The sky is made of lava: how lava worlds reveal their interiors through their atmospheres

Buchem, C.P.A. van

1 Introduction

1.1 Hell on (super-) Earth

You are a photon. A light particle released from the core of a star during the nuclear fusion of two hydrogen atoms forming a helium atom. You have spent the last one hundred thousand years[1] bouncing your way around the tightly packed interior of the star, slowly making your way towards the surface. A hundred thousand years to cover seven hundred thousand kilometres. Suddenly, you are free. Out into space. Moving at an unimaginable speed - the speed of light in a vacuum. The next seven hundred thousand kilometres take you a bit more than two seconds to cover.

Unlike the large majority of your fellow photons, however, your freedom is short-lived. Just under eight seconds after your escape from the surface of the star from which you were born, you reach a hellish place. As you approach what at first appears to be nothing more than a tiny glowing rock in the distance, quickly reveals itself to be a dramatic world.

Your progenitor, the star, is shining down unbiddingly from behind you upon the eternal day-side of this planet, appearing sixty times larger in the sky than the sun does on earth (see Figure 1.1). This has turned the entire dayside of the planet into a huge, churning, and bubbling ocean of lava. The centre of this ocean is so hot that even the heaviest of elements melt. Enormous amounts of vaporized rock and metal are carried up into the atmosphere and blown outward toward the planet's night side by super-sonic winds. As the star drops towards the horizon and you near the night side, clouds start to form. Clouds of molten glass that rain down onto cliffs made of metal.

Suddenly, you bounce off of a silicon monoxide molecule that is floating in the atmosphere of the planet. Before you know it, you are leaving the inferno behind you. You skim the surface of the star to say goodbye one final time, and then you are truly on your way. Free, out into interstellar space.

Forty years later, you notice that a star, very similar to the one you were born in, is growing in size. You also notice a tiny blue speck appearing next to it. Growing ever larger, but right as you think you might hit it, you suddenly bounce off of a golden mirror instead, into a dizzying network of countless other tiny mirrors, until you reach your final resting place in one of the charge-coupled detectors of the James Webb Space Telescope. Do not worry though, your legacy does not end here.

The energy that you release as you are absorbed by the charged-coupled detector is turned into an electric current. This current is used to magnetically charge a tiny component of a solid-state drive. About twenty-four hours later,

[1] Although technically everything happens simultaneously from the perspective of a photon, let's say that you understand the passing of time.

Figure 1.1: Apparent size comparison of the Sun and 55 Cnc A: As seen from Earth and 55 Cnc e respectively. 55 Cnc A appears about sixty times larger in the sky of 55 Cnc e than the Sun appears on Earth. The left side image was taken by the author. The right side image is composed of an image of the sun (taken by instagram.com/cosmic_background/) and an image generated using ChatGPT by OpenAI.

this magnetic charge is read out and emitted as a radio signal. The radio photon takes five seconds to travel to the nearby blueish planet. There, on its surface, it is caught by a large antenna.

Finally, after months of being sent around, copied, merged, and sent around again across a network of fibre optical cables spanning the entire planet, the evidence of your existence as a photon enters a worn down laptop on ten percent charge. After all this, the fact of your existence and the journey that you have made caused a chain of events that is now telling a few of the pixels on an LCD screen to emit less light than the neighbouring pixels (making it appear black). Photons emitted from this screen reach the eyes of a highly caffeinated bipedal primate - whose digestive system is sending signals that a far greater range of nutrients are required than the "boterham" with peanut butter and bean juice that have been offered to it in the past twenty four hours.

"There is no way that this data is real, this doesn't look anything like my models!"

You did your best. Maybe one of the other quadrillions of photons that have made the same journey as you will manage to convince this hangry monkey of your existence.

Much of modern (exoplanet) astronomy relies on maximizing the amount of information that we can draw from a single point of unresolved light. As the (somewhat dramatized) story above aims to illustrate, a great number of different things happen to a photon before it reaches the Earth, and once it has, a great many steps must be taken to draw information from it. In order to be able to appropriately interpret this information, we must understand the journey which the photon has taken and the subsequent steps taken to analyse the information from this photon as best as we can.

This is achieved by making models. Broadly speaking, astronomers that work on theoretical modelling focus on the physical and chemical processes that influence the photons before reaching a telescope. Astronomers that are more oriented towards observations model the physical processes that influence photons once they have reached the telescope and pass through all the detection instruments. In short, you could say that one focuses on what you "should" see (according to models of astronomical objects) and the other focuses on what you "do" see (according to models of telescopes and their constituent parts). It is at the conjunction of these two disciplines[2] that we try and decipher the truth behind what is out there in the universe. The work presented in this

[2] It should be noted here that none of this would be possible without the monumental amount of work done by astronomers, physicists, and engineers to ensure that working telescopes exist in the first place.

thesis starts off from a theoretical basis and then attempts to build a bridge towards what we could potentially observe.

We apply this to the niche field of hot rocky exoplanets, which are planets so close to their host star that they have a molten surface - as illustrated in the opening of this chapter and shown in Figure 1.1. Through the modelling done for this work, we try to answer one question:

“What do lava planet atmospheres tell us about their interiors?”

This introduction serves to give context as to why this is a relevant question and to provide the knowledge necessary to understand the work done in the following chapters. We end this chapter with a broad overview of the main conclusions that we draw from the work presented in this thesis.

1.2 Why do we care?

1.2.1 (exo)Planets

If you dig down to the core of astronomical research, most of it boils down to some variation of the following questions: "Why do we exist?" and "Are we alone?". Within the niche of (exo)planet astronomy, these questions lead us to ask: "How did the Earth form?" and "Are there other planets like it?". For the majority of the history of modern astronomy, the solar system bodies were the only (non-stellar) worlds besides Earth that we could observe. At first glance, this may not seem problematic. We have a large moon in our backyard, three nearby rocky planets, further out we have four gas giants with plenty of large moons, and if studying those would not suffice, we could visit the huge number of other tiny planetoids orbiting the Sun. In the grand scheme of the universe, however, this is a trivially small sample size that leaves many questions unanswered. Do planets outside the solar system exist? If so, are they similar to the planets we have here? What about their size and distance from their host stars? Are they made of similar material? Could they support life? And so on. It wasn’t until a bit over 30 years ago that we could start answering some of these questions.

