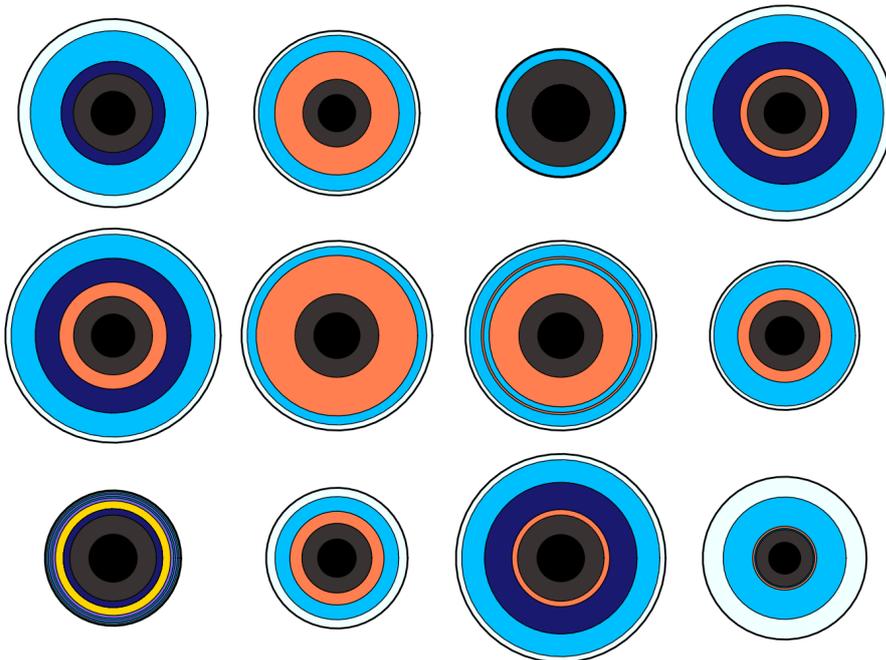


Scott W. Thomas

Internal and atmospheric structures of heated watery super-Earths



This dissertation is submitted for the degree of Doctor of Philosophy.

TRINITY COLLEGE · UNIVERSITY OF CAMBRIDGE · DECEMBER 2016

Internal and atmospheric
structures of heated watery
super-Earths

Copyright © 2016 S. W. Thomas



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

This dissertation is typeset with the `tufte-book` \LaTeX class. It uses Markdown source, the `Pandoc` document converter, and a custom template. The source code, including Jupyter notebooks containing all my figures, is available at www.github.com/swt30/thesis.

Title image: Water can exist in any of several different phases. As well as its familiar vapour, liquid and solid forms, it becomes a supercritical fluid, exotic ice or plasma when heated or compressed. These colourful circles represent a variety of potential structures for planets that are covered in thick layers of water. Essentially, they are models of giant snowballs.

First submitted December 2016. Final version deposited April 2019.

*For those who heard my dreams
and said, "Go for it."*

Internal and atmospheric structures of heated watery super-Earths

Scott W. Thomas

Astronomers are discovering more and more super-Earths, planets around other stars whose sizes and masses lie somewhere between those of Earth and Neptune. We would like constraints on their composition to investigate whether they are more similar to rocky Earth or gaseous Neptune. To do this we need numerical models of their interiors. Such models often exclude any thermal effects, a choice justified by noting that a heated rocky planet expands by only a small amount. But this is not necessarily true for planets with thick oceans or watery atmospheres. Water has a rich and interesting thermal behaviour: at high pressure and temperature it can be in any of several exotic plasma and ice phases. Planets with thick water layers, known as waterworlds, cannot therefore be accurately represented by models that treat them as cold spheres. But understanding how waterworlds vary in size and structure is important as we seek to interpret new observations of super-Earths.

I developed temperature-dependent structure models of waterworlds, treating both the interior structure and the atmosphere and including both internal and external heating. In doing so, I synthesized an improved equation of state for water to better capture how it behaves when heated or pressurised. Using these models, I show the following: heat can significantly affect a watery planet's size and structure; these planets can have large and diffuse yet opaque atmospheres; and a planet can have a hot extended steam atmosphere even if only moderately heated from the inside. My models are simpler than those based on energy transfer codes, yet are fast to evaluate and still capture thermal behaviour trends appropriately. I therefore suggest that they would be ideally suited to use in statistical models of planetary systems. I also explore how a planet might change size if it migrates or exists in an elliptical orbit, consider the astrobiological implications of heating a watery planet, and present the results of applying these models to a recently discovered potential waterworld.

Contents

<i>1</i>	<i>Introduction</i>	<i>13</i>
<i>1.1</i>	<i>Background</i>	<i>14</i>
<i>1.2</i>	<i>Internal structure models</i>	<i>22</i>
<i>1.3</i>	<i>Heated watery super-Earths</i>	<i>30</i>
<i>1.4</i>	<i>This dissertation</i>	<i>35</i>
<i>2</i>	<i>An improved water equation of state</i>	<i>37</i>
<i>2.1</i>	<i>About the EOS</i>	<i>37</i>
<i>2.2</i>	<i>How I collected and synthesized the data</i>	<i>44</i>
<i>2.3</i>	<i>The final equation of state</i>	<i>55</i>
<i>2.4</i>	<i>Making the data freely available</i>	<i>57</i>
<i>3</i>	<i>Watery planet interiors</i>	<i>59</i>
<i>3.1</i>	<i>Theory</i>	<i>59</i>
<i>3.2</i>	<i>Models</i>	<i>64</i>
<i>3.3</i>	<i>Results</i>	<i>71</i>
<i>3.4</i>	<i>Discussion</i>	<i>77</i>
<i>4</i>	<i>Heating and the atmosphere</i>	<i>81</i>
<i>4.1</i>	<i>Heated planets</i>	<i>81</i>
<i>4.2</i>	<i>The atmospheric treatment</i>	<i>90</i>

4.3	<i>Results</i>	95
4.4	<i>Summary and discussion</i>	105
5	<i>Phase structure and migration</i>	109
5.1	<i>The phase structure of a planet</i>	110
5.2	<i>Structural changes of migrating planets</i>	117
5.3	<i>Results</i>	118
5.4	<i>Discussion</i>	123
6	<i>A water-rich super-Earth?</i>	129
6.1	<i>Gliese 1132 b</i>	129
6.2	<i>A potential water detection</i>	130
6.3	<i>Results</i>	131
7	<i>Summary and conclusion</i>	137
	<i>Bibliography</i>	143

Preface

This dissertation is the result of my work at the Institute of Astronomy between October 2013 and December 2016 under the guidance of Nikku Madhusudhan. Though all the text is mine, some of it is based on work for which he, Ian Parry and Christopher Tout have provided advice or editing. All external sources of data are referenced in the text; all figures are my own except where indicated; and any errors or omissions are mine.

Chapters 2 and 3 were published as “In hot water: effects of temperature-dependent interiors on the radii of water-rich super-Earths”.¹ Chapter 6 is based on work carried out jointly with Oliver Shorttle. I contributed the interior structure models, he used them to produce the figures shown, and I wrote the chapter.

¹ Thomas & Madhusudhan (2016)

I HEREBY DECLARE that this dissertation entitled “Internal and atmospheric structures of heated watery super-Earths” is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in this Preface and specified in the text. It is not substantially the same as any that I have submitted or is being concurrently submitted for a degree or diploma or other qualification at the University of Cambridge or any other University or similar institution.

I further state that no substantial part of my dissertation has already been submitted or is being concurrently submitted for any such degree, diploma or other qualification at the University of Cambridge or any other University or similar institution.

This dissertation does not exceed 60,000 words.

Hello, babies. Welcome to Earth. It's hot in the summer and cold in the winter. It's round and wet and crowded. At the outside, babies, you've got about a hundred years here. There's only one rule that I know of, babies—God damn it, you've got to be kind.

Kurt Vonnegut, “God Bless You, Mr. Rosewater”

*Within your lifetime will, perhaps,
As souvenirs from distant suns
Be carried back to earth some maps
Of planets and you'll find that one's
So hard to color that you've got
To use five crayons. Maybe, not.*

Marlow Sholander, “Maybe”

Sometimes I think the surest sign that intelligent life exists elsewhere in the universe is that none of it has tried to contact us.

Bill Watterson, “Calvin and Hobbes”

I Introduction

Not many discoveries have ever captured the imagination of scientists and non-scientists alike as much as exoplanets. Until 1992, we had only the ground under our feet and a select few lights in our sky to guide our understanding of planetary systems. Now we are confronted with a plethora of planets: almost every star is expected to host one, and it seems that papers and press releases trumpet a new detection every week.

Astronomers who work in the young field of exoplanet astronomy are blessed¹ with a rapidly increasing amount of data. In the twenty-odd years since the first exoplanets were found, ground- and space-based missions have provided thousands of exoplanetary detections. This prodigious increase in the number of known extrasolar systems owes its existence to improved light-gathering ability, instrumentation, and analysis techniques. From minute fluctuations in light and tiny velocity wobbles, a whole industry has sprung up devoted to simply finding exoplanets—a task which is certainly not easy on its own.

¹ Or is it cursed?

There is another task which is parallel to the push to find exoplanets: the community has turned its gaze towards understanding the ones that it has already found. With huge amounts of telescope time being devoted to observing them and a flurry of new missions and instruments on the way, it has quickly become a cliché in the field to say that we are moving from an era of exoplanet *detection* to one of *characterisation*. But this is only a cliché because it is absolutely true.

Image removed for
copyright reasons

This field is seeing increasing interest from both scientists and non-scientists. And why wouldn't it? New discoveries abound: huge Earths! Planets in the habitable zone! Lava planets, diamond planets, all sorts of sizes and masses! Exoplanet characterisation has increased public interest in astronomy; renewed interest in complementary fields such as stellar activity modelling; brought astronomers, planetary scientists, geologists and atmospheric scientists closer; and interested more people in big questions about our place in the universe.

xkcd.com/1371

As with any field of astronomy, studying exoplanets depends on reliable modelling and correctly understanding physical processes just as much as it depends on observational techniques. What can we understand about exoplanets from the information we have? How much can we trust our models? And how useful can these models actually be? In this chapter I take scope of recent work in a field that has only just seen its twentieth birthday, leading to the motivation for the primary focus of the dissertation: models of watery super-Earths.

1.1 Background

Why do exoplanets appeal to so many people? Once we made the connection between our Sun and the other stars we see, finding planets around them was never going to be a surprise. Yet since the discovery of the first exoplanets using pulsar timing² and radial velocity,³ the field has become very popular indeed.

Many authors would claim that the ultimate goal of exoplanetary research is to detect the bio-signatures that would indicate a habitable planet.⁴ Others are interested in characterising exoplanets for the insights this may offer us about our own solar system and its place in our universe. For example, [Raymond et al. \(2013\)](#) notes that “planet formation models lag behind observations”, so observing exoplanetary systems provides an opportunity to expand and test our formation models. Because the Earth is just a planet, the reasoning goes, any model which correctly explains how planets form will tell us more about our origin as well.⁵

To do this, a wide range of skills are needed. Perhaps this is why the field is so popular: it demands clever observations with new technology, improved data reduction techniques, and a solid theoretical background. So in order to know what problems are worth pursuing, it will help us to briefly review what is currently achievable in exoplanet science.

1.1.1 The field is driven by observations

Exoplanetary science, like most fields of astronomy, is driven by observations. This should come as no surprise and it is a point I will not dwell on. And although much of the formalism for understanding exoplanets initially carried over from our understanding of binary stars, it was the detection of the first extrasolar planets that drove theoretical work to improve detection techniques: better transit modelling, dynamical studies

² Wolszczan & Frail (1992)

³ Mayor & Queloz (1995)

⁴ For example, see the review of Seager & Deming (2010).

⁵ See, for example, Mordasini et al. (2015).

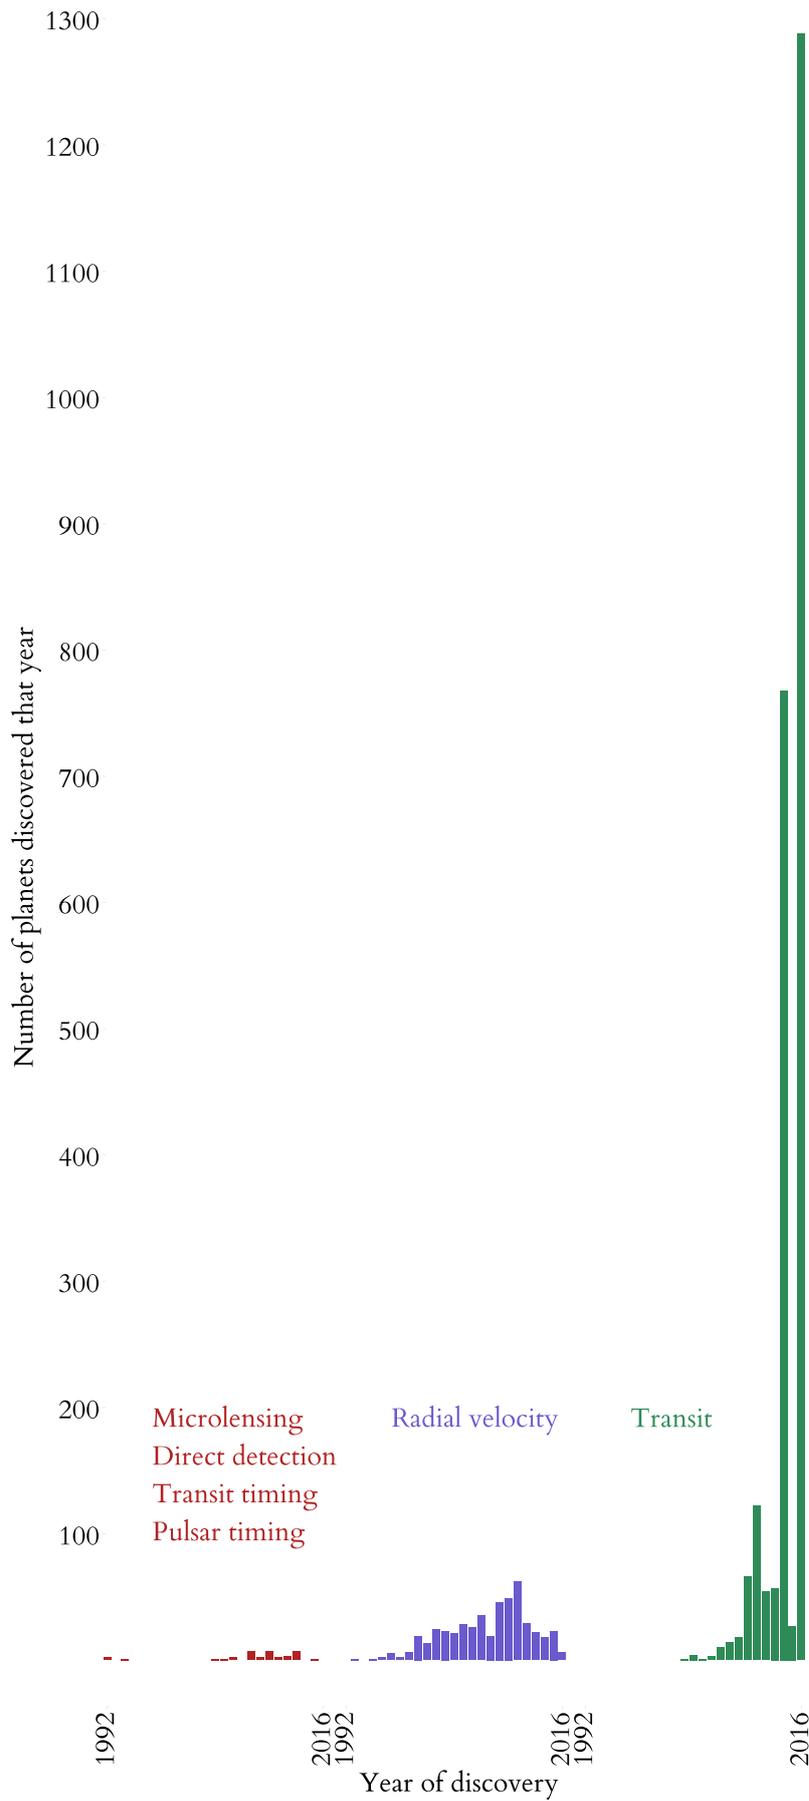


Figure 1.1: Several methods for finding exoplanets exist. Comparatively few exoplanets are found using microlensing, direct imaging or timing methods. The radial velocity method produced many early discoveries of hot Jupiters. But thousands of planets have been found by the transit method, largely thanks to the Kepler mission's data releases in 2014 and 2016. — Data from www.exoplanets.org

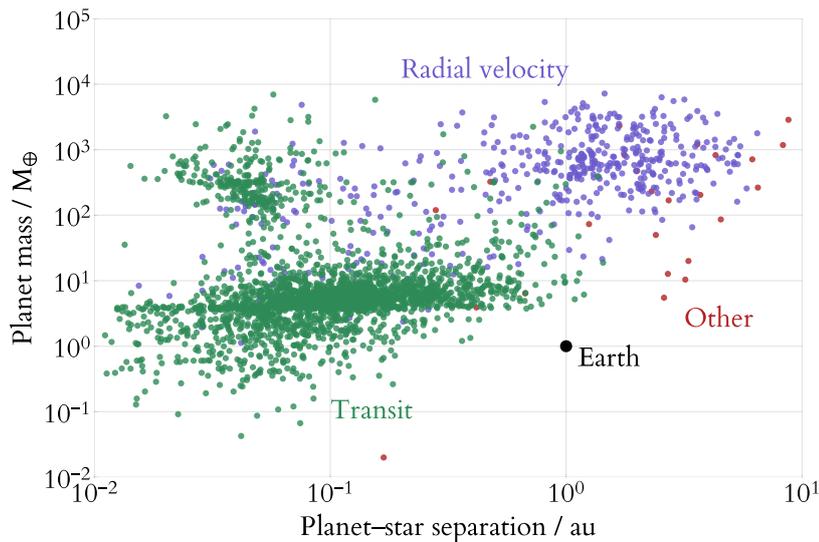


Figure 1.2: Different detection techniques find planets with different properties. Radial velocity planets are often more massive and farther from the star than transiting planets. This is because different methods are biased in different ways: short-period planets are favoured by the transit method while massive planets are favoured by the radial velocity method. — Data from www.exoplanets.org

of systems, and instrumental design that is allowing improvements like high-contrast direct imaging.

In particular, observations of exoplanets have raised a host of questions, most of which we could not have predicted based on our own solar system. These questions span atmospheric science, chemistry, geology, and solar system dynamics. Why do we see so many Jupiter-mass planets with inflated radii?⁶ Might we expect to see systems with binary planets, where two planets have captured each other through tidal interactions?⁷ How do hot Jupiters—gas giants orbiting very close to their host star—form and end up where they are?⁸ How do the systems we see fit with our understanding of planet formation, evolution, and dynamics?⁹ What are the internal structures of these planets like?¹⁰ And finally, what do we need to do to find an Earth-like planet, how will we know when we've done so... and how many might we expect to find?¹¹

1.1.2 Detection methods

A number of different techniques are used to detect exoplanets (figures 1.1 and 1.2). After claims in the early 90s of planet detections around pulsars,¹² the first extrasolar planet around a Sun-like star was found using the radial velocity method described below.¹³ However, in the past few years, the transit method has produced the largest number of new discoveries. Here I summarise different methods of exoplanet detection and the observable information that can be gained from each. This will be important shortly when we consider what directions the field is moving based on these available data.

⁶ Baraffe et al. (2010); see also figure 1.7

⁷ Ochiai et al. (2014)

⁸ Seager & Deming (2010)

⁹ Mordasini et al. (2012a)

¹⁰ Baraffe et al. (2014)

¹¹ Petigura et al. (2013)

¹² Wolszczan & Frail (1992)

¹³ Mayor & Queloz (1995)

THE RADIAL VELOCITY (RV) or Doppler method entails using spectral Doppler shifts to measure the recoil motion of a star as a planet orbits it. The formalism for finding planets in this way has carried over from the older field of binary star research, and the information that can be gained about a star–planet system is very similar. Using Doppler measurements of the recoil motion of the star, we can calculate the *mass function*,¹⁴ which relates the planet’s mass to the star’s mass.

The RV method has the advantage that it works with a wide range of inclinations, not just those near $i = 90^\circ$. This is counterbalanced by the fact that there is an inherent degeneracy between inclination and planetary mass in the $M_p \sin i$ term of the mass function. It is therefore only possible to measure the minimum mass unless some constraint can be placed on i .

The Doppler method is also useful for detecting multiple planet systems. Once the signal of one planet has been found, it can be removed and the system re-analysed to search for more planets. In the case of certain large, bright planets, the emission lines from the planet itself may be used to directly measure the planet’s motion and obtain its mass.¹⁵

Although this method is distance-independent and easily detects large planets close to their host stars, it is not as suited for large surveys because a spectroscopic observation is required. A number of sources of contamination—stellar and magnetic activity, stellar rotation, telluric lines, moonlight, faint companions—must also be accounted for.¹⁶

As can be seen in figure 1.1, the Doppler method was responsible for the majority of exoplanet detections until 2010. However, the advent of a new detection method resulted in an influx of new planets. This method was the transit method.

THE TRANSIT METHOD—where the planet moves between us and its host star (figure 1.3)—is the most intuitive way of finding exoplanets. I will not elaborate on the mechanics of a transiting system except to point out that the probability of observing a transit and the expected decrease in the observed flux are both very small.¹⁷ The greatest gains in the number of known transiting planets have therefore come from systematic surveys of many stars. These surveys are capable of yielding planets on the order of $1 R_\oplus$ in size.¹⁸ This lower limit is constantly being pushed downwards. For example, Kepler-37b was discovered by the transit method; later observations and asteroseismological analysis suggest that its radius is around $0.3 R_\oplus$.¹⁹

¹⁴ The mass function is defined as

$$\frac{(M_p \sin i)^3}{(M_p + M_\star)^2}. \quad (1.1)$$

Here, M_p is the planet’s mass, M_\star is the star’s mass, and i is the orbital inclination. (Wright & Gaudi, 2013)

¹⁵ Snellen et al. (2010); Rodler et al. (2012)

¹⁶ Seager (2010); Haswell (2010)

¹⁷ The transit probability for a circular orbit is approximately R_\star/a , where R_\star is the stellar radius and a is the orbital separation. Ground-based surveys are able to detect flux decreases on the order of less than 1% while space-based surveys can push this to at least 10^{-4} .

¹⁸ Perryman (2011), p. 103

¹⁹ Barclay et al. (2013)

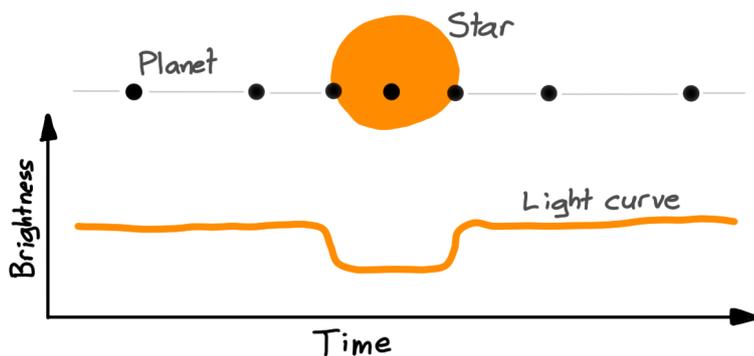


Figure 1.3: In this schematic of a primary transit, we see a dip in the light curve of the star as the planet passes in front of it. Secondary transits, where reflected light from the planet is instead blocked by the star, are smaller.

The transit method also provides a significant amount of information about the transiting planet. From the observables of period, ingress and egress time, equatorial crossing time, and flux depth, the transit method provides an immediate planet–star radius ratio. When combined with radial velocity measurements, the mass and radius of the star and planet can be measured. There is a one–parameter degeneracy,²⁰ which can be broken if more is known about the host star through parallax, spectroscopy, or theoretical mass–radius constraints. Transit information may also inform estimates of planetary surface gravity²¹ and host star density.²²

However, there are several sources of uncertainty. In addition to the expected systematic and random errors in photometry (instrumentation, telluric variability, observing conditions, and stellar properties like surface brightness variations), transit surveys have a high false positive rate.²³ This is because eclipsing binaries, brown dwarf companions, background eclipsing binaries and background exoplanet transits can all produce similar photometric signals.

To handle this possibility, the detections must be validated. A detection can be considered good if a wide variety of possible alternate events are also modelled and the probability of each is deemed low enough. Follow-up observations using other methods are key to rejecting false positives and further constraining the system. For example, the validation technique used by the Kepler team is called `BLENDER`²⁴ and makes use of such observations; these could include high-resolution spectroscopy, infrared photometry, or speckle/adaptive optics imaging to constrain or rule out the possibility of a multiple stellar system. Finally, radial velocity measurements can probe the transiting object’s mass, and can confirm that it is comparable with that of a planet.

²⁰ Wright & Gaudi (2013)

²¹ Southworth et al. (2007)

²² Wright & Gaudi (2013)

²³ Cameron (2012)

²⁴ Fressin et al. (2011)

A NUMBER OF OTHER METHODS for detecting planets are also listed in figure 1.1. The first of these uses timing information, a method which resulted in some of the earliest exoplanet findings.²⁵ This was possible because the exoplanets in question orbit the millisecond pulsar PSR1257+12. The precise period of pulsars means that the gravitational effect of a planet is measurable, even for planets down to $0.1 M_{\oplus}$. “Timing” can also be used in another sense: if a transiting exoplanet is known, variations in its time of transit can indicate an unseen planet elsewhere in the system. Pulsar timing can be used to place a lower limit on the planet’s mass and estimate the semi-major axis of its orbit. Transit timing variations provide the $M \sin i$ of the unseen planet but little else.

²⁵ Wolszczan & Frail (1992)

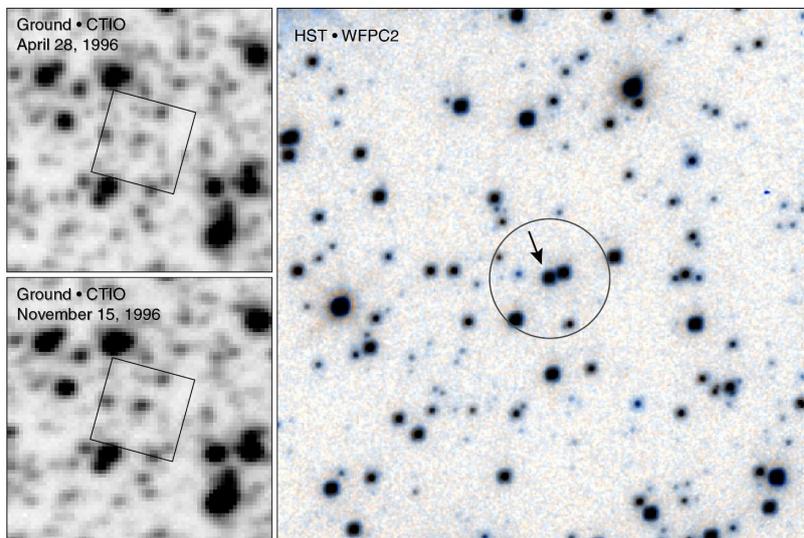


Figure 1.4: Microlensing can be used to find planets. In a microlensing event, a background source star (top left) brightens for a short time as its light is focused on us (bottom left). Here, Hubble’s Wide Field and Planetary Camera has been used to follow up the observation to resolve the source star (right). Planets may be detected around the lensing star if the additional effect of their gravitational field changes the shape of the light curve. The effect can be relatively large if the light path happens to cross a *caustic*, a discontinuity or rapid change in the lensing pattern. — NASA HST/WFPC2

Planets may also be found during microlensing events (figure 1.4). If a planet lies in the right region during the lensing event, the planet’s gravity will further perturb the incoming light and cause a detectable deviation from the single-lens case. In most cases, microlensing detections do not provide much additional information. The single lens event itself provides a weak constraint on the lens mass, and the addition of a planet gives a mass ratio and a projected separation. By themselves these provide little constraint on the orbital elements of the system.²⁶ Extra observations may not be useful because the lensing star is typically not visible and microlensing events are extremely rare. But in the past few years some authors have used higher-order effects to improve the amount of information obtainable from single microlensing events. For example, Bennett et al. (2010) claim strong constraints on 4 of the 6

²⁶ Wright & Gaudi (2013)

orbital elements.

The final method I mention here is direct detection. It is possible to spatially resolve an exoplanet and its host star if the planet is large, bright, and orbits sufficiently far from the star. A number of instruments like the Gemini Planet Imager²⁷, Project 1640²⁸ and SPHERE²⁹ are currently active and attempting to directly image exoplanets. This method is sensitive to young massive self-luminous planets in wide orbits. About twenty such exoplanets have already been found.³⁰ In this case, it may also be possible to take a direct spectrum of the planet.³¹ Others have suggested the use of external occulters, also known as “star shades”, which would fly between a telescope and its target to occult the light from the star.³²

²⁷ Macintosh et al. (2014)

²⁸ Crepp et al. (2011)

²⁹ Beuzit et al. (2008)

³⁰ Fischer et al. (2014)

³¹ Oppenheimer et al. (2013)

³² Soummer et al. (2009)

1.1.3 *Complementary information is required for characterisation*

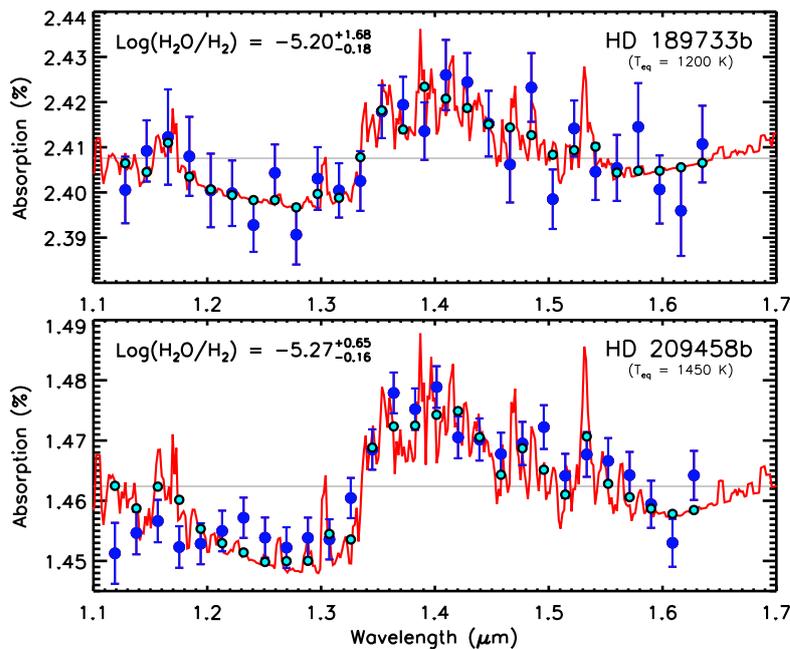
All these methods provide us with different information, and none of them is perfect. In many cases, the information gained from one of these methods is only enough to identify the presence of the planet. But we saw above that this is often not enough—Kepler requires a statistical validation method due to the possibility of contamination from events that look similar to a transit. As well as demonstrating the relative difficulty of identifying clear transit events in data that can often be quite noisy, this also highlights the difference between *validating* and *observationally confirming* the presence of exoplanets. The second is more direct, but requires a radial velocity or transit timing variation measurement. As Fressin et al. (2011) note, “many of Kepler’s smaller candidates will receive a similar validation treatment since dynamical confirmation may be difficult or impractical with the sensitivity of current instrumentation.” It therefore became apparent as large scale surveys were being constructed that clever strategies for validation were needed.³³

³³ For example, see Tingley & Sackett (2005).

Identifying the presence of a planet is one thing, but characterising it is quite another, and this is where the field has been moving. A primary driver of this work is the quest to find a habitable planet, but the definition of habitability is fraught with difficulty.³⁴ Do we accept any Earth-like planet in the habitable zone of any star—that is, the region where we might expect liquid water to exist on the surface? Do we restrict ourselves to Earth-like planets around Sun-like stars? What observational signatures might we expect from a habitable planet? Here, we begin to see a strong overlap with other areas of science such

³⁴ Gaidos et al. (2005)

as atmospheric physics, planetary science, geology, and even biology. In order to understand more about the planets themselves, we turn to other methods to allow us to retrieve more information. The key development in this area has been the use of transmission and emission spectroscopy (figure 1.5), which attempt to separate the star's light and the contribution from the planet.³⁵



³⁵ Charbonneau et al. (2002); Charbonneau et al. (2005); Deming et al. (2005)
 Figure 1.5: Model transmission spectra (red lines) and observations (blue points) for two hot Jupiters. We are increasingly able to measure high-resolution spectra like these for exoplanets, and this allows us to use models of their atmospheres to constrain their atmospheric chemical compositions. — From Madhusudhan et al. (2014b); data from Deming et al. (2013) and McCullough et al. (2014)

1.1.4 Observations motivate models

Recently, much effort has gone into finding ways to retrieve more information about specific exoplanets. The increasing possibilities to directly image or spectrally probe these planets mean that we have seen huge interest in modelling their atmospheres in an attempt to determine likely compositions.³⁶ Not only this, but the need to distinguish the star's light from the planet's light also results in a need to better understand the stellar spectrum. And stellar variability (figure 1.6) becomes important too when attempting to discover or characterise smaller exoplanets; it is perhaps even more critical when dealing with stellar noise in radial velocity measurements.³⁷ In short, good models have never been in more demand.

When I began writing this chapter three years ago, I commented that I was of mixed opinion on the overall success of atmospheric modelling because of the poor quality of available spectra. Today, such

³⁶ Madhusudhan et al. (2014a)

³⁷ Wright & Gaudi (2013)

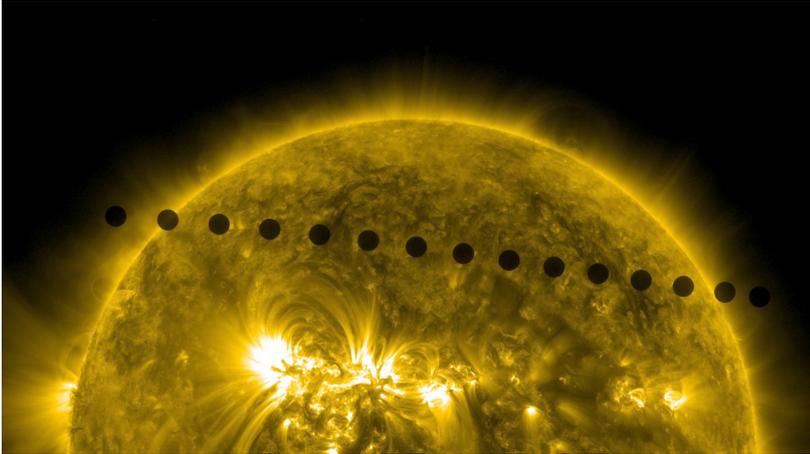


Figure 1.6: In an image of the 2012 transit of Venus, the planet is opaque at this wavelength. Stellar surface activity, clearly visible in this image, is a key complication when disentangling information from a planet and star. — NASA SDO/AIA

pessimism is entirely unjustified. High-quality spectra are now available³⁸ and there has been a tremendous advancement in the detection of molecular species. Giant planets are now routinely observed in transmission to establish the chemical compositions of their atmospheres, and these detections are just beginning to happen for super-Earths as well.³⁹ So, with atmospheric modelling and observation an established and thriving field of research, what next?

³⁸ See for example figure 1.5.

³⁹ e.g. Demory et al. (2016); Southworth et al. (2017)

1.2 *Internal structure models*

As we have seen, a relatively large number of methods are available to detect exoplanets, many of which are well-suited to large surveys. I have discussed the ways in which planets are found, and briefly reviewed the efforts to follow them up and understand their atmospheres. I now turn to their interiors.

Exoplanetary interior modelling is generally based on our knowledge of one set of planets: our Solar System. However, a highlight over the past two decades has been that the planets we have discovered are unexpectedly diverse. This can be seen in a simple *mass–radius diagram*, where the mass and radius of discovered planets are plotted against each other (figure 1.7). This unexpected variety of planets leads us to consider the question: can we come up with self-consistent models to explain this diversity? And what other observables might we be able to tap into to provide further constraints? In this section I will briefly introduce the attempts to do so.

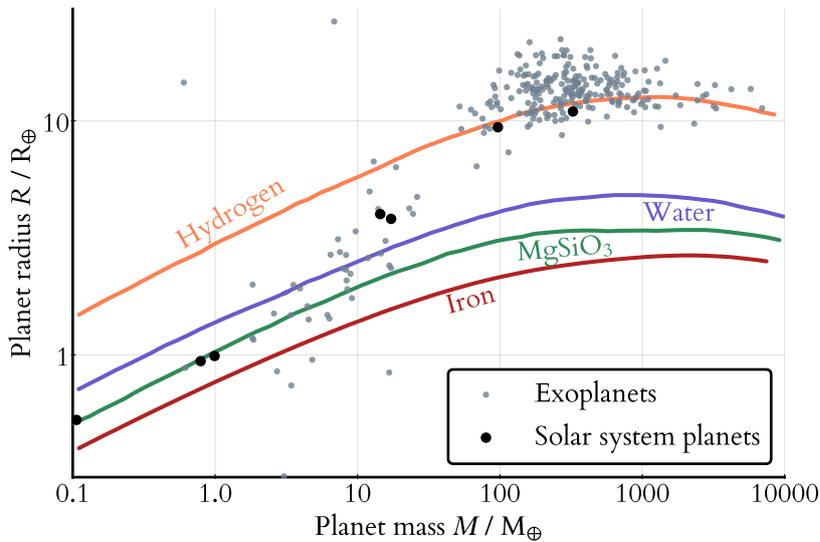


Figure 1.7: Mass–radius diagram of exoplanets and Solar system planets. Many Jupiter–mass planets lie above the predicted relationship for hydrogen spheres (coral lines). They are the so-called *puffy Jupiters*. The blue, green and red lines show the relationships for spheres of water, MgSiO_3 perovskite, and iron. By comparing a planet’s mass and radius to these, we can rule out certain classes of compositions—for example, if a planet has a larger radius than would be permitted for a sphere of rock or ice, it is likely to have a gaseous atmosphere. From this we see already that the exoplanetary inventory includes planets unlike anything in our solar system: hot Jupiters, puffy gas giants, mini Neptunes and super Earths. — *Mass–radius curves adapted from Rogers (2012), data from exoplanets.org*

1.2.1 Interior structures are the new frontier

As atmospheric characterisation techniques improve, a new question has naturally arisen: what lies beneath the atmosphere? We care about planetary interiors because they are linked to the formation history of the planet, because they are shaped by and shape the planetary atmosphere and because they are key to answering questions about habitability.⁴⁰ Understanding these exoplanets also allows us to place our own Earth into context: how unique are we? We therefore seek to understand, if not the interiors of individual exoplanets, at least something about broad classes of planets. But it is here that we are confronted by a lack of data, because we have very little ability to directly probe the interiors of exoplanets.

This lack of a rich source of observational data for planetary interiors means that we rely strongly on models. Even inside our solar system, our knowledge of planetary interiors is limited by the indirect ways in which we can probe them. On Earth we have the advantage of seismic measurements, and in our solar system we have various gravitational moments to constrain interior structures. Outside the solar system we have only the masses and radii of planets to work with. Models from first principles (numerical or analytical models based on the physics of solid and liquid spheres) therefore dominate the field.

Planetary interior models are a worthwhile starting point to make sense of the limited observational data we have. These models are inspired by earlier successes with stellar structure models, which are key to interpreting observations of stars. The ever-increasing number of

⁴⁰ Sotin et al. (2010)

known exoplanets, many of which have both mass and radius measurements, are a diverse and interesting set of objects to which to apply these models.

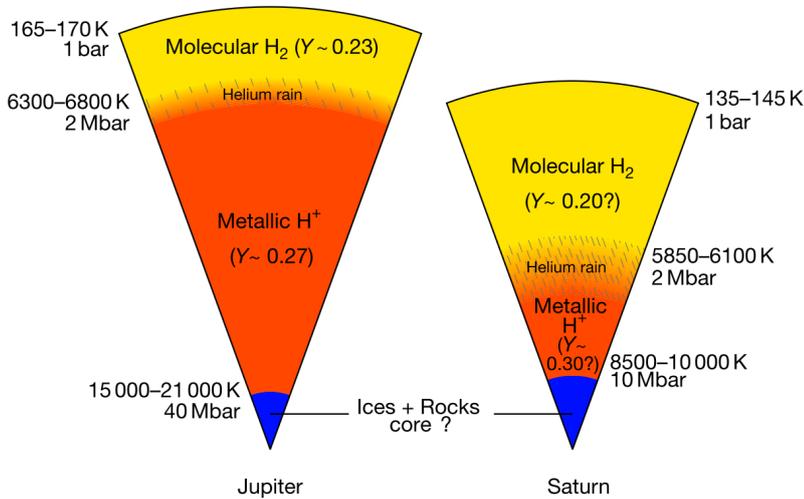


Figure 1.8: Example of some internal structure models for Solar System planets, Jupiter and Saturn in this case. — From Guillot & Gautier (2015)

Others had previously considered the internal structures of planets in our solar system,⁴¹ but the seminal study in the exoplanetary community is often considered to be the work of Zolpolsky & Salpeter (1969), who constructed mass–radius relations for large homogeneous isothermal spheres. Since then a number of internal structure models have been developed specifically for exoplanets,⁴² with work on super-Earths beginning about a decade ago.⁴³ In a nutshell, the basic assumption is that the internal structure may be modelled in one dimension by using similar equations to those of stellar structure.⁴⁴

Such one-dimensional modelling has been carried out for a number of bodies, including rocky/icy planets.⁴⁵ Planets with a gas envelope can also be modelled, although additional equations come into play in the atmosphere.⁴⁶ This method aims to build a self-consistent model based on some known constraints: for example, given a mass and radius, the mass fractions of a two-layer model may be determined exactly if the compositions of each layer are assumed.⁴⁷ Other authors have used this approach to scan the parameter space of possible compositions, often using *ternary diagrams* (figure 1.9) to show the ranges of materials permitted for a given planet,⁴⁸ or producing grids of models.⁴⁹

In addition to one-dimensional models, numerical simulations of processes like convection may also be useful for understanding the interiors of planets. Here there are very strong links to the geological sciences. But it is not necessarily possible to directly apply methods from

⁴¹ For example, see Hubbard & MacFarlane (1980) or Guillot et al. (1995).

⁴² Guillot (1999)

⁴³ e.g. Valencia et al. (2006); Fortney et al. (2007); Sotin et al. (2007); Seager et al. (2007); a recent review for giant planets is provided by Guillot & Gautier (2015)

⁴⁴ See chapter 3 for more on these models.

⁴⁵ Valencia et al. (2006); Valencia et al. (2007); Seager et al. (2007); Nettelmann et al. (2011)

⁴⁶ Vazan et al. (2013); Fortney et al. (2010); Baraffe et al. (2008)

⁴⁷ Nettelmann et al. (2011)

⁴⁸ Valencia et al. (2007)

⁴⁹ Zeng & Sasselov (2013)

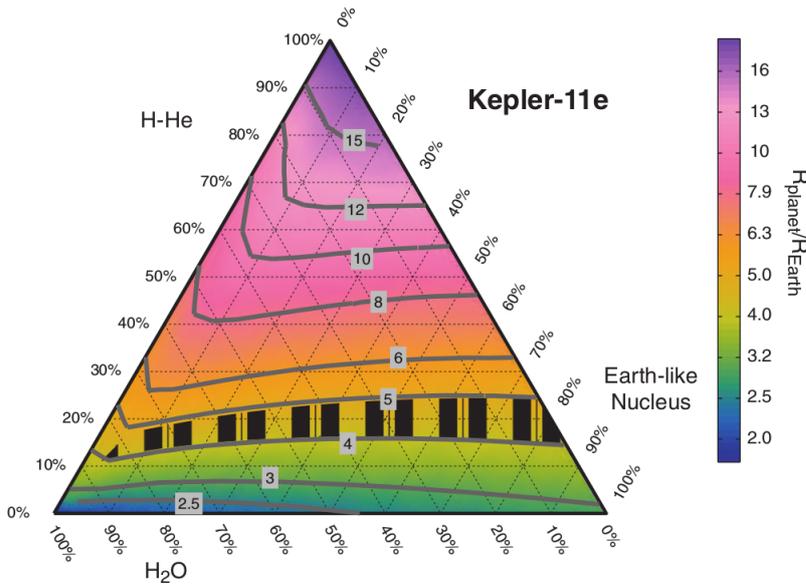


Figure 1.9: This is an example of a ternary diagram. It shows the link between the modelled radius of an exoplanet (Kepler-11e) and its composition, represented here by the three sides of the triangle. The triangle’s sides correspond to three mass fractions: an Earth-like nucleus, a water layer, and a H-He atmosphere. The black stripes show the constraint placed by the mass of the planet. Ternary diagrams like this are useful when three numbers, like these mass fractions, sum to one. — From Valencia et al. (2013)

the earth sciences because many of them depend on a more detailed knowledge of the Earth than we can achieve for exoplanets—for example, seismological data and gravitational moment measurements. An immediate first check on the validity of exoplanetary models does come from applying them to the Earth, however, and the results from this are generally promising, as shown in figure 1.10.

