what form an “uninformative” prior should take is not always obvious - asked to pick a uniformly random number between 1 and 10, do you pick a linearly uniform or a logarithmically uniform value? A more physical example would be the difference between a uniform prior in magnitude and in flux, see example in Figure 2.12. Users of MCMC should consider the priors very carefully, as even not explicitly declaring a prior (i.e., it effectively takes a value of 1 for any value of θ) *implicitly* assumes that the value is linearly distributed, which may not be true. With some information on the form of the parameter, a Jeffry prior can be constructed, which is known to be uninformative (Jeffreys, 1946).




Figure 2.12: An example of an incautiously applied prior. A uniform distribution in Magnitude is strongly biased to lower fluxes.

$P(D)$ it is the *total* probability of observing the data for all possible values of the model parameters. It serves as a normalising constant in Bayes theorem, such that the sum of $P(M(\theta)|D)$ for all values of θ is equal to 1. It can be expressed as

$$P(D) = \int P(M(\theta'))P(D|M(\theta'))d\theta' \quad (2.13)$$

This is not a trivial quantity to calculate, and rarely has an analytic solution. It must be solved computationally. In fact, a useful property of the MCMC algorithms is that the evidence does not need to be calculated, as will be seen.

Now we come to the “likelihood” term (given the shorthand L), or the probability of getting the data given the model. This is the easiest term to understand intuitively. For practical applications of statistics, our assumption is typically that observations can be expressed as a set of independent Gaussians with a variance which represents the uncertainty in the measurement, and a mean value equal to the underlying model. The probability of recovering a set of known observations given a model and these assumptions is:

$$P(D|M(\theta), \sigma) = L = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{(D_i - M(\theta)_i)^2}{2\sigma_i^2}} \propto e^{-\chi^2} \quad (2.14)$$

So in this situation, minimizing the χ^2 is equivalent to maximizing the likelihood. Note that Likelihood functions are very general, and do not need to assume inde-

pendent Gaussian statistics. *This is a major advantage of maximizing the likelihood over minimizing χ^2 .*

An example of induction within the Bayesian framework, where all these components are seen to work, can be given by a simple thought experiment. Imagine a rare genetic defect that from epidemiological studies is known to effect 1 in 1000 people. A scientist develops a test for the disease that will correctly detect the genetic defect in 99% of cases of people who have it, but will also generate a false positive on healthy people 1% of the time. For a given person, if the test comes back positive, what is the probability that they actually have the disease?

To put this in Bayesian terms, the *prior* is 0.001, the chance of the person having the disease if no test results were known. The *likelihood* is the probability of getting a positive result if they had the disease, 0.99. The normalising constant $P(D)$ is the total probability of receiving a positive result in any scenario, so the sum of the probability that the person is in fact healthy and received a false positive, and the chance that they have the disease and received a correct detection. Writing this out in full:

$$P(\text{Disease}|\text{Positive}) = \frac{P(\text{Disease})P(\text{Positive}|\text{Disease})}{P(\text{Positive})} \quad (2.15)$$

Where

$$P(\text{Positive}) = P(\text{Disease})P(\text{Positive}|\text{Disease}) + P(\text{Healthy})P(\text{Positive}|\text{Healthy}) \quad (2.16)$$

Entering the numbers we find

$$P(\text{Disease}|\text{Positive}) = \frac{0.001 \times 0.99}{(0.001 \times 0.99) + (0.999 \times 0.01)} \quad (2.17)$$

Which gives the probability of having the disease given a positive test of only 9%, since given the rarity of the disease it is much more likely that they are healthy and only received a false positive. This simple example demonstrates the utility of the Bayesian formalism.

2.8.2 MCMC operation

The basic principle of operation of an MCMC is quite simple. The chain sequentially randomly samples possible values of θ (the exact process of choosing a random value, such as whether it is constrained to be near the current value, is algorithm dependent). The key is, at each step in the chain, the algorithm decides whether

to accept the new value of θ , or reject it and stay in the same position, based on their relative values of $P(M(\theta)|D)$. Once again, the process by which the algorithm decides to accept or reject a new chain value is algorithm dependent, but generally uses a ratio of $P(M(\theta)|D)$ values. So the comparison between two positions is

$$\frac{P(D|M(\theta_1))}{P(D|M(\theta_2))} = \frac{P(M(\theta_1))P(D|M(\theta_1))}{P(M(\theta_2))P(D|M(\theta_2))} = \frac{P(M(\theta_1)) e^{-\chi_1^2}}{P(M(\theta_2)) e^{-\chi_2^2}} \quad (2.18)$$

note that the evidence term has been lost in the ratio, since it has no dependence on the value of theta. For computational reasons, the logarithm of these values are often used. This means that in a typical case, the proportion of time a chain spends in an area of parameter space is proportional to its χ^2 , moderated by its prior probability. This allows the recovery of the underlying probability distribution.

Credible intervals⁵ are defined for the parameters from the frequency that they appear in the chain. It is most common to report the median and the 18th and 84th percentiles (equivalent to 1 standard deviation if the distribution is Gaussian.)

2.8.3 Bayesian Information Criterion

The Bayesian Information Criterion, (BIC)

$$BIC = -2 \ln L + k \ln(n) \quad (2.19)$$

where k is the number of degrees of freedom in the model, n is the number of datapoints and L is the model Likelihood. This can be rewritten in terms of a χ^2 value as:

$$BIC = \chi^2 + k \ln(n) \quad (2.20)$$

As can be seen, the BIC severely penalizes increasing the complexity of the model unless it provides a significant improvement to the goodness of fit. BIC is a popular method for testing whether more parameters are justified.

2.9 Gaussian processes

Gaussian processes are a natural extension of Bayesian methods, and fit neatly into an MCMC scheme. The statement of their operation is simple, but the practical consequences are rich and interesting, as shall be shown.

⁵subtly different to the concept of the confidence interval in traditional statistics, mainly due to the presence of the prior

The likelihood function given in Equation 2.14 assumes a set of N independent measurements with a Gaussian scatter; it is a multivariate Gaussian with N dimensions. If instead we assume that the measurements are not completely independent, but have a co-variance matrix K with non-zero off axis elements, then the likelihood function is written as

$$\log p(\mathbf{y}|X, \theta) = -\frac{1}{2}\mathbf{y}^\top K_y^{-1}\mathbf{y} - \frac{1}{2}\log |K_y| - \frac{n}{2}\log 2\pi \quad (2.21)$$

The key difference here is that models of systematics in the data that we are not interested in, or do not know how to model properly, can be absorbed into the covariance matrix K . This likelihood function can be understood by breaking down into its three components. The third,

$$-\frac{n}{2}\log 2\pi \quad (2.22)$$

is a normalizing factor, and plays no particularly special role in the fit.

$$-\frac{1}{2}\mathbf{y}^\top K_y^{-1}\mathbf{y} \quad (2.23)$$

is the “model fit” - it is the probability that a multivariate Gaussian with mean zero (the prior) defined by the covariance matrix will pass through the points of the training set. This can be seen trivially by using the covariance matrix for n independent Gaussians, where it reduces to

$$\log p(\mathbf{y}|X, \theta) \propto -\frac{1}{2}\sum_{i=1}^n \frac{y_i^2}{\sigma_i^2} \quad (2.24)$$

or equivalently

$$\begin{aligned} p(\mathbf{y}|X, \theta) &\propto e^{-\frac{1}{2}\frac{y^2}{\sigma^2}} \\ &\propto \chi_n^2 \end{aligned} \quad (2.25)$$

If we consider a typical covariance function, such as the squared exponential, this function will always be maximized by picking a high variance and a short timescale (see Figure 2.13 for the form of some common kernels). To put it another way, you can in principle reproduce any set of datapoints by assuming they’re generated by white noise! This tendency is tempered by the second term,

$$-\frac{1}{2}\log |K_y| \quad (2.26)$$

which functions as a complexity penalty. The more diagonally condensed (i.e, the closer to white noise) the covariance matrix becomes, the larger this penalty term becomes. Precisely because white noise is capable of replicating any set of observations, it is a bad model by Bayesian standards as it has very low specificity and very low predictive power. Underlying parameter values and extrapolated points will have very large error bars due to the enormous range of equally good functions that can replicate the training set. This term rewards specificity and simplicity, or covariance matrices with higher long-range covariance and smaller white noise, as this class of functions is narrower so has greater explanatory and predictive power. This term will tend to shrink white noise to a minimum and increase length scales as much as possible - for a squared exponential kernel, a minimum complexity model would tend to a flat line with the mean of the training set. This of course, most likely, would be a terrible fit to the data.

Another way of looking at this term is *If you generated a random Gaussian process with these parameters, how likely is it that it would be the “correct” one that fits your data points?* For a white noise kernel the parameter space is obviously huge, but a minimum complexity model making specific predictions will be very similar draw to draw, so is preferred.

The power of Gaussian process regression is in playing these terms off together, as the sum of these terms will be maximized by finding the simplest model that adequately fits the data.

Considering the flexibility of the Gaussian process, it is natural to worry that when attempting to fit a model to data, the GP may compete with your model to fit out the signal as if it were a component of the noise. The extent to which this is true will depend on how much your proposed signal “looks like” the noise in the data. If the correlated noise in your data set is of a similar characteristic timescale and amplitude to your proposed signal, it is justified to expect the GP to increase the uncertainty on your proposed model fit, as this represents a more accurate picture of the true uncertainty of your measurements in the presence of correlated noise. If there are spurious signals in the dataset of similar size to the one you are attempting to measure, your confidence in any derived quantities *should* decrease! The inflation of errorbars that can arise from including a GP in a model fit can be viewed as a consequence of being more honest about all the sources of uncertainty in the measurement.

Gaussian Processes are used in Chapter 6 of this work to generate a model of the correlated noise in the lightcurve of the transit of WASP-52b on two nights, and is also found to be suitable for removing the contribution of a star-spotted region.

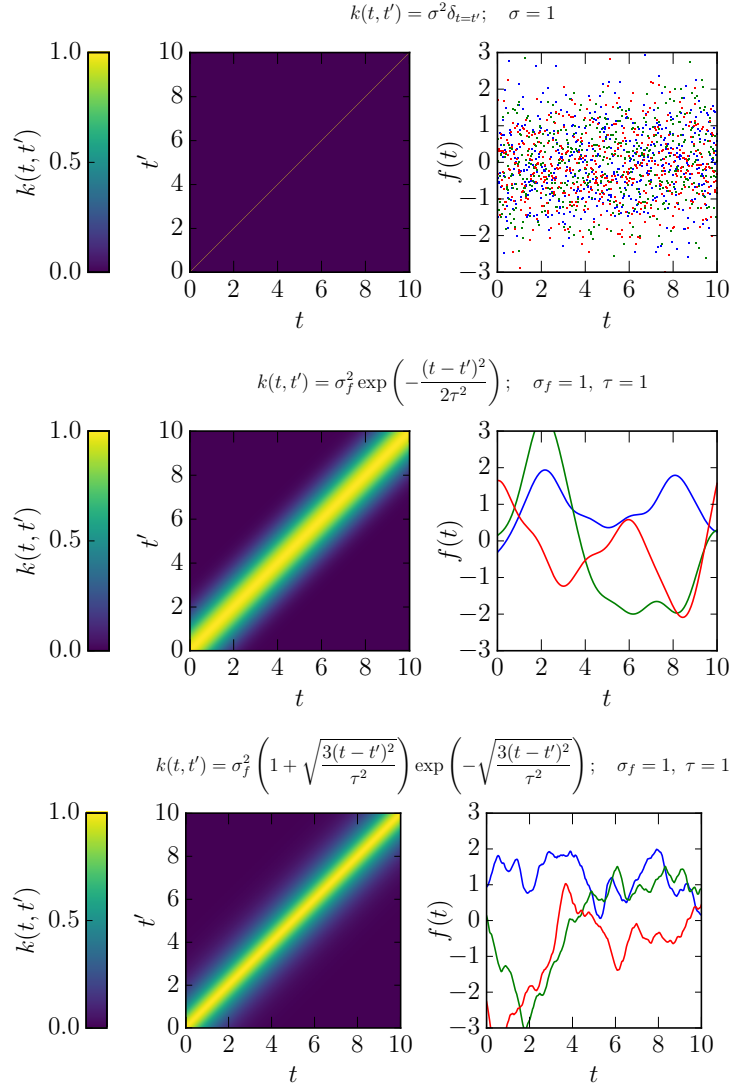


Figure 2.13: A white noise kernel, squared exponential kernel and Matern kernel. Left side is a visual representation of the covariance matrix, on the right are three random draws from a Gaussian process with these noise properties.

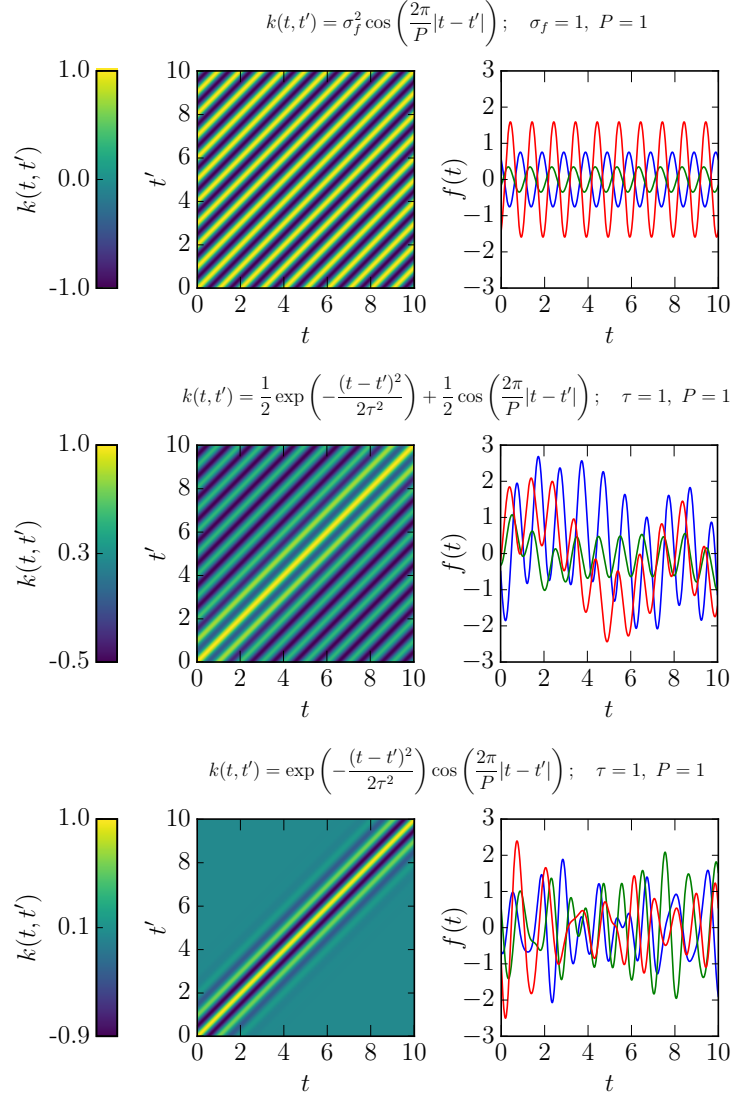


Figure 2.14: Three periodic kernels, same as Figure 2.13

2.10 Conclusions

In this Chapter I have described some of the key instruments and techniques used in the following chapters. I have given a description of Bayesian statistics, and how this leads into more advanced techniques such as MCMC and Gaussian processes, which can be a powerful techniques to model systematic noise, which is used in Chapter 6 to analyse the spectroscopic lightcurves of WASP-52b.

Chapter 3

The hunt for hot Neptunes with NGTS

There are ways of approach to unknown territory which lead surprisingly far, and repay their followers richly. There is probably no better example of this than eclipses of heavenly bodies.

Henry Norris Russell, 1948

3.1 Introduction

In this Chapter I will describe some of the involvement I have had in the Next Generation Transit Survey (NGTS) consortium. The NGTS project had official first light in Paranal on 14/01/15, and the search for new planets is ongoing.

There are two primary science goals for the NGTS project. Firstly, to detect a *statistical sample* of well characterized super-Earths and Neptunes. The currently known planets in this size range are few in number, and do not have well constrained masses. Constraining the masses of planets this size with Radial Velocity measurements is important to be able to place constraints on their bulk composition, and is an essential input to models of planet formation. The majority of transiting planets in this size range were detected with *Kepler*, and, due to the relatively small field of view transit on average quite faint stars. NGTS is designed to have the sensitivity to detect super-Earth and Neptune sized transits, but will cover a larger sky area, leading to a larger sample of these planets around bright stars suitable for this

characterization.

Secondly, to detect a small number of exceptional targets around bright stars that are suitable for atmospheric characterization. Currently, nearly all of the exciting transmission spectrum results have been from hot Jupiters and hot Saturns, with a few notable exceptions (e.g. Knutson et al., 2014; Kreidberg et al., 2014; de Wit et al., 2016). A sample of transiting small planets orbiting bright ($I_{mag} < 11$) would open the potential to exciting characterisation of a potentially very different atmosphere class. Studying the transmission spectra of exoplanets can allow one to break some of the degeneracies in bulk composition from mass and radius alone, for example, a planet with a given mass could be constructed from a rocky core with a hydrogen/helium envelope, or have a very large ice and volatile fraction, and have the same radius. The detection of spectral features in the atmosphere can inform which of these scenarios is more likely.

In this Chapter I will describe a simulation of the expected planet yield of NGTS Section 3.2, The planet catch simulations were done in collaboration with Peter J. Wheatley and Simon Walker, and are also described in Walker (2013). I will present analysis of an NGTS test campaign on the *Kepler* field taken in the summer of 2013 in Section 3.3, the software and workflow I used to conduct this analysis has become part of the current NGTS pipeline operations in Warwick, reducing data taken in Paranal, I will also briefly describe several data quality assessment tasks that I wrote for the current version of the pipeline. I will provide examples of their output from recent NGTS data, generated by the current version of the NGTS pipeline. The code to produce them was written by myself, but small cosmetic and structural modifications were made by Simon Walker during implementation for consistency with the rest of the pipeline.

3.2 Planet catch simulations

In order to characterize the expected performance of NGTS, we wanted to estimate the number and size distribution of planets that the survey could be expected to detect. Information from such calculations can be used to further refine the survey design, or influence follow-up strategy by additionally estimating the occurrence of false positive blend scenarios.

Peter Wheatley used the Besançon galaxy model (Robin et al., 2003) and model stellar spectra (Pickles, 1998) to generate simulated fields of stars at several different galactic latitudes. Combining this list of stellar types and brightnesses with the expected total system throughput, including the filter, CCD, and atmospheric

response for a typical airmass of 1.5, gave the expected NGTS signal levels for each star (Wheatley et al., 2013). Working in collaboration with a colleague, Simon Walker, we used this information to estimate the predicted planet yield of the NGTS survey.

To calculate the detectability of planet transits, a Signal to Noise estimate is required. The star signal provides \sqrt{N} photon noise, and we added a red noise component in quadrature. The red noise component captures the baseline systematic error in the observations that cannot be improved by adding additional measurements. The NGTS survey design required this value to be less than 1 mmag, and early observations in Geneva (see Section 3.4.5) indicate that the instrument is routinely performing at least this well. The value of the red noise component was selected to be 0.3 mmag based on the noise model described in Walker (2013). A scintillation component is also added, to account for additional white noise from atmospheric fluctuations. Scintillation noise depends on properties of the instrument and the observing site, and is calculated using equations from Dravins et al. (1998). The scintillation term can be dominant for bright stars in short exposures, but since it is a white source it will decrease with longer exposures.

Simon Walker wrote a Monte Carlo simulation that would, for each star in the simulated field, have the chance to generate a planet with probabilities based on the best available occurrence rates at the time (Howard et al., 2012). These authors find that for orbits shorter than 50 days the occurrence rate around FGK stars is 0.165 ± 0.008 , for planets between 2 and $32 R_{\oplus}$. Since these simulations were carried out we now know that the occurrence rate is even higher for low mass stars (Dressing and Charbonneau, 2015), which would likely boost the number of detected super-Earths and Neptunes even higher.

After determining whether a planet exists, the radius is drawn from the occurrence distribution, also given by Howard et al. (2012). These rates are complete to ~ 2 Earth radii, so are appropriate for the kinds of planets that NGTS will be capable of detecting.

Once a planet had been generated, its inclination can be selected from a uniform random distribution (there is no reason to suppose there is a preferred inclination). From this, it is trivial to calculate whether the planet transits or not, using the equations in Section 1.2.1, repeated here.

$$b = \frac{a \cos i}{R_*} \left(\frac{1 - e^2}{1 + e \sin \omega} \right) \quad (3.1)$$

We assumed an eccentricity of 0 for all targets. If the impact parameter is

less than $1 - \frac{R_p}{R_*}$ then the full disc of the planet transits the star, between $1 - \frac{R_p}{R_*}$ and $1 + \frac{R_p}{R_*}$ the planet transit is grazing, and beyond $1 + \frac{R_p}{R_*}$ no transit occurs. We included both grazing and non-grazing transits.

The signal to noise ratio achieved on the transit will scale with the transit duration. The transit duration, from 1.2.1, is:

$$T_{tot} = \frac{P}{\pi} \sin^{-1} \left(\frac{R_*}{a} \frac{\sqrt{\left(1 + \frac{R_p}{R_*}\right)^2 - b^2}}{\sin i} \right) \quad (3.2)$$

The *white noise* components of the signal to noise ratio, the photon noise and the scintillation, will be decreased by a factor $\sqrt{N_{exp}}$, where the number of exposures, N_{exp} , is $T_{tot}/(\text{exposure time} + \text{readout time})$ for each simulated planet. The signals calculated by Peter Wheatley assumed exposure times of 10 minutes. The red noise component does not reduce with exposure time.

Armed with this signal to noise ratio, the total detected significance, σ_{det} , of a transit is simple to quantify:

$$\sigma_{det} = \frac{\left(\frac{R_p}{R_*}\right)}{\left(\frac{S}{N}\right)} \quad (3.3)$$

Of course, multiple transits will be observed in a typical campaign. Using simulations of NGTS field observability in Bento (2012), Simon Walker estimated a typical NGTS target would be observed for 800 hours (Walker, 2013). From this, a maximum number of observable transits, N_{trans} can be calculated using the planet period. For each target, the detection significance is increased by the number of transits observed. To estimate the significance from combining multiple independent transit events we used Fisher's combined probability test (Fisher, 1925):

$$X_{2k}^2 \sim -2 \sum_{i=1}^k \ln(p_i) \quad (3.4)$$

Here X_{2k}^2 is the combined test statistic, k is the number of observations, in this case the transits to combine, and p_i is the p-value of each individual transit, which is calculated simply from σ_{det} assuming a gaussian distribution (e.g. the p value for $\sigma_{det} = 3$ would be 0.9973)

Therefore, the combined test statistic of multiple transits in our simulation is:

$$X_{2N_{trans}}^2 \sim -2N_{trans} \ln(p) \quad (3.5)$$

The test statistic has a chi-squared distribution with $2N_{trans}$ degrees of freedom. We chose a threshold value for this significance test that is equivalent to 6σ , in order to represent a very robust detection of a transiting system.

In reality, not all transits in a given observation window will be observed, due to effects like bad weather. This factor is quantified by a “window function”. A simplified window function based on the work of Bento (2012) was used to randomly reject a fraction of planets based on their orbital period. The functional form of the window function we chose was to be 100% efficient for periods less than 7 days, and then decay monotonically to be zero by 25 days.

To summarize, planets are generated for each star in the simulated NGTS fields with the occurrence rate, and period and radius distributions from Howard et al. (2012). The planet must then pass three checks:

- Given a random inclination, does the planet present at least a grazing transit?
- Is the total expected significance for the number of transits in the campaign $> 6\sigma$?
- Given the period, the planet has a probability of being missed due to poor conditions given by the window function

If a planet passes all of these steps, it is considered “detected” and stored.

The simulation was repeated for the total estimated number of fields in the survey. We assumed, for a baseline survey strategy, that each of the 12 telescopes would each observe 4 fields a year for 4 years in order to calculate the total planet catch of the survey. This is approximately 400 square degrees per year, for a total of 1600 square degrees. This is 16 times higher sky coverage than the *kepler* telescope, and represents approximately 4% of the sky.

Detecting the planet in transit data is only half the story. In order for a planet candidate to be considered a true detection, its planetary nature must be confirmed by a second method, typically through Radial Velocity measurements. In some cases it is possible for planets to be confirmed by techniques such as Transit Timing Variations, as demonstrated by *Kepler* (Ford et al., 2012), but this requires quite fortuitous alignment of multiple planets in a system. In some cases, the multiplicity alone makes it exceedingly unlikely to be a blend scenario (a recent example is the TRAPPIST system (Gillon et al., 2016)). For the purposes of this simulation, we

consider a planet to be “confirmable” only if it can be convincingly detected with a modest amount of follow-up time (~ 10 hours total). Larger programs are of course capable of beating down the noise to levels where very small planets can be detected for particularly interesting objects (e.g. CoRoT-7b, Queloz et al., 2009), but this is not practical when attempting to survey populations of exoplanets, as NGTS intends to, and particularly when a large fraction of objects for which time is requested may turn out to be false positives.