In 1992, Wolszczan & Frail (1992) discovered two planets orbiting the Pulsar PSR B1257+12. Their discovery was confirmed two years later along with the discovery of a third planet in the same system (Wolszczan 1994). The exact timing of the pulses coming from this millisecond radio pulsar could only be explained with the presence of these planets. However, these systems are thought to be quite rare. At the time of writing of this thesis, only eight confirmed exoplanet detections were made through pulsar timing variations[3]. Furthermore, besides their mass, it is hard to derive any other properties of these planets

[3] exoplanetarchive.ipac.caltech.edu

using this observational technique. Hence, the discovery of the first exoplanet around a main-sequence star by Mayor & Queloz (1995) is generally considered to be the true starting shot of the era of exoplanets.

By measuring variations in the radial velocity of the sun-like star 51-Peg, Mayor & Queloz (1995) were able to derive the existence of a Jupiter-mass planet: 51-Peg b. This immediately conflicted with the classical view of planetary systems that existed based on our solar system. A Jupiter sized planet was located at a distance from its star smaller than the distance from Mercury to the Sun. This was only the start. In the three decades that followed, the number of discovered exoplanets grew exponentially, especially with the advent of dedicated exoplanet space-telescopes such as Kepler and the Transiting Exoplanet Survey Satellite (TESS) (see top panel in Figure 1.2). As of the 28th of November 2024, 5788 exoplanets have been discovered[4]. A number which will continue to increase in the foreseeable future, especially with new dedicated space missions that are already online or soon will be such as Plato, Ariel, and Cheops, the promise of exoplanet discoveries in Gaia data (Panahi et al. 2022), and a long list of ongoing and planned ground-based observing programs (e.g. CARMENES, HARPS, WASP, and TRAPPIST).

As of today, one thing has become clear: the diversity of star-planet system configurations is enormous. The bottom panel of Figure 1.2 illustrates the broad range of masses and orbital periods for all confirmed planet discoveries to date. Only a small fraction of these planets have masses and orbital periods comparable to those of the planets in our solar system (marked in purple). This is because planets are easier to detect when they are closer to their host star and have greater mass. As a result, there is a strong observational bias favouring larger, close-in planets.

If most planetary systems were analogous to our solar system - with smaller rocky planets near the star and larger gas giants farther out - we would expect to detect mostly smaller rocky planets close to their star. Instead, there is a dizzying array of different systems, sometimes referred to as the "exoplanet zoo" (Naeye 2017). Rather than conforming to the traditional binary classification of rocky planets and gas giants, many planets fall into intermediate categories, such as super-Earths and mini-Neptunes (Bean et al. 2021).

As was clear early on, thanks to 51-Peg b, gas giants are not confined to the outer parts of planetary systems. These 'hot Jupiters' - broadly defined as gas giants with an orbital period of less than 10 days - have become some of the most well-studied objects in exoplanet astronomy (Dawson & Johnson 2018). Beyond that, we've uncovered planets with highly eccentric orbits, planets around binary- (Armstrong et al. 2014) and triple-star systems (Busetti et al. 2018), and even planets orbiting stellar remains such as neutron stars (as we saw above) and white dwarfs (Vanderburg et al. 2020; Blackman et al. 2021; Mullally et al. 2024).

These discoveries make one thing increasingly apparent: planets are ubiq-

[4] https://exoplanetarchive.ipac.caltech.edu/

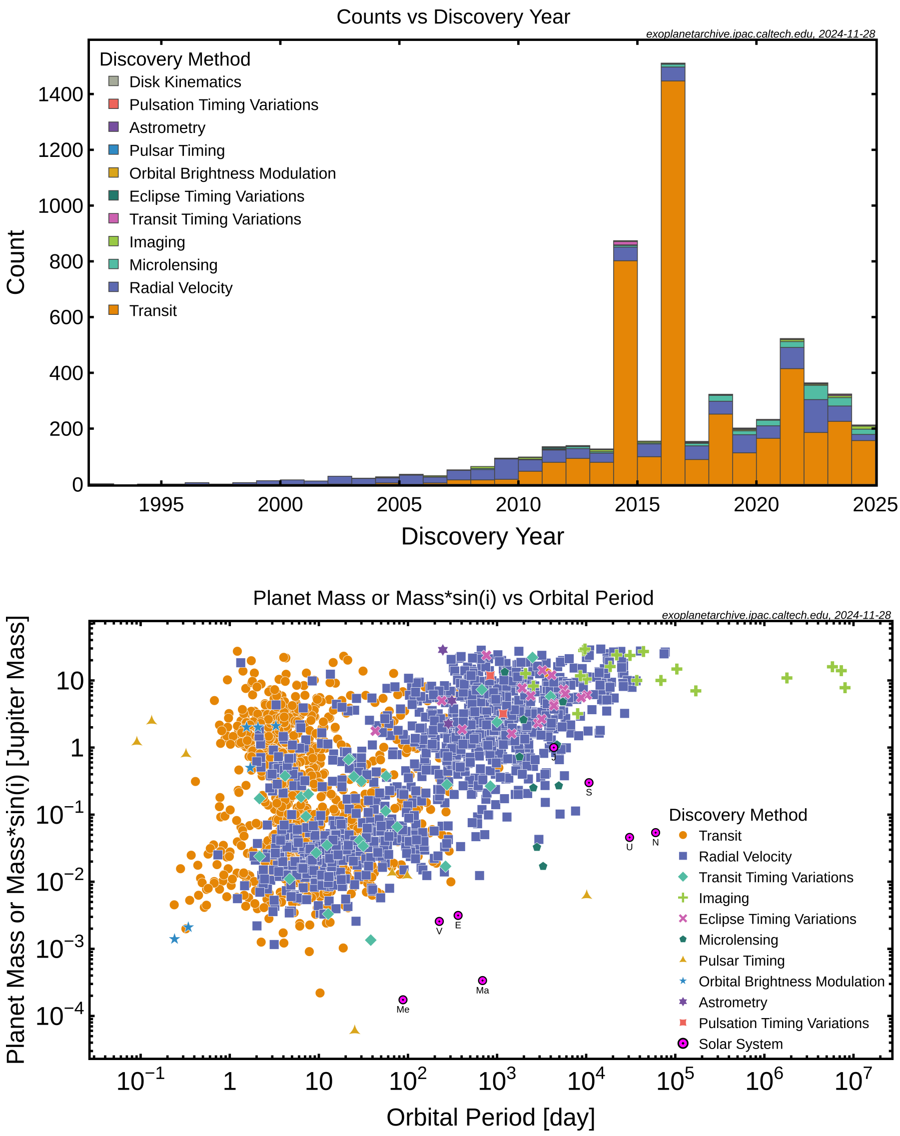

Figure 1.2: Exoplanet archive statistics as of 28th of November 2024: The top panel shows the discovered planet count per year since the first discovery in 1992. The bottom panel plots all discovered exoplanets as a function of mass and orbital period. Figures were made using the pre-generated exoplanet plots available on: `exoplanetarchive.ipac.caltech.edu/exoplanetplots/`

uitous. As detection methods become more sophisticated, pushing the limits of detection out to longer periods and down to smaller masses, the diversity of planetary systems will likely continue to expand, revealing even more of the remarkable variety in the universe.