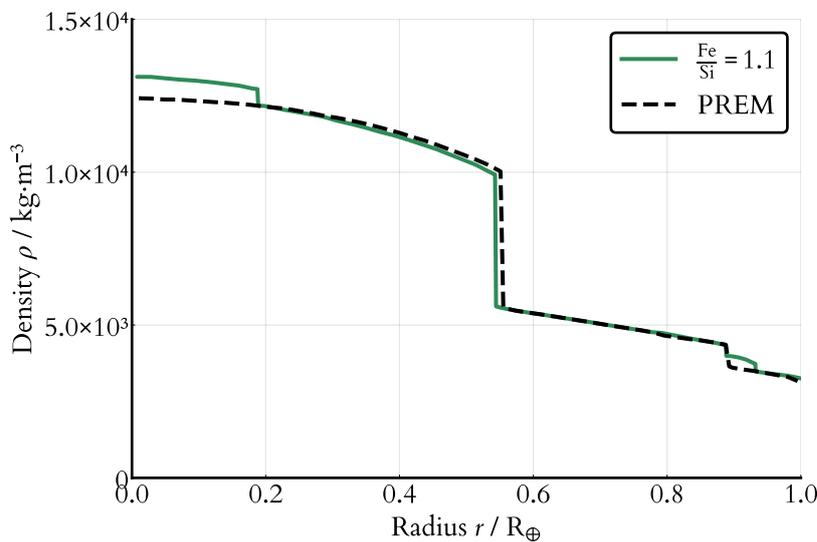


Figure 1.10: Comparison of a self-consistent density profile to an Earth model obtained through seismological data. Modelling the density profile of the Earth as if we were modelling an exoplanet produces results that compare well with the Preliminary Reference Earth Model (PREM), which is determined by inverting measurements of seismic wave propagation to solve for the internal structure. — Adapted from Baraffe et al. (2014)

1.2.2 These interior structure models are useful in several ways

The first way in which planetary interior models can be useful is to make broad inferences about the structure of a planet. There is some

information available about any planet despite an inherent degeneracy between different compositions. We can immediately exclude certain classes of models: for example, small planets with large radii must almost certainly have large hydrogen envelopes. We can also take more sophisticated approaches. For example, additional information such as the host star's properties can be included in the models.⁵⁰ Madhusudhan et al. (2012) argued for a carbon-rich interior in the exoplanet 55 Cnc e based on the stellar abundance and on the planet's density matching that of pure carbon. Using stellar abundances known at the time for the nebular composition, they determined that a high molar ratio of carbon during planet formation could lead to a carbon-rich interior, an interpretation supported by mass and radius measurements. And Dorn et al. (2015) showed that mass and radius alone can constrain the size of a planet's core if we assume it is pure iron. They also found that knowledge of stellar elemental abundances can reduce degeneracy in interior structure models and constrain the mantle composition.

⁵⁰ Sotin et al. (2007)

We can also hope to make progress in a statistical sense by examining populations of planets. Such progress is possible even if we are unable to pin down the precise structure of an individual planet. There are promising advances in this direction already. These usually involve inverse Bayesian analyses. For example, Rogers (2015) investigated the size demographics of planetary populations and set an approximate boundary of $1.6 R_{\oplus}$ beyond which planets are likely to have gaseous envelopes.

Finally, interior structure models may be useful when combined with prescriptions for planetary formation. We would like to link internal structures with simulations of planet formation and use our knowledge of planetary systems to generate synthetic populations of exoplanets. From a statistical point of view, this makes sense: because we only have one solar system to compare our observations to, our best bet for understanding other systems may be to statistically evaluate the likelihood of observing them based on formation models. Though this approach is undoubtedly complicated, it is a promising one for assessing the statistical properties of planetary systems. Mordasini et al. (2012b) did this, combining interior structure calculations with models of planet formation to produce synthetic populations of planets. In 2014 they updated their model, building an extensive hierarchical model for planet formation and evolution in an approach termed "planet population synthesis" with the goal of both predicting statistical properties of

observed planetary populations and testing the constituent models.⁵¹ Lopez et al. (2012) have also made model planets and explored how they evolve and lose mass through time.⁵²

If we are to use mass and radius to constrain the interior structure of a planet, we should ensure that our models are precise and accurate. But more importantly, we should understand where our models need to be precise and accurate and where such effort is wasted. We therefore require a thorough understanding of which factors can affect the mass–radius relation. We also need to know to what extent we are able to invert the relation to determine a composition.

1.2.3 Accuracy is difficult because of structural degeneracies

The internal structure of a planet is not well-constrained by its mass and radius alone.⁵³ However, we know that we can obtain some compositional constraints from observations of the planet and its host star. Above, I mentioned works by Sotin et al. (2007) and Madhusudhan et al. (2012), who used host star information in this way. Dorn et al. (2015) also used probabilistic models, incorporating the host star chemical abundances, to conclude that “uncertainties on mass, radius, and stellar abundance constraints appear to be equally important.” Grasset et al. (2009) indicated the need for good radius measurements, especially for dry silicate-rich planets for which numerical models can provide radius estimates to precisions of less than 5%. And Unterborn et al. (2016) used a mineral physics toolkit to perform a sensitivity analysis for rocky super-Earths, concluding that the mass–radius relationship is most strongly altered by the core radius and the presence of light elements in the core.

The presence of an atmosphere could also contribute significantly to the observed radius. Rogers & Seager (2010b) have modelled isothermal super-Earth interiors overlaid by a volatile atmosphere. And Valencia et al. (2013) considered coupled atmosphere–interior models, which also included atmospheric mass loss, and explored the dependence of radii on various model parameters such as the irradiation, water content and metallicity. The effect of an atmosphere is important, especially given that observations can probe spectral ranges where atmospheric absorption could be significant.⁵⁴

Though the factors above are all important, the effect of temperature on the mass–radius relation has not been thoroughly explored. This is for several reasons. First, its effects are thought to be relatively minor

⁵¹ Mordasini et al. (2012b); Mordasini et al. (2012a); Mordasini et al. (2015)

⁵² See also Owen & Wu (2016).

⁵³ Rogers & Seager (2010a)

⁵⁴ Madhusudhan & Redfield (2015)

in the first place: [Howe et al. \(2014\)](#) estimate that the effect of thermal corrections on an iron-silicate planet's radius is approximately 5%. [Grasset et al. \(2009\)](#) also describe how the radius of an Earth-like planet is not strongly affected by temperature changes. If the effect is small compared with current observational uncertainties, it is not necessarily relevant. Second, modelling is easier if we assume zero-temperature or isothermal spheres of material, because we do not have to deal with energy transport within the planet. Finally, the data on thermal expansion of heavy elements are sparse at the high temperatures and pressures characteristic of planetary interiors.⁵⁵ Therefore mass-radius relations or models of individual planets traditionally had no temperature dependence at all⁵⁶—but it is increasingly being included and thermal effects on radii are being explored.⁵⁷

⁵⁵ [Baraffe et al. \(2008\)](#)

⁵⁶ [Zapolsky & Salpeter \(1969\)](#); [Seager et al. \(2007\)](#)

⁵⁷ For example, see [Valencia et al. \(2013\)](#).

1.2.4 *Observational constraints are increasingly strong*

To what degree can interior structure models be linked with observations? We have already seen in [figure 1.9](#) that there is necessarily some degeneracy in the interior structure of any given planet. Unless the planet is taken to consist of only two layers of known material with mass fractions that are allowed to vary freely, a given mass and radius cannot be used to uniquely determine the internal structure of a planet. More encouraging results may come from a Bayesian approach such as that taken by [Rogers & Seager \(2010a\)](#), who quantify the constraints that transit and radial velocity observations can place on a planet's interior structure ([figure 1.11](#)), or [Dorn et al. \(2015\)](#), who characterise how the model variability depends on the data and its uncertainties.

Unfortunately, these constraints are still not ideal—there is an inherent degeneracy which cannot be overcome with this method, “originating from the fact that planets of differing compositions can have identical masses and radii.”⁵⁸ This is further exacerbated when the models are extended to include four-layer structures with the aim of handling a possible H/He gas layer. This approach, however, is completely independent of any constraints from planet formation hypotheses.

⁵⁸ [Rogers & Seager \(2010a\)](#)

Perhaps a better chance of correctly modelling these planets may come from incorporating knowledge of their host stars, or of formation conditions in the protoplanetary disk. For example, giant planet occurrence is known to correlate with stellar metallicity,⁵⁹ indicating that the host star's chemistry is important. Other works are increasingly

⁵⁹ [Fischer & Valenti \(2005\)](#)

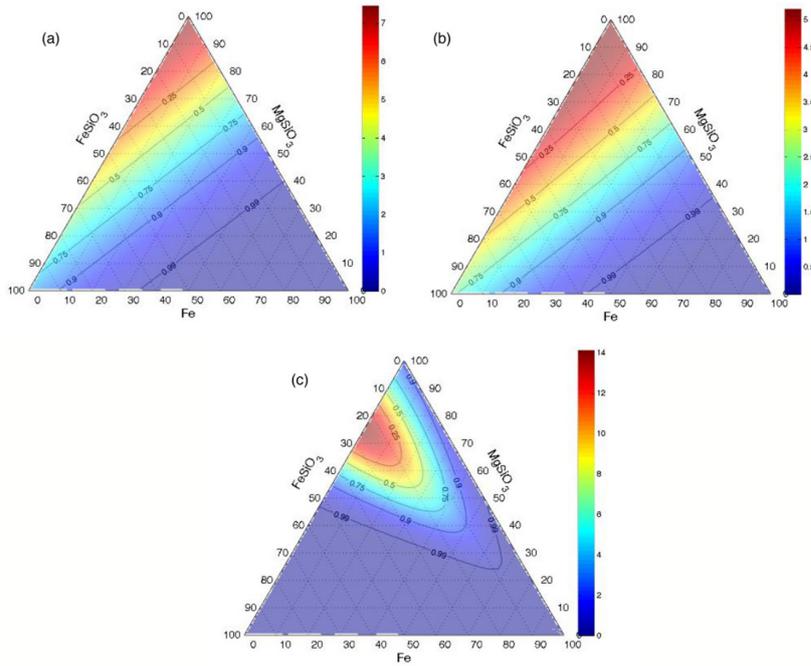


Figure 1.11: Bayesian techniques can be used to quantify the constraints placed on the interior composition of a planet, and they also account for uncertainties in the mass and radius and the inherent degeneracy in structure. Here, the posterior likelihood distribution for the interior composition of CoRoT-7b shows that a range of structures are permitted. The three different panels are for three different priors. — *From Rogers & Seager (2010a)*

addressing this link from a theoretical perspective.⁶⁰

Next-generation surveys promise improved measurements of mass and radius. Many of the approximations made in interior models were based on the fact that the uncertainties in mass and radius were large enough that the temperature variation could be discarded. However, current observations can already constrain these values to a few per cent for the best-studied targets, and new surveys like TESS, CHEOPS, and PLATO will provide high-quality observations for many more. For example, TESS is expected to detect “hundreds of super-Earths ($1.25\text{--}2 R_{\oplus}$)” and provide asteroseismological measurements for thousands of stars.⁶¹

Finally, with new observations comes the potential to use data in new ways. For example, transit observations could be used to constrain the mass of a planet.⁶² Polarimetry measurements could provide more information about particulate matter in the atmosphere or surface properties.⁶³ And by taking measurements of repeated transits, we could gain information on the planetary Love number, a measure of central concentration.⁶⁴ From this we should take away one key point: the observations are getting better, and so our modelling capabilities must match them. This, therefore, is the direction I take, focusing on one important category of planets: super-Earths.

⁶⁰ Bond et al. (2010); Carter-Bond et al. (2012); Moriarty et al. (2014)

⁶¹ Ricker et al. (2014)

⁶² de Wit & Seager (2013)

⁶³ Madhusudhan et al. (2014a)

⁶⁴ Maxwell et al. (2011)

1.3 Heated watery super-Earths

One of the most interesting classes of planets today is that of *super-Earths*, planets with masses between 1 and $10 M_{\oplus}$. With no analogues in the solar system, it is not known whether they are scaled up rocky planets or scaled down Neptunes. About forty super-Earths with measured masses and radii are currently known. Their radii range from 1 to $7 R_{\oplus}$.⁶⁵ With the potential to have moderate atmospheres and plate tectonics, super-Earths represent an important class of planets in the broader context of planetary diversity and planetary habitability.⁶⁶

Above, I highlighted recent observational advancements that have led to increasingly precise measurements of masses and radii of these small planets. Such measurements are being used with internal structure models to place constraints on the interior compositions of super-Earths. Many planets are well-described by multi-layer models consisting of iron, silicates, and water⁶⁷ and others have included layers of hydrogen or other volatiles to explain the inflated radii of some super-Earths.⁶⁸ Given the high-precision radii measurements, it is now important to quantify the dependence of predicted radii of super-Earth models on the various model assumptions—in particular, the temperature and energy transport.

1.3.1 The equation of state is a key component of these models

To correctly determine the interior structure and evolution of planets, we rely on a strong knowledge of the equations of state⁶⁹ of the relevant materials. This is not an easy task because the temperatures and pressures inside planets are often inaccessible to experiments. We must therefore “rely on simulations to infer the iron melting curve for Earth, super-Earth and giant planet conditions”.⁷⁰ But such simulations are computationally expensive and so this is an area that is still developing. For this reason, many authors choose to use simplified forms of the equations of state in calculating their models.

The most common simplifications of equations of state are semi-analytic forms such as those presented by [Vinet et al. \(1987\)](#) (the “Vinet EOS”), [Birch \(1947\)](#) (the “Birch-Murnaghan EOS” or “BME”), and [Poirier & Tarantola \(1998\)](#) (the “logarithmic EOS”). These are generally parametrised by values such as the bulk modulus of the material and its derivative. The parameters are then fitted to experimental measurements of the pressure and volume, and so are described as “universal”

⁶⁵ This number is taken from the [exoplanets.org](#) database (confirmed planets only).

⁶⁶ [Haghighipour \(2011\)](#); [Baraffe et al. \(2014\)](#)

⁶⁷ e.g. [Valencia et al. \(2006\)](#); [Fortney et al. \(2007\)](#); [Sotin et al. \(2007\)](#); [Seager et al. \(2007\)](#)

⁶⁸ e.g. [Rogers & Seager \(2010a\)](#); [Lopez et al. \(2012\)](#)

⁶⁹ The equation of state is introduced fully in chapter 2.

⁷⁰ [Baraffe et al. \(2014\)](#)

equations of state.⁷¹ These simple prescriptions cannot capture all the behaviour of every material, although they are generally good approximations for most.

⁷¹ Vinet et al. (1987)

In addition to being approximations to experimental data, most of the prescriptions used are temperature-independent for two reasons. First, the temperature dependence can be difficult to measure experimentally, especially at thousands of degrees. Second, including temperature dependence in the EOS generally does not result in differences greater than a few per cent in the bulk composition (mass/radius) for rocky planets.⁷² Therefore, an important question to ask ourselves is what we can gain from investigating the equation of state.

⁷² Seager et al. (2007)

1.3.2 *Water's thermal properties may lead to interesting structure*

The degree to which thermal structure may affect the properties of a water-rich planet has not yet been well studied.⁷³ Super-Earth planets with significant water layers, sometimes called waterworlds, provide an interesting testbed for our investigation. They may display more significant variation in their properties, both observable and internal, than purely Earth-like (iron and silicate) planets. They are therefore a worthwhile target for study.

⁷³ The closest work is perhaps that of Valencia et al. (2013), who include some temperature information in their models of water-rich sub-Neptunes.

Water presents an opportunity to assess thermal effects in a material that has a rich and interesting phase structure across a large temperature and pressure range (figure 1.12). At low temperature and pressure, water exists as a liquid, vapour, or solid (Ice Ih). At high pressure, it takes on a number of alternate ice forms (Ice V, VI, VII, X, etc.)⁷⁴ It can also exist as a low-density supercritical fluid or superheated vapour. This all means that the behaviour of water layers is thermally interesting. The behaviour of water is also strongly linked to questions of habitability because Earth-sized solid planets with oceans provide the best approximation to our Earth, the one planet known to harbour life.

⁷⁴ Choukroun & Grasset (2007)

Others have previously investigated the structures of planets containing a significant water component. For example, Ehrenreich et al. (2006) studied the internal structure of the exoplanet OGLE 2005-BLG-390Lb, modelling the phase changes throughout. Zeng & Sasselov (2014) chose to explore evolutionary effects, following the phase transitions within model water-rich planets. They comment that “[phase] transformations may have a significant effect on the interior convective pattern and also the magnetic field of such a planet, but they may only affect the overall radius slightly.” It is therefore appropriate to address

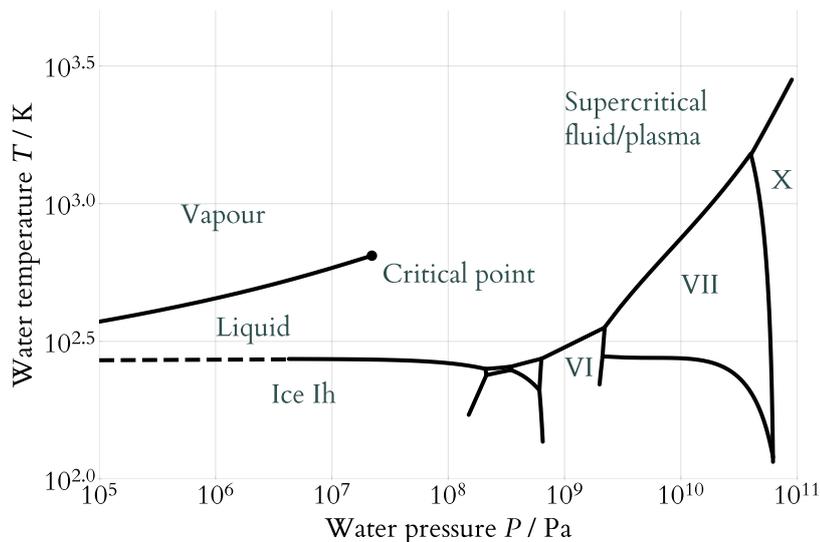


Figure 1.12: Phase diagram of water. Water has a rich and interesting phase structure. Here I show some of the key phases that are relevant when modelling a watery planet: liquid, vapour, and solid ice Ih, but also more exotic phases such as the high-pressure ices. Lines mark the boundaries of each phase. — Data from Choukroun & Grasset (2007) and Wagner & Pruß (2002)

the question of exactly how much temperature variations affect the structure and radius of water-rich planets and whether such effects are observable.

In summary, water is an interesting test case for assessing how the thermal behaviour of the equation of state may affect the observable properties of a planet. And although it is certainly not the only component of interest, the potential for planets with large water layers means that it is worthwhile investigating.

1.3.3 Cold icy satellites in our solar system have been well-studied for decades

Before we continue, it will be worthwhile to review previous work on icy bodies in our solar system. Although the diversity of exoplanetary mass–radius relationships suggests that planets have a broad range of water contents, from dry planets to waterworlds, the latter are conspicuously absent from our solar system—except for the icy moons of Jupiter outwards, plus planetary objects beyond Neptune. Water exists in these moons as icy shells and deep oceans (Hussmann et al., 2006) and so these satellites and minor planets are natural low-temperature analogues to the kinds of planets we will be considering in this dissertation. Due to their relative accessibility, these icy bodies have been studied for several decades. And although many of these techniques used to do so are not available to us when observing exoplanets, previous work on cold icy moons provides a valuable starting point.

EVEN DECADES AGO, the thermal behaviour of water was considered

absolutely key to understanding the potential structures of icy Solar System bodies. Lewis (1971) built steady-state models for thermally conductive icy bodies, showing that they could have extensive melted interiors. And Consolmagno & Lewis (1978) studied the thermal history of icy bodies of various sizes, considering how they might be heated internally and how this heat might alter their properties. In these early days, these authors were already highlighting the need for basic research into the behaviour of ices and rock–ice mixtures at high pressures. This is a refrain repeated throughout the literature today.

The reason for this interest in water is simply that there is so much of it in the satellites of Jupiter and beyond. I will not be reviewing the various methods used to show this⁷⁵ but, for example, spectroscopic observations show high proportions of water ice at the surface of moons like Callisto (20–45%), Ganymede (90%) and Europa (almost 100%).⁷⁶ Bulk densities alone suggest that, while some satellites like Io are almost entirely silicate, others like Tethys are almost entirely water.⁷⁷ With these and other indirect methods available to determine the internal structures of these objects, we know or strongly suspect that many of them contain a lot of water. Furthermore, this may often be in the form of liquid reservoirs or oceans. Cassini’s observations of regular watery jets emanating from the south pole of Enceladus⁷⁸ provide direct evidence for this.⁷⁹ More indirect evidence also comes in the form of magnetic fields: Europa and Callisto are expected to have salty subsurface oceans, and we infer this because Jupiter’s magnetic field appears to induce a secondary field in these moons via magnetic induction.⁸⁰

How can liquid water exist within these frigid icy bodies? The answer is because water’s unusual phase structure allows for an ice I shell on top of a liquid ocean (figure 1.13). Ice I’s melting temperature decreases with increasing pressure, and ice I is less dense than the adjoining liquid phase—this is not true for rock, and also not true for higher-pressure water ices. Larger bodies such as Ganymede, Callisto and Titan have internal pressures sufficient to cause water phase transitions, and internal heat sources sufficient to maintain subsurface liquid water oceans.⁸¹ And such liquid layers are more likely if antifreezes like ammonia are present.⁸² For this reason, icy satellites are often modelled with two or three water layers (ice I, liquid ocean, and possibly ice V or VI for larger satellites) over rock. Such differentiation is assumed to be common in icy bodies because rock may easily sink through liquid

⁷⁵ They rely on a host of measurements from imaging or probe flybys: density, gravity and gravitational moments, rotation, shape, magnetic field, temperature, surface/atmospheric composition, surface activity, geology and tectonics, and orbital dynamics. Nimmo & Pappalardo (2016) provides a comprehensive overview.

⁷⁶ Hussmann et al. (2015)

⁷⁷ Nimmo & Pappalardo (2016)

⁷⁸ These geysers appear to be caused by water trapped in pressurised subsurface reservoirs. They open on a regular basis due to tidal forces, and the trapped water boils off into space.

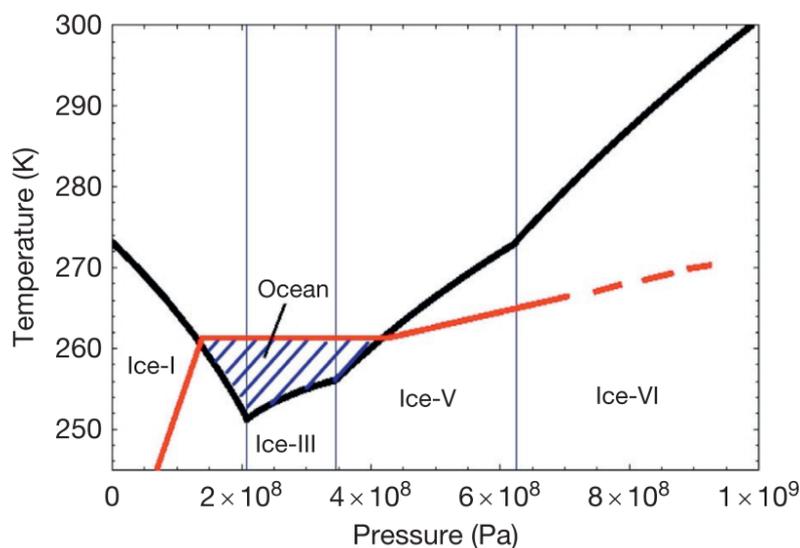
⁷⁹ Porco et al. (2006)

⁸⁰ Nimmo (2018)

⁸¹ Hussmann et al. (2015)

⁸² Lewis (1971)

water during formation.⁸³



Closely linked to the phase structure of water is the thermal structure of these icy bodies. Though their internal heating processes are not fully understood,⁸⁴ a couple of things are clear. First, these icy bodies generally have radiogenic heating as a primary heat source, with short-lived nuclides like ^{26}Al and long-lived nuclides like K, U and Th helping to keep larger satellites warm enough to maintain liquid oceans.⁸⁵ But even in our solar system we also expect some cases where tidal heating plays a significant role and others where different accretion and thermal histories have produced the diverse structures that we see.⁸⁶ Second, it is likely that larger bodies have convection as the dominant mode of heat transfer under a conductive ice I “stagnant lid”. The more internal heating, the thinner this lid—but thanks to the lid, an icy body’s internal heat can be maintained over a long period of time despite the efficiency of convective heat transfer deeper inside.⁸⁷ For example, Pluto, Titan, Ganymede and Calisto are thought to be able to maintain subsurface water reservoirs despite the lack of any appreciable tidal heating.⁸⁸ There is also increasing use of global circulation models to examine energy transport in more detail, such as by Soderlund et al. (2014) who modelled energy transport in Europa’s icy shell.

This previous work on icy bodies in our solar system suggests that water-rich super-Earths are also likely to be rich in thermal behaviour and strongly influenced by internal heating. With new data on the behaviour of water at high pressures and temperatures becoming available, they are therefore ideal targets for further study.

⁸³ Nimmo & Pappalardo (2016); Nimmo (2018)

Figure 1.13: Sketch of the phase diagram of water near the ice I/III/liquid triple point. The phase structure of water and the fact that ice I is less dense than liquid means that cold ice I may float on warmer water, which itself may sit on warmer high-pressure ice. Here we see the melting curves of water ices (black) and a representative temperature profile (red) that crosses a liquid region between two different ice phases. An icy body with this temperature profile would have an ice I crust above a liquid ocean. The effect is exaggerated if antifreezing agents like ammonia are present. — Image from Hussmann et al. (2015)

⁸⁴ For example, it is unclear how Enceladus is so strongly heated as to produce the cryovolcanic jets of water vapour seen at the south pole (Porco et al., 2006). And even Pluto’s interior was discovered to be more active than previously expected when New Horizons observed mounds possibly of cryovolcanic origin (Monteux et al., 2018).

⁸⁵ Hussmann et al. (2006)

⁸⁶ Hussmann et al. (2015); Monteux et al. (2018); tidal heating is thought to be responsible for keeping Europa’s subsurface ocean liquid and for Io’s volcanism (Nimmo & Pappalardo, 2016).

⁸⁷ Hussmann et al. (2015); Nimmo (2018)

⁸⁸ Nimmo (2018)

1.4 *This dissertation*

In this chapter I have given an overview of recent work in exoplanet detection and characterisation. I have shown that, although the field is driven by observations, we need good models to interpret these observations. These considerations therefore naturally motivate the development of interior structure models. I have identified objects with thick water layers as being interesting targets for questions about heating and thermal effects. And finally, I have given an overview of cold icy satellites in our solar system, a well-studied class of objects with unique and intriguing interiors thanks to the rich phase structure of water. We now push onwards to examine the structures of objects that are similar, yet not found in our solar system: hot water-rich planets.

In this dissertation I investigate a range of topics relating to the internal structure of water-rich super-Earths. I address the following key questions:

1. What is the water equation of state at the temperatures and pressures expected in heated water-rich super-Earths?
2. Under what conditions, and how strongly, does temperature affect the observable properties of such a planet?
3. How does this change when a vapour atmosphere is included in the models?
4. How do different modes of heating and energy transport affect the interpretation of a watery planet?
5. Do water-rich planets undergo significant changes during migration?
6. Can we use these models to understand more about the geology of watery planets?

The remainder of this dissertation presents, in five chapters, my contributions to these questions.

I begin by presenting a newly synthesized equation of state for water that includes a full treatment of temperature dependence. I explain where I obtained the data from, how I constructed the equation of state, and the features that make it suitable for planetary modelling. By comparing this to the equations of state used in other studies, I show that I have produced a comprehensive and up-to-date coverage of the

Chapter 2
"An improved water equation of state"

behaviour of water in the temperature and pressure ranges relevant to planetary interiors.

Next I explain the theory of planetary interior models. I describe my code, `ONION`, which solves the equations of planetary structure in a self-consistent fashion and treats the temperature structure of the planet appropriately. After incorporating my water equation of state into these models, I produce multi-layered planet structures for super-Earths with large water layers. I also validate these models against some previous results. This chapter concludes by exploring how strongly the radius of a watery super-Earth is linked to its water fraction, surface temperature, and other model parameters.

I then present an approach for understanding how different heating modes may alter a planet's size. I extend the models from the previous chapter to include a diffuse atmospheric layer, using two different approaches for comparison. By linking the external boundary conditions to the planet's internal and external heating, I show that the amount of energy required to significantly inflate these watery planets is modest.

To wrap up my theoretical development of planetary structure, I turn to the question of how a planet's internal structure is altered by this heating. Using the information on the phases of water contained in my equation of state, I produce structural diagrams for watery planets. I show how different heating scenarios may alter this structure, and construct some toy migration scenarios to investigate how it may change over time. This chapter finishes by considering some potential astrophysical implications of different phase structures for these ocean planets.

Finally, we have recently had the opportunity to apply these models to a super-Earth with a potential water detection. I present some initial results, for which my models were used to analyse the planet. These include constraints on its structure and some geophysical properties.

Chapter 3
"Watery planet interiors"

Chapter 4
"Heating and the atmosphere"

Chapter 5
"Phase structure and migration"

Chapter 6
"A water-rich super-Earth"

2 *An improved water equation of state*

The most important piece of information that we require when making a planetary model is the *equation of state* (EOS). We have a choice of several approaches for the planetary model itself, which I will discuss in chapter 3. But no matter how we go about modelling the planet, we require some information about the properties of the materials that comprise it. More specifically, we require a relationship that specifies how the density of a material varies with pressure and temperature. Such a relationship is the equation of state.

As the goal of this study was to investigate thermal effects in watery planets, I required a temperature-dependent equation of state for water. This allowed me to treat thermal expansion self-consistently in my models. However, no available equations of state met my requirements; I therefore synthesized a new equation of state for water from the best available experimental and theoretical data over a wide range of pressure and temperature.

In this chapter I provide an overview of the equations of state used in interior modelling so far. After explaining why none of them are suitable for the work I did, I then present my improved equation of state for water. I explain the process of compiling this equation of state, show how it is an improvement on previous work, and discuss the difficulties that arose when dealing with several disparate sources of data.

2.1 *About the EOS*

The equation of state is a relationship between pressure P , temperature T and density ρ (equation 2.1). The definition of an equation of state allows that it might also specify relations between other thermodynamic state variables; these include the specific heat capacities c_P and c_V and specific energies like the internal energy u or entropy s . But in this context I use it to refer to the relationship between the set of values (P, T, ρ) , which are the key values required for a one-dimensional

$$\rho = f(P, T) \quad (2.1)$$

The equation of state. The function f may be analytic or it may require looking up or interpolating between values in a table. The equation may also be written in an inverted form:

$$P = g(\rho, T) \quad (2.2)$$

depending on how the data is presented.

planetary structure model.

To calculate the temperature gradient also requires the specific isobaric heat capacity c_P , and so associated with this equation of state I include the heat capacity as a function of pressure and temperature.

Many potential expressions for the equation of state exist, spanning all phases from gas to plasma. Depending on its source and the material it represents, the equation of state may be expressed in a simple functional form or it may be more complicated than this. So-called universal equations of state take advantage of the fact that different materials behave in a common way under certain conditions. The best-known one is probably the ideal gas law, but others exist too: for example, there exist cubic equations of state which can partially account for the liquid–vapour transition, and the behaviour of solids when heated is often parameterised by including the material’s thermal expansivity in its equation of state. The ideal gas law is one of the simplest possible equations of state, and the complexity of the equation of state can be increased indefinitely.

Above, I used the word “universal” to describe equations of state, but in reality there is no such thing. No equation of state accurately captures the behaviour of any material in all situations. Though some approximations are good for certain phases or limits, a planetary model spans many orders of magnitude in temperature and pressure from the top of the atmosphere to the centre of the core. Real materials also have phase transitions between different states of matter, and these manifest themselves as discontinuities in the density surface.

2.1.1 Where do EOS data come from?

Because no true universal equation of state exists, it is inappropriate to use just one equation of state in a planetary model. And despite progress in both experimental measurements and *ab initio* theoretical calculations, there is still no all-encompassing equation of state database for all of the minerals of interest in planetary interiors. Perhaps the closest to our intended goal is SESAME,¹ the Los Alamos National Laboratory equation of state database. But US regulations preclude the distribution of SESAME to foreign nationals. Instead we must maintain several different equations of state for each different material, choosing appropriately from among them depending on the chemical composition, pressure and temperature.

Previous studies have approached the problem of there being no

$$c_P = h(P, T) \quad (2.3)$$

The specific isobaric heat capacity can also be represented in a similar form to equation 2.1. It is required to calculate the temperature gradient in the planet.

The ideal gas law,

$$P = \rho RT \quad (2.4)$$

links three state variables: pressure P , density ρ and temperature T . The material’s properties are encoded by the specific ideal gas constant R , defined as the ratio \bar{R}/M where \bar{R} is the universal gas constant ($\bar{R} = 8.314 \text{ J} \cdot \text{mol}^{-1} \cdot \text{K}^{-1}$) and M is the molar mass of the material.

¹ Lyon & Johnson (1992)

centralised database for equations of state by stitching together other equations of state which are valid for different pressure regimes. Although the choice of the exact equations has varied as new experimental data were released, few of these studies included thermal expansion. They included several different materials of interest for planetary interiors: water ice, iron, and silicates. For example, Seager et al. (2007) took this approach with water, combining three temperature-independent equations of state for ice VII:

- the Birch–Murnaghan equation (BME) of state at low pressures,
- density functional theory (DFT) calculations at intermediate pressures and
- the Thomas–Fermi–Dirac (TFD) model at very high pressures.

The piecewise function defined in this way is therefore appropriate over a wide pressure range.

This pressure piecewise approach neglects temperature dependence in the equation of state but provides a robust approximation that is easy to evaluate. In some cases, stitching the data in this fashion has revealed that a simpler functional form works just as well. For example, this is the case for the *polytropic equation of state*.² Such simple functional forms for the equation of state have been used successfully to model planets as cold spheres since the work of Zapolsky & Salpeter (1969). In other cases, a more detailed functional form is needed to capture the behaviour of the material fully; this is especially true if it undergoes phase transitions. For example, the IAPWS formulation³ uses a complicated series of equations fitted to various sources of experimental data for the behaviour of water in the vapour and liquid phases.

IN AN IDEAL WORLD we would simply measure the density (and other state variables) of the material for all relevant pressures and temperatures. But there are two issues that make this approach untenable. First, many of the temperatures and pressures relevant to planetary interiors are outside the reach of experiments for now. Second, even in the cases where they are accessible, it is often only via dynamic shock experiments rather than static compression experiments. Such dynamic experiments may not be appropriate for understanding the behaviour of materials under static pressure. In practice, we therefore choose experimental data for our equation of state where appropriate, and use different sources of data to fill in other regions.

Where do these other sources of data come from? I surveyed previ-

Where abbreviations are used in this chapter, table 2.1 indicates their origin.

² “[T]he common building blocks of solid planets all have equations of state that are well approximated by a modified polytrope of the form $\rho = \rho_0 + cP^n$ ” (Seager et al., 2007)

³ Wagner & Pruß (2002)

Equation of state	Source
ANEOS	Thompson & Lauson (1972)
BME	Birch–Murnaghan equation of state; see Poirier (2000)
DFT	Density functional theory; refers to theoretical calculations which multiple authors have performed
H ₂ O-REOS	Nettelmann et al. (2011); includes IAPWS, SESAME, French et al. (2009), Feistel & Wagner (2006)
IAPWS	Wagner & Pruß (2002)
LM-REOS	Nettelmann et al. (2008) (precursor to H ₂ O-REOS)
MGD	Mie–Grüneisen–Debye thermal pressure expansion; described in Sotin et al. (2007)
SESAME	Lyon & Johnson (1992)
TFD	Thomas–Fermi–Dirac; described in Salpeter & Zapolsky (1967)
Vinet	Vinet et al. (1987)

Table 2.1: Sources for the abbreviated equation of state designations used in this paper.

ous works⁴ that deal with internal structure, focusing exclusively on the water equations of state across all its phases (table 2.2). In order of complexity (and therefore their potential ability to capture the behaviour of the mineral), these equations of state are:

- Simple semi-empirical analytic⁵ prescriptions like the Birch–Murnaghan equation of state (BME), Vinet equation of state, or power laws;
- Theoretical high-pressure equations of state like the Thomas–Fermi–Dirac formulation (TFD), which produce results applicable in the high pressure regime;⁶
- Quantum molecular dynamics simulations from first principles;
- Complicated empirical analytic prescriptions, such as the work of the International Association for the Properties of Water and Steam (IAPWS) or the US national laboratory databases such as ANEOS and SESAME, plus tables of values taken from these; and
- Fully experimental data sets generated from shock wave compression or diamond anvil experiments.⁷

⁴ Howe et al. (2014) completed a similar exercise, providing an overview of previous equations of state chosen to model planetary interiors. But they dealt mostly with equations of state for other materials like iron and silicates.

⁵ By this I mean that the equation of state is of a fixed functional form, but the function is parameterised by one or more parameters that are tuned based on measurements of the material.

⁶ The ideal gas equation of state is similar in that it can be derived based directly on statistical mechanics, rather than being measured for each gas.

⁷ Diamond anvils are used to compress samples of material, measuring among other properties their density, up to pressures in the GPa range and higher. By adding resistive coils to the anvil or heating the sample with lasers, temperatures up to several thousand K can also be probed (Mao & Mao, 2007). Dynamic compression experiments, including shock compression, are also used to measure the equation of state at high temperatures and pressures (Asimow, 2015).

Work(s)	Water equation of state used
Baraffe et al. (2008); Baraffe et al. (2014) Fortney et al. (2007)	TFD, BME, MGD Simple power law from Hubbard & MacFarlane (1980)
Fortney & Nettelmann (2009)	H ₂ O-REOS
Grasset et al. (2009)	MGD; Vinet; BME; TFD; ANEOS; Belonoshko & Saxena (1991)
Guillot (1999)	Hubbard & Marley (1989)
Howe et al. (2014)	Vinet
Hubbard & MacFarlane (1980)	Simple power law
Hubbard & Marley (1989)	Exponential polynomial EOS without temperature dependence
Lopez et al. (2012)	H ₂ O-REOS
Madhusudhan (2012)	BME
More et al. (1988)	Quotidian EOS (ion EOS with Thomas–Fermi model)
Nettelmann et al. (2008)	LM-REOS
Nettelmann et al. (2011)	H ₂ O-REOS
Redmer et al. (2011)	French et al. (2009)
Rogers & Seager (2010b)	IAPWS; IAPWS extrapolations; TFD
Seager et al. (2007); Rogers & Seager (2010a); Zeng & Sasselov (2014)	Low-temperature polytropic EOS
Senft & Stewart (2008)	IAPWS; Feistel & Wagner (2006); Stewart & Ahrens (2005); BME
Sotin et al. (2007); Sotin et al. (2010)	BME with thermal expansion (MGD)
Valencia et al. (2006)	BME with thermal expansion
Valencia et al. (2010)	French et al. (2009); SESAME
Vazan et al. (2013)	Quotidian EOS; TFD
Wilson & Militzer (2012); Wilson et al. (2013)	DFT
Zeng & Sasselov (2013)	Frank et al. (2004); French et al. (2009); TFD

Table 2.2: Previous studies on super-Earth interior structures use a variety of equations of state for water.

2.1.2 The role of the water equation of state in planetary models

For the one-dimensional models that we use, the equation of state enters the structural equations in several places. The first is in the equation of mass continuity, which depends on the density ρ ; the second is in the equation for the temperature gradient, which depends on several quantities: thermal expansivity α , density ρ and specific isobaric heat capacity c_P . All of these quantities are provided by the equation of state.

To capture the thermal behaviour of a planet appropriately, an equation of state therefore requires the following properties:

- It must cover a wide range of pressures and temperatures relevant to both a planet's atmosphere and its interior.
- It must provide the heat capacity c_P , at least at pressures relevant to the envelope region. In the atmosphere the heat capacity is not required.
- Because the adiabatic temperature gradient lies on an isentrope, the EOS must include enough information to evaluate this isentrope. This could mean providing the entropy directly; however, as we will see, it is sufficient to include the heat capacity.
- It should include temperature dependent density changes wherever possible.

None of the individual equations of state in table 2.2 meet all these criteria. We therefore require an improved equation of state if we wish to model the thermal behaviour of a water-rich planet.

2.1.3 Motivation for an improved water EOS

Why do we care about thermal effects in the equation of state? Although the expression for the equation of state⁸ nominally includes temperature dependence, only a few studies have attempted to calculate self-consistent planetary models for super-Earths that include a temperature component.⁹ Many simply use zero-temperature equations of state instead. And in cases where the temperature has been included, it has often been treated in a simplified manner. For example, Zeng & Sasselov (2014) use a scaled Neptune adiabat for their internal temperature profile and also use a simplified equation of state that does not include all the thermal behaviour.

The lack of a complete treatment of temperature has been justified by the fact that the thermal expansion effect on an equation of state is

The equation of mass continuity is

$$\frac{dr}{dm} = \frac{1}{4\pi r^2 \rho} \quad (2.5)$$

and the equation for an adiabatic temperature gradient is

$$\frac{dT}{dm} = -\frac{T\alpha Gm}{4\pi r^4 \rho c_P} \quad (2.6)$$

where I have highlighted in red the terms which require the equation of state to calculate. These equations are introduced more fully in chapter 3.

⁸ Equation 2.1

⁹ For example, Valencia et al. (2013)

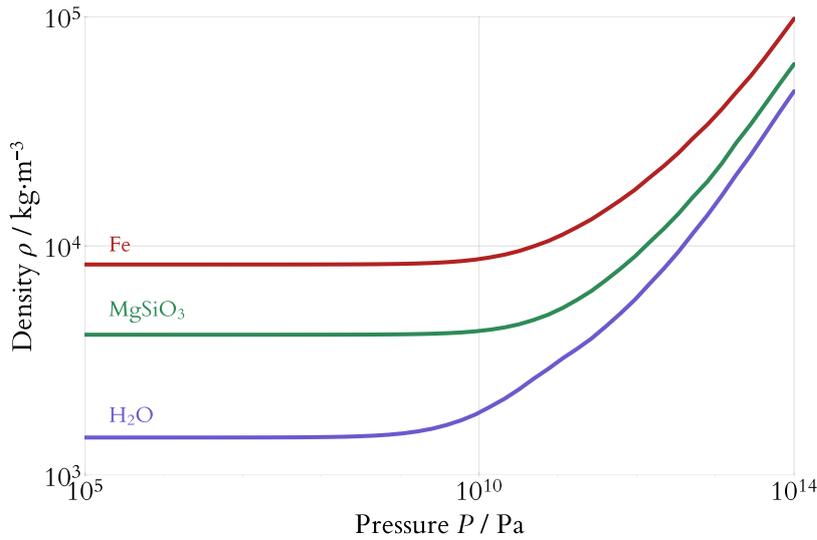


Figure 2.1: Simple zero-temperature equations of state used in Seager et al. (2007).

minor compared to the pressure. Seager et al. (2007) find that “water ice VII density changes by less than a few percent” up to 800 K. However, they do not include the high-pressure ionic and plasma fluid phases in this, and say that “[m]ore work needs to be done to quantify the thermal pressure effects above 250 GPa and in the ionic phase”.