The detectability of an RV follow up observation will depend on the brightness of the star and the expected semi-velocity amplitude. Recalling the equation for minimum mass from 1.2.1, and assuming zero eccentricity, the semi-velocity amplitude induced by the planet can be written as

$$K_p = \frac{M_p \sin i}{M_*^{2/3}} \left(\frac{2\pi G}{P_{orb}} \right)^{1/3}. \quad (3.6)$$

All of these parameters are generated by the Monte Carlo simulation, except for mass. As discussed in 1.3, the mass-radius relationship for exoplanets is complex, but for a first order approximation, a power law can be fit between objects roughly the mass of the Earth and of Saturn, which is the mass range of most interest to NGTS. In this regime, $M \propto R^{1/2}$ (this powerlaw also works for the solar system planets). With this additional piece of information, a semi-velocity amplitude can be estimated for each detected planet in the simulation. Assuming that the followup observations are carried out with HARPS, and given a maximum of 10 hours per object, an estimate can be made of the faintest magnitude that the signal could be detected at. If the host star is brighter than this, the planet can be considered a good follow up target.

An issue that is not considered by these simulations is the issue of false positives, however, most false positive scenarios, such as blended Eclipsing Binaries can be uncovered by much cheaper instruments before the candidate is submitted to HARPS. For example, a suspected background blend could be detected rapidly by high contrast imaging. Initially, smaller spectrographs such as FEROS will be used to weed out the very large radial velocity signatures that would indicate a binary, a non detection with this instrument is a good indication that the eclipsing object is sub-stellar. HARPS is required for its metre level precision which allows the detection of very small planets, not because the stars are too faint. Only planet candidates that are not ruled out by these previous steps will be passed to HARPS, therefore, the total ammount of HARPS time per bona fide planet should not be too much greater than assumed here.

The simulated population of NGTS detected planets is plotted in Figure 3.1, as well as the currently confirmed planets, and the planets from *Kepler*. With the RV detectability limits overplotted, it is immediately obvious that the vast majority of *Kepler* planets cannot be confirmed through radial velocity without a significant investment of observing time, due to their faint magnitudes and low semi-velocity amplitudes. Indeed, many *Kepler* planets have velocity amplitudes lower than 1 m/s, which is the practical noise floor for HARPS. Alternatively, the confirmed transiting planets are mostly relatively bright, but with high velocity amplitudes, as most of this population is hot Jupiters from projects such as WASP.

The simulated NGTS planets cover an interesting middle-ground. Due to the larger precision than other ground based surveys, small planets, producing relatively small semi-velocity amplitudes can be detected, but unlike *Kepler*, these smaller planets are detected around bright stars, so a large number fall comfortably below the HARPS S/N cut.

The planets that pass this cut are plotted in 3.2, now as a function of planet radius. A large population of super-Earths ($2-4R_{\oplus}$) and super-Neptune ($4-8R_{\oplus}$) planets are detected around bright stars in our simulation that cannot be reached by other surveys. With next generation spectrographs such as ESPRESSO, our simulation indicates that of order 200 planets of this size could be confirmed with NGTS. Of these, around 40 are around stars brighter than $J \sim 11$, which would make good targets for atmospheric follow up, and would be excellent potential targets for JWST.

3.3 Zero level Pipeline design

In late August 2012 a completed NGTS unit was installed on the roof of the Geneva observatory to characterize the performance of the CCD and optics. Many of the early observations were intended to test the mount alignment and tracking, I participated in these tests in early 2013.

The most stringent test of the performance of the instrument came with a 2 month long campaign to observe a single field on every available night, to test the stability and quality of the photometry over long timescales. The *Kepler* field was chosen, as it had good visibility at the time, and the stars in it are particularly well characterized. Comparison photometry from *Kepler* could also be accessed and analysed.

From 04/06/13 to 02/08/13, the telescope observed for a total of 17 nights, taking over 26000 images. To assess the quality of the observations, I constructed a




Figure 3.1: The potential for following up planetary targets with Radial-Velocity measurements. The thick dashed line indicates 10 hours of HARPS time, and the thin dashed line indicates what will be achievable with ESPRESSO time on important targets. In this plot, existing planets are marked as green, planet candidates from *Kepler* are marked in blue, and NGTS simulated planets are in orange. Note that the majority of *Kepler* planets are not feasible to follow up with RV, whilst many NGTS planets are.




Figure 3.2: The sample of planets that can be confirmed by RV detections, color scheme the same as Figure 3.1. Systems brighter than I magnitude 11 are particularly valuable for atmospheric follow up. NGTS will significantly increase the number of small planets detected in this brightness range.




Figure 3.3: An NGTS unit assembled in Geneva.

basic reduction and aperture photometry pipeline, mainly using software from the Cambridge Astronomical Survey Unit (CASU) developed by Irwin et al. (2004)¹, as well as my own scripts. The basic steps of the pipeline reduction are outlined below.

3.3.1 Reduction

An example of a raw NGTS frame can be seen in Figure 3.4. I wrote simple scripts to reduce the images in Python, using standard techniques.

Dark frames with the same exposure time as the science frames (10 seconds) were taken during the day with the CCD set to the same temperature and subtracted from the science frames to remove both the dark and bias noise.

Flat fields were calculated by median combining well exposed twilight and dawn flats. Dark corrected images are divided by the dark corrected flatfield to remove pixel to pixel variations and instrument vignetting.

The flat-field is also used to construct a *confidence map* which is a required input for the CASU photometry tools. The confidence map is simply the information needed such that the photometry apertures can be variance weighted sums, which gives a more precise measure of the true flux. Without a confidence map, pixels with low illumination, and so a higher variance, would be given the same weight

¹source code available at <http://casu.ast.cam.ac.uk/surveys-projects/software-release>

in an aperture as a well illuminated pixel. The confidence map can also contain a bad-pixel mask, which will be given a weight of 0, so not included in aperture summation.

3.4 Astrometry

The pipeline is built around the CASU photometry package, which is designed to extract photometry from fixed apertures in RA and DEC. For this to work, the astrometric solution of every frame must be extremely accurate and stable across the whole field. CASU provides a tool, **wcsfit**, for fitting the World Coordinate System (WCS) solution that translates from pixel space to RA and DEC.

wcsfit uses the Zenithal Polynomial projection (ZPN) WCS format, which is described in Calabretta and Greisen (2002). The advantage of using this system is that it is a natural way to deal with distortions that are radial about a point on the CCD, which is common result of the design in (mostly) axially symmetric telescopes. The NGTS astrographs are fast and wide field, and are expected to suffer from high order radial distortions. By default, CASUTOOLS fits the odd-only (to preserve axial symmetry) terms of a 5th order polynomial. I found that this was unable to correct for the distortions across the whole chip, so added a 7th order term. This performed very well across the whole chip, and on a well solved frame, the RMS error on the positioning can be as low as 0.4 arcseconds, which is less than a 10th of a pixel (see Figure 3.5).

The CASU **wcsfit** routine does not perform well unless the initial solution to the frame is already quite close to correct, and in particular, the code is not able to fit for radial distortion terms, these must be supplied manually. It is capable of fitting for the coordinates of the field center, plate scale, and field rotation using the CRVAL and CD matrix keywords described in Calabretta and Greisen (2002).

This makes sense for the large, single instrument surveys that CASU is designed for, a single characterisation of the optical distortion of the system will be sufficient for very long stretches of time. Unfortunately, for NGTS, particularly in the commissioning phase, there are 12 telescopes and 12 cameras, frequently being removed, and in some cases taken apart and resembled, and cameras switched from one telescope to another. It would be time consuming and impractical to attempt to manually solve the distortion terms on a regular basis, so I wrote a script that would be capable of solving for them.

The issue is that with wide chips and high order polynomial solutions, the “Goodness of fit space” is highly non linear, as a small perturbation could push the




Figure 3.4: A raw NGTS frame taken in Paranal on 06/09/15. The exposure time was 5 seconds and the FOV is 9 square degrees.




Figure 3.5: A solved astrometry frame. The points indicate the detected sources, lines indicate the direction and distance to the catalogue position for this source, multiplied by a factor of 1 million. The plate scale is $5''$ per pixel, so the astrometric residuals here are better than 1/10th of a pixel across the whole frame. The red point marks the center of the radial distortion term.

solution in one region of the chip to a region where the error in the solution is greater than the average distance between the stars, the fit-goodness space then becomes on average flat and noisy, so gradient descent methods do not work, getting trapped in local minima that are extremely poor globally. MCMC are very useful in these kinds of situations, as the stochastic nature allows them to move around parameter space without getting trapped, and given enough time the solution will converge to the global solution.

I wrote an MCMC based on the EMCEE package (Foreman-Mackey et al., 2013). The code fits for the CD matrix (a 2x2 matrix that parameterizes the plate scale and field rotation), the pixel coordinates of the center of the distortion, the RA and DEC coordinates at this position, and the odd co-efficients of the 7th order ZPN radial distortion, for a total of 12 model parameters. The goodness of fit at each step is evaluated as the median of the separation between the calculated WCS positions of the detected objects in the field and their closest match in the 2MASS all sky catalog (Skrutskie et al., 2006). The code is run until the solution stops improving, which is typically at a median error of ~ 0.4 arcseconds.

This MCMC only needs to be run once for each telescope/camera combination, after instrument setup. The 7th order ZPN coefficients, and the pixel coordinates of the center of the distortion are then fixed in the header files of all future images, until the instrument is modified again, when the MCMC must be re-run.

After the radial distortion terms are found and stamped in the header, only **wcsfit** is required to be run to fine tune the WCS solution on individual image frames, by updating the CD matrix and reference RA and DEC values. This keeps the WCS solution stable in the face of, for example, errors in the guiding, or field rotation caused by a misaligned north axis.

Initial guess solution

During the initial stages of the project, such as the test run in Geneva and commissioning in Paranal, the mount pointing model was not sufficiently accurate, and the difference between the internally recorded and actual position on the sky could be several degrees. This was not a problem for the photometry, as the excellent performance of the autoguider kept the stars centered on the same pixels, with a typical RMS of 0.1 pixels (comparable to the precision of the mounts themselves). The autoguider code is described in McCormac et al. (2013), and cross correlates the image in 2 dimensions to calculate the required offsets.

This made finding an initial astrometric solution difficult with the software previously described, since the initial offset was many orders of magnitude larger

than the average separation between stars. **wcsfit** will not work in these situations, and the previously described MCMC script would take an impractically long time to run (though it would eventually find the correct fit).

The solution was to use *astrometry.net* (Lang et al., 2010) for the initial solution, before the further fitting routine described above. *astrometry.net* is a powerful and robust software tool that has the capability to astrometrically solve a field blind. The code works by matching “asterisms” (geometrically hashed sets of 4 or 5 stars). *astrometry.net* requires no other input but a raw image frame, but the speed of this step can be improved by providing prior knowledge, e.g., the plate scale of the detector, and an approximate region of sky.

3.4.1 Aperture photometry

Aperture photometry is performed on the science frames using the CASUtools **imcore_list**. **imcore_list** is a list driven photometry code, it will extract aperture fluxes for a list of RA and DEC coordinates. This could in principle be supplied with positional data from an external survey (e.g. 2MASS, or GAIA), but many of these sources would be blended within the large pixels of NGTS. For simplicity, I chose to generate the input catalogue from the data themselves.

For each new field to be analysed, the first 10 frames were reduced and combined with **imstack**, which stacks images in RA and DEC coordinates. A source detection algorithm, **imcore**, is then run on this high signal to noise frame to produce a catalogue of source positions.

For operations in Paranal, the reference image is generated by observing the field while *dithering* with a sub-pixel dither pattern before stacking. This can be used to effectively increase the resolution of the instrument, and so measure more precisely the position of detected sources. However, this does have the effect that some of the faintest sources will not be picked up in the input catalogue (Simon Walker, personal communication).

Aperture photometry requires a background removal step to remove the contamination of the sky. Typically this is done by using a nearby “sky apperture”, but **imcore** instead generates a background map, which gives every pixel in the image an estimated sky value. The background map is generated by splitting the image into a grid (with a default cell size of 64 pixels) each cell has its background value estimated with an iterated sigma clipped median. Individual pixels have their background value estimated by a bilinear interpolation between these grid values. An advantage of this approach is that it uses information from a larger number of pixels, so in principle it can provide more precise sky background estimates, if the

sky background varies on scales larger than the smoothing box. If the sky varies on scales smaller than the smoothing box, then the estimate would be significantly worse. An independent analysis of the same dataset by Simon Walker using a sky-aperture approach found similar data quality, implying that the background image approach works at least as well as sky apertures on the data considered.

Each science frame has the fluxes for each catalog source measured by **im-core**, which measures the fluxes in several different user defined apertures. The best results for the Geneva dataset used an aperture with a radius of 3 pixels. The individual result files are then combined into a single FITS file for further processing and analysis.

3.4.2 Frame rejection

Before constructing light curves and detrending, the data were filtered for frames that showed evidence of clouds, or errors in tracking or focus stability.

A proxy for the presence of clouds is the median pixel value of the image divided by the standard deviation of the image. This is essentially a test for the presence of structure in the image, which clouds will present. If this metric is greater than 2, the image is considered cloudy and is rejected. Frames were also rejected if the tracking shifted by more than one pixel, or if the average FWHM of stars in the frame became larger than 5 pixels.

26000 frames were taken in the 2 month *Kepler* field test in Geneva, and of these, 19000 were accepted, and 7000 were rejected, most often for failing the cloudiness test.

The frame rejection was a step carried out for my analysis of the Geneva dataset, and has not carried over to the production pipeline in use at Paranal.

3.4.3 Detrending

Several different schemes exist for removing correlated noise from survey data. For my analysis of the Geneva fields I used the SYSREM algorithm (Tamuz et al., 2005), as implemented for the WASP project by Richard West.

The SYSREM algorithm works by constructing “basis functions” from the lightcurves of a sub-sample of bright non-variable stars. Each object in the frame is then iteratively fit to find a weighting coefficient to give the basis function, such that the residuals of the object are minimised when the weighted basis function is removed. In practice, this means that global trends in the data that might be induced by changing transparency or airmass are removed without having to explicitly




Figure 3.6: A flux rms diagram from the Geneva run. The raw data is in blue, the SYSREM detrended data are in red.

model the effects.

I found that when combining the data from the 17 nights, multiple SYSREM basis functions were required. I found that there were improvements to the median RMS's of the bright stars until SYSREM had been run a total of 8 times. An example of the improvement that SYSREM can have on the RMS can be seen in Figure 3.6. In this example, from 02/08/13, the RMS of the bright stars is improved from ~ 30 mmag to ~ 7 mmag. Note that not all of this reduction is due to removing systematic noise sources, a major contributor is the airmass trend, the origin of which is well understood.

The production pipeline in Warwick is capable of using either SYSREM, or a separate detrending technique developed for CASU, which fits flux offsets with a 2d surface in the detector space. The Trend Filtering Algorithm (TFA) can also be applied to the data (Kovács et al., 2005).

3.4.4 Light curves

In this Section I present a few example lightcurves using the detrended data from the 17 observed nights between 04/06/13 and 02/08/13.

As an acid test of the data quality and the measured S/N, I decided to check that known variable sources could be recovered reliably from NGTS photometry. The field observed was the *Kepler* field, for which a variable star catalogue was constructed early on, in anticipation of the planet finding survey. (Uytterhoeven et al., 2011) I matched sources by their catalog positions, and folded them on the ephemeris and period given in the database. Two examples of binary stars with deep eclipses are shown in Figure 3.7. For comparison, I accessed *Kepler* photometry for the same sources through the MAST archive² and folded on the same ephemeris. Excellent agreement is found, which is a clear indication of the stability of NGTS photometry over timescales greater than a month. It is also an important demonstration that the SYSREM algorithm is not over correcting and removing astrophysical signals from the data.

Taking this a step further, to confirm that NGTS could definitely reach its design goal of red noise significantly lower than 1 mmag, I selected low amplitude variable stars with high frequency from the database. If red noise were important on timescales of order an hour, similar to that of a typical transit, it would be very clear in these lightcurves. In Figure 3.8, I show that a delta scuti with an amplitude of only 1 mmag and a period of an hour is clearly detected in the data.

To confirm that this frequency is distinctly present in the data, I use a Lomb-Scargle periodogram, and find that the peak in the power spectrum does correspond to the expected frequency³. There is no natural alias with a timescale of an hour, and the amplitude and phase are too close to that found in the *Kepler* photometry to be a coincidence. The small difference in amplitude between the *Kepler* and NGTS bandpasses is not unexpected, as they cover different wavelength ranges.

The fact that binning NGTS photometry to this timescale allows this detection clearly shows that NGTS is capable of reaching the levels of precision required to detect mmag level transits. These are the highest precision lightcurves ever achieved in a wide-field ground-based survey.

TrES-2b was the first transiting planet detected in the *Kepler* field (O’Donovan et al., 2006). Several transits occurred on nights which NGTS was unable to observe, but fortunately, a single partial transit was observed in the 17 nights. This was the first transit of an exoplanet observed with a complete NGTS unit. In Figure 3.9 the

²<https://archive.stsci.edu/Kepler/>

³<https://www.aavso.org/vsx/index.php?view=detail.top&oid=269501>





Figure 3.7: Two examples of eclipsing binaries recorded with NGTS

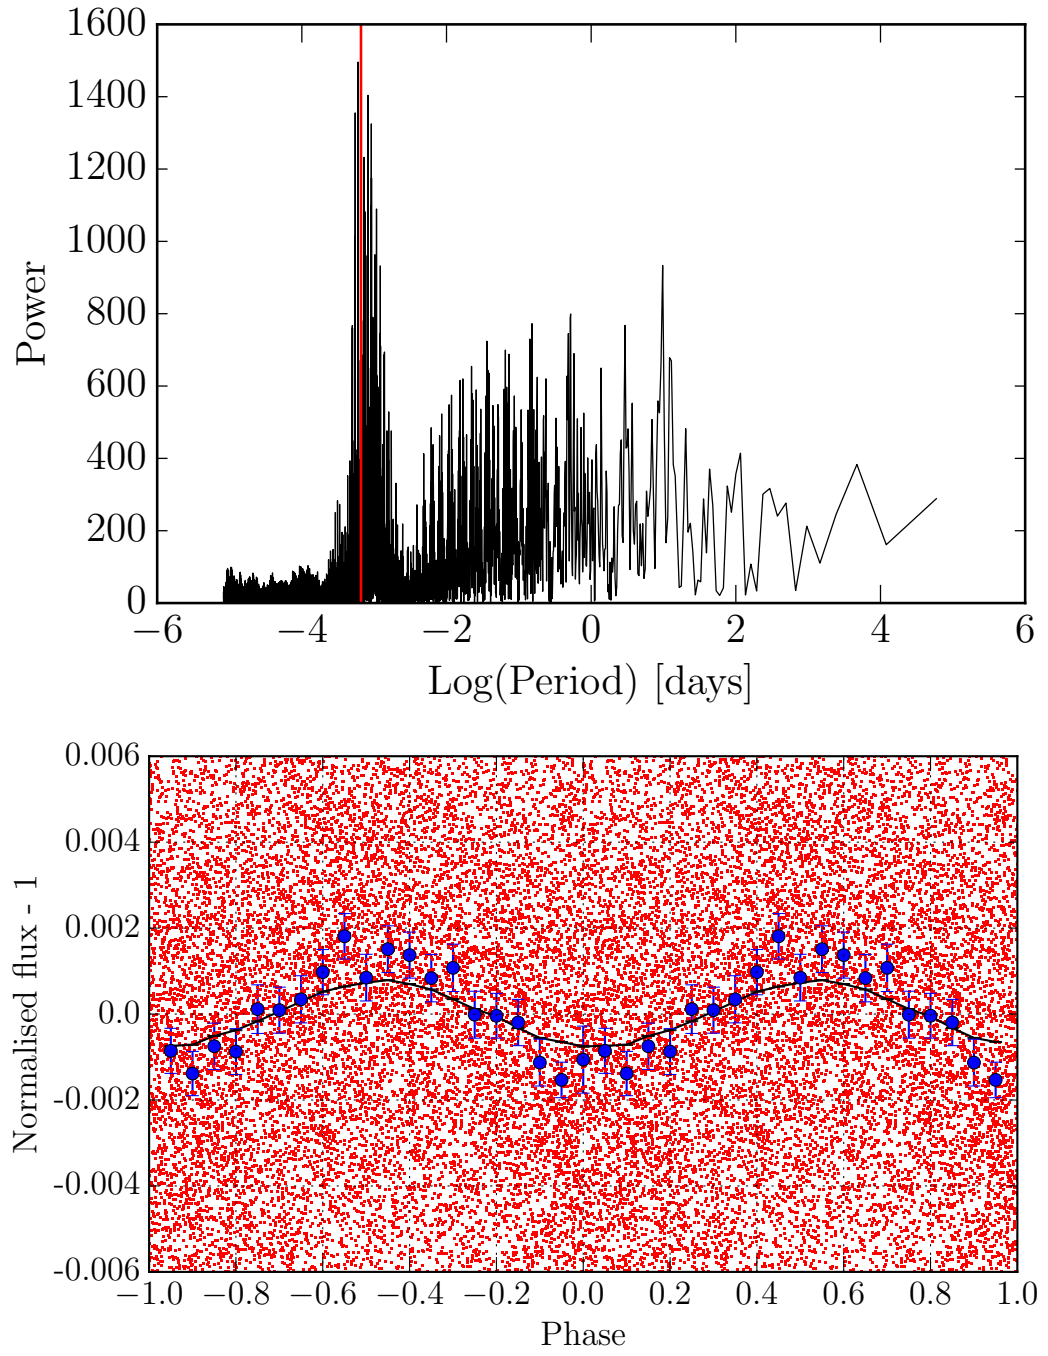


Figure 3.8: A 1 mmag delta scuti recorded with NGTS top: A periodogram of the delta scuti detected with NGTS

transit is very clear in both the individual 10 second frames and data binned by the planetary phase. The high precision of NGTS is compared to WASP in the same figure, here, several seasons of data, totaling over 40 transits are combined to reach a similar level of Signal to Noise. This demonstrates that even a single transit of a Jupiter sized planet could be convincingly detected, which could in principle lead to the detection of much longer period planets than have been detected from the ground before.

3.4.5 Quality Control plots for the Current pipeline

The pipeline in Warwick currently used to analyse NGTS data from Paranal is similar to the basic version I designed to analyze the data from Geneva, as described above. It has been extensively updated by Simon Walker, Richard West, and other members of the NGTS consortium. However, the same basic software tasks and workflow were followed for the production of early NGTS science data.

I have produced several quality control plots for the pipeline to track its performance.

Vector astrometry

In order to track the quality of the astrometric solution over time, the pipeline produces vector astrometry plots. The detected positions of the sources are marked with black points, and the vector to the catalog position of the source, multiplied by a factor of 10^6 to make it visible, is also plotted. The center of the distortion terms are also indicated with a red mark. An example of this pipeline product is given in Figure 3.5.

Focus performance

My Point Spread Function (PSF) recovery code utilises the extremely precise astrometric solution described in the previous subsection.

The image is divided into 9 sections, and in each Section the regions surrounding the brightest 100 stars are *super sampled* by a factor of ten. These regions are centered on the catalogue position for that source. The local sky background is removed, and the pixel counts are rescaled, so the peak of the star psf is at 1. These 100 super sampled regions are combined, which produces a PSF sampled with higher resolution. This works in a similar way to the Drizzle algorithm, which is often used to improve the resolution in HST data (Fruchter and Hook, 2002). A 2 dimensional gaussian is fit to the oversampled PSF in each of the 9 sectors, and the




Figure 3.9: A single partial transit of TrES-2b recorded with NGTS (top) is comparable in significance to multiple seasons of WASP data, which combine 39 transits (bottom).

semi-minor, semi-major, and rotation angle of these fits are stored for each image frame. The purpose is to check that the PSF remains stable over time, but these data could also be used to test for, e.g., the tilt in the focal plane. An example output plot is shown in Figure 3.10.

Photometric performance

If the noise in the lightcurves behaves as white noise, then binning together multiple points will decrease the RMS as \sqrt{N} , where N is the number of points. At some point, the red noise limit will be reached. Red noise does not have a frequency dependence, so will not be reduced by combining more frames. Quantifying the levels of white and red noise in the data is an important diagnostic, so a useful plot to generate for a set of data is the RMS level as a function of the number of binned points.

Due to the instrument response, it is possible that there are different red and white noise levels as a function of star brightness. An example of this would be stars bright enough to be saturated, which will have an additional source of red noise from non conservation of charge. Therefore, the diagnostic plot can be made more useful by separating the stars into brightness bins. An example of this diagnostic plot from recent NGTS data is shown in 3.11, both before and after detrending.

