

ABSTRACT

Title of Dissertation: LEAVING FLATLAND: TOWARD 3D INFERENCE
OF EXOPLANET ATMOSPHERES
VIA SPECTROSCOPIC CHARACTERIZATION

Arjun B. McDanal-Savel
Doctor of Philosophy, 2026

Dissertation Directed by: Professor Eliza Miller-Ricci Kempton
Department of Astronomy

The vast majority of known transiting exoplanets are extreme worlds with no Solar System analogs. Take “hot Jupiters,” inflated gas giants orbiting their stars in less than 10 Earth days. Tidally locked, these planets feature one hemisphere that is constantly irradiated and another that is constantly dark. Complex, physics-rich simulations predict intertwined physics and chemistry in hot Jupiter atmospheres (cloud feedback, magnetic drag, and more), spanning micro to planetary scales. Yet the models currently used for interpreting spectra are comparatively limited. Such inference tools (“retrievals”) generally assume that hot Jupiters vary radially (if at all), neglecting spatial inhomogeneities such as stark day–night thermal contrasts. The rapidly increasing volume of high-quality spectra from both ground and space presents an opportunity: we can now shed our “1D assumption” and directly measure the 3D processes that set atmospheric inhomogeneity.

After reviewing the motivating physics and typical observational strategies relevant to my work, this dissertation begins with a set of forward modeling investigations. Chapter 2 focuses on

the case study of WASP-76 b, a now-canonical hot Jupiter. By computing spectra corresponding to a suite of 3D general circulation models (GCMs), I showed that the previously favored hypothesis, iron condensation, cannot explain the existing data. Rather, those high-resolution, cross-correlation spectroscopy (HRCCS) observations are better explained by a combination of clouds and a slight orbital eccentricity. The next chapter takes a broader view of 3D effects in HRCCS. I demonstrated that a chemically stable tracer can establish a “baseline” to disentangle the various potential causes of time-variable HRCCS data. I closed this portion of my work by illustrating how differences in measured gas velocity — a signature of 3D processes in planetary atmospheres — may also implicate magnetic fields, an important but rarely measured planetary property.

The middle portion of my dissertation is concerned with retrieval ingredients — namely, their computational burden, uncertainty budgets, and information content. I first showed that compressing gaseous absorption tables reduces simulation memory consumption by more than an order of magnitude. The subsequent chapter presents an HRCCS simulator that incorporates spectral contributions from the exoplanet, star, the Earth’s atmosphere, and the instrument. Using this tool, I found that variations in the Earth’s atmosphere can bias retrieved properties of exoplanet atmospheres. Next, I adapted HRCCS retrieval techniques to JWST spectra. Doing so allowed me to measure the abundance of CO, in turn enabling inference of robust C/O ratios and metallicities from exoplanet atmospheres using only a single instrument mode.

Finally, I established which types of spatial variation most strongly sculpt HRCCS retrievals. By simulating 1D retrievals on a hierarchy of 3D atmospheres, Chapter 8 demonstrates that 1D retrievals imprecisely approximate the distorted spectral lines of more realistic, 3D atmospheres. This chapter in particular — and the dissertation in general — paves the way for targeted retrieval upgrades, advancing efforts to measure the leading-order 3D effects in exoplanet atmospheres.

LEAVING FLATLAND: TOWARD 3D INFERENCE
OF EXOPLANET ATMOSPHERES VIA SPECTROSCOPIC
CHARACTERIZATION

by

Arjun B. McDanal-Savel

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2026

Advisory Committee:

Professor M. Coleman Miller, Chair
Professor Eliza M.-R. Kempton, Advisor
Professor Megan Weiner Mansfield, Committee Member
Professor Kayo Ide, Dean's Representative
Dr. Joshua D. Lothringer, External Examiner

© Copyright by
Arjun B. McDanal-Savel
2026

Preface

The research presented in Chapters 2–6 has been published as first-authored works in peer-reviewed journals; the text and figures appear here with only minor modifications. The work in Chapters 7 and 8 will be submitted to AAS Journals in the coming weeks as first-author articles.

Chapter 2 appeared in *The Astrophysical Journal* as “No Umbrella Needed: Confronting the Hypothesis of Iron Rain on WASP-76b with Post-processed General Circulation Models” ([Savel et al. 2022](#)). For this work, I spectrally post-processed general circulation models, analyzed the model spectra, and performed the orbital eccentricity analysis. My advisor provided the initial topic, radiative transfer code, and guidance throughout the project. Other coauthors provided general circulation models, opacity files, and feedback on the text.

Chapter 3 was published in *The Astrophysical Journal* as “Diagnosing Limb Asymmetries in Hot and Ultrahot Jupiters with High-resolution Transmission Spectroscopy” ([Savel et al. 2023](#)). I defined the asymmetry metric that this work relies on, generated the parameterized atmospheric models, and performed the chemical calculations. My advisor provided the topic and guidance throughout the project. Additional coauthors provided the climate models and feedback on the analysis methods and text.

The work in Chapter 4 was published in *The Astrophysical Journal Letters* as “A New Lever on Exoplanetary B Fields: Measuring Heavy Ion Velocities” ([Savel et al. 2024b](#)). For this study, I performed the magnetohydrodynamics analysis, 1D radiative transfer, and photochemistry. Coauthors provided feedback on the text.

Chapter 5 was published in *The Journal of Open Source Software* as “cortecs: A Python package for compressing opacities” ([Savel et al. 2024a](#)). I developed the underlying methodology, code, and documentation. Coauthors provided feedback on the text.

Chapter 6 was published in *The Astronomical Journal* as “Peering into the Black Box: Forward Modeling of the Uncertainty Budget of High-resolution Spectroscopy of Exoplanet Atmospheres” (Savel et al. 2025). I developed the mock observation simulator and performed the mock parameter inference (“retrievals”). My advisors provided the topic and guidance throughout the project. Coauthors provided the retrieval and data-cleaning framework, guidance on how to use it, and feedback on the text.

The work in Chapter 7 is adapted from a manuscript submitted to the *Astrophysical Journal Letters*. I devised the methodology and performed the retrievals presented in this chapter. My advisor provided guidance throughout the project. Coauthors provided re-reductions of the existing JWST data and feedback on the text.

Finally, Chapter 8 is a work in preparation. I generated the parameterized atmospheres and performed the mock retrievals. My advisor collaborated with me on identifying the topic and provided guidance throughout the project. Coauthors provided general circulation models and feedback on the text.

Dedication

To my partner. To my family.

Acknowledgments

Thank you, Eliza — for encouraging the rabbit holes, for supporting ambitious proposals, for kind course-correction when needed. You taught me to think about leading-order effects, to break down problems (terrestrial and otherwise) into their major constituents.

I am deeply grateful to my dissertation committee for their time and feedback. I appreciate that I have had such a wonderful learning opportunity right before I embark on my next steps as a scientist.

Thank you to Megan Bedell, whose advising at Flatiron brought me closer to the delights and horrors of real data. Your love of dogs is inspiring.

Thank you as well to the Flatiron Institute for providing me a research home away from research home. And to the scientists — particularly Julianne Delcanton and the Astro Data Group — who warmly welcomed me.

Thank you to the business office (especially Mona, Susan, and Dorinda) for helping me through all my silly questions. Travel has been such an extraordinary component of my time as a student, and you have made that possible.

Thank you to Lauren for shepherding me through the dissertation process.

And to my wonderful mentors, formal and informal, over the years: Tad, Courtney, Julia, Adam, Matt, Kaze, Wolfgang, Matej, Alex, Quang, and Emily. You showed me the myriad ways to be a scientist, to be rigorous and kind.

Thank you to those who started grad school with me in Fall 2020. Great people made odd times palatable.

Thank you, Jegug, for modeling appropriate skepticism, good taste, and the possibility of graduation.

I would be remiss to not acknowledge the spaces and places I've loved these six years — Art Bar, Kingfisher, Aroqa, Rock Creek Park, the scary steps between Brooklyn and Queens, Greenacre Park, the secluded woods in Hyattsville. And the in-betweens, in which so much of this dissertation was composed: in car (emails on the Jersey Turnpike, Riley's vat of kombucha strapped into the backseat), and bus (top floor, front row, with the floor-to-ceiling windows), and subway (squinting at papers on my phone as I zip past Fort Totten), and train (with bleary commuters on the Northeast Corridor; with jubilant tourists out west), and plane (furiously debugging in economy; Vegas at night trundles along a mile below). We'll get "boat" on the next go-around.

I thank my dear friends for suffering my complaints and constraints over the years. You have made me a better and better-rounded person.

suypachi and Jason Bappa and Sathya: Thank you for all the joy and curiosity you've injected into my life. Where would any of us be without "The Structure of Scientific Revolutions"?

Taara: The remarkable thing about having a younger sibling is that you grow up, and you wake up one day, and suddenly there is a thoughtful person in your life who knows your entire life. I appreciate all the wisdom you've given me.

Amma: Thank you for all the years of support, of academic enrichment, of lessons in empathy and exuberance. I'm inspired to live a life even a quarter as full of love and mirth as yours.

Riley, of course. Words won't do justice, of course. I'll write some anyway. Thank you for reading and substantially improving nearly everything I've written — but also for laughing with me, for humoring me, for thinking with me. I might still, to this day, be a business major if not for you. I've learned so much from you this past decade, from credit cards to dialectics. I strive to be the scientist you are, **and** I strive to be the person you are.