What exoplanets lack in resolution and possibilities for in situ measurements, they make up for in sheer numbers. As mentioned earlier, we are already up to 5788 confirmed exoplanet detections, enabling the first statistical studies of their properties. One of the more straightforward properties to derive with current detection methods is planet mass, which has been the focus of early studies on occurrence rates (e.g. Wittenmyer et al. 2011a,b). By combining mass and orbital distance as selection criteria, studies have been able to look at the frequency of hot Jupiters (Wright et al. 2012) and the likelihood of finding Earth-like rocky planets in the habitable zone of their star (Dressing & Charbonneau 2015; Bryson et al. 2021).

As the catalogue of discovered exoplanets grows and our knowledge of their properties becomes increasingly detailed, the scope of statistical studies continues to expand. For example, the ongoing ESO SupJup Survey (Regt et al. 2024; Picos et al. 2024; Zhang et al. 2024a; Gandhi et al. 2025) leverages recent advancements in high-resolution spectroscopic instrumentation, such as CRIRES^{+}, to measure $^{12}\mathrm{C}/^{13}\mathrm{C}$ isotope ratios in super-Jupiters. These measurements provide new insights into the formation pathways that can lead to the birth of this class of giant exoplanets.

Among the many unanswered questions in exoplanet science is the bulk composition of rocky planets. These planets are both harder to detect than gas giants and, due to them having a surface, have atmospheres that may or may not be connected to their interior compositions. This makes it challenging to infer their bulk compositions from the composition of their atmospheres. However, there is one type of rocky planet that might hold the key to unravelling this mystery.

1.2.2 Lava Planets

'Hot rocky exoplanets' (HREs), interchangeably called 'lava planets', are rocky planets (typically defined as having a radius $\lesssim 2R_{\oplus}$), that orbit their host star at such a short distance that their day-side surface temperature reaches beyond the melting point of rock ($\gtrsim 1500$ K). Due to the proximity to their host-star, these planets experience strong tidal forces leading to a decrease in their rotation speed down to the point that they become tidally locked. This means that the time it takes for a planet to rotate about itself once is equal to the orbital period of the planet around the star (analogous to the Moon with respect to Earth), effectively creating a permanent day and night side (Barnes 2017; Pierrehumbert & Hammond 2019). Due to this and the extremely high surface temperatures, it is likely that HREs support large lava oceans on their daysides (e.g. Miguel et al. 2011; Demory et al. 2011; Kite et al. 2016). If these planets are able to support a volatile atmosphere or some other effective medium through which to transport heat over to the night-side, they might even be able to support global lava oceans (Zilinskas et al. 2023; Meier et al. 2023).

Observing these planets offers several significant advantages compared to their more distant and colder rocky counterparts. First, they are easier to detect and study. Their short orbital radii result in extremely short orbital periods, sometimes just a few hours. The most extreme example currently known is K2-137 b with an orbital period of only 4.3 hours (Smith et al. 2018). Such short periods greatly increase the likelihood of observing a planet during a transit (when it passes in front of its host star from Earth's perspective) or an eclipse (when it passes behind the star). Additionally, the high surface temperature of these close-in planets results in stronger infrared emission compared to colder planets. This increased emission leads to a more pronounced dip in the system's apparent brightness during eclipses (more on this later).

Another unique advantage of HREs is that their lava oceans may provide a direct interface between the interior and atmosphere of these planets (Léger et al. 2009; Henning et al. 2018; Boukaré et al. 2022). The temperature and chemical composition of the surface lava determine the composition of the vapour released into the atmosphere (Fegley & Cameron 1987; Schaefer & Fegley Jr. 2004; van Buchem et al. 2023; Wolf et al. 2023; Seidler et al. 2024). This rock vapour, in turn, fully or partially governs the atmospheric composition, depending on the abundance of atmospheric volatile elements (such as hydrogen). The composition of the surface lava is likely closely related to the composition of the planet's mantle (Kite et al. 2016; Boukaré et al. 2023; Lichtenberg & Miguel 2024). As a result, analysing the chemical makeup of an HRE's atmosphere could provide valuable insights into the planet's interior composition, information that is difficult to obtain even for solar system planets.

Gaining a novel approach to derive rocky-planet interior compositions, even if limited to a specific population type, would see a wide range of applications. To name but a few, it could help break the degeneracy currently faced when deriving interior compositions based on the mass and radius of a planet (Rogers & Seager 2010), provide insight into the relation between the composition of rocky planets and that of their host star (Bond et al. 2010; Dorn et al. 2017; Putirka & Rarick 2019; Wang et al. 2019; Putirka et al. 2021), and help constrain how differentiation within the interior of rocky planets affects the composition of the crust and mantle (Dyck et al. 2021). Furthermore, there is mounting evidence that most, if not all, rocky bodies such as planetoids, moons, and rocky planets had a lava ocean phase at some point in their evolutionary history (Greenwood et al. 2005; Elkins-Tanton 2012; Hin et al. 2017; Norris & Wood 2017; Schaefer & Elkins-Tanton 2018). Being able to constrain the composition of atmospheric products of a lava ocean may therefore also inform us about what the atmosphere of early Earth may have looked like.

Until recently, the study of HRE atmospheres was mostly confined to the realm of theory. Observational data were limited, with the best results coming from the Spitzer Space Telescope observations of 55-Cnc e[5] (Demory et al. 2016a,b) and K2-141 b (Zieba et al. 2022). However, Spitzer was not specifically designed for the observation of exoplanets, and reducing the data for these objects was notoriously difficult. This has now changed with the advent of JWST.