The apparently very high and consistent RMS level in the first plot is due to airmass trends in the data, as this task covered an entire nights worth of data. As a result, every star has an airmass curve, with an amplitude of several 10s of mmags that cannot be reduced by binning, so appears as a red noise source. The airmass curve is one of the first trends to be removed by SYSREM, so in the second plot the RMS is allowed to decrease much further with binning, in this case most brightness bins reaching levels of a few millimag on timescales of ~ 100 minutes, except for the brightest bin, which flattens off earlier due to saturation effects.

3.5 Conclusions

In this Section I have summarized some of my contributions to the ongoing NGTS project.

I demonstrated with my analysis of the Geneva data gathered with an NGTS unit in June-August of 2013 that the instrument was performing well, capable of regularly reaching precisions better than 10 mmag in 10 second exposures on bright stars. The data are white noise dominated and stable over long timescales, as demonstrated by the lightcurves of variable stars, including a short period delta-

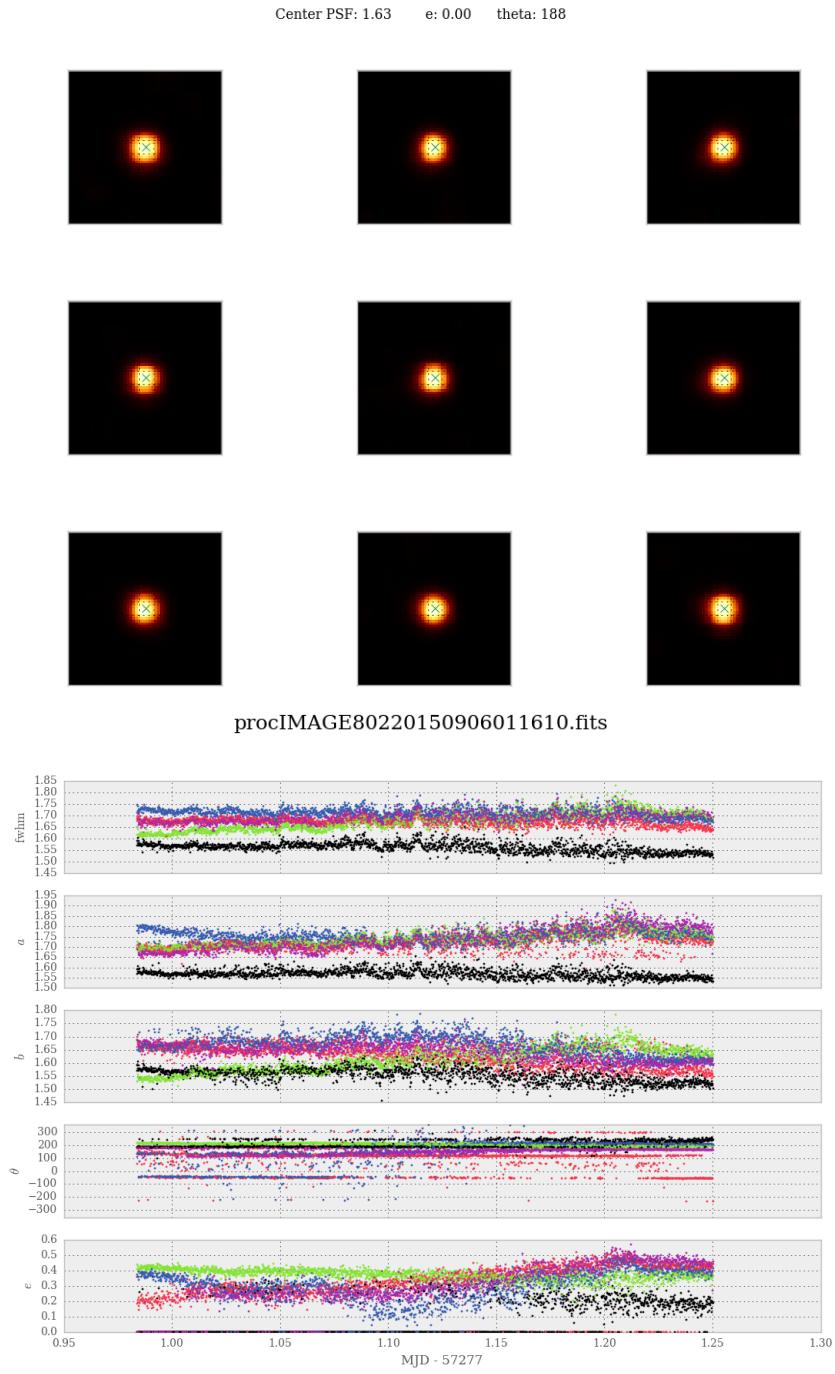


Figure 3.10: Top: Recovered PSFs for the 9 image sections. Bottom: evolution of the PSF fit parameters over a typical observing night.




Figure 3.11: An RMS binning plot divided by magnitude bins, from red (faintest), to yellow (brightest). Top: raw data Bottom: after detrending by SYSREM. Data is from one night in Paranal on 06/09/15. Stars near saturation (yellow) do not improve with binning. Dashed lines indicate how data would improve if noise was purely white. The precision on this night was limited to a few mmag by red noise from weather conditions.

scuti with an amplitude of 1 mmag. I demonstrated that the detector could detect an exoplanet transit at very high significance, with a partial transit of TrES-2b being comparable to several seasons of WASP data. These are the highest precision lightcurves ever achieved in a wide-field ground-based survey. My validation of the CASU photometry tools led to them forming the core of the current development version of the NGTS production pipeline.

Chapter 4

Reconstructing the high energy irradiation of the evaporating hot Jupiter HD 209458b

So what happened to the atmosphere of Mars? I'll quote Bob Dylan: 'The answer my friend is blowing in the wind'

Michael Meyer, 2015

The atmosphere of the exoplanet HD 209458b is undergoing sustained mass loss, believed to be caused by X-ray and extreme-ultraviolet (XUV) irradiation from its star. The majority of this flux is not directly observable due to interstellar absorption, but is required in order to correctly model the photo-evaporation of the planet and photo-ionisation of the outflow. In this Chapter I present a recovered high energy spectrum for HD 209458 using a Differential Emission Measure (DEM) retrieval technique. I construct a model of the stellar corona and transition region for temperatures between $10^{4.1}$ and 10^8 K which is constrained jointly by ultraviolet line strengths measured with the Cosmic Origins Spectrograph (COS) on the *Hubble Space Telescope* (*HST*) and X-ray flux measurements from *XMM-Newton*. The total hydrogen ionising luminosity ($\lambda < 912$ Å) is found to be $10^{28.26}$ erg s⁻¹, which is similar to the value for the mean activity level of the Sun. This luminosity is incompatible with energy limited planetary mass loss rates estimated from the same COS dataset, even the lower bound requires an uncomfortably high energetic efficiency of >40%. However, my luminosity is compatible with early estimates of

the mass loss rate of HD 209458b based on results from the *HST* Space Telescope Imaging Spectrograph (STIS). Precisely reconstructed XUV irradiation is a key input to determining mass loss rates and efficiencies for exoplanet atmospheres.

4.1 Introduction

Observations of the transit of HD 209458b in the ultraviolet show that the planet occults a much larger fraction of the star than in the optical, with a 15% loss of flux reported in the Ly α line with *Hubble Space Telescope* (*HST*) Space Telescope Imaging Spectrograph (STIS) (Vidal-Madjar et al., 2003). This implies a cloud of optically thick neutral hydrogen that extends further than the Roche lobe of the planet, indicating a significant exosphere of escaping hydrogen. The presence of such an extended atmosphere, where it would be vulnerable to radiation pressure and stellar wind, implies a mass-loss rate on the order of 10^{10} g s^{-1} . Further strengthening this conclusion, the occulted Ly α line showed evidence of velocity structure, with velocities between -130 and +100 km s $^{-1}$ having the strongest signal. These observations have been repeated with lower resolution instruments (Vidal-Madjar et al., 2004; Ehrenreich et al., 2008; Vidal-Madjar et al., 2008), whose results are consistent with this conclusion, though they were unable to improve the mass-loss estimates or the velocity structure of the wind.

The exosphere has also been shown to contain heavier elements, O I and C II absorption have both been detected in transit (Vidal-Madjar et al., 2004). This indicates a much more dramatic mass-loss, since heavier elements should not be present in detectable quantities this high in the exosphere through molecular diffusion alone. Rather than Jeans escape, a more plausible scenario is a hydrodynamic blow off, sweeping the heavier elements with it into space. Using the then newly installed Cosmic Origins Spectrograph (COS) instrument on the *HST* Linsky et al. (2010) also detected a significant outflow of ionized carbon and silicon, with transit depths of $\sim 8\%$. Using a spherically-symmetric evaporation model, they inferred a mass-loss rate between 8 and $40 \times 10^{10} \text{ g s}^{-1}$. This higher evaporation rate uses several independent lines and avoids Ly α which suffers from both heavy interstellar medium absorption at low velocities and strong time varying geocoronal contamination, particularly with the slitless COS instrument. Whilst their carbon transit depth was consistent with that of Vidal-Madjar et al. (2004), their detection of silicon absorption was inconsistent at the 2σ level. Revisiting the system in the FUV, Ballester and Ben-Jaffel (2015) find variability in the stellar emission of Si III and C II which was not accounted for in Linsky et al. (2010). They find no evidence

for Si III absorption, and slightly smaller absorption of C II, so the mass-loss rate is most likely smaller than that reported by Linsky et al. (2010).

There is a consensus that HD 209458b is undergoing large-scale atmospheric evaporation, believed to be caused by X-ray and EUV (XUV) heating of the planetary atmosphere (Lammer et al., 2003). The optical flux of the star is not capable of heating the atmosphere to the $\sim 10,000$ K temperature required to outflow from the planet, but XUV photons from the stellar corona and transition region have enough energy to ionize a hydrogen atom ($\lambda < 912$ Å), releasing cascades of electrons that quickly heat the gas to extremely high temperatures. Unfortunately the neutral and molecular hydrogen of the ISM is very efficient at blocking radiation at these wavelengths, so it is very difficult to observe on all but the very closest of stars. There is also a current lack of space based observatories capable of imaging in this wavelength region.

X-ray observations can characterize the hotter regions of a stellar corona, but for solar type stars over 90% of the ionising flux is emitted in the EUV ($100 < \lambda < 912$ Å) which is mainly emitted from cooler regions, between $10^{4.5}$ and $10^{6.5}$ K. In the absence of direct observations of the EUV flux of the star, a model of the corona is necessary for the full high energy spectrum to be recovered. Fortunately, high excitation ultraviolet lines exist that are formed in the same temperature regions of the corona as the EUV flux and can be used as a proxy for ionising flux. These lines have high ionisation states, so are only minimally effected by ISM absorption. I combine measurements of these UV lines with *XMM-Newton* observations to construct a model of the corona which I use to recover the entire high energy spectrum of HD 209458.

Linsky et al. (2014) present a technique for calculating the EUV luminosity of stars by scaling Ly α luminosity. Since for many exoplanet host stars the majority of the Ly α flux will be absorbed by the interstellar medium, it has to be reconstructed, which could introduce uncertainties to this technique. Although their relationship is based on a small sample that includes only 4 G class stars, and despite the uncertainties of Ly α flux reconstruction, this method can be a useful tool for estimating the EUV fluxes of planet-hosting stars on a population level, but the $\sim 30\%$ dispersion of stars around their fit means DEM recovery is more appropriate for detailed analysis of individual evaporating systems such as HD 209458b. An accurate determination of the ionizing flux is required to calculate the heating of the planet exosphere, and also more accurately determine the mass-loss rate, since the photo-ionisation rate and mass loss rate are highly correlated in transit observations (Bourrier and Lecavelier des Etangs, 2013).




Figure 4.1: Relative emission of ionising radiation as a function of plasma temperature. Dotted line - X-ray flux (5–100 Å), Dashed line - EUV flux (100–920 Å), solid line - Total/XUV (5–920 Å)

Accurate calculations of atmosphere heating and mass loss rates are required in order to constrain the energetic efficiency of exoplanet evaporation, which is important, as evaporation may have effects on the observed population of exoplanets (e.g. Davis and Wheatley, 2009; Jackson et al., 2012).

4.2 Observations

Ultraviolet data is obtained from France et al. (2010), who use data from the *HST* COS instrument (*HST* GTO program 11534). Lines have been selected from their line list which are uncontaminated, have formation temperatures $> 10^{4.1}$ K, and which appear in the CHIANTI (v8.0) spectral database (Dere et al., 1997; Del Zanna et al., 2015). The resulting line list is shown in table 4.1.

HD 209458 is listed as a detected source in the 3XMM DR5 catalogue (detid 104044501010110). The detection is based on observations taken in 2006 (Obs. ID 0404450101, PI Wheatley). The observations were made with the MEDIUM filters on all cameras, the PN was in FULL FRAME mode and the MOS cameras were in LARGE WINDOW mode to eliminate optical loading. The source detection is significant at over 3σ in the 0.2 - 0.5 keV band, but is not significantly detected at higher energies, this is not an unexpected characteristic for a stellar source, since the flux will be dominated by soft X-rays. As I report in Section 6.4, the corresponding

Table 4.1: line strengths used in this work, obtained from France et al. (2010), the max formation temperature is calculated using CHIANTI (v8.0)

Species	λ_{rest} (Å)	Line Flux (10^{-16} erg cm $^{-2}$ s $^{-1}$)	T max
Si III	1206.50	12.89 ± 0.35	4.8
O V	1218.34	3.97 ± 0.15	5.4
N V	1238.82	1.08 ± 0.58	5.3
N V	1242.80	0.26 ± 0.04	5.3
N V	1242.80	1.30 ± 1.07	5.3
S II	1259.52	1.68 ± 0.54	4.4
Si II	1264.74	2.11 ± 0.11	4.1
Si III	1298.95	0.63 ± 0.09	4.8
Si II	1309.28	3.15 ± 0.48	4.1
C II	1334.53	8.04 ± 0.25	4.3
C II	1335.71	16.06 ± 0.14	4.3
Si IV	1393.76	9.47 ± 0.40	4.9
Si IV	1402.77	5.11 ± 0.49	4.9
Si II	1526.71	1.19 ± 0.072	4.1
Si II	1533.43	2.31 ± 0.14	4.1
C IV	1548.19	10.54 ± 0.38	5.1
C IV	1550.77	7.07 ± 0.96	5.1
Fe II	1559.08	1.98 ± 0.29	4.1
Fe II	1563.79	1.38 ± 0.74	4.1
Fe II	1570.24	0.64 ± 0.40	4.1

flux is significantly higher than the upper limits reported by Sanz-Forcada et al. (2011), who analysed the same data. I am unsure of the cause of the discrepancy.

Previous claims of the detection of any flux of HD 209458 (Penz et al., 2008; Kashyap et al., 2008) are based on earlier *XMM-Newton* data that were taken with the THIN optical blocking filters (Obs. ID 0130920101, PI Bertaux). For a bright target like HD 209458b this can lead to contamination of the X-ray image by optical loading. I have analysed the spectrum of the earlier data and confirmed that the X-ray spectrum is unphysically soft, which is characteristic of optical contamination.

I use the reported count rates for the PN and MOS instruments in the five energy bands defined in the 3XMM catalogue.

4.3 Method

4.3.1 Differential Emission Measure

The Differential Emission Measure (DEM) is a powerful diagnostic for an optically thin emitting plasma, such as a stellar corona. Essentially, as a measure of electron density along all lines of sight as a function of temperature, it can be used to characterize the atmosphere of a star. The DEM is defined as

$$\varphi(T) = N_e^2 \frac{dh}{dT} \quad (4.1)$$

Where N_e is the electron number density, T is the plasma temperature and h is the distance taken along the line of sight. The DEM is directly proportional to the strength of an emission line at a given temperature, so with a wide enough baseline of high temperature lines a near continuous DEM can be constructed. Combining a DEM with a suitable plasma model can recover the entire high energy spectrum, including continuum processes. In this work I use CHIANTI (v8.0) (Dere et al., 1997; Del Zanna et al., 2015) to model the plasma emission. HD 209458 is a G class star with an $[\text{Fe}/\text{H}]$ of 0.02 ± 0.03 (Santos et al., 2004), so I assume the corona has solar abundances (Schmelz et al., 2012).

Figure 4.1 shows the relative strength of hydrogen ionising flux ($\lambda < 912 \text{ \AA}$) emitted from the plasma as a function of temperature. It demonstrates the temperature regions that are most important to the level of ionizing radiation, and hence exospheric heating. The most potent temperatures are between $\approx 10^{4.75}$ and $10^{6.5}$ K, with about half the integrated area $< \sim 10^{5.7}$ K. This is the temperature region which can be well characterized and constrained by lines in the UV, such as those measured by France et al. (2010), listed in Table 4.1. Above $10^{5.7}$ K the strongest constraint is the soft X-ray flux. Combining these two constraints can allow the whole temperature range to be characterized.

Model spectra are folded through the instrument Response Matrix calculated using the SCIENCE ANALYSIS SYSTEM (SAS) (Gabriel et al., 2004). At the low fluxes under consideration pile-up is not an issue, so the relation between absolute flux level and count level should be linear. To simulate the detectable flux from the Earth I simulate interstellar medium absorption using the Tübingen-Boulder absorption model (Wilms et al., 2000) generated in XSPEC (Arnaud, 1996). I use a H I column density $N_H = 2.3 \times 10^{18} \text{ cm}^{-2}$, which was calculated by fitting the Ly α profile by Wood et al. (2005).




Figure 4.2: The best-fitting DEM (red) with 3σ limits (red dotted). At the top left, the coloured lines are the emission loci, representing the constraints from the individual species contributing to the UV spectrum. The red squares are the average of the emission loci over a 0.3 dex temperature range. The emission loci and X-ray flux limits are the emission measure required for a single temperature model, and represent hard limits to the DEM. Band 1 is 0.2–0.5 keV. Since no X-ray flux was detected in *XMM-Newton* band 2 (0.5–1.0 keV) or band 3 (1.0–2.0 keV) the 3σ upper limit is plotted. The higher energy bands 4 and 5 do not impose any additional constraints, so are not plotted. A quiet solar DEM is plotted for comparison (grey dashed) (Kretzschmar et al., 2004)

4.3.2 Fitting procedure

I assume that the DEM of the star can be described by a smoothly varying function, which appears to be an adequate description for the Sun and other well studied stars. For my model I use a 4th order Chebyshev polynomial. The DEM is used to calculate the UV line strengths and scale the contribution from 40 pre-calculated model spectra, log evenly spaced in the temperature range $10^{4.1}$ - 10^8 K. The model spectra and *XMM-Newton* Response Matrixes are used to give a simulated *XMM-Newton* response for each instrument and each temperature. The count rates in five energy bands are calculated: 0.2–0.5 keV, 0.5–1.0 keV, 1.0–2.0 keV, 2.0–4.5 keV and 4.5–12 keV.

The line strengths and *XMM-Newton* count rates are compared to the measured values and errors to produce a residuals vector. In X-ray bands where no detection was made I use a value of 0 and the count errors from the background rate.

A Levenbert Marquardt process is initially used to solve the minimization problem and find a best fit solution. I use an MCMC to calculate the formal errors on the DEM and luminosity in all energy bands of interest. I use the emcee implementation in Python (Foreman-Mackey et al., 2012). I use a burn in of 10^6 steps to ensure the solution has converged and is well behaved, and use a further 10^7 steps for parameter estimation.

I use the stellar and planetary parameters reported by Southworth (2010). I choose not to use parameters derived from the interferometrically measured stellar radius from Boyajian et al. (2014), as the authors caution that the angular size of HD 209458 is close to the resolution limit of the instrument, so may be effected by systematics. I use the distance of 49.63 ± 1.97 pc from van Leeuwen (2007), and include the distance errors in my reported luminosities.

4.4 Results

The best-fitting DEM is shown in Fig. 4.2, with loci for each of the lines included in the fit, and the X-ray flux limit. The loci represent the value required to reproduce the line strength with a single temperature, so form an upper limit to the possible DEM. The spectrum in Fig. 6.5 is generated from the best fit DEM. The XUV luminosity (defined as the integrated flux between 5 and 920 Å) is $28.26^{+0.05}_{-0.05} \log(\text{erg s}^{-1})$.

A 4th order polynomial was judged to provide the best fit to the data, higher

Table 4.2: Numerical values for my recovered differential emission measure for HD 209458b

Log(T)	Log(DEM)
4.1	25.7
4.2	24.8
4.3	24.1
4.4	23.5
4.5	23.0
4.6	22.5
4.7	22.2
4.8	22.0
4.9	21.8
5.0	21.6
5.1	21.5
5.2	21.5
5.3	21.4
5.4	21.4
5.5	21.4
5.6	21.4
5.7	21.4
5.8	21.4
5.9	21.3
6.0	21.2
6.1	21.1
6.2	20.9
6.3	20.6
6.4	20.3
6.5	19.8
6.6	19.3
6.7	18.7
6.8	17.9
6.9	17.0
7.0	16.0
7.1	14.9
7.2	13.6
7.3	12.1
7.4	10.5
7.5	8.6
7.6	6.6
7.7	4.4
7.8	2.0
7.9	-0.6
8.0	-3.5

orders became flexible enough to favour unphysical plasma distributions with multiple sharp peaks and drop-offs with temperature that did not resemble typical DEMs for the Sun or other solar type stars.

The electron density of the chromosphere/corona is uncertain; however, the major UV lines and broadband X-ray flux in this temperature range are only weakly dependent on plasma density. Some lines are susceptible to collisional de-excitation, but they are not important to the broadband fluxes. For my best-fitting model I assumed an electron density typical to the corona of the Sun, $\sim 10^8 \text{ cm}^{-3}$ (e.g. Fludra et al., 1999). I find that varying the electron density over 5 orders of magnitude, between 10^6 and 10^{11} cm^{-3} has no significant effect on the total XUV flux of the model, all values agreeing to within 1σ with the reported fluxes for 10^8 cm^{-3} .

I also investigate the sensitivity of my results to errors in abundances. The metal abundances of the corona of the Sun are significantly different to the abundances of the photosphere. Different atmospheric regions of the Sun have enhanced abundances of some elements due to the First Ionisation Potential (FIP) effect (Feldman and Widing, 2002). In the Sun, atoms with $\text{FIP} > 11 \text{ eV}$ have approximately the same abundance throughout the atmosphere, while low FIP elements with $\text{FIP} < 10 \text{ eV}$ can be enhanced by factors of 3–5 in the hotter regions of the corona. HD 209458 has a similar spectral type and activity level to the Sun, so would be expected to show a FIP effect of comparable strength (Laming, 2004) but spatially resolved measurements of the atmospheric composition are unavailable. For my recovered spectrum I find that 43% of the emitted XUV flux is from ions of iron alone, which has a measured photospheric abundance that is within 1σ of the solar value (Santos et al., 2004), giving confidence that the important coronal abundances are likely to be very similar to those of the Sun. The next highest contributor to XUV flux is oxygen, with 20% of the total flux. Oxygen is a high FIP element and has little scatter in abundance for solar type stars with solar $[\text{Fe}/\text{H}]$ (Mena et al., 2015). All other elements contribute significantly less than oxygen, so any errors in abundances either intrinsic or due to the FIP effect are unlikely to alter my total XUV flux by a significant amount.

4.5 Discussion

4.5.1 Comparison to previous work

The recovered DEM appears qualitatively similar to DEM measured for the quiet Sun, to demonstrate this I overplot the DEM found by (Kretzschmar et al., 2004)




Figure 4.3: A recovered high energy spectrum for HD 209458, scaled to the orbital distance of 0.047 AU.

Table 4.3: The flux of my recovered EUV spectrum in several bands, compared to the Sun. F_{XUV} is the flux received by HD 209458b at 0.047 AU

Spectrum	Log L_X (erg s ⁻¹)	Log L_{EUV} (erg s ⁻¹)	Log F_{XUV} (erg s ⁻¹ cm ⁻²)
Mean Sun	27.38	27.93	3.28
Quiet Sun	26.73	27.70	2.99
HD 209458 ¹	<26.40	<27.74	<3.06
HD 209458 ²	27.08 ^{+0.07} _{-0.07}	28.23 ^{+0.05} _{-0.05}	3.50 ^{+0.05} _{-0.05}

¹ Sanz-Forcada et al. (2011)

² This work.

in Figure 4.2. The resemblance is quite striking, particularly in the low and high-temperature regions. In order to make a direct comparison of the fluxes, I generate a synthetic spectrum from this quiet Sun DEM using the same plasma model as for HD 209458. Integrating the synthetic quiet Sun spectrum in the ROSAT bandpass (5–125 Å) gives a log flux of 26.8, which matches the solar minimum activity value of 26.8 reported for the ROSAT band reported by Judge et al. (2003). I also compare my fluxes to those of the mean solar irradiance spectrum reported by Ribas et al. (2005). I find that the X-ray flux falls between the quiet and mean Sun, but the total ionising flux (5–920 Å) flux of HD 209458 is slightly higher than that of the mean Sun, by 0.2 dex.