Table of Contents

Preface	ii
Dedication	iv
Acknowledgments	v
Table of Contents	vii
List of Tables	xi
List of Figures	xii
List of Abbreviations	xxiii
Chapter 1: Introduction	1
1.1 The first known exoplanets: hot Jupiters	1
1.2 Why do we study hot Jupiter atmospheres?	3
1.2.1 Hot Jupiter formation	3
1.2.2 Hot Jupiter chemistry	7
1.2.3 Hot Jupiter climate	9
1.3 Inferring the properties of exoplanet atmospheres	17
1.3.1 Staring at exoplanetary systems	18
1.3.2 Spectroscopy	20
1.3.3 Statistics	24
1.4 Spectroscopy, now at higher resolution	27
1.4.1 Isolating the planetary signal	28
1.4.2 Sensitivity to exo-atmospheric winds	31
1.5 Popping the spherical cow: Limits of the 1D retrieval paradigm	31
Chapter 2: No umbrella needed: Confronting the hypothesis of iron rain on WASP-76 b with post-processed general circulation models	34
2.1 Introduction	34
2.2 Model Description	38
2.2.1 GCM	38
2.2.2 Radiative transfer model	42
2.2.3 Suite of models	52
2.3 Analysis	52
2.3.1 Cross-correlation	52
2.3.2 Individual line-fitting	54
2.4 Results	55
2.4.1 Cross-correlation	55
2.4.2 Line-fitting	64

2.5	Discussion	67
2.5.1	Comparison to other work	67
2.5.2	Alternative explanations for the anomalous blueshift	72
2.5.3	Limitations of our model	76
2.6	Conclusions	79
Chapter 3: Diagnosing Limb Asymmetries in Hot and Ultrahot Jupiters with High-resolution Transmission Spectroscopy		83
3.1	Introduction	83
3.2	Selected drivers	84
3.2.1	Asymmetry metric	85
3.2.2	Model atmospheres	86
3.2.3	Equilibrium chemistry	92
3.2.4	Asymmetry metric: application	92
3.3	Selected diagnostics	95
3.3.1	Diagnostics for specific mechanisms	95
3.3.2	Phase bins	104
3.3.3	Full phase-resolved spectra	112
3.4	Conclusions	116
Chapter 4: A New Lever on Exoplanetary B Fields: Measuring Heavy Ion Velocities		119
4.1	Introduction	119
4.2	Magnetohydrodynamics	121
4.3	Light ions are not preferable probes of B fields	126
4.4	Proof of concept: Barium ion in an ultra-hot Jupiter	129
4.5	Conclusions	134
Chapter 5: cortecs: A Python package for compressing opacities		139
5.1	Introduction	139
5.2	Methods	142
5.3	Workflow	144
5.4	Benchmark: High-resolution retrieval of WASP-77Ab	144
Chapter 6: Peering into the Black Box: Forward Modeling of the Uncertainty Budget of High-resolution Spectroscopy of Exoplanet Atmospheres		148
6.1	Introduction	148
6.2	Methods	150
6.2.1	Baseline observation simulation pipeline and data processing	150
6.2.2	Perturbations to baseline pipeline	160
6.3	Results	167
6.3.1	K_P – V_{sys} maps	167
6.3.2	Likelihood sampling	172
6.3.3	Retrievals	173
6.4	Discussion	177
6.4.1	Confirmation of HRCCS principles	177

6.4.2	Noise and time-varying tellurics can strongly affect HRCCS inference . . .	177
6.4.3	Effects of PCA on the planetary signal	180
6.4.4	On spurious results	181
6.4.5	Caveats	183
6.5	Conclusions	185
Chapter 7: Precise Determination of the Metallicity and C/O of WASP-39b From a Single JWST Instrument Mode with Phase-Resolved Cross Correlation Retrievals 187		
7.1	Introduction	187
7.2	Methods	191
7.2.1	The data	192
7.2.2	The forward model	193
7.2.3	The likelihood function	197
7.2.4	The sampling	198
7.3	Results	201
7.3.1	Gas abundances and velocity parameters	201
7.3.2	C/O and metallicity	204
7.4	Discussion	205
7.4.1	Comparison to other work	205
7.4.2	What drives the signal quality?	207
7.5	Conclusions	208
Chapter 8: 1D on 3D: The impact of inhomogeneous atmospheric properties on high-resolution transmission spectroscopy retrievals 210		
8.1	Introduction	210
8.2	Methods	213
8.2.1	3D atmospheres	213
8.2.2	3D modeled spectra	222
8.2.3	Simulated data pipeline	224
8.2.4	1D forward model	226
8.2.5	Cross-correlation	227
8.2.6	High-resolution retrievals	228
8.3	Results	230
8.3.1	3D forward models	230
8.3.2	K_P – V_{sys} maps	232
8.3.3	Retrievals	236
8.4	Discussion	242
8.4.1	What are the 3D biases in high-resolution retrievals?	242
8.4.2	What is <i>safe</i> from 3D biases in high-resolution retrievals?	244
8.4.3	Are there any indications that a 1D retrieval is biased?	245
8.4.4	How should we update retrievals?	246
8.4.5	How do these “3D biases” compare to those from low-resolution spectroscopy?	248
8.4.6	How sensitive are K_P – V_{sys} maps to 3D effects?	249
8.4.7	What are the implications for the Extremely Large Telescopes?	252

8.4.8	Caveats	253
8.5	Conclusions	254
Chapter 9:	Conclusions and future work	256
9.1	Summary	256
9.2	Future directions	258
9.2.1	The 3D processes of exoplanet atmospheres with high-resolution spec- troscopy	258
9.2.2	High-resolution spectroscopy	261
9.2.3	Inferring the properties of exoplanet atmospheres	262
9.2.4	The study of exoplanet atmospheres	263
Appendix A:	Software used in this dissertation	265
Appendix B:	Facilities used in this dissertation	267
Appendix C:	Appendix from Chapter 2	268
Appendix D:	Appendix from Chapter 6	271
D.1	Optimization	271
D.1.1	Optimization with JAX	272
D.1.2	Using <code>flowMC</code>	273
Appendix E:	Appendix from Chapter 7	276
E.1	Benchmarking nested sampling	276
E.1.1	Posterior precision	278
E.1.2	Evidence accuracy	278
E.1.3	Runtime / efficiency	279
E.2	Full retrieval outputs	279
Appendix F:	Appendix from Chapter 8	282
F.1	A derivation of a transmission spectrum broadening kernel	282
F.2	Benchmarking against CHIMERA	284
F.3	Opacity sources	287
Bibliography		288

List of Tables

3.1	Example drivers of phase asymmetries of ultra-hot Jupiters	118
4.1	Ba II detections in exoplanet atmospheres	119
4.2	Parameters of the 1D ultra-hot Jupiter model presented in Section 4.4	127
7.1	Prior distributions for our retrieval parameters. $\mathcal{U}(a, b)$ denotes a uniform distribution between a and b . Samples in which the gas abundances sum to greater than unity are rejected.	200
8.1	Planetary Parameters	216
8.2	Summary of experiments performed	218
C.1	MCMC values	270
F.1	Opacity Sources and Citations	287

List of Figures

1.1	Exoplanets with measured masses and orbital periods, colored by equilibrium temperature if available (retrieved from the NASA Exoplanet Archive on April 8, 2026). 51 Pegasi b, the first exoplanet found around a Sun-like star, is marked with a golden star. Detection biases skew the distribution of known planets toward hot Jupiters (Dawson & Johnson 2018). These exoplanets drastically differ from the Solar System planets (shown, not to scale). It remains to be seen how well our Solar System-centric atmospheric theory extends to these extreme worlds.	2
1.2	Changing disk chemistry as a function of radius, from Öberg et al. (2011). As distance from the disk center increases, the disk becomes colder, and different gaseous species freeze out (water, carbon dioxide, and carbon monoxide “ice lines” are indicated here). As these molecules are removed from the gas phase, they deplete the remaining gas of those atoms and alter the proportions available as solids. Exoplanets forming at different positions in the protoplanetary disk, then, may retain compositional clues as to their origins.	6
1.3	Cloud-free hot Jupiter dayside thermal structures calculated with the HELIOS (Malik et al. 2017, 2019) 1D radiative–convective equilibrium code. The modeled planet is 51 Pegasi b-like but orbits a Sun-like star at $a = 0.018$ AU. Given the astrophysical setting, these plausible thermal structures are determined by gas opacity; including gases that strongly absorb optical light creates a thermal inversion (red). The hot deep atmosphere is due to the substantial interior entropy inferred for hot exoplanets (Thorngren et al. 2019a). The exact physics and chemistry that controls the transition between these two classes of thermal structures remains unknown.	10
1.4	The 3D abundance of H_2O (left) and thermal structure of WASP-76 b (right) predicted by the R–M GCM (Beltz et al. 2021) and FastChem (Stock et al. 2018). 0° is the substellar longitude; the center of each panel is the antistellar point (on the planetary nightside). Overlaid in gray in the left column are the longitudes and pressures probed by transmission spectroscopy (Section 1.3; Wardenier et al. 2022). These regions bracket sharp compositional gradients and are notably distinct from one another. This rich structure cannot be predicted by 1D climate models (Fig. 1.3) and motivates the search for 3D effects in hot Jupiters.	13
1.5	An outline of high-resolution cross-correlation spectroscopy.	19
1.6	Three model transmission spectra calculated with petitRADTRANS (Mollière et al. 2019). Varying one parameter at a time (i.e., increasing the atmospheric temperature) leads to stark changes, but other parameters complicate the picture. Here, an increase in the metal content of the atmosphere could masquerade as a moderate increase in temperature. Degeneracies such as this motivate the use of Bayesian computation.	25