With the knowledge that precisions on the order of a few per cent in mass and radius are attainable by current and upcoming observational campaigns, I argue that the role of thermal effects in these types of planetary models needs to be further investigated. This need is especially apparent when we consider the phase structure of water and its volatile nature. Not only can water vaporise easily, changing density suddenly across the liquid–vapour boundary, but it also has a rich phase structure at higher pressures and temperatures too. The volatile nature of water also means that it is also difficult to separate it into “atmosphere” and “envelope” layers and draw an artificial boundary in terms of pressure. See chapter 4 for further discussion on this: in that chapter I incorporate an atmospheric layer into my models and present an approach for resolving this difficulty.

In the introduction I argued that watery planets are interesting; here I have argued that the equations of state used to model them are insufficient. The next step is therefore to compile an improved equation of state for water in the range of temperatures and pressures relevant to planetary interiors. It is this task that I turn to in the remainder of this chapter.

2.2 How I collected and synthesized the data

I compiled a pressure–density–temperature relation for water across a wide range of pressures and temperatures, drawing from a number of the sources of data listed in table 2.2. I extended the piecewise approach described above to include temperature as a second dimension in parameter space. My stitched equation of state is valid over a wide domain: its temperature domain is from 275 K to 24 000 K, and its pressure domain is from 10^5 Pa (1 bar) upwards. My approach was similar to that of Senft & Stewart (2008), who generated a “5-Phase” equation of state across different liquid, vapour, and ice phases.¹⁰ However, their work focused on the lower temperatures needed to model impact craters.¹¹ I have explicitly included much higher temperatures so as to capture the behaviour of large super-Earth planets: we expect the cores of these to reach thousands of Kelvin. The sources I used are detailed in table 2.3.

Equation of state	Type*	Region of validity
IAPWS	Tabular	Vapour and liquid phases; 0.05 to 1000 MPa and 252.462 to 1273 K
French et al. (2009)	Tabular	Superionic, plasma and high-pressure ice phases; 79 to 9.87×10^6 MPa and 1000 to 24 000 K. I did not use table VIII from this work, as these low-density data disagree with the IAPWS formulation.
Feistel & Wagner (2006)	Tabular	Ice Ih; 0 to 200 MPa and 0 to 273 K
Sugimura et al. (2010)	Tabular	Ice VII; 18880 to 50250 MPa and 431 to 881 K
Vinet + MGD correction using parameters from Fei et al. (1993)	Functional	Ice VII
TFD	Functional	Ice X
Seager et al. (2007)	Functional	Ice VIII–X transition
Choukroun & Grasset (2007)	Functional	Ices I, III, V, VI; phase boundaries as specified by Dunaeva et al. (2010)
IAPWS extrapolations	Functional	Remaining regions

Some of these sources provided data in the form of tables, either in a machine-readable format or inserted directly into a journal article. Where available, I used the machine-readable information, but other-

The final equation of state database can be found online at github.com/swt30/WaterData.jl

¹⁰ Valencia et al. (2010) also constructed a similar equation of state, though using only data from SESAME and the IAPWS formulation.

¹¹ S. Stewart, personal communication.

Table 2.3: I used a variety of equations of state in compiling my final EOS. *“Tabular” indicates that I interpolated between values specified in the paper. “Functional” indicates that I used the functional form given in the paper.

wise I extracted values directly from the tables in the journal articles. Other sources provided equations of state using a functional description. In these cases I implemented the function as given and used the parameters specified in the originating article. When tabular values were given and the phase of water was specified, I also collected information on the phase and recorded this. I converted all units to SI units and produced a series of tables, standardising them for easy input into the next stage of the modelling process.

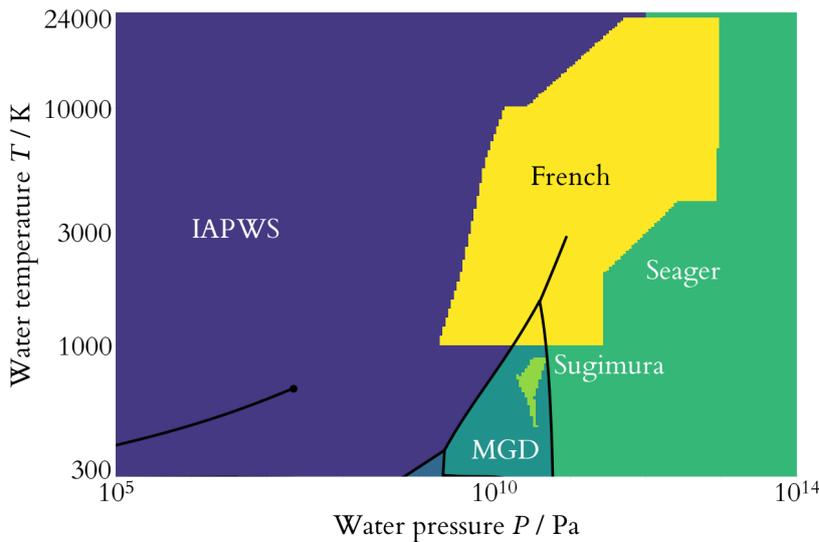


Figure 2.2: Data sources in my water equation of state. The equation of state covers a wide range of temperature–pressure space. Here I show some of the key data sources used and their regions of validity: the IAPWS formulation (Wagner & Pruß, 2002); theoretical calculations (French et al., 2009); the piecewise equation of state described by Seager et al. (2007); the Mie–Grüneisen–Debye (MGD) thermal correction approach for ice VII (Sotin et al., 2007); and measurements for a small region of ice VII (Sugimura et al., 2010). I also show the relevant phase boundaries.

My equation of state is for pure water only. Others have investigated how impurities may affect the equation of state and the planet’s properties. For example, Levi et al. (2014) included a methane component in their models, resulting in a new phase of water (filled ice) which changes the planet’s thermal profile. They note that, while neglecting volatiles is an impediment to understanding the planet’s atmosphere, pure water models may be sufficient for planetary mass–radius relations.

In selecting the equations of state I was often faced with choices between different sources of data. The exact behaviour of water at very high pressures is still uncertain and experimental and theoretical results are sometimes in conflict.¹² Ensuring absolute accuracy of the chosen equations of state was therefore a secondary priority. In general, I preferred more recent data to older data, I prioritised measured and tabulated values over functional approximations, and I chose representations that included temperature dependence over those that did not. In the following paragraphs, I describe my equation of state choices; these choices are summarised in table 2.3.

¹² Baraffe et al. (2014)

Liquid and vapour: The behaviour of water in the liquid and vapour phases is well understood and plenty of data are available. I was unable to gain access to the SESAME database¹³ because there are restrictions on the distribution of these data to non-US nationals. Instead, to represent water liquid and vapour, I selected the IAPWS (International Association for the Properties of Water and Steam) formulation,¹⁴ which provides both tabular and functional data for water in these phases. These are well-tested and validated by years of experiments.

I implemented the IAPWS functional relationships between temperature, density and pressure. Where appropriate, I numerically inverted these to give a relation of the form $\rho = f(P, T)$.¹⁵ I then tested these against the tables to verify that we had replicated them correctly. [Wagner & Pruß \(2002\)](#) claim that these functional forms can be extrapolated outside the range of the tables.

Ice VII: I explicitly chose a temperature-dependent formulation because I expect ice VII to form a significant fraction of the planet in the cases where the water layer is large. This temperature-dependent formulation is in contrast to other studies which have assumed that the ice VII layer is isothermal: for example, [Rogers & Seager \(2010b\)](#) assumed no expansion in any solid layers, choosing to include temperature effects only in the gaseous and liquid phases.

The best temperature-dependent formulation I found for ice VII was the Mie-Grüneisen-Debye (MGD) thermal correction approach.¹⁶ I used a Vinet equation of state with this thermal correction, taking the coefficients of [Fei et al. \(1993\)](#), within the ice VII region delimited by the phase boundaries of [Dunaeva et al. \(2010\)](#). However, I preferred the more recent tabulated measurements of [Sugimura et al. \(2010\)](#) wherever these were applicable; these are shown within the ice VII region in figure 2.2.

Supercritical fluid and plasma: [French et al. \(2009\)](#) presented quantum molecular dynamics simulations of high-temperature and high-pressure plasma, ice, and superionic fluid phases of water. I used their tables in the region beyond 1000 K and 1.86×10^9 Pa. This region has also been probed by laboratory experiments thanks to [Knudson et al. \(2012\)](#), who strongly advocate “that [the French equation of state] be the standard in modeling water in Neptune, Uranus, and ‘hot Neptune’ exoplanets.”. These temperatures and pressures are also relevant to the interiors of super-Earths. I did not use the low-density tables that they presented

¹³ [Lyon & Johnson \(1992\)](#)

¹⁴ [Wagner & Pruß \(2002\)](#)

¹⁵ This numerical inversion is described further in section 2.2.1.

¹⁶ [Sotin et al. \(2007\)](#)

separately because these differ significantly from the IAPWS results in the same temperature and pressure range. The IAPWS tables are better tested and a more appropriate choice in this region of phase space.

Low-temperature ices: For completeness, my equation of state also includes low-pressure ice Ih¹⁷ as well as higher-pressure ices such as ice III, V and VI. For the latter, I took the phase boundaries from Dunaeva et al. (2010) and used the temperature-dependent formulations for these ices given by Choukroun & Grasset (2007).

¹⁷ Feistel & Wagner (2006)

Ice X and beyond: I adopted a temperature-independent piecewise equation of state¹⁸ to describe the transition from ice VII to ice X and beyond. This does not include any temperature dependence: any behaviour of ice at these high pressures is increasingly theoretical and unconfirmed by experiment. Temperature effects approach zero at these high pressures anyway (figure 2.7), so I used the Thomas–Fermi–Dirac equation of state for all regions beyond 7686 GPa which were not covered by one of the other regions listed above.

¹⁸ Seager et al. (2007)

Other regions: Finally, I filled in all other regions according to the IAPWS formulation or extrapolations thereof. In practice, the only regions not covered above were low-pressure and high-temperature vapour regions, which I do not expect to be relevant for super-Earth interior models.

2.2.1 Numerical inversion

The functional forms of the EOS often specify a relationship between variables that is inverted compared to the one that I required.¹⁹ I wrote numerical inversion routines to handle these cases, parallelizing them for greater speed when evaluating large numbers of points in the state space at the same time.

¹⁹ Compare equation 2.2 to equation 2.1.

As an example of where this is necessary, let us take the IAPWS data set.²⁰ The paper in which it is introduced provides a functional description of the behaviour of water and steam that is applicable over a wide range of temperatures and pressures. It also provides a narrower set of tabulated values. Interpolating between the tabulated values is simple, but I also used the functional form of the EOS outside the range of the table. I based this decision on the claim of Wagner & Pruß (2002) that the IAPWS formulation extrapolates appropriately outside the

²⁰ Wagner & Pruß (2002)

range of validity of the table. In this case I needed to numerically invert the IAPWS equation of state outside the tabulated range.

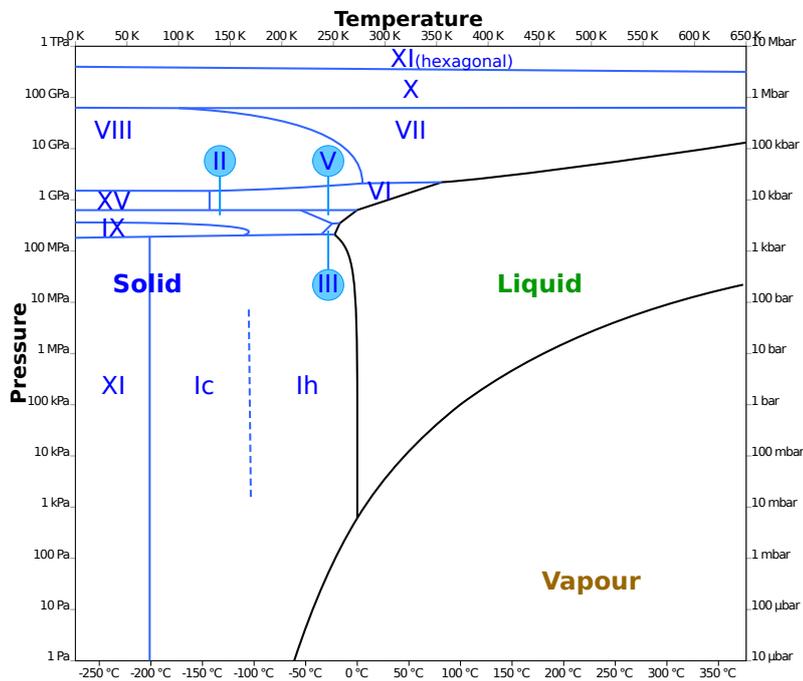


Figure 2.3: The phase diagram of water around room temperature. My final EOS includes the ice phases Ih, III, V, VI, VII, XI and X (both through the high-pressure TFD limit), liquid, vapour, and two theoretical high-temperature phases (plasma and super-ionic) not shown in this figure. — Adapted from a figure by cmglee on Wikimedia Commons / CC-BY-SA-3.0

Numerical inversion can be complicated by phase transitions. In the region covered by the IAPWS data, there is a phase transition (and therefore a density discontinuity) between liquid and vapour. This phase transition occurs along a line in (P, T) space called the saturation curve (shown on figure 2.3). This can result in numerical trouble when performing the inversion near the curve.

I handled numerical problems like this near phase boundaries by using a bounded root-finding algorithm when solving the inverse equations. By pinning one end of the bounds to the phase curve and fixing the temperature, I solve a one-to-one inversion problem on either the higher-density or lower-density side of the phase curve. This prevents the solver from stepping across the curve and obtaining the wrong solution. I did not expect this vapour-liquid transition region to be of much importance in my interior structure models, at least for the case when the surface is solid and the interior consists mostly of high-pressure ice. However, accurately handling the behaviour here was important when I later extended the models to include an atmospheric layer.

I specified an order of priority for which EOS to prefer in the case of conflict between sources. I rely firstly on tabular sources, then

functional sources, in the order specified in table 2.3. That is, the first valid EOS on that list determines the density at a particular (P, T) pair.

I defined the limits of each EOS patch through the notion of a *convex set*, which is the largest possible region of the state space fully enclosed by a set of points (figure 2.4). The convex set gives a natural interpretation of the area covered by a table of values. In the case of the low-pressure ices, I used phase boundaries provided in Dunaeva et al. (2010) and assigned points within each region an appropriate EOS from Choukroun & Grasset (2007). For the regions based on tabulated values, I applied the EOS only in the region covered by the data points where interpolation can be carried out. Outside these bounds, I revert to the next highest priority EOS from table 2.3.

2.2.2 Validation

Wherever I implemented a functional EOS, I wrote extensive tests to ensure that these functional forms correctly reproduced any values provided in the paper. I compared the implemented EOS to test values in the published paper (in the case of Wagner & Pruß (2002)), a subset of experimental data points (in the case of experimental papers, in order to check that the interpolation was valid), or against plotted isotherms (for other papers).²¹ I did this for papers that provided EOSes in the form $\rho = f(P, T)$, as well as the inverse form $P = g(\rho, T)$. I also verified that the IAPWS tabulated values matched up with the functional values at the boundary of the table. Figure 2.6 shows the match between these two data sets.

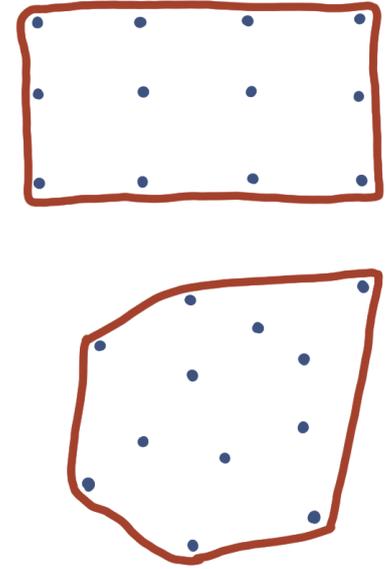
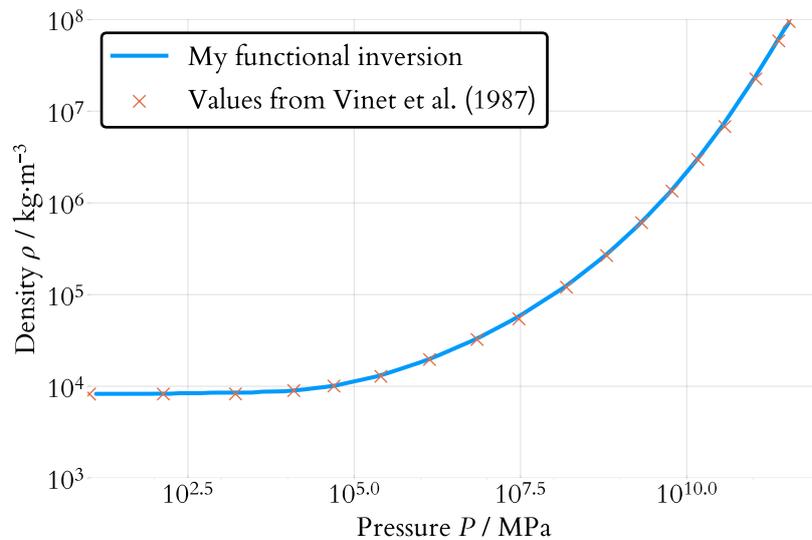


Figure 2.4: A convex set is the largest possible region of a space that is fully enclosed by a set of points. An intuitive way to visualise it is to imagine stretching a rubber band around the points; the area enclosed by the rubber band is the convex set. Above, the convex set for a regular grid; below, the convex set for an unstructured set of points.

²¹ For an example of one of these validations, see figure 2.5.

Figure 2.5: Example of a validation plot showing that my EOS matches the literature values. This is an example of one of the validation plots I produced to check that my implemented EOS matches the literature values. In this case, a numerically inverted Vinet equation of state is compared against values from Vinet et al. (1987), showing good agreement.

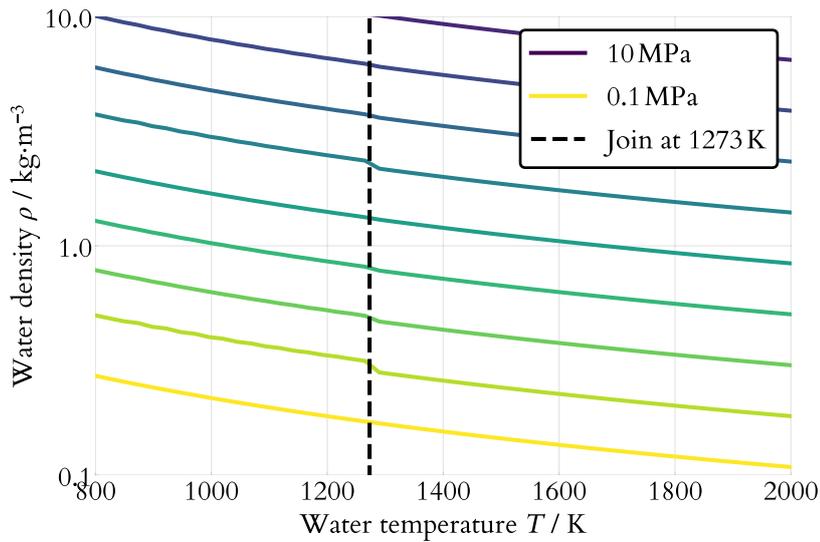


Figure 2.6: I verified that the tabular IAPWS data (left) joined smoothly to the numerically inverted form (right) of the IAPWS equations.

I confirmed that my chosen EOS approaches the TFD (the high-pressure limit EOS) as pressure increases. Figure 2.7 shows this. I made no attempt to match or otherwise smooth the boundaries between each region, trusting the published equations of state as being accurate within their range of validity. This leads to some false density discontinuities which are visible in the final output. I have chosen to leave these as they are, because there is no justification for artificially smoothing out the density profile.

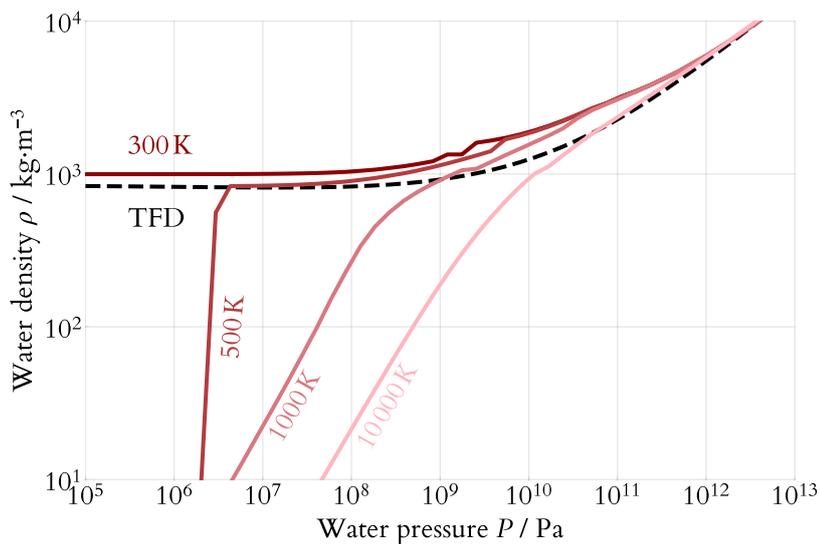


Figure 2.7: Comparison of my equation of state with the high-pressure limit. The TFD (Thomas-Fermi-Dirac) equation of state is increasingly accurate in the high-pressure limit, where temperature effects on the water density disappear. I also show temperature contours of my water equation of state. The TFD, which has no temperature correction, is a poor approximation of the behaviour of water at low pressures, especially across the liquid-vapour phase boundary (e.g. the near-vertical line at 500 K). But all other choices of equation of state approach the TFD at high pressures, and so it is appropriate in the TPa region and beyond.

I can generate a final table of the EOS at any resolution, because it uses functional forms or interpolation to give the output density for a given pressure and temperature. In generating a final grid of densities, I

used logarithmic spacing on both axes to reflect the fact that the density surface has its interesting features at lower temperatures and pressures.

A representation of the final water EOS is shown in figure 2.8.

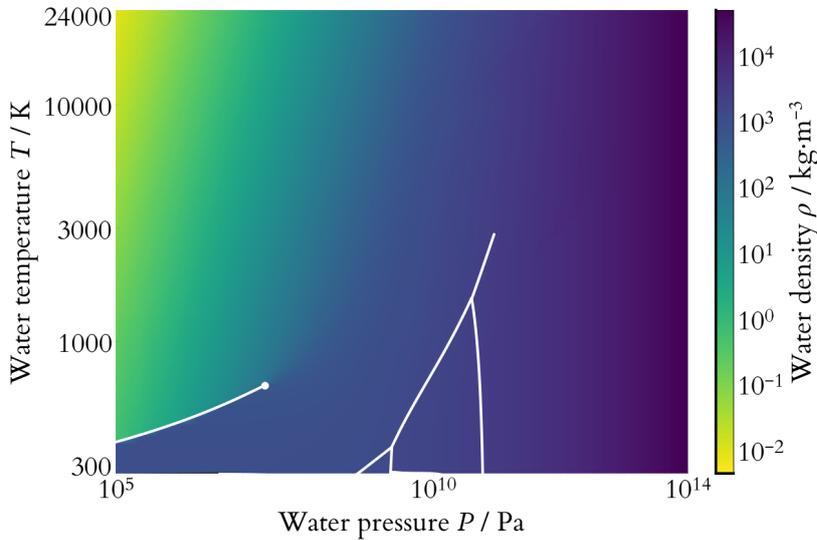


Figure 2.8: Density of my final water equation of state. Here I show the density variation across the entire pressure–temperature range. The density of water is more strongly affected by pressure across the range I consider, but temperature also affects its density too, especially across the liquid–vapour phase boundary and in the supercritical region.

2.2.3 *The phase structure of water*

The key phases of water relevant to the interiors of watery planets include liquid, supercritical and superionic fluid, ice VII, ice X, and plasma. I defer further discussion of phase transitions in the equation of state until chapter 5, in which I examine the phase structure of watery planets.

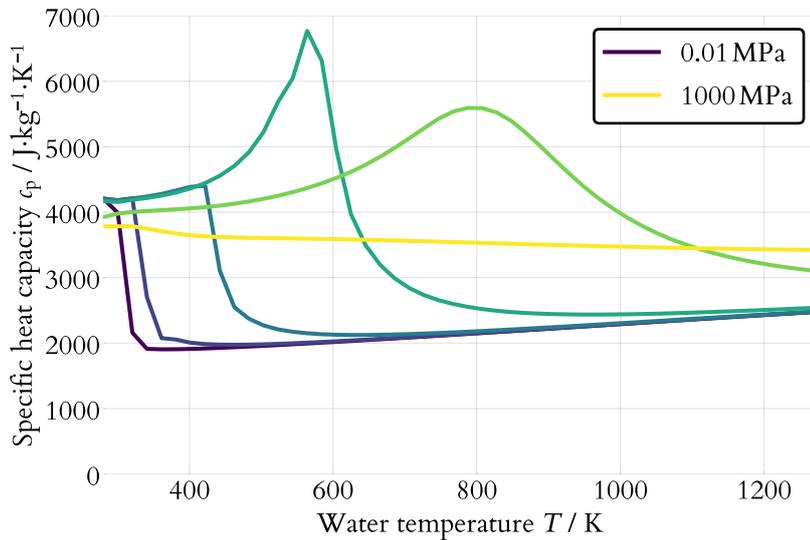
2.2.4 *Thermal effects*

I endeavoured to choose equations of state that were most representative of the thermal behaviour of water across this temperature and pressure domain. I was guided by two principles in doing so. First, we expect thermal expansion effects to approach zero as the pressure increases (figure 2.7): this is a consequence of the equations of state approaching the high-pressure TFD limit. There are significant temperature effects at lower pressures, and it is these effects that I expected to be most important in this study. Second, I aimed for a full treatment of density changes over phase boundaries. Accordingly, I used the phase boundaries specified by Dunaeva et al. (2010) to divide the temperature–pressure phase space into regions corresponding to different phases of water. I then chose an appropriate equation of state to represent each phase.

2.2.5 Thermal expansion and heat capacity of water

Equation 2.6 requires both a heat capacity c_p and a thermal expansion coefficient α . Following my goal of handling temperature effects appropriately, I explicitly sought out temperature-dependent forms for these.

I used the IAPWS tables for heat capacity in the liquid–vapour range, then took the nearest available data point from these tables for all other pressure–temperature points. This is because I could not find readily available heat capacity data across the full range of phases in our equation of state. This approach therefore does not reflect any changes in heat capacity between the high–pressure ice phases. But the most significant effect is the change in heat capacity across the liquid–vapour phase boundary, which is captured appropriately (figure 2.9).



As well as the discontinuities observed in the density of water (figure 2.7), figure 2.9 shows that discontinuities in heat capacity are also a feature of my equation of state. As I aimed for a full treatment of density discontinuities, it is also appropriate that the heat capacity be allowed to be discontinuous too. As we will see in later chapters, the heat capacity is a key parameter in the temperature profile of a waterworld; it therefore benefits us to treat it as variable and not simply fix it to a constant. The heat capacity of liquid water is approximately twice that of both vapour and ice at standard atmospheric pressures and temperatures, and this difference becomes even more pronounced when comparing superheated liquid and vapour (the intermediate green lines in figure 2.9). In that figure we also see that the steep discontinuities

The specific isobaric heat capacity c_p is the heat capacity of water held at a constant pressure.

The volumetric thermal expansion coefficient α , sometimes denoted α_V , is the fractional change in volume as the temperature is increased at constant pressure:

$$\alpha_V = \frac{1}{V} \left(\frac{\partial V}{\partial T} \right)_p = -\frac{1}{\rho} \left(\frac{\partial \rho}{\partial T} \right)_p \quad (2.7)$$

Figure 2.9: The heat capacity across the region of my EOS covered by the IAPWS data set. The heat capacity approaches a constant value near the edge of the domain, indicating that pegging values outside the domain to the nearest known value is likely to be a suitable approximation. Only in the volatile region around the phase boundaries and critical point of water does the heat capacity vary significantly.

between the liquid and vapour phases, a feature of the liquid–vapour transition at lower pressures, give way to a smooth and continuous region of higher heat capacity when pressure is increased.

I drew the thermal expansion coefficient α directly from the equation of state by evaluating equation 2.7. Because the equation of state provides the density ρ as a function of P and T , we can evaluate the thermal expansion at any (P, T) co-ordinate by taking a partial derivative of the density with respect to temperature. I used automatic differentiation²² where possible to evaluate this derivative. In some cases this was not possible²³ so I used finite differencing.²⁴ As well as pre-computing the equation of state itself, I pre-computed and tabulated the thermal expansion coefficient on the same pressure–temperature grid. Some previous works have assumed a fixed thermal expansion coefficient²⁵ but I believe that my approach is more appropriate for understanding how the temperature gradient and physical properties of a watery planet are affected by the thermal properties of water.

²² I used forward-mode automatic differentiation provided by `ForwardDiff.jl`.

²³ The Delaunay triangulation method in the library I used incorporates a method called floating-point filtering, which relies on the specific properties of floating point numbers. It could not be used with the automatic differentiation approach I used, which evaluates functions as usual but replaces the inputs with a special numeric type.

²⁴ I used the package `Calculus.jl`.

²⁵ For example, Ehrenreich et al. (2006) took a fixed value for α in their models.

2.2.6 *Dealing with disparate or conflicting data sources*

I made no attempt to smooth or otherwise interpolate between the different sources of data described above. This approach means that sharp density changes across phase boundaries are well-represented in the final equation of state. This is desirable so that we may examine the differentiation that results solely from phase transitions. It also results in some artificial density discontinuities at the boundaries between different data sets. I believe that this has not affected the results: these discontinuities are minor compared with the density variations within each phase of water and, in most cases, I also bounded the domain of each data set to that of a particular phase.

2.2.7 *A custom interpolation routine for data on a non-uniform grid*

Because I used disparate sources of data, I evaluated the density at a given temperature and pressure in different ways depending on the data source. Although I did not smooth or interpolate *between* data sets, I needed to interpolate some data sources *within* the data set.

Where data were published in tabulated form on a structured rectangular grid, I used simple two-dimensional linear interpolation.²⁶ This interpolation method constructs an interpolating function by taking a weighted average of the density four surrounding points. The weighting depends on the relative (P, T) value of the point being evaluated. In

²⁶ For multidimensional linear interpolation I used `Dierckx.jl`.

this way we can calculate the density at points within the domain of the table but off the rectangular grid.

Not all data can be represented on a structured grid in this way.

This is the case where the equation of state data are provided as sets of (P, T, ρ) values. In this case, we require a different interpolation scheme.

I tested many different interpolation routines written in Python but found them to be very slow and more complicated than necessary. Most of these interpolators use some form of spline-fitting, or perhaps a functional approximation by way of radial basis functions. But I did not require this degree of smoothness in my final interpolation: it was sufficient to instead use the equivalent of linear interpolation on an unstructured grid.

I therefore moved to a custom interpolator written in Julia. Where data were published as unstructured points, I used barycentric interpolation on the mesh of Delaunay triangles defined by these points. Barycentric interpolation can be thought of the triangular analogue of linear interpolation on a regular grid, and Delaunay tessellation is a method for generating a triangular tessellation on a plane between pre-specified points.

To calculate the density at a pressure-temperature point, my interpolator performs the following steps (figure 2.10):

- It takes the points in (P, T) space to define the vertices of a Delaunay tessellation of the plane.²⁷ It then caches the Delaunay mesh created so that it does not have to be re-created on every interpolation.
- It associates each of these (P, T) pairs with its corresponding density value. The result is a weighted Delaunay mesh where each vertex has an attached density value.
- For each interpolation query, it finds the triangle that encloses the (P, T) point of interest. The interpolated value of the density ρ therefore lies between the values at each vertex of this triangle, which we denote (ρ_1, ρ_2, ρ_3) .
- It performs a co-ordinate transform from the Cartesian (P, T) value to a barycentric coordinate system $(\lambda_1, \lambda_2, \lambda_3)$ (where $\lambda_1 + \lambda_2 + \lambda_3 = 1$). In this system, the corners of the triangles correspond to the basis vectors $(1, 0, 0)$, $(0, 1, 0)$ and $(0, 0, 1)$ and the barycentre of the triangle corresponds to the point $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.
- Finally, it calculates the interpolated value: $\rho(P, T) = \sum_{i=1}^3 \rho_i \lambda_i$.

²⁷ To construct a Delaunay tessellation I used the Julia package `VoronoiDelaunay.jl`.

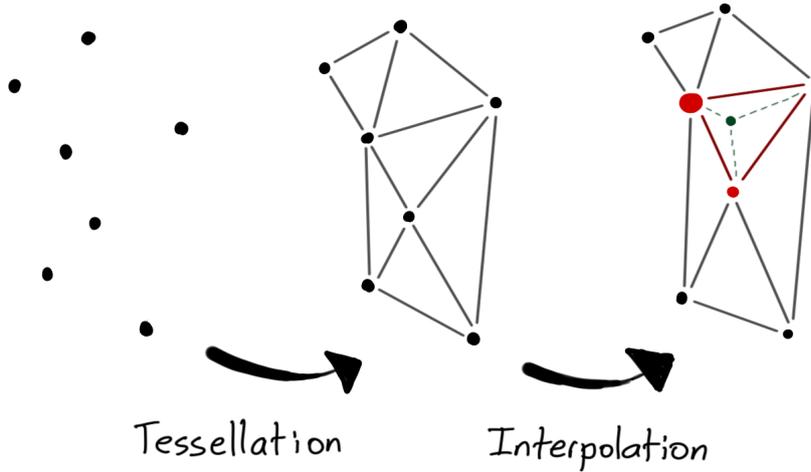


Figure 2.10: This schematic shows my linear interpolation scheme for unstructured data. First the unstructured (P, T) points are used to define a Delaunay tessellation where the points lie on the vertices of the tessellating triangles. Then the density at any given point can be found by taking a weighted sum of the density values at the enclosing vertices.

This custom routine takes advantage of a method called “floating point filtering”, which is also used in some cosmological simulations.²⁸

The other benefit of this approach is that it naturally defines the domain of the table of interest: only points which are contained in one of the Delaunay triangles lie within the domain of the table. I used this Delaunay mesh to determine if a given (P, T) pair lay within the domain of a particular equation of state, allowing me to fall back to another equation of state if necessary.

Finally, I evaluated functional forms of the equation of state as is, defining their domain by means of a bounding box or a polygon in (P, T) space taken from the phase boundaries of [Dunaeva et al. \(2010\)](#).

²⁸ For example, the Illustris simulation uses this method.

2.3 The final equation of state

Some of the equations of state used in this final synthesized version were much simpler than others. This meant that the evaluation time varied from point to point, from very quick table lookups and interpolation to the slower IAPWS formulae. In addition, any equation of state that was specified in the inverse form $P = P(\rho, T)$ needed to be numerically inverted to give the canonical form $\rho = \rho(P, T)$ used in my final EOS. To avoid duplicating this calculation unnecessarily, I re-sampled the final equation of state on to a 500 by 500 pressure–temperature grid. Pre-computing and tabulating the data in this way saved significant time.

In my trials, the resolution of the grid barely altered the properties of the planetary models. I checked that altering the number of points has a negligible effect on the final planet structure. This suggests that

the density behaviour within a single phase region was more important than any effects across phase boundaries that might be lost by sampling from this discrete grid.

2.3.1 *Uncertainty*

The equation of state that I used necessarily has some uncertainty in it, especially in regions near the critical point of water²⁹ and at high temperatures and pressures where there are sometimes conflicting experimental and theoretical data.³⁰ The error in the equation of state varies depending on the original data source. For the region encompassed by the IAPWS data,³¹ the density uncertainty is approximately 0.01% (liquid and solid), 0.03 to 0.1% (vapour), and up to 0.5% in the region around and beyond the critical point. [Wagner & Pruß \(2002\)](#) give a more detailed breakdown of these errors in their section 6.3.2, in particular figure 6.1 in that work. I estimate that the error beyond these regions is closer to 1% if we extrapolate beyond the table and assume that the uncertainty continues to increase at higher temperatures and pressures. For the supercritical fluid, plasma and superionic phases in the data of [French et al. \(2009\)](#), they state that “the QMD EOS is accurate up to 1% for the conditions relevant for the giant planet’s interiors of our solar system.” For the ice VII phase, the measurements of [Sugimura et al. \(2010\)](#) have errors of between 0.003% and 0.5%. Finally, it is not possible to give a meaningful uncertainty estimate at higher pressures where no measurements exist, but I do not treat the temperature dependence there anyway.

²⁹ [Wagner & Pruß \(2002\)](#)

³⁰ [Baraffe et al. \(2008\)](#)

³¹ [Wagner & Pruß \(2002\)](#)

2.3.2 *How my EOS improves on previous work*

The final software package, `WaterData.jl`, has the following advantages.

- It spans a wide range of pressure–temperature space relevant for the structures of planetary interiors. Most other equations of state are intended for use only in an appropriate limit, or contain no temperature dependence, or are only applicable for colder planets.
- It is based on the latest available experimental data, and it uses temperature–dependent formulations for the pressure–temperature–density relation where possible.
- It comprises both functional and tabular data sources and unifies them under a single interface.
- Its structure is modular: it is easy to swap out part of the represen-

tation of the EOS or to add a new region of phase space not yet covered. For example, if new high-pressure experiments produce results for the heat capacity and density of water in the ice X phase, they can easily be added to the EOS.

- It provides tools to export the data in a raw format or to interpolate and grid the data.
- It contains routines for numerical inversion with sensible defaults, which facilitate the addition of functional equations of state.
- It provides heat capacity information in the vapour/gas phase covered by the IAPWS formulation, which is the region of phase space with the greatest heat capacity variation, and extrapolates the heat capacity sensibly outside this region.

2.4 *Making the data freely available*

I have made the equation of state freely available online.³² This chapter was also published as part of our paper “In hot water: effects of temperature-dependent interiors on the radii of water-rich super-Earths”.³³

³² <http://www.github.com/swt30/WaterData.jl>

³³ Thomas & Madhusudhan (2016)

3 *Watery planet interiors*

Having created an improved equation of state for water, the next step was to build structural models of watery planets and ask the question: does temperature matter?

In this chapter I develop the framework to apply my equation of state to interior structure models. After outlining the theory of exoplanet interior structural models, I present my code, `ONION`, which solves the system of structural equations for a one-dimensional planet. Then I show the results of building models with this code. First I examine the effects of my equation of state on homogeneous isothermal spheres. Next I add in temperature dependence and model the interiors of inhomogeneous layered planets. Finally I discuss how, and under what circumstances, my equation of state produces results that differ from other models.

3.1 *Theory*

Guided by the motivation presented in the introduction, I constructed models of heated watery planets with the aim of taking full account of the thermal effects in these models.

By *thermal effects* I mean several things in particular. First, I mean the thermal expansion of a heated water layer on the surface or within the planet's interior: this contrasts with those models that treat the planets as cold spheres. Second, I mean any temperature gradient established within the planet: this is in contrast to the isothermal case. And third, we expect phase transitions within the water layer if the temperature and pressure cross one of the boundaries between different phases of water seen in figure 1.12. This is treated later in chapter 5.

3.1.1 *Model assumptions*

In constructing planetary models, I have assumed the following.

First, the planet is spherically symmetric. This reduces the struc-

tural equations to one dimension. It is possible to construct higher-dimensional models of planets. In fact this is particularly useful when investigating atmospheric behaviour: full three-dimensional global circulation models can reveal interesting details about the behaviour of the atmospheres of large planets, such as the existence of zonal winds or hot spots, that one-dimensional models will miss.¹ But such models are beyond the scope of this work.

¹ e.g. Kataria et al. (2016)

Second, the planet is in a state of hydrostatic equilibrium. That is, pressure and gravity are balanced throughout, allowing us to use the equation of hydrostatic equilibrium as a structural equation.

Third, the planet is in thermal equilibrium or quasi-equilibrium. By this I mean simply that the temperature equations are time-independent: there is no time evolution of the structure, and both external and internal heating are assumed to be constant. Although the planet may still be generating some internal energy due to gravitational contraction, the cooling time is assumed to be very long. Similarly, any radiogenic heating is assumed to be produced by elements with a very long half-life.

Finally, the planet is non-rotating and non-magnetic. These situations only complicate the models; I have ignored them for the time being.

With these assumptions, the structure of a planetary interior is governed by a linked set of ordinary differential equations.

3.1.2 Structural equations

We first consider the isothermal case where the temperature T is constant throughout the planet. In this case the structural equations are the mass continuity equation,

$$\frac{dr}{dm} = \frac{1}{4\pi r^2 \rho}, \quad (3.1)$$

Here r is the radius of a spherical shell, m is the mass interior to that shell and ρ is the density of the shell.

and the equation of hydrostatic equilibrium,

$$\frac{dP}{dm} = -\frac{Gm}{4\pi r^4}. \quad (3.2)$$

Here P is the pressure at the shell and G is the gravitational constant.

The equation of state,

$$\rho = \rho(P, T), \quad (3.3)$$

is also used to calculate the density of the material in question from its pressure and temperature.

Together, these equations define a structural model: a linked set of

ordinary differential equations. To generate a model of a homogeneous isothermal planet, it is sufficient to solve this system of equations. But this is not the whole story. We require a temperature treatment for the planet to take advantage of the improved temperature-dependent equation of state.

3.1.3 The temperature gradient

Previous models of super-Earth interiors have treated temperature gradients within the planet in a number of ways. From simple to complex, these include:

1. Isothermal models, which use only the equations above and an isothermal equation of state of the form $\rho = \rho(P)$.²
2. Simple temperature prescriptions:
 - a) One may assume a fixed temperature–pressure relation $T = T(P)$ so as to reduce the equation of state to the form $\rho = \rho(P)$ and then use the equations above.³
 - b) Or one may choose a temperature profile $T = T(r)$ for the planet (perhaps scaled appropriately to an internal or external boundary temperature) and then use the equations above.
3. An adiabatic or conductive temperature gradient or some combination of the two.⁴
4. A full treatment, which adds an energy transport equation to the equations above then self-consistently solves this with a prescription for luminosity.⁵

I did not explicitly handle energy transport in the manner of the fourth option. Instead I chose the third approach and assumed an adiabatic (isentropic) temperature gradient throughout the planet. The equation for the adiabatic temperature gradient⁶ is

$$\frac{dT}{dr} = -\frac{T\alpha g}{c_P}. \quad (3.4)$$

My sources for the thermal coefficients α and c_P are detailed in section 2.2.5.

Equation 3.4 combined with equation 3.1 gives the temperature gradient in terms of the mass co-ordinate,

$$\frac{dT}{dm} = -\frac{T\alpha Gm}{4\pi r^4 \rho c_P}. \quad (3.5)$$

² This is the approach taken by Seager et al. (2007).

³ For example, Zeng & Sasselov (2013) chose the melting curve of water for this purpose.

⁴ For example, Valencia et al. (2010) used a convective interior with conductive boundary layers.

⁵ For example, Wagner et al. (2011) modelled an adiabatic core underneath a radiogenically heated mantle.

⁶ Milone & Wilson (2014)
Here $g = Gm/r^2$ is the gravity at the shell, c_P is the isobaric heat capacity and α is the volumetric thermal expansion coefficient defined in equation 2.7.

Is the assumption of an adiabatic temperature gradient appropriate? In order to answer this, we need to show that convective heat transfer is dominant over conductive and radiative heat transfer within these heated watery super-Earths.