Sanz-Forcada et al. (2011) estimate the XUV flux of HD 209458 by fitting an emission measure model to an X-ray upper limit measured from the same dataset, but with a broader energy band, and extrapolating to lower temperatures. I find that their results for total high energy flux are not consistent with mine. They estimate L_X (5–100 Å) and L_{EUV} (100–920 Å) upper limits of 26.40 and 27.74, which are significantly lower than the results from my recovered spectra, with logarithmic fluxes of $27.08^{+0.07}_{-0.07}$ and $28.23^{+0.05}_{-0.05}$. My results agree on the ratio of X-ray to EUV, which suggests my coronal models are compatible. The absolute ~ 0.6 dex offset is due to their very low upper limit on the X-ray flux, which is significantly fainter than the flux reported in the 3XMM catalogue. Estimating the luminosity directly from the reported 3XMM flux gives a 0.2–12 keV luminosity of 26.8. This bandpass does not include all the soft X-ray flux of this very soft source, which is why my final recovered flux is slightly higher. The 3XMM source is at the expected position of HD 209458b, and there are no other spurious sources detected nearby, so I believe this detection is genuine and that the X-ray flux is significantly higher than reported by Sanz-Forcada et al. (2011).

In my model 93% of the XUV radiation is in the EUV band. Owen and Jackson (2012) compare the effects of X-ray and EUV radiation on the evaporation of exoplanets, and find that for planets like HD 209458b, EUV radiation is dominant for atmosphere erosion, though the location and efficiency of absorption is different compared to the X-ray flux. Ehrenreich and Désert (2011) calculate mass-loss efficiencies for HD 209458b and HD 189733b, however, they consider only X-ray flux in their calculations, which neglects the majority of the total ionising flux.

The STIS measurements of deep metal absorption presented in Vidal-Madjar et al. (2004) are unresolved, so there is no velocity information, but the thermal broadening of gas at 10,000 K would not be sufficient to absorb such a high fraction of the relatively broad stellar line profiles. Koskinen et al. (2013a) suggest that this

can be explained by a metallicity enhancement of a factor of 3–5 in the atmosphere of HD 209458b, accompanied by the XUV flux being greater than a factor of 2 higher than the mean solar flux, which is not consistent with my recovered XUV flux. This discrepancy may be a sign that superthermal processes are at work in the atmosphere of HD 209458b (Ballester and Ben-Jaffel, 2015).

Guo and Ben-Jaffel (2015) study the effects of spectral energy distribution of EUV radiation on the structure of hot Jupiter atmospheres and their evaporation. They define an EUV spectral index β which is the ratio of the flux in a 50–400 Å band to the flux in a 50–900 Å band. The value of β for my recovered stellar spectrum is $0.64^{+0.02}_{-0.02}$. My synthetic quiet Sun spectrum gives a value of beta of 0.74, which matches observed spectra of the Sun during periods of low activity. As a G type star, it is perhaps not surprising that my recovered spectrum for HD 209458 has a qualitatively similar spectral shape to that of the quiet Sun. Guo and Ben-Jaffel (2015) show that the location of the H/H⁺ transition in the atmosphere outflow of HD 209458b depends strongly on the spectral index β . The location of this transition region itself is important to calculating the structure of the wind. My measurement of this value is an important input to any future models of the evaporation and structure of HD 209458b’s atmosphere, both for driving escape and constraining the photo-ionisation of the wind.

4.5.2 Energy-limited mass-loss

Energy-limited mass-loss can be used to calculate an upper limit to the evaporation rate of an exoplanet. The only assumption is that some fraction of incoming high-energy flux, η , is converted into work done against the gravitational potential of the planet (Watson et al., 1981). Using the formulation from Erkaev et al. (2007), the equation for energy-limited mass-loss is:

$$\dot{M} = \frac{\eta \pi F_{XUV} \alpha^2 R_P^3}{G M_p K} \quad (4.2)$$

$$K = \left(1 - \frac{3}{2\xi} + \frac{1}{2\xi^3} \right) \quad (4.3)$$

$$\xi = \frac{R_{Hill}}{R_p} \quad (4.4)$$

where F_{XUV} is the X-ray and EUV flux incident on the planet, G is the gravitational constant and M_p is the mass of the planet. The formula includes a K factor to account for the reduced energy required to reach the Hill sphere, R_{Hill}

where the gas is effectively unbound from the planet. For a close-in hot Jupiter like HD 209458b, this factor can increase the predicted mass-loss rate by upwards of 50%. The α factor accounts for the increased cross section of the planet to XUV radiation, since it is typically absorbed very high in the atmosphere. This value can be derived using a simple analytical structure model of the atmosphere (Murray-Clay et al., 2009). The optical depth of 1 for incoming EUV radiation will occur at approximately 1 nanobar, which is 20 scale heights above the 1 bar optical radius of the planet. Assuming the dayside atmosphere is approximately isothermal puts the XUV absorption radius at $\sim 1.1R_p$, making $\alpha \sim 1.1$. X-rays are absorbed deeper in the atmosphere, but they make up a comparatively small contribution to the flux.

The mass-loss efficiency, η , is unknown, which is a major limitation for models of mass loss. There are calculations of the expected value of η , but they are far from precise. Based on modelling of Hot Neptunes, Owen and Jackson (2012) suggest that η should be between 0.05 and 0.2. At very high fluxes, the energy transfer transitions to a recombination limited regime where mass-loss scales with $F_{XUV}^{1/2}$, and the efficiency decreases, since a greater proportion of energy is being lost to recombination Murray-Clay et al. (2009). Shematovich et al. (2014) simulated absorption of XUV radiation in the lower thermosphere of HD 209458b and find $0.10 < \eta < 0.15$. Salz et al. (2016) perform hydrodynamic simulations of mass-loss and XUV absorption up to the Roche-lobe height, and find $\eta = 0.21$ for the case of HD 209458b, though the exact efficiency depends on the input spectrum. Koskinen et al. (2013b) find higher heating efficiencies of 0.4–0.6 for XUV fluxes ~ 450 times today’s solar flux, however, for planets with lower incident XUV flux, IR cooling from photochemically produced H_3^+ becomes important (Koskinen et al., 2007), and can decrease heating efficiency to closer to 0.15. An efficiency of 0.15 has also been found to be appropriate for the case of Venus (Chassefière, 1996). The planet’s magnetic field could also effect the efficiency, and will require magnetohydrodynamic simulations to model (e.g. Owen and Adams, 2014). A recovered high energy spectrum will allow better estimates both of the input energy and mass loss rates of the evaporating atmosphere, improving the observational constraints on η .

Assuming $\eta = 0.2$ my best-fitting DEM incurs a mass-loss rate of $3.8 \pm 0.2 \times 10^{10} \text{ g s}^{-1}$, compatible with the reported values of $\sim 10^{10} \text{ g s}^{-1}$ in Vidal-Madjar et al. (2003, 2004) and the models of Owen and Jackson (2012), but significantly lower than the $8\text{--}40 \times 10^{10} \text{ g s}^{-1}$ value reported by Linsky et al. (2010). The efficiency would have to be at least 0.42 to reach the lower end of this range, and the higher value cannot be reached without exceeding 100% efficiency, which is clearly unphysical. This implies that the mass-loss rates have perhaps been systemati-

cally overestimated, as Ballester and Ben-Jaffel (2015) suggest. I emphasise that a strength of my analysis is that the ultraviolet line strengths I use in my analysis are taken from the same dataset that is used to calculate the planet mass-loss rates, so stellar variability should not be a concern.

In reality, energy-limited mass-loss ignores a number of important physical effects, and should be regarded as an upper limit. 3D models are required to properly model higher order features such as the shape of the outflow and self-shielding effects, but require an accurate XUV spectrum as a model input. Bourrier and Lecavelier des Etangs (2013) show that the estimated hydrogen escape rate and the level of ionising flux are highly related. Since only neutral hydrogen is visible in the Ly α absorption, it is important to understand the ionising flux in order to correctly infer the true mass-loss rate. The ionizing flux I measure for HD 209458 is $\sim 3\times$ higher than the solar minimum value, so in their model is most consistent with a mass-loss rate of 10^9 g s^{-1} , and is inconsistent at the 2σ level with higher mass-loss rates.

4.6 Conclusions

I have constructed a synthetic coronal spectrum for HD 209458 using ultraviolet line strengths and X-ray fluxes. I find that the star is comparable in activity levels to the Sun, with a total XUV flux of $10^{28.26} \text{ ergs s}^{-1}$. Assuming a mass loss efficiency of 0.2, this could cause the atmosphere of HD 209458b to evaporate at a rate of $3.8 \pm 0.2 \times 10^{10} \text{ g s}^{-1}$, which is significantly lower than the estimated value of $8\text{--}40 \times 10^{10} \text{ g s}^{-1}$ of Linsky et al. (2010). The higher end of this range cannot be reached without exceeding 100% efficiency, which is not physical, and the efficiency would have to be at least 0.42 to reach the low end. Models of atmospheric efficiency typically do not find efficiencies greater than 0.2 (e.g. Owen and Jackson, 2012), and it is often lower. This suggests that evaporation rates may have been overestimated, due to not accounting for stellar variability Ballester and Ben-Jaffel (2015), and likely compounded by degeneracies between outflow structure/ionization and total mass (Bourrier and Lecavelier des Etangs, 2013). Modelling the outflow correctly requires an input spectrum, as the neutral hydrogen that is observed through Ly α absorption is sculpted by radiation pressure and photoionisation rates, which are dependent on both the total XUV flux, but also its spectral shape Guo and Ben-Jaffel (2015). My DEM recovery provides these important values for future models of the evaporation of HD 209458b.

Chapter 5

Rotation and Eastward winds on HD 189733b

O WILD West Wind, thou breath of Autumn's being
Thou from whose unseen presence the leaves dead
Are driven like ghosts from an enchanter fleeing

Percy Bysshe Shelley – Ode to the west wind

In this Chapter I measure wind velocities on opposite sides of the hot Jupiter HD 189733b by modeling sodium absorption in high-resolution HARPS transmission spectra. My model implicitly accounts for the Rossiter-McLaughlin effect, which I show can explain the high wind velocities suggested by previous studies. My results reveal a strong eastward motion of the atmosphere of HD 189733b, with a redshift of $2.3^{+1.3}_{-1.5}$ km s⁻¹ on the leading limb of the planet and a blueshift of $5.3^{+1.0}_{-1.4}$ km s⁻¹ on the trailing limb. These velocities can be understood as a combination of tidally locked planetary rotation and an eastward equatorial jet; closely matching the predictions of atmospheric circulation models. My results show that the sodium absorption of HD 189733b is intrinsically velocity broadened and so previous studies of the average transmission spectrum are likely to have overestimated the role of pressure broadening.

5.1 Introduction

Detections of exoplanetary winds to date have used high-resolution transmission spectra constructed from entire planetary transits to measure the average Doppler shift of atmospheric absorption around the planetary limb. In the hot Jupiter

HD 209458b a study with the Cryogenic high-resolution InfraRed Echelle Spectrograph (CRIRES) revealed a blue-shift of $2 \pm 1 \text{ km s}^{-1}$ in the average line profile of carbon monoxide absorption (Snellen et al., 2010). While in HD 189733b a high-resolution optical study with the High Accuracy Radial Velocity Planet Searcher (HARPS) found a blue-shift of $8 \pm 2 \text{ km s}^{-1}$ in the atomic sodium doublet (Wyttenbach et al., 2015) (although in this Chapter I show this value is too high). Blue-shifted absorption in the average transmission spectra indicates net wind flows from the dayside to the night side of the planets.

Measurements of the average Doppler shift, however, cannot test atmospheric circulation models of hot Jupiters, which predict very different velocities on either side of the planet. A combination of tidally locked planetary rotation and an eastward equatorial jet (Showman and Polvani, 2011; Rauscher and Menou, 2010; Perna et al., 2010) are expected to result in red-shifted absorption on the leading limb of the planet and blue-shifted absorption on the trailing limb (Showman et al., 2013; Miller-Ricci Kempton and Rauscher, 2012). The equatorial jet is believed to arise from planetary scale Rossby waves interacting with the mean atmospheric flow (Showman and Polvani, 2011), and it is likely to be a common feature of heavily irradiated close-in planets (Showman et al., 2015), controlling the redistribution of heat from the permanent dayside to the night side of the planet. While not detected directly, the existence of an eastward jet is supported by the observation that the hottest point of the atmosphere of HD 189733b is offset east of the sub-stellar point by 30° (Knutson et al., 2007b, 2009).

In this Chapter I spatially resolve the rotation of the planetary atmosphere using a time resolved model of sodium absorption during the planet transit.

5.2 Observations

I analysed data from a transit observation of HD 189733b made with the HARPS instrument (Mayor et al., 2003) mounted on the European Southern Observatory’s 3.6m telescope at La Silla, Chile. I used data from ESO programme 079.C-0127(A) taken on the night of 28 August 2007, for which the planetary sodium lines have been shown to be free of systematic contamination by Wyttenbach et al. (2015) (false-alarm probability of $< 0.01\%$). HD 189733 was observed for a period of 4 hours, with 40 spectra taken in total, covering the two-hour transit with an additional one hour on either side for baseline comparison. I accessed the reduced spectra and data products through the ESO Archive and extracted the order from the echelle spectrum that contains the sodium feature (585.0 - 591.6 nm). The spectral resolution is

115,000 (Mayor et al., 2003). As well as the sodium absorption (Wyttenbach et al., 2015), these data have been used previously to measure the stellar rotation (Collier Cameron et al., 2010), the Rossiter-McLaughlin effect (Triaud et al., 2009) and to detect Rayleigh scattering in the planetary atmosphere (Gloria et al., 2015).

The spectra were aligned in a common velocity reference frame by removing both the barycentric velocity at the time of each spectrum as recorded by the HARPS pipeline, and the line-of-sight velocity of the star using the system orbital parameters (Triaud et al., 2009). This ensured that misalignment of the stellar line profiles did not introduce spurious features to the transmission spectra.

Weak absorption features in the Earth’s atmosphere were apparent in the raw spectra, and I took steps to remove the time varying component of the telluric absorption. All the science frames were combined into a single very high signal to noise spectrum, and this master spectrum was then subtracted from each individual frame, exposing the sources of variability. I created a model telluric spectrum from a detailed line list in the region surrounding the sodium doublet, which includes the relative strengths of all lines (Lundstrum et al., 1991). The telluric template spectrum was then fitted to each difference spectrum, with the FWHM, velocity offset and absolute strength of the spectrum as the free parameters. It was found that the quality of the fit could be improved by allowing atmospheric water lines to have a separate scaling factor. This model was then subtracted from the science frames, removing the variability in the telluric lines.

The sodium absorption features in the Earth’s atmosphere are weak in comparison to water features, and are offset from the stellar spectrum by the barycentric velocity of 9 km s^{-1} during the observation window. They are therefore unable to mimic a planetary transmission signal in my data.

The individual spectra were then normalised in order to remove the absolute depth of the planetary transit, leaving only the relative transmission as a function of wavelength.

I confirmed the presence of sodium absorption during transit by repeating the bootstrapping analysis of Wyttenbach et al. (2015), finding a fractional increase in transit depth in a narrow (1.2 nm) band around the sodium doublet of 0.00077 ± 0.00013 , a highly significant ($> 5\sigma$) detection that is consistent with the previous findings of Wyttenbach et al. (2015), Huitson et al. (2012), Jensen et al. (2011), and Redfield et al. (2008). As a false positive check, I repeated this technique on the hydrogen alpha line at 656.28 nm. Similar to sodium, hydrogen alpha has a strong stellar feature, however, planetary atmosphere models do not predict that this feature should show a significant excess in transit. Indeed, I do not detect a

significant transit excess using the same technique on this line.

5.3 Transit model

I search for differences in the Doppler shift of absorption from opposite sides of the planet by modeling the planetary absorption line profiles as a function of time through the transit. The power of my model to distinguish between atmospheric velocities at different points on the planetary limb is illustrated in Figure 5.1 where I compare the stellar flux available for absorption by atmospheric sodium on opposite sides of the planet. At each phase the stellar flux was sampled at the wavelengths corresponding to the sodium lines at the orbital velocity of the planet. During the transit, the orbital velocity of the planet causes the planetary lines to shift from blue to red by $\pm 16 \text{ km s}^{-1}$, while the rotational velocity of the star as sampled by the planet varies from blue to red by $\pm 3 \text{ km s}^{-1}$. The planetary sodium lines overtake the stellar lines at mid-transit, accounting for the relatively low weighting of this phase in Figure 5.1 (when the planetary and stellar velocity align there is little stellar flux for the planetary atmosphere to absorb). It is apparent from Figure 5.1 that with sufficiently high signal-to-noise data it is possible to model variations in the planetary absorption line profile to spatially resolve velocities in the planetary atmosphere.

The star-planet system was simulated on a 2D pixel grid, which implicitly accounts for both the atmospheric absorption and the Rossiter-McLaughlin effect. A quadratic limb darkening profile was used to give each pixel on the star a relative intensity (Claret and Bloemen, 2011). I adopted a stellar rotational velocity of 3.1 km s^{-1} , which was measured using Doppler tomography from the same set of HARPS observations (Collier Cameron et al., 2010). Solid body rotation of the star was assumed in order to derive the rotation speed, but if this assumption is not correct the rotation velocity of the transit chord will not be affected.

The absorption by the planet is modeled as an opaque occulting disk with the white light radius ratio of HD 189733b. The atmospheric absorption is represented as an additional layer with absorption of the sodium doublet modeled by Voigt profiles. The atmospheric dynamics are parameterized by the equatorial velocities on the leading and trailing limbs of the planet together with the assumption of constant angular velocity (rigid body rotation). This allows continuous values of velocity to be defined at all grid points representing the planetary atmosphere with just two free parameters, while allowing higher velocities at the equator than the poles. This parameterization was selected in order to allow a reasonable approximation to




Figure 5.1: *Top Panel:* A scale diagram of the star-planet geometry, with the thickness of the planetary atmosphere doubled for clarity. The planet moves from left to right as it transits the star, with the leading limb coloured red and the trailing limb coloured blue. *Bottom Panel:* My model illustrates how the background illumination of the planetary limb changes during transit, allowing absorption by different parts of the planetary atmosphere to be separated. The red/blue lines represent the stellar flux available for sodium absorption at the planetary equator on the leading/trailing limbs (the model accounts for stellar rotation and limb darkening, and the planetary orbit). The asymmetry in illumination is strongest during ingress and egress, where absorption is also most heavily weighted (due to the relative Doppler shifts of the planetary and stellar sodium lines).

Table 5.1: Best fit parameters for the stellar model used for HD 189733b. T_{eff} is the effective temperature, $\log g$ is the surface gravity, v_{mic} and v_{mac} are the micro and macro turbulence velocities, $v \sin i$ is the line of sight velocity of the star, ld is a linear limb darkening parameter, and R is the spectral resolving power of the instrument.

Parameter	Value	Units
T_{eff}	4871.67	K
$\log g$	4.47	$\log(\text{cm/s}^2)$
MH	-0.26	-
v_{mic}	2.22	km/s
v_{mac}	5.47	km/s
$v \sin i$	3.10	km/s
ld	0.6	-
R	115000	-

the predictions of planetary rotation and an equatorial jet (Showman et al., 2013; Miller-Ricci Kempton and Rauscher, 2012). The orbital velocity of the planet has been measured precisely from water (de Kok et al., 2013) and carbon monoxide (Birkby et al., 2013) absorption lines in the dayside spectrum of the planet, so the orbital velocity in the model is fixed at the measured value of $154 \pm 4 \text{ km s}^{-1}$. During transit the line-of-sight component of this orbital velocity is relatively small, varying from -16 to 16 km s^{-1} , because the planet is moving almost perpendicular to our line of sight.

The stellar line profile observed with HARPS is integrated over the entire disk of the star, and the natural line profile is broadened by macroturbulence and the rotational velocity of the star. The light blocked by the planet and its atmosphere as it transits the star will not be broadened by rotation, and this could potentially affect the transmission signal. To test for the importance of this effect, I attempted to recover the natural line profile of the sodium doublet through stellar atmosphere modeling. The Python package iSpec (Blanco-Cuaresma et al., 2014) was used to find a best fitting model for the stellar sodium absorption line. As the broadening parameters can be degenerate, I fixed the rotation velocity at the value of 3.1 km s^{-1} (Collier Cameron et al., 2010). The model was then fit for effective temperature, microturbulence, macroturbulence, surface gravity and limb darkening parameter. The best fitting model parameters, listed in table 5.1, were then used to synthesise a stellar profile that does not have the 3.1 km s^{-1} rotational broadening applied. In practice, I found that my final velocity measurements are not sensitive to the difference between this synthetic line profile and the integrated stellar line profile.

Yan et al. (2015a) find that the Fraunhofer lines of the Sun display a centre-to-limb variation which is not currently accounted for in my model, and Czesla et al. (2015) detect a similar variation in spectra of HD 189733. Since the effect is primarily to alter the strength of the stellar lines, and the line variations will be symmetric on either side of the transit, I believe it is unlikely that this effect could lead to the asymmetric velocities I report in Section 5.4. However, further modelling of this effect will be the subject of future work.

At each epoch the position of the planet on the disk of the star is calculated using the system parameters. The unbroadened stellar spectrum is Doppler shifted to the velocity of each of the grid points occulted by the planet. Planetary sodium absorption is calculated individually for each grid point after correcting for the orbital velocity and the modeled velocity of the planet atmosphere. The simulated absorption spectrum of the planet is normalised by a simulated out-of-transit spectrum to model the relative transmission. My model accounts implicitly for the Rossiter-McLaughlin effect.

I attempted to spatially resolve the atmospheric velocities by applying my model simultaneously to all twenty in-transit spectra. I use a bootstrap analysis to assess the significance of the velocity measurements. By resampling with replacement from the set of twenty in-transit spectra many hundreds of times the analysis is robust to sources of systematic noise and individual outlier spectra.

5.4 Results

My average transmission spectrum and best fitting model is plotted in Figure 5.2, together with the stellar line profile and example planetary absorption line profiles at different transit phases showing the contributions of the opposite sides of the planet. The signal to noise ratio is too low for the absorption to be individual in the data from individual exposures. The average transmission spectrum displays an asymmetric double peaked structure which is an artefact of the Rossiter-McLaughlin effect (accounted for in my model and discussed in Section 5.5).

My bootstrap distributions of atmospheric velocities are plotted in Figure 5.3, revealing a significant difference in the equatorial velocities on opposite sides of the planet. At the leading limb, I measure a redshift of $2.3_{-1.5}^{+1.3}$ km s⁻¹ and at the trailing limb I find a blue shift of $5.3_{-1.4}^{+1.0}$ km s⁻¹. The two velocities differ by $7.6_{-2.6}^{+2.0}$ km s⁻¹ and the bootstrap distributions differ with more than 99.7% confidence. For the first time I have spatially resolved the atmospheric dynamics of an exoplanet.

Assuming the rotation period of HD 189733b is locked to its 2.22 day orbit




Figure 5.2: Planetary absorption line profiles. (a), the average sodium transmission spectrum of HD 189733b calculated in the frame of the planet (binned by a factor of 10 for clarity) with the average of my time-dependent model overlaid. The asymmetric double-peaked line profile is caused by poor cancellation of the Rossiter-McLaughlin effect in the frame of the planet (accounted for implicitly in my model). (b), the stellar sodium line profile in the frame of the star. (c-e), my modeled planetary absorption line profiles in the stellar frame at ingress, mid-transit and egress respectively (with the Rossiter-McLaughlin effect removed for clarity). The contributions from the trailing (blue) and leading (red) hemispheres can be compared.




Figure 5.3: The posterior distributions of atmospheric velocities from my bootstrap analysis. On the leading limb of the planet (left) a red shift of $2.3^{+1.3}_{-1.5}$ km s $^{-1}$ is found. The trailing limb is blue shifted by $5.3^{+1.0}_{-1.4}$ km s $^{-1}$ (right). The average velocity (middle) is found to be blue shifted with a velocity of $1.9^{+0.7}_{-0.6}$ km s $^{-1}$. The strength and direction of the velocity offsets are consistent with a combination of tidally locked rotation and an eastward equatorial jet that is seen crossing from the dayside to the night side of the planet on the trailing limb.




Figure 5.4: Model calculations of transmission profile around the sodium doublet for a planet without an atmosphere to highlight the Rossiter-McLaughlin effect. *Top Panel:* The transit spectra are aligned in the reference frame of the star. Here the Rossiter-McLaughlin effect is symmetric, and cancels except for a small residual caused by uneven phase coverage. *Bottom Panel:* In this case the spectra are aligned in the planetary reference frame. This breaks the symmetry of the Rossiter-McLaughlin effect and leads to a spurious signal.

symmetric red and blue shifts of 2.9 km s^{-1} would be expected on the opposite equators. Subtracting these expected rotational velocities from my measured equatorial velocities I find the velocity of the leading limb is consistent with tidally locked rotation (super-rotation of $-0.6^{+1.3}_{-1.5} \text{ km s}^{-1}$), but that the trailing limb has an eastward excess velocity of $2.4^{+1.0}_{-1.4} \text{ km s}^{-1}$ (68% confidence errors). The bootstrap posterior distributions are not Gaussian, and I find that this super-rotation at the trailing limb is significant with 95% confidence.