2.1	A schematic diagram of the “toy model” introduced in Ehrenreich et al. (2020). WASP-76 b is pictured at three different phases (left to right: ingress, center of transit, and egress). Colder regions of WASP-76 b’s atmosphere, where Fe condenses and produces no absorption, are the primary regions probed at ingress. Over the course of its transit, the planet rotates and brings hotter regions, which contain gas-phase iron, into view for the observer. This geometry makes the strong combined velocity of the trailing limb ($\approx -5.3 \text{ km s}^{-1}$ from winds, $\approx -5.3 \text{ km s}^{-1}$ from rotation) observable as a Doppler blueshift by mid-transit. Atmosphere size not to scale.	38
2.2	1D temperature-pressure (T-P) profiles of WASP-76 b as output by our 3D GCMs, with the condensation curves of Mg_2SiO_4 , Fe, and Al_2O_3 from Mbarek & Kempton (2016) and terminators overplotted. Though the GCMs’ outputs extend to 100 bars, our radiative transfer is limited to the (plotted) region at pressures less than 1 bar. Every 20th profile is plotted for ease of visualization, resulting in 902 samples of 18048 total latitude-longitude points. Regions near the substellar point ($0^\circ / 360^\circ$) tend to be hotter on average than the rest of the planet, and regions near the antistellar point (180°) tend to be cooler on average than the rest of the planet. Significant scatter at a given longitude is controlled by latitudinal variation. The GCMs with deeper RCBs tend to have cooler nightsides than those with shallower RCBs, though all sets have regions in which Fe, Mg_2SiO_4 , and Al_2O_3 condense (assuming equilibrium chemistry). Models with weaker drag (i.e., longer drag timescales) tend to produce much greater asymmetry between the east and west planetary limbs.	39
2.3	Maps of atmospheric temperature at the terminator of WASP-76 b for weak ($1 \times 10^7 \text{ s}$) and strong ($1 \times 10^3 \text{ s}$) drag timescales and three phases (ingress, center of transit, and egress). The planet core is not shown to scale, and we restrict our visualization to the region that is probed by high-resolution transmission spectroscopy ($\lesssim 1 \text{ bar}$ — our model domain slightly exceeds 1 bar in some regions, as our post-processed atmosphere is interpolated onto a grid evenly spaced in altitude). Black dashed lines correspond to the millibar isobar. As also shown in Figure 2.5, the strong drag case enforces more overall east-west symmetry in the atmosphere of WASP-76 b. The temperature inversion that can be seen in Figure 2.2 is also present here — as would be expected for an ultra-hot Jupiter with high-altitude optical absorbers.	40
2.4	Representative model spectra computed in this study. The bottommost, gray spectrum is computed without including Doppler shifts from planetary winds and rotation and includes Fe I and Fe II as the sole opacity sources. The middle, gold spectrum is the same as the bottom, but now including Doppler effects. The top-most, maroon spectrum includes Doppler effects and all opacity sources discussed in Section 2.2.2.2. The inset axis is centered near the Ca II H line (393.4 nm); the vertical dashed lines indicate the peak of a representative Fe feature for the gray and gold spectra, respectively. In the optical range of WASP-76 b, the inclusion of Doppler effects broadens and blueshifts strong lines, such as those depicted.	44

2.5	Column-integrated surface density of Fe abundance (Σ_{Fe}) of WASP-76 b. Each figure column corresponds to an orbital phase (left to right: ingress, center of transit, and egress), and each row corresponds to a model run (top to bottom: strong drag, weak drag, strong drag with condensation, weak drag with condensation, weak drag with Fe clouds). All models shown have a deep RCB. One scale height clouds are drawn with a darker gray, and 10 scale height clouds are drawn with a lighter gray. (In practice, the latter do not achieve a full thickness of 10 scale heights because the upper atmosphere of WASP-76 b presents conditions incompatible with iron condensation — see text for more information.) The strong drag case is more symmetric, especially at center of transit, than the weak drag case.	51
2.6	Normalized cross-correlation functions of VO (gray), Fe I (black), and Ca II (gold). Note that while the CCFs of Fe I and Ca II have clearly defined peaks (indicated by vertical dashed lines), the VO CCF is very noisy, with no discernible peak value. Most of our computed CCFs are similarly bimodal in quality.	54
2.7	Representative planet-frame Doppler shifts as a function of orbital phase for our Fe-only models, as computed by our cross-correlations of forward-modeled spectra. Overplotted are data from the Ehrenreich et al. (2020) WASP-76 b observations and their corresponding 1σ error bars (corresponding to the top panel of that work’s Extended Data Figure 7). Our worst-fitting model consistent with our GCM has strong drag ($\tau_{\text{drag}} = 1e3$ s) and a shallow RCB (solid gray line). Conversely, our best-fitting model consistent with our GCM has weak drag ($\tau_{\text{drag}} = 1e7$ s) and a deep RCB (dashed teal line). Finally, our best-fitting model with 1 scale height clouds inclusive of condensation effects has weak drag, a deep RCB, and Al_2O_3 composition (dotted gray line).	56
2.8	Terminator line-of-sight velocity maps of our deep RCB WASP-76 b GCMs, including contributions from both winds and rotation, plotted in the style of Figure 2.3 (two drag timescales and three phases). Our (Fe) modeled clouds are overplotted as a gray band in the bottom row, for the 10 scale height implementation. Our clouds, as opposed to our wind field, is not a terminator slice, but rather a depiction of every (projected) sightline that intersects a cloud in the full 3D geometry. As in maps depicting other quantities, the line-of-sight velocities are more symmetric for the high-drag case, while they are more asymmetric for the weak-drag case. Stronger blueshifting is further present at egress than ingress for the weak-drag case.	57

2.9	Center-of-transit Doppler shifts for all models explored in this work. From left to right, we show the effect of increasing the drag timescale of our GCM (i.e., decreasing the effect of drag within the atmosphere). Within a given row, we include gas-phase condensation for some models (“True” vs. “False”), as described in Section 2.2.2.4. For the weakest drag timescales, we also explore the effect of placing optically thick clouds above different species’ condensation curves (Mg_2SiO_4 , Fe, and Al_2O_3), as described in Section 2.2.2.4. Cells representing models that are not computed are filled in black. The upper sets of models presented here only includes iron as a spectral opacity source (“Fe-only”), whereas the lower sets of models include all gas opacity sources discussed in Section 2.2.2.2 (“Full Species”). For our full species models, we explore the effect of cross-correlating against different single-opacity template spectra (e.g., “Full species: Fe” convolves against a full species spectrum against an Fe template) and increasing our clouds’ scale heights (left to right: 1, 5, 10 scale heights). The strongest blueshift is detected for the all-species model / Fe cross-correlation mask with condensation on, and 10 scale height thick Al_2O_3 clouds. Aside from the top row, all models are computed with a deep RCB. For reference, the observed transit-averaged Doppler shift identified by Tabernero et al. (2021) via Fe cross-correlation is $-8.25 \pm 0.25 \text{ km s}^{-1}$ for the first analyzed transit (T1) and $-8.75 \pm 0.56 \text{ km s}^{-1}$ for the second analyzed transit (T2). Line analysis by Tabernero et al. (2021) in T1 yielded Doppler shifts of $4.1 \pm 5.1 \text{ km s}^{-1}$ and $-4.4 \pm 2.5 \text{ km s}^{-1}$ for the Ca II H&K lines, respectively. T2 analysis produced Doppler shifts of $1.0 \pm 3.0 \text{ km s}^{-1}$ and $-2.1 \pm 1.9 \text{ km s}^{-1}$ for these lines.	61
2.10	Phase-resolved CCFs of our spectra including all opacity sources. Subplot (a) is the strong drag case, subplot (b) is the weak drag case, subplot (c) showcases the weak drag case for varying cloud treatments (only convolving with an Fe template), and subplot (d) compares the condensation on/off cases. Most species with detectable CCFs follow the Doppler trend of Fe I, with the notable exception of Ca II. Note that the CCF of Ca II is too noisy to fit during ingress in subplot (c). As in the Fe-only models, introducing thick clouds to our full-species spectrum reproduces the sharp jump in Doppler shift noted by Ehrenreich et al. (2020) and Kesseli & Snellen (2021), and the condensation effect alters the fit for condensed species depending on whether clouds are included; non-condensed species are unaffected by condensation treatments. Our best-fitting model (up to 10 scale height Al_2O_3 clouds, weak drag, condensation on, deep RCB) is shown in black in panel (d). . . .	65
2.11	Doppler shifts investigated by Tabernero et al. (2021) in comparison with the lines from the forward models in our study. Including 10 scale height, gray Fe clouds in our model (purple squares) blueshifts our lines to varying degrees across species. The weak drag models are slightly preferred over the strong drag case. Although endmembers of drag (teal circle for strong drag, teal diamond for weak drag) are plotted for clarity, intermediate drag seems to be slightly more consistent with the observed Doppler shift data overall.	66