Convection vs conduction We can assess whether a planetary layer is likely to be convective or conductive by calculating its Rayleigh number,

$$Ra = \frac{g\alpha\Delta Th^3}{\nu\kappa}, \quad (3.6)$$

which is the ratio of buoyant and viscous drag forces (Milone & Wilson, 2014). Rayleigh numbers over the critical threshold for convection Ra_c , typically taken as 1000–2000, indicate that buoyancy forces dominate and so the layer will convect.

I will not be assessing the above equation for every model constructed in this dissertation, and an order-of-magnitude calculation will quickly show why: the Rayleigh number of a icy super-Earth far exceeds the critical threshold for convection. Taking as a representative example the case of an icy planet with a similar temperature structure to the Earth, here is a rough sketch of the Rayleigh number of its water layer.

- For an Earth-sized and Earth-mass planet, the acceleration due to gravity g is $\sim 10 \text{ m}\cdot\text{s}^{-2}$.
- Ice VII is a good representative choice for the high-pressure phase of water inside a solid Earth-sized planet. The volumetric thermal expansion coefficient α of ice VII is $\sim 10^{-4} \text{ K}^{-1}$.⁷
- Assume that our representative waterworld has a thick ice layer similar to the Earth’s mantle. Then the temperature difference ΔT is $\sim 3000 \text{ K}$ across a characteristic layer height h of $\sim 3 \times 10^6 \text{ m}$.
- The coefficient of kinematic viscosity ν can also be expressed as $\frac{\eta}{\rho}$ where η is called the dynamic viscosity and ρ is the density.⁸ In figure 2.1 we see that a representative density of ice is 10^3 – 10^4 Pa over the GPa–TPa range; I take $\rho \sim 10^3$ to account for the planet’s increased temperature over the zero-temperature EOS shown in that figure. Then taking the dynamic viscosity of ice VII as $10^{14} \text{ Pa}\cdot\text{s}$,⁹ this gives $\nu \sim 10^{11} \text{ m}^2\cdot\text{s}^{-1}$.
- The thermal conductivity of ice VII is $\sim 10 \text{ W}\cdot\text{m}^{-1}\cdot\text{K}^{-1}$.¹⁰
- Putting this together with equation 3.6 we find that $Ra \sim 10^8$, five orders of magnitude higher than the critical Rayleigh number of 10^3 .

Here, g is the acceleration due to gravity, α is the volumetric thermal expansion coefficient, ΔT is the temperature differential across the layer, h is the layer thickness, ν is the coefficient of kinematic viscosity and κ is the thermal diffusivity. This is the Rayleigh number for a fluid layer heated from below; later in chapter 4 we will indeed assume that internal heating comes from deep in the planet and that the water layer is not itself luminous.

⁷ Fei et al. (1993)

⁸ Milone & Wilson (2014)

⁹ Poirier et al. (1981)

¹⁰ Chen et al. (2011)

Convective energy transport therefore almost certainly dominates within the bulk of the icy planet we have just considered. In order for viscosity effects to suppress convection, the kinematic viscosity of ice ν would have to be five orders of magnitude larger.¹¹ Therefore I proceed by assuming that convection holds throughout the planet. Although this is a reasonable assumption, convection could potentially be suppressed if the planet were much more weakly heated or if the ice layer were much thinner. Therefore, in a more detailed model it might also be sensible to allow thin conductive layers at key transition points like at the planet's crust. Such layers would be expected to sustain conductive heat transfer because they are thin and occur at points where the composition of the planet is changing. This is the approach taken by Valencia et al. (2006) when modelling super-Earth and super-Mercury planets and it is inspired by the boundary layers used in Earth models and the “stagnant lid” models for solar system icy bodies.

Phase transitions may also affect convection. As a rule of thumb, exothermic phase transitions¹² intensify convection and endothermic ones suppress it.¹³ The key phase transitions seen in my models—from high-pressure or superionic fluid to ice VI or VII—are exothermic and therefore unlikely to suppress convection.¹⁴

Convection vs radiation In contrast to the above, a single representative order-of-magnitude calculation is not appropriate for showing whether radiation is likely to be a dominant mode of energy transport throughout the entire planet. This is because radiative transfer is almost certain to be dominant at some point in the outer layers of every waterworld: as the planet's solid/liquid interior transitions into a gaseous atmosphere, radiation becomes a viable method of carrying increasing amounts of energy. Though this chapter does not yet treat the atmosphere of the planet, we will see in chapter 4 that a radiative temperature gradient is needed to represent the atmosphere. We will therefore treat this radiative atmosphere appropriately when we come to it.

For the icy planetary interior models in this chapter, however, radiation is unimportant to the energy transport of the planet. This is by analogy with the Earth, where only long-wavelength (≥ 1 km) radio waves can propagate, transporting negligible amounts of heat compared to conduction and convection.¹⁵ We expect similar behaviour to hold with ice/liquid waterworlds because these materials are effectively opaque to shorter-wavelength radiation.

¹¹ Within the icy moons of Jupiter and beyond, kinematic viscosity decreases with depth (that is, with increasing temperature and pressure) and only varies by about an order of magnitude anyway (Hussmann et al., 2006; Nimmo, 2018).

¹² Here these are defined for the transition from the less dense phase i.e. down into the planet.

¹³ Milone & Wilson (2014)

¹⁴ See chapter 5 for further discussion of phase transitions within my models.

¹⁵ Milone & Wilson (2014)

3.2 Models

Together, equations 3.1 to 3.3 and 3.5 define a structural model: three ordinary differential equations and an equation of state linking pressure, temperature and density.

The choice of how to solve this system depends on one's aim. A common approach has been to treat it as a boundary value problem; that is, to integrate the structural equations from initial conditions at the surface or centre of the planet. For example, Seager et al. (2007) approached the isothermal problem (equations 3.1 to 3.3 only) from the inside out, choosing appropriate central pressures at the ($r = 0, m = 0$) boundary and building their models outward from there.

3.2.1 Solving the boundary value problem

Boundary value problems are common in stellar structure codes and can be solved in several ways. The approach of Seager et al. (2007) just described is called a *shooting method*.¹⁶ Building on this idea, the method of *double shooting* or *shooting to a fitting point* integrates both inwards and outwards from assumed boundary conditions at the core and exterior of the planet. The solutions must match where they meet, and so after each trial solution the boundary conditions are adjusted based on the discrepancy.

This double shooting method is increasingly slow to converge as the number of structural equations increase. The correct values for the inner boundary (at the centre of the planet) must be arrived at iteratively based on an initial guess. With more structural equations, the iterative procedure must span more dimensions, so performing gradient descent to find the correct solution is slower. It is therefore only useful for simple models, and more sophisticated approaches may be needed.

The relaxation method is one such approach. I mentioned above the idea of shooting to a fitting point. We could also imagine increasing the number of fitting points, using two or three or more points at which the solution must match, and supplying a predicted value at each of these. If we have a pre-existing structure model on which to base the values at these fitting points, then our initial guesses will be relatively good and we can attempt to iterate toward a solution that satisfies the matching condition at each of these points. Taking this idea to its extreme, we might then use every point in the discretised mesh as a fitting point. This transforms the problem: instead of solving a coupled

¹⁶ For a general implementation of a shooting method, see chapter 18 of Press (2007).

system of ordinary differential equations, we must now solve a much larger system of non-differential equations. We can use standard linear algebra techniques to do this, and we can also use a shooting method to provide the initial guess for the structure. The relaxation method comes into its own when we want to perform evolutionary calculations because we can use the previous structural model as the initial guess.

I INVESTIGATED SEVERAL boundary value problem (BVP) and more general ordinary differential equation (ODE) solvers to handle the problem defined by equations 3.1 and 3.2. There are several reasons to desire a generic method of solving these equations. Generic ODE solvers are well-tested, they are optimised for performance, and using one would save development time.¹⁷

To handle this problem, any ODE solver needs to meet the following requirements:

- It needed to solve boundary value problems with systems of equations.
- It needed to handle the numerical singularity at ($r = 0, m = 0$) where the denominators of equations 3.1, 3.2 and 3.5 approach 0.
- It needed to be able to switch between different equations of state depending on the planet's structure.¹⁸

Most of the solvers I evaluated fell short in one of two ways: either they could not handle the singularity at the centre of the core where $r = 0, m = 0$, or they could not handle the transition between equations of state at layer boundaries. I also desired to retain some control over the solver itself rather than treating the solution of the ODE as a black box. The differential equation solver I found that best met my needs was `DifferentialEquations.jl`.¹⁹ Working from this differential equation solver as a base, I developed a package that solves systems of equations for planetary structures.

3.2.2 *A boundary value solver in Julia*

My solver is called `ONION`.²⁰ I prototyped it in Python and later ported it to Julia, a new scientific programming language that offers much better numerical performance.

I used the shooting method, described above, to solve for the planet's structure. At first I was concerned that this method would prove insufficient for the coupled system of ODEs due to the numerical singularity

¹⁷ Perhaps even more importantly, they are likely written by someone more competent at scientific programming than us.

¹⁸ This is sometimes referred to as the solver having support for events. An “event” is any situation in which the system of equations is altered, such as the transition between layers.

¹⁹ Rackauckas & Nie (2017)

²⁰ Because planets, like onions and ogres, have layers.

at the centre of the planet. I trialled both single and double shooting approaches, but the single shooting method proved to work acceptably. In fact, the singularity proved to be a useful signal for failed convergence (see section 3.2.3 below).

Compared to Seager et al. (2007), who integrate from the centre of the planet outwards, my code integrates from the outside in.²¹ This has the advantage of allowing us to specify the surface temperature and pressure as boundary conditions rather than specifying a central pressure and temperature. These surface boundary conditions are more closely linked to observable parameters than the boundary conditions at the core.

I used a Lagrangian system, where the mass interior to a given shell is the independent variable; this is already reflected in equations 3.1, 3.2 and 3.5, which are written in the form $dx/dm = \dots$. It is in contrast to the Eulerian co-ordinate system used by Seager et al. (2007), who take the radius r as the independent variable. The Lagrangian formulation is more stable under numerical integration;²² I also found it more convenient to be able to specify differentiated planets in terms of mass fractions rather than radial distances.

²¹ Several others take this approach too (e.g. Rogers & Seager, 2010a; Madhusudhan et al., 2012)

²² Rogers (2012)

3.2.3 Conditions for convergence

The shooting method uses a series of trial solutions. For the initial trial solution, I specify the surface boundary conditions: total planetary mass M , surface pressure $P(M)$, and surface temperature $T(M)$. I further allow for multi-layer planets by specifying the composition and mass fractions of each layer, $\{x_i\}$. I also specify a search bracket for the radius $[R_1(M), R_2(M)]$. The code uses a Runge–Kutta integrator²³ to solve the system of differential equations above.

²³ Tsitouras (2011)

Specifying the parameters $\{M, R, \{x_i\}\}$ gives an overdetermined system. The trial solution will therefore fail to meet the inner boundary condition, which is that $r = 0$ where $m = 0$. After creating one trial solution by integrating the system of ODEs from the outside, we need to adjust one of the parameters to approach a self-consistent solution.

I therefore chose to iterate the initial radius R until I obtain a value which is consistent with the given M and $\{x_i\}$. In each successive trial, my code iteratively adjusts the radius boundary condition $R(M)$, using a bisection root-finding method to ensure that the radius approaches zero as the mass approaches zero. In this way it converges on the correct value for the total radius, giving a self-consistent solution. I

further required that r remains positive: this avoids any numerical difficulty arising from the behaviour of the equations at $r = 0$. Figure 3.1 summarises this process.

I was unable to use a faster method, such as the secant method, to solve for the correct radius. This is because the equations break down at the very centre and are not valid for $r < 0$ or $m < 0$, so I am unable to integrate past the central point and use this information to derive a gradient to correct the next iteration. Instead, I terminate the integration if $r < 0$ before $m = 0$, using this as a signal that I have chosen an initial radius that is too small—that is, we have hit the centre before we have accounted for all the mass in the model. I then stop the integration, bisect the search region, and try again. In contrast, if the initial radius is too large, we will run out of mass before we reach $r = 0$, at which point I also terminate and refine the initial guess. I consider the model to have converged successfully if the final value of r is less than 1000 m. Figure 3.1 summarises how the solver works.

It is trivial to use this method to solve for another parameter instead of the radius. Because my intention was to investigate the change in observable parameters, I choose to leave the planet’s radius as the free parameter. But it is also possible to leave another parameter free, such as the core mass, layer thickness or even the composition. For an example of this, see chapter 6, in which my models are used to analyse the potential composition of a water-rich super-Earth.

3.2.4 *Mass–radius relations for differentiated planets*

I used my models to produce mass–radius relations for homogeneous spheres of water as well as differentiated multi-layer models. I did this first for the homogeneous isothermal case (in the manner of [Zapolsky & Salpeter, 1969](#)) and then extended my models to include an adiabatic temperature gradient. My differentiated multi-layer models include a water layer on top of a silicate mantle and an iron core. To do this, they treat the equation of state, equation 3.3, as piecewise in the mass co-ordinate. For example, consider a model which has a 5% (by mass) water layer on top of a silicate mantle. For this model, my code begins evaluating equation 3.3 using the water equation of state. It then switches to the silicate equation of state once m , the mass interior to the spherical shell in equations 3.1, 3.2 and 3.5, drops below 95% of the planetary mass. It is possible to choose the integration grid such that this occurs exactly at the end of an integration step. However, in practice a

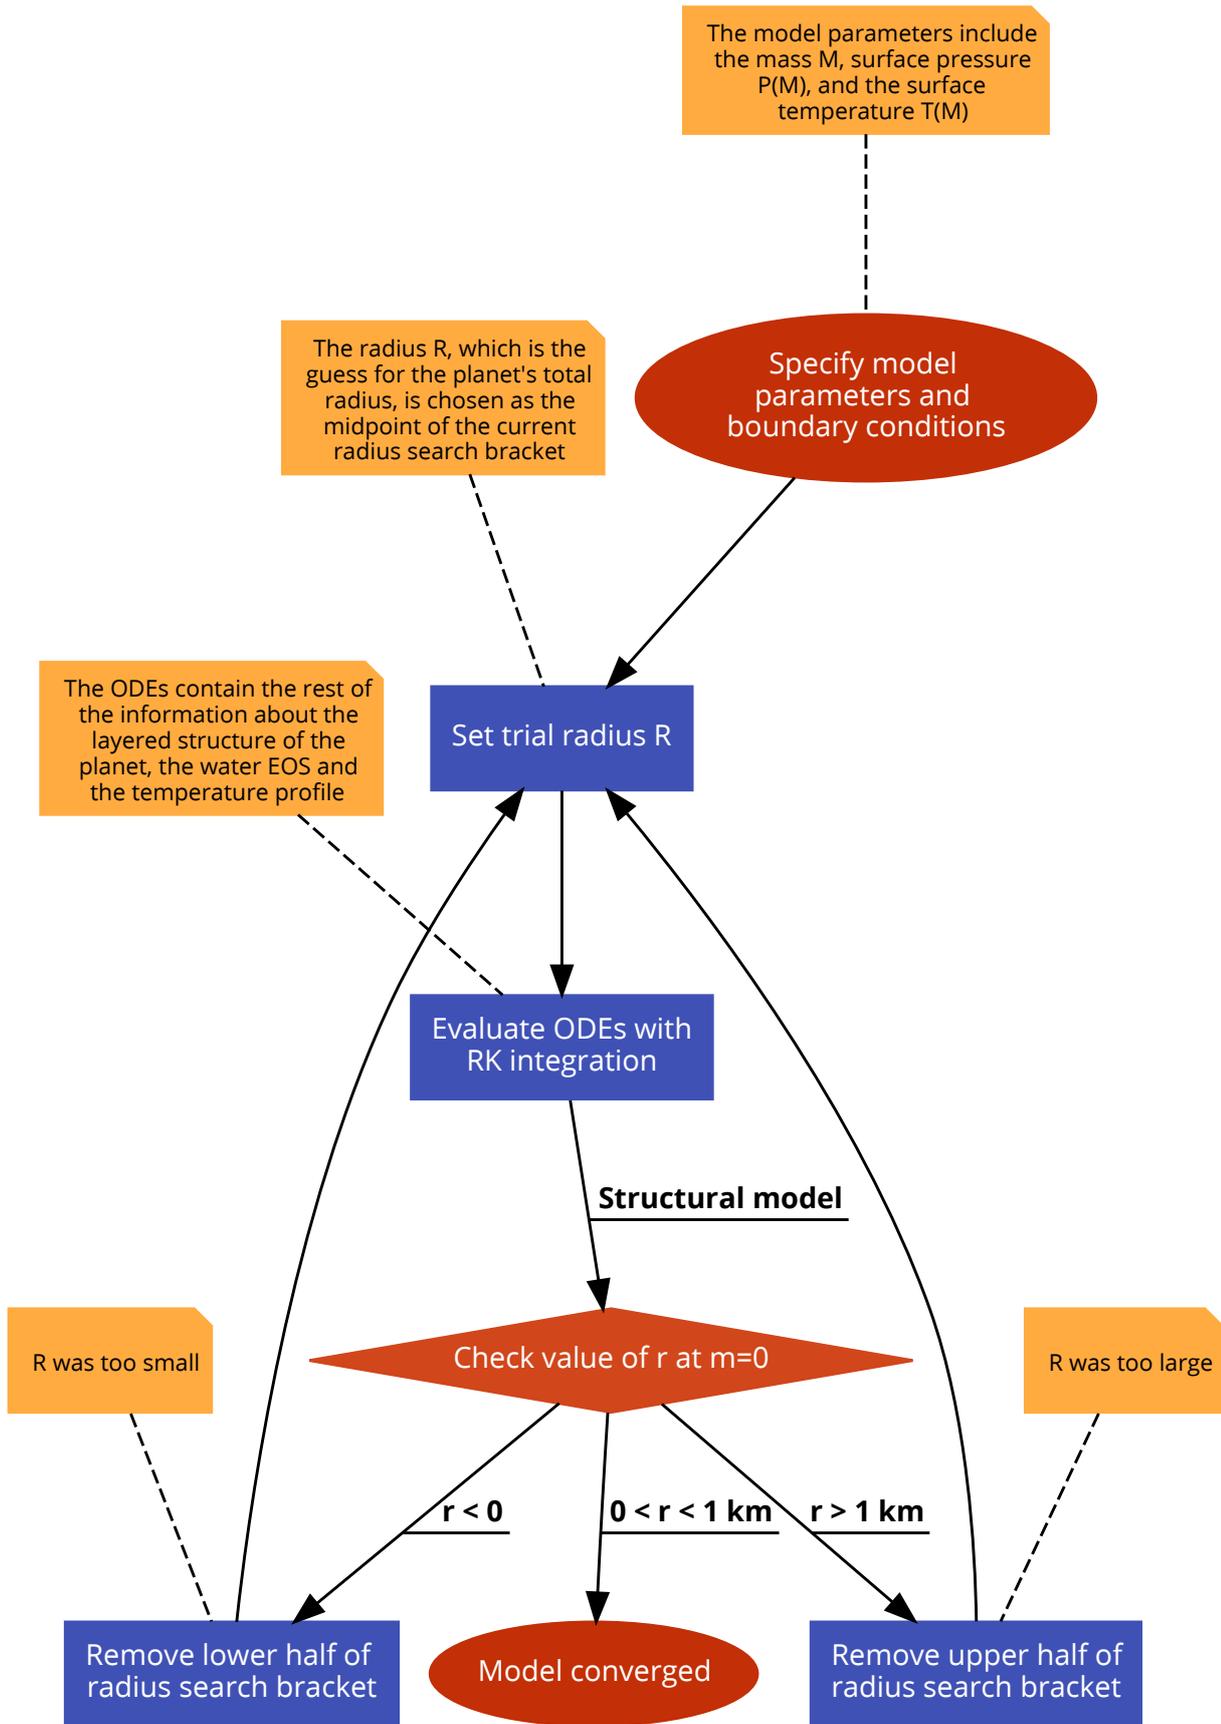


Figure 3.1: A flowchart summarising how my structural solver ONION works.

sufficiently fine grid is also acceptable.

I ignored thermal effects within the iron and silicate layers. The effect of thermal expansion in these solids is thought to be low.²⁴ It would be simple to include the expansion effects of these materials, but I did not collate the equation of state data to enable me to do so. Because I ignored thermal expansion in these layers, I modelled them as isothermal.²⁵

I aimed to accurately capture the density change of water at its phase boundaries. My equation of state for water therefore included its phase transitions, which appear as density discontinuities in pressure–temperature space. When calculating the adiabatic temperature profile, I enforced temperature and pressure continuity at these phase boundaries. In practice, I did this by ensuring that the equation for the adiabatic temperature gradient, equation 3.5, was finite and continuous. This effectively split the adiabatic temperature profile into several different sections, consisting of one separate adiabat for each phase and meeting at the phase boundaries of water. By handling each phase separately, I avoided the numerical difficulty of taking a derivative (equation 2.7) across a density discontinuity. I explain this procedure more in chapter 5 where I consider the phase structure of the final models.

3.2.5 Model verification

Before proceeding further, I first verified these models by making mass–radius diagrams as described in the previous section and comparing them with previous works.

The isothermal case I checked that my models work in the isothermal case by replicating the mass–radius relations of Seager et al. (2007). My models reproduce the mass–radius relations when I constructed homogeneous isothermal 300 K planets using the equations of state specified in their paper, as shown in figure 3.2. I set the surface pressure of my models to zero, following the boundary condition they used.²⁶ This identical mass–radius relation verified that my integrator works correctly, and I therefore began to investigate where the differences lie upon including temperature effects.

The adiabatic case I verified my adiabatic multi-layer models by comparing them with those of Valencia et al. (2007), who constructed similar models using the ice VII equation of state for water (figure 3.3).

²⁴ Seager et al. (2007); Grasset et al. (2009)

²⁵ This follows from setting $\alpha = 0$ in equation 3.5 so that $dT/dm = 0$.

²⁶ The surface pressure hardly affects the results because the equations of state are for the solid phase only.

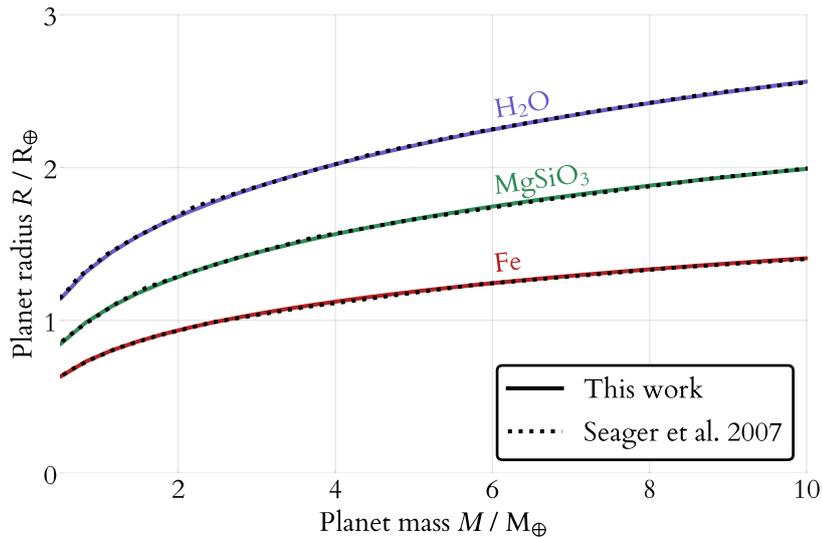


Figure 3.2: Validation of isothermal models. My structural models exactly reproduce previous results in the isothermal case. Here I show mass–radius relations for homogeneous isothermal spheres. If I adopt identical equations of state to those used by Seager et al. (2007), I obtain the same result. This serves as a verification that my code correctly solves the structural equations. These models used zero surface pressure and have no temperature dependence: the equations of state are isothermal and are taken at 300 K.

When I set high surface pressures (10^{10} Pa) I force the surface layer of water to begin as ice VII or close to it and therefore produce a very similar mass–radius relation. However, I predict inflated radii at lower surface pressures and therefore conclude that surface temperature and surface pressure are both important factors for determining the radius of a planet with a water layer.

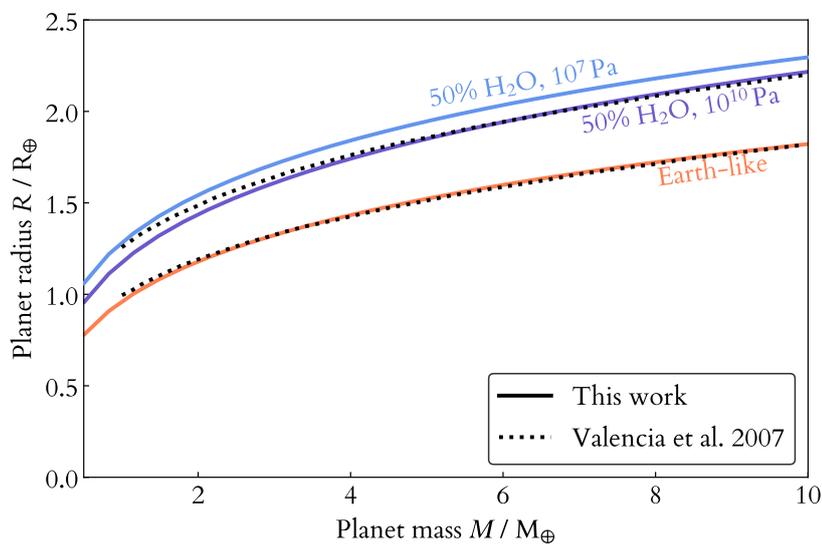


Figure 3.3: Validation of adiabatic models. My mass–radius relations reproduce those for dry planets well, and predict inflated radii for planets with water layers. Here I show mass–radius relations for two classes of models: dry planets (33% Fe and 67% MgSiO_3 by mass), and wet planets (17% Fe, 33% MgSiO_3 , and 50% water). I compared the mass–radius relations with the work of Valencia et al. (2007) who constructed models with ice VII layers. At a surface pressure of 10^{10} Pa the water layer in the wet planets is mostly ice VII and so my results are similar in this case. Small differences are likely due to my different equation of state choice for ice VII. However, at lower surface pressures, water can have an extended lower density shell that results in a larger planet than otherwise expected. The surface temperature in these models is 550 K, matching the characteristic temperature used by Valencia et al. (2007) in their models.

There are minor differences between my mass–radius relations and the mass–radius relations presented by Valencia et al. (2007). I slightly underpredict the radii of lower-mass planets in models with surface pressures of 10^{10} Pa. These differences are likely due to my choice of equation of state: I use only simple isothermal prescriptions for iron and

magnesium silicate and include more phases of water than just ice VII. I also did not include any treatment of conductive boundary layers in my models. In general, however, my results agree well with theirs.

I also compared my results with the evolutionary models of Lopez et al. (2012). I was able to reproduce their mass–radius relation for Earth-like planets, and also when adding extended water layers (figure 3.4). And although I can match the radius of an arbitrary planet by choosing an appropriate surface pressure, I slightly underpredict the radii of small planets compared with their results. This may be a result of different equation of state choices or different temperature gradients during the course of their evolutionary calculations.

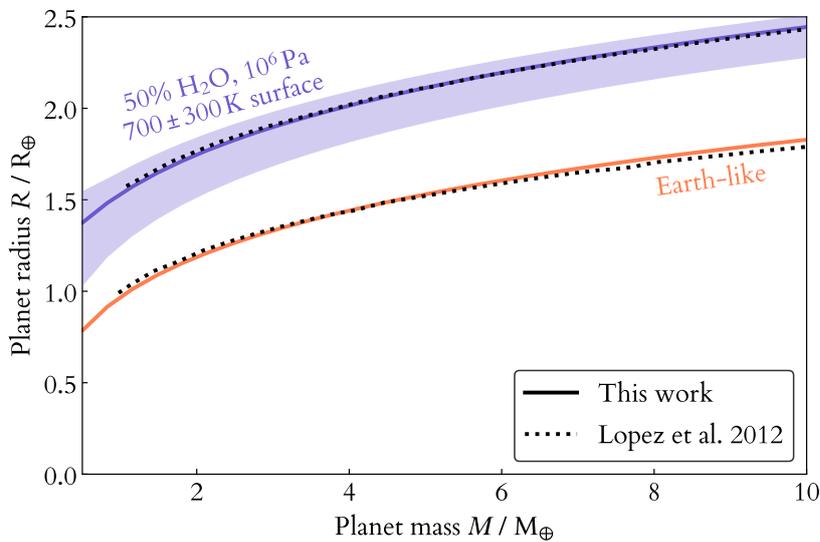


Figure 3.4: Comparison with evolutionary models. Here I plot dry (Earth-like) and wet (50% water on an Earth-ratio core/mantle) mass–radius relations. Shown for comparison are models by Lopez et al. (2012), who build on work by Fortney et al. (2007) and Nettelmann et al. (2011) by using a thermal evolution approach to track the entropy within each planet as it cools. Surface temperature significantly alters the mass–radius relation in my models. The surface temperature in these models is 700 K but the shaded band shows models with surface temperatures from 500 to 900 K, a significant spread, which is caused by temperature-dependent density changes of water at lower pressures. I chose a surface pressure of 10^7 Pa to approximately match the radii of Lopez et al. (2012). Their method does not begin from an explicit surface pressure, as ours does.

Figure 3.4 also provides a first indication of how changes in surface temperature can affect the mass–radius relation. I highlight the magnitude of these differences and note that they are still significant at pressures of 10^6 Pa (10 bar) and up, well into the pressure region where many atmospheric models terminate. I explore the effects on our models of changing surface temperature, surface pressure and composition in the next section.

3.3 Results

Having selected an appropriate temperature–dependent equation of state and built planetary interior structure models that included it, we are now in a position to investigate how this changes the mass–radius relationship of watery planets. I explored the model parameter space, comparing the mass–radius relationships for these watery super-

Earths across a range of surface pressures, surface temperatures and interior compositions. In particular, I investigated the dependence on the following parameters:

1. Planet surface temperature, with the water layer temperature profile taken as
 - a) isothermal, or
 - b) adiabatic.
2. Planet surface pressure.
3. Planet composition, i.e. water mass fraction.

3.3.1 *Effect of surface temperature on isothermal and adiabatic interiors*

I found that thermal expansion can lead to significant changes in the radii of water-rich super-Earths. I constructed super-Earths in two different ways. First I modelled them as isothermal spheres containing an Earth-like core (33% Fe and 67% MgSiO₃) underneath a water layer of 30% of the planet's mass. Then I instead allowed the temperature to increase adiabatically into the water layer.

Figure 3.5 shows that the assumption that thermal expansion effects are negligible, which was made in some previous studies, is not the case. This is true in two senses. First, a significant temperature dependence exists when we adopt an adiabatic interior temperature profile compared with an isothermal one. Second, the surface temperature also affects the radius of a planet within both types of models.

The adiabatic models have a larger radius for a given mass when compared with the isothermal case. This is to be expected: the average temperature is higher along an adiabat than an isotherm fixed at the surface temperature, and the density of water generally decreases with temperature. The increase in radius is significant at higher surface temperatures, as shown in figure 3.5. For example, a 4 M_⊕ 30% water planet with a 600 K surface has a radius of 1.68 R_⊕ if its water layer is isothermal, but 1.79 R_⊕ if it is adiabatic. Across the super-Earth mass range I considered, the adiabatic radii increases by up to 0.11 R_⊕ when compared with the isothermal case. The difference becomes particularly pronounced at higher surface temperatures, at which point the water layer may consist of supercritical fluid rather than liquid, solid, or vapour (figure 1.12).

A significant dependence on surface temperature also exists when using the adiabatic models. That is, changing the surface temperature

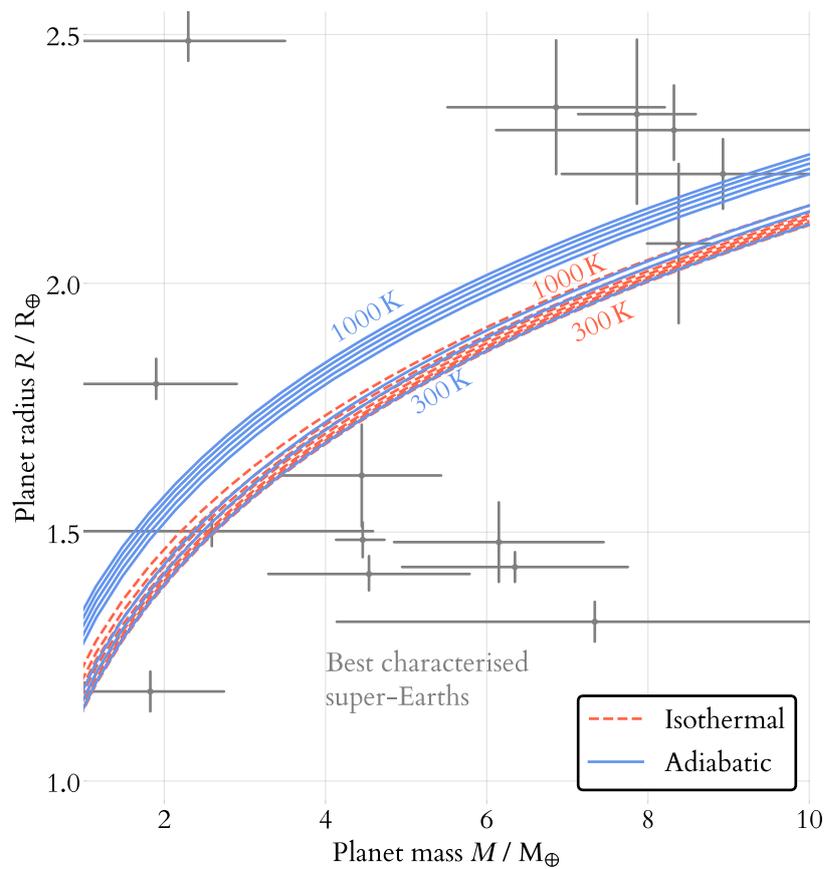


Figure 3.5: Dependence of watery super-Earth radii on surface temperature and internal temperature profile. An increased surface temperature results in an increased planetary radius. This effect is especially pronounced in the full adiabatic temperature treatment. Here I show super-Earths with an Earth-like core under a 30% water layer by mass. I treated the temperature in two different ways: an isothermal treatment with a fixed constant temperature and an adiabatic treatment where I fixed the surface temperature but allowed the temperature to increase inwards according to the adiabatic relation (equation 3.5). The adiabatic models are warmer and therefore significantly larger overall, but even the isothermal planets display some radius change due to temperature. The effects of this temperature dependence are comparable to current uncertainties on measured masses and radii for some of the best-characterised exoplanets. The surface pressure in these models is 10^7 Pa (100 bar), and the temperature increases in steps of 100 K. The larger gap between 500 and 600 K in the adiabatic case is due to a density discontinuity between the liquid and vapour phases.

affects the radius of a model water super-Earth even when its temperature profile is already being treated as adiabatic. In the case of a $10 M_{\oplus}$ planet, increasing the surface temperature from 300 to 1000 K gave a radius increase of $0.15 R_{\oplus}$. For an Earth-mass planet the increase was approximately $0.2 R_{\oplus}$ for the same temperature range.

I have highlighted above the change in the adiabatic models, which I claim are a more realistic representation of the actual temperature structure within the planet. But even the isothermal models show a significant increase in radius with the planet's temperature. For a $10 M_{\oplus}$ planet, the change in radius is $0.04 R_{\oplus}$ from 300 to 1000 K. For a $1 M_{\oplus}$ planet it is $0.1 R_{\oplus}$, nearly 10% of the planet's radius. Because these models are isothermal, this is due to the thermal expansion of the planet as a whole rather than of one small part of the water layer near the surface.

I do not necessarily expect an adiabatic temperature gradient throughout the whole planet because the entire interior may not all be convective. For example, Valencia et al. (2007) included conductive boundary layers in their models. In that case, the true temperature-dependent behaviour of the mass-radius diagram might lie between the adiabatic and isothermal cases. Despite this, figure 3.5 shows that the surface temperature can still play an important role in determining the radius of a planet if it has a substantial water layer. This is true even in the extreme isothermal case where there is no temperature gradient at all within the planet.

These models have a surface pressure of 10^7 Pa (100 bar) so this effect is not due to the strong liquid-vapour transition at 1 bar. In fact, we still see these effects past the critical pressure of water (2.206×10^7 Pa). The critical point, which is visible in figures 1.12 and 2.2, is the point in temperature-pressure space beyond which there is no distinct phase transition from liquid to vapour. This indicates that a liquid-vapour transition is not required to produce a significantly inflated radius when the water layer is heated. I discuss the effect of pressure on these models further in the next section.

3.3.2 *Effect of surface pressure*

The surface pressure can strongly affect the temperature-dependent thickness of the water layer (figure 3.6). For example, at high temperatures (1000 K), increasing the surface pressure of a 30% water and $1 M_{\oplus}$ planet from 1 bar to 1000 bar compresses the water layer signif-

icantly, decreasing the planet's radius from $1.8 R_{\oplus}$ to $1.3 R_{\oplus}$. And at low pressures we see a bifurcation in the temperature–radius contours where a surface temperature increase of 100 K or less can inflate the radius of a watery super–Earth by more than 20%. This is the result of a transition across the liquid–vapour phase boundary, which exists at pressures up to the critical pressure of water (2.206×10^7 Pa). My interior structure code is most likely not the best choice for modelling such a quasi–atmospheric layer: I did not handle radiative energy transfer at this stage. They require a more sophisticated treatment of the temperature profile than the adiabatic assumption in this chapter. See chapter 4 for more detail on the behaviour of these vapour layers.

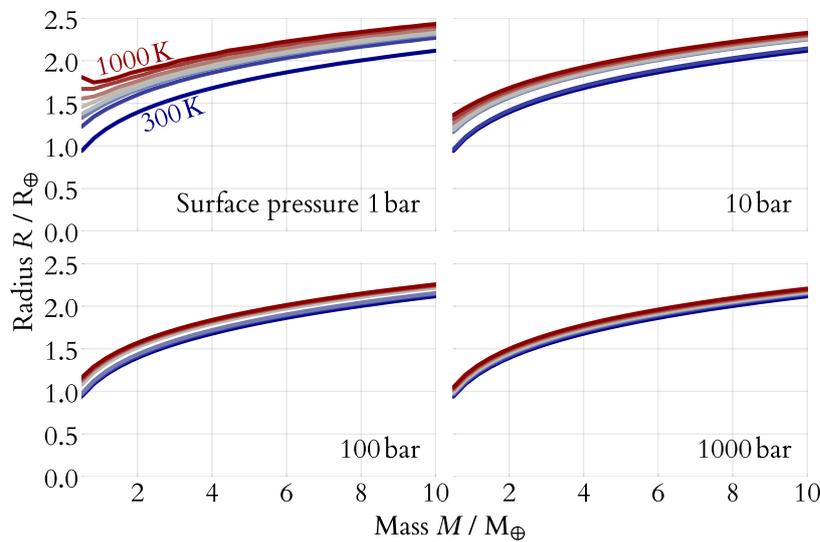


Figure 3.6: Dependence of radii on surface pressure. The effect of temperature on the radius of watery planets decreases with increasing surface pressure, but remains significant (greater than about $0.1 R_{\oplus}$) for pressures below 1000 bar. Here I show mass–radius relations for spheres with an Earth–like core under a 30% water layer, changing only the surface pressure each time. The temperature dependence remains even beyond the critical pressure of water (2.206×10^7 Pa), at which point the surface water exists as a supercritical fluid. Only at very high pressures (10^9 or 10^{10} Pa; 10 000 or 100 000 bar) does this temperature dependence vanish.

Despite observing highly inflated radii when the temperature is increased across the liquid–vapour phase boundary, we still see temperature–dependent variation in the planet's radius past the critical pressure of water. This is because the density of water is still strongly temperature–dependent in the super–critical regime. In fact, we might reasonably expect the same inflated radii in any situation where the pressure of the water layer places it in a region of the water phase diagram that has significant temperature dependence. If the water layer is heated to thousands of Kelvin, this temperature dependence may only begin to disappear around 10^{10} Pa (100 000 bar, figure 2.7). With a surface pressure of 10^8 Pa (1000 bar), a watery super–Earth with a surface temperature of 1000 K still has a radius that is up to $0.1 R_{\oplus}$ larger than one with a surface temperature of 300 K. This is comparable to or greater than the best current uncertainties on measured super–

Earth radii (figure 3.5), and indicates that the surface temperature is a key parameter to consider when one attempts to model planets with significant water mass.

3.3.3 Effect of water content

I find that changing the water content affects the temperature-dependent behaviour discussed in earlier sections (figure 3.7). I constructed planets with water, silicate, and iron layers, fixing the silicate:iron mass ratio to the Earth value of 2:1 and allowing the water shell to vary in mass. These models correspond to an Earth-like nucleus with an extended water layer at the surface.

The effects of surface temperature on radius vary in magnitude across all my models with water layers, but still exist even when I set the water layer mass to just 1% of the mass of the entire planet. For a $1 M_{\oplus}$ super-Earth with a surface pressure of 10^7 Pa (100 bar), the radial change when the surface temperature increases from 300 to 1000 K is $0.3 R_{\oplus}$ (for a 50% water planet) and $0.1 R_{\oplus}$ (for a 1% water planet). This similarity holds across the entire range of planetary masses I considered.

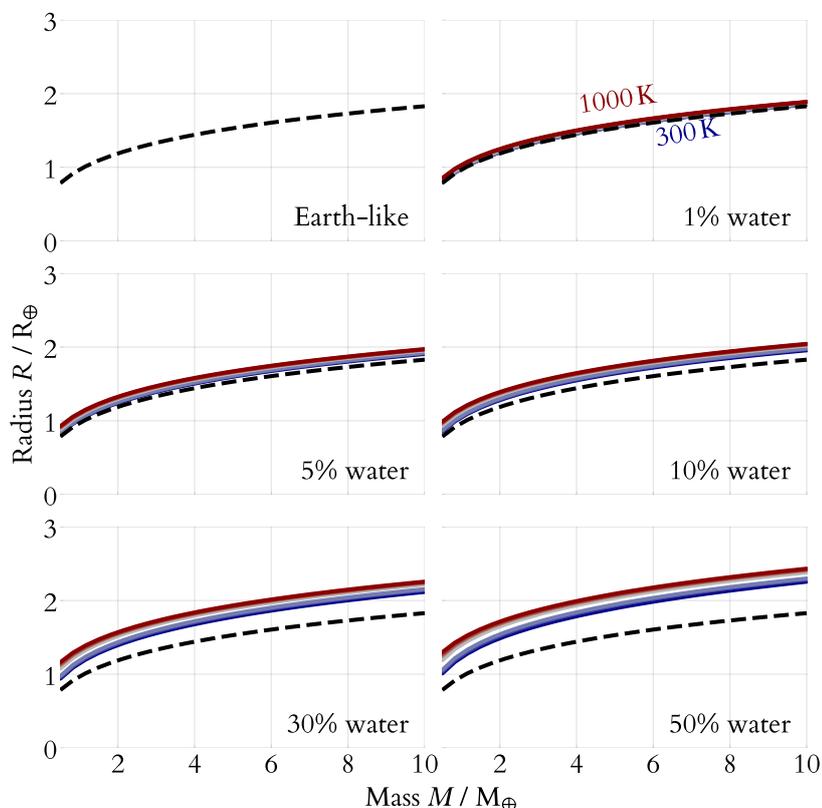


Figure 3.7: Dependence of radii on water mass fraction. Even low-mass water layers result in planets that are strongly affected by temperature changes, especially when water on the surface is hot enough to be in the vapour or supercritical phase. Here I show mass-radius relations for multi-layer planets: an iron core with silicate and (in all but the first panel) water layers. I show the Earth-like iron-silicate core in each panel for comparison. All the watery planets are larger than the dry case owing to the lower density of water. Surface temperature variation affects the radius of a watery planet by a different amount in each case, and it can nearly double the radius for low-mass planets. Because the iron and silicate layers are isothermal, this variation is due solely to temperature effects in the water layer. I fixed the silicate:iron mass ratio at 2:1 and set the surface pressure to 10^7 Pa (100 bar). The temperature contours are in steps of 100 K from 300–1000 K.