5.5 Discussion and conclusions

My results show that the atmosphere of HD 189733b has a strong eastward motion that is the same direction and has a similar amplitude to that expected from tidally locked planetary rotation. For a close-in planet such as HD 189733b it is almost inevitable that tidal forces will synchronise the rotation to the orbital period in just a few million years (Guillot et al., 1996).

The super-rotating material seen absorbing at the trailing limb of HD 189733b is separated by a longitude of only 60° from the atmospheric hotspot observed to be offset eastward from the substellar point with Spitzer (Knutson et al., 2007b, 2009). It is likely, therefore, that the measured velocity excess is a direct detection of the eastward equatorial jet that is predicted in atmospheric circulation models and has been invoked to explain the offset of the atmospheric hotspot in HD 189733b. The observed jet velocity and the lower velocity on the leading limb of the planet are remarkably consistent with predictions from state-of-the-art circulation models for HD 189733b (Showman et al., 2013).

My detection of this excess velocity is in line with the previous detections of net blue shift in the average transmission spectra of HD 209458b (Snellen et al., 2010) and HD 189733b (Wyttenbach et al., 2015). However, by spatially resolving the atmospheric dynamics I have been able to locate the excess blue shift to the trailing limb of the planet. For comparison with the unresolved studies, my measured limb-averaged velocity offset is $1.9^{+0.7}_{-0.6} \text{ km s}^{-1}$ (Figure 5.3), which is consistent with the average velocity offset of carbon monoxide absorption measured for HD 209458b (Snellen et al., 2010), but significantly smaller than the value of $8 \pm 2 \text{ km s}^{-1}$ found for sodium in HD 189733b (Wyttenbach et al., 2015). I believe this discrepancy is due to the Rossiter-McLaughlin effect, which is accounted for implicitly in my time-resolved model.

The Rossiter-McLaughlin effect is usually assumed to be symmetric and to cancel out in average transmission spectra. This is true when transmission spectra

are constructed in the frame of the star, however, in the frame of the planet the symmetry is broken and the Rossiter-McLaughlin effect causes a spurious signal. This is illustrated in Figure 5.4, where the effect on transmission spectra when analysed in both the reference frame of the star and of the planet is shown. In the stellar frame the Rossiter-McLaughlin effect cancels except for residuals caused by slightly uneven phase coverage in my data. In the planetary reference frame the effect no longer cancels and there is a strong artefact, even with ideal phase coverage. This is accounted for implicitly in my model and is apparent in the asymmetric double-peaked line profile of my best-fitting model in Figure 5.2. Fitting this asymmetric line profile with a symmetric model recovers the larger (spurious) velocity offset found by Wyttenbach et al. (2015), who analyse their spectra in the reference frame of the planet. Note that even in the stellar frame the Rossiter-McLaughlin effect only cancels if the phase coverage is symmetric, otherwise leading to large systematics that likely account for the very high blue shifted velocity of 37 km s^{-1} reported by Redfield et al. (2008) and Jensen et al. (2011) for sodium absorption in HD 189733b.

Brogi et al. (2016) have analyzed high resolution infra-red spectroscopy of CO and H₂O absorption in HD 189733b and find a global blue shift of $-1.7^{+1.1}_{-1.2} \text{ km s}^{-1}$ which is consistent with my measurement.

My resolved velocities show that absorption by the planetary atmosphere is intrinsically velocity broadened, suggesting that previous attempts to model the line profiles in terms of atmospheric temperature and pressure are likely to have over-estimated the role of pressure broadening (Huitson et al., 2012; Wyttenbach et al., 2015; Heng et al., 2015). Similarly, the extent of the planetary atmosphere inferred from the depth of the average line will tend to have been under-estimated (Redfield et al., 2008; Jensen et al., 2011; Huitson et al., 2012; Wyttenbach et al., 2015).

As can be seen in Figure 5.1, the asymmetry in background illumination extends across the planetary limb, so with sufficiently high signal-to-noise objects it should be possible to extend my model to recover wind maps at higher resolution.

Chapter 6

Probing the atmosphere of WASP-52b

Philosophy will clip an Angels wings,
Conquer all mysteries by rule and line,
Empty the haunted air, and gnomed mine
Unweave a rainbow, as it erewhile made
The tender-persond Lamia melt into a shade.

John Keats – Lamia

Transmission spectroscopy currently presents the best opportunity to catalogue the diversity of exoplanet atmospheres. I have measured a precise transmission spectra for WASP-52b, a highly inflated hot jupiter with an equilibrium temperature of 1300 K. Two transits of the planet were observed spectroscopically at low resolution with ACAM on the WHT, covering a wide range of 4000–8750 Å.

I use a Gaussian Process approach to model the trends and systematics in the data. Most systematics are shared between spectral bins, so a common noise model is first calculated from the white light transit curves and removed from each individual lightcurve. This removes the correlated noise within the spectrum, and provides a highly precise relative transmission spectrum with errors on the order of a pressure scale height.

I attempt to fit a variety of different representative model atmospheres to the transmission spectrum, but it was not possible to find a satisfactory fit to the entire spectral range. For the majority of the spectrum (4000–7750 Å) a cloud layer at 0.1 mbar provides an acceptable fit to the data, but is inconsistent with a slightly deeper transit at wavelengths > 7750 Å. This additional opacity may be the result of an additional unknown absorber.

6.1 Introduction

Exoplanet science is rapidly developing into a mature discipline where detailed characterisation of individual planets is beginning to build a picture of the population of planets as a whole. The characterisation of exoplanet atmospheres, primarily through transmission and emission spectroscopy, has become a key topic. Hot jupiter atmospheres have been shown to display a wide variety (Sing et al., 2015), but broadly can be categorised into atmospheres that are “clear”, displaying large Sodium and sodium features (e.g. HD 209458b), “cloudy”, which display muted, or no features at all at optical wavelengths, and “hazy” atmospheres, such as HD 189733b where the dominant feature is a Rayleigh scattering slope caused by μm scale scattering particles high in the atmosphere of the planet. While many high profile discoveries have been space based (e.g. Vidal-Madjar et al., 2003; Sing et al., 2008a; Kreidberg et al., 2014), increasingly, large ground based telescopes have also been found to be suitable to the task for the detection and characterization of both narrow (e.g. Redfield et al., 2008; Louden and Wheatley, 2015) and broadband features (e.g. Bento et al., 2014; Gibson et al., 2013; Gloria et al., 2015; Borsa et al., 2016). While susceptible to their own systematic issues, progress has been made in identifying and improving the major issues, such as differential slit loss (Sing et al., 2012) and upgrading the Linear Atmospheric Dispersion Corrector (LADC) at FORS2 (Moehler et al., 2010; Sedaghati et al., 2015).

Many exoplanets have been shown to present a flat, cloud dominated spectrum (e.g. Lendl et al., 2016; Gibson et al., 2013), data from HST shows that broad wavelength coverage is required to characterize a hot Jupiter, as most appear flat over ranges of several hundred angstrom without the aid of high resolution measurements (Sing et al., 2015). So far, there is no obvious correlation between planetary parameters, e.g.; equilibrium temperature, surface gravity, and host spectral type, though there are broad predictions about what kinds of condensates are available at different temperature ranges and their impact on observational properties (e.g Wakeford and Sing, 2015; Parmentier et al., 2016). The transmission spectrum of a planet is an emergent property of many branches of atmospheric physics, that likely depends quite sensitively on the specifics of heat transport and vertical mixing in each individual planet. It will probably be necessary to analyse many more systems before the underlying trends begin to emerge.

Ground based observations have the advantage of being relatively cheap to perform in large number, compared to space based observations, and also fill an important role in validating the results from *HST* and *Spitzer*, where exoplanet

observations are known to suffer from the idiosyncratic systematic errors of these instruments (e.g. Evans et al., 2015; Wakeford et al., 2016). The way these systematics are treated can have a large impact on the resulting transmission spectra, and many early results are controversial (e.g., the detection of a thermal inversion on HD 209458b Knutson et al., 2008; Line and Yung, 2013). Independent detections with a wider range of instruments would remove some of the uncertainty stemming from these systematics concerns.

Wavelength dependent trends in the data injected by, for example, stellar activity could be mistaken for real atmospheric features on the planet. The repeatability of a purported signal over multiple nights is a powerful diagnostic, giving immense value to have multiple nights on target.

WASP-52b is a Saturn mass planet with an inflated radius that orbits a K type star on a 1.75 day orbit (Hebrard et al., 2012). WASP-52b is an excellent target for transmission spectroscopy due to its large scale height, relatively deep transit, and the presence of a nearby bright comparison star. The low density of WASP-52b gives it a scale height of 700 km, which is over three times greater than HD 189733b, despite having a similar equilibrium temperature of 1300 K. Although it orbits a comparatively faint star ($V_{mag} = 12$), this high scale height means that the Signal-to-Noise achievable on atmospheric features is comparable, assuming that white noise is the dominant noise component (see 1.11). The similarity of the planets equilibrium temperature and parent star spectral type to HD 189733b make WASP-52b an interesting test of the variations in exoplanet atmospheres.

I present two high precision spectral transits of WASP-52b, this Chapter is organised as follows: the reduction of the data is described in Section 6.3.1, and a short description of the Common mode Gaussian process detrending analysis is presented in Section 6.3.3. The system parameters are refined in Section 6.4.2 and a low resolution spectrum is presented with a discussion of its features in Section 6.5.

6.2 Observations

WASP 52 was observed for 2 full nights on August 22 and 29th 2014 using the WHT on La Palma (The instrument is discussed in detail in Section 2.4). ACAM was used in single slit spectroscopy mode with a custom made extra wide slit of $27''$. This was to minimise slit losses, which have been found to be a limiting factor in past observations (Sing et al., 2012), though a slit is still of course necessary to avoid contamination from other stars in the field of view and to lower sky background




Figure 6.1: Spectral trace for target and comparison star. Top: The spectrum image after background subtraction. For display purposes, the image has been normalised by dividing each pixel column by its standard deviation, such that the colour scale represents the significance above the background. This is only for visual clarity, and is not a part of the actual data reduction. The extraction region for each trace is represented by red lines. The central line indicates the distance between the two traces, as a function of pixel position. Bottom: Normalised extracted spectra for target (blue) and comparison (green). The ratio between in the two is plotted in red

contribution. An exposure time of 100 seconds was used to keep the peak flux well below saturation at all times. The windowed observing mode reduced the readout time to 11 seconds, giving a duty cycle of 90%. In order to keep the stars centered in the slit, the data were reduced during the observations and the resulting spectra were cross-correlated to produce guiding offsets in X and Y which were entered manually at regular intervals throughout the observations. A small defocus was applied to the telescope ($\sim 2''$) to spread the light over a larger number of pixels to reduce the contribution of flat fielding errors, and to create a more stable PSF that is less effected by variable seeing.

6.3 Method

6.3.1 Extraction

The reduction was carried out with custom scripts written in Python. Standard bias corrections were made, but it was not possible to flat field the data satisfactorily with the sky flats taken due to differential vignetting between the imaging and spectral modes of ACAM. Spectral dome flats did not contain enough blue light to be useful across the whole spectral range. Due to the large number of pixels used during the analysis the flat fielding errors are not expected to be an important error source.

The combination of vignetting towards the edge of the chip and the broad wings of the brighter comparison star made it difficult to judge where to place sky background traces. However, the background was found to vary smoothly and predictably over the rest of the chip, so I decided to fit this function directly using a third order iterative sigma clipped polynomial and subtract it from the image on a column by column basis (here columns are perpendicular to the spectral axis). This approach has the advantage of leveraging information from a larger number of pixels, which increases the accuracy of the background estimate.

I fit the spectral trace for the target and reference star with an iterative cubic spline. The adjustments to the guiding kept the brightest region of the spectrum dispersed over the same set of pixels, but the spectral trace was found to rotate slowly throughout the night due to differential refraction, meaning the extreme blue and red ends of the spectrum traverse a larger number of pixels throughout the observation and are more likely to suffer an increased noise budget due to the lack of flat-fielding.

An example spectral extraction is shown in Figure 6.1. The colour scale has been selected to emphasise the target and reference star trails. It is clear that the resulting spectra and their ratios are well behaved with high S/N between pixels ~ 600 – 2200 , which after wavelength calibration, are found to correspond to wavelengths ~ 4000 – 9000 Å.

Persistent artifacts are visible in the images, which I suspect is caused by fringing. These become visible beyond ~ 8750 Å, and make it more difficult to trust the spectra extracted in this region. Below 4000 Å a combination of low intrinsic flux, detector response, and instrument vignetting drop the count rate below the background, making a total usable range of 4000 – 8750 Å.

An accurate and stable wavelength calibration is important, errors in wavelength calibration can cause systematic bluewards and redwards slopes that can be mistaken for spectral features in the atmosphere of the planet. Due to the low reso-

lution of the measurements, arc calibration frames taken with the same slit produce very few useful features, particularly at the blue ends of the spectrum where the (lamp type) available at the WHT produces very few lines. Using a smaller slit and removing the offset in slit positions is possible, but I found that the wavelength solution drifted noticeably over the 8 hour observation period. I elected instead to self-calibrate using telluric and strong stellar features. At the resolution of these observations, the velocity difference between the stellar and telluric features is unimportant. The line positions are measured by fitting a Gaussian profile, which well describes the instrumentally broadened profiles. The central pixel co-ordinates of these features are used to generate a wavelength solution using a 3rd order polynomial fit. Each spectrum is then resampled onto a common wavelength grid using (pysynphot¹), which conserves flux. This produces a stable and consistent wavelength solution for the entire set of observations.

Obvious cosmic ray hits are removed from the individual spectra at this stage. A sliding Gaussian weighted average with the FWHM of the instrument is run along the spectra, and points that are 5σ or greater discrepant with this value are replaced by it. Typically only a small number of points per spectra are effected.

To generate a light curve I take the sum of the flux in the desired wavelength range at each time point for both target and comparison, and then take the ratio. To reduce computing time in the Gaussian Process analysis I perform a basic fit to remove extreme outliers. I fit a simple transit model (Mandel and Agol, 2002) with a polynomial baseline fit, and remove any points greater than 5 sigma from the model. Most lightcurves did not require more than one point to be removed. The lightcurve is then normalised such that the median of the out-of-transit data is 1.

6.3.2 Fitting procedure

In order to generate the final spectrum for each night I follow several steps, each of which will be expanded on in detail. First, the transit and systematics model (described in 6.3.3) is fit to the whole of the usable spectral range (4000–8750 Å). I then divide the spectrum into 250 Å spectral bins and fit the transit model to each separately, after removing the best-fit systematics model for that night, and fixing the system parameters to the values found for the white light curve. This removes the correlated noise between spectral bins, which significantly lowers the errorbars.

A Markov Chain Monte Carlo approach is used to maximise model parameter likelihood and estimate their errors. To perform the MCMC integration, I use the

¹distributed as part of Space Telescope Science Institute stsci.python
<http://www.stsci.edu/institute/software.hardware/pyraf/stsci.python/>




Figure 6.2: top: A periodogram for the most recent season of WASP data on WASP-52. The stellar rotation period is clearly detected at 17.79 days. The one, two and three σ FAP limits are indicated. Bottom: Photometric data from WASP folded on the detected period of 17.79 days.

emcee code of (Foreman-Mackey et al., 2013)², which implements the affine-invariant ensemble sampler described in (Goodman and Weare, 2010).

To fit the transit light curve, I use the analytic models described in Mandel and Agol (2002). The eccentricity is fixed to 0, and at each step draw the period from a distribution with mean and standard deviation values from Hebrard et al. (2012). The inclination, system scale (a/R_*), radius ratio (R_p/R_*) and limb darkening co-efficients are free parameters, as well as the parameters describing the Gaussian Process systematic model. LDTK (Limb Darkening Toolkit) (Parviainen and Aigrain, 2015) was used to calculate the priors on the limb darkening parameters for each wavelength range.

Each fit has 5,000,000 steps, with the first 1,000,000 discarded as burn in to ensure that the model is converged.

6.3.3 Gaussian process and white noise model

Gaussian processes are a generalized class of functions that can be used to model correlated noise in time series data (see Section 2.9 for a more detailed introduction). They have a wide range of applications, but have increasingly become popular for accurate parameter estimation for exoplanet transits (e.g. Gibson et al., 2013; Parviainen et al., 2016; Evans et al., 2015). The covariance between data points is modelled with “kernel”, which has a simplified functional form. Including a Gaussian Process (GP) noise model allows one to effectively marginalise over the sets of functions that can represent the systematic noise in the data. This approach has the advantage of implicitly including the increased level of uncertainty due to correlated noise in the parameter estimates, a method similar to, but more mathematically rigorous than, the common practice of inflating errorbars to achieve a reduced χ^2 of 1.

GP’s are useful as they allow one to fit with only a couple of parameters the nuisance trends that do not hold any scientific interest, or are difficult to characterise, such as simultaneous changes in atmospheric transparency, seeing or instrument stability. They are also intrinsically Bayesian, effectively sampling all possible systematic trends that explain the data and folding this uncertainty into the model parameter estimates. In my case, I use the Gaussian Process to model everything in the data that is not explained by the transit model.

The covariance matrix K is modelled as

$$K_{ij} = \sigma_i^2 \delta_{ij} + k(t_i, t_j) \quad (6.1)$$

²available at <https://github.com/dfm/emcee>

where σ_i is the stationary white noise on each data point, δ_{ij} is the Kronecker delta function, and $k(t_i, t_j)$ is the covariance matrix that parameterises the relationships between the data points at time t_i and t_j . In this work the Matern 3/2 kernel is used, which is similar to the commonly used squared exponential function, but includes higher short term variance.

$$k(t_i, t_j) = \exp\left(\frac{t_i - t_j}{\tau}\right) \quad (6.2)$$

Fitting with a Gaussian process model requires a Maximum Likelihood estimation, which is solved with a MCMC. The log likelihood function for Gaussian process regression is

$$\log p(\mathbf{r}|X, \theta) = -\frac{1}{2}\mathbf{r}^\top K^{-1}\mathbf{r} - \frac{1}{2}\log |K| - \frac{1}{2}n \log 2\pi \quad (6.3)$$

Where n is the number of datapoints and r is the residuals of the physical model to the data with the parameter vector X , in this case representing the transit model. K is the previously described covariance matrix and θ is the set of parameters that describe the form of the covariance matrix, the so-called “hyperparameters”, these represent entire *classes* of functions with desired properties, and not an individual model as such.

The likelihood function can be understood in a straightforward way by inspecting the 3 terms separately. The first term is effectively a goodness of fit, and is analogous to a χ^2 . It represents how well the signal *and* the residuals can be fit by the model and the covariance matrix. The second term acts as a complexity penalty which prefers simpler covariance kernels and prevents the GP from over-fitting the data. This term is effectively an “Occams razor”, and is what prevents the Gaussian Process from competing with the model to fit signal in the data.

If the first term was taken in isolation it would be maximised by a GP with high white noise, i.e, if you assume there is enough high-frequency noise in the data you can find a GP that will pass through every data point. The second term would be maximised by a kernel with 100% correlation between all points, which is just a flat line with the mean value of the data. The interplay of these two terms is what gives the Gaussian Process its power.

The third term is just a normalisation factor, and plays no special role in the fit.

George³ (Ambikasaran et al., 2016) is used to efficiently calculate the kernels and Likelihood function at each Markov chain step.

³available at <https://github.com/dfm/george>

Kernels such as the squared exponential and the Matern do well at fitting noise sources with characteristic length scales (see Section 2.9), if there are multiple systematic sources with very different length scales, it may be necessary to use a more complex kernel such as a rational quartic, or combinations of simpler kernels. In this case I found that in all wavelength channels a single component was sufficient, and had long characteristic timescales, primarily tracking second order airmass effects and transparency.

However for some channels, particularly the bluer ones, I found that the white noise was systematically underestimated. Adding an extra parameter to increase the calculated errorbars by a fixed amount produced improved fits, but in practice I chose to take a similar approach to Gibson et al. (2013) and model the white noise as a smoothly time varying function in all fits. This makes sense, since an enhancement to the white noise, most likely due to changing atmospheric conditions, is unlikely to be constant. The white noise term in equation 6.1 is now modelled as:

$$\sigma(t_i) = \sigma_a \exp\left(-\frac{t_i}{\sigma_b}\right) + \sigma_c \quad (6.4)$$

where σ_a , σ_b and σ_c are now additional terms in the model fit.

I find that for most wavelength regions the white noise model makes no difference to the parameter posteriors, but in the bluest channels, where scintillation is expected to have the greatest effect, I find that the more accurate characterisation of the noise properties slightly narrows the credible intervals of the transit depth.

6.3.4 Application of the model

Due to the presence of the spot-crossing event on the first night, I chose to initially fit the two nights separately to ensure that the Gaussian process systematic model was capable of fitting out the event without introducing bias into the other parameter values. Comparison of the parameter values satisfied me that, apart from a small difference in absolute depth which is expected from changes in activity levels, the parameters are consistent on both nights, see table 6.2. The final parameter values are calculated by fitting the two nights simultaneously.

Ultimately, what I am interested in is the *relative* depth of the wavelength channels, or their conditional probability distribution with respect to each other in order to build a transmission spectrum. Systematic noise sources, such as guiding instability or changes in atmospheric transparency will be the same, or subject to a small scaling factors between bands. It is therefore incorrect to treat a resulting transmission spectrum as a set of independent data-points, and their calculated

errorbars will be higher than the intrinsic scatter between them. This could also lead to unrealistic model fits, since most goodness of fit statistics, such as χ^2 , assume that the datapoints are independent.

I use the white light curves for each night to remove the common noise trends

6.3.5 Spot crossing event

Due to the changing activity levels of the star, it would not a priori be expected that the depth, and hence inferred radius should be consistent between the two nights. Indeed, I find that on the night with the spot crossing event the implied transit depth is slightly deeper than the other. The difference in absolute depths between the two nights is not formally significant, at 0.0102 ± 0.0067 . Taking these results at face value, they would imply that the stellar surface was on average 0.4% dimmer on that night. Whilst I unfortunately did not have access to contemporaneous monitoring of WASP-52 during the observation period, Hebrard et al. (2012) found that during their second observing season WASP-52 displayed a 16.4 ± 0.04 day period with an amplitude of 0.89%, though there was no significant periodicity found in the first season. The separation between the two transits I observed is 7 days, or roughly half of the rotational period of the star.

Since Hebrard et al. (2012) carried out their analysis, an additional season of WASP data on WASP-52 has become available. I carry out a harmonic analysis to assess how the activity levels vary between seasons. I used a generalized weighted Lomb Scargle periodogram (Zechmeister and Kürster, 2009) to search for significant periods in the data, which were taken between reduced HJD 2455400 - 2455800. I generate False Alarm Probability power levels by using the night shuffling technique presented in Collier Cameron et al. (2009) and Maxted et al. (2011). Randomly sampling the nights instead of individual datapoints preserves correlated noise in the data, whilst removing any periodicities greater than 1 day, this gives a more realistic and conservative estimate of the FAP. I find a highly significant peak at 17.79 days in the most recent season. To estimate the errors on the period and the amplitude I ran an MCMC of 10,000 steps, with a burn in of 1000 steps. I find that the amplitude is $0.0015^{+0.04}_{-0.04}$ and the period is $17.79^{+0.04}_{-0.04}$. The amplitude is significantly higher than during the first and second seasons observed with WASP, which may be indicative of activity cycles. The period is 1.4 days longer than the period found in the second season, which is likely caused by the active regions occurring at different latitudes and differential rotation.

Despite the lack of contemporaneous coverage, the difference in absolute

Table 6.1: Derived system parameters for WASP-52b.

Parameter	Symbol	Night 1	Night 2	Combined	Unit
Epoch (JD _{UTC} - 2450000)	T_0	6892.53861 ± 0.00029	6899.53749 ± 0.00016	5793.676184 ± 0.000033	days
Epoch (BJD _{TBD} - 2450000)	T_0	6892.54464 ± 0.00029	6899.54374 ± 0.00016	5793.682045 ± 0.000033	days
Orbital period	P	-	-	$1.74978089 \pm 0.000000013$	days
Inclination	i	$85.33^{+0.22}_{-0.23}$	$85.28^{+0.18}_{-0.17}$	$85.32^{+0.14}_{-0.14}$	degrees
System Scale	a/R_*	$7.23^{+0.12}_{-0.13}$	$7.176^{+0.089}_{-0.087}$	$7.200^{+0.076}_{-0.071}$	-
Radius Ratio	R_P/R_*	$0.1741^{+0.0063}_{-0.0054}$	$0.1639^{+0.0030}_{-0.0030}$	-	-

depth between the two nights is therefore consistent with the historical activity levels of the star. Note that the depth on the uncontaminated night is consistent with both the original published value, and the faculae model in Kirk et al. (2016), but is inconsistent with their spot model. All other parameter values agree to within one sigma, and there is no sign of correlation between the GP hyper-parameters and transit parameters (except depth) in the MCMC corner plot. The significant difference between the spot model in Kirk et al. (2016) and both the discovery paper and this work suggest that their faculae model is correct.