2.12	Height of lines investigated by Tabernero et al. (2021) in comparison with the lines from representative forward models in our study. Our strong drag case is plotted with teal circles, whereas our weak drag case is plotted with teal diamonds. Most of our lines are consistent with the Tabernero et al. (2021) line heights — with the exception of the Ca II lines, which have observed heights consistent with an escaping atmosphere that extends beyond our model domain.	66
2.13	Similar to Figure 2.7, but now including the effect of a small eccentricity, both with (maroon lines) and without (teal lines) clouds. Circular orbit models are plotted with dashed lines; best-fit, eccentricity-inclusive models are plotted with solid lines; and MCMC posterior draws are plotted in faded lines. Our best-fit combination of eccentricity and longitude of periastron (solid teal) is consistent with observations (West et al. 2016; Fu et al. 2021). This work’s best explanation of the Ehrenreich et al. (2020) Doppler shift signature includes 10 scale height Al_2O_3 clouds, condensation, and a small eccentricity. The models shown here both include weak drag and a deep RCB.	73
3.1	Asymmetry (as defined in Equation 3.1) of all chemical species considered in this study in our parameterized atmospheres at an equilibrium temperature of 2200 K. These models do not self-consistently inflate the hotter limb of the parameterized model (i.e., they do not observe the “scale height effect”). The shading of each species represents the normalized temperature difference, $\tilde{\Delta T}$, across the two limbs of our parameterized atmospheres; the lightest boxes have $\tilde{\Delta T} = 0.1$, whereas the darkest have $\tilde{\Delta T} = 0.6$. For illustrative purposes, we color in green tick marks for species with detections noted in Guillot et al. (2022) (and including the recent CO_2 detection; Ahrer et al. 2022). We also draw a vertical line denoting 0 asymmetry. Without taking the scale height effect into account, positive ions form much more predominantly on the warmer limb (i.e., have negative asymmetry) than other species and reach the greatest asymmetry values.	89
3.2	Similar to Figure 3.1, but now including the scale height effect (inflating the hotter limb in our parameterized models). Now, all species have asymmetries that favor the hotter limb (negative asymmetry)—simply because the hotter limb subtends more solid angle on the sky. However, there still exists inter-species variability in asymmetry, implying that the scale height effect does not entirely swamp genuine differences in equilibrium chemistry across limbs. Furthermore, negative ions still have larger asymmetries than positive ions or neutral species.	90
3.3	Asymmetry (per Equation 3.1) for Sr II, Al, H_2O , and CO in our parameterized atmospheres. Our grid sweeps over equilibrium equilibrium temperature and normalized temperature difference across limbs, and includes models that observe the scale height effect (circles) and do not (squares). We find that species with strong temperature-dependent abundances (e.g., Sr II) are less dominated by the scale height effect than species with weaker temperature-dependent abundances.	91

3.4	Volume mixing ratio of CO as a function of pressure and temperature as calculated by <code>FastChem</code> . Overplotted are the onset of ultra-hot Jupiters (as defined by their dayside temperature; Parmentier et al. 2018), the CO/CH ₄ equivalency curve from Visscher (2012) as a function of pressure, the Fe condensation curve from Mbarek & Kempton (2016), and 1D temperature–pressure profiles for a hot Jupiter (WASP-39b) and an ultra-hot Jupiter (WASP-18b) as computed with <code>HELIOS</code> (Malik et al. 2017). Both the condensation curve and the equivalency curve are computed at solar metallicity. Considering the regime of ultra-hot Jupiter atmospheres, CO is a relatively stable chemical species.	97
3.5	Atmospheric Doppler shifts, which should remain in the HRCCS signal after the orbital motion is subtracted, as a function of orbital phase for our forward models. Shown are representative species that span Doppler shifts and are noted as potentially observable by Kesseli et al. (2022): Fe, Sr II, and Sc. Cloud-free models are represented with solid lines, whereas models with fully cloudy morning limbs are represented with dashed lines. The first half of transit (RV ₁) and second half of transit (RV ₂) Doppler shifts for Fe from Kesseli et al. (2022) are overplotted as horizontal lines, with width determined by observational errors. Our cloudy models are much more strongly blueshifted than their cloud-free counterparts, become less blueshifted over transit, and do not have significant CCF peaks at early phases.	102
3.6	Single-species (OH, CO, H ₂ O, HCN) 3D forward-modeled spectra of WASP-76 b. These spectra are simulated over the CARMENES waveband and resolution. Doppler effects are not included in these spectra, which are modeled at center of transit. H ₂ O is the dominant absorber in this bandpass, followed by OH. HCN exhibits no spectral features above the continuum for WASP-76 b in this bandpass.	106
3.7	CCFs of individual species averaged over two phase bins. Each column corresponds to different species (OH, CO, H ₂ O), and each row corresponds to different bin selection: without including ingress and egress, including the full transit, and only including ingress and egress. Central bars between the CCFs are colored blue if the difference between the CCFs is greater than optimal Doppler shift errors (1.55 km/s, in black; Kesseli et al. 2022); otherwise, they are colored red. In our models, CO only displays detectable CCF differences when only including ingress and egress. The SNR in each plot refers to the difference between the two phase bins’ CCF peaks relative to the optimal Doppler shift errors.	107
3.8	The net Doppler shifts of Kesseli et al. (2022) (error bars) as compared to this work’s models (crosses). The first phase bin is drawn thinner than the second phase bin; observed phase bins are connected by a dotted line for visibility’s sake. The species are ordered and colored by total observed detection SNR. Rows without crosses correspond to species that we could not recover via cross-correlation in our models. Our models are able to explain some species (e.g., Na), fail to explain others (e.g., Cr) and fail to detect yet others (e.g., K).	109

3.9	Modeled phase-resolved Doppler shifts for select NIR-absorbing species, with representative error bars (Ehrenreich et al. 2020) drawn on. We find that OH and H ₂ O have distinct Doppler signatures from CO; however, OH and H ₂ O have Doppler shifts that are indistinguishable from one another with current best-case error bars (e.g., Ehrenreich et al. 2020). Considering CO as a “baseline species” here allows one to better understand how H ₂ O and OH may change through the atmosphere.	111
3.10	Results of an investigation into anomalous Ca II blueshift between different model runs. In panel (a), it can be seen that forward models that include absorption due to Sc opacity yield a larger Ca II blueshift than models that lack Sc (Fe Doppler shift is included for comparison). Panel (b) illustrates the cause of this anomalous blueshift: a Sc line overlapping one line in the optical Ca II doublet. These results imply that overlapping line profiles can subtly contaminate calculated Doppler shifts.	113
4.1	Calculated properties of the model atmosphere described in Table 4.2. Left: the limb temperature–pressure structure calculated with HELIOS (Malik et al. 2017, 2019). Regions where the atmosphere is convective are highlighted in orange. Pressures generally probed in high-resolution transmission spectroscopy (e.g., Kempton et al. 2014; Gandhi et al. 2020; Hood et al. 2020; Wardenier et al. 2023) are filled in green. Middle: free electron volume mixing ratio calculated by FastChem equilibrium (in black; Kitzmann et al. 2024) and VULCAN disequilibrium (shaded blue; Tsai et al. 2021) chemistry. We run VULCAN over a wide range of parameters—hence, there is a width to this region. Right: ion–neutral velocity difference (Equation 4.11), calculated under equilibrium and disequilibrium chemistry assuming two B field strengths. The gold line indicates the approximate minimum detectable velocity difference via high-resolution spectroscopy (e.g., Ehrenreich et al. 2020). The width of the disequilibrium chemistry grid is due almost entirely to the changing input stellar flux.	126
4.2	The primary electron donors in our chemical calculations. The thick solid lines represent equilibrium chemistry. Thin solid lines represent VULCAN disequilibrium chemistry with the F5V stellar spectrum, and thin dotted lines represent disequilibrium chemistry with the HD 189733 stellar spectrum.	131
4.3	The stellar UV spectra input to VULCAN.	132
4.4	An R = 145,000 CHIMERA optical transmission spectrum of an ultra-hot Jupiter under chemical equilibrium. A spectrum with Ba II as the only spectrally active trace species is overplotted (gold), indicating the location and strength of Ba II absorption lines in our model atmosphere. Ba II lines in this model probe pressures near a microbar. We generate an analogous spectrum for neutral iron, as well (black).	133
5.1	The results of a simple toy model demonstrating a cortecs failure mode. The toy model is a single column, with the number of grid cells within the column varied (along the x-axis). For each number of cells, we calculate the difference in the optical depth τ as calculated using cortecs-compressed opacities (using the polynomial compression method) and the second-to-last wavelength point of native PLATON methane opacity.	141