3.4 Discussion

In this chapter I have presented planetary interior structure models of water-rich super-Earths. The models incorporate a temperature-dependent water equation of state and use an adiabatic treatment for the temperature gradient. I directly calculated the thermal expansion coefficient α from the equation of state, rather than treating it as a constant, and I used a variable heat capacity based on experimental data. My conclusions are as follows.

First, when one models a solid planet, adding a water layer comes with a substantial thermal dependence. By this I mean that the temperature of the planet may substantially alter the radius of the planet as the water layer expands and contracts. Previous studies have shown that including a temperature gradient in Earth-like planets produces a minimal change in its radius.²⁷ I showed that this assumption no longer holds once large water layers are considered, even setting aside the unrealistic case of a 100% water planet. For example, consider the case of a $2 M_{\oplus}$ planet with an Earth-like core underneath a water layer of 5% of the planet's total mass. If the surface pressure is 10^7 Pa (100 bar), the difference in the planet's radius when the surface is heated from 300 K to 1000 K is approximately $0.3 R_{\oplus}$ (figure 3.7). This effect is on top of any thermal expansion of iron and silicate: my models treated the rocky layers as isothermal. It is also in addition to any uncertainty in the equation of state itself. Such changes in radii are significant considering that current observations can already measure super-Earth radii to precisions better than $0.1 R_{\oplus}$ (e.g. figure 3.5).

The strength of the planet radius–temperature relation also depends on the surface pressure. This is a result of the decreasing thermal expansion of water with pressure: the coefficient of thermal expansion is much smaller in high-pressure ice than in the liquid, vapour, or supercritical fluid phases. At pressures of more than about 10^{10} Pa (100 000 bar) any temperature change becomes irrelevant: the uncertainty in current planetary radius measurements is larger than any conceivable radial change owing to temperature effects, so more precise structural models may not be useful. However, there is still a significant radial dependence on temperature at lower surface pressures. At 10^8 Pa (1000 bar), a watery super-Earth with a surface temperature of 1000 K can be up to $0.1 R_{\oplus}$ larger than one with a surface temperature of 300 K. It is therefore important to include temperature effects in the interior

²⁷ Howe et al. (2014); Grasset et al. (2009); Seager et al. (2007)

models if an accurate radius is required as part of the model.

This pressure dependence manifests itself most strongly below the critical point of water. At pressures below this critical pressure, a phase transition still exists between liquid and vapour. There is therefore a bifurcation in the mass–radius diagram: a small increase in surface temperature causes a larger change in radius as the surface water vaporises. This is visible as a gap in the contours in figures 3.6 and 3.7. I caution that it is likely not appropriate to attempt to treat such vapour layers using the approach in this chapter, which is intended for interior structures.²⁸ However, a lesser version of this effect is still visible at higher pressures.

I therefore conclude that, in some cases, planetary heating may alter the interpretation of a planet’s radius if a water layer is part of the model. This is especially true if the planet consists entirely of water, but this is an unlikely physical scenario. More importantly, the result is still significant even if the surface of the water layer is at moderately high pressures and lies underneath a heavy atmosphere. All that is required for the water layer’s density to change significantly from the isothermal case is for a temperature increase of a few hundred Kelvin. Moreover, even isothermal watery planets have some degree of radial temperature dependence: up to $0.3 R_{\oplus}$ across the mass range of super-Earths and in the temperature range of 300 to 1000 K.

The temperature dependence is also important to take into account in approaches such as that of Kipping et al. (2013) or Madhusudhan & Redfield (2015), where a watery interior model is used to place a lower bound on the atmospheric height of an observed planet. I have shown that the radius of an adiabatic watery planet may be significantly higher than the zero-temperature or isothermal case. Incorporating a surface temperature estimate into this approach should therefore give better constraints.

From an observational perspective, these results are most interesting at intermediate pressures. At low pressures (10^5 Pa or 1 bar) we cannot yet claim that we accurately capture the behaviour of what is now essentially an atmosphere, because we include no prescription for radiative energy transport in our models. At high pressures (10^{10} Pa or 100 000 bar) any temperature dependence in the water equation of state disappears. The physical scenario most relevant for these models is therefore that of a water layer (ocean, ice or supercritical fluid) underneath a thin or moderate atmosphere. Others have already included

²⁸ The next chapter will treat the atmosphere of the planet.

volatile layers on top of interior structure models.²⁹ Adding more complete temperature dependence to the interior portion of these planetary models is a worthwhile future direction if we wish to treat them as water-rich.

I have included no atmospheric layers in these models. Other studies have provided more complete treatments of atmospheric layers. For example, [Rogers & Seager \(2010b\)](#) included a gas layer on top of an isothermal interior structure model in order to interpret the structure of the planet GJ 1214b. And [Valencia et al. \(2013\)](#) used internal structure models coupled with an atmospheric layer, exploring the dependence of radii on various model parameters including equilibrium temperature and water content. Given that we set the surface pressure to between 10^5 and 10^{10} Pa (1 and 100 000 bar), these models must therefore represent the layers interior to an atmosphere of some sort. In the next chapter, I consider how an atmosphere changes these models and ask: how can we link this to the heating of the planet?

²⁹ [Rogers & Seager \(2010b\)](#)

4 *Heating and the atmosphere*

In the previous chapter, I presented the initial results for the interior structures of heated planets and showed that the interior radius can change significantly with heating. But upon reading that chapter it should quickly have become apparent that these models are limited. What about the atmosphere? Isn't it inappropriate to leave the surface pressure as a free parameter? How do we make the link to a planet's observable radius? And instead of setting the surface temperature explicitly, shouldn't we be linking this to the planet's host star and its environment?

This chapter addresses these questions. Here I present an extension to my models that takes into account the planet's atmosphere. I incorporate an analytic two-stream radiative transfer model to give an atmospheric temperature profile. Then I compare the effects of heating via two different modes, internal and external, to assess whether watery planets are preferentially inflated by one mode of heating.

4.1 *Heated planets*

A planet may be heated in several ways, both internal and external.

The Earth's net surface energy flux from internal sources is about 40 TW.¹ Of this, 5 TW is generated within its crust and conducted to the surface. The remaining 35 TW is generated within the interior of the planet and transported convectively. This internal energy generation in planets arises primarily from three sources: the decay of radioactive elements; secular cooling and contraction; and tidal forces.

Tidal heating occurs when a planet is tidally distorted by its gravitational interaction with a host star. For the Earth, tidal heating currently plays little role in internal energy generation. It is generally neglected in thermal calculations. However, rocky exoplanets that have eccentric orbits and pass close to their host star can be much more strongly tidally heated.² This mode of heating is also important for moons, especially

¹ Davies (1980)

² Barnes et al. (2010); Driscoll & Barnes (2015)

those around giant planets that are only weakly illuminated. This could be the dominant form of heating in such moons.³

The decay of radioactive elements in a planet’s interior results in heating. In the Earth this value is about half of the total internal energy generation⁴ while in icy Solar System bodies it tends to be the dominant source of energy.⁵ The so-called “chondritic coincidence”⁶ refers to the fact the Earth’s energy generation rate per unit mass is about the same as that observed in chondrites.

Cooling and contraction supply the remainder of a planet’s internal energy budget. In the Earth, the other half of the internal energy is supplied by cooling.

At the same time, a planet is heated externally by its host star. One measure of this heating is the *equilibrium temperature*, which is calculated by assuming that the planet is in thermal equilibrium and that incident and outgoing radiation are balanced. The equilibrium temperature T_{eq} is

$$T_{\text{eq}} = T_{\star} (1 - a)^{1/4} \sqrt{\frac{R_{\star}}{2D}}. \quad (4.1)$$

In this chapter I will also follow Guillot (2010) by defining the irradiation temperature T_{irr} , which is an effective temperature characterising the irradiation at the substellar point:

$$T_{\text{irr}} = T_{\star} (1 - a)^{1/4} \sqrt{\frac{R_{\star}}{D}}. \quad (4.2)$$

Because the temperature at the substellar point is hotter than the averaged temperature across the entire planet, the irradiation temperature is higher by a factor $\sqrt{2}$ than the equilibrium temperature.⁷

In the previous chapter we saw that the temperature of the planet affects its radius. The planet’s temperature is controlled by both internal and external energy sources, but they may have different effects on the radius. We must therefore ask: which mode of heating is dominant? We know that for the Earth, the balance is in favour of irradiation: internal heating only contributes a small amount to the energy balance of the planet.⁸ But this is not necessarily the case for all planets. For example, half of Jupiter’s energy budget comes from internal heating. Planets with thick opaque atmospheres might trap internal energy more effectively, exaggerating the heating effects of internal sources. And water-rich icy moons could also be heated by tidal interactions while receiving very little irradiation. It is therefore useful to understand whether internal and external heating can significantly contribute to the

³ Ross & Schubert (1989)

⁴ Davies (1980)

⁵ Hussmann et al. (2006)

⁶ Suess & Urey (1956); Urey (1956)

Here T_{\star} is the stellar temperature, a is the planet’s albedo, R_{\star} is the stellar radius and D is the star–planet distance (Burrows & Orton, 2010). For the remainder of this chapter I assume zero albedo ($a = 0$).

⁷ This corresponds to a factor of 4 in incident flux.

⁸ 40 TW of energy is produced internally but more than 150 000 TW is supplied by irradiation (Wengenmayr, 2012).

observed radius of a planet. But incorporating heating also means that we must take into account the contribution from the atmosphere, which was not included in the models in the previous chapter.

4.1.1 *Incorporating the atmosphere*

The models in the previous chapter have surface pressures of tens or hundreds of bars. Madhusudhan & Redfield (2015) discussed planets with water-rich atmospheres, describing the use of measurements both in and out of opacity windows to determine the atmospheric thickness. The pressure to which these measurements probe can be 0.1 bar or lower in regions of high opacity. My interior models therefore begin deeper within the planet, and we must include atmospheric layers if we are to correctly treat the region above that.

In the case of a volatile layer such as water, the line between interior and atmosphere can become blurred. The picture is complicated by atmospheric effects that can increase the opacity. If a cloud layer forms in the atmosphere, the opacity surface may not necessarily be at the same depth or pressure as any solid surface of the planet. Turbidity effects around the critical point may also affect the opacity. It is for this reason that very high-temperature super-Earths are interesting: at higher temperatures, a cloud deck is less likely to occur⁹ and atmospheric measurements are therefore able to probe deeper. The previously mentioned opacity windows may therefore be able to provide a view through the atmosphere to the planet's surface, or at least to a point where the assumption of interior convective mixing is more likely to hold.

⁹ Some high-temperature condensates are possible but in general we expect clouds to be less of a problem at these higher temperatures.

WE SAW EARLIER that heating a planet can moderately inflate its interior. The same will be true for its atmosphere, but to what extent? To answer this question completely, we need atmospheric models. An ideal atmospheric model would include line-by-line opacities, clouds and chemical networks, and it would then be combined with an interior model like the one in the previous chapter. But though a full atmospheric treatment is desirable, it is also difficult: it introduces many free parameters, it is computationally expensive, and ultimately it is beyond the scope of this dissertation. I therefore conducted an initial exploration of this problem with an analytic atmospheric temperature model instead and investigated the effect on the observed radius.

4.1.2 An improved structural model

To achieve the goal of incorporating atmospheric and heating information into the models, I did several things. I extended the equation of state to lower pressures. I added an atmospheric layer to the models. And I treated the temperature gradient within this layer, which necessitated adding information on the optical depth and therefore the opacity of water vapour. Finally, I added realistic boundary conditions to the models that are linked to the physical properties of the planet (its temperature and its transit radius).

My final model (figure 4.1) comprises the interior structure from chapter 3 (an iron core, a silicate mantle, and a water envelope) plus a water atmosphere. Here I outline the full structural treatment, including my updated prescription for the boundary conditions.

Before we continue, some definitions are in order. Normally we would understand the *surface* of a terrestrial planet as being defined by a sharp transition from a thin atmospheric layer to a rocky centre. But for volatile-rich planets like these waterworlds, such a definition is no longer useful. In this chapter I therefore use the following nomenclature to describe the different layers and boundaries in my model:

- The **photospheric surface** of a water-rich planet is defined by the depth below which light can no longer penetrate, which for this work I take to be a transverse optical depth of $\tau_t = 1$. It defines the outer boundary of my models (section 4.1.5). The radius of the planet R_p is therefore the radius to the photosphere.
- The **atmosphere** is the outer portion of the planet that follows a radiative temperature profile below the photosphere.
- The **atmosphere above the photosphere** refers to the very outer layer of the atmosphere. It is not included in the models, which begin at the photosphere, but I later make some assumptions about the atmosphere above the photosphere to derive the model boundary conditions at the photosphere.
- The **radiative-convective boundary** is the transition between the radiative and convective regimes. It defines the base of the atmosphere and the top of the envelope.
- The **envelope** of my waterworld models refers to the water layer at depths greater than that of the radiative-convective boundary. It follows an adiabatic (convective) temperature profile.
- The **mantle** and **core** refer respectively to the silicate and iron layers,

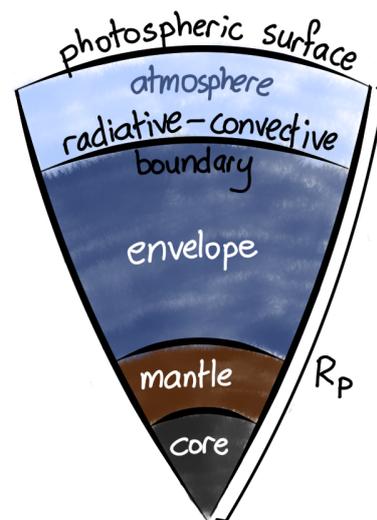


Figure 4.1: Sketch of my final planetary models for this chapter, showing the layered structure. The planetary radius R_p is taken as the radius to the photospheric surface, which is also where the external boundary conditions are defined.

which are taken to have a mass ratio of 2:1.

- The **Earth-like nucleus** refers to the mantle and the core.
- Finally, the **interior** of the planet includes the watery envelope, the mantle, and the core.

4.1.3 Structural equations

In the previous chapter I modelled planets using three structural equations: equations 3.1, 3.2 and 3.5. In this chapter I left the first two of these equations unmodified. They are reproduced below as equations 4.3 and 4.4 for completeness. I then modified equation 3.5 to a more generic form (equation 4.5), allowing for different temperature structures. Finally, I added equation 4.6 to track the optical depth within the atmosphere.

Together, equations 4.3 to 4.6 form a system of ordinary differential equations that I solve in the same manner as before. The equations are:

- the mass continuity equation,

$$\frac{dr}{dm} = \frac{1}{4\pi r^2 \rho}, \quad (4.3)$$

Here r is the radius to a spherical shell, m is the mass enclosed by that shell and ρ is the density of the shell.

- the equation of hydrostatic equilibrium,

$$\frac{dP}{dm} = -\frac{Gm}{4\pi r^4}, \quad (4.4)$$

Here P is the pressure at the shell and G is the gravitational constant.

- the temperature gradient,

$$\frac{dT}{dm} = \nabla, \quad (4.5)$$

Here ∇ describes how the temperature varies inside the planet; it varies depending on the mode of energy transport (see section 4.2).

- and the optical depth relation,

$$\frac{d\tau}{dm} = -\frac{\kappa}{4\pi r^2}. \quad (4.6)$$

Here τ is the radial optical depth and κ is the opacity.

4.1.4 The optical depth and opacity

The fourth integration variable, optical depth, is a measure of attenuation in the atmosphere. Within the convective envelope, I ignore any further increase of the optical depth because it is unimportant for calculating the convective temperature profile.

The optical depth gradient depends directly on the density of the gas (equation 4.6), which as we have seen is related to its temperature and pressure through the equation of state. It also depends indirectly on

The definition of the optical depth is

$$d\tau = -\kappa \rho dr \quad (4.7)$$

where τ is the dimensionless optical depth, κ is the opacity (units of $\text{m}^2 \cdot \text{kg}^{-1}$), ρ is the density and r is the radius. Combining this with equation 4.3 gives equation 4.6.

the gas's temperature and pressure through the opacity $\kappa(P, T)$, which is another state function of the gas.

The meaning of “opacity” can vary from author to author. Here, for consistency with the other authors I have cited in this chapter,¹⁰ I denote the opacity κ to be the quantity with units of $\text{m}^2 \cdot \text{kg}^{-1}$.¹¹ This opacity κ is related to the absorption cross-section of the material by the formula $n\sigma = \kappa\rho$ where n is the number density of absorbers, σ is the absorption cross-section per absorber and ρ is the density of the material. I took the thermal opacity, κ_{th} , as the representative opacity for the atmosphere.¹²

The optical depth relation can take either a constant opacity κ or allow the opacity to vary throughout the atmosphere. Some previous studies have used a constant κ_{th} : for example, the optical depth–temperature curves presented in Guillot (2010) are for a constant opacity. In this study I improve on these by allowing κ_{th} to change throughout the atmosphere, but I also examine the effects of using a constant κ_{th} to see how great the difference is. This choice is likely to be significant because the pressure varies over several orders of magnitude within the atmosphere; the opacity should then also vary. I defer further discussion of the opacity to section 4.3.1, where I describe my source for the opacity data used in this study.

4.1.5 Boundary conditions

A key requirement for building a coupled interior–atmosphere model is the knowledge of realistic boundary conditions. The previous chapter used an arbitrary pressure for the surface of the models, but in reality the pressure is set by the depth beyond which atmospheric measurements can no longer probe. And in principle this pressure could be observationally constrained through spectroscopic observations of the planetary atmosphere, though such observations are currently difficult for super-Earths.

Following the definitions in section 4.1.2, I use the following boundary conditions in my models.

Centre of the planet For the internal boundary condition, I again take $r(m = 0) = 0$ so that the mass enclosed is zero at the physical centre of the planet. In practice I specify this by requiring that $0 \leq r(m = 0) < 1000 \text{ m}$. The final radius is therefore precise to within 1 km.

¹⁰ Rogers & Seager (2010a); Kurosaki et al. (2014); Guillot (2010)

¹¹ It is sometimes called the *mass attenuation coefficient*, as distinct from the *attenuation coefficient* whose units are m^{-1} .

¹² Rogers & Seager (2010a)

Layer boundaries For the transitions between layers of different materials—for example, between water and silicate or silicate and iron—I enforce continuity of the pressure P , temperature T , radius r and mass enclosed m . The optical depth τ is also an integration variable, but its value is unimportant below the radiative–convective boundary because it is not used in the convective region. I therefore do not continue to track the optical depth within the envelope; in practice I achieve this by simply setting $\frac{d\tau}{dm} = 0$ once the integrator reaches the envelope.

Photosphere The external boundary condition is more complicated. Because my models begin at the photosphere, we require a prescription that links the observable radius of the planet to the temperature, pressure and optical depth at that radius. For the external boundary condition I therefore use the following derivation.¹³

I begin with the planet’s mass M_p and radius R_p . I take $r = R_p$ at $m = M_p$. This definition therefore assumes that the atmosphere above the photosphere has negligible mass.

I next assume that the atmosphere above the photosphere is isothermal and that the gravity is constant. In the atmosphere the equation of state is an ideal gas and therefore

$$P = \frac{\rho k_B T}{\mu_{\text{eff}}}. \quad (4.8)$$

I assume that the opacity of the atmosphere can be written as a power law of the form $\kappa = CP^\alpha T^\beta$. Then taking the equation of hydrostatic equilibrium $dP/dr = -\rho g$ and the definition of optical depth¹⁴ $d\tau/dr = -\kappa\rho$ gives expressions for P and τ in the atmosphere above the photosphere that are exponential in the radius r :

$$P(r) = P_R e^{-(r-R_p)/H} \quad (4.9)$$

and

$$\tau(r) = \tau_R e^{-(\alpha+1)(r-R_p)/H} \quad (4.10)$$

where the pressure scale height H is

$$H = \frac{R_p^2 k_B T}{GM_p \mu_{\text{eff}}}. \quad (4.11)$$

It can also be shown that the relationship between P_R and τ_R is

$$P_R = \left(\frac{GM_p (\alpha + 1) \tau_R}{R_p^2 C T^\beta} \right)^{\frac{1}{\alpha+1}}. \quad (4.12)$$

¹³ This follows Rogers & Seager (2010a).

¹⁴ See equation 4.6.

Here P_R and τ_R denote the pressure and optical depth at the planet’s radius i.e. the photospheric pressure and the optical depth to the photosphere.

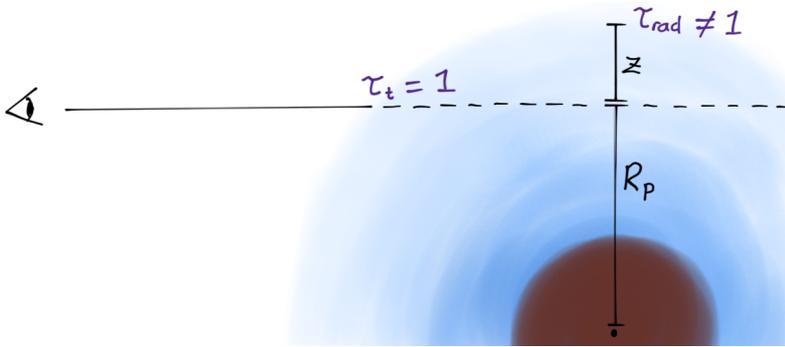
I define the photosphere to be at an optical depth of unity ($\tau_t = 1$). Optical depths lower than this lie above the photosphere; I treat this region as transparent and ignore it when I construct the final model.

We must also make a correction to this optical depth. When an observer views a planet in transit, their line of sight at the photosphere passes through the limb of the atmosphere (figure 4.2). When I define the photosphere to be an optical depth of one, I therefore do so for the *transverse optical depth*¹⁵ τ_t only. But my models are spherically symmetric and require the integration variables to be expressed in a radial form instead. The corresponding *radial optical depth* τ at the photosphere is less than one; the exact value depends on the opacity of the atmosphere. The relationship between these two optical depths¹⁶ is

$$\tau_t \approx \gamma \tau_R \sqrt{\frac{2\pi(\alpha+1)r}{H}} e^{-(\alpha+1)(r-R_p)/H}. \quad (4.13)$$

Setting $\tau_t = 1$ at $r = R_p$ gives the boundary condition on τ :

$$\tau_R = \frac{1}{\gamma} \sqrt{\frac{H}{2\pi(\alpha+1)R_p}}. \quad (4.14)$$



¹⁵ This optical depth is also called a chord or slant optical depth.

¹⁶ I do not include this derivation here but it is from Hansen et al. (2008) and was used by Rogers & Seager (2010a). Figure 4.2 shows the difference between the two optical depths. The relationship is approximate because it is truncated from an exponentially decaying integrand.

Figure 4.2: When an observer looks through the limb of a planet's atmosphere, the optical depth they see at the planet's radius R_p is a *transverse* depth τ_t . In general, this is not the same as the *radial* optical depth, notated as τ in the text and τ_{rad} here for clarity. This schematic diagram shows how the transverse optical depth τ_t through the atmosphere is different from the radial optical depth τ_{rad} down into the atmosphere. Because the atmospheric temperature profile uses a radial optical depth, we need to correct for this discrepancy. The relation between the two is equation 4.13.

Finally, I obtain the temperature at the photosphere simultaneously with the radial optical depth from equation 4.14. The photospheric temperature T_R differs from the irradiation temperature T_{irr} , which is a model parameter. The photospheric temperature therefore needs to be evaluated from the radiative temperature profile, which I will describe in section 4.2. But the temperature profile depends on the optical depth, while the optical depth depends on the scale height and so also on the temperature. I therefore solve simultaneously for T_R and τ_R at the photosphere. I evaluate the temperature at the photosphere from the radiative temperature gradient described in section 4.2. At the same time, I evaluate the optical depth at the photosphere from equation 4.14.

After finding the (T_R, τ_R) pair that satisfies these equations, I set the external pressure boundary condition using equation 4.12.

In almost all other respects, the solver works as described in chapter 3 and shown in figure 3.1. The only additional step is this calculation of the pressure, temperature and optical depth boundary conditions as described above. This process is repeated after each trial solution. In this fashion my models converge on the correct values of P_R , T_R and τ_R as the radius of the planet converges.

4.1.6 The atmosphere–envelope transition

The boundary between the irradiated atmosphere and the adiabatic interior is defined by the Schwarzschild criterion for convective instability.¹⁷ This criterion depends on the thermal behaviour of water and requires assessing whether the radiative or adiabatic temperature gradient is steeper. My code chooses appropriately between convective transport as described in section 3.1.3 and radiative transport as described in section 4.2.1: if $\nabla_{\text{rad}} < \nabla_{\text{ad}}$ then $\nabla = \nabla_{\text{rad}}$, otherwise $\nabla = \nabla_{\text{ad}}$. In practice, this means that the water layer consists of a radiative outer region on top of a convective envelope such that the temperature gradient matches at the boundary. The transition pressure at the top of the envelope, where $\nabla_{\text{rad}} = \nabla_{\text{ad}}$, I denote P_{tr} .

I also investigated how sensitive the final radius was to the choice of this boundary. In section 4.3.1 I assess several fixed values for this transition pressure and show that the choice has a negligible effect on the final radius. In practice, it might therefore be acceptable to choose a fixed pressure for the transition if we were only interested in calculating the final radius.

4.1.7 The updated equation of state

For this work, I added the ideal-gas behaviour of water (equation 2.4) to the equation of state. The domain of my EOS now extends from 1 Pa up to 10^{14} Pa. Previously I had extrapolated upon leaving the edge of the EOS domain; this was not a problem because I started our models at high pressure, but in order to represent the atmosphere appropriately I needed to extend the EOS to low pressures. In order that the atmosphere is treated entirely as an ideal gas, I use the ideal gas EOS for all pressures less than 100 bar.

All other EOS behaviour is as described in chapter 2. In particular, I maintained the same data sources and grid resolution for the base EOS.

¹⁷ A material for which

$$\left| \frac{d \ln T}{d \ln P} \right|_{\text{ad}} < \left| \frac{d \ln T}{d \ln P} \right|_{\text{rad}}$$

is unstable against convection. That is, the convective–radiative boundary occurs at the point where the adiabatic temperature gradient matches the radiative temperature gradient.

Only in the region where $P < 100$ bar does the EOS differ from that chapter.

4.2 The atmospheric treatment

Here I outline the atmospheric temperature profile and how I link it to the internal and external heating. I use three different prescriptions for the temperature gradient ∇ in equation 4.5, depending on the mode of energy transport.

Within the core, where temperature effects are negligible, I set $\nabla = 0$ to make the core isothermal. The edge of the core is specified as a fraction of the planet mass and fixed as a model parameter. When solving the structural equations, I therefore switch to this isothermal treatment once $m < fM_p$, where m is the mass co-ordinate, f is the core fraction and M_p is the total mass of the planet.

For convective transport within the watery envelope, I use the adiabatic temperature gradient from equation 3.5 as explained in the previous chapter.

Finally, for the atmosphere, I compared two different treatments. The first a two-stream grey analytic temperature profile; the second is an isothermal fixed-height atmosphere.

4.2.1 Two-stream temperature profile

I added a radiative temperature profile in the atmosphere by using an analytic two-stream temperature profile. The profile I used is obtained from Guillot (2010), given by

$$T^4(\tau) = \frac{3T_{\text{int}}^4}{4} \left[\frac{2}{3} + \tau \right] + \frac{3T_{\text{irr}}^4}{4} \mu_* \left[\frac{2}{3} + \frac{\mu_*}{\gamma} + \left(\frac{\gamma}{3\mu_*} - \frac{\mu_*}{\gamma} \right) e^{-\gamma\tau/\mu_*} \right] \quad (4.15)$$

where I set $\mu_* = 1/\sqrt{3}$, the appropriate level for isotropic irradiation.

This temperature profile is defined in terms of the radial optical depth τ . Because I use the thermal opacity κ_{th} as the representative opacity for the atmosphere, the optical depth is for thermal radiation (equation 4.7); this matches the definition of Guillot (2010).

There are two ways in which we could incorporate this temperature profile. The first is to remove the temperature as an integration variable within the atmosphere. Instead, the temperature can be calculated directly from the optical depth using equation 4.15. But this approach

Here $T(\tau)$ is the temperature at an optical depth of τ ; T_{int} is the effective internal temperature; T_{irr} is the planet's irradiation temperature; $\mu_* = \cos \theta_*$ is the angle of the normal at the surface to the incident radiation ($\mu_* = 1$ at the substellar point); $\gamma = \frac{\kappa_{\text{th}}}{\kappa_{\text{vis}}}$ is the ratio of the thermal and visual opacities.

is inconsistent with the way the temperature profile is treated in the envelope: there, the temperature is evaluated as the solution of an ODE (equation 4.5). The second option is therefore to instead take the derivative of equation 4.15 and use it as the right-hand-side of equation 4.5.

For convenience, I chose this second approach: the way my code `ONION` is structured means that it is easy to add a new differential equation to the model but harder to add an explicit relation. Although this is more prone to numerical error than evaluating the temperature directly for a given τ , it is no more so than any of the structural equations, all of which require solving an ODE in the same way. Nevertheless, I minimised the risk of this by using automatic differentiation to provide fast exact derivatives of the temperature profile¹⁸ and by using a 5th-order integrator. For the two-stream atmosphere treatment, I therefore evaluate the temperature gradient ∇ as

$$\nabla = \frac{dT}{dm} = \frac{dT}{d\tau} \frac{d\tau}{dm}. \quad (4.16)$$

where $\frac{dT}{d\tau}$ is the derivative of the temperature profile given above (equation 4.15) and $\frac{d\tau}{dm}$ is equation 4.6.

4.2.2 Isothermal fixed-height atmosphere

I also compared the approach above to a simpler one where we instead take an isothermal fixed-height atmosphere over a convective core.¹⁹ But I construct this isothermal atmosphere in a different way. First I model the planet as if it had a full atmosphere, then strip it of its atmosphere to leave a bare envelope: a waterworld with a convective envelope but no atmosphere on top, as we saw in chapter 3. On top of this bare envelope, I then assume a constant-gravity isothermal layer of gas. I take the thickness of this layer to be eight scale heights,²⁰ where the scale height is calculated for a given gravity and temperature. The equation for the pressure scale height of a plane-parallel water vapour atmosphere is

$$H = \frac{R_{\text{H}_2\text{O}} T}{g}, \quad (4.17)$$

I calculate the scale height H , taking the characteristic temperature T as the temperature at the top of the bare envelope. Then by adding an eight-scale-height thick ($8H$) isothermal atmospheric layer on top of the bare interior model, I obtain a planetary model with a fixed-height atmosphere. This can be compared to the radius of a planet obtained

¹⁸ Automatic differentiation is a technique for obtaining exact derivatives of explicitly specified functions without using finite-differencing. I used the Julia package `ForwardDiff.jl`.

¹⁹ I use the term fixed-height because the thickness of the atmosphere is constant for a given planetary model, depending only on the temperature and gravity, rather than generated by solving for a self-consistent temperature profile.

²⁰ Typical atmospheres are 5–10 scale heights thick (Madhusudhan & Redfield, 2015).

$R_{\text{H}_2\text{O}}$ is the specific gas constant for water ($R_{\text{H}_2\text{O}} = 461 \text{ J} \cdot \text{K}^{-1} \cdot \text{kg}^{-1}$).

with the full treatment to see whether the fixed-height atmosphere is a reasonable approximation or not.

4.2.3 Source of opacity

My source for the opacity data was Kurosaki et al. (2014). In their paper they simulated the evolution of water-rich sub- and super-Earths. They provided simple power-law expressions for the opacity of water vapour based on a fit to HITRAN²¹ data. I have taken these opacities and used them in the two-stream analytic atmosphere.

²¹ Rothman et al. (2013)

The opacity I used is

$$\kappa_{\text{th}} = 3.07 \times 10^2 \left(\frac{P}{1 \text{ bar}} \right)^{0.9} \left(\frac{T}{1000 \text{ K}} \right)^{-4.0} \text{ cm}^2 \cdot \text{g}^{-1}. \quad (4.18)$$

which is a power-law fit to the Rosseland mean opacity. In SI base units this becomes

$$\kappa = CP^\alpha T^\beta \quad (4.19)$$

where $C = 9.71 \times 10^8 \text{ m}^2 \cdot \text{kg}^{-1}$, $\alpha = 0.9$, $\beta = -4.0$, and P and T are in Pa and K respectively. This is the same form as was used to derive the boundary conditions in section 4.1.5; consequently, I use the values of α and β here in equations 4.12 and 4.14.

Table 4.1 shows sample values of the thermal opacity κ_{th} evaluated from equation 4.18. It also shows the visual opacity κ_{vis} from the same paper²² and the ratio of the two, $\gamma = \frac{\kappa_{\text{vis}}}{\kappa_{\text{th}}}$.

²² Kurosaki et al. (2014)

P / bar	T / K	$\kappa_{\text{th}} / \text{m}^2 \cdot \text{kg}^{-1}$	$\kappa_{\text{vis}} / \text{m}^2 \cdot \text{kg}^{-1}$	$\gamma = \frac{\kappa_{\text{vis}}}{\kappa_{\text{th}}}$
0.01	300	6.01×10^1	3.56×10^{-3}	5.93×10^{-5}
0.01	1000	4.87×10^{-1}	2.20×10^{-3}	4.52×10^{-3}
0.01	2000	3.04×10^{-2}	1.67×10^{-3}	5.48×10^{-2}
1	300	3.79×10^3	3.56×10^{-1}	9.39×10^{-5}
1	1000	3.07×10^2	2.20×10^{-1}	7.17×10^{-3}
1	2000	1.92	1.67×10^{-1}	8.69×10^{-2}
100	300	2.39×10^5	3.56×10^1	1.49×10^{-4}
100	1000	1.94×10^3	2.20×10^1	1.14×10^{-2}
100	2000	1.21×10^2	1.67×10^1	1.38×10^{-1}

Table 4.1: Opacities (κ) and opacity ratios (γ) calculated from the power-law forms in Kurosaki et al. (2014) across a range of pressures and temperatures relevant to the atmosphere of a heated waterworld.

I caution that these opacities may be too simple for the following reasons:

- They appear to include no treatment of H₂O–H₂O collisionally induced absorption.
- They are a least-squares fit to mean line opacities only.
- They use the HITRAN database, albeit a recent version, rather than

the more comprehensive HITEMP database which contains many more transitions for water.

- They are fitted to just three different temperatures, the lowest of which is 1000 K, and three different pressures, the lowest of which is 1 bar. At the top of the atmosphere we expect the pressure and temperature to be lower than both of these.

Nevertheless, choosing a power-law opacity allows us to simplify the expressions for the surface optical depth and pressure to the forms in section 4.1.5. We therefore have reasonable physically motivated boundary conditions for the external boundary, and hence a plausible link between the transit radius of the planet and its internal structure model.

The opacity enters the structural equations through the optical depth gradient (equation 4.6). I compute κ_{th} using equation 4.18 and the pressure–temperature profile of the model. My models therefore take into account the varying opacity of water vapour within the atmosphere, rather than assuming a constant κ_{th} .

In contrast to the variable κ_{th} , I fixed the visual–thermal opacity ratio γ to 0.01 based on representative values at 1000 K from Kurosaki et al. (2014) (table 4.1). My justification for doing so is that if I instead allow γ to vary, my models become unphysically large.²³ I believe that this is because the power-law in equation 4.18 is made by fitting to only three temperatures: 1000 K, 2000 K and 3000 K. Outside this range there is no guarantee that either κ_{th} or κ_{vis} is valid, and accurately assessing the value of γ is difficult. This is especially problematic near the photosphere because the density of the gas becomes very low (10^{-4} bar), far outside the range of the power-law relationship. I concluded that the ratio of the two in the upper atmosphere at low pressure and temperature is unreliable, at least with my current opacity prescription. I therefore fixed $\gamma = 0.01$; table 4.1 shows that this is an appropriate choice for 1000 K water vapour, which is the lowest temperature for which the opacities are valid.

I also assessed what happened when I changed the fixed value of γ and when I fixed κ_{th} . Rogers & Seager (2010a) tested values of $\gamma = [0.1, 1, 10]$ in their models of hydrogen-atmosphere planets, so I also tested these values. And I trialled various fixed values for κ_{th} across the range in table 4.1. I defer these results to section 4.3.1, as I must first wrap up my theoretical development of the atmospheric models by explaining how I link the planet’s heating to its structure.

²³ For example, I was able to produce $1 M_{\oplus}$ planets larger than $50 R_{\oplus}$ by setting $\gamma = 10^{-6}$.

4.2.4 Internal heating

Because my code does not solve any equations of internal energy transport, I represent the internal energy generation as a temperature boundary condition.

To do this, I assume that the nucleus of the planet (but not the watery envelope) generates energy that is transported through the envelope to the base of the atmosphere.²⁴ I represent this energy by a black-body temperature at the radiative–convective boundary. That is, I set the energy generation rate per unit mass in the Earth-like nucleus, ϵ , which has units of $\text{W}\cdot\text{kg}^{-1}$. The luminosity of the nucleus is then

$$L = M_p(1 - f)\epsilon. \quad (4.20)$$

I assume that the planet radiates this energy isotropically; for convenience I take the radius at which the energy is radiated to be R_p . I then use the Stefan–Boltzmann black-body relation,

$$T_{\text{blackbody}}^4 = \frac{L}{4\pi R_p^2 \sigma_B}, \quad (4.21)$$

to determine the black-body temperature of the planet based solely on its internal heating.

This black-body temperature is used as an input to the atmospheric temperature profile. The temperature itself can be relatively low but thanks to the insulating effect of the atmosphere this results in a stronger heating effect than otherwise would be expected. This prescription therefore provides a simple link between the internal energy generation and the internal heating of the atmospheric model, and hence the total radius of the planet.

In reality, because the internal heat is assumed to be deposited at the base of the atmosphere, equation 4.21 should use the radiative–convective boundary radius rather than the planet’s radius R_p . This prescription therefore assumes that the atmosphere is relatively thin.²⁵

As in the previous chapter, my internal heating treatment also assumes that the energy generated within the planet is sufficient to drive convection through the planet’s envelope. This results in a temperature gradient that increases downward into the planet until we reach the isothermal nucleus.

²⁴ The Earth’s internal energy generation rate is on the order of 1 to 10 pW·kg^{−1} if we assume that energy generation occurs only in the mantle (Davies, 1980).

L is the luminosity of the Earth-like nucleus, M_p is the planet’s mass, ϵ is the luminosity of the nucleus per unit mass, and f is the water content as a fraction of the planet’s mass ($f_{\text{nucleus}} = 1 - f$).

σ_B is the Stefan–Boltzmann constant.

²⁵ We can show that this effect is minimal for Earth-like planets if we assume an isothermal fixed-height atmosphere. The scale height H of a water atmosphere on an Earth-like nucleus at 500 K is approximately 20 km. Therefore a 10 H atmosphere is approximately 200 km thick (Madhusudhan & Redfield, 2015), only 3% of the total radius.

4.2.5 External heating

The irradiation temperature T_{irr} of a planet (equation 4.2) provides a proxy through we can characterise the radial change due to external heating from the star. I have defined the irradiation temperature in the same way as Guillot (2010), and can therefore use it as a parameter in the two-stream analytic atmospheric model (equation 4.15). We can also use the irradiation temperature in the isothermal fixed-height atmosphere by equating it with the temperature of the entire atmosphere.

The definition of the irradiation temperature, equation 4.2, is based on the star's radius and temperature. But in this chapter I do not evaluate the irradiation temperature from the properties of a host star, choosing instead to set it directly. That is, I vary over different values of T_{irr} in my models.

4.3 Results

Using the models that I have just described, I investigated the effects of the various model parameters on the mass–radius relation and internal structure of watery planets. Below, I first summarise the key results, then elaborate on these by detailing which parameters I explored. I also show a comparison intended to assess the relative effects of internal and external heating on an irradiated waterworld, and I compare the two-stream analytic atmosphere to an isothermal fixed-height atmosphere.

In this chapter I will show five key findings:

1. The presence of an atmosphere makes an observable contribution to the mass–radius relation due to the opacity κ_{th} .
2. An Earth-like degree of internal heating is enough to significantly inflate a watery super–Earth when compared to a planet with no internal heating.
3. These watery super–Earths can be inflated by both internal or external sources of energy. But, all other things being equal, changing the external heating appears to affect their radius more than changing the internal heating.
4. An isothermal fixed–height atmosphere is a good approximation to an isothermal two–stream analytic atmosphere for planets more massive than $2 M_{\oplus}$. But the fixed–height atmosphere is not a good approximation when there is a strong temperature gradient, such as that due to a strong opacity.
5. The water fraction of these planets also makes a significant contribu-

tion to their total radius.

4.3.1 Parameter explorations

Here I present the results of changing various parameters. Table 4.2 shows the default values used unless specified otherwise. Each figure in this chapter is also labelled with the parameters used to generate those models.

Parameter		Value
Planet mass range	M_p	1–10 M_\oplus
Internal energy generation rate	ϵ	$10^{-12} \text{ W}\cdot\text{kg}^{-1}$
Irradiation temperature	T_{irr}	300 K
Water mass fraction	f	0.3
Thermal opacity	κ_{th}	variable*
Visual/thermal opacity ratio	γ	0.01

Table 4.2: Default values used in the parameter exploration. *By “variable” I mean that the opacity is allowed to increase into the atmosphere according to equation 4.18 and the atmospheric temperature and pressure, rather than kept fixed.

Opacity and opacity ratio The opacity of the atmosphere is an important parameter affecting the radius of the planet (figure 4.3). I generated mass–radius relations for planets with two different opacity treatments: a variable opacity calculated from the pressure–temperature profile, and a fixed opacity. The radius of the planet changes significantly if the opacity is increased. For example, for a $2 M_\oplus$ planet, the radius is $2 R_\oplus$ if the opacity is fixed to $100 \text{ m}^2\cdot\text{kg}^{-1}$. But increasing it to $1000 \text{ m}^2\cdot\text{kg}^{-1}$ gives a planetary radius of $2.5 R_\oplus$.

If I fix the opacity to an appropriate value, I can reproduce the mass–radius relation for the variable–opacity case. This value is approximately $100 \text{ m}^2\cdot\text{kg}^{-1}$ for the models in figure 4.3.

This change in radius is because the opacity controls the rate at which the optical depth increases into the atmosphere, and therefore affects the degree to which heat can be trapped. Compared to H/He atmospheres, H_2O absorbs more strongly in the infrared²⁶ and so the optical depth at the radiative–convective boundary is often on the order of 10^4 or 10^5 . This means that energy is readily absorbed in the water layers and converted into a temperature difference.

Setting $\kappa_{\text{th}} = 0$ has the effect of making the atmosphere isothermal because τ is constant and the temperature is a function of τ equation 4.15. Figure 4.3 shows that the radius of a planet is smaller for $\kappa_{\text{th}} = 0$ than when the atmosphere has non-zero opacity.

²⁶ The opacity that Rogers & Seager (2010a) use for hydrogen atmospheres gives an opacity $\kappa_{\text{th}} = 3 \times 10^{-3} \text{ m}^2\cdot\text{kg}^{-1}$; the water opacity from equation 4.18 under the same conditions is $3 \times 10^2 \text{ m}^2\cdot\text{kg}^{-1}$, five orders of magnitude higher.