Since the spot crossing event appears to be grey to within the level of precision, the approach of removing the common Gaussian process noise model is preferred over attempting to model the spot directly. This would add additional model parameters and degeneracies, and would not improve the level of precision attained in the recovered spectrum. Inspection of the GP model for night 1 (Figure 6.3) reveals at least 2 distinct components which indicate that the idealisation of a single “spot crossing” is unsuitable to this case, and it is more likely to be the crossing of an active region, with multiple spots simultaneously transited. Given the success of the GP model I feel it would be unnecessary to attempt to model this explicitly.

This approach is not strictly correct, as the implicit assumption of the Gaussian process approach is that the noise properties are the same for the whole observation window. In principle it would be possible to specify a compound GP that had an additional kernel active only during the transiting portion of the dataset in order to simulate the additional correlated noise component from inhomogeneities of the stellar surface. However, since the event is not visible in the individual spectral bins after common mode removal (see figures 6.9 and 6.11), I chose not to add additional complexity to the model.

6.4 Results

6.4.1 Transmission spectra

For each 250 Å bin, the median value of the transit depth and the 68% credible region are calculated from the MCMC posteriors. I present the resulting transmission spectra for the two nights in table 6.3 and Figure 6.4.

Having measurements on two night proves valuable, as I was able to use them as an independent test of the assigned errorbars. Assuming that the same transmission spectrum is present on both nights, it is expected that the error-normalised deviation of the points from their average will behave like a set of independent Gaussian measurements, with standard deviation proportional to $1/\sqrt{2}$. Performing this

test on the two nights, I find that they do indeed behave as expected for a set of independent measurements of the same distribution, once an absolute depth difference between the two nights has been removed. This indicates that the GP/MCMC method has not underestimated the relative uncertainty in each wavelength bin, and that little common-mode uncertainty remains.

Since the spectra of the two nights do appear to be describing the same underlying spectrum, and the errorbars are reasonable, I use the variance weighted average of the two nights as my final result, shown in Figure 6.5

It is worth emphasising that since the common-noise has been removed, these error bars are only correct relative to each other, and are not a valid measurement of uncertainty of the absolute depth.

As an additional test for sodium absorption, I repeated the process with narrower 50 Å bins around the 5950 feature, but found no significant difference to the broader bands. Narrower bands had significantly larger noise, so could not provide further constraints. Whilst the narrow line core of sodium may be present in the spectrum of WASP-52b it is not possible to detect it at the resolution of the instrument. Higher resolution observations are capable of detecting the line core of sodium even above a cloud/haze deck (e.g. Redfield et al., 2008; Louden and Wheatley, 2015).

6.4.2 Updated system parameters

I fit the white light curves simultaneously for both nights, using the MCMC GP model, the results are shown in Figure 6.3. I included the transit epoch data from the discovery paper and Kirk et al. (2016) in order to improve the global system parameter estimates. Since the mean stellar surface brightness, and the noise properties are not expected to be the same on both nights, they are allowed to fit separately, but the other system parameters (a/R_* , limb darkening coefficients and inclination) are fit simultaneously. My final system parameters are listed in table 6.2, and agree to within 1σ with the values in the discovery paper.

6.5 Discussion

6.5.1 Transit timing

Baluev et al. (2015) search for periodic Transit Timing Variations in 10 exoplanet systems, including WASP-52b and find no evidence of periodic variations. I leverage the high precision of the white noise lightcurves and the longer time baseline between

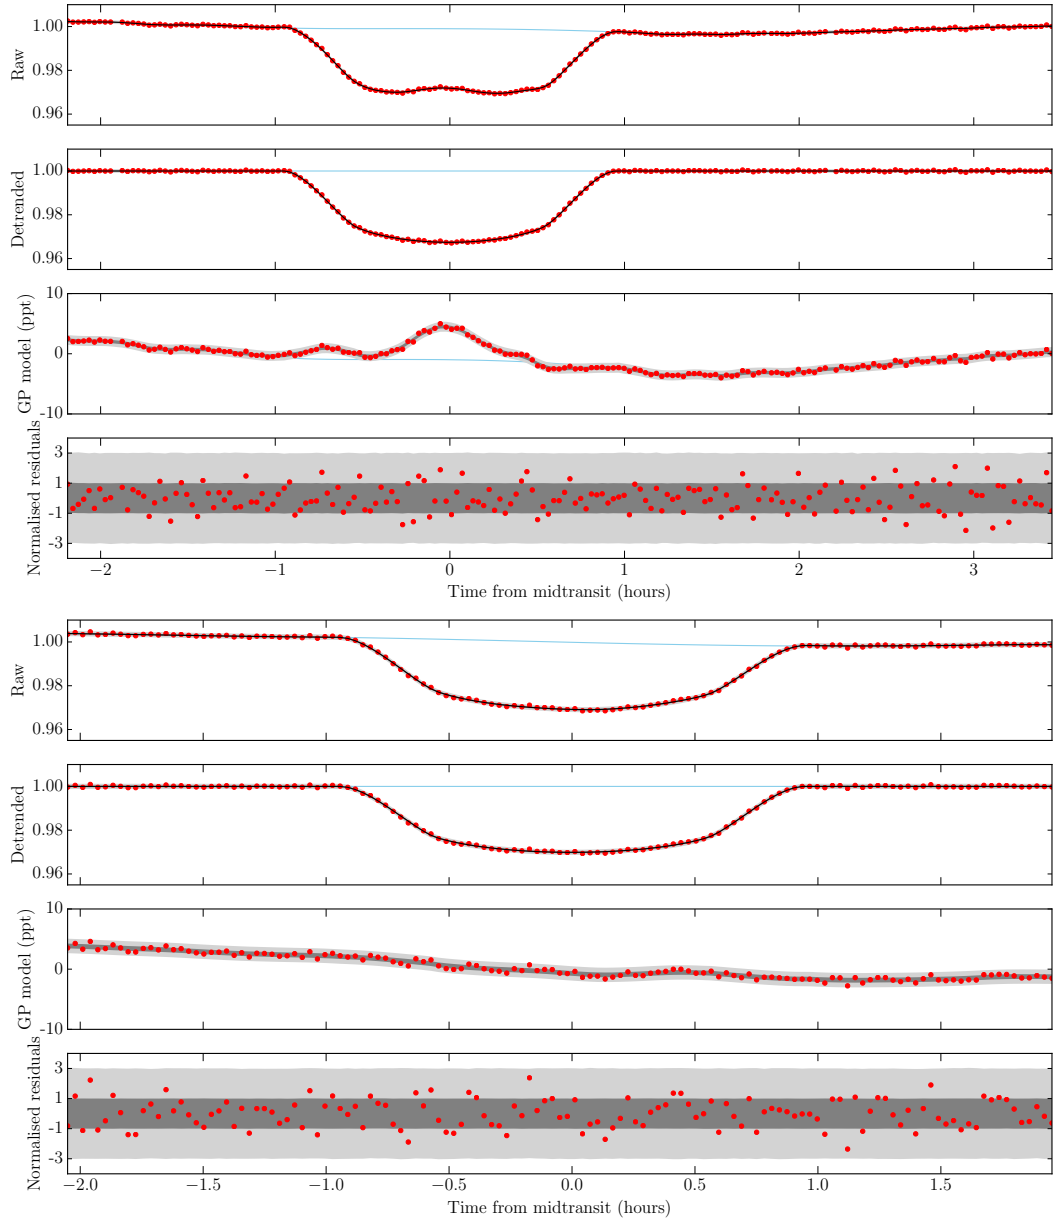


Figure 6.3: White light model fits for night 1 (top) and night 2 (bottom). In each, the first panel shows the Gaussian process model fit to the raw data, the second panel shows the model fit to the data with the systematics model removed, the third panel shows the systematic model, and the fourth panel shows the normalised residuals. The dark and light shaded grey regions are the 1 and 3σ confidence regions, respectively.




Figure 6.4: Top: Spectrum of the target and comparison stars and their ratio (orange). The 250 Å bins used to generate the spectrum are indicated with dashed lines. Bottom: Transmission spectra from night one (pink) and night two (blue). The absolute offset between the two nights has been removed from the second night. The error bars are the 68% credible intervals from the MCMC posteriors for each spectral bin. Night 2 has been offset slightly in wavelength for clarity.




Figure 6.5: Final combined transmission spectrum for the two nights. Over plotted are a variety of atmosphere classes. Clear, with solar metallicity and no TiO (pink), a HD 189733b style Rayleigh scattering slope that covers the entire optical range, with a temperature inversion (blue). A clear atmosphere with a solar abundance of T and V, assuming 100% is locked up in TiO and VO (orange). A Rayleigh slope with settling (green) representing a hazy atmosphere with larger particles. A grey absorbing cloudbase at 0.1 mbar (yellow).




Figure 6.6: Temperature-Pressure profiles used to calculate the model atmospheres in Figure 6.5. Both were generated using the semi-analytical models of Heng et al. (2012). The green line represents a clear, zero-albedo atmosphere with no significant shortwave scattering, the blue line is the high short-wave scattering model used to generate the hot-haze spectrum.

the ephemeris reported in the discovery paper (Hebrard et al., 2012) in order to test for secular drifts in the ephemeris of WASP-52b. The errorbars on my ephemerides are higher than those reported in Kirk et al. (2016) and Hebrard et al. (2012) due to the lower cadence of the observations and the greater model flexibility allowed by my Gaussian Process analysis.

I convert the transit dates to BJD in order to calculate the ephemeris using an online tool⁴ (Eastman et al., 2010).

Hebrard et al. (2012) report their central transit time in HJD, which can introduce errors at the ~ 1 second level compared to calculating BJD_{TBD} directly, so I add this uncertainty in quadrature to their reported errors.

I calculated the errors using an MCMC with a chain length 100,000 steps and a burn in chain of 1000 steps. The best fitting ephemeris is shown in Figure 6.7.

A linear ephemeris is formally a poor fit, which could be interpreted as evidence of a non-linear ephemeris in this system. However, I believe it is more likely that previous work may have systematically underestimated the size of their errors by not properly accounting for the presence of correlated noise through Gaussian processes or some other technique. Future measurements of the ephemeris of this system may shed more light on the possibility of additional companions.

The large time baseline (4 years) between the discovery paper and this work allows a significant improvement in the uncertainty of the period. The aforementioned concern over underestimated errors is not particularly important, as the error budget is dominated by the gradient over time periods this long.

6.5.2 Model atmospheres

Due to the quality and resolution of the available data, a full atmospheric retrieval is not appropriate, especially considering the degeneracies involved in reproducing the flat region of the spectrum ($\sim 4000\text{--}7750\text{\AA}$). Instead I generate a sample of representative atmosphere classes to compare to the data, and qualitatively discuss their applicability to the dataset.

All of the model atmospheres are generated using the NEMESIS radiative transfer code (Irwin et al., 1997, 2008), and simple parametrized assumptions about the Temperature-Pressure profile using the equations presented in Heng et al. (2012). I assume the zero albedo equilibrium temperature for WASP-52b in the case of uniform redistribution, which is 1300 K. (using $f=0.25$ from Seager (2010)). Using the

⁴<http://astrutils.astronomy.ohio-state.edu/time/utc2bjd.html>

equations in Heng et al. (2012), I calculate a Temperature-Pressure profile appropriate for a clear atmosphere with low levels of short-wave scattering $k_s = 6 \times 10^{-4} \text{ cm}^{-2} \text{ g}^{-1}$ and $\gamma_0 = 0.06$ (Figure 6.6).

Pont et al. (2013) find that for HD 189733b, the entire optical and near infrared transmission spectrum is dominated by a Rayleigh scattering slope. The majority of the ACS data is well fit by a Rayleigh slope model with an equilibrium temperature of $\sim 1300 \text{ K}$ (matching the expected equilibrium temperature of the planet), but bluewards of 6000 \AA with data taken with STIS at a different epoch, they find that the slope is somewhat steeper, having a best fitting temperature of $2100 \pm 500 \text{ K}$ (Sing et al., 2011). The difference in slope between the ACS and STIS data is only significant at the 1.5σ level, with the majority of the error coming from corrections due to spot coverage. The authors speculate that the difference in slope could be interpreted as weak evidence for a thermosphere. A thermal inversion is a natural consequence of an atmosphere with a strong shortwave scattering component, such as a photochemical haze at high altitudes. This is demonstrated in the semi-analytical models of Heng et al. (2012), which are based on a generalisation of the formalism of Guillot (2010). Using these equations, Pont et al. (2013) generate an example Temperature-Pressure profile for HD 189733b with a strong thermal inversion layer with a thermosphere that reaches 2000 K . This would be capable of reproducing the strong Rayleigh feature seen in the ACS data. The analysis of the Sodium doublet by Huitson et al. (2012) and Wyttenbach et al. (2015) suggest that a temperature inversion does exist on HD 189733b, characteristic of a thermosphere, which could explain the slightly steeper gradient at the bluer end of the Rayleigh slope, which naturally probes higher altitudes and thus higher temperatures. I note that the results I find in Chapter 5 suggest that the line profile of hot Jupiters are rotationally broadened. This will likely have an impact on attempting to recover temperature-pressure profiles from line profile fitting.

As WASP-52b has a similar equilibrium temperature and parent star spectral type to HD 189733b, it is plausible for the atmosphere to sustain a similar scattering haze layer. I assume the same parameters as Pont et al. (2013) in their example T-P profile ($k_s = 0.0670$, $\gamma_0 = 10.0$, $\eta = 0.1$) which is equivalent to a geometric albedo of 0.5. For WASP-52b, this results in a thermosphere which reaches 2100 K at 0.1 mbar (see Figure 6.6). The steep Rayleigh slope a haze produces in this hot thermosphere is clearly a very poor fit to the data (χ^2 of 331.07 for 19 DOF) but I include it in order to demonstrate the sensitivity my measurements would have to an atmospheric feature of this type- HD 189733b class atmospheres are relatively easy to detect from the ground therefore, and even a substantially lower S/N spectrum

would be sufficient to detect one.

Of the models tested, the flat, cloud dominated spectrum provided the best fit to the data, with a χ^2 of 40, but as I point out in Section 6.5.3, the error in this fit is dominated by the last 4 points, and may be indicative of an unknown absorber. The cloud layer at 0.1 mbar is consistent with cloud formation theories for HD 189733b, which has the same equilibrium temperature as WASP-52b (Lee and Juncher, 2015). These clouds would likely be composed of a range of chemical species, with magnesium silicates as important contributors.

A clear atmosphere with solar metallicity without TiO and VO absorption is found to be a very poor fit, with χ^2 90.06 for 19 DOF, giving $\chi_R^2 = 4.74$. The 99.7% upper limit for the χ_R^2 distribution with 19 DOF is 2.29. This poor fit is driven by the clear lack of a broadened sodium feature and Rayleigh slope. These features could be masked by high altitude clouds, or by adding additional molecular absorption, such as titanium oxide and vanadium oxide, which can introduce a large number of spectral features that are mostly unresolved at this resolution and serve to de-emphasise the Rayleigh slope and sodium and potassium features by raising the average optical depth of the atmosphere.

I assume solar abundances of T and V (Asplund et al., 2006), in hot Jupiter atmospheres T and V in the gas phase are expected to be practically entirely in the oxide forms from chemical equilibrium, but below 1600K condensation and cold trapping remove significant fractions from the atmosphere (Fortney et al., 2008). The details of TiO and VO chemistry in hot Jupiter atmospheres are not fully understood, and depend sensitively on the T-P profile, vertical mixing, and availability of condensation compounds, so a full treatment is beyond the scope of this work. Instead, I test the effect of several different un-condensed fractions on the spectrum.

A fraction of 1 part per million is indistinguishable from the clear spectrum, but increasing the fraction above this level begins to improve the fit by de-emphasising the Rayleigh slope and sodium feature. I find that an abundance fraction of 0.01 decreases the χ_R^2 to 2.88, which is still poor, but produces an acceptable fit in the 4000–7750 Å range, with a χ_R^2 of 1.43 (The 68% and 95% upper limit for the χ_R^2 distribution with 15 DOF are 1.12 and 1.59, respectively).

The spectrum could also be flattened by high altitude clouds, which seem to be a common feature of exoplanet atmospheres (e.g. Kreidberg et al., 2014). It is not clear what condensed species are involved, but their properties may make it possible to determine this (Wakeford and Sing, 2015). I insert opaque layers to the radiative transfer code at 0.1, 1 and 10 mbar. All provide significant improvement over the clear spectrum. The flat spectrum provided by the 0.1 mbar cloud layer is

Table 6.2: Model atmosphere goodness of fits. The 4000–7750 band has 15 DOF and the 4000–7750 has 19 DOF. No models provide satisfactory fits to the whole range, but both a cloud dominated atmosphere and a clear atmosphere with TiO absorption can provide acceptable fits to the 4000–7750 range.

Model	χ^2 (χ_R^2)	
	4000–7750 Å	4000–8750 Å
Clear	41.01 (2.73)	90.06 (4.74)
Hazy (2100 K)	159.85(10.66)	331.07(17.42)
Hazy (1300 K)	59.9(3.99)	132.79(6.99)
Hazy (with settling)	26.52(1.77)	60.43(3.18)
TiO and VO (1e-6)	40.98 (2.73)	90.0(4.74)
TiO and VO (1e-4)	38.69 (2.58)	85.28(4.49)
TiO and VO (1e-2)	21.45 (1.43)	54.7(2.88)
Cloud (10 mbar)	29.75 (1.98)	65.57(3.45)
Cloud (1 mbar)	20.13 (1.34)	44.07(2.32)
Cloud (0.1 mbar)	19.36 (1.29)	41.18(2.17)

the best fit to the data, with a χ_R^2 of 2.17 in the full range, and 1.29 for 4000–7750 Å. I am unable to distinguish between clouds at higher altitude than 0.1 mbar, as the spectrum is already completely flattened.

Kirk et al. (2016) also find a flat spectrum compatible with a cloudy interpretation in their analysis of WASP-52b with high precision photometry.

At this level of precision, I am formally unable to significantly distinguish between a cloudy atmosphere and an atmosphere with low level TiO and VO absorption. The goodness of fit is best in the cloud dominated case, but the improvement over a clear atmosphere with some TiO contribution, $\Delta BIC = \Delta\chi^2 = 2.1$, is marginal. However, since TiO signatures in exoplanet atmospheres appear to be rare, even in the hotter atmospheres where they are predicted to be much stronger, the cloudy interpretation is preferred. A higher resolution search for the presence of the sodium feature (which is still present in the TiO case, but mostly obscured in the cloud-dominated case) could potentially help break the degeneracy. Detecting the water feature in the infra-red would also be a useful diagnostic, as Sing et al. (2015) shows this is typically obscured in cloudy atmospheres.




Figure 6.7: Transit timing residuals from a linear ephemeris. The 1σ contour for a linear ephemeris drift model is plotted in grey. No significant deviation from a linear ephemeris is found.

6.5.3 An unknown absorber beyond 7750 Å?

There is a statistically significant excess in transit depth beyond 7750 Å which I was unable to replicate with my sample of models (see 6.5).

In the combined spectrum, the best fitting model is cloud dominated flat spectrum, which gives a reduced χ^2 is 19.35 for 15 dof when excluding the last 4 points (7750-8750), but for 32.55 for 19 dof when including them. This is a highly significant difference.

The feature appears to be present on both nights so it is unlikely to be a systematic error, but it may be a common artefact of the reduction process. Several possibilities are explored here.

The feature is not close to the edge of the chip, and by visual inspection it is not on the region of the chip that is effected by fringing (y pixel > 2200). Additionally, the regions of the chip effected by fringing show a dramatically increased high frequency noise, which is distinct to a persistent, systematic increase in radius ratio. The errorbars in the 7750–8750 Å region are comparable to the errors elsewhere in the spectrum.

An optical blocking filter was not used when taking these data, so redwards of 6000 Å the spectrum is weakly contaminated by the blue end of the second order spectrum. However, this effect is low order (see Section 2.4) and is not able to increase the observed transit depth above the blue end of the spectrum.

Incorrect limb darkening parameters can cause errors in radius retrieval (e.g. Csizmadia et al., 2012; Espinoza and Jordan, 2015) and could introduce a correlated bias to a transmission spectrum. The limb darkening values and errors were generated with Limb Darkening Tool Kit. LDTK fits a parametrised limb darkening profile to PHOENIX-calculated specific intensity spectra (Husser et al., 2013), propagating the errors in the provided stellar parameters to the intensity profile. It is useful to incorporate this additional knowledge into the model fit, if for no other reason than it implicitly and self-consistently disallows unphysical limb darkening profiles.

To test that these model profiles were not influencing the shape of the resulting transmission spectrum the fit was repeated without the priors on the limb darkening parameters. To improve convergence, the formalism of Holman et al. (2006) and Burke et al. (2008) was followed, adopting $a1 = 2u1 + u2$ and $a2 = u1 - 2u2$ as the parameters used internally in the MCMC, since these parameters are uncorrelated. To avoid the limb darkening parameters taking unphysical values, the following parameter constraints were imposed; that the highest surface brightness is located at the disk center ($u1 \geq 0$), the specific intensity remains above zero

($u1 + u2 \leq 1$) and limb brightened profiles are not allowed ($u1 + 2u2 > 0$).

I find that without the additional limb darkening constraints the errors on the resulting spectrum are slightly higher, but the shape of the spectrum is preserved. Therefore it is not likely that limb darkening can account for the 4 outlier points.

I can find no obvious systematic reason for this feature in my analysis, this may be a hint of a high altitude absorbing species that is not included in current models.

(Nikolov et al., 2013) find a similar redward feature in their STIS G750L low-resolution transmission spectrum of HAT-P-1b. A similar effect can also be seen on HD 209733b (Knutson et al., 2007b), WASP-19b with FORS2 (Sedaghati et al., 2015) and TrES-3b with OSIRIS (Parviainen et al., 2016), suggesting that this may be a common feature of hot Jupiter atmospheres. It seems unlikely that such a similar systematic trend would occur independently on multiple instruments with multiple different reduction techniques, and for only some planets. If this is indeed the signature of an unknown, high altitude absorber and not a common systematic across multiple instruments, detailed non-equilibrium chemical models with experimental validation such as Venot et al. (2012) will likely be necessary to ascertain what the absorber is likely to be.

The flat cloud model with cloud decks > 1 mbar provide the best fit to the data, and is a good fit below 7750 Å. Reproducing the feature above 7750 Å would require the absorber to be higher than > 1 mbar. A strong, broad feature at such low pressures is likely to be difficult to replicate. It is also possible that the atmosphere may be deficient in sodium, as WASP-31b appears to be (Sing et al., 2014). This would allow the cloud deck to be lower in the atmosphere, where a strong molecular feature is more plausible.

Alternatively, the majority of the opacity may be provided by molecules instead of clouds. TiO and VO are expected to be important absorbers in exoplanet atmospheres due to their ubiquity in cool star and brown-dwarf spectra (Fortney et al., 2008), but have so far proved difficult to detect in exoplanet atmospheres, it has been hypothesised that this is due to cold trapping or rain-out on the night side of the planets (Parmentier et al., 2013), the lack of hemispherical symmetry not being a consideration of note in stars. A recent detection of TiO absorption on WASP-121b (Evans et al., 2016) suggests that TiO may indeed be an important absorber in very hot Jupiters, but in cooler atmospheres the impact will be reduced as it condenses out.

Our clear atmosphere models including low levels of TiO absorption were a significant improvement over a clear atmosphere, but could not replicate the higher

Table 6.3: Transmission spectra for WASP-52b. To produce the combined values the absolute offset between Night 1 and Night 2 was removed and a weighted average of the two nights is taken.

Wavelength (Å)	R_p/R_*		
	Night 1	Night 2	Combined
4000–4250	0.1771 ± 0.0018	0.1647 ± 0.0022	0.1643 ± 0.0014
4250–4500	0.1750 ± 0.0014	0.1672 ± 0.0026	0.1632 ± 0.0013
4500–4750	0.1780 ± 0.0020	0.1660 ± 0.0014	0.1656 ± 0.0012
4750–5000	0.1762 ± 0.0016	0.1646 ± 0.0012	0.1641 ± 0.0010
5000–5250	0.1739 ± 0.0015	0.1641 ± 0.0011	0.1629 ± 0.0009
5250–5500	0.1752 ± 0.0013	0.1618 ± 0.0009	0.1619 ± 0.0008
5500–5750	0.1762 ± 0.0015	0.1653 ± 0.0017	0.1641 ± 0.0011
5750–6000	0.1764 ± 0.0013	0.1645 ± 0.0015	0.1639 ± 0.0010
6000–6250	0.1769 ± 0.0012	0.1637 ± 0.0009	0.1638 ± 0.0007
6250–6500	0.1779 ± 0.0013	0.1633 ± 0.0012	0.1640 ± 0.0009
6500–6750	0.1777 ± 0.0012	0.1625 ± 0.0007	0.1630 ± 0.0006
6750–7000	0.1806 ± 0.0016	0.1635 ± 0.0009	0.1645 ± 0.0007
7000–7250	0.1782 ± 0.0017	0.1644 ± 0.0009	0.1646 ± 0.0008
7250–7500	0.1786 ± 0.0018	0.1652 ± 0.0009	0.1652 ± 0.0008
7500–7750	0.1789 ± 0.0015	0.1649 ± 0.0009	0.1651 ± 0.0008
7750–8000	0.1791 ± 0.0007	0.1659 ± 0.0009	0.1660 ± 0.0006
8000–8250	0.1778 ± 0.0010	0.1669 ± 0.0010	0.1658 ± 0.0007
8250–8500	0.1780 ± 0.0047	0.1667 ± 0.0017	0.1664 ± 0.0016
8500–8750	0.1798 ± 0.0024	0.1681 ± 0.0022	0.1675 ± 0.0016

depth $> 7750 \text{ \AA}$, as the strongest features are centred around 6000 \AA . The line-lists of TiO are known to be incomplete, and improved TiO linelists are under construction by projects such as ExoMol (Tennyson and Yurchenko, 2012), but inspection of cool dwarf stars do not show an unexpectedly strong features in this wavelength region. The combination of TiO and an additional absorber may be the most parsimonious solution, as the unknown absorber is not confined to heights above 1mbar as it is in a cloud dominated model.