5.2	Top panel: The original opacity function of CO (Rothman et al., 2010) (solid lines) and its <code>cortecs</code> reconstruction (transparent lines) over a large wavelength range and at multiple temperatures and pressures. Bottom panel: the absolute residuals between the opacity function and its <code>cortecs</code> reconstruction. Note that opacities less than 10^{-60} are not generally relevant for the benchmark presented here; an opacity of $\sigma_\lambda = 10^{-60}$ would require a column nearly 10^{35}m long to become optically thick at a pressure of 1 bar and temperature of 1000 K.	143
5.3	The posterior distributions for our baseline WASP-77Ab retrieval (teal) and our retrieval using opacities compressed by <code>cortecs</code> (gold).	146
6.1	Schematic of <code>scope</code> pipeline. This example uses a planetary emission spectrum calculated with the CHIMERA (Line et al. 2013) atmospheric model as input. This simulation is meant to reconstruct the IGRINS dataset used to identify the emission of WASP-77A b by Line et al. (2021). The slight discontinuity in the planetary emission tracks is due to a gap in exposures corresponding to these observations. The SNR of this simulation matches the SNR of the Line et al. 2021 observations (SNR=12.8); 9 other noise realizations span SNRs from 11.0 to 19.0.	151
6.2	The eigenweights of our PCA decomposition for IGRINS standard star observations for a single order (order 5 in the H band). In this case, the first two components clearly covary with airmass, and multiple components exhibit subtle sub-structure as as function of date.	161
6.3	Simulated telluric spectra, zoomed in on a narrow spectral region to display the effect of changing airmass (top) and date (bottom). Airmass is clearly the first-order effect, but some lines are particularly time-variable.	163
6.4	A comparison of our data-driven telluric reconstruction against ATRAN (Lord 1992) radiative transfer simulations over one IGRINS order.	165
6.5	The effect of removing different numbers of PCA components from HRCCS data on the CCF detection significance in our baseline results. This is calculated in noiseless simulations (top left), SNR = 250 simulations (top right), and SNR = 60 simulations (bottom left) with 10 different noise realizations. Including more spectral components (e.g., the stellar spectrum) severely decreases the planet’s detection significance, but a small number of PCA components can recover the planet detection about the detection threshold. These significance values are compared to a commonly accepted HRCCS detection threshold (SNR=4.5; e.g., Cabot et al. 2019).	168
6.6	The spectral components learned by four PCA components from a simulated HRCCS dataset that only includes the planetary signal and instrumental blaze function. The PCA algorithm is able to approximate the diagonal trace of the planetary atmosphere with a superposition of a few strong stationary components. When other spectral components are present (including photon noise), PCA is not able to approximate the planetary signal, as it is most sensitive to the time-invariant signals with greater variance across wavelengths.	169

6.7	A set of spurious HCN signals in our baseline simulations. The maxima of these K_P-V_{sys} are denoted with a white “x.” While HCN is not included in this set of models, its signal can be found in cross-correlation when using ATRAN tellurics (top left), data-driven tellurics (top right), with low photon noise (top row), with high photon noise (bottom row), and in different noise realizations (bottom left).	170
6.8	Results of an MCMC fit to the velocity parameters (in km s^{-1}) and scale factor for simulated HRCCS data at two photon noise levels. This model used four PCA components to remove the stellar, telluric, and instrumental features in the simulated data. Using PCA in this experiment did not bias the scale factor toward smaller values, indicating that using PCA alone does not scale the planetary signal in HRCCS data.	174
6.9	Gas and velocity constraints from our <code>PyMultiNest</code> high-resolution atmospheric retrieval. Including strongly time-dependent tellurics (maroon) yields distinct posteriors from our time-dependent telluric (teal) model.	175
6.10	Gas constraints from our baseline retrieval on simulated data of WASP-77A b. We are able to recover gas constraints about as well as Line et al. (2021), with bounded constraints on H_2O , CO, and the CO isotope ratio in the lower-noise case (SNR=250). The higher-noise case (SNR=60) substantially changes the likelihood surface for the gas constraints shown, with less of an impact on velocity and temperature–pressure parameters.	176
7.1	Schematic comparing our retrievals (left) and standard JWST retrieval methodology (right). The retrievals that we present in this work begin with a different data product and fit for additional parameters with a different likelihood function.	194
7.2	Our native pixel resolution cross-correlation constraints on gases in WASP-39 b using G395H data vs. those acquired with standard methods ($R = 100$, phase-integrated). Our method is sensitive to previously hidden gases’ abundance. Specifically, we obtain a bounded abundance for CO. Vertical dashed lines mark the 16th and 84th percentiles for each distribution.	199
7.3	The (volatile-proxy) envelope metallicity and C/O ratio of WASP-39 b, as inferred by our native pixel resolution, cross-correlation, phase-resolved free retrieval (purple). With a single NIRSpec/G395H transit, we are able to simultaneously constrain both chemical parameters. As points of comparison, we also show the posterior distributions obtained without including the contribution of CO abundance (yellow) and using standard JWST retrieval techniques (gray). 1σ and 2σ contours are outlined in black for all three posterior distributions. The orange star indicates the solar values. The ticks on the marginals mark the 16th and 84th quantiles of each distribution.	206

8.1	Schematic of our simulation pipeline. We begin with a 3D model of the atmosphere (either a GCM output or a simplified atmospheric structure), then post-process it with 3D ray-striking radiative transfer to yield a transmission spectrum. We then simulate an observation of this spectrum using the <code>scope</code> package (Savel et al. 2025). Next, we generate a 1D forward model and cross-correlate it against our simulated dataset. Finally, we perform parameter inference on the simulated dataset with a full HRCCS retrieval.	214
8.2	The zonal-mean zonal wind plots of the two GCMs used in this work. The hot Jupiter (left) has a faster and more spatially extended jet than the ultra-hot Jupiter (right), in large part because the latter is damped via coupling to an intrinsic magnetic field (Showman & Polvani 2011; Beltz et al. 2021).	215
8.3	Simulated transmission spectra corresponding to a subset of the model grid described in Table 8.2. These spectra are calculated at a resolution of $R = 10^6$ for visualization purposes.	229
8.4	K_P – V_{sys} maps from our grid of models. The text in the top left of each panel refers to K_P offsets from the known orbit of the planet and proper motion of its host star, the top right analogous offsets in V_{sys} , and bottom left signal-to-noise ratios. For each planet, we consider three families of spatial variation: temperature, chemistry, and atmospheric dynamics (see Table 8.2). We additionally include cross-correlation against a fully 3D model — i.e., the output from a GCM.	233
8.5	The retrieved and input temperature–pressure profiles. Retrieved isotherms (mean and 1σ distributions) are shown in vertical bars. The input temperature–pressure profiles are color-coded by morning and evening limb where applicable. For the GCM-input temperature structure, only limb profiles are shown.	238
8.6	Select retrieved parameters from the HD 209458b chemical equilibrium model. Input values are indicated with solid black lines. We show the biases induced in the retrieved water, rotational velocity, and cloud-top pressure parameters. This experiment is repeated for two different noise instances (Retrievals 1 & 2).	240
8.7	Select retrieved parameters from the inhomogeneous winds and rotation models. Input values are shown with a solid black line. We show the biases induced in the K_P , V_{sys} , and rotational velocity parameters.	251
C.1	Posterior corner plot corresponding to our 10 scale height Al_2O_3 cloud MCMC run. Both priors result in similar posteriors, implying that the posterior is heavily influenced by the data. Quantities are reported for the asymmetric normal prior (drawn from the Fu et al. 2021 constraints).	270
D.1	The learned normalizing flow topology as compared with <code>PyMultiNest</code> equally weighted samples from a baseline retrieval. The normalizing flow is able to learn the topology of the posterior distribution. Note that this exercise does not constitute actual parameter inference—rather, we are demonstrating the flexibility of the global sampling scheme to approximate the final global posterior distribution. . . .	275
E.1	The posterior distributions output by <code>nautilus</code> and <code>MultiNest</code> for our PLATON benchmarking exercise.	277

E.2	The result of benchmarking the <code>nautilus</code> code against <code>pymultinest</code> . The estimated natural log of the Bayesian evidence is shown as a function of normalized runtime ($\tau_{\text{norm}} = 2000\tau_{\text{runtime}}/n_{\text{livepoints}}$). The nested sampling (NS) evidence and nested importance sampling (NIS) evidence as calculated by <code>MultiNest</code> are both shown.	278
E.3	Posterior distribution of our cross-correlation retrievals, with gases and cloud deck pressure fit in $\log_{10}$. We label the relative log abundance of the minor CO isotopologue as “ISO.”	280
E.4	Posterior distribution of our traditional retrievals.	281
F.1	A cosine transform of a uniform distribution, with the analytic solution from Eq. F.7 overplotted.	285
F.2	Benchmarking CHIMERA and our 3D forward model.	286

List of Abbreviations

1D	One-dimensional
3D	Three-dimensional
CCF	Cross-correlation function
C/O	Carbon-to-oxygen ratio
ELT	Extremely Large Telescope
ESPRESSO	Echelle SPectrograph for Rocky Exoplanets and Stable Spectroscopic Observations
GCM	General Circulation Model
HD	Henry Draper (catalog of stars)
HJ	Hot Jupiter
HRCCS	High-resolution, cross-correlation spectroscopy
HST	Hubble Space Telescope
IGRINS	Immersion GRating INfrared Spectrometer
JWST	James Webb Space Telescope
LOS	Line-Of-Sight
NASA	National Aeronautics and Space Administration
NIRSpec	Near-Infrared Spectrograph
PCA	Principal components analysis
RAM	Random access memory
RCB	Radiative-Convective Boundary
RV	Radial Velocity
SVD	Singular Vector Decomposition
TESS	Transiting Exoplanet Survey Satellite
UHJ	Ultra-Hot Jupiter
VLT	Very Large Telescope
WASP	Wide Angle Search for Planets (catalog of stars)

Chapter 1: Introduction

1.1 The first known exoplanets: hot Jupiters

Planetary scientists in 1994 had ample reasons to be pleased. After decades of work, a hard-won theory of the Solar System’s formation had emerged. There were still details to work out — how to make Mercury, for instance — but the field could elegantly explain a vast array of observations. An example is one of the most obvious features in the Solar System: models explained how gaseous planets formed far from the Sun, and rocky planets formed closer ([Mizuno 1980](#); [Bodenheimer & Pollack 1986](#)). Some even called this growing consensus a “Standard Model of Planet Formation” ([Wetherill 1990](#)).