IN CONTRAST, the visual–thermal opacity ratio γ appears to have a

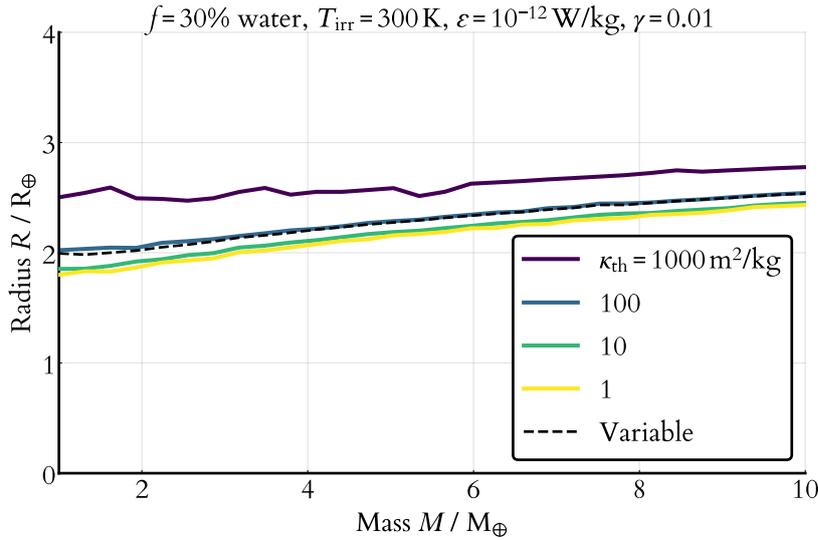


Figure 4.3: The choice of opacity κ_{th} for a watery atmosphere can strongly affect the final radius of the planet. Allowing the opacity (equation 4.18) to increase into the atmosphere with the pressure and temperature profile produces inflated radii over the zero-opacity case. But if we fix κ_{th} to an appropriate value ($100 \text{ m}^2 \cdot \text{kg}^{-1}$ in this figure), we can also reproduce the same mass–radius relation as the variable- κ_{th} case.

relatively weak effect on the final radius (figure 4.4). My models use a fixed γ of 0.01, which is typical for hot (1000–2000 K) regions deeper within the atmosphere (1–100 bar). I also assessed the mass–radius relation for the values [0.1, 1, 10] which covers the range of γ used by Rogers & Seager (2010a) in their models of hydrogen envelopes. The mass–radius relation does not change significantly over this range.

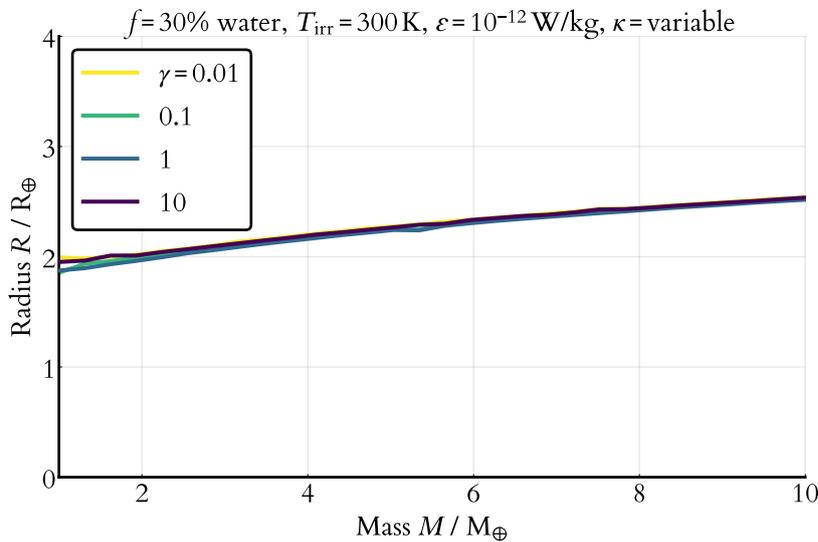


Figure 4.4: The visual/thermal opacity ratio γ does not strongly affect the mass–radius relation. Here I test several values on both sides of the $\gamma = 1$ case, which is where the visual and infra-red opacity are identical. I caution that an accurate determination of γ in cooler regions of the upper atmosphere ($P < 1 \text{ bar}$, $T < 1000 \text{ K}$) is difficult due to the domain of the opacity data used to generate the power-law opacities from Kurosaki et al. (2014). For this reason I choose a fixed value of $\gamma = 0.01$ in my models. Here I show that testing several other values around $\gamma = 1$ does not significantly alter the final radius of the planet.

Internal heating I assessed the degree of internal heating needed to significantly change the radius of a water-rich planet (figure 4.5). I make two observations based on this figure.

The first thing I observe from figure 4.5 is that even an Earth-like level of internal heating ($\epsilon = 1 \text{ pW} \cdot \text{kg}^{-1}$; that is, $10^{-12} \text{ W} \cdot \text{kg}^{-1}$) is

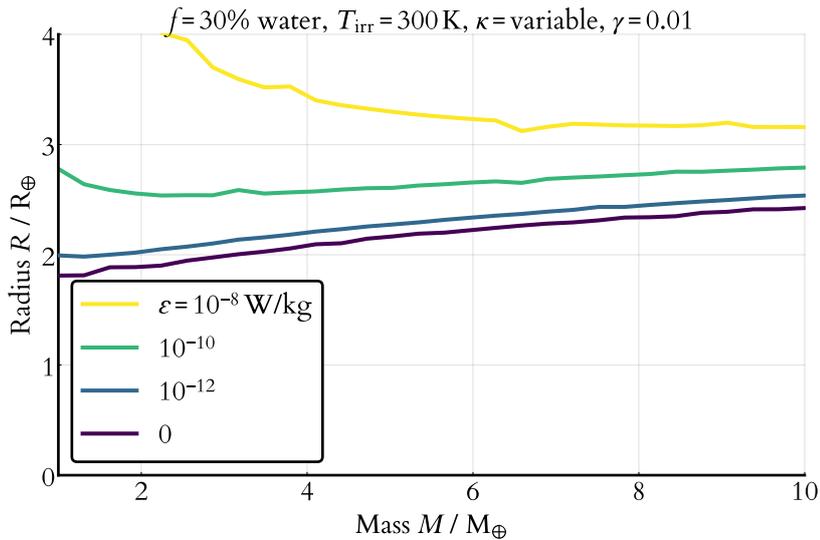


Figure 4.5: A watery planet’s atmosphere is significantly inflated with relatively low levels of internal heat. Here I show the mass–radius relations for different levels of core heating ϵ . Unlike the case of a bare planet, we do not see a sharp transition in radius associated with the liquid–vapour transition.

enough to grow the planet significantly. Compared to the zero-heating case, a $2 M_{\oplus}$ planet with $\epsilon = 1 \text{ pW}\cdot\text{kg}^{-1}$ is $0.2 R_{\oplus}$ larger. This is because the energy generated internally is trapped in the lower levels of the atmosphere, leading to a large temperature change and therefore inflating the planet’s radius.

The second point is that the dynamic range of ϵ required to significantly change the planet’s radius is quite large. The contours shown in figure 4.5 span several orders of magnitude in energy output, but the change in radius is more modest. This is a consequence of the fact that $T_{\text{blackbody}} \propto \epsilon^{1/4}$ (equation 4.21) and so the temperature change is much less.²⁷

External heating External irradiation can have a significant effect on the radius. Figure 4.6 shows the extent of this inflation. For example, above I described a $2 M_{\oplus}$ planet with $\epsilon = 1 \text{ pW}\cdot\text{kg}^{-1}$. If we also increase this planet’s irradiation temperature from 300 to 1000 K, its radius increases by another $0.4 R_{\oplus}$. This is in addition to the increase in radius already included in the model when compared to a planet with no internal energy generation.

Pressure at the radiative–convective boundary I compared the mass–radius relation for models with different fixed values for the atmosphere–envelope transition pressure P_{tr} . I did this to assess the sensitivity of my models to the exact value of the transition pressure P_{tr} . I find that varying it produces very little change in the planet’s size (figure 4.7), suggesting that the matching between atmosphere and interior need

²⁷ For example, an hundred–fold increase in the internal energy generation rate ϵ corresponds to a black–body temperature change of $100^{1/4} = 3.2\times$.

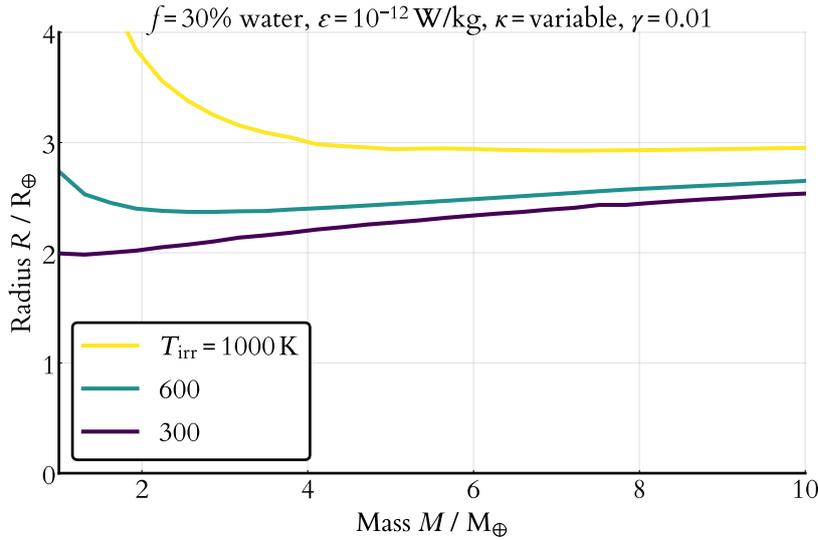


Figure 4.6: Externally irradiated planets are also inflated above the bare planet case. Here I show the inflation due to external heating on top of an Earth-like rate of core energy generation.

only be approximate if we are only interested in the final radius.

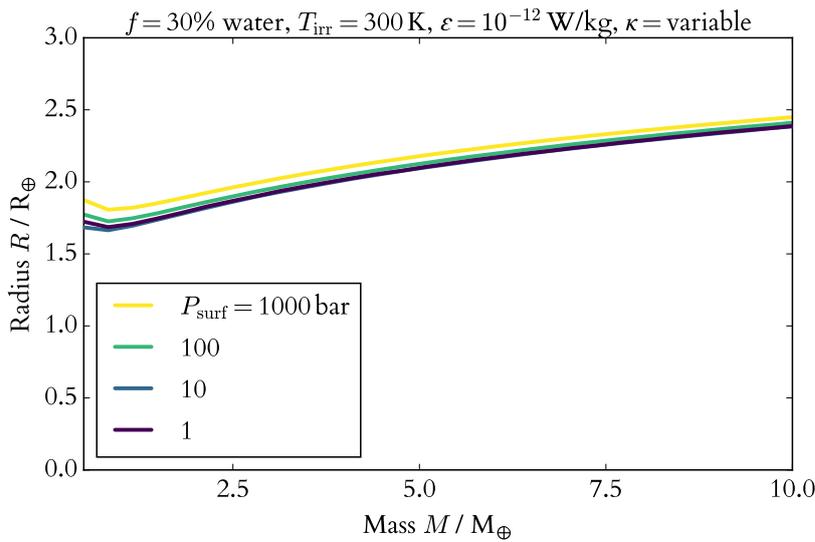


Figure 4.7: The pressure at the radiative-convective boundary only marginally affects the final radius. Here I show the effect of changing the atmosphere-envelope transition pressure P_{tr} . Although a higher P_{tr} gives a larger planet because more of the planet's mass is in the atmosphere, the effect is minimal.

Water fraction I compared the mass-radius curves for planets with varying water fractions (figure 4.8). Here I find a strong link between the planet's water fraction and its radius. For example, a $1 M_{\oplus}$ planet with an Earth-like irradiation temperature and internal heating more than doubles in size if 50% of its mass is instead taken to consist of a water envelope and atmosphere. This is larger than the difference seen in section 3.3.3 because I now treat the atmospheric layer correctly, rather than ignoring it as in chapter 3.

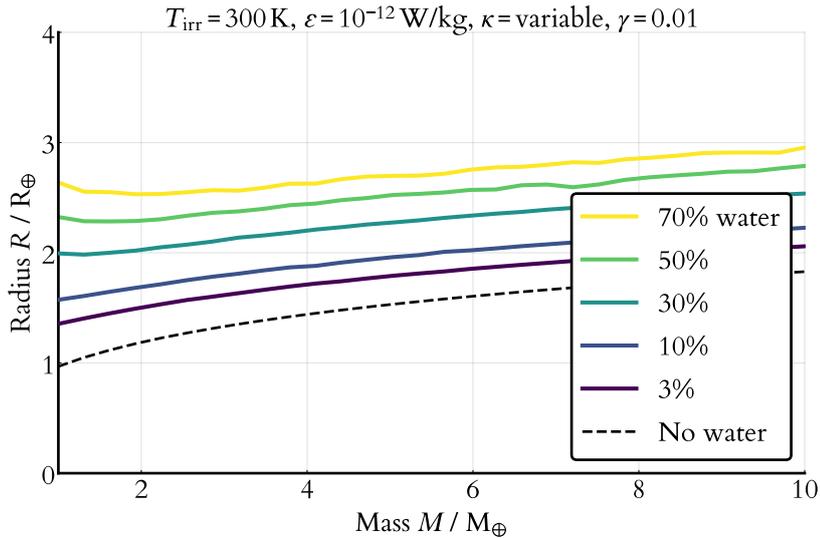


Figure 4.8: When an atmospheric layer is included, the water fraction becomes much more important in predicting the planet’s final radius. Here I show mass–radius curves for planets with increasing water fractions. The extended steam atmosphere is larger for a planet with a higher water fraction.

4.3.2 Comparing internal and external heating

Consider two planets of the same mass and composition. One of them is heated internally and one is heated externally, but the degree of heating is the same. How does the radius compare between the two cases? To address this question, I have used two different methods intended to evaluate the relative effectiveness of internal and external heating. To begin, we must model a planet that is heated by an internal source only; then we model a planet heated by an equivalent source of heat applied externally. The two methods differ in how we choose what an “equivalent source of heat” is.

A planet heated solely by internal energy generation For both methods, I begin by creating a planetary model with no irradiation; that is, I set $T_{\text{irr}} = 0$. I fix the energy generation rate per unit mass in the nucleus ϵ . Then I evaluate the radius of this planet as described in section 4.1. This model represents a planet heated only from within.

A planet heated solely through externally applied flux I take the internally heated planet as above and calculate its black-body temperature (equation 4.21) by using its surface flux. In general this value is quite low.²⁸ To compare internal and external heating, we can then equate this black-body temperature to the black-body temperature of a source whose energy input is solely irradiative. In practice this means creating a second planetary model with no internal heating and this irradiation temperature just calculated; that is, I set $T_{\text{int}} = 0$ and $T_{\text{irr}} = T_{\text{blackbody}}$,

²⁸ For example, the black-body temperature of the Earth calculated this way is about 35 K.

where $T_{\text{blackbody}}$ is the black-body temperature of the equivalent internally heated planet. Finally, I calculate the radius of this irradiated planet and compare it to the radius of the internally heated planet.

This method invariably produces planets that are smaller than the equivalent internally heated model (the dotted line in figure 4.9). This is because, with no internal heating and with an irradiation temperature much less than 300 K, the irradiated planets have virtually negligible atmospheres. Changing the internal energy generation rate used in the internally heated model over four orders of magnitude produced no discernible increase in the radii of equivalently irradiated planets.²⁹

This indicates that perhaps the black-body temperature is not a good choice for the characteristic temperature of such an internally heated planet. In reality, any internal energy input is trapped at the base of the opaque atmosphere, so the temperature through the atmosphere is much higher than $T_{\text{blackbody}}$. I therefore produced a second set of irradiated models using a different reference temperature. Here, instead of setting the irradiation temperature equal to the black-body temperature from an internally heated model, I use the temperature at the radiative-convective boundary T_{tr} . It is a better representative temperature for the interior of the planet: it is hotter than the planet's black-body temperature owing to the insulating effect of the thick, opaque atmosphere.

This second method produces planets that are larger than the first method, but still smaller than the equivalent internally heated planets. Figure 4.9 shows this for some different heating scenarios (dashed lines). For example, if we take a $2M_{\oplus}$ planet with 30% water by mass and an Earth-like level of internal energy generation ($10^{-12} \text{ W}\cdot\text{kg}^{-1}$), its radius is approximately $2R_{\oplus}$. A planet with the same temperature applied *externally* is slightly smaller. But at higher energy generation rates we see that the internal heating more strongly affects the radius: the same planet would be almost twice the size if heated internally.

These two methods produce very different results. The first method is a more direct comparison, because it effectively compares energy fluxes. But I argue that the second method better accounts for the fact that energy is trapped in the lower layers of the atmosphere when it comes from the interior. And both methods give the same conclusion: from figure 4.9, all other things being equal, it appears that a watery super-Earth is more easily inflated by internal heating than by an equivalent external heating.

²⁹ In the example above, increasing the black-body temperature of the Earth from 35 K to 350 K needs a thousand-fold increase in Earth's energy generation rate ϵ .

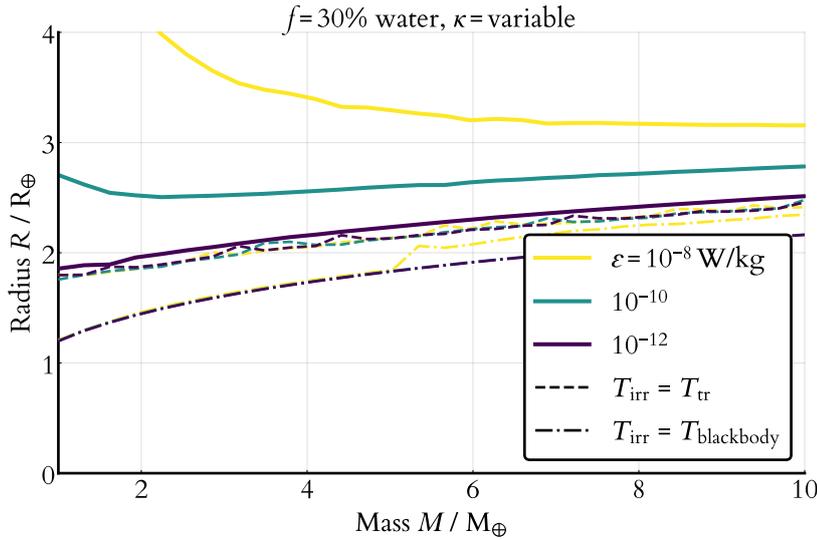


Figure 4.9: A planet’s radius increases more when it is internally heated than when it is externally heated to an equivalent temperature. Here I compare planets where the energy source is solely internal (solid lines) to planets that are otherwise equivalent except that the heating is applied externally (dashed/dotted lines). I do this by setting the irradiation temperature of the *externally* heated planet to a reference temperature from the *internally* heated planet. This reference temperature is either the radiative–convective boundary temperature (dashed line) or the blackbody temperature calculated from the planet’s surface flux (dotted line). But the solid lines always lie above the dashed/dotted lines—so the degree of inflation for external heating is less than when the planet is heated solely by an internal source of energy.

4.3.3 Atmospheric profiles

I examined the temperature profiles of heated watery planets to see how their temperature structures change with heating.

Figures 4.10 and 4.11 show the atmospheric pressure–temperature profiles for $5 M_{\oplus}$ planets with varying degrees of internal heating and irradiation. The top panel in each figure shows the profile when the opacity κ_{th} is allowed to vary through the atmosphere. The bottom panel shows fixed- κ cases where $\kappa_{\text{th}} = 100 \text{ m}^2 \cdot \text{kg}^{-1}$.

We see from these figures that the two modes of heating produce atmospheres with very different structures. When heated internally, the temperature at the base of the atmosphere increases but the bulk of the atmosphere remains at the same temperature (figure 4.10). This is because all the energy is deposited at the base of the atmosphere where the pressure, optical depth and opacity are all high. On the other hand, externally heating the models results in a profile where the bulk of the atmosphere’s temperature is increased. It also increases the photospheric pressure, because this boundary condition depends on the temperature profile at the photosphere.

We also see that, if we assume constant opacity, we may produce profiles that are markedly different from the variable-opacity case. For example, if we assume $\kappa_{\text{th}} = 100 \text{ m}^2 \cdot \text{kg}^{-1}$ as in section 4.3.1, we produce models where the upper atmosphere heats faster than in the variable-opacity case. And for internally heated planets, we can obtain significantly differently temperatures at the base of the atmosphere. Because these temperatures set the upper boundary condition for

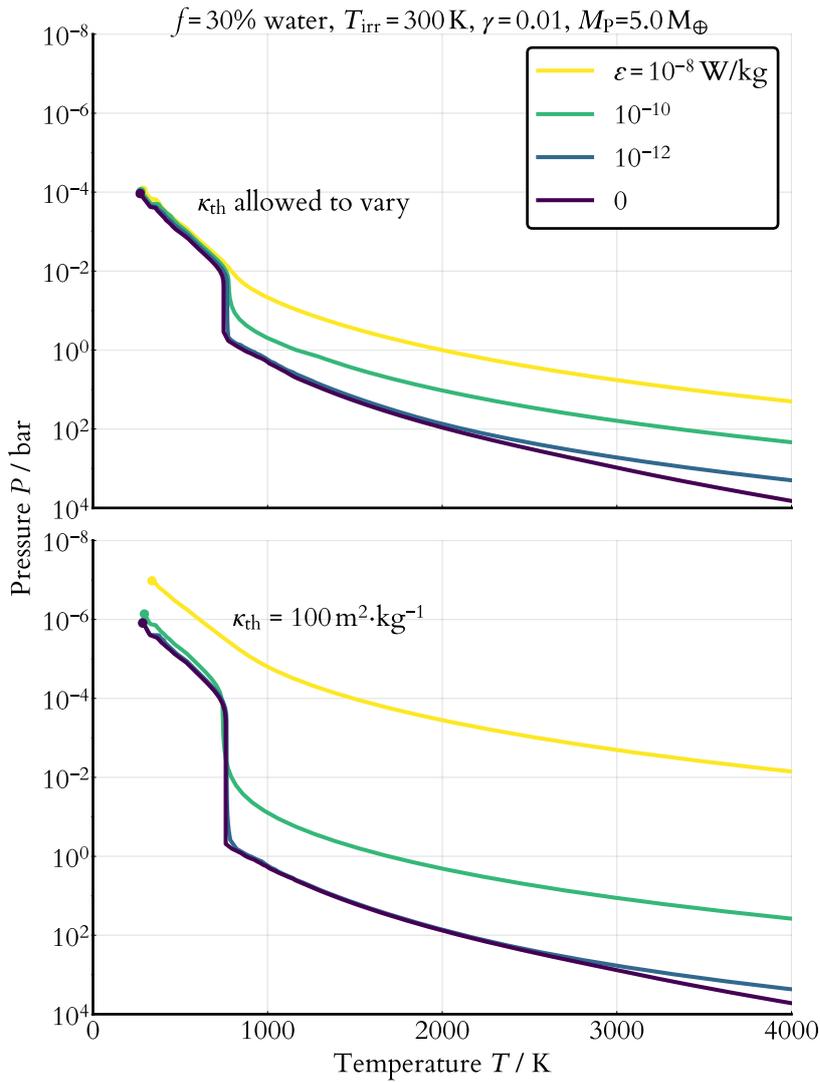


Figure 4.10: Internally heated planets have temperature profiles that consist of a rising temperature in the outer layers, an isothermal region in the middle, and a temperature that continues rising as we cross the radiative–convective boundary into the convective interior. Here I display different temperature–pressure profiles for the atmospheres of $5 M_{\oplus}$ planets with increasing core energy generation rates (and therefore internal temperatures T_{int}). The top panel shows the case where the opacity κ_{th} is allowed to vary throughout the atmosphere; the bottom panel shows fixed- κ cases. As I increase the core energy generation rate, the temperature at the base of the atmosphere increases—but this does not correspond to a change in the outer atmosphere.

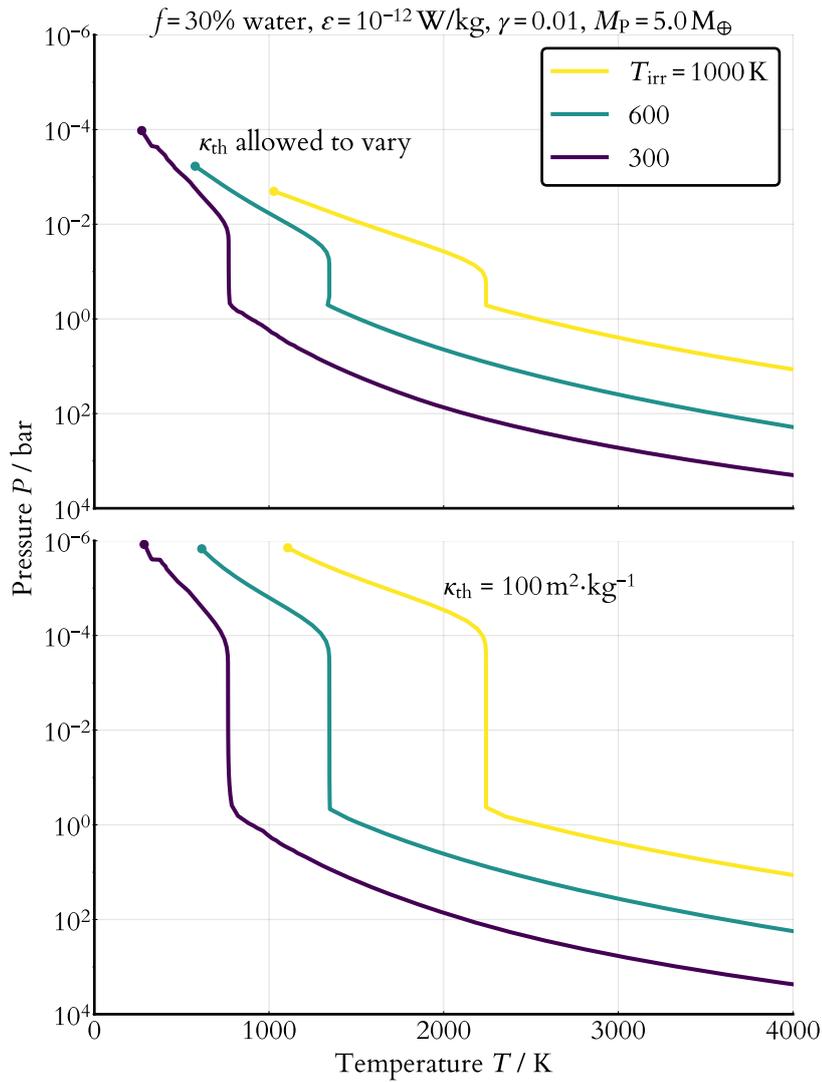


Figure 4.11: In contrast to figure 4.10, increasing a planet's irradiation temperature increases the bulk temperature of the atmosphere due to the greenhouse effect of the steam atmosphere. The photospheric pressure also increases with increasing temperature. The temperature of the isothermal region in the middle atmosphere is approximately twice as high as the irradiation temperature. These models also use a non-zero core luminosity, which manifests itself as the departure from an isotherm near the base of the atmosphere, but this does not affect the temperature of the bulk of the atmosphere.

the interior part of my model, the interior may also be at a different temperature if we assume a fixed opacity.

4.3.4 Comparison to a fixed-height atmosphere

Finally, I compared the results from the two-stream analytic atmosphere to an isothermal fixed-height atmosphere. The fixed-height atmosphere (section 4.2.2) comes from assuming that the atmosphere is isothermal and hence exponential in pressure with depth. My two-stream atmosphere becomes exponential for constant g and constant T at $\tau \ll 1.3^{\circ}$. Therefore, if we force the two-stream atmosphere to be isothermal by setting $\kappa_{\text{th}} = 0$ (and therefore $dT/dm = 0$, by equation 4.5), we should approach the limit of the fixed-height atmosphere.

³⁰ Guillot (2010)

When I set $\kappa_{\text{th}} = 0$ in this fashion, my models roughly reproduce the extent of the fixed-height atmosphere (figure 4.12). This is promising as it suggests that they behave appropriately in the limit of zero opacity. It does not mean that the zero-opacity limit or the fixed-scale-height atmosphere are appropriate choices for modelling such planets, however: as we saw in figure 4.3, the radius depends strongly on the opacity of the atmospheric water vapour.

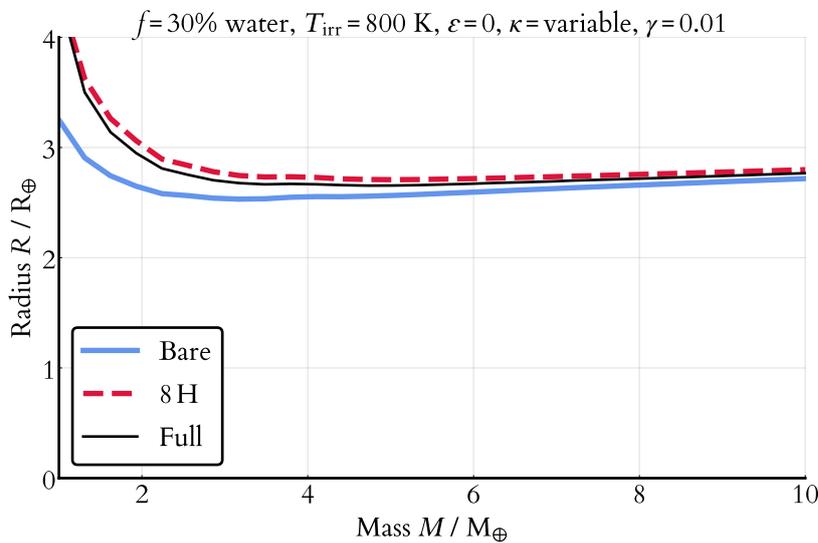


Figure 4.12: When we set the opacity κ_{th} to zero, forcing the atmosphere to become isothermal, this gives an atmosphere that is similar in thickness to an **isothermal fixed-height atmosphere**. For comparison, I also show the radius of the "bare" planet (the portion of the planet interior to the radiative-convective boundary).

4.4 Summary and discussion

I have examined the mass-radius relation for watery planets consisting of three distinct layers: an isothermal iron/silicate core, an adiabatic watery envelope and a radiative atmosphere. I employed a two-stream

analytic atmospheric temperature profile from Guillot (2010) to simplify the atmospheric structure and avoid a full radiative calculation, reducing the number of key parameters to two: the irradiation temperature and the internal temperature. I used my improved temperature-dependent equation of state to fully capture the thermal behaviour of water in the convective region, as well as temperature-dependent prescriptions for opacity. I then linked the internal heating of the planet to the temperature at the base of the atmosphere to test whether these planets could be significantly inflated by internal sources of heating.

I found that:

- The presence of an atmosphere can contribute significantly to the observed radius.
- With internal energy generation per unit mass equal to that of the Earth, a waterworld can grow sizeably compared to the case with no internal energy generation.
- For equivalent amounts of heating, internal energy generation is a more influential parameter for determining the change in planetary radius.

I also found the following:

- The use of an isothermal fixed-height approach to estimate the extent of a heated steam atmosphere gives radii that are similar to those from a full zero-opacity calculation, but too small compared to a full calculation where opacity is included. This is because the opacity of water results in heat being trapped in the atmosphere, meaning that the assumptions of constant temperature and gravity that underlie the fixed-height approach are no longer applicable.
- The thermal opacity κ_{th} is a key driver of trends in the mass-radius relation. Because my water opacity model consists of a simple power-law fitted to few data points, I cannot necessarily claim that my models necessarily represent the full opacity-dependent behaviour of water. But they do indicate that there is a significant relation to the planet's temperature that should not be ignored.
- The assumption of constant atmospheric opacity, which has been made in some previous studies, is likely to produce incorrect radii for water-rich planets. Opacity can drastically increase towards the base of the atmosphere as the optical depth and pressure become higher.

As expected, the presence of an irradiated atmosphere introduces an additional complication over the bare planet model of the previous chapter. In particular, I observe a strong dependence on the water fraction of the planet. This is a fundamental degeneracy with interior structure models that cannot be avoided. However, this means that we can, with some reasonable assumptions on the structural composition of Earth-like planets, at least assess which layered structures are compatible with a given mass and radius. See chapter 6 for an example of this application.

4.4.1 *Limitations of the atmospheric model*

The two-stream analytic atmosphere that I used has several limitations and assumptions.³¹ The two-stream approximation is valid in the limit where the incoming radiation is mostly in the visible and the outgoing radiation is mostly in the infrared, and these two wavebands are well-separated. This is because of the simplifying assumption that the thermal emission from the atmosphere at visible wavelengths contributes negligible flux. This may not be the case for heavily irradiated or strongly heated planets because a hotter atmosphere will emit more light in the visible. But even in the extreme cases shown in figure 4.10, the bulk of the atmosphere is not much hotter than 1000–2000 K. And it is apparent from figures 4.10 and 4.11 that the region in which incident radiation is absorbed, near the top of the atmosphere, is characterised by a much lower temperature than the typical temperature of a Sun-like star. For this reason, the two radiation fields used to calculate the two-stream grey atmosphere are mostly decoupled. Some caution might be warranted if we were to use these models for planets around cooler stars, however.

Guillot (2010) notes that the two-stream model diverges from numerical calculations away from temperatures significantly lower or higher than 1000 K due to changes in the mean opacities. This highlights the importance of the mean opacity in these calculations. The model also assumes that the irradiation is distributed isotropically over the surface; a tidally locked planet³² could have very different atmospheric dynamics and would be better served by a three-dimensional atmospheric model.

4.4.2 *How are these models useful?*

A simple model for irradiated planets could be useful in population synthesis models of planetary populations, or in Bayesian analysis of

³¹ Guillot (2010)

The peak wavelength for a black body of temperature 1500 K is 2 μm .

He also notes in that paper that a useful value for the convective–radiative boundary pressure is 10 bar; compare to our choice of 100 bar and see figure 4.7.

³² These tidally locked planets are sometimes called “eyeball planets” because they could have very different climates and surfaces on opposite sides with a transition ring around the edge, possibly giving the planet the appearance of a large eyeball.

large numbers of planets. Such approaches benefit from being able to quickly evaluate the radius of a planet, as they involve assessing many models at once. Until we have detailed spectral measurements of a planet's atmosphere, a full temperature model of the atmosphere may be unnecessarily complicated. These analytic two-stream models both account for temperature structure and include parameters necessary to incorporate at least the planet's irradiation temperature. With knowledge of the host star and the planet's orbit, this is an easily calculable value. These models could therefore be useful in studies of entire planetary populations to address questions such as if there are dividing lines beyond which planets are mostly gaseous.³³

³³ Rogers (2015)

The remainder of this dissertation deals with applications of these models. In the next chapter, I examine the phase structures of these heated planets and assess whether migration could significantly change their size. And I wrap up with an application of my model code to a super-Earth with a possible watery layer.

5 *Phase structure and migration*

Having constructed models of watery planets and examined how they change size with different modes of heating, what comes next? I have continued my development of watery planet structure by investigating two key questions. First, I would like to understand what the interiors of these planets actually look like.¹ Second, I would like to understand how this internal structure (and the planet's observable properties) dynamically change as the planet is moved within the star–planet system. This is a particularly relevant question in the context of planetary migration, but planets in elliptical orbits also vary in their star–planet separation, albeit on a much shorter timescale.

WHY IS IT DESIRABLE that we know the phase structure of a planet? The planet's structure could be important for several reasons. The phase of water affects its bulk geophysical properties: water behaves very differently when it is solid than when it is liquid. This consequently affects heat transfer within the planet. And different ice phases may also adopt configurations in which other molecules such as methane are trapped within the water matrix, known as filled ice, and this can also alter the temperature of the interior.² This may all have astrobiological consequences. It is therefore useful to have some grasp on how varied the internal structures of these planets can be.

The phase structure of these planets is useful to know for other reasons too. It provides a sanity check that these models are producing appropriate structures: we expect to see a gaseous atmosphere over layers of liquid, supercritical fluid/plasma, or high-pressure ices. It allows us to see whether the equation of state had its detail in the right regions: we would prefer that the final planetary models consist of phases like gas and supercritical fluid rather than ice X because the latter's equation of state is not well-known. And it provides an at-a-glance overview of the similarities and differences between planets with different heating parameters.

¹ By this I mean that I would like to understand the *phase structure* of my models; that is, the layered structure of water. For example, the phase structure of the water layer might consist of gas over liquid over high-pressure ice.

² Levi et al. (2014)

TO INVESTIGATE THESE two questions, I began by linking the planet's surface temperature to its physical context in a planet–star system. I assumed a Sun-like star and fixed the orbital separation between this star and a water-rich planet of a given mass. Then I calculated the irradiation temperature (equation 4.2) of the planet and solved for its radius and structure as described in previous chapters. Within the planetary structure, I tracked the pressure–temperature profile and used this to determine the phase structure of water within the planet. By changing the star–planet separation, I can observe how this phase structure varies. This approach therefore provides an idea of what structural changes the planet might undergo as it migrates.

In this chapter I first give some background of how I calculated the phase structure of water in my models. Then I present visualisations of the interior structure of a watery planet at different distances from a host star. I show the results for a few different scenarios and highlight the interesting features of these results.

5.1 *The phase structure of a planet*

A phase transition is a discrete change in the properties of a material. This could be a dramatic change such as the transition between a liquid and a vapour, or it could be more subtle such as a change in the crystalline structure of ice.

Phase transitions occur within my planetary models when the temperature and pressure cross a phase boundary. For example, figure 5.1 shows how several different planetary temperature–pressure structures each span different phases within pressure–temperature space.³

Because I construct my models by assuming a differentiated one-dimensional planet, the water layer in the final structure will itself consist of several different layers. I explicitly treat the atmospheric layer differently with the two-stream radiative temperature gradient, but even the adiabatic interior may have several sub-layers of increasing density towards the centre. This does not mean that any given phase may only be found in a single layer within the planet. But in practice the structure of the equation of state and the path of the adiabat in my models is such that this is almost always the case.

My solver traces an adiabat through the envelope of the planet, and this adiabat's path depends on the equation of state: it is calculated from the thermal expansion coefficient α , which I evaluate directly from the EOS by taking a partial derivative in the temperature direction.⁴ But

³ Recall that these boundaries are not necessarily well-defined at high temperatures and pressures because the behaviour of water has yet to be measured in these regions.

⁴ See equation 2.7.

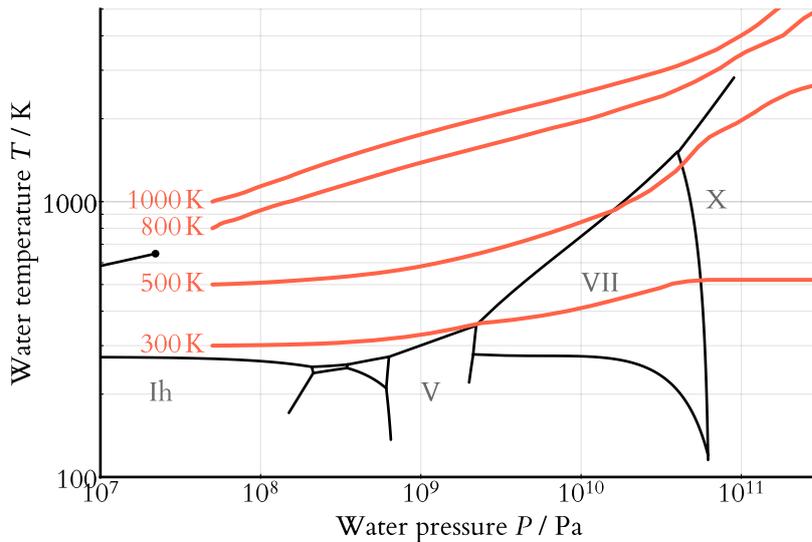


Figure 5.1: Example pressure–temperature profiles for watery planet interiors. The orange lines are the P–T profile for planets with the given explicit surface temperatures. Each orange line follows an adiabat (isentropes) within a particular phase, and the pressure and temperature are matched at phase boundaries. The planets are $3 M_{\oplus}$, consist entirely of water, and I have ignored their atmospheres by setting the surface pressure to 5×10^7 Pa, which is beyond the critical pressure. Pressure and temperature increase towards the top right of the phase space (deeper into the planet). Increasing the surface temperature means that more of the planet consists of superionic or plasma phases of water, with the transition to high-pressure ice happening deeper within the interior or not at all.

this approach breaks down at phase boundaries because they feature density discontinuities at which the thermal expansion coefficient becomes undefined.

This poses the problem of how to treat the boundaries between layers. Is the adiabat continuous and smooth in P–T space, continuous and non-smooth, or discontinuous at the phase boundary? And how should this be implemented in a way that can be used in a differential equation solver? The answer depends on the properties of the material in question, in particular whether there is latent heat involved in the phase transition and what the relative thermal expansivities and heat capacities are in the different phases.

5.1.1 The adiabatic temperature gradient at phase transitions

In the Earth, we see “seismic discontinuities”, significant changes in seismic wave speeds at different depths. These seismic discontinuities are caused by density discontinuities in the mantle minerals. For example, one component of the mantle is a silicate called olivine ($(\text{Mg}_{0.9}\text{Si}_{0.1})_2\text{SiO}_4$) which has a phase transition at a depth of about 400 km or 14 GPa.⁵ This is not a change in the chemical composition of the mantle, only in its molecular configuration—but it is enough to cause a density change and deflect seismic waves at that depth. We therefore expect to see similar behaviour in water layers as they transition between ice phases.

First, it is worth pointing out that phase transitions do not actually happen at a single depth. Because heat transport within the planet is

⁵ Milone & Wilson (2014)

driven by convection, the precise depth must vary within the planet depending on whether material is rising or falling. Rising parcels of material within the mantle will be hotter than the surrounding material. They will transition to the cooler phase only once they have exchanged heat with their surroundings, meaning that the phase transition occurs closer to the surface of the planet. But, as the exact opposite effect occurs within falling parcels of cool material, it is reasonable to use the temperature and pressure of the surrounding material when calculating the depth at which a phase transition occurs on average.⁶

HOW DO SUCH transitions affect the path of the adiabat in a convecting water ice layer? Within a single phase, the adiabat follows the P–T path defined by the solution to equations 3.4 and 3.5. Between phases there may be a discontinuity, which may be accounted for in two ways. If we know the specific entropy across the P–T range of interest, we can continue the adiabat into the new phase by matching the specific entropy at the boundary.⁷ Or if we know the latent heat of transition at the boundary but not the specific entropy on either side, we can use this to calculate the deflection of the adiabat, which is on the order of

$$\Delta T \approx \frac{-\Delta H}{C_p}. \quad (5.1)$$

But it is better to know the entropy directly as it can be used to calculate a more accurate deflection in both temperature and pressure, known as the Verhoogen effect.⁸

In summary, phase transitions cause a discontinuity and/or deflection of the adiabat into the planet, and this effect can only be correctly characterised with knowledge of latent heat or entropy across the phase boundary in question.

5.1.2 *My treatment of the adiabat*

From the above section, a complete treatment of the adiabat across water phase changes requires one of two things. Either we need the latent heat of transition between the two phases, or we need the specific entropy on either side of the phase boundary. Furthermore, these measurements need to be complete and continuous in P–T space to allow for any potential adiabat to be calculated for any potential path.

However, neither of these pieces of information are available across the range of temperatures and pressures seen in my planetary models. Even recent state-of-the-art equations of state do not treat entropy in a

⁶ Such an assumption might no longer be appropriate if we were not considering the planet as a whole—for example if we were considering a subducting slab i.e. material moving unidirectionally.

⁷ This is because the adiabat is also an isentrope.

Here ΔT is temperature change, ΔH is latent heat and C_p is the isobaric heat capacity in the high-pressure phase.

⁸ Bina (1998); Verhoogen (1965)

fully consistent way between phases,⁹ and the data sources I drew my equation of state from in chapter 2 did not provide this information. Although this information is readily available for water, ice and steam, there is a paucity of such measurements in the high-pressure ice phases.

Because we do not have measurements of latent heat or entropy to guide us, we lack a principled way to determine where to start the adiabat as we cross into the next phase. As discussed above, this would otherwise be provided by the Verhoogen effect or by taking isoentropes on an entropy surface. So in the absence of this information, my implementation of the adiabat enforces continuity of temperature and pressure across phase boundaries. I do not require the thermal expansivity α to be continuous, so I therefore allow for the planetary adiabat to be non-smooth where it crosses boundaries. But because I match the pressure and temperature of the adiabat at phase boundaries, this treatment is equivalent to assuming that there is no latent heat of phase transition.

THE WAY IN WHICH I achieve continuity of temperature and pressure across each phase boundary is as follows. It is not sufficient to simply calculate the thermal expansion coefficient α using equation 2.7 across all of P–T space, because this approach only yields the correct value for α within each phase. The density discontinuity at the phase boundaries (figure 5.2) means that we would obtain a peak in the value of α across the boundary (figure 5.3).