My results highlight the need for both low and high resolution observations to break the degeneracies of exoplanet atmosphere retrieval.

6.6 Conclusions

I have performed transmission spectroscopy on WASP-52b. My Gaussian Process and common noise model approach to modelling resulted in a precise transmission spectrum that was consistent across two nights despite the presence of a spot, with errors on the order of a single scale height.

The “spot-crossing” feature on the first night appears to be grey, and is not visible in the spectrally resolved light curves after common-noise term removal. This implies that the temperature contrast of the spot to the rest of the photosphere is high, which is an expected feature for cooler stars such as WASP-52.

I attempted to fit representative model atmospheres to the transmission spectrum, but was unable to find a satisfactory fit to the entire spectral range. For the majority of the spectrum ($4000\text{--}7750 \text{ \AA}$), a cloud layer at 0.1 mbar provides an acceptable fit to the data, but is inconsistent with a slightly deeper transit at wavelengths $> 7750 \text{ \AA}$. I explore several different systematic reasons for this excess depth, but believe it is most likely a real feature of the atmosphere, and may be the result of an additional unknown absorber.

The cloud layer at 0.1 mbar is consistent with cloud formation theories for HD 189733b, which has the same equilibrium temperature as WASP-52b (Lee and Juncher, 2015). These clouds would likely be composed of a range of chemical species, with magnesium silicates as important contributors.

The transmission spectrum is dramatically different to HD 189733b, which has the same equilibrium temperature and is dominated by a Rayleigh slope caused by high altitude haze. The difference between these transmission spectra provides further evidence of the rich and as yet unexplained variety within exoplanet atmospheres.




Figure 6.8: Night 1 individual lightcurves




Figure 6.9: Night 1 residuals




Figure 6.10: night 2 individual lightcurves




Figure 6.11: night 2 residuals

Chapter 7

Conclusions and Future prospects

We find ourselves on the threshold of a vast and awesome Universe that utterly dwarfs-in time, in space, and in potential-the tidy anthropocentric proscenium of our ancestors.

Carl Sagan – Pale blue dot

In this thesis I have presented new results from transmission spectroscopy at high and low resolution, I will present the main conclusions from each chapter, with avenues for future research, followed by concluding remark.

7.1 The hunt for hot Neptunes with NGTS

In Chapter 3 I described some of the contributions I have made to the Next Generation Transit Survey.

I participated in the design of a code to simulate the yield of planets that could be expected to be detected by NGTS over the course of a 4 year survey. An important component of this simulation was to estimate the fraction of planets that, once detected by NGTS, could be followed up by radial velocity observations within a reasonable time. Observing time on systems such as HARPS and HARPS-N, and ESPRESSO in the future, is valuable, so not every planet can be followed up. Additionally, in reality a large fraction of false positives will also be detected. I estimated that approximately 200 hot-Neptunes and super earths could be expected to be detected with modest follow-up time requirements. This would provide a

statistical sample of intermediate mass planets, which would form a vital input into theories of planet formation and evolution. The majority of transiting planets in this size range currently known are from *Kepler*, and most do not have well constrained masses, and therefore bulk densities, due to their faint host stars.

Of the detectable planets, 40 orbit bright stars. This would provide opportunities for rich atmospheric characterization, such as described in the other Chapters of this thesis. In particular, transmission spectroscopy of these planets may help break the degeneracy between rocky worlds with hydrogen envelopes, and more ice and volatile rich planets.

I analysed 17 nights of data taken of the *Kepler* field with an NGTS unit in Geneva. I showed that the telescope was performing in line with the design, and was capable of detecting a hot-Jupiter in a single transit. Periodic variables with mmag level amplitudes were found in the data, making these the highest precision lightcurves ever achieved in a wide-field ground-based survey.

The NGTS project is in the early stages, and there is much still to be done to characterise the instrument performance and validate early planet candidates. Of particular interest is the stability of the instrument over long timeframes, as this will have an impact on the ability to recover weak signals that must be combined. My analysis of the data taken in Geneva showed that certainly over the 17 days studied the instrument was stable enough to detect low amplitude variables. A detailed study of a field over a whole observing season (~3 months) would still be valuable. The diagnostics that I generate, such as the quality and stability of the astrometric solution and the focuser stability could be tested for correlation with any systematics that do appear in the data.

Additionally, with limited telescope time for follow up observations, and a large number of candidate planets expected, extensions to the planet catch simulation to include false positives from, e.g., blending scenarios could help informing the prioritisation of candidates, or selection of future fields.

NGTS over the coming years will provide a wealth of new and interesting targets to study, and the design of the study is intended to insure that a sample of these are bright enough that the atmospheric studies described in my other chapters would be applicable to them.

7.2 Reconstructing the high energy irradiation of the evaporating hot Jupiter HD 209458b

Planets are not static objects over their lifetime, but can be sculpted by the extreme radiation of their parent stars. This is known to be the case for planets in the solar system, and has been observed on several exoplanets (Vidal-Madjar et al., 2003; Ehrenreich et al., 2015; Lecavelier des Etangs et al., 2010; Haswell et al., 2012). One of these evaporating worlds is HD 209458b. It is shrouded by an expanding exosphere of evaporating material, and may lose as much as 10^{10} g s^{-1} . This will only result in the loss of a few % of its mass over Gigayear timescales, but it is possible that the same mechanism has significant consequences for smaller planets. It is therefore important to develop an understanding of the mass loss process, but this has previously not been possible for HD 209458b because a key component is missing - the high energy flux of the star. In Chapter 4, I showed that it is possible to recover this flux by building a coronal and transition region model for the star, using constraints for different temperatures of plasma from UV and X-ray sources.

I constructed a “Differential Emission Measure” model using individual line strengths in the UV that were measured by *HST* COS, and broad bands in the X-ray as recorded in the 3XMM DR5 catalogue. The model of the corona and transition region covered temperatures between $10^{4.1}$ and 10^8 K. Using an emission model (CHIANTI (v8.0)) I showed that the majority of EUV flux is emitted between $10^{4.5}$ and $10^{6.5}$ K, which is well sampled by the formation temperature of the high excitation UV lines used in this study. The advantage of using these lines is that they are not significantly absorbed by the ISM.

I found that HD 209458 is a relatively quiet star, with a total hydrogen-ionising flux of $10^{28.26} \text{ erg s}^{-1}$, which is similar to that produced by the Sun. The spectral index is softer than the Sun, with a higher percentage of this flux coming from EUV.

The mass loss of HD 189733b was estimated to be between 8 and $40 \times 10^{10} \text{ g s}^{-1}$ (Linsky et al., 2010), but I showed that this high mass loss rate is likely incorrect, as it would require a mass loss efficiency approaching 100%. Most physical models of mass loss suggest that the efficiency should not exceed 20%.

In fact, there are significant degeneracies between inferred mass loss rates and the outflow structure and ionization levels Bourrier and Lecavelier des Etangs (2013). Modelling the outflow correctly requires a recovered spectrum, not only a single flux value, as the spectral shape effects the physics of the outflow (Guo and Ben-Jaffel, 2015). Recovering the full spectrum in the manner I have shown in

therefore an essential input to models of mass loss.

Understanding and correctly modeling mass loss will be extremely important in the characterisation of smaller and potentially habitable worlds in the future. The majority of potentially habitable worlds found to date are around low mass stars that are highly active (Armstrong et al., 2015). A minimum level of mass loss may be required to make planets habitable in the first place by removing the thick hydrogen/helium primary atmosphere (Owen and Mohanty, 2016), but too high levels could destroy the planets water inventory (e.g. Wheatley et al., 2016). Characterizing the high energy emission of a star through the Differential Emission Measure is therefore extremely important to the search for life.

7.3 Rotation and Eastward winds on HD 189733b

In Chapter 5 I demonstrated a new technique that is capable of *spatially resolving* the atmospheres of exoplanets. The technique uses the non-uniform surface brightness of the star, which the planet samples during the transit, to separate the contributions from the two sides of the planet to the absorption profile.

I showed that when measuring line profiles at high resolution, the Rossiter-McLaughlin effect can inject significant systematics into a transmission spectrum if the phase coverage isn't even, or if the data are combined in the wrong reference frame. This likely explains systematics seen by previous authors, and an average atmosphere velocity measured by Wyttenbach et al. (2015) which is significantly higher than both my result, and an independent detection by Brogi et al. (2016).

What sets my analysis apart from that of previous studies is the ability to separate the contributions from the two sides of the planet, instead of only giving a single average velocity value for the atmosphere. My results reveal a strong eastward motion of the atmosphere of HD 189733b, with a redshift of $2.3^{+1.3}_{-1.5}$ km s⁻¹ on the leading limb of the planet and a blueshift of $5.3^{+1.0}_{-1.4}$ km s⁻¹ on the trailing limb. The velocity on the leading limb is consistent with the expectation for a tidally locked planet, but the trailing limb displays a $2.4^{+1.0}_{-1.4}$ km s⁻¹ excess velocity, which I interpret as evidence for an eastward equatorial jet.

A superrotating jet is predicted by atmospheric models, and is strongly suggested by the offset hotspot observed in *spitzer* phase curves of HD 189733b.

An implication of this is that the sodium line in the atmosphere of HD 189733b is rotationally broadened. This likely means that temperature pressure profiles constructed by fitting to this profile that do not account for the rotational broadening are incorrect.

Combining the data from multiple nights would increase the potential signal to noise. There is no reason why higher resolution maps of the atmospheres could not then be made, or a different parameterisation of the rotation velocity could provide information on the width of the equatorial jet.

Since the sodium absorption is expected to be produced higher in the atmosphere than the carbon monoxide and water lines that are observable in the infrared with instruments such as CRIRES (e.g. Snellen et al., 2010), it would be valuable to compare the rotation speeds measured with different absorption lines, as this may provide information on the vertical structure of weather systems.

Measuring the strength of equatorial jets on other planets and comparing to this case would allow specific predictions made by computational models: for example, that the jet should increase in strength as the equilibrium temperature is increased (Showman et al., 2013).

Perhaps at some point, with the next generation of instrumentation, a technique like this could provide information on the atmosphere of a planet in or close to the habitable zone of an M dwarf. It would be a particularly interesting test of the heat transport on a tidally locked Earth-like planet, which could have major implications for the habitability of such an environment.

7.4 Probing the atmosphere of WASP-52b

Sing et al. (2015) showed that hot Jupiters exist on a continuum from cloudy to clear, but there is so far no clear correlation between a planet's parameters (e.g., equilibrium temperature) and its emergent spectrum. It will be necessary to measure the spectra of many planets from across the discovery space to begin to tease out correlations. In Chapter 6 I presented an optical transmission spectrum for the inflated hot Jupiter WASP-52b.

I gathered the data using ACAM on WHT, which is advantageous for its optical simplicity, and the wavelength coverage achievable in a single shot ($\sim 4000\text{--}9000\text{ \AA}$). Steps were taken to avoid the problem of differential slit losses - the data were taken with a very wide $27''$ slit, and were manually guided throughout the observations.

The lightcurves were detrended with a Gaussian process method. The common noise model recovered from the white light curves was removed from the individual wavelength bins in order to remove the correlated systematics, and the system parameters were fixed before fitting. Stellar activity led to a marginally different depth between the two nights, but the effect appeared to be grey, as apart from an

offset the spectra from the two nights matched.

I found that the resulting transmission spectrum was particularly flat, which is suggestive of a cloud dominated atmosphere. The interpretation was complicated by a significant feature redwards of 7750 \AA for which I do not have a firm explanation. I found that several systematic scenarios were unlikely, so it may be that this is a signal of an additional absorbing species in the atmosphere. Shortwards of 7750 \AA the best fitting atmosphere is a opaque cloud layer at or above 1 mbar, which produces a flat spectrum.

Despite having the same equilibrium temperature and a similar host star to HD 189733b, WASP-52b has a significantly different transmission spectrum, the large Rayleigh slope that dominates on HD 189733b can be ruled out with very high confidence, though more modest slopes where the haze settles and reduces the effective scale height cannot be ruled out to the same level of confidence. Other classes of atmosphere can be ruled out with high confidence.

Ultimately, the spectra of many more exoplanets will have to be gathered before we understand the underlying trends. My results show that accuracies comparable to those of *HST* can be achieved for broad band spectroscopy from the ground.

7.5 Concluding remarks

The study of exoplanets over the last 20 years has been rich and exciting, with surprises at nearly every turn. There are major questions still unanswered on the formation, evolution, composition and structure of exoplanet atmospheres. Projects like NGTS will allow the detection of bright, observable planets in areas of parameter space previously poorly explored. Further development in the instrumentation and application of transmission spectroscopy will be required to make progress with the complex problems of atmosphere retrieval. It is impossible to guess what the next 20 years will bring, but it is almost certain that it will be just as thrilling and unexpected.

Bibliography

- Agol, E., Cowan, N. B., et al. 2010, *ApJ*, 721, 1861
- Alcock, C., Allsman, R. A., et al. 1995, *ApJ*, 454, 1
- Alonso, R., Brown, T. M., et al. 2004, *ApJ*, 613, L153
- Alsubai, K. A., Parley, N. R., et al. 2013, *Acta Astron.*, 63, 465
- Ambikasaran, S., Foreman-Mackey, D., et al. 2016, *IEEE Trans. Pattern Anal. Mach. Intell.*, 38, 252
- Armstrong, D. J., Pugh, C. E., et al. 2015, *MNRAS*, 455, 3110
- Arnaud, K. A. 1996, *ASPC*, 101
- Arrhenius, S. *Das Werden der Welten*. Leipzig, 1913
- Asplund, M. 2005, *Annu. Rev. Astron. Astrophys.*, 43, 481
- Asplund, M., Grevesse, N., and Jacques Sauval, A. 2006, *Nucl. Phys. A*, 777, 1
- Atreya, S., Mahaffy, P., et al. 2003, *Planet. Space Sci.*, 51, 105
- Bahcall, J. N. and Soneira, R. M. 1984, *ApJS*, 55, 67
- Bakos, G. Á., Lázár, J., et al. 2002, *PASP*, 114, 974
- Ballester, G. E. and Ben-Jaffel, L. 2015, *ApJ*, 804, 116
- Baluev, R. V., Sokov, E. N., et al. 2015, *MNRAS*, 450, 3101
- Baraffe, I., Chabrier, G., et al. 2003, *A&A*, 402, 701
- Barge, P., Baglin, A., et al. 2008, *A&A*, 482, L17
- Batalha, N. M., Rowe, J. F., et al. 2013, *ApJS*, 204, 24

- Batygin, K. and Stevenson, D. J. 2010, *ApJ*, 714, L238
- Bayliss, D. D. R. and Sackett, P. D. 2011, *ApJ*, 743, 103
- Beaulieu, T., Kerins, E., et al. 2008, *eprint arXiv*, 0808, 5
- Beichl, I. and Sullivan, F. 2000, *Comput. Sci. Eng.*, 2, 65
- Bento, J., Wheatley, P. J., et al. 2014, *MNRAS*, 437, 1511
- Bento, J. *Research and development of ground-based transiting extrasolar planet projects*. Ph.d thesis, 2012
- Berta-Thompson, Z. K., Irwin, J., et al. 2015, *Nature*, 527, 204
- Beuzit, J.-L., Feldt, M., et al. page 701418, 2008
- Bickerton, A. W. 1878, *Trans N.Z. Inst*, 11, 128
- Birkby, J., Hodgkin, S., Pinfield, D., and Consortium, W. In *ASP*, page 803, 2011
- Birkby, J. L., Kok, R. J.de, et al. 2013, *MNRAS*, 436, L35
- Birney, D. S., Gonzalez, G., and Oesper, D. *Observational Astronomy*. Cambridge University Press, 2 edition, 2006. ISBN 9780521853705
- Blanco-Cuaresma, S., Soubiran, C., Heiter, U., and Jofré, P. 2014, *A&A*, 569, A111
- Bodenheimer, P., Lin, D. N. C., and Mardling, R. A. 2001, *ApJ*, 548, 466
- Bonfils, X., Lo Curto, G., et al. 2013, *A&A*, 556, A110
- Borsa, F., Rainer, M., and Poretti, E. 2016, *A&A*, 590, A84
- Borucki, W. J., Koch, D., et al. 2009, *Science*, 325, 709
- Bouchy, F., Udry, S., et al. 2005, *A&A*, 19, 5
- Bourrier, V. and Lecavelier des Etangs, A. 2013, *A&A*, 557, A124
- Boyajian, T., Braun, K.von, et al. 2014, *MNRAS*, 447, 846
- Brace, L. H., Kasprzak, W. T., et al. 1987, *J. Geophys. Res.*, 92, 15
- Brahm, R., Jordán, A., et al. 2016, *Astron. J.*, 151, 89
- Broggi, M., Kok, R. J.de, et al. 2016, *ApJ*, 817, 106

- Brogi, M., Snellen, I. a. G., et al. 2012, *Nature*, 486, 502
- Brown, T. M. 2001, *ApJ*, 553, 1006
- Buchhave, L. a., Bizzarro, M., et al. 2014, *Nature*, 509, 593
- Buffon, G. L. L. *De la Formation des Planetes*. Paris, 1745
- Bundy, K. and Marcy, G. 2000, *PASP*, 112, 1421
- Burke, C. J., McCullough, P. R., et al. 2008, *ApJ*, 686, 1331
- Calabretta, M. R. and Greisen, E. W. 2002, *A&A*, 395, 1077
- Campbell, B., Walker, G. A. H., and Yang, S. 1988, *ApJ*, 331, 902
- Charbonneau, D. B., Allen, L. E., et al. 2005, *ApJ*, 626, 523
- Charbonneau, D., Brown, T. M., Latham, D. W., and Mayor, M. 2000, *ApJ*, 529, L45
- Charbonneau, D., Brown, T. M., Noyes, R. W., and Gilliland, R. L. 2002, *ApJ*, 568, 377
- Charbonneau, D., Berta, Z. K., et al. 2009, *Nature*, 462, 891
- Chassefière, E. 1996, *J. Geophys. Res.*, 101, 26039
- Chatterjee, S., Ford, E. B., Matsumura, S., and Rasio, F. a. 2007, *ApJ*, 686, 25
- Claret, A. and Bloemen, S. 2011, *A&A*, 529, A75
- Collier Cameron, A., Davidson, V. A., et al. 2009, *MNRAS*, 400, 451
- Collier Cameron, A., Bruce, V. A., et al. 2010, *MNRAS*, 403, 151
- Cosentino, R., Lovis, C., et al. 2014, *Proc. SPIE*, 9147, 91478C
- Coustenis, A., Schneider, J., et al. In *ASP Conf. Ser. Vol. 119*, page 101, 1997. ISBN 9788578110796
- Coustenis, A., Schneider, J., et al. In *ASP Conf. Ser. vol 134*, page 296, 1998. ISBN 9788578110796
- Cowan, N. B., Agol, E., and Charbonneau, D. 2007, *MNRAS*, 379, 641
- Csizmadia, S., Pasternacki, T., et al. 2012, *A&A*, 549, A9

- Czesla, S., Klocová, T., et al. 2015, *A&A*, 582, A51
- Davis, T. a. and Wheatley, P. J. 2009, *MNRAS*, 396, 1012
- Kok, R.de, Brogi, M., et al. 2013, *A&A*, 554, A82
- Lalande, J. J. L. F.de. *Astronomie*. Chez la Veuve Desaint, Paris, 1771
- Wit, J.de, Wakeford, H. R., et al. 2016, pages 1–27
- Del Zanna, G., Dere, K. P., et al. 2015, *A&A*, 582, A56
- Deming, D., Seager, S., and Richardson, L. J. 2005, *Nature*, 434, 740
- Deming, D., Wilkins, A., et al. 2013, *ApJ*, 774, 95
- Demory, B.-O., Gillon, M., et al. 2016, *Nature*, 532, 207
- Demory, B.-O., Gillon, M., Madhusudhan, N., and Queloz, D. 2016, *MNRAS*, 455, 2018
- Dere, K. P., Landi, E., et al. 1997, *A&AS*, 125, 149
- Désert, J.-M., Vidal-Madjar, A., et al. 2008, *A&A*, 492, 585
- Désert, J.-M., Sing, D., et al. 2011, *A&A*, 526, A12
- Désert, J.-M., Lecavelier des Etangs, A., et al. 2009, *ApJ*, 699, 478
- Dhillon, V. S., Marsh, T. R., et al. 2007, *MNRAS*, 378, 825
- Diamond-Lowe, H., Stevenson, K. B., et al. 2014, *ApJ*, 796, 66
- Doyle, L. R., Carter, J. A., et al. 2011, *Science*, 333, 1602
- Dravins, D., Lindegren, L., Mezey, E., and Young, A. . 1998, *PASP*, 110, 1118
- Dressing, C. D. and Charbonneau, D. 2015, *ApJ*, 807, 45
- Dyer, M., Frieze, A., and Kannan, R. 1991, *J. ACM*, 38, 1
- Eastman, J., Siverd, R., and Gaudi, B. S. 2010, *PASP*, 122, 935
- Ehrenreich, D. and Désert, J.-M. 2011, *A&A*, 529, A136
- Ehrenreich, D., Lecavelier Des Etangs, A., et al. 2008, *AA*, 483, 933
- Ehrenreich, D., Bourrier, V., et al. 2015, *Nature*, 522, 459

- Einstein, A. 1936, *Science*, 84, 506
- Enoch, B., Collier Cameron, A., and Horne, K. 2012, *A&A*, 540, A99
- Erkaev, N. V., Kulikov, Y. N., et al. 2007, *A&A*, 472, 329
- Eshleman, V. R. 1979, *Science*, 205, 1133
- Espinoza, N. and Jordan, A. 2015, *Mnras*, 450, 1879
- Evans, T. M., Aigrain, S., et al. 2015, *MNRAS*, 451, 680
- Evans, T. M., Sing, D. K., et al. 2016, *ApJ*, 822, L4
- Faria, J. P., Santos, N. C., et al. 2016, 25, 18
- Farihi, J. 2016, *New Astron. Rev.*, 71, 9
- Feldman, U. and Widing, K. G. 2002, *Phys. Plasmas*, 9, 629
- Fischer, D. a. and Valenti, J. 2005, *ApJ*, 622, 1102
- Fisher, R. A. *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh, 1925. ISBN ISBN 0-05-002170-2
- Fludra, A., Zanna, G. D., Alexander, D., and Bromage, B. 1999, *J. Geophys. Res. Sp. Phys.*, 104, 9709
- Ford, E. B., Ragozzine, D., et al. 2012, *ApJ*, 756, 185
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. 2012, *arXiv Prepr. arXiv ...*, 125, 1
- Foreman-Mackey, D., Hogg, D. W., Lang, D., and Goodman, J. 2013, *PASP*, 125, 306
- Fortney, J. J., Marley, M. S., and Barnes, J. W. 2007, *ApJ*, 659, 1661
- Fortney, J. J., Lodders, K., Marley, M. S., and Freedman, R. S. 2008, *ApJ*, 678, 1419
- Fortney, J. J. 2004, *Science*, 305, 1414
- Fortney, J. J., Sudarsky, D., et al. 2003, *ApJ*, 589, 615
- France, K., Stocke, J. T., et al. 2010, *ApJ*, 712, 1277