Parallel to the effort to understand our Solar System was an effort to find other such systems — i.e., planets orbiting other stars. Astronomers had cataloged stars for thousands of years, galaxies for (arguably) centuries, and the cosmic microwave background for decades, yet by the early 1990s we had found only planets that orbited our Sun. Such searches were powered by a strong intuition: why would we be so lucky to be the only Solar System? Given our understanding of planet formation at this point, there was no reason to think that the processes that form our system would not be universal.

Intuition does not necessarily translate to detection, and this dissonance was reflected in the difficulty of finding exoplanets. Given that we knew from firsthand experience that the Solar System could form, many initial searches explicitly targeted Solar System analogs.¹ The observational signatures of planets like those in our Solar System are quite small; this fact, in addition to the long

¹By 1994, we had already found the first planetary-mass objects outside the Solar System ([Wolszczan & Frail 1992](#)). But these objects are quite odd. They orbit a low-mass stellar corpse — a “pulsar.” This discovery, while especially exciting to those in the burgeoning exoplanet community, was not so shattering as later detections, as that the pulsar planets provided little contextual information about the Solar System. That is, they did not tell us whether or not the Solar System is unique. The search continued for planets that orbited stars similar to our Sun.




Figure 1.1: Exoplanets with measured masses and orbital periods, colored by equilibrium temperature if available (retrieved from the NASA Exoplanet Archive on April 8, 2026). 51 Pegasi b, the first exoplanet found around a Sun-like star, is marked with a golden star. Detection biases skew the distribution of known planets toward hot Jupiters (Dawson & Johnson 2018). These exoplanets drastically differ from the Solar System planets (shown, not to scale). It remains to be seen how well our Solar System-centric atmospheric theory extends to these extreme worlds.

temporal baselines required to detect Solar System analogs, led to decades of non-detection and the occasional spurious signal.

1995 marked the discovery of the first exoplanet orbiting a Sun-like star (Mayor & Queloz 1995). This object was not an ice giant like Neptune, a barren rock like Mars, or a water-rich, temperate world like Earth. It was a **hot Jupiter** — hotter than Venus, larger than Jupiter, on a tighter orbit than Mercury, and a direct challenge to our understanding of planetary bodies (Fig. 1.1). That first object, 51 Pegasi b, orbits its star in just 4.2 Earth days and maintains a staggering temperature of ≈ 1300 Kelvin (1880° Fahrenheit). It is “spin-synchronized” — the gravity is so strong in this tight orbit that one side of the planet experiences permanent day, and the other permanent night.²

Decades later, much of exoplanet science still seeks to understand how such extreme planets

²This state is also referred to as “tidal locking,” which more generally refers to spin–orbit resonances such as 3:2.

form, what their chemical makeup may be, and how energy is redistributed within them. The nature of these astrophysical objects has broad ramifications. Perhaps most pressingly, hot Jupiters strongly test our Solar System models, helping us determine whether our theories are genuinely predictive or simply tuned to a small sample of 8 well-measured objects. My dissertation focuses on the models that we use to measure the properties of hot Jupiter atmospheres: their temperatures, compositions, and wind speeds. I argue that these models are incomplete. Without improving them, we will not be able to understand either the exoplanet population or our Solar System in context. To understand these connections, we must examine three key domains: planetary formation, atmospheric chemistry, and atmospheric dynamics.

1.2 Why do we study hot Jupiter atmospheres?

A mildly opportunistic reason to examine hot Jupiter atmospheres is simply because we can; these objects are quite observationally favorable with current technology. However, they are also very interesting from physics and chemistry perspectives.

1.2.1 Hot Jupiter formation

51 Pegasi b faced substantial scrutiny upon its discovery. But as subsequent observations identified many more such planets ([Butler et al. 1997, 1998](#); [Fischer et al. 1999](#); [Mazeh et al. 2000](#)), the scientific community had to confront the implications of the existence of such odd objects. A primary reason that the existence of hot Jupiters initially seemed implausible to some scientists — and exciting to others — was that their formation seemed unlikely. To appreciate this point, let us examine how planets are expected to form.

Imagine a large, roughly spherical ball of gas — a molecular cloud in the Milky Way. It rotates slowly in some direction due to its previous history of interactions in the galaxy. Once it cools enough, thermal pressure in its interior can no longer counteract gravity, and the cloud begins to collapse. As it collapses, it rotates faster; think of a figure skater pulling in their arms. And as it rotates faster, it collapses into a plane; in the radial direction, the disk is centrifugally supported, whereas it can collapse more freely along the rotation axis.³ We then arrive at a protoplanetary *disk* orbiting a young star, which ignites once the density at the core is great enough to kick-start nuclear fusion. This disk is sustained for up to 10 Myr before dissipating (Ribas et al. 2015).

In this disk of gas and dust, two main pathways have been proposed for forming planets: gravitational instability (e.g., Boss 2001) and core accretion (e.g., Pollack et al. 1996). In the gravitational instability paradigm, planets form much like stars do. This picture maintains that overdense clumps of gas collapse upon themselves due to self-gravity. In contrast, under core accretion, planets are built “bottom-up”: solids collide, stick, accrete, and grow until they reach enough mass to gravitationally capture gas from the disk (“runaway gas accretion”).

Let us now assess why in-situ formation appeared so unlikely for hot Jupiters. Gravitational instability may proceed if Toomre’s Q parameter dips below 1 (Safronov 1960; Toomre 1964):

$$Q = \frac{c_s \Omega}{\pi G \Sigma}, \quad (1.1)$$

where Σ is the surface density of the disk, G is the gravitational constant, c_s is the sound speed, and Ω is the orbital angular velocity. Plugging in typical values near 51 Pegasi b’s orbit and following expectations from the Solar System’s planet-forming disk (Weidenschilling 1977) yields $Q \approx 200$

³Real protoplanetary disks do not exist in isolation. For instance, nearby gaseous material can substantially torque disks, leading to warped morphologies (Thies et al. 2011; Kuffmeier et al. 2021; Gupta et al. 2023).

— well into the stable regime. The problem with forming planets here is that the gas in this region is very hot (stabilizing the gas against collapse), and the orbital rate around the star is very fast (also stabilizing the gas against collapse). It therefore seems unlikely that gravitational instability works well at hot Jupiters’ current orbits.

Now for core accretion, let us consider how much mass can be accreted near 51 Pegasi b’s current orbit. The isolation mass, M_{iso} , describes how much mass a forming planetary body can gather from the protoplanetary disk near its orbit ([Pollack et al. 1996](#)):

$$M_{iso} \sim \frac{a^3 \Sigma^{3/2}}{M_*^{1/2}}. \quad (1.2)$$

Plugging in appropriate values yields an isolation mass just 10% of Earth’s mass — not close to the ~ 10 Earth masses required for runaway gas accretion ([Rafikov 2006](#)). In this case, the forming planet is too close to its host star, so the annulus in which it resides contains too little planet-forming material. This exercise demonstrates that it would be difficult to build a hot Jupiter via core accretion at current-day semimajor axes.

Caveats to the above arguments certainly do exist, and in-situ formation is not entirely ruled out for all hot Jupiters (e.g., [Bodenheimer et al. 2000](#); [Batygin et al. 2016](#); [Boley et al. 2016](#)). But the prevalence of hot Jupiters provided ample motivation to find other, more efficient ways to create such objects. The trick seems to lie in planetary *migration* ([Lin et al. 1996](#)); rather than being formed where they are, hot Jupiters are formed farther from their host stars (presumably at distances to where our own giant planets are found), are delivered to their current locations by some mechanism, and are parked there by some (perhaps related) mechanism. The details are still under debate, but there is substantial evidence that at least one migration pathway is at play




Figure 1.2: Changing disk chemistry as a function of radius, from [Öberg et al. \(2011\)](#). As distance from the disk center increases, the disk becomes colder, and different gaseous species freeze out (water, carbon dioxide, and carbon monoxide “ice lines” are indicated here). As these molecules are removed from the gas phase, they deplete the remaining gas of those atoms and alter the proportions available as solids. Exoplanets forming at different positions in the protoplanetary disk, then, may retain compositional clues as to their origins.

([Dawson & Johnson 2018](#)).