In the past year, Hu et al. (2024) have found signs pointing toward 55-Cnc e hosting a volatile atmosphere, potentially containing CO. In addition to this work, many currently ongoing studies are working on extracting chemical compositions from HRE atmospheres using JWST data. Although many uncertainties and degeneracies remain, JWST is now bringing us tantalizingly close to being able to characterize the chemistry of HREs. As reduction techniques improve (thanks to a growing understanding of JWST's instruments), as new observations accumulate, and as new telescopes such as Ariel, the Habitable World Observatory, and the Extremely Large Telescope (ELT) join the ranks, it is likely that we can start characterizing the atmospheric composition of these planets in the near future.

To interpret these observations and make progress towards deriving interior compositions from atmospheric data, comprehensive and self-consistent models are essential. In the following section of this introduction, we provide the necessary context for understanding our approach to this challenge.

1.3 How did we do it?

Having gained a better understanding of *why* we are interested in modelling and observing HREs, we will now take a look at *how* HREs are modelled and observed. Although the work presented in this thesis is mainly focussed on the modelling of HREs, we start by explaining how HREs are observed so as to have a clear idea of what model spectra aim to replicate. We then give an overview of the modelling approach that we took.

1.3.1 Observing hot-rocky exoplanets

The majority of star-planet systems are too small to be able to spatially resolve with current telescope resolution, especially if one is observing close in exoplanets such as HREs. Therefore, to observe the flux being emitted from an exoplanet one needs to make use of the geometry of star planet systems.

When observing a star that hosts a planet which is aligned with our line of sight as seen from Earth, the planet appears to pass in front of and behind the star as it moves through its orbit (as illustrated in Figure 1.3). When a planet passes in front of the star, it is called a transit, and when it passes behind the star, it is called an eclipse. These events affect the amount of light that we

[5] Officially also named 'Janssen' in honour of the Dutch spectacle maker Zacharias Janssen.

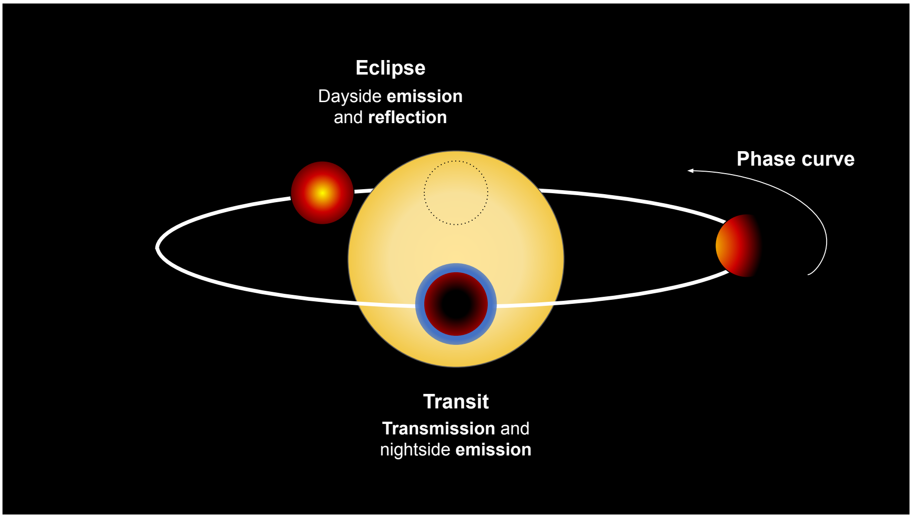

Figure 1.3: Different views of a planet: Depending on its position relative to its host star, we observe different aspects of a planet. HRE spectra produced in Chapters 3 and 5 represent what we may observe during an eclipse.

receive from the system due to the fact that when the planet passes behind the star, no light from the planet reaches us, and when the planet passes in front of the star, less light from the system reaches us. An example of an observation where this is very clearly visible is shown in Figure 1.4, where the light curve of the hot-Jupiter WASP-43 b is shown as observed with JWST (Bell et al. 2023, 2024). When the planet passes behind the star during the eclipse we see a small dip in the light received from the star and when the planet passes in front of the star we see a relatively larger dip in the amount of starlight. Both of these events have different advantages and drawbacks when trying to characterize an exoplanet atmosphere.

Transit spectroscopy, which was used to characterize solar systems bodies before the era of exoplanets (Smith & Hunten 1990), allows for the analysis of light after it has passed through the atmosphere of a planet. Depending on the composition, temperature, and size of the atmosphere, varying amounts of light are let through at different wavelengths, essentially providing a fingerprint of the planet's atmosphere. This has enabled the detection of a wide range of different chemical species such as H_2, H_2O, CO, CO_2, NH_3, Na, Fe, and many more[6].

This approach works best for large gaseous planets that support large atmospheres, such as WASP-43 b. For smaller planets with less prominent atmospheres, however, transit spectroscopy is harder to apply effectively. Heavy

[6] A fairly recent and comprehensive list including all sources is given in Guillot et al. (2022) in Table 4 of their appendix.