The “spike” seen in figure 5.3 is a partially a reflection of a physical process: a sharp density change at the phase boundary will be reflected as a spike in the value of α , which is a directional derivative of density. This change in density is associated with the release or absorption of latent heat. But it is also a reflection of a numerical process: the height of this artifact depends on the EOS grid resolution because α is undefined at the boundary. To put it another way, the value of the density derivative ceases to be well-defined at phase boundaries due to a discontinuity in density itself. It therefore cannot be used to infer properties of the adiabat across the boundary, and we cannot continue the adiabat across the phase boundary without a different approach.

To resolve this and produce an adiabat that crosses the boundary smoothly, I still calculate α from equation 3.5 as above—but I prevent the value of α from spiking at the boundary, forcing it to switch immediately to its value in the adjacent phase. In practice I achieve this

⁹ For example, Mazevet et al. (2018) produced a comprehensive water equation of state that includes specific entropy within each phase. Yet they acknowledge in that paper that an adiabat calculated with their EOS is only valid within a single phase because the entropy measures are only consistent within phases and not between them.

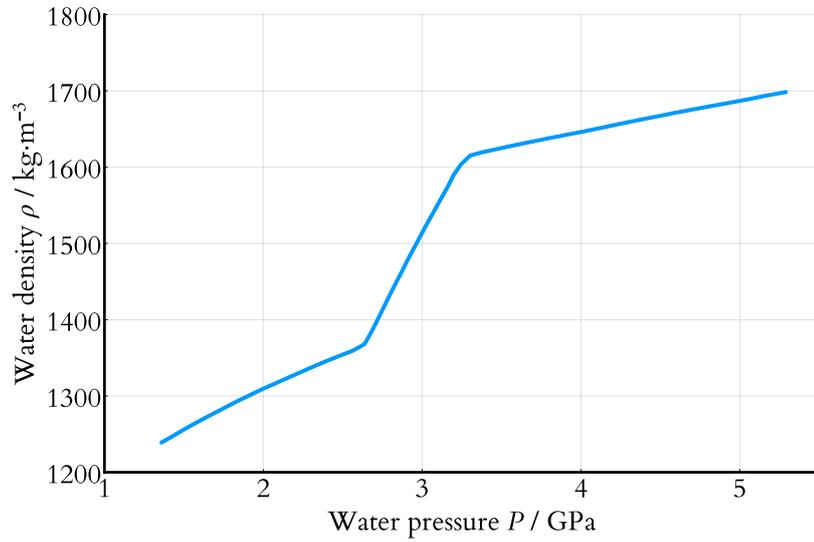


Figure 5.2: In this slice of constant temperature across the vapour–ice VII phase boundary, we see that the density increases rapidly—a result of interpolating between discontinuous values.

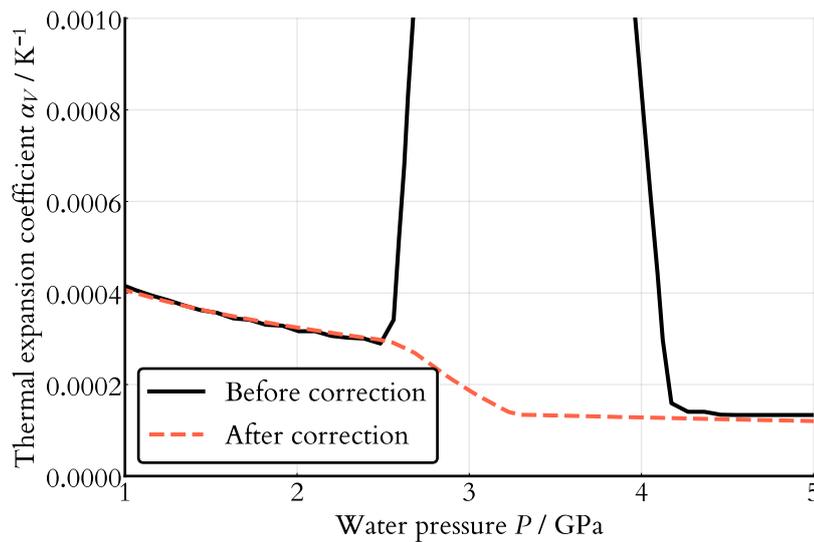


Figure 5.3: The thermal expansivity α is the partial derivative of the density in the temperature direction. Here we see that a simple treatment of α leads to a large peak at the boundary. This is numerically problematic because it means we can no longer use α in equation 3.5 to trace the path of the adiabat across the phase boundary, and a different approach is needed.

as follows. First I pre-calculate the values of α across the pressure–temperature domain of the equation of state.¹⁰ Within each phase I calculate α as normal, then stitch the phases together by joining them directly in the same fashion as the EOS stitching in chapter 2. In this way, I avoid generating artificial spikes in α at the phase boundaries, retaining the behaviour of α within each phase and producing a change in its slope, rather than a spike, at the boundary. This yields adiabats that remain continuous at phase boundaries, like those shown in figure 5.1. In contrast, there is no such need to repeat this process to deal with heat capacity discontinuities because the heat capacity is available directly and not as a directional derivative of the EOS; see section 2.2.5 for more information about heat capacity discontinuities in my EOS.

¹⁰ I tried a thresholding procedure, in which α is set to zero if it exceeds some threshold value. But the value of α varies strongly across the parameter space, and some of the phase boundaries have relatively small density changes, so I could not find a suitable threshold value for α that left all the normal values of α within the phases untouched. I therefore instead re-generated the table of α phase by phase.

ALTHOUGH THE ASSUMPTION of no latent heat of phase transition is not ideal, I claim three reasons why it is not an irresponsible choice for the planets considered in this chapter. First, theoretical studies of high-pressure ices suggest that their configurational entropies are quite similar. That is, the portion of the specific entropy change that arises from the crystalline configuration of the ice phase does not appear to differ much from phase to phase—at most 5%.¹¹ This then suggests that the latent heat difference between adjacent ice phases is also small. Second, this is supported by studies that trace specific water adiabats through high-pressure ice. For example, Dolan et al. (2007) show an isentrope calculated from the liquid phase through ice VII; there is no perceivable discontinuity at the phase transition in that paper. Finally, as we will see in this chapter, the bulk radius of a planet is driven primarily by changes to its atmospheric thickness and not by the thickness of deeper layers. A planet’s size is therefore insensitive to the precise position of phase boundaries in its water layers.

¹¹ Herrero & Ramírez (2014)

This assumption may be appropriate for the high-pressure ice–liquid phase transitions seen in this chapter, but there are two scenarios that warrant further caution. The first scenario is where the planet contains materials other than pure water: the inclusion of other minerals can substantially alter the thermochemical properties of ice. For example, adding methane results in a structure called methane clathrate or “filled ice” which has a very different thermal conductivity;¹² we might also expect its entropy to differ more between phases, producing a larger adiabatic discontinuity. A case-by-case analysis of the latent heat of transitions in these altered water ices would be needed if we were to

¹² Levi et al. (2014)

attempt to calculate an adiabat through their P - T phase space. The second scenario is where a pure water planet undergoes a liquid-gas transition, as this has the largest latent heat of any water transition. But the surrounding environment during a transition to gas is unlikely to be governed by an isentropic temperature profile any more, as we are likely in an atmospheric (radiative) regime; at this point a proper atmospheric treatment is a better choice. In any case, I did not see any solid-liquid or liquid-gas transitions in the models in this chapter (figures 5.5 to 5.8).

5.1.3 Extracting phase information

My equation of state includes phase information for water¹³ across its entire pressure and temperature range. This information is either implicitly included in the definition of that region of the EOS, or it is explicitly included as part of the data tables that the EOS is drawn from. For example, the phase ice-six (Ice VI) appears as one contiguous region in my equation of state and shares one functional form for density. The domain of this region therefore defines the pressure-temperature region in which water is ice-six. In comparison, the data tables from French et al. (2009) cover several different phases of water. For each P - T pair in their tables, they label the corresponding phase, and I use these labels to determine which phase water is in at a given temperature and pressure.

I used this information to allow me to label the layers of my planetary models with the appropriate phases. I label the atmosphere layer and everything at higher temperature than the condensation curve as gaseous.¹⁴ Within the interior, I match pressure-temperature pairs to phases in a similar manner to the procedure described in section 2.2. I define the boundaries of the ice phases using the boundaries from Choukroun & Grasset (2007). For non-ice phases, I extract phase information directly from the EOS tables where this is provided.¹⁵ For points off the unstructured grid, I take the phase of the nearest grid point in barycentric coordinates. And for any phase not yet covered by any of these regions, I treat anything with greater pressure and temperature than the critical point as supercritical fluid, and anything with $P > P_c$ but $T < T_c$ as liquid.

The end result is a mapping between (P, T) and phase across the entire range of validity of my EOS. I use this mapping to determine the phase at each mass step within the final model. By linking it to the corresponding radial distance, I can therefore produce the radial phase

¹³ *The phases of water*: The phases included in my EOS are: gas, liquid, supercritical fluid, plasma, superionic fluid; and ices Ih, II, III, V, VI, VII, VIII, and X.

¹⁴ Technically an additional distinction may be made between *vapour* and a *gas*, the first being compressible to a liquid, but I treat these under the same label for the purpose of determining phase structure.

¹⁵ French et al. (2009); Wagner & Pruß (2002)

P_c and T_c are the pressure and temperature of the critical point of water (22 MPa, 647 K).

structure of the watery layers in any planet model. Finally, I include labels for both iron and silicate, though I do not specify the phases of these materials in more detail. I label these materials based on the mass co-ordinate m rather than pressure and temperature because the core and mantle are defined in terms of the mass fraction of the entire planet.

5.2 Structural changes of migrating planets

When I investigate the phase structures of heated waterworlds in this chapter, I take the physical context to be that of a migrating planet. Up to half of Sun-like stars have a super-Earth with a period less than 100 days, and one way to explain these short periods is that the planets formed farther out and migrated inwards.¹⁶ Moreover, a third of confirmed exoplanets lie in multi-planet systems, and in such systems there are resonant processes at play that also affect migration.¹⁷ If these planets undergo migration, how are they affected by their changing proximity to the host star? My key goal here is therefore to understand how changing the orbital separation of a water-rich planet can affect its radius and/or phase structure.

I generated internal structure models for watery planets undergoing migration. I did this by moving the planet closer to its host star and recalculating its structure based on the changed surface boundary conditions. Below, I have assumed a Sun-like star, but this method is applicable to any star and the distances will be scaled appropriately.¹⁸

5.2.1 Irradiation and internal heating

To place the planet in its physical context, I use its irradiation temperature (equation 4.2), which is repeated to the right as equation 5.2 for convenience. The irradiation temperature is a parameter to the two-stream irradiated atmospheric profile from chapter 4. I calculated this temperature based on a Sun-like star ($T_{\star} = 5800$ K, $R_{\star} = 7 \times 10^8$ m). I then evaluated the planet's structure at different distances from the host star to simulate the effect of migration.

As well as setting the irradiation temperature according to the star-planet distance, I also fix the internal heating parameter ϵ . In the previous chapter I showed how I link this parameter to an internal temperature for two-stream grey atmospheric model: the luminosity of the core is assumed to produce a black-body surface temperature that is then used as the internal temperature T_{int} in the atmospheric model. I

¹⁶ Izidoro et al. (2015)

¹⁷ Baruteau et al. (2014)

¹⁸ One caveat: if the star is much cooler than the Sun, the two-stream irradiated atmospheric temperature profile may not be appropriate. See section 4.4.1.

The irradiation temperature is

$$T_{\text{irr}} = T_{\star} \sqrt{\frac{R_{\star}}{D}} \quad (5.2)$$

where the planet orbits at a distance D from a star of radius R_{\star} and temperature T_{\star} . For simplicity, I assume zero albedo.

do not change this value as the planet is moved closer to its host star.

5.3 Results

Here I present the structural changes of migrating planets. Figures 5.5 to 5.8 show two different groups of results. In the top panels, I show how the radius of a watery planet changes with distance to its host star. I also show the associated irradiation temperature at this distance. In the bottom panels, I show the phase structure of these planets at selected points in this migration.

I have chosen to represent the phase structures of the planets as a colour-coded cross-section through the planet's interior. Therefore I also provide figure 5.4 as a key to the colours and scale of figures 5.5 to 5.8.

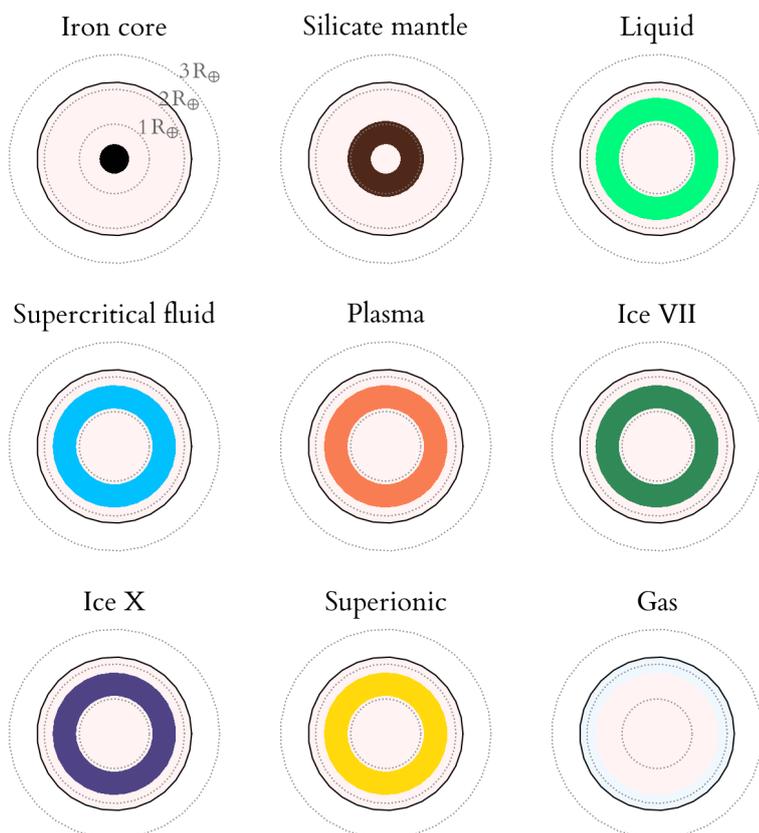


Figure 5.4: Key to figures 5.5 to 5.8. The concentric dotted circles are at integer multiples of Earth's radius and the outermost solid black line is the planet's photospheric radius. Coloured layers show the phase structure of the planet.

These figures show several values and trends at a glance:

- how much of a planet's radius is contributed by its atmosphere;
- the relative size of heated watery planets; and
- the similarity or otherwise of the structure of their water layers.

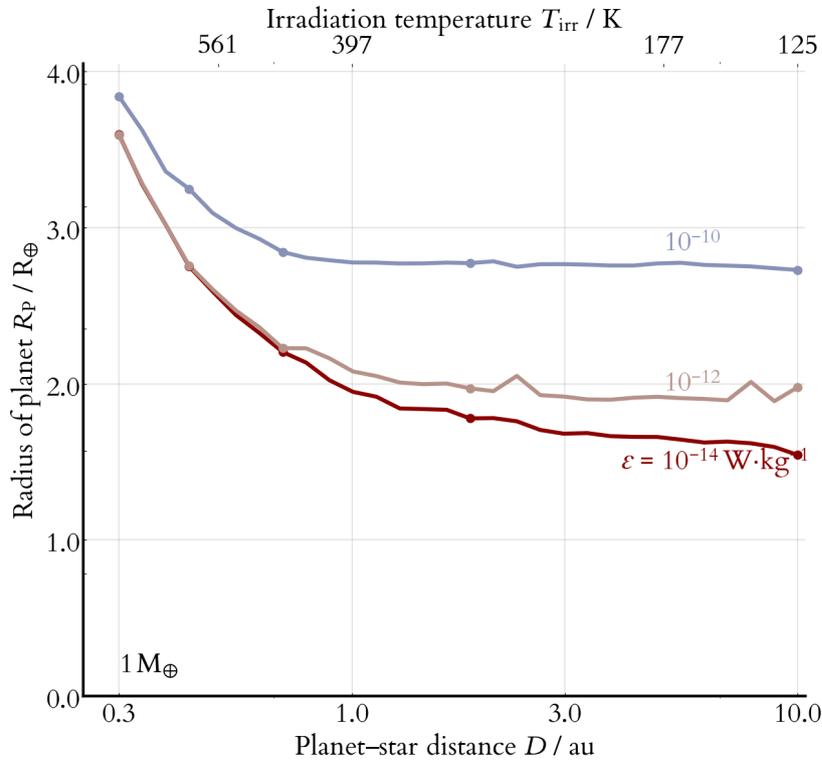
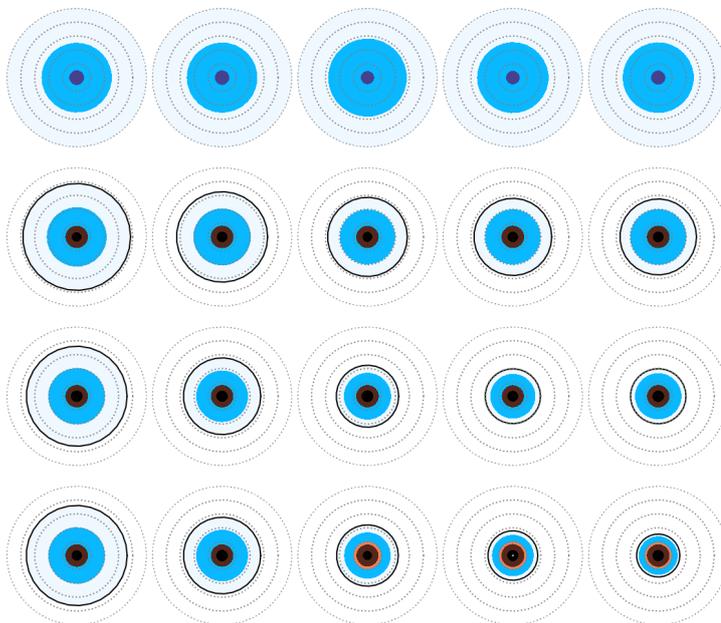


Figure 5.5: *Top panel:* Heated planets are larger, but how close does a watery planet need to be to its host star to grow significantly? Here I show the radius of a $1 M_{\oplus}$ watery planet (30% water by mass) at several different distances from a Sun-like star and with several different core energy generation rates. Such a planet can more than double in size if it is moved from 3 au to 0.3 au. *Bottom panel:* What do the interiors of these planets look like? Here I show structural diagrams for the models indicated with circles in the top panel. See figure 5.4 for the colour key. Most of the diagrams show a thick water vapour atmosphere, transitioning to supercritical fluid and possibly plasma deeper within the planet. The relative positions of the models are maintained. For example, the bottom red line in the top panel corresponds to the bottom row of models in the bottom panel. The exception is for the top row of diagrams, which show planetary structures with $\epsilon = 10^{-8} \text{ W} \cdot \text{kg}^{-1}$ that have not converged and are consequently not shown in the top panel.



The key results from these figures are as follows.

1. Small ($\lesssim 3 M_{\oplus}$) watery planets can have extended low-density atmospheres, and these can be caused by either internal or external heating. But these models may only be reliable to a point: at high levels of heating, we see what seems to be unphysical runaway atmospheric growth.
2. For the most part, waterworlds in orbits comparable to Earth contain two or three phases of water: a gaseous atmosphere, a supercritical fluid layer, and possibly a plasma or ice X shell nearer the core. There is no liquid, ice VII or superionic fluid present in any of the planets modelled in this chapter.
3. For larger planets, such as those $\approx 10 M_{\oplus}$, the size of the planet is relatively insensitive to its position within its orbit. This is because the massive planet is more easily able to hold onto its atmosphere, so external heating effects are weak when compared to the potential size change from increasing the internal energy generation rate instead.

5.3.1 *Small watery planets can have extended low-density atmospheres*

Figure 5.5 reveals that the atmospheric component of these models is key to their inflated size. When comparing phase structures of these identical-mass planets, most of the difference can be explained by the atmosphere itself. There is some expansion in the heated internal layers too, but most of the expansion is caused by the atmosphere itself, at least for the lower-mass case ($M_p \approx M_{\oplus}$). For example, compare the second row of models in figure 5.5. As the planet is moved closer to the star, the atmosphere expands while the other layers remain virtually static.

Comparing figure 5.8 to figure 5.5 then shows that this effect is exaggerated for low-mass planets ($\lesssim 3 M_{\oplus}$). This is a consequence of lower-mass planets having lower gravity. Together with high temperatures, this leads to large atmospheric scale heights and hence extended atmospheres.¹⁹

By linking the temperature of the planet to its orbital separation, we can also see that large radial changes occur across a reasonable range of orbital distances and are not just confined to planets orbiting very close to their host star. The range of orbital separations shown in figures 5.5 to 5.8 is similar to the range from Mercury to Saturn in our solar system.

¹⁹ However, here I have assumed no photoevaporation, which I do not treat in this study (see Lopez et al. (2012) and Kurosaki et al. (2014)).

These results are only reliable to a point, however. Compare the top panels of figure 5.5 and figure 5.6. The latter shows orbital separation–radius curves for four values of internal heating ε ; the former shows only three. What happened to the $10^{-8} \text{W}\cdot\text{kg}^{-1}$ case? The answer is that these models suffered from runaway atmospheric growth and did not converge.²⁰ The upper row of models in the lower panel of figure 5.5 show these unconverged models: they differ from the converged models by their lack of an Earth-like core. This is because the core is defined as the region where $m < 0.7M_{\text{planet}}$, but this region is not reached before the integrator terminates at $r = 0$.

I believe that this runaway growth is a result of incomplete opacity information in the outer atmosphere. In chapter 4 I explained that our information about water opacity in the outer atmosphere is represented as a power-law fit to a grid of a few pressure–temperature pairs. But the lowest of these temperatures is 1000 K and the lowest of these pressures is 1 bar. Further, we saw that the outer edge of the atmosphere is defined as the solution to the equation equation 4.14. This equation depends on the water opacity and, as we heat the planet, the pressure at which it is satisfied becomes very low.

To refine our estimates of planet sizes, the most important thing we could therefore do is to obtain better estimates for the opacity of water (and corresponding quantities like the opacity ratio γ) in this low-pressure regime. But at these low pressures other considerations like photoevaporation become important too. Because of the restrictions above, I am not confident that the planetary radii shown in this chapter are robust for strongly heated low-mass planets (i.e. figure 5.5). I am more confident in the later figures and in the other observations about planetary structure, which I describe below.

5.3.2 *These planets mostly contain three phases of water*

At the irradiation temperatures shown in figures 5.5 to 5.8, the bulk of these planetary models contain only three phases of water: gas, supercritical fluid and either plasma or ice X. There is no evidence that other high-pressure ice phases like ice VII are present.²¹

These graphs do not explore the trends for planets with water mass fractions other than 30%. However, we might expect that planets with larger water mass fractions are more likely to reach other high-pressure water phases near their core. They will also be less dense and have a larger radius overall than planets with iron and silicate layers. But the

²⁰ By this, I mean that the trial ODE solution failed to satisfy the inner boundary condition $r = 0, m = 0$ and instead reached $r = 0$ first. Normally this would then result in us increasing our radius guess until the boundary condition was met (figure 3.1) but no such increase in the radius could be found to produce a consistent solution.

²¹ It is certainly possible to generate planets with more varied phase structures. For example, see the images on the title page, which I generated by taking random temperatures, planet masses and water fractions. But most of the interesting structures arise with planets that are several times more massive than Earth and consist mostly of water (water fractions of 50% or more).

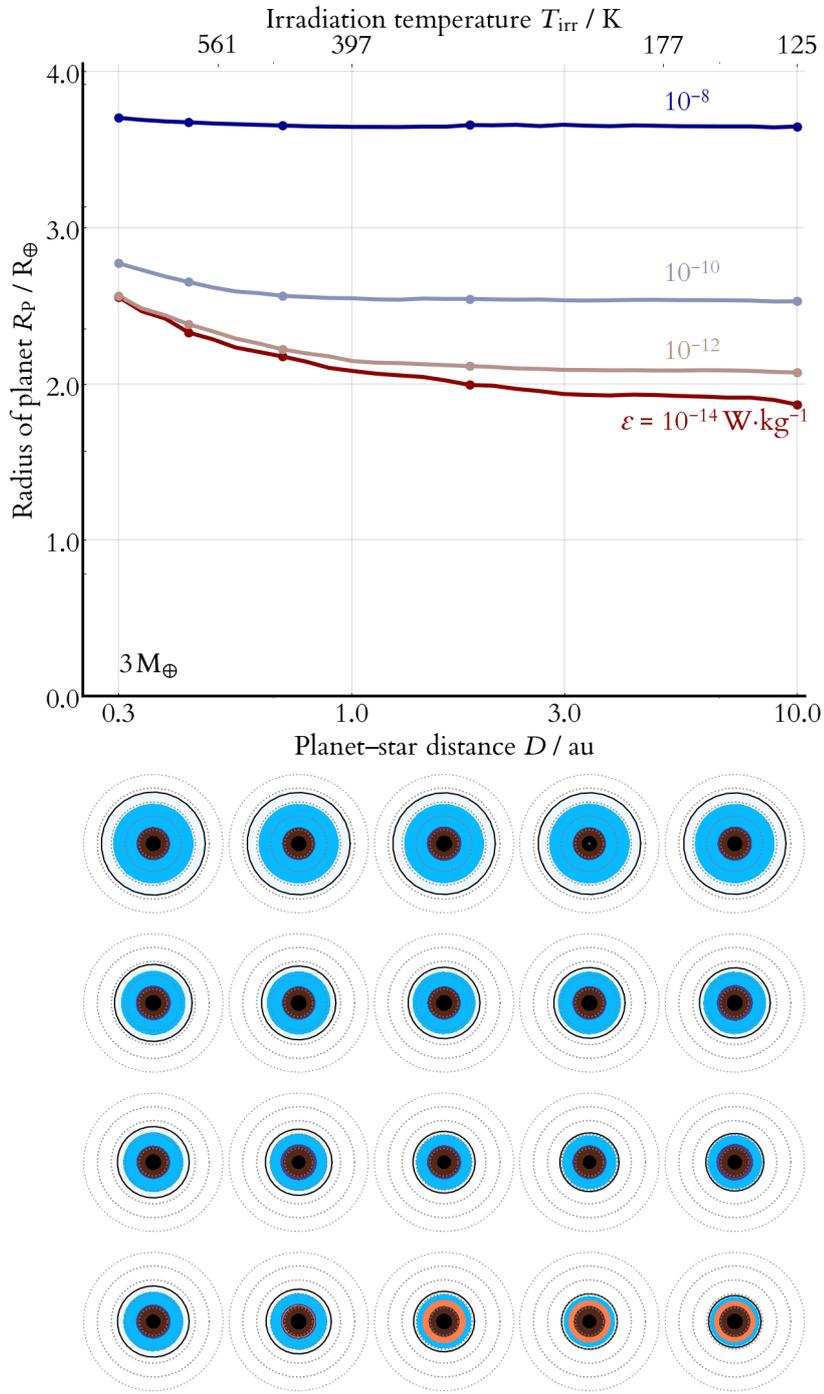


Figure 5.6: As in figure 5.5, but for a $3 M_\oplus$ planet. Comparing to figure 5.5, we see that the $3 M_\oplus$ planets can actually be smaller than the $1 M_\oplus$ planets, at least where there is moderate internal or external heating.

way in which my internal heating models are constructed, in which the internal heating comes from the core of the planet, means that the internal heating approaches zero as the core mass approaches zero. We would need a way of representing energy generation from the water envelope itself (by a tidal heating model, for example) in order to continue modelling these high-water-mass planets as being internally heated.

5.3.3 *Larger planets have a weak relationship between orbital radius and planetary radius*

Comparing figure 5.5 to figure 5.8 shows how much flatter the orbital–planetary radius curve is for the more massive planet. Earlier, we saw that low-mass planets can become highly inflated when heated, and that this inflation is due almost exclusively to the atmospheric layer of the planet expanding. With higher-mass planets, we see the opposite: the planet’s size is almost constant no matter whether it is positioned at 0.3 or 10 au, the gaseous atmosphere occupying only a small portion of the planet’s total radius in either case.

In the previous chapter we asked the question: given a planet that could be heated internally and externally, which mode of heating is dominant? There we concluded that internal heating is more capable of increasing a planet’s radius because it can affect the temperature throughout the entire planet rather than simply inflating the outer atmospheric layer. In the flat orbital–planetary radius curves of figure 5.8 we can see this effect strongly. It is easier to inflate these massive water-worlds by heating them from the inside than by moving them closer to their host star.

This finding has implications for assessing the internal structure of an observed super-Earth, because it means that information about the planet’s environment is not enough to determine whether its radius is inflated. Instead, we must take into account that any *internal* energy generation—even if only comparable to the Earth’s—can increase its radius and therefore lead to a different interpretation of the planet’s structure.

5.4 *Discussion*

To wrap up the chapter, here I provide a short discussion on the limitations of these models and some thought on their astrobiological

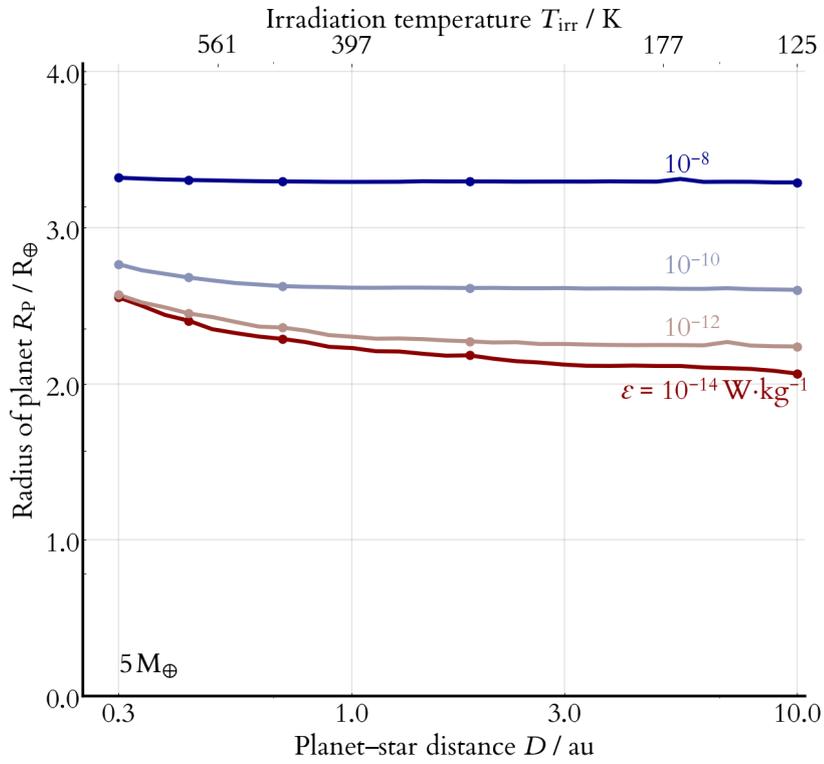
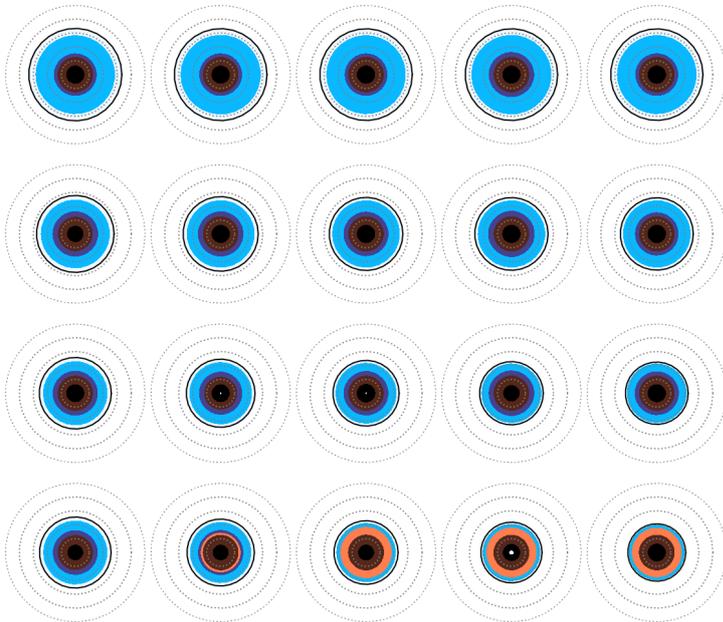


Figure 5.7: As in figures 5.5 and 5.6, but for a $5 M_\oplus$ planet.



implications.

5.4.1 *Limitations of the phase structural models*

The key limitation of labelling phases like this comes from the domain of the equation of state. The EOS only extends down to 275 K.²² Below this temperature or otherwise outside its domain, I extrapolate by simply taking the closest point within the domain of the EOS. At very high pressures, this gives results that are correct (because ice X occupies most of the phase space there), but it neglects the finer structure of various low-temperature ices. It would not be appropriate to consider wholly solid planets with this approach because my models lack any conductive heat transfer; others have developed models of cold icy bodies that are better suited to these applications.²³

²² See section 2.2 for full details on its domain.

The actual positions of the phase boundaries are still a matter of debate too, at least at high pressures and temperatures. However, I argue that the exact position of the phase boundaries is less important for gaining what is better understood as a broad overview of trends in planetary structure. It would be very easy to repeat this analysis in future with improved equation of state data, or to add in different mineral phases of the iron and silicate layers if we wanted to examine their structure from a geological perspective.

²³ See for example Ehrenreich et al. (2006).

I used these models to simulate the effects of migration. In order for this to be reasonable, we must assume that any movement occurs on timescales that are longer than the thermal equilibrium timescale of the planet (so that it remains in thermal quasi-equilibrium) and that there is no significant mass loss over this timescale.²⁴ Further, my atmospheric models assume a one-dimensional atmosphere and so therefore correspond to a well-mixed atmosphere without any latitudinal or longitudinal structure. Such an atmosphere could be found around a planet that rotates rapidly compared to the orbital timescale. Though my models also assume no rotation, in practice only a very rapidly rotating planet would affect the interior structure significantly.

²⁴ Kurosaki et al. (2014) show that super-Earths with watery envelopes and radii between 1.5 and 3 R_{\oplus} can sustain these envelopes against photo-evaporation, while lighter components like hydrogen and helium are lost from the atmosphere on shorter timescales. (Lopez et al., 2012; Owen & Wu, 2016) This is promising support for the stability of such water-rich super-Earths.

5.4.2 *Astrobiological implications*

I find the most interesting feature of these results to be what they indicate about the habitability properties of water-rich planets. I include a short discussion of this below, but have not had time to pursue these ideas any further than these brief initial thoughts.

When I initially began to study water-rich planets, I expected that

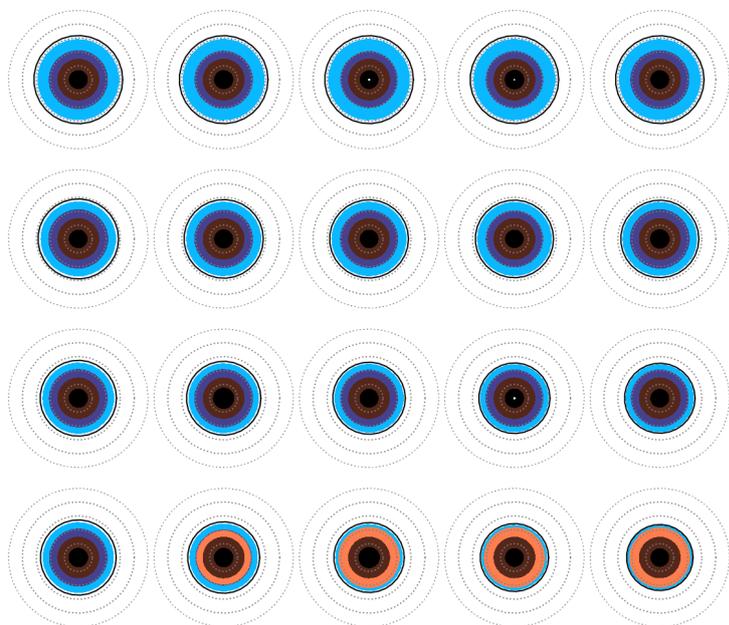
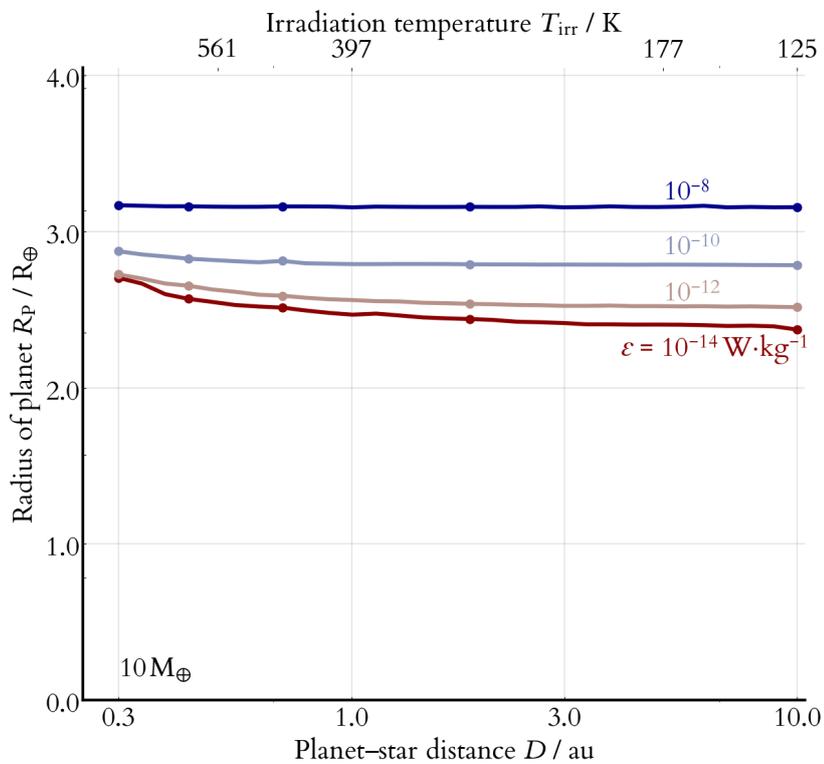


Figure 5.8: As in figures 5.5 to 5.7, but for a $10 M_{\oplus}$ planet. The radius–distance relation is much flatter for these higher-mass planets. Looking at the phase structures, we can see why: they have much thinner atmospheres. As it is the low-density atmosphere that is most responsive to temperature changes, these planets grow less when heated than a lower-mass planet would.

the most relevant phases of water would be liquid and ice. But this chapter has shown that my models, at least, result mostly in diffuse gaseous layers when we consider a relatively broad range of heating parameters. As well as finding that the radial changes due to temperature are significant, this project has also shown that the phase structure varies surprisingly little across this parameter space. With larger water fractions, such layers mostly consist of gas and supercritical fluid. What does this mean for the hospitability of such a planet?

If we accept that planets can be formed with significant water layers—and this is not particularly far-fetched, especially when we consider that even low water mass fractions can result in extended atmospheres—then these models show that liquid water may exist in a relatively narrow region of phase space. This is an important consideration in calculations of the habitable zone. For example, the phase structure of a planet with a 30% water layer allows for liquid water at the surface only when the planet is quite cold (irradiation temperature less than that of Earth's) and has very little internal heating ($10^{-18} \text{ W}\cdot\text{kg}^{-1}$). But such a planet's atmosphere is diffuse and gaseous even if the internal heating is up to two orders of magnitude lower than the Earth.

I did not have time to investigate the effect of the water fraction further. We could consider checking the models by comparing them to the Earth's atmosphere. But this is complicated by the fact that the Earth's atmosphere contains gases other than water. It also terminates at a surface pressure of 1 bar, at which point my EOS has not yet even transitioned from the ideal gas equation of state. Enhancing the EOS by adding a more appropriate treatment at the low-pressure end would be a useful next step. Ultimately, my models currently lack the capacity to handle such a thin atmosphere; they are better suited to modelling planets with much heavier water layers. Having said that, we can still use the Earth's water fraction as a useful reference point: its water fraction by mass is about $10^{-3.6}$ (see the next chapter), much lower than any of the water fractions considered in this study so far.

Depending on the orbital timescale and the timescales for convective turnover and atmospheric redistribution, we might expect a planet's structure to vary over time. Investigating this would be another interesting extension to this work: might we expect the physical properties of a planet on an eccentric orbit to change over an orbital timescale? Or would heat redistribution in the thick watery atmosphere and envelope smooth out the changes in internal²⁵ and external heating that would

²⁵ Tidal forces might become significant if the planet comes close to its host star.

arise from a highly eccentric orbit?

What about the higher-pressure and higher-temperature phases such as supercritical fluid? Are these environments inhospitable to life? This is a deeper question than can be answered here but it is worth noting that thick ocean layers (and thick watery atmospheres) can easily shield the interior from radiation that would be problematic for life on other planets. However, there are fundamental changes in the properties of water at these temperatures and pressures. Hydrogen bonds are destroyed at high temperatures and water ceases to be a good solvent for electrolytes, becoming instead a non-polar fluid.²⁶ And there are sudden changes in the chemical properties of liquid water past 0.2 GPa as well.²⁷ As much life on Earth is dependent on the peculiar chemical properties of water, it is unlikely that life as we know it could survive under such conditions. Of course this does not rule out other conceptions of life based on non-polar solvents. It does mean that the layered phase structure of a thick ocean planet could result in a relatively thin layer in which life as we know it could survive.²⁸

²⁶ Ansimov et al. (2004)

²⁷ Kruse & Dinjus (2007)

²⁸ I would like to acknowledge that these interesting points about the properties of supercritical water were raised by the anonymous reviewer for our first paper (Thomas & Madhusudhan, 2016).

6 *A water-rich super-Earth?*

The previous chapters in this dissertation have consisted of a theoretical development of the structures of watery planets. Though my models are based mostly on the measured physical properties of water, and though our motivation has always been to explain current and future observations, the topics I have explored so far have been theoretical in nature. Until just a few weeks before its submission, this was to be the entire content of the dissertation. However, we recently received news of an exciting new result: a potential water detection in the atmosphere of a rocky super-Earth, Gliese 1132 b.¹ This heated watery super-Earth is nearly a perfect candidate for the application of my models.

In this chapter, I present the results of applying my models to this super-Earth. These results were obtained by Oliver Shorttle, a post-doctoral fellow in our group working with my supervisor Nikku Madhusudhan. Oliver carried out the analysis in this chapter using an earlier version of my code `ONION` to analyse the potential composition of this planet. I have included this work as an example of the first direct application of my models to observational data.

6.1 *Gliese 1132 b*

The planet Gliese 1132 b, hereafter GJ 1132b, was discovered by MEarth² in 2015. It orbits a M4V-type star at a distance of 12 pc, considerably closer than the closest previously known Earth-sized M-dwarf planets. With an equilibrium temperature of up to 580 K depending on albedo, it is cool enough to have retained an atmosphere yet warm enough that any atmosphere is probably depleted of hydrogen.³ Table 6.1 shows some key parameters of the system.

The radius given in table 6.1 is the continuum or “surface” radius of GJ 1132b; the radius in the z and K bands are higher, suggesting that the planet has an atmosphere.⁴

¹ Southworth et al. (2017)

² MEarth is a survey that searches for Earth-like planets around M dwarfs. It takes advantage of the “small star opportunity”: small faint stars such as M dwarfs have proportionally larger transits because the ratio of stellar to planetary radius is lower. Because small planets are common around these cooler stars, they represent a strong potential source of super-Earth detections.