Planetary migration is notoriously difficult to model in detail ([Morbidelli & Raymond 2016](#)). Studying the atmospheres of planets may help; the considerable information content of atmospheric observations ([Batalha & Line 2017](#)) implies that they could play a role in resolving questions about formation. In particular, the formation location of hot Jupiters may leave imprints on their current chemical compositions. Static protoplanetary disk models predict that planets that form beyond the so-called “ice lines” — regions in the disk where different gases freeze out — will more easily accrete those elements as solid material. A planet that forms beyond the water ice line, for instance, may be enriched in oxygen (Fig. 1.2; [Öberg et al. 2011](#)). Hot Jupiters are especially privileged among exoplanets in this regard. Many host atmospheres hot enough to sustain refractory species in the gas phase, and the comparison of refractory to volatile species can yield

additional hints as to planet formation ([Lothringer et al. 2021](#); [Chachan et al. 2023](#)). This is a very active area of research, and a great deal of effort is being spent to make these connections clearer in light of the complex and sometimes stochastic nature of planet formation (e.g., [Mordasini et al. 2016](#); [Mollière et al. 2022](#); [Feinstein et al. 2025](#); [Boss 2026](#); [Li et al. 2026](#)).

The formation of hot Jupiters is therefore still far from solved. Importantly, this problem is related to how planetary systems form *in general*. Giant planets are the “movers and shakers”⁴ of planet formation: they are massive enough to carve gaps in planet-forming disks, block inward-streaming dust, scatter other planetary-mass bodies, and migrate inward and outward ([Raymond & Morbidelli 2022](#)). Indeed, Jupiter’s migration may have sculpted our own Solar System’s architecture and enabled the delivery of Earth’s oceans ([Walsh et al. 2011](#); [O’Brien et al. 2014](#); [Brasser et al. 2018](#)). Given their elevated importance in planet formation, it behooves us to understand how giant (exo)planets form and migrate.

1.2.2 Hot Jupiter chemistry

If we accept that hot Jupiters do form, what might they be made of? Smaller planets can be more complicated,⁵ but we generally assume that hot Jupiters are balls made of predominantly hydrogen and helium. One can appeal to a variety of formation arguments to arrive at this conclusion in detail, but it suffices to say that these two gases are far more cosmically abundant than others. So, assembling a large volume of gas will necessarily gather quite a bit of hydrogen and helium.

What about the other, trace species? The simplest approach would be to assume that the material exists in the same proportions as in the host star, as both planet and star formed from

⁴I first heard this framing from Dr. Marta Bryan.

⁵Higher metal fractions from formation, outgassing, atmospheric escape, and bombardment by icy bodies are some muddling factors, to name a few.

the same cloud of material. The processes involved in planet formation may severely complicate this picture. Under the core accretion paradigm, there should exist an inverse correlation between a planet’s mass and its *metallicity*, the abundance of elements heavier than helium ([Ida & Lin 2005](#)); a simple argument for this correlation is that higher-mass objects dilute their metal mass with nebular hydrogen in greater quantities.

While formation sets the initial atomic abundances of the system, atmospheric chemistry controls what molecules, atoms, and ions we observe today. The equilibrium chemistry approach is reasonable for a first-order estimate of the proportions of these different chemical species. Each chemical reaction rate has a change in Gibbs free energy (the maximum “usable” work a system can perform) associated with it. We can therefore balance the reaction rates in a way that minimizes the whole system’s Gibbs free energy. This is the equilibrium state: if perturbed, the system would seek to return to its energy minimum ([Gibbs 1878](#)).

This equilibrium state is modified by *disequilibrium* chemical processes. Vertical motions in the atmosphere can “quench” chemistry — if the mixing rate is faster than the chemical reaction rate, the composition of gas at some pressure can actually be controlled by the composition of gas at some far deeper pressure, instead of local conditions (e.g., [Prinn & Barshay 1977](#); [Moses et al. 2011](#); [Miles et al. 2020](#)). Photochemical processes can further modify the lowest pressures in the atmosphere; for instance, high-energy photons can dissociate molecules, whose constituents then rapidly form chains that they otherwise would not. Through vertical mixing, these species can then impact lower regions of the atmosphere than where the photolysis reactions actually take place ([Yung et al. 1984](#); [Gautier & Owen 1989](#); [Zahnle et al. 2009](#); [Miller-Ricci Kempton et al. 2012](#); [Morley et al. 2013](#); [Tsai et al. 2021](#)). Although equilibrium chemistry is a sensible first guess of an atmosphere’s composition, this diversity of atmospheric processes can pull an

atmosphere’s true chemical state far from those predictions. Hot Jupiters may be less affected by these processes, though. Because their atmospheres are so hot, and because chemical reactions are faster at higher temperatures, these are the least likely exoplanets to be strongly influenced by disequilibrium effects.

Hot Jupiters therefore host rich chemical environments, offering an opportunity to test our chemical models in an extreme regime. Chemistry is a very stiff problem (e.g., [Hu et al. 2012](#); [Wogan et al. 2025](#)) — missing even a few reaction rates can produce greatly different predictions. Furthermore, it is not yet clear how our understanding of photochemistry and mixing translates to planets orbiting stars with different energy outputs from the Sun and driven into distinct circulation regimes. It is therefore crucial that we take advantage of opportunities such as these exoplanets to refine our chemical understanding of atmospheres in general.

1.2.3 Hot Jupiter climate

Relatively simple physical arguments can demonstrate the possible complexity of hot Jupiter atmospheric structure. Let us begin from the most straightforward possible setup: an isolated, isothermal blob of gas.

1.2.3.1 1D Thermal structure

Imagine a planetary-mass ball of gas at the same temperature (i.e., isothermal) alone in a vacuum. Such a self-gravitating sphere would eventually settle into hydrostatic equilibrium: the force due to a gradient of pressure in the atmosphere being equally opposed by the gravitational force. The description of this body’s large-scale climate would essentially end here.




Figure 1.3: Cloud-free hot Jupiter dayside thermal structures calculated with the HELIOS (Malik et al. 2017, 2019) 1D radiative-convective equilibrium code. The modeled planet is 51 Pegasi b-like but orbits a Sun-like star at $a = 0.018$ AU. Given the astrophysical setting, these plausible thermal structures are determined by gas opacity; including gases that strongly absorb optical light creates a thermal inversion (red). The hot deep atmosphere is due to the substantial interior entropy inferred for hot exoplanets (Thorngren et al. 2019a). The exact physics and chemistry that controls the transition between these two classes of thermal structures remains unknown.

Planets orbiting stars, however, are not kept isothermal in a vacuum. Rather, they are radiatively forced by their host stars' light. From pressures between roughly 1 bar and 1 μ bar in the planet's atmosphere,⁶ a range of possible thermal structures is allowed, depending on where this stellar energy is deposited (e.g., [Hubeny et al. 2003](#); [Fortney et al. 2008](#); [Malik et al. 2019](#)). If the atmosphere is composed of gas that very strongly absorbs stellar radiation, the energy will be deposited high in the atmosphere; this leads to a thermal inversion, or stratosphere. Energy balance dictates that such upper-atmospheric heating is accompanied by lower-atmospheric cooling. This feature is analogous to Titan's anti-greenhouse effect, wherein thick hazes heat the upper atmosphere and cool the surface ([McKay et al. 1991](#)). If there are no such optical absorbers in the atmosphere, the thermal structure will (roughly) monotonically decrease with increasing altitude (Fig. 1.3) This is the greenhouse effect: infrared absorbers prevent outgoing thermal radiation from escaping to space, warming the lower layers relative to the upper atmosphere.

Stellar radiation is not the only source of heat for planetary bodies; energy upwells from deep within the planet, as well. Firstly, planet formation is a violent process ([Canup 2004](#)), and heat from those initial collisions is slow to dissipate from the deepest regions of planets. Furthermore, radioactive decay of elements sequestered in the interior can provide energy over long timescales ([Davies & Davies 2010](#)). Hot gas giants feature an additional, mysterious source of interior heat flux. While its effect on the planetary population is evident (e.g., [Thorngren et al. 2019a](#)), its exact mechanism remains unclear to this day ([Thorngren 2024](#)). Compelling options include delayed contraction (e.g., [Chabrier & Baraffe 2007](#)), Ohmic dissipation (e.g., [Batygin & Stevenson 2010](#)), hydrodynamic mechanisms (e.g., [Sainsbury-Martinez et al. 2019](#)), and tidal effects (e.g., [Leconte et al. 2010](#)). Regardless of its physical source, high interior flux produces an interior thermal profile

⁶I will return in the final chapter of this dissertation to the question of lower pressures.

that increases in temperature as pressure increases. And if the conditions for convection are met, for instance the Schwarzschild Criterion ([Schwarzschild 1958](#)):

$$-\frac{dT}{dz} < \frac{g}{c_p}, \quad (1.3)$$

for temperature T as a function of altitude z , gravity g , and specific heat capacity c_p , then large-scale convective motions are driven within the atmosphere. These motions pin the temperature structure along that region to an *adiabat*, in which parcels in the atmosphere move without exchanging heat with their surroundings. Within the convective regions, the thermal structure follows the dry adiabatic lapse rate as a function of pressure, P , and gas mass density ρ :

$$\frac{dT}{dP} = \frac{T}{\rho c_p P}. \quad (1.4)$$

We therefore arrive at a typical predicted temperature–pressure structure for a hot gas giant atmosphere (Fig 1.3): a deep convective adiabat that smoothly connects to an upper atmosphere, which is largely controlled by the interplay of opacity (i.e., chemical mixtures) and radiation.