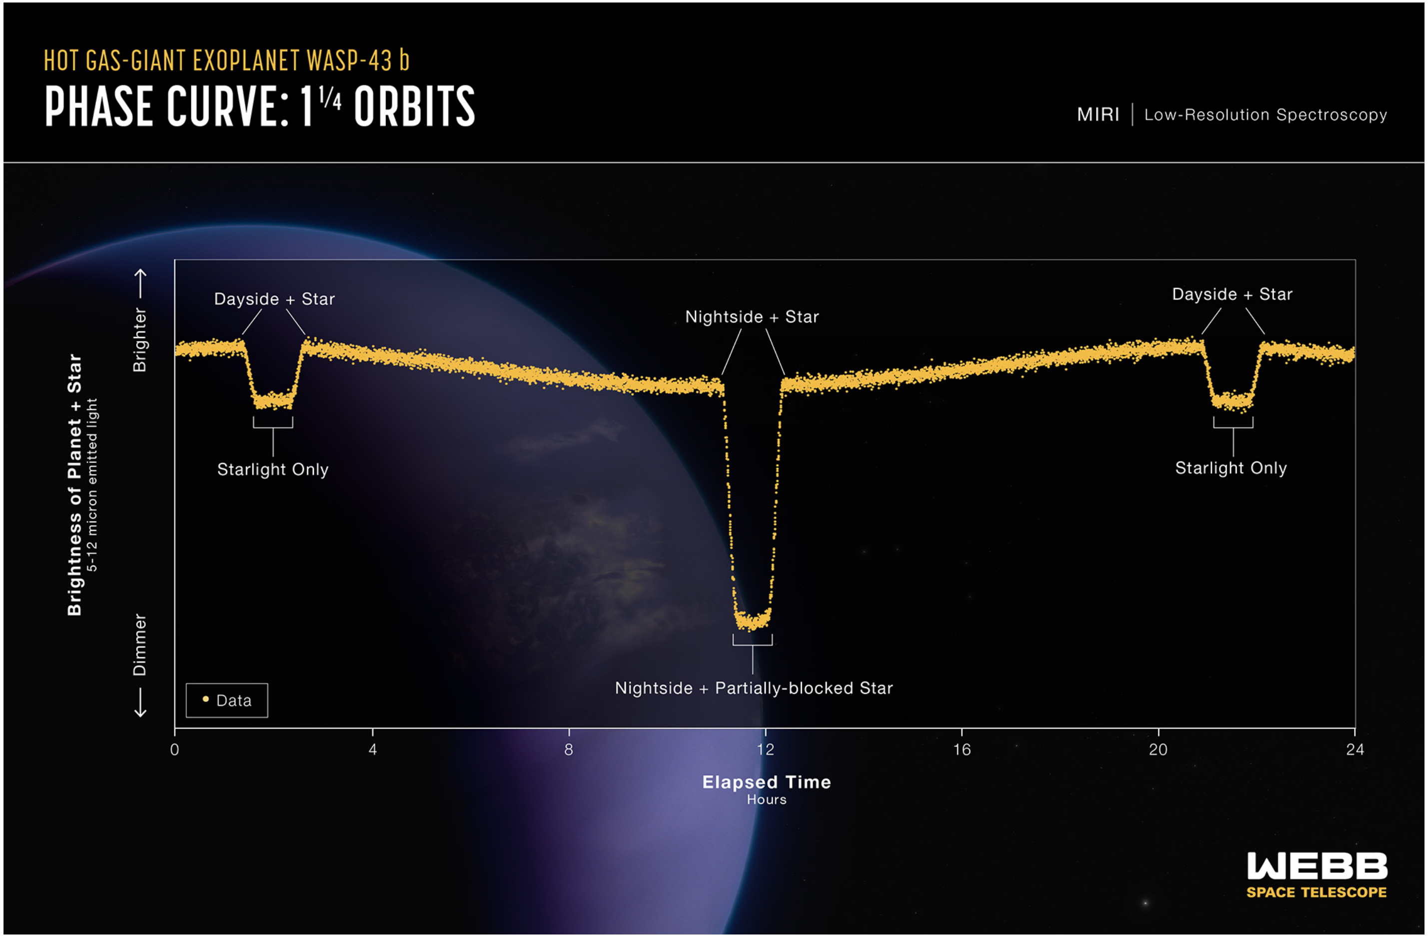

Figure 1.4: Phase Curve of WASP-43b: A clear example of how the flux received from a star changes with a function of time as a planet passes behind it during the eclipse (small dips at the beginning and end) and as the planet passes in front of the star (large dip in the middle). Data published in Bell et al. (2023, 2024). Source: NASA, ESA, CSA, Ralf Crawford (STScI).

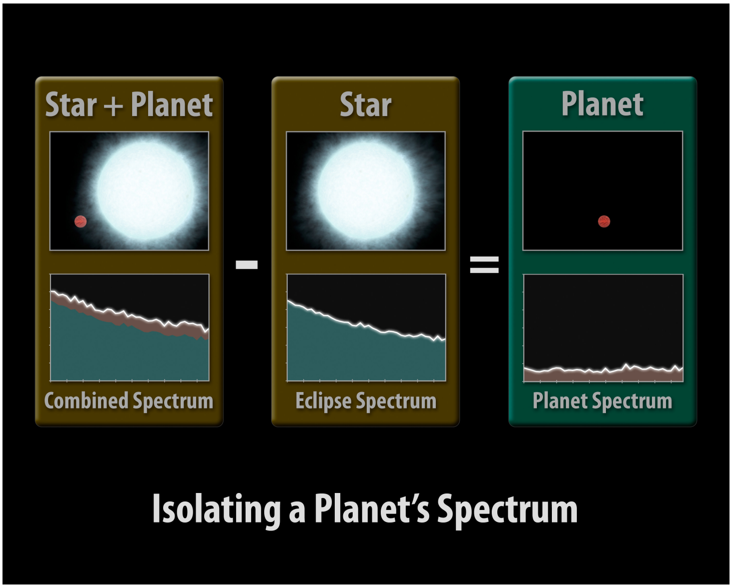

Figure 1.5: Isolating a planet's spectrum: The emission spectrum of an exoplanet can be extracted by subtracting the light received from the star only (during an eclipse) from that of the combined light of both the star and the planet (before and after the eclipse). Source: ESA

lava-vapour rich atmospheres, such as we expect to be present on HREs, could potentially have the additional issues of clouds and hazes potentially blocking light, leading to flat featureless spectra.

That is why for rocky planets (including HREs) one of the most common ways to analyse their atmospheres is by taking eclipse observations. One of the advantages of eclipses is that it allows for the isolation of the light radiated by the planet. As illustrated in Figure 1.5, subtracting the light from the star (gathered during the eclipse, when the planet is behind the star with respect to Earth) from the combined star and planet light, allows for the isolation of the planetary light.

Emission spectra are most easily taken for systems where the ratio of the flux from the planet to the flux from the star is as high as possible. This flux ratio is determined by the relative areas of the star and planet and their respective temperatures. This makes planets orbiting smaller stars from K- down to M-type more desirable targets (as shown in Chapter 3). Another way to increase the ratio of planet to star light is by observing in the infrared, a wavelength range in which planets are relatively bright thanks to thermal radiation. Hot planets, such as HREs, especially so.