³ Berta-Thompson et al. (2015)

⁴ Southworth et al. (2017)

Property		Value
Right ascension	α	10 ^h 14 ^m 51.846 ^s
Declination	δ	-47° 09' 24.47''
Distance	D_{\star}	12.0 ± 0.2 pc
Mass	M_p	1.6 ± 0.5 M_{\oplus}
Radius	R_p	1.35 ± 0.21 R_{\oplus}
Orbital period	P	1.628930 ± (3.1 × 10 ⁻⁵) day
Equilibrium temperature	T_{eq}	Up to 580 K depending on albedo

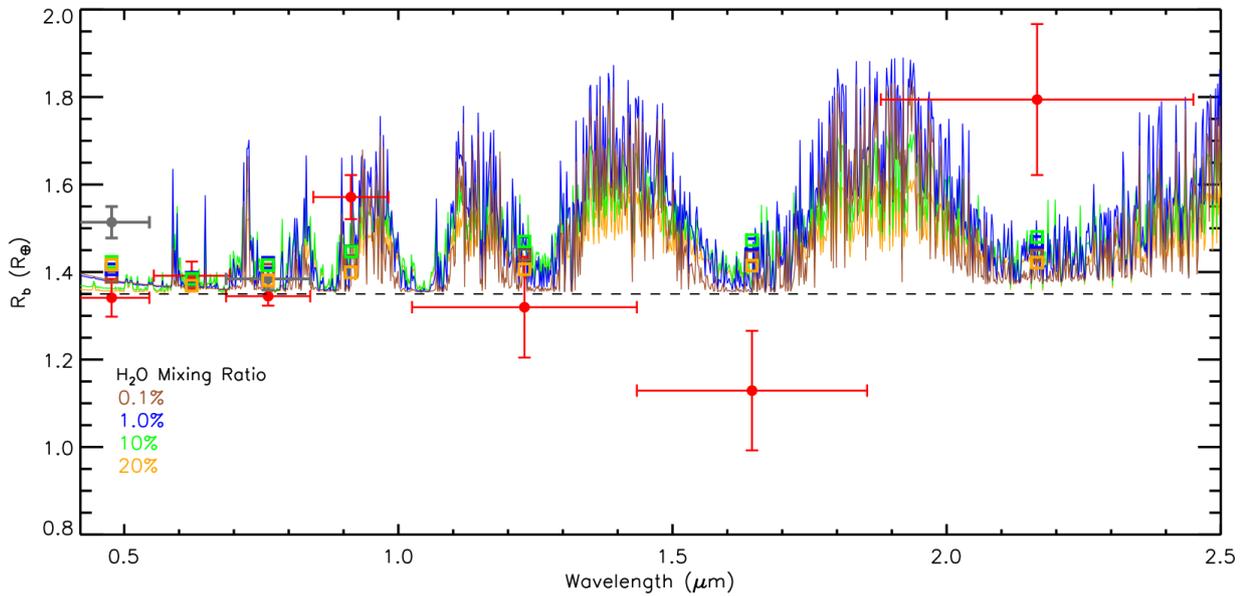
Table 6.1: Key properties of GJ 1132b. All values come from Berta-Thompson et al. (2015) except the radius, which has been revised upward from an initial estimate of $1.16 \pm 0.11 R_{\oplus}$ (Southworth et al., 2017).

6.2 A potential water detection

GJ 1132b has been observed in several photometric bands.⁵ The most recent measurement in the z-band differs from a flat spectrum at the 4σ level. This band overlaps with a strong H₂O absorption window whereas all the other bands are windows in the H₂O opacity; it has previously been suggested that this is an ideal situation in which to make simultaneous observations of the atmosphere and interior.⁶ This planet is therefore an excellent choice to apply my models to: it is a heated super-Earth ($T_{\text{eq}} \approx 600$ K) with a potential water detection (figure 6.1).

⁵ Southworth et al. (2017)

⁶ Madhusudhan & Redfield (2015)



The detection paper used simple temperature-independent models of the interior structure. In their paper announcing the atmospheric detection, Southworth et al. (2017) note that “In principle, consideration of temperature-dependent internal structure models would lead to larger model radii for the same composition (Thomas & Madhusudhan,

Figure 6.1: The observed transmission spectrum of GJ 1132b (red points) compared to theoretical spectra of H₂/H₂O atmospheres (coloured lines, with squares showing band-integrated values). The feature at 0.9 μm (the z-band) in particular is indicative of a water detection. The dashed line is the baseline radius of the planet inside opacity windows. — From Nikku Madhusudhan in Southworth et al. (2017)

2016) and therefore could lower the upper limit on the water mass fraction.” I made an earlier version of my ONION code, called OGRE, available to Oliver Shorttle to do just that: he has used it to investigate the structure of GJ 1132b.

6.3 Results

My models have been used in two ways so far to provide insight into GJ 1132b’s structure. First I show the results of comparing the measured mass and radius of the planet to models with different compositional parameters. In particular, we can vary the mass fractions of water, silicate and iron and compare the model and observed radii. Then I show how the phase structure of the planet would be expected to vary as we change the water mass fraction, and how the predicted temperature at the base of the water envelope relates to the melting properties of a silicate mantle.

In these results, unless stated otherwise, all other parameters were fixed to the fiducial values specified in table 4.2, except for the equilibrium temperature of the planet, which is taken as 580 K from table 6.1.

6.3.1 Constraints on interior composition

The mass and radius of GJ 1132b are consistent with a planet that has a watery layer of some kind. Figure 6.2 shows the one-sigma error ellipse for the mass and radius measurement of GJ 1132b.⁷ Overlaid are some mass–radius curves for different planet compositions, the largest of which includes a 10% layer of water by mass on an Earth-like nucleus. These models use my full atmosphere–interior structural code from chapters 3 and 4 so that the effect of temperature is included in the water layer. While a planet with a large iron fraction is ruled out by these mass–radius relations, GJ 1132b could have a sizeable water layer.

⁷ Southworth et al. (2017)

If we assume that GJ 1132b has an Earth-like nucleus, we can vary the water mass fraction to see what range of water fractions are consistent with its radius. Figure 6.3 shows a one-sigma band around GJ 1132b’s measured radius. Overlaid on this, we show the predicted radius for planets with varying water mass fractions over an Earth-like nucleus. The water fraction could be up to about 10% of the mass of the planet. But the radial uncertainty of the planet also allows for a hydrosphere with the same mass fraction as Earth, which is about $10^{-3.6}$.

Using a ternary diagram, we can change all three different mass

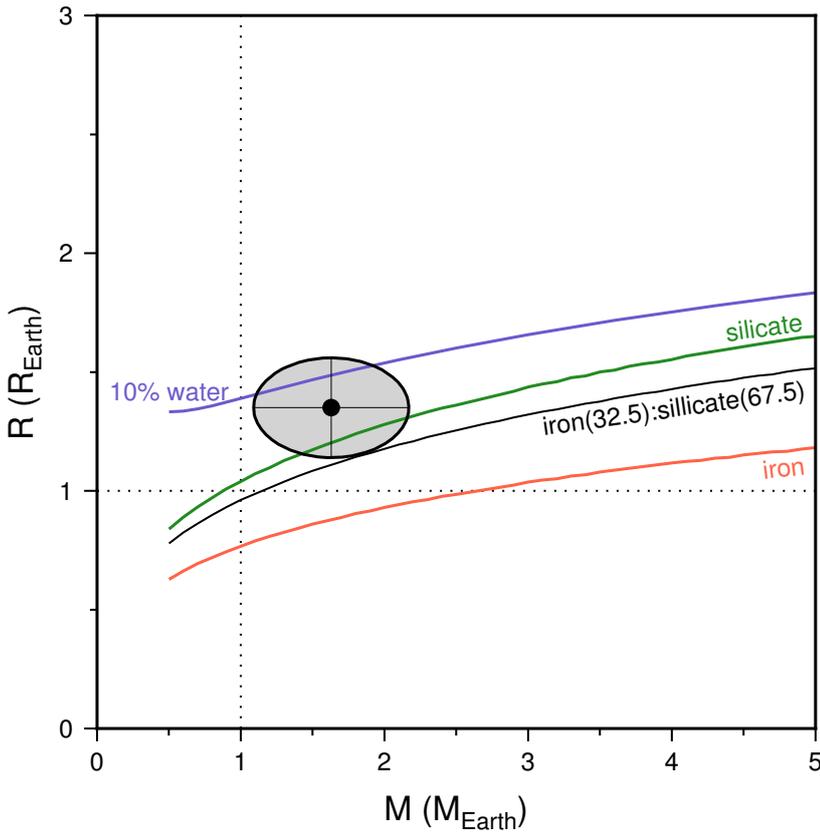


Figure 6.2: The recently revised estimate for GJ 1132b’s radius means that its density is only just consistent at the one-sigma level with an Earth-like planet. This figure shows mass–radius curves for pure silicate and pure iron planets, as well as for an Earth-like mixture of 1:2 iron:silicate by mass (in black). It also shows the mass–radius curve for a planet with a layer of 10% water by mass on an Earth-like core. GJ 1132b’s density is consistent with the planet having an extended water atmosphere. — Figure provided by Oliver Shorttle

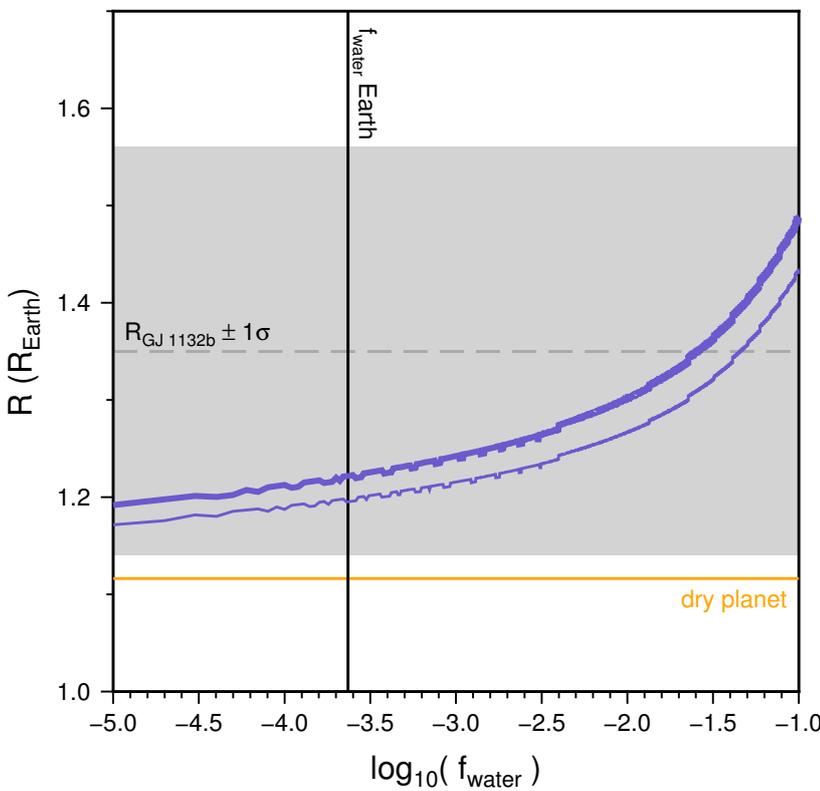


Figure 6.3: When we compare the radius of GJ 1132b with the radius of watery planets, we see that GJ 1132b could have as much water as Earth or even several orders of magnitude more. This figure shows the radius of a dry planet (1:2 iron:silicate by mass) with GJ 1132b’s mass, as well as how the planet’s radius increases when part of the mass is taken to consist of a water envelope. The lower line represents a planet with a Bond albedo of 0.75, while the upper line is the zero-albedo case. The Earth’s water fraction is also shown for comparison. A wide variety of water mass fractions are permitted. — Figure provided by Oliver Shorttle

fractions (iron, silicate and water) and determine which of these are most suitable for explaining the observations. Figure 6.4 shows the range of compositions that are consistent with GJ 1132b's observed mass and radius. Compositions that are very low in iron are not consistent with its observed radius given its mass. But planets with up to 30% water are consistent if the planet has a more massive iron core; that is, if its nucleus is more like Mercury in structure than the Earth or Mars.

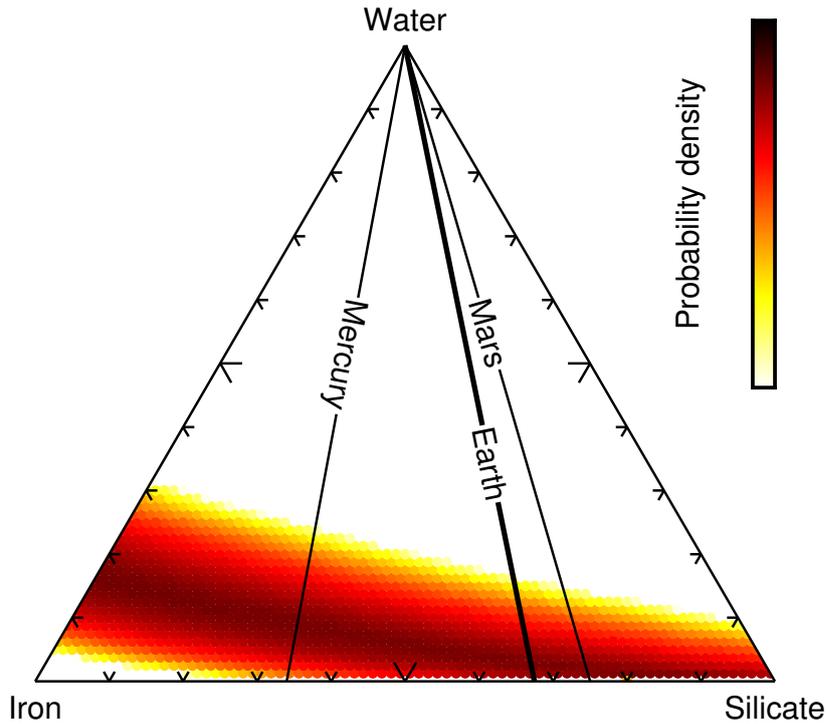


Figure 6.4: This ternary diagram shows the range of compositions permitted in a three-component iron-silicate-water model for GJ 1132b. The axes are the mass fractions of the given material. The planet is assumed to be fully differentiated and with Earth-like internal energy generation. The shading shows the probability that a model with the given layered composition and a mass drawn randomly from the error distribution can match the observed radius. The lines show the observed iron-silicate fractions of Mercury, Earth and Mars for comparison; the overlaps between these lines and the shaded region represent the range of water GJ 1132b would need to be consistent with observations if it had the same core/mantle composition as the named planet. — Figure provided by Oliver Shorttle

6.3.2 Phase structure and geology

We can also explore the phase structure of the planet using the method I described in chapter 5. Figure 6.5 shows how the phase structure and temperature of the planet varies as the water fraction is increased. If GJ 1132b's water layer is the same mass fraction as the Earth's, then the water layer is most gaseous. But if it has more than the Earth's water fraction then it could transition to supercritical fluid or plasma deeper within the planet. The water-silicate interface is also pushed to higher pressures down into the planet as the water mass fraction increases.

Finally, using information on the melting properties of silicate, we are able to take the temperature information at the base of the water layer and determine whether the Earth-like nucleus would be molten. Figure 6.6 shows the temperature differential from the solidus/liquidus⁸

⁸ At temperatures lower than the solidus, a material is completely solid. At temperatures higher than the liquidus it is completely liquid. In between, crystals form in the melt and the material exists as a slurry.

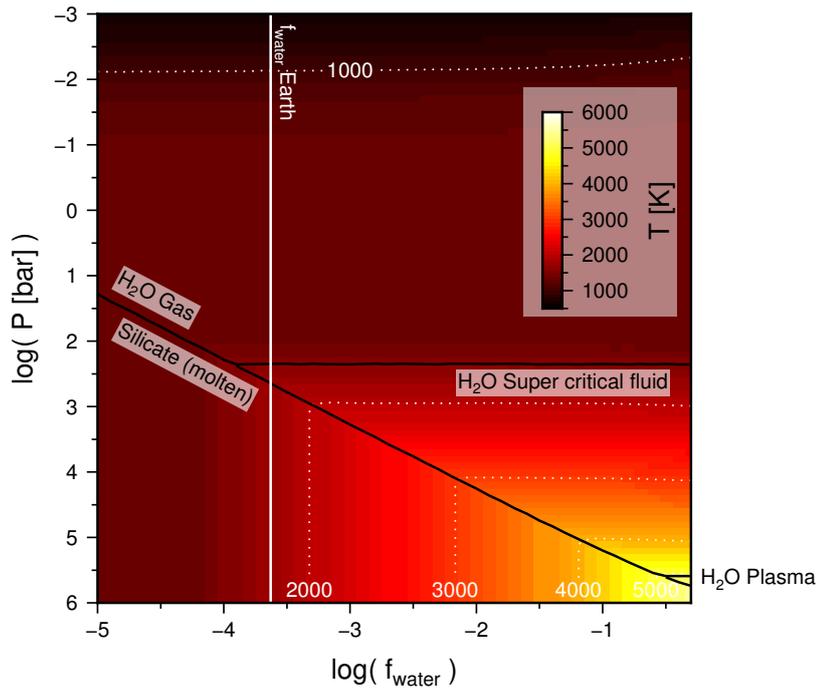


Figure 6.5: Changes in the structure and temperature of the GJ 1132b planetary model. At low water fractions, the model comprises a gaseous water layer on molten silicate. At water mass fractions greater than that of the Earth, the water layer transitions to supercritical fluid or plasma. Recall that the silicate and iron core are treated as isothermal, so the trend of increasing temperature at high pressure with increasing water fraction is an artifact of the model assumptions. — Figure provided by Oliver Shorttle

at the top of the silicate layer in the GJ 1132b models. We vary the water fraction as before to investigate if this changes the melt state of the mantle. We see that, at the temperatures at the base of a water layer, any silicate mantle is most likely molten. This is also found to be true even under a very modest water layer (up to an order of magnitude less in mass fraction than the Earth).

WITH THIS GEOLOGICAL information, the models raise some questions about the geological stability and properties of water-rich planets. In particular, the existence of a gaseous atmosphere over molten silicate is unusual and may not be physically plausible depending on the nature of the atmosphere–interior interactions. We might also ask how much water can be stored in its interior, whether it is outgassed, and what time scale this happens over.⁹ Ultimately, the interior structures of planets like GJ1132b are an exciting avenue of investigation and I look forward to seeing what advances the future brings.

⁹ Komacek & Abbot (2016)

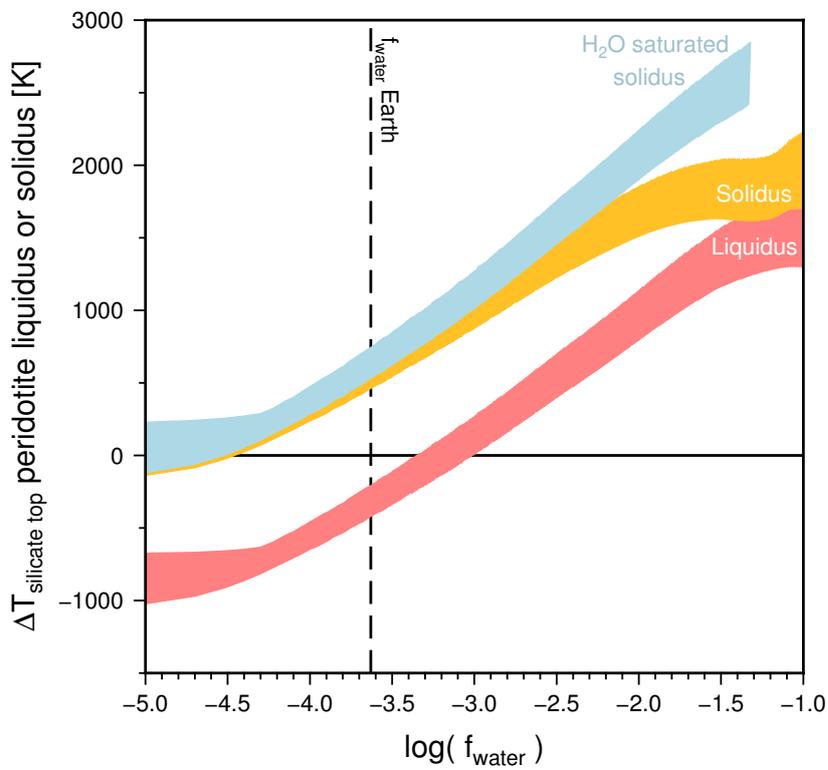


Figure 6.6: Temperature difference between the melting curve of peridotite, a common Earth mantle rock, and the modelled temperature at the base of the water layer in GJ 1132b. The coloured bands show three different mineralogical melting curves, and the width of each band represents models with Bond albedos ranging from 0–0.75. When these curves pass above the horizontal line at $\Delta T = 0$, the rock is melted. The first and second curves are the solidi for water-saturated peridotite and dry peridotite respectively. At temperatures beyond the solidus the rock begins to melt, but initially the melting is non-total and the rock is in a slurry. The third curve is the liquidus: beyond this point the rock is completely molten. Across the entire range of water fractions considered, almost all provide for some degree of melting at the top of the silicate layer. — Figure provided by Oliver Shorttle

7 *Summary and conclusion*

Water-rich super-Earths present a fascinating target for theoretical studies of planetary interiors and atmospheres. As we have just seen, these theories are becoming reality, with new observations already appearing and more on their way.

Why do we care about watery super-Earths? We care because they represent an important class of planets—not quite like anything in our solar system, but not alien to us. We care because water is such an important, life-giving substance on Earth that we seek it out throughout the Universe to understand if it could have given rise to life elsewhere. And therefore we care because these planets represent one of the greatest unsolved mysteries of our time: are we alone? Planets and moons with watery layers are currently our best hope for answering this question.

This dissertation has focused on modelling and understanding water-rich super-Earths. In it I have asked how the presence of water affects both the observable properties and the internal structures of these planets. Because our information on exoplanet interiors is limited in the best case to their masses and radii, we must necessarily work in an information-poor environment when we do this. This poses a problem because it leads to a fundamental degeneracy in potential interior structures. But it is also an opportunity for us to put our theories to work and explore which parameters are important to the final structure.

IN THIS DISSERTATION I have explored several different areas relating to the internal and atmospheric structures of water-rich super-Earths.

I began by constructing a comprehensive water equation of state that includes a full treatment of temperature dependence. This temperature-dependent equation of state is important because of water's rich and varied phase structure, and because planets with thick water layers can have internal pressures and temperatures that span several orders of magnitude.

Image removed for
copyright reasons

danodonoghue.com

Next, I set out to understand how the physical properties of a water-rich planet are affected by temperature. I included an adiabatic temperature gradient in my models, which is the expected temperature gradient for a planet that is transporting energy via convection. I showed that planets with watery layers are sensitive to temperature changes, especially when this new equation of state is included. The degree of change depends on a number of factors, but it is clear that the temperature dependence is an important component of these models that should not be neglected.

I then included a treatment for the planet's atmosphere to make the link to a planet's physical transit radius. I used an analytic two-stream atmospheric temperature profile. It is physically motivated yet simple enough for our needs; and it lets us examine the effects of internal and external heating on the planet's structure. Though this atmospheric treatment makes a number of simplifying assumptions, it has proved to be very useful for understanding the structure of the planet's atmosphere. By including this atmospheric profile, I showed that the atmosphere can have a significant temperature-dependent contribution to the observable radius of a water-rich planet. I also showed that the water opacity is a critically important parameter: it controls the atmospheric boundary conditions and therefore affects the structure of the entire model.

Next, I applied my models to the context of a migrating planet. I asked how the phase structure of a watery planet can vary with internal and external heating; more specifically, I investigated this by modelling planets moving closer to their host star. By modelling their heated interiors, I was able to show that the phase structure of these planets remains roughly similar during migration, though the radius can be significantly inflated by internal heating and to a lesser degree when the planet approaches the star.

Finally, I presented the results of using my models to analyse the interior structure and geology of the planet GJ 1132b. GJ 1132b's mass and radius can be explained by a planet with a water layer over a molten silicate mantle. Though the water mass fraction is not well-constrained, this fits with new spectral observations of the planet that suggest it has a watery atmosphere.

WHAT OF THE FUTURE? Where is the field headed, and how do contributions like these fit into the picture? By understanding how the

mass–radius relation can be affected by temperature, we can detect and characterise water-rich planets while including their surface temperatures in our models. This is an important precursor to narrow the search to planets that would be considered more classically habitable. It will be especially useful in the context of the next generation of super-Earths expected to be found orbiting bright stars by missions such as PLATO¹, TESS² and CHEOPS³. My approach is promising because it is linked to the characteristic equilibrium temperature, which can be determined from observations of the planet, and so can be included in analyses of populations of planets. Through this we might better understand what proportion of planets include substantial water content.

I look forward to two developments in particular. The first is improved atmospheric characterisation and modelling, which will provide useful pressure and temperature boundary conditions at the base of the atmosphere. The question of interior–atmospheric interactions is a rich one that is only starting to be explored. Integrating atmospheric and interior models promises progress on questions about surface chemistry, outgassing and other processes that can shape the atmosphere of a planet. The second development that will make use of this work is improved spectral resolution of atmospheric observations, and in particular the ability to seek out atmospheric windows.⁴ By observing at wavelengths which pass through the atmosphere, we can in principle directly measure the radius of any solid interior underneath that atmosphere and thus have a better starting point for interpreting the interior structure. The case of GJ 1132b in chapter 6 is an excellent early example of this sort of work. With models like these combined with stunning new observations, our future understanding of these intriguing planets looks bright indeed.

¹ Rauer et al. (2014)

² Ricker et al. (2014)

³ Broeg et al. (2013)

⁴ Madhusudhan & Redfield (2015)

My supervisor
NIKKU MADHUSUDHAN
and other academic staff
DR. IAN PARRY
PROF. CHRISTOPHER TOUT
DR. VASILY BELOKUROV
for advice and guidance.

(For financial support & accommodation)
TRINITY COLLEGE
Walking through the courtyards never
fails to make me think "Wow."

My financial backers
**THE RUTHERFORD
FOUNDATION TRUST
UNIVERSITIES NEW
ZEALAND**

It has been a pleasure
to have had your support.

For being a wonderful place to work
THE INSTITUTE OF ASTRONOMY
and for being oh-so-helpful
ITS FANTASTIC SUPPORT STAFF

Acknowledgements

For brightening my life in diverse and unexpected ways

MY FELLOW STUDENTS and a special thanks to **CHRISTINA HEDGES**
for fun and friendship.

There are too many of you to
name so let's just say that you
guys are the best. My thanks also go to your lovely

ASSORTED PARTNERS, HOUSEMATES, FAMILIES AND FRIENDS

for making me so welcome in the UK.

For never failing to
brighten my day

**THE SQUIRRELS
OUTSIDE MY
WINDOW**

Your tails are so
twitchy and your little
paws are adorable.

Finally, for so many things: for always
believing in education, for being there from
afar, for letting me find my future, for
your love.

RACHEL AND HER FAMILY

MY FAMILY

Bibliography

- Ansimov M. A., Sengers J. V., Sengers J. M. L., 2004, in Fernandez-Prini R., Harvey A., Palmer D., eds, , Aqueous Syst. Elev. Temp. Press.. Elsevier, Chapt. 2, p. 29
- Asimow P., 2015, in Schubert G., ed., , Treatise Geophys., 2nd edn, Elsevier, Chapt. 2.16, pp 393–416, doi:10.1016/B978-0-444-53802-4.00050-6
- Baraffe I., Chabrier G., Barman T. S., 2008, *A&A*, 482, 315
- Baraffe I., Chabrier G., Barman T. S., 2010, *Reports Prog. Phys.*, 73, 016901
- Baraffe I., Chabrier G., Fortney J., Sotin C., 2014, in Beuther H., Klessen R., Dullemond C., Henning T., eds, , Protostars Planets VI. University of Arizona Press, Tucson, pp 763–786, doi:10.2458/azu_uapress_9780816531240-cho33
- Barclay T., et al., 2013, *Nature*, 494, 452
- Barnes R., Raymond S. N., Greenberg R., Jackson B., Kaib N. A., 2010, *ApJ*, 709, L95
- Baruteau C., et al., 2014, in Beuther H., Klessen R., Dullemond C., Henning T., eds, , Protostars Planets VI. University of Arizona Press, Tucson, pp 667–690, doi:10.2458/azu_uapress_9780816531240-cho29
- Belonoshko A., Saxena S., 1991, *Geochim. Cosmochim. Acta*, 55, 381
- Bennett D. P., et al., 2010, *ApJ*, 713, 837
- Berta-Thompson Z. K., et al., 2015, *Nature*, 527, 204
- Beuzit J.-L., et al., 2008, in McLean I. S., Casali M. M., eds, Proc. SPIE 7014. SPIE, p. 701418, doi:10.1117/12.790120
- Bina C. R., 1998, *Earth, Planets Sp.*, 50, 1029
- Birch F., 1947, *PR*, 71, 809
- Bond J. C., O'Brien D. P., Lauretta D. S., 2010, *ApJ*, 715, 1050
- Broeg C., et al., 2013, *EPJ Web Conf.*, 47, 03005
- Burrows A., Orton G., 2010, in Seager S., ed., , Exoplanets. University of Arizona Press, Tucson, p. 419

- Cameron A. C., 2012, *Nature*, 492, 48
- Carter-Bond J. C., O'Brien D. P., Delgado Mena E., Israelian G., Santos N. C., González Hernández J. I., 2012, *ApJ*, 747, L2
- Charbonneau D., Brown T. M., Noyes R. W., Gilliland R. L., 2002, *ApJ*, 568, 377
- Charbonneau D., et al., 2005, *ApJ*, 626, 523
- Chen B., Hsieh W. P., Cahill D. G., Trinkle D. R., Li J., 2011, *PRB*, 83, 13201
- Choukroun M., Grasset O., 2007, *J. Chem. Phys.*, 127, 124506
- Consolmagno G. J., Lewis J. S., 1978, *Icarus*, 34, 280
- Crepp J. R., et al., 2011, *ApJ*, 729, 132
- Davies G. F., 1980, *J. Geophys. Res.*, 85, 2517
- Deming D., Seager S., Richardson L. J., Harrington J., 2005, *Nature*, 434, 740
- Deming D., et al., 2013, *ApJ*, 774, 95
- Demory B.-O., Gillon M., Madhusudhan N., Queloz D., 2016, *MNRAS*, 455, 2018
- Dolan D. H., Knudson M. D., Hall C. A., Deeney C., 2007, *Nat. Phys.*, 3, 339
- Dorn C., Khan A., Heng K., Connolly J. A. D., Alibert Y., Benz W., Tackley P., 2015, *A&A*, 577, A83
- Driscoll P. E., Barnes R., 2015, *Astrobiology*, 15, 739
- Dunaeva A. N., Antsyshkin D. V., Kuskov O. L., 2010, *Sol. Syst. Res.*, 44, 202
- Ehrenreich D., Lecavelier des Etangs A., Beaulieu J., Grasset O., 2006, *ApJ*, 651, 535
- Fei Y., Mao H.-k., Hemley R. J., 1993, *J. Chem. Phys.*, 99, 5369
- Feistel R., Wagner W., 2006, *J. Phys. Chem. Ref. Data*, 35, 1021
- Fischer D. A., Valenti J., 2005, *ApJ*, 622, 1102
- Fischer D. A., Howard A. W., Laughlin G. P., Macintosh B., Mahadevan S., Sahlmann J., Yee J. C., 2014, in Beuther H., Klessen R., Dullemond C., Henning T., eds, , *Protostars Planets VI*. University of Arizona Press, Tucson, pp 715–738, doi:10.2458/azu_uapress_9780816531240-cho31
- Fortney J. J., Nettelmann N., 2009, *Space Sci. Rev.*, 152, 423
- Fortney J. J., Marley M. S., Barnes J. W., 2007, *ApJ*, 659, 1661
- Fortney J. J., Baraffe I., Militzer B., 2010, in Seager S., ed., , *Exoplanets*. University of Arizona Press, Chapt. 17, pp 397–418
- Frank M. R., Fei Y., Hu J., 2004, *Geochim. Cosmochim. Acta*, 68, 2781

- French M., Mattsson T., Nettelmann N., Redmer R., 2009, *PRB*, 79, 054107
- Fressin F., et al., 2011, *ApJ SS*, 197, 5
- Gaidos E., Deschenes B., Dundon L., Fagan K., Menviel-Hessler L., Moskovitz N., Workman M., 2005, *Astrobiology*, 5, 100
- Grasset O., Schneider J., Sotin C., 2009, *ApJ*, 693, 722
- Guillot T., 1999, *Planet. Space Sci.*, 47, 1183
- Guillot T., 2010, *A&A*, 520, A27
- Guillot T., Gautier D., 2015, in Schubert G., ed., , *Treatise Geophys.*, 2nd edn, Elsevier, Chapt. 10.16, pp 529–557, doi:10.1016/B978-0-444-53802-4.00176-7
- Guillot T., Chabrier G., Gautier D., Morel P., 1995, *ApJ*, 450, 463
- Haghighipour N., 2011, *Contemp. Phys.*, 52, 403
- Hansen B. M. S., et al., 2008, *ApJ SS*, 179, 484
- Haswell C. A., 2010, *Transiting exoplanets*. Cambridge University Press
- Herrero C. P., Ramírez R., 2014, *J. Chem. Phys.*, 140, 234502
- Howe A. R., Burrows A. S., Verne W., 2014, *ApJ*, 787, 26
- Hubbard W. B., MacFarlane J. J., 1980, *J. Geophys. Res.*, 85, 225
- Hubbard W., Marley M. S., 1989, *Icarus*, 78, 102
- Husmann H., Sohl F., Spohn T., 2006, *Icarus*, 185, 258
- Husmann H., Sotin C., Lunine J., 2015, in Schubert G., ed., , *Treatise Geophys.*, 2nd edn, Elsevier, Chapt. 10.18, pp 605–635, doi:10.1016/B978-0-444-53802-4.00178-0
- Izidoro A., Raymond S. N., Morbidelli A., Hersant F., Pierens A., 2015, *ApJ*, 800, L22
- Kataria T., Sing D. K., Lewis N. K., Visscher C., Showman A. P., Fortney J. J., Marley M. S., 2016, *ApJ*, 821, 9
- Kipping D. M., Spiegel D. S., Sasselov D. D., 2013, *MNRAS*, 434, 1883
- Knudson M. D., Desjarlais M. P., Lemke R. W., Mattsson T. R., French M., Nettelmann N., Redmer R., 2012, *PRL*, 108, 091102
- Komacek T. D., Abbot D. S., 2016, *ApJ*, 832, 54
- Kruse A., Dinjus E., 2007, *J. Supercrit. Fluids*, 39, 362
- Kurosaki K., Ikoma M., Hori Y., 2014, *A&A*, 562, A80
- Levi A., Sasselov D., Podolak M., 2014, *ApJ*, 792, 125

- Lewis J. S., 1971, *Icarus*, 15, 174
- Lopez E. D., Fortney J. J., Miller N., 2012, *ApJ*, 761, 59
- Lyon S. P., Johnson J. D., 1992, Technical report, SESAME: The Los Alamos National Laboratory equation of state database. LA-UR-92 3407, Los Alamos National Lab
- Macintosh B., et al., 2014, *PNAS*, 111, 12661
- Madhusudhan N., 2012, *ApJ*, 758, 36
- Madhusudhan N., Redfield S., 2015, *Int. J. Astrobiol.*, 14, 177
- Madhusudhan N., Lee K. K. M., Mousis O., 2012, *ApJ*, 759, L40
- Madhusudhan N., Knutson H., Fortney J. J., Barman T., 2014a, in Beuther H., Klessen R., Dullemond C., Henning T., eds, , *Protostars Planets VI*. University of Arizona Press, Tucson, pp 739–762, doi:10.2458/azu_uapress_9780816531240-cho32
- Madhusudhan N., Crouzet N., McCullough P. R., Deming D., Hedges C., 2014b, *ApJ*, 791, L9
- Mao H.-K., Mao W., 2007, in Schubert G., ed., , *Treatise Geophys.*, 2nd edn, Elsevier, Chapt. 2.11, pp 263–291, doi:10.1016/B978-0-444-53802-4.00036-1
- Maxwell A., Ragozzine D., Rogers L. A., Seager S., Zeng L., 2011, *Bull. Am. Astron. Soc.*, 43
- Mayor M., Queloz D., 1995, *Nature*, 378, 355
- Mazevet S., Licari A., Chabrier G., Potekhin A. Y., 2018, preprint (arXiv:1810.05658)
- McCullough P. R., Crouzet N., Deming D., Madhusudhan N., 2014, *ApJ*, 791, 55
- Milone E. F., Wilson W. J. F., 2014, *Solar System Astrophysics: Background Science and the Inner Solar System*, 2nd edn. Springer, New York, doi:10.1007/978-1-4614-8848-4
- Monteux J., Golabek G. J., Rubie D. C., Tobie G., Young E. D., 2018, *Space Sci. Rev.*, 214
- Mordasini C., Alibert Y., Klahr H., Henning T., 2012a, *A&A*, 111, 1
- Mordasini C., Alibert Y., Georgy C., Dittkrist K.-M., Klahr H., Henning T., 2012b, *A&A*, 547, A112
- Mordasini C., Mollière P., Dittkrist K.-M., Jin S., Alibert Y., 2015, *Int. J. Astrobiol.*, 14, 201
- More R. M., Warren K. H., Young D. A., Zimmerman G. B., 1988, *Phys. Fluids*, 31, 3059
- Moriarty J., Madhusudhan N., Fischer D., 2014, *ApJ*, 787, 81
- Nettelmann N., Holst B., Kietzmann A., French M., Redmer R., Blaschke D., 2008, *ApJ*, 683, 1217
- Nettelmann N., Fortney J. J., Kramm U., Redmer R., 2011, *ApJ*, 733, 2

- Nimmo F., 2018, in , Vol. 1, Oxford Res. Encycl. Planet. Sci., pp 1–44,
doi:10.1093/acrefore/9780190647926.013.29
- Nimmo F., Pappalardo R. T., 2016, *J. Geophys. Res. Planets*, 121, 1378
- Ochiai H., Nagasawa M., Ida S., 2014, *ApJ*, 790, 92
- Oppenheimer B. R., et al., 2013, *ApJ*, 768, 24
- Owen J. E., Wu Y., 2016, *ApJ*, 817, 107
- Perryman M., 2011, *The Exoplanet Handbook*. Cambridge University Press, Cambridge,
doi:10.1017/CBO9780511994852
- Petigura E. A., Howard A. W., Marcy G. W., 2013, *PNAS*, 110, 19273
- Poirier J.-P., 2000, *Introduction to the Physics of the Earth's Interior*, 2nd edn. Cambridge University Press,
Cambridge, doi:10.1017/CBO9781139164467
- Poirier J.-P., Tarantola A., 1998, *Phys. Earth Planet. Inter.*, 109, 1
- Poirier J. P., Sotin C., Peyronneau J., 1981, *Nature*, 292, 225
- Porco C. C., et al., 2006, *Science*, 311, 1393
- Press W. H., 2007, *Numerical Recipes: The Art of Scientific Computing*, 3rd edn. Cambridge University
Press, Cambridge
- Rackauckas C., Nie Q., 2017, *J. Open Res. Softw.*, 5, 15
- Rauer H., et al., 2014, *Exp. Astron.*, 38, 249
- Raymond S. N., Kokubo E., Morbidelli A., Morishima R., Walsh K. J., 2013, in Beuther H., Klessen R.,
Dullemond C., Henning T., eds, , *Protostars Planets VI*. University of Arizona Press, Tucson, pp 595–618,
doi:10.2458/azu_uapress_9780816531240-cho26
- Redmer R., Mattsson T. R., Nettelmann N., French M., 2011, *Icarus*, 211, 798
- Ricker G. R., et al., 2014, in Oschmann J. M., Clampin M., Fazio G. G., MacEwen H. A., eds, *Proc. SPIE 9143*.
p. 914320, doi:10.1117/12.2063489
- Rodler F., Lopez-Morales M., Ribas I., 2012, *ApJ*, 753, L25
- Rogers L. A., 2012, Doctoral thesis, MIT
- Rogers L. A., 2015, *ApJ*, 801, 41
- Rogers L. A., Seager S., 2010a, *ApJ*, 712, 974
- Rogers L. A., Seager S., 2010b, *ApJ*, 716, 1208
- Ross M., Schubert G., 1989, *Icarus*, 78, 90

- Rothman L., et al., 2013, *J. Quant. Spectrosc. Radiat. Transf.*, 130, 4
- Salpeter E., Zanolzy H., 1967, *PR*, 158, 876
- Seager S., ed. 2010, *Exoplanets*. University of Arizona Press, Tucson
- Seager S., Deming D., 2010, *ARAA*, 48, 631
- Seager S., Kuchner M., Hier-Majumder C. A., Militzer B., 2007, *ApJ*, 669, 1279
- Senft L. E., Stewart S. T., 2008, *Meteorit. Planet. Sci.*, 43, 1993
- Snellen I. A. G., de Kok R. J., de Mooij E. J. W., Albrecht S., 2010, *Nature*, 465, 1049
- Soderlund K. M., Schmidt B. E., Wicht J., Blankenship D. D., 2014, *Nat. Geosci.*, 7, 16
- Sotin C., Grasset O., Mocquet A., 2007, *Icarus*, 191, 337
- Sotin C., Jackson J. M., Seager S., 2010, in Seager S., ed., , *Exoplanets*. University of Arizona Press, Chapt. 16, pp 375-395
- Soummer R., et al., 2009, in Shaklan S. B., ed., *Proc. SPIE 7440*. SPIE, p. 74400A, doi:10.1117/12.826702
- Southworth J., Wheatley P. J., Sams G., 2007, *MNRAS*, 379, L11
- Southworth J., Mancini L., Madhusudhan N., Mollière P., Ciceri S., Henning T., 2017, *AJ*, 153, 191
- Stewart S. T., Ahrens T. J., 2005, *J. Geophys. Res.*, 110, E03005
- Suess H. E., Urey H. C., 1956, *Rev. Mod. Phys.*, 28, 53
- Sugimura E., Komabayashi T., Hirose K., Sata N., Ohishi Y., Dubrovinsky L. S., 2010, *PRB*, 82, 134103
- Thomas S. W., Madhusudhan N., 2016, *MNRAS*, 458, 1330
- Thompson S. L., Lauson H. S., 1972, Technical report, Improvements in the Chart B radiation hydrodynamic code III: revised analytic equations of state. SC-RR-71 0714, Sandia National Lab
- Tingley B., Sackett P. D., 2005, *ApJ*, 627, 1011
- Tsitouras C., 2011, *Comput. Math. with Appl.*, 62, 770
- Unterborn C. T., Dismukes E. E., Panero W. R., 2016, *ApJ*, 819, 32
- Urey H. C., 1956, *PNAS*, 42, 889
- Valencia D., O'Connell R. J., Sasselov D. D., 2006, *Icarus*, 181, 545
- Valencia D., Sasselov D. D., O'Connell R. J., 2007, *ApJ*, 665, 1413
- Valencia D., Ikoma M., Guillot T., Nettelmann N., 2010, *A&A*, 516, A20
- Valencia D., Guillot T., Parmentier V., Freedman R. S., 2013, *ApJ*, 775, 10

- Vazan A., Kovetz A., Podolak M., Helled R., 2013, *MNRAS*, 434, 3283
- Verhoogen J., 1965, *Phil. Trans. R. Soc. A*, 258, 276
- Vinet P., Smith J., Ferrante J., Rose J., 1987, *PRB*, 35, 1945
- Wagner W., Pruß A., 2002, *J. Phys. Chem. Ref. Data*, 31, 387
- Wagner F., Sohl F., Hussmann H., Grott M., Rauer H., 2011, *Icarus*, 214, 366
- Wengenmayr R., 2012, *Max Planck Res.*, 3, 71
- Wilson H. F., Militzer B., 2012, *ApJ*, 745, 54
- Wilson H., Wong M., Militzer B., 2013, *PRL*, 110, 151102
- Wolszczan A., Frail D. A., 1992, *Nature*, 355, 145
- Wright J. T., Gaudi B. S., 2013, in Oswald T., French L., Kalas P., eds, , *Planets, Stars Stellar Syst.* Chapt. 59
- Zapolsky H. S., Salpeter E. E., 1969, *ApJ*, 158, 809
- Zeng L., Sasselov D. D., 2013, *PASP*, 125, 227
- Zeng L., Sasselov D., 2014, *ApJ*, 784, 96
- de Wit J., Seager S., 2013, *Science*, 342, 1473