- Fressin, F., Torres, G., et al. 2013, *ApJ*, 766, 81
- Fressin, F., Ballard, S., et al. 2014, *ApJ*, 790, 12
- Fruchter, A. S. and Hook, R. N. 2002, *PASP*, 114, 144
- Gabriel, C., Denby, M., et al. 2004, *Astron. Data Anal. Softw. Syst. XIII*, 314, 759
- Gaudi, B. S. and Winn, J. N. 2007, *ApJ*, 655, 550
- Gibson, N. P., Pont, F., and Aigrain, S. 2011, *MNRAS*, 411, 2199
- Gibson, N. P., Aigrain, S., et al. 2012, *MNRAS*, 422, 753
- Gibson, N. P., Aigrain, S., et al. 2013, *MNRAS*, 436, 2974
- Gierasch, P. J. and Conrath, B. J. 1987, *J. Geophys. Res.*, 92, 15019
- Gilliland, R. L., Jenkins, J. M., et al. 2010, *ApJ*, 713, L160
- Gillon, M., Jehin, E., et al. 2016, *Nature*, 533, 221
- Gilmore, G. and Reid, N. 1983, *MNRAS*, 202, 1025
- Gloria, E. D., Snellen, I. A. G., and Albrecht, S. 2015, 84, 1
- Gonzalez, G. 1997, *MNRAS*, 285, 403
- Goodman, J. and Weare, J. 2010, *Commun. Appl. Math. Comput. Sci.*, 5, 65
- Gould, A., Dorsher, S., Gaudi, B. S., and Udalski, A. 2006, *Acta Astron.*, 56, 1
- Grange, A. and Goad, M. Poster presented at the Exoplanet UK Community Meeting Cambridge, 2014, April 14-16, 2014
- Grether, D. and Lineweaver, C. H. 2007, *ApJ*, 669, 1220
- Guillot, T., Burrows, A., et al. 1996, *ApJ*, 459, 1689
- Guillot, T. *Saturn from Cassini-Huygens*. Springer Science & Business Media, 2009. ISBN 1402092172, 9781402092176
- Guillot, T. 2010, *A&A*, 520, 27
- Guillot, T. and Showman, A. P. 2002, *A&A*, 165, 11
- Guo, J. H. and Ben-Jaffel, L. 2015, pages 1–52

- Harrington, J., Hansen, B. M., et al. 2006, *Science*, 314, 623
- Haswell, C. a., Fossati, L., et al. 2012, *ApJ*, 760, 79
- Hatzes, A. P., Cochran, W. D., et al. 2003, *ApJ*, 599, 1383
- Hebrard, G., Cameron, a. C., et al. 2012, *A&A*, 134, 13
- Heng, K., Hayek, W., Pont, F., and Sing, D. K. 2012, *MNRAS*, 420, 20
- Heng, K., Wytttenbach, A., et al. 2015, *ApJL*, 803, L9
- Hinkley, S., Oppenheimer, B. R., et al. 2011, *PASP*, 123, 74
- Holman, M. J., Winn, J. N., et al. 2006, *ApJ*, 652, 1715
- Holt, J. R. 1893, *A&A*, 12, 646
- Horne, K. 2003, *Sci. Front. Res. Extrasolar Planets*, 294, 361
- Howard, A. W., Marcy, G. W., et al. 2012, *ApJS*, 201, 15
- Howell, S. B. *Handbook of CCD Astronomy*. Cambridge University Press, 2 edition, 2006. ISBN 113945076X
- Howell, S. B., Sobeck, C., et al. 2014, *PASP*, 126, 398
- Hubbard, W. B., Fortney, J. J., et al. 2001, *ApJ*, 560, 413
- Huber, D., Chaplin, W. J., et al. 2013, *ApJ*, 767, 127
- Huitson, C. M., Sing, D. K., et al. 2012, *MNRAS*, 422, 2477
- Husser, T.-O., Wende-von Berg, S., et al. 2013, *A&A*, 553, A6
- Ingraham, P., Marley, M. S., et al. 2014, *ApJ*, 794, L15
- Ingrosso, G., Novati, S. C., et al. 2011, *Gen. Relativ. Gravit.*, 43, 1047
- Irwin, M. J., Lewis, J., et al. In Quinn, P. J. and Bridger, A., editors, *Optim. Sci. Return Astron. through Inf. Technol. Ed. by Quinn*, volume 5493, pages 411–422, 2004
- Irwin, P., Barstow, J. K., et al. 2014, pages 1–53
- Irwin, P., Calcutt, S., and Taylor, F. 1997, *Adv. Sp. Res.*, 19, 1149
- Irwin, P., Teanby, N., et al. 2008, *J. Quant. Spectrosc. Radiat. Transf.*, 109, 1136

- Jackson, A. P., Davis, T. a., and Wheatley, P. J. 2012, *MNRAS*, 422, 2024
- Janesick, J. and Blouke, M. 1987, *Sky Telescope*, September, 238
- JEANS, J. *THE DYNAMICAL THEORY OF GASES*. Cambridge University Press, 1916
- Jeffreys, H. 1946, *Proc. R. Soc. A Math. Phys. Eng. Sci.*, 186, 453
- Jensen, A. G., Redfield, S., et al. 2011, *ApJ*, 743, 203
- Jha, A. *A Textbook Of Applied Physics, Volume 1*. I. K. International Pvt Ltd, 2009. ISBN 9380026773, 9789380026770
- Jovanovic, N., Martinache, F., et al. 2015, *PASP*, 127, 890
- Judge, P. G., Solomon, S. C., and Ayres, T. R. 2003, *ApJ*, 593, 534
- Karkoschka, E. 1994, *Icarus*, 111, 174
- Kashyap, V. L., Drake, J. J., and Saar, S. H. 2008, *ApJ*, 687, 1339
- Kataria, T., Sing, D. K., et al. 2016, *ApJ*, 821, 9
- Kirk, J., Wheatley, P. J., and Loudon, T. 2016, *MNRAS*, *submitted*
- Kirkpatrick, J. D., Henry, T. J., and Irwin, M. J. 1997, *Astron. J.*, 113, 1421
- Knutson, H. A., Charbonneau, D., et al. 2008, *ApJ*, 673, 526
- Knutson, H. A., Charbonneau, D., et al. 2007, *Nature*, 447, 183
- Knutson, H. A., Charbonneau, D., et al. 2007, *ApJ*, 655, 564
- Knutson, H. A., Charbonneau, D., et al. 2009, *ApJ*, 690, 822
- Knutson, H. A., Lewis, N., et al. 2012, *ApJ*, 754, 22
- Knutson, H. a., Benneke, B., Deming, D., and Homeier, D. 2014, *Nature*, 505, 66
- Koch, D. G., Borucki, W. J., et al. 2010, *ApJ*, 713, L79
- Koester, D., Gänsicke, B. T., and Farihi, J. 2014, *A&A*, 566, A34
- Koppenhoefer, J., Saglia, R. P., et al. 2013, *MNRAS*, 435, 3133
- Koskinen, T. T., Yelle, R. V., Harris, M. J., and Lavvas, P. 2013, *Icarus*, 226, 1695

- Koskinen, T. T., Aylward, A. D., and Miller, S. 2007, *Nature*, 450, 845
- Koskinen, T., Harris, M., Yelle, R., and Lavvas, P. 2013, *Icarus*, 226, 1678
- Kovács, G., Bakos, G., and Noyes, R. W. 2005, *MNRAS*, 356, 557
- Kreidberg, L., Bean, J. L., et al. 2014, *Nature*, 505, 69
- Kreidberg, L., Line, M. R., et al. 2015, *ApJ*, 814, 66
- Kretzschmar, M., Lilensten, J., and Aboudarham, J. 2004, *A&A*, 419, 345
- Kurucz, R. *ATLAS9 Stellar Atmosphere Programs and 2 km/s grid*. Kurucz CD-ROM No. 13. Cambridge, 1993
- Lagrange, A. M., Bonnefoy, M., et al. 2010, *Science*, 329, 57
- Laming, J. M. 2004, *ApJ*, 614, 1063
- Lammer, H., Selsis, F., et al. 2003, *ApJ*, 598, L121
- Lammer, H., Erkaev, N. V., et al. 2016, 6, 1
- Lang, D., Hogg, D. W., et al. 2010, *Astron. J.*, 139, 1782
- Langland-Shula, L. E., Vogt, S. S., et al. 2009, *ApJ*, 696, 1355
- Latham, D. W., Stefanik, R. P., et al. 1989
- Laughlin, G., Crismani, M., and Adams, F. C. 2011, *ApJ*, 7, 3
- Lecavelier des Etangs, A., Ehrenreich, D., et al. 2010, *A&A*, 514, A72
- Leconte, J. and Chabrier, G. 2012, *A&A*, 540, A20
- Lee, E. J. and Chiang, E. 2016, *ApJ*, 817, 90
- Lee, G. and Juncher, D. 2015, 12, 1
- Lendl, M., Delrez, L., et al. 2016, *A&A*, 587, A67
- Li, C.-H., Benedick, A. J., et al. 2008, *Nature*, 452, 610
- Lin, D. N. C. and Papaloizou, J. 1986, *ApJ*, 309, 846
- Line, M. R. and Yung, Y. L. 2013, *ApJ*, 779, 3
- Linsky, J. L., Yang, H., et al. 2010, *ApJ*, 717, 1291

- Linsky, J. L., Fontenla, J., and France, K. 2014, *ApJ*, 780, 61
- Lopez, E. and Fortney, J. 2013, *ApJ*, 11, 12
- Lopez, E. D. and Fortney, J. J. 2014, *ApJ*, 792, 1
- Lopez, E. D. and Fortney, J. J. 2015, *ApJ*, 818, 4
- Lopez, E. D., Fortney, J. J., and Miller, N. 2012, *ApJ*, 761, 59
- Louden, T. and Wheatley, P. J. 2015, *ApJ*, 814, L24
- Louden, T., Wheatley, P. J., and Briggs, K. 2016, *eprint arXiv*, 1605.07987
- Lovis, C. and Pepe, F. 2007, *A&A*, 468, 1115
- Lovis, C., Pepe, F., et al. In McLean, I. S. and Iye, M., editors, *Proc. SPIE*, volume 6269, page 62690P, 2006. ISBN 0819463345
- Lundstrum, I., Ardeberg, A., Maurice, E., and Lindgren, H. 1991, *A&AS*, 91, 199
- Macintosh, B., Graham, J. R., et al. 2014, *Proc. Natl. Acad. Sci.*, 111, 12661
- Madhusudhan, N., Harrington, J., et al. 2011, *Nature*, 469, 64
- Mandel, K. and Agol, E. 2002, *ApJ*, 580, L171
- Martins, J. H. C., Santos, N. C., et al. 2015, *A&A*, 576, A134
- Matsuo, T., Shibai, H., Ootsubo, T., and Tamura, M. 2007, *ApJ*, 662, 1282
- Maxted, P. F. L., Anderson, D. R., et al. 2011, *PASP*, 123, 547
- Mayor, M., Pepe, F., et al. 2003, *The Messenger*, pages 20–24
- Mayor, M., Marmier, M., et al. 2011, *eprint arXiv*, 1109, 2497
- Mayor, M. and Queloz, D. 1995, *Nature*, 378, 355
- Mazeh, T., Naef, D., et al. 2000, *ApJ*, 532, L55
- McCormac, J., Pollacco, D., et al. 2013, *PASP*, 125, 548
- McCullough, P. R., Stys, J. E., et al. 2005, *PASP*, 117, 783
- McLaughlin, D. B. 1924, *ApJ*, 60, 22
- McLaughlin, S. F. 1983, *Astron. J.*, 88, 1633

- McNeil, W. J., Murad, E., and Lai, S. T. 1995, *J. Geophys. Res.*, 100, 16847
- Mena, E. D., Santos, N. C., and Sousa, S. G. 2015
- Miller-Ricci Kempton, E. and Rauscher, E. 2012, *ApJ*, 751, 117
- Mocquet, A., Grasset, O., and Sotin, C. 2014, *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.*, 372, 20130164
- Moehler, S., Freudling, W., et al. 2010, *Pasp*, 122, 93
- Molaro, P., Monaco, L., Barbieri, M., and Zaggia, S. 2013, *MNRAS*, 429
- Morton, T. D., Bryson, S. T., et al. 2016
- Moutou, C., Coustenis, A., et al. 2001, *A&A*, 371, 260
- Murray-Clay, R. a., Chiang, E. I., and Murray, N. 2009, *ApJ*, 693, 23
- Nakajima, T., Oppenheimer, B. R., et al. 1995, *Nature*, 378, 463
- Naoz, S., Farr, W. M., et al. 2011, *Nature*, 473, 187
- Nikolov, N., Sing, D. K., et al. 2013, *MNRAS*, 437, 46
- O’Donovan, F. T., Charbonneau, D., et al. 2006, *ApJ*, 651, L61
- Oppenheimer, B. 1998, *Astrophys. . . .*, 1, 932
- Osterman, S., Green, J., et al. 2011, *Astrophys. Space Sci.*, 335, 257
- Owen, J. E. and Adams, F. C. 2014, *MNRAS*, 444, 3761
- Owen, J. E. and Jackson, A. P. 2012, *MNRAS*, 425, 2931
- Owen, J. E. and Mohanty, S. 2016, *MNRAS*, 459, 4088
- Owen, J. E. and Wu, Y. 2013, *ApJ*, 775, 105
- Owen, J. E. and Wu, Y. 2016, *ApJ*, 817, 107
- Pallé, E., Osorio, M. R. Z., et al. 2009, *Nature*, 459, 814
- Paquette, C., Pelletier, C., Fontaine, G., and Michaud, G. 1986, *ApJS*, 61, 197
- Parmentier, V., Showman, A. P., and Lian, Y. 2013, *A&A*, 558, A91
- Parmentier, V., Fortney, J. J., et al. 2016

- Parviainen, H. and Aigrain, S. 2015, *MNRAS*, 453, 3822
- Parviainen, H., Pallé, E., et al. 2016, *A&A*, 585, A114
- Penz, T., Micela, G., and Lammer, H. 2008, *A&A*, 477, 309
- Pepe, F. A., Mayor, M., et al. 2004, *A&A*, 423, 385
- Pepper, J., Pogge, R. W., et al. 2007, *PASP*, 119, 21
- Pepper, J., Kuhn, R. B., et al. 2012, *PASP*, 124, 230
- Perna, R., Menou, K., and Rauscher, E. 2010, *ApJ*, 724, 313
- Petrie, R. M. 1938, *J. R. Astron. Soc. Canada*, 32, 257
- Pickles, A. J. 1998, *PASP*, 110, 863
- Plane, J. M. C. 1991, *Int. Rev. Phys. Chem.*, 10, 55
- Pollacco, D., Skillen, I., et al. 2006, *Astrophys. Space Sci.*, 304, 253
- Pollack, J. B., Hubickyj, O., et al. 1996, *Icarus*, 124, 62
- Pont, F., Knutson, H., et al. 2008, *MNRAS*, 385, 109
- Pont, F., Sing, D. K., et al. 2013, *MNRAS*, 432, 2917
- Postman, M. 2009, *Astro2010 Decad. Comm.*, page 22
- Queloz, D., Eggenberger, A., et al. 2000, *A&A*, 359, L13
- Queloz, D., Bouchy, F., et al. 2009, *A&A*, 506, 303
- Raghavan, D., McAlister, H. A., et al. 2010, *ApJS*, 190, 1
- Rauer, H., Bockelée-Morvan, D., et al. 2000, *A&A*, 355, 573
- Rauer, H., Catala, C., et al. 2014, *Exp. Astron.*, 38, 249
- Rauscher, E. and Menou, K. 2010, *ApJ*, 714, 1334
- Redfield, S., Endl, M., Cochran, W. D., and Koesterke, L. 2008, *ApJ*, 673, L87
- Ribas, I., Guinan, E. F., Gudel, M., and Audard, M. 2005, *ApJ*, 622, 680
- Richer, H. B., Ibata, R., Fahlman, G. G., and Huber, M. 2003, *ApJ*, 597, L45
- Ricker, G. R., Winn, J. N., et al. 2014, *arXiv*, 9143, 914320

- Robin, a. C., Reyl, C., Derrire, S., and Picaud, S. 2003, *A&A*, 409, 523
- Rodler, F., Lopez-Morales, M., and Ribas, I. 2012, *ApJ*, 753, L25
- Rogers, L. A. 2015, *ApJ*, 801, 41
- Rosen, S. R., Webb, N. a., et al. 2016, *A&A*, 590, A1
- Rossiter, R. A. 1924, *ApJ*, 60, 15
- Sahu, K. C., Casertano, S., et al. 2006, *Nature*, 443, 534
- Salz, M., Schneider, P. C., Czesla, S., and Schmitt, J. H. M. M. 2016, *A&A*, 585, L2
- Santos, N. C., Israelian, G., and Mayor, M. 2004, *A&A*, 415, 1153
- Sanz-Forcada, J., Micela, G., et al. 2011, *A&A*, 532, A6
- Saumon, D. S. and Guillot, T. 2004, *ApJ*, 609, 1170
- Schatzman, E. 1962, *Ann. d'Astrophysique*, 25, 18
- Schlesinger, F. 1910, *Publ. Allegheny Obs.*, 1, 123
- Schmelz, J. T., Reames, D. V., Steiger, R.von, and Basu, S. 2012, *ApJ*, 755, 33
- Seager, S. and Sasselov, D. D. 2000, *ApJ*, 537, 916
- Seager, S. *Exoplanet Atmospheres: Physical Processes*. Princeton University Press, 2010. ISBN 978-1-4008-3530-0
- Seager, S. and Mallen-Ornelas, G. 2003, *ApJ*, 585, 1038
- Sedaghati, E., Boffin, H. M. J., et al. 2015, *A&A*, 576, L11
- Selby, M. J. and Mampaso, A. 1991, *Astrophys. Lett. Commun.*, 28
- Shematovich, V. I., Ionov, D. E., and Lammer, H. 2014, *A&A*, 571, A94
- Showman, A. P. and Guillot, T. 2002, *A&A*, 180, 14
- Showman, A. P. and Polvani, L. M. 2011, *ApJ*, 738, 71
- Showman, A. P., Fortney, J. J., Lewis, N. K., and Shabram, M. 2013, *ApJ*, 762, 24
- Showman, A. P., Lewis, N. K., and Fortney, J. J. 2015, *ApJ*, 801, 95

- Silva Aguirre, V., Davies, G. R., et al. 2015, *MNRAS*, 452, 2127
- Simons, D. 1995, *Tech. note no. 30. Gemini Obs.*
- Sing, D. K. 2010, *A&A*, 510, A21
- Sing, D. K., Désert, J.-M., et al. 2009, *A&A*, 505, 891
- Sing, D. K., Pont, F., et al. 2011, *MNRAS*, 416, 1443
- Sing, D. K., Huitson, C. M., et al. 2012, *MNRAS*, 426, 1663
- Sing, D. K., VidalMadjar, A., et al. 2008, *ApJ*, 686, 658
- Sing, D. K., VidalMadjar, A., et al. 2008, *ApJ*, 686, 667
- Sing, D. K., Wakeford, H. R., et al. 2014, *MNRAS*, 446, 2428
- Sing, D. K., Fortney, J. J., et al. 2015, *Nature*, 529, 59
- Skrutskie, M. F., Cutri, R. M., et al. 2006, *Astron. J.*, 131, 1163
- Snellen, I. A. G., Albrecht, S., Mooij, E. J. W.de, and Le Poole, R. S. 2008, *A&A*, 487, 357
- Snellen, I. a. G., Kok, R. J.de, Mooij, E. J. W.de, and Albrecht, S. 2010, *Nature*, 465, 1049
- Snellen, I. a. G., Brandl, B. R., et al. 2014, *Nature*, 509, 63
- Snellen, I., Stuik, R., et al. 2013, *EPJ Web Conf.*, 47, 03008
- Sousa, S. G., Santos, N. C., et al. 2008, *A&A*, 487, 373
- Southworth, J. 2010, *MNRAS*, 408, 1689
- Southworth, J., Wheatley, P. J., and Sams, G. 2007, *MNRAS*, 379, L11
- Spiegel, D. S., Burrows, A., and Milsom, J. A. 2011, *ApJ*, 727, 57
- Spitzer, Lyman, J. 1939, *ApJ*, 90, 675
- Steffen, J. H. 2016, *MNRAS*, 457, 4384
- Struve, O. 1952, *Obs.*, 72, 199
- Sumi, T., Kamiya, K., et al. 2011, *Nature*, 473, 349

- Swedenborg, E. *Opera Philosophica et Mineralia*. Leipzig, 1734
- Tamuz, O., Mazeh, T., and Zucker, S. 2005, *MNRAS*, 356, 1466
- Tennyson, J. and Yurchenko, S. N. 2012, *MNRAS*, 425, 21
- Triaud, A. H. M. J., Queloz, D., et al. 2009, *A&A*, 506, 377
- Tuomi, M., Jones, H. R. A., et al. 2014, *MNRAS*, 441, 1545
- Udalski, A., Zebruń, K., et al. 2002, *Acta Astron.*, 52, 115
- Uytterhoeven, K., Moya, A., et al. 2011, *A&A*, 534, 125
- Leeuwen, F.van. *Hipparcos, the New Reduction of the Raw Data*, volume 350 of *Astrophysics and Space Science Library*. Springer Netherlands, 2007. ISBN 978-1-4020-6341-1
- Maanen, A.van. 1917, *PASP*, 29, 258
- Vanderburg, A., Johnson, J. A., et al. 2015, *Nature*, 526, 546
- Venot, O., Hébrard, E., et al. 2012, *A&A*, 546, A43
- Veras, D. and Raymond, S. N. 2012, *MNRAS*, 421, L117
- Vidal-Madjar, A., Des Etangs, A. L., et al. 2003, *Nature*, 422, 143
- Vidal-Madjar, A., Dsert, J.-M., et al. 2004, *ApJ*, 604, L69
- Vidal-Madjar, A., Lecavelier des Etangs, A., et al. 2008, *ApJ*, 676, L57
- Vidal-Madjar, a., M Huitson, C., et al. 2011, *A&A*, 533, C4
- Vidal-Madjar, A., Sing, D. K., et al. 2011, *A&A*, 527, A110
- Vidal-Madjar, a., Huitson, C. M., et al. 2013, *A&A*, 54, 1
- Wakeford, H. R. and Sing, D. K. 2015, *A&A*, 573, A122
- Wakeford, H. R., Sing, D. K., et al. 2013, *MNRAS*, 435, 3481
- Wakeford, H. R., Sing, D. K., et al. 2016, *ApJ*, 819, 10
- Walker, S. R. *Analysis and optimisation of ground based transiting exoplanet surveys*. PhD thesis, 2013
- Wang, J. and Fischer, D. A. 2014, *Astron. J.*, 149, 14

- Wang, J., Fischer, D. A., Horch, E. P., and Huang, X. 2015, *ApJ*, 799, 229
- Watson, A. J., Donahue, T. M., and Walker, J. C. 1981, *Icarus*, 48, 150
- Weber, E. J. and Davis, Leverett, J. 1967, *ApJ*, 148, 217
- Weldrake, D. T. F., Bayliss, D. D. R., et al. 2008, *ApJ*, 675, L37
- Wheatley, P. J., Pollacco, D. L., et al. 2013, *EPJ Web Conf.*, 47, 13002
- Wheatley, P. J., Loudon, T., et al. 2016, 5, 1
- Williams, I. P. and Cremin, A. W. 1968, *Q. J. R. Astron. Soc.*, 9, 40
- Williams, J. P. and Cieza, L. A. 2011, *Annu. Rev. Astron. Astrophys.*, 49, 67
- Wilms, J., Allen, A., and McCray, R. 2000, *ApJ*, 542, 914
- Wilson, D. J., Gänsicke, B. T., Farihi, J., and Koester, D. 2016, *MNRAS*, 459, 3282
- Winn, J. N. 2010, *arXiv:1001.2010*
- Winn, J. N. and Fabrycky, D. C. 2014, *Annu. Rev. Astron. Astrophys.*, 53, 409
- Winn, J. N., Fabrycky, D., Albrecht, S., and Johnson, J. A. 2010, *ApJL*, 718, 145
- Wolszczan, a. and Frail, D. a. 1992, *Nature*, 355, 145
- Wolter, H. 1952, *Ann. Phys.*, 445, 94
- Wolter, H. 1952, *Ann. Phys.*, 445, 286
- Wood, B. E., Redfield, S., et al. 2005, *ApJS*, 159, 118
- Wright, J. T., Marcy, G. W., et al. 2012, *ApJ*, 753, 160
- Wytttenbach, A., Ehrenreich, D., et al. 2015, *A&A*, 577, A62
- Yan, F., Fosbury, R. a. E., et al. 2015, *ApJ*, 806, L23
- Yan, F., Fosbury, R. a. E., et al. 2015, *Int. J. Astrobiol.*, 14, 255
- Yang, F., Sbaiz, L., et al. In *2009 IEEE 12th Int. Conf. Comput. Vis. Work. ICCV Work.*, pages 2212–2219. IEEE, 2009. ISBN 978-1-4244-4442-7
- Young, R. E. 2003, *New Astron. Rev.*, 47, 1
- Zacharias, N., Finch, C. T., et al. 2013, *Astron. J.*, 145, 44

- Zapolsky, H. S. and Salpeter, E. E. 1969, *ApJ*, 158, 809
- Zechmeister, M. and Kürster, M. 2009, *eprint arXiv*, 0901, 2573
- Zellem, R. T., Lewis, N. K., et al. 2014, *ApJ*, 790, 53
- Zhu, W., Wang, J., and Huang, C. 2016, *eprint arXiv*, 1605.04310
- Zuckerman, B. In *19th Eur. Work. White Dwarfs, ASP Conf. Ser. Vol. 493*, page 291, San Francisco, 2015