Even when leveraging these advantages, extremely high photometric sensitivity is required to be able to discern the thermal emission of an exoplanet

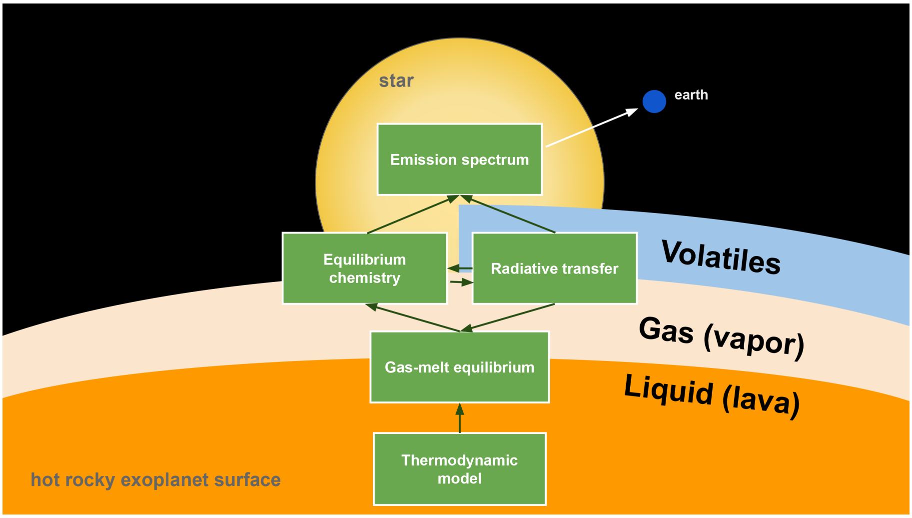

Figure 1.6: Overview of a forward model structure: Used for the self-consistent modelling of emission spectra for HREs.

from the light of its host star (planet to star flux ratios are most commonly denoted in parts per million). Hence, it wasn't until the advent of JWST that infrared emission spectra of HREs have been observed at a sensitivity high enough to start attempting to effectively characterise their atmospheres. This entails that the models of HRE atmospheres can finally be put to the test.

1.3.2 Modelling hot-rocky exoplanets

In Figure 1.6 we show a schematic that broadly outlines the way in which we have structured our models. Starting at the bottom, we use a thermodynamic model to calculate the thermodynamic properties of silicate melts. Using these properties, we then model the lava-atmosphere interaction using the LavAtmos code. The development of this code played a central role in the work done for this thesis, as can be read in Chapters 2 and 4. This code allows us to calculate the composition of rock vapour coming from lava at a given temperature, pressure, melt composition, and (optionally) with the presence of given volatile elements.

Once we know the composition of the vapour, we use this to determine the chemical composition of our atmosphere - including any volatile elements if needed. This information is then passed on to a code that determines the temperature pressure structure of the atmosphere. Finally we used a code that uses all this information to calculate the emission spectrum of the planet. This section is structured in such a way that we follow the model upward, explaining each step along the way.

Interior-atmosphere interactions

Throughout this work, we assume that the size of the lava oceans present on HREs are such that they dominate the atmosphere in terms of reservoir size. This is a common assumption made by most work that concerns itself with lava oceans on HREs (e.g. Schaefer & Fegley Jr. 2004; Zilinskas et al. 2022; Wolf et al. 2023; Seidler et al. 2024). This allows us to treat the interior-atmosphere interactions as one directional, with the lava ocean affecting the atmosphere and not vice versa.

The way in which a lava ocean affects atmospheric compositions is through vaporization reactions through which melt oxide species, which the lava is made of, transition from the liquid to the gas phase. Analogous to water evaporation at room temperature. To understand how we can calculate the partial pressure of different gas species due to this process, we have to understand the following equation derived from the law of mass action[7]:

$$P_{ij} = K_{r_{ij}}(T, P) a_j^{c_{ij}} P_{\mathrm{O_2}}^{d_{ij}} \tag{1.1}$$

P_{ij} is the partial pressure of a vapour species i resulting from the vaporization reaction of melt species j. To calculate this value we need to know the value of the chemical equilibrium constant of reaction ij $K_{r_{ij}}$, the activity of the melt oxide j a_j, and the O_2 partial pressure in the atmosphere P_{O_2}. The exponents $c_{i,j}$ and $d_{i,j}$ are the stoichiometric coefficients which ensure that the reaction is balanced. Taking the vaporization of SiO_2 (liquid) in the melt to SiO (gas) in the atmosphere as an example, equation 1.1 would be re-written as follows:

$$P_{\mathrm{SiO}} = K_r a_{\mathrm{SiO_2}} P_{\mathrm{O_2}}^{-1/2} \tag{1.2}$$

We start by finding the value of the chemical equilibrium constant of this reaction. Although named a constant, this value is dependent on both temperature and pressure. It being named a constant refers to the fact that it relates the concentration of a reactant to the concentration of a product at chemical equilibrium - this relation being constant at a given temperature and pressure. The value of K_r for each reaction ij can be derived from tabulated values of thermochemical properties of chemical species. The most widely used of these are the JANAF-NIST tables (Chase 1998), which is what we use throughout the majority of this work.

The next variable that we need to know is a_{SiO_2}, which is the activity of the melt oxide SiO_2. In essence, this value indicates how much of the SiO_2 in the melt is "available" for taking part in the vaporization reaction. The activity of melt oxide j is given by

$$a_j = y_j X_j \tag{1.3}$$

Where y_j is the so-called Raoultian activity coefficient, and X_j is the mole fraction of the given melt oxide. Using SiO_2 again as an example, if it were an

[7] The full derivation of this equation is given in section 2.2.

ideal system where all of the SiO_2 in the melt was available, y would equal 1 and hence the activity would be equal to the mole fraction X of SiO_2 in the melt. For non-ideal systems, the y may have any value greater than 0.

Due to the complexity of the behaviour of oxide species in melts, there are currently no comprehensive analytical models derived from first principles that are able to accurately calculate the activities of oxide species in complex melts (melts containing more than two different oxides). Instead, models are built based on empirical measurements of melt-oxide activities. In Chapter 2 we detail how we use the thermochemical code MELTS developed by Ghiorso & Sack (1995) to calculate the activity of melt oxides for vaporization reactions - which is what we use for the work done throughout the rest of the chapters as well.

When the work for this thesis was originally started, most studies of HREs assumed that they were too hot to support any volatile elements (such as H, C, N, S, and P) in their atmospheres. However, theoretical models have since shown that HREs could potentially host large amounts of water in their lava oceans (Hirschmann 2012; Lebrun et al. 2013; Dorn & Lichtenberg 2021; Kite & Schaefer 2021) and that volatile atmospheres could be supported even when considering the extremely strong irradiation from their host stars (Herbort et al. 2022; Charnoz et al. 2023; Maurice et al. 2024). This has prompted modelling work on what the emission spectra of HREs with volatile species in their atmospheres could look like (Piette et al. 2023; Zilinskas et al. 2023; Falco et al. 2024). In addition, the recent JWST observations of the HRE 55-Cnc e indicate the possibility of the presence of a significant volatile atmosphere, based on a significantly lower probed temperature than expected as well as strong infrared absorption features (possibly due to CO or CO_2). This amounts to a growing body of work pointing at the importance of including volatile elements when modelling HREs.

The first order approach to including volatile elements in forward models of HREs, is to calculate the vapour composition coming from a melt while assuming no volatile elements are involved and to then add the abundances of the vaporised elements (Si, Al, Ti, Fe, Mg, Ca, Na, K, and depending on the approach also O) to predetermined abundances of volatile elements. This is the approach taken in Piette et al. (2023) and (Zilinskas et al. 2023) and gives a good initial idea of how volatile species in an HRE atmosphere affect the emission spectrum. However, this approach does not take into account what the effect is of volatile elements on the vaporisation process itself.

In Charnoz et al. (2023), it is shown how even low H abundances can have a significant effect on the vaporisation of species from the surface lava ocean, leading to higher abundances of the vaporised elements than previously thought. Based on this work and with a similar approach, we expanded LavAtmos (the vaporisation model from chapter 2) to also include H, C, N, S, and P in the vaporisation reactions, leading to the development of LavAtmos 2.0. This work is covered in detail in chapter 4.

With these vaporisation codes, we are able to calculate the elemental composition of the vapour released into the atmosphere of an HRE based on the surface temperature, pressure, and composition of the surface lava ocean. The next step in building a full model of the atmosphere is to calculate the atmospheric gas chemistry.

Atmospheric chemistry

We make use of vaporisation codes to calculate the composition of an HRE atmosphere at the surface of the lava ocean. Although the output of these codes is in terms of chemical species, we tend to convert this to abundances and pass it on to a dedicated gas-chemistry code to calculate the chemical speciation of the atmosphere. This is because 1. dedicated gas-equilibrium chemistry codes often include a greater number of gas species than our vaporisation and 2. we want to know the speciation throughout the entire atmosphere (at varying temperatures and pressures) and not just at the surface of the planet.

Due to the high temperatures of HREs, it is common to assume chemical equilibrium throughout the entirety of the atmosphere. This allows the problem of chemical speciation to be approached through the method of minimizing the Gibbs free energy of a system (e.g. Denbigh 1955; Aris 1969), an approach through which it is possible to calculate the equilibrium composition of a system without needing to write out every possible equilibrium reaction. To do this, we make use of the code FastChem (Stock et al. 2018, 2022; Kitzmann et al. 2024) which is semi-analytical code which is able to quickly and accurately calculate chemical speciation over a wide range of pressures and temperatures.

Radiative Transfer

Once we have a grid of chemical species for a range of temperatures and pressures, we can use this to build a model of how light propagates through the atmosphere. This is done through the use of 1-D radiative transfer climate models. Initially developed for the modelling of atmospheres of solar-system bodies, these are now also commonly applied to exoplanet atmospheres (Manabe & Strickler 1964; Sudarsky et al. 2003; Fortney et al. 2005). The aim of these models is to solve the propagation of stellar flux through the planetary atmosphere and, in doing so, derive its temperature-pressure structure.

One of the most influential factors in this process is the chemical composition of the atmosphere. Since different species absorb and emit radiation differently depending on their abundance, temperature, and pressure, an accurate representation of the chemical makeup across a broad temperature-pressure (TP) grid is essential. Once the composition is known, opacities of each chemical species dictate how radiation interacts with the atmosphere. The absorption properties of various atomic, ionic, and molecular species vary across wavelengths, making opacity calculations a key ingredient in atmospheric modelling.

Determining these opacities is an entire research field of its own, with dedicated collaborations (such as DACE[8] and ExoMol[9]) for the compiling and refining of spectroscopic line lists based on both experimental measurements and theoretical calculations. Table 5.2 in the appendix of Chapter 5 illustrates the vast number of species required for accurate modelling and the collaborative effort behind acquiring the data necessary to calculate opacities.

Due to the complexity of solving radiative transfer from first principles, an approximation was developed to enable more efficient computations. The two-stream approximation, first introduced by Schuster (1905), simplifies the radiative transfer equation by considering only two directional fluxes: upward from the atmosphere's base and downward from the top. To achieve radiative equilibrium, the atmospheric temperature structure is iteratively adjusted until the net energy flux at each layer is zero, ensuring a balance between incoming and outgoing radiation. In this work, we utilize the HELIOS code (Malik et al. 2017), which implements this approach with additional optimizations for exoplanetary atmospheres.

Once the radiative code converges on a temperature-pressure (TP) profile of the atmosphere, the chemical equilibrium code is used one more time to determine the chemical composition throughout the atmosphere[10], providing us with all of the information that we need to produce an emission spectrum. This last step is done using a different radiative transfer code called petitRADTRANS (Mollière et al. 2019), the output of which we can then compare to real observed emission spectra of HREs.

Even with this simplified 1-D approach, constructing a self-consistent model of an exoplanet atmosphere requires integrating diverse datasets, including chemical abundances, opacity sources, and radiative transfer methods. The complexity of these models highlights the extensive groundwork necessary to accurately simulate (exo)planetary atmospheres and interpret observational data.

1.4 What did we find?

The overarching goal of the work presented in this thesis is to understand how the surface lava oceans of HREs influence their atmospheric composition and how this in turn affects their emission spectra. The forward modelling technique that we use is based on combining a range of different numerical models which each simulate a different aspect of an HRE atmosphere. A large part of this thesis is focussed on improving and expanding the code dealing with one specific aspect of the forward modelling: the vaporisation from a lava ocean surface. As a result of this, the chapters included in this thesis can be broadly

[8] https://dace.unige.ch/opacityDatabase/

[9] https://www.exomol.com/

[10] We also include an additional iteration to ensure that the surface temperature derived from the TP profile is consistent with the temperature used for the vaporisation calculations. See Chapters 3 a 5 for a detailed explanation.

categorized in two types of works: 1. the development of vaporisation codes (Chapters 2 and 4) and 2. the integration of these codes into the greater forward modelling framework to assess their impact on synthetic emission spectra of HREs (Chapters 3 and 5).

In short, we start in Chapter 2 with the development of a melt-vaporisation code (LavAtmos) that assumes that no volatile elements are present. In Chapter 3, we then apply this code within the greater forward modelling framework to understand how varying melt compositions affects HRE temperature-pressure structures and emission spectra. Chapter 4 focusses on expanding the vaporisation model to also include volatile elements in the vaporisation reactions (LavAtmos 2) and on understanding how this affects the abundance of certain melt-vapour species. Finally, Chapter 5 uses this expanded version of the melt-vaporisation code within the full forward modelling framework to investigate what the emission spectra of HREs may look like if they support an atmosphere containing volatile elements. Below is a concise overview of the approach we took for each chapter and the most important findings.

Chapter 2: Developing LavAtmos

The first challenge of this project was to develop a melt-vaporisation mode. When we started there were two melt-vaporisation codes in use in literature. The first of these - called MAGMA - was a code originally developed for the study of lava-vaporisation on solar system bodies (Fegley & Cameron 1987; Schaefer & Fegley Jr. 2004), but was subsequently applied to the study of exoplanets (Schaefer & Fegley 2009; Miguel et al. 2011; Kite et al. 2016). The reasons that an alternative to this code was desirable were that it is not open source, so it had limited availability, and it used a (potentially) outdated model for calculating melt activities[11] called 'ideal mixing of complex components' (IMCC) developed by Hastie et al. (1982b); Hastie & Bonnell (1985, 1986).

The second available code was VapoRock (Wolf et al. 2023). This code is open-source and uses a more comprehensive model for estimating melt activities called MELTS (Ghiorso & Sack 1995; Asimow & Ghiorso 1998; Ghiorso et al. 2002; Gualda et al. 2012; Ghiorso & Gualda 2015). The reason we opted to still develop our own melt-vaporisation code was because VapoRock treats O_2 partial pressure as a user input instead of calculating it self-consistently, which was being done by MAGMA, a feature that we desired as well.

As such, we wrote a new open-source melt-vaporisation code called LavAtmos which combined the (in our opinion) best aspects of the two other existing codes: making use of the MELTS thermochemical model for calculating the activities of the melt species (as used by VapoRock) and using the mass-balance and mass-action constraints for calculating the O_2 partial pressure of the system (as used by MAGMA).

[11] See section 1.3.2 for an explanation of the role of melt activities.

Chapter 3: Applying LavAtmos to different melt compositions

Having developed a melt-vaporisation code, the next logical step was to apply it. One of the key questions of HREs is whether or not variation in the composition of surface lava oceans could have an observable effect on their emission spectra. As of yet, we have little information on what compositions we should expect for rocky exoplanets and so the majority of work done on HREs assumes that the composition of the surface lava ocean is equivalent to that of the bulk silicate earth (BSE) (Palme & O'Neill 2003).

Using this standard composition as a starting point, we investigated how deviating from this composition for each of the main melt oxides would affect the emission spectra of HREs. We found that HRE emission spectra are sensitive to the amounts of TiO_2, SiO_2, and to a lesser extent Na_2O and K_2O in their surface lava oceans. The resulting changes in the emission spectra due to varying TiO_2 and SiO_2 melt abundances may potentially be large enough to be detectable on the some of the brightest HREs given sufficient JWST observing time.

Chapter 4: Developing LavAtmos 2.0

A key question to answer when modelling melt-vaporisation is how the process is affected by the presence of volatile elements such as H, C, N, S, and P. In Charnoz et al. (2023) it was shown how including H in vaporisation reactions could have a significant impact on the abundances of melt-vapour species. We realised that with some work we could expand the capabilities of LavAtmos to also include volatile elements using a similar approach. In addition tos adding H, we were also able to include C, N, S, and P.

Unlike the first version of LavAtmos, which was an iteration upon other codes with output that was generally well understood, LavAtmos 2 required a much more in-depth exploration of how including volatile elements could affect the chemistry of a melt-vapour system. We found that the increase in vapour-species abundance due to the presence of volatile elements varies depending on the abundance of the different volatile elements, with elements that are prone to forming oxidised molecules such as C having a stronger effect on the vaporisation than N for example.

Since the O abundance in the atmosphere is also dictated by the vaporisation reactions and the presence of volatile elements increases the abundance of vaporised species, we find that this leads to a more oxidised atmospheric composition. As such, we find that a low C/O ratio in an HRE atmosphere could potentially serve as an indication that a surface lava ocean is present.

Chapter 5: Applying LavAtmos 2 to volatile HRE atmospheres

After developing LavAtmos 2, we integrated it into the full forward model, allowing for a more self-consistent approach to modelling the thermal emission spectra of volatile bearing atmospheres of HREs.

We found that the presence of both a surface lava ocean and atmospheric volatiles leads to very strong H^- continuum opacity. Melt-vapour from the lava ocean enriches the atmosphere with metals, which, when ionized, provides an abundance of free-electrons. Adding but a small amount of H to the atmosphere ($\simeq$ 1e-6 bar) yields a H^- continuum that is so strong that the spectral features of other major atmospheric species are no longer detectable.

Additionally, we found that H^- continuum opacity shifts the photosphere of an HRE to higher altitudes compared to an atmosphere without H^- continuum. In the inverted atmospheres typical of most HREs, this means the photosphere radiates from a much hotter region, making the emission spectrum significantly brighter.

Based on these findings, we conclude that if an H^- continuum is detected in an HRE, it could potentially point to the presence of a surface lava ocean.

Future work

It lies in the nature of scientific research to never truly be done. There is always something more to look into, one last thing to check, or a different direction from which to approach an issue. This is especially true in the field of exoplanet research. As our research methods improve, more and more different scientific disciplines need to be involved to gain a comprehensive understanding of the objects we study. Taking this single thesis as an example, we need to understand: geochemistry, gas-chemistry, thermodynamics and radiative transfer, each of which are fully fledged research fields of their own. However, to quote a professor from our institute "The best thesis, is a finished thesis".

Although the work for this thesis may be done, the work in this field is anything but. We are only at the very start of being able to characterise exoplanet atmospheres. Besides the promising results already coming in from JWST, there is a whole suite of telescopes under development that will push the sensitivity of our observational data to new limits. Increasingly information rich data will also force the theoretical aspects of exoplanet research to be greatly expanded. Something which is also already happening with the development of increasingly efficient numerical methods, ever more computation power, and the promise of the effective applications of artificial intelligence.

I am very excited to see what the future holds and grateful that I've had the chance to play a small part in furthering our understanding of these neighbouring yet terribly far off worlds.