1.2.3.2 Atmospheric dynamics

The pure radiative–convective solution described above can be substantially modified by atmospheric dynamics. In spin-synchronized planets, the symmetry of the system is most notably broken by two factors: day–night thermal contrast and rotation ([Showman et al. 2020](#)). The thermal forcing sets up standing planet-scale waves. Planetary rotation introduces vorticity on the largest scales, organizing the flow into complex large-scale features (Fig. 1.4).

Why do such large-scale motions arise? Although explaining such motions in detail requires




Figure 1.4: The 3D abundance of H_2O (left) and thermal structure of WASP-76 b (right) predicted by the R-M GCM (Beltz et al. 2021) and FastChem (Stock et al. 2018). 0° is the substellar longitude; the center of each panel is the antistellar point (on the planetary nightside). Overlaid in gray in the left column are the longitudes and pressures probed by transmission spectroscopy (Section 1.3; Wardenier et al. 2022). These regions bracket sharp compositional gradients and are notably distinct from one another. This rich structure cannot be predicted by 1D climate models (Fig. 1.3) and motivates the search for 3D effects in hot Jupiters.

solving complex nonlinear equations, we can motivate why this problem is interesting with a simple comparison of relevant scales.

First, we have the radiative timescale (Showman & Guillot 2002):

$$\tau_{rad} \approx \frac{P c_p}{4g\sigma_B T^3}, \quad (1.5)$$

for the Stefan-Boltzmann constant σ_B . This quantity describes how quickly a parcel of gas in the atmosphere will radiate away heat in excess of radiative equilibrium. At 1 bar on 51 Pegasi b, such a process would take more than 2 days; at 1 μ bar, it would take less than a second.

Next, consider the amount of time that winds moving at some characteristic velocity u could redistribute heat around a planet of radius R_p :

$$\tau_{advect} \approx \frac{R_p}{u}. \quad (1.6)$$

Comparing these timescales illustrates the regimes in which heat is effectively or ineffectively transported throughout the atmosphere (Showman & Guillot 2002).⁷ For 51 Pegasi b, if we assume that wind moves approximately at the sound speed, we arrive at roughly 3 km s⁻¹. So at 1 bar, $\tau_{advect} \ll \tau_{rad}$, whereas the converse is true at 1 μ bar.

Deepest down, as mentioned previously, the entire atmosphere is convective. At pressures deeper than the radiative-convective boundary (RCB), the planet can efficiently mix entropy across large spatial scales. Therefore, all latitude–longitude points are expected to conform to roughly the same adiabat in a time-averaged sense.

⁷While there are issues with this timescale comparison (Perez-Becker & Showman 2013; Komacek & Showman 2016), it is suitable for building intuition for this problem.

Now let us briefly examine some of the relevant circulation regimes. The Rossby number indicates the ratio of Coriolis to inertial forces in an atmosphere:

$$Ro = \frac{u}{Lf}, \quad (1.7)$$

for some length scale L and Coriolis parameter $f = 2\Omega \sin(\phi)$, for rotation rate Ω and latitude ϕ . At the mid-latitudes of 51 Pegasi b, $Ro \approx 2$; it is not a planet whose circulation is dominated by rotation. Other, faster-rotating planets exhibit values of Ro that approach unity. The nonlinear advective accelerations and Coriolis accelerations compete in this regime.

As a final quantity, we consider the Rhines scale. This value describes the scale over which a coherent, horizontal jet (like Jupiter’s bands) can be stable (Rhines 1975; Cho & Polvani 1996; Lian & Showman 2008):

$$L_r = \pi \sqrt{\frac{2u}{\beta}}, \quad (1.8)$$

where β is the meridional gradient in the Coriolis parameter. For 51 Pegasi b, $L_r > R_p$, which implies that only one jet dominates the flow.

Taken together, these scales imply that there is a large thermal contrast in hot Jupiter atmospheres that can be (partially) regulated by atmospheric flow. Such flow may be controlled by planetary-scale waves, and there may exist planetary-scale jets. Winds effectively homogenize the deep atmosphere (near the infrared photosphere), whereas the atmosphere is closer to radiative equilibrium at lower pressures (near 1 μ bar).

What is the nature of these waves and jets? It turns out that the two phenomena are linked. The spin-synchronized state of a hot Jupiter sets up two dominant, standing waves. The first is

Kelvin waves, which straddle the equator, propagate eastward, and are restored by the gravitational force. The second major planetary wave is Rossby waves. These waves move as gyres in the mid-latitudes and propagate westward, and they are restored by gradients in the Coriolis force.

While motivating the extent of jets is not too difficult, the mechanism that regulates such flows in hot Jupiters is deceptively complex. In brief: because the Rossby waves propagate west and the Kelvin waves propagate east, the Rossby gyres develop tilts in the mid-latitudes relative to the radiative equilibrium field.⁸ This tilted wave pattern preferentially pumps eastward momentum to the equator, accelerating the mean flow there (Fig. 1.4). This jet can be disrupted when the radiative timescale is very short (and waves cannot efficiently deliver energy to the equator) or if fluid drag is very strong. In the absence of a jet, flow can be predominantly divergent, following the pressure gradient and moving from the substellar point on the planet to the antistellar point (day–night) (Showman & Polvani 2011).

In practice, we do not predict exoplanet atmospheric dynamics with these timescale arguments and length scales; even first-order jet structure is a complex mean flow–eddy flow interaction. Analytical methods can capture day–night temperature differences (Komacek & Showman 2016), but the full spatial distribution of the velocity and thermal fields requires more. We often use general circulation models (GCMs; Fig. 1.4), most of which obey the “primitive equations” of meteorology. While the full Navier-Stokes equations describe fluid motions more generally, the primitive equations apply a number of physically motivated approximations to the Navier-Stokes equations that are relevant for planetary atmospheres. The primitive equations have been validated in the Solar System context, so continuing in our Occam’s Razor argument, it is sensible to begin

⁸A second reason is that the three-way force balance between the Coriolis, drag, and pressure-gradient forces is only satisfied if the velocity tilts northwest–southeast in the northern hemisphere and southwest–northeast in the southern.

from known frameworks that are at least validated in certain contexts.

These main features of exoplanet atmospheric circulation that I have described — day–night winds and jets — are expected to vary across the exoplanetary population ([Komacek et al. 2017](#)). Pure hydrodynamics arguments find that wind speeds should increase as planets get hotter, but we have yet to find decisive evidence that this is the case. Therefore, unsurprisingly, there remains substantial work to understand the underlying physics. In general: what physics and/or chemistry is missing from our models that is required to explain exoplanetary climate?

1.3 Inferring the properties of exoplanet atmospheres

The above theoretical arguments may paint an overconfident picture of our understanding of exoplanet atmospheres. While much of the early theory of hot Jupiter atmospheres has proven prescient ([Burrows et al. 1997](#); [Seager & Sasselov 1998, 2000](#); [Brown 2001](#); [Showman & Guillot 2002](#)), I will emphasize that the study of exoplanets has been driven by observational surprise — going back to the discovery of 51 Pegasi b itself. Theory can help guide observations, interpret what we see, and provide hints as to the latent structure of the Universe, but we rely on astrophysical observation to show us which possible theories are reflected in *our* Universe. This is especially true when competing theories, built on differing assumptions, make testable predictions.⁹


⁹Assumptions must be made in astrophysics. The reader may wonder: why not simply simulate these bodies in full, given that we have excellently verified the laws of physics in terrestrial laboratories? Perhaps most pressing is the computational intractability of fully modeling even an infinitesimally small patch of the Universe. Cutting-edge simulations model up to 500,000 atoms at once ([Hospital et al. 2015](#)); simulating a cubic meter of exoplanet atmosphere requires at least $10^{20} \times$ as many atoms. Self-consistently modeling chemistry from first-principles physics would be even worse — it would require simultaneously solving the Schrödinger equation for every atomic nucleus *and all of the electrons*. This approach may be excessive, because we could simply model the motions in bulk as a fluid using the long-held techniques of hydrodynamics and statistical mechanics. But such a model has assumptions with theoretical and observational ramifications — and it should therefore be tested!

1.3.1 Staring at exoplanetary systems


Observing exoplanet atmospheres, it turns out, is nontrivial. Compared to other commonly studied, self-luminous astrophysical objects (e.g., stars and galaxies), most exoplanets are intrinsically faint, intrinsically small, and physically close to their much-brighter host stars. Generally speaking, we cannot distinguish a planet from its host star by its position on the sky. The classic comparison is to a lightning bug fluttering next to a lighthouse — the lighthouse’s light overwhelms any signal you might try to measure from the little bug.¹⁰ We therefore rely on a few tricks to make this problem easier: differential measurement and time-resolved measurement. For instance, if an exoplanet, host star, and observer are aligned just so, the planet will pass between the star and observer at some point along the orbit. During this primary occultation, or “transit,” the host star’s light is dimmed from the observer’s perspective, as a fraction of starlight is absorbed and scattered by the planet (Fig. 1.5a). This technique enables the indirect measurement of an exoplanet’s size.

We can now appreciate why the early years of studying exoplanet atmospheres were so difficult. In sum, to characterize an exoplanet atmosphere, we must measure very small differential signals, and we must maintain the stability of that measurement over time. In some cases, the theory of atmospheres was sufficiently unconstrained to allow a wide range of possible physical scenarios. Worse yet, the cutting-edge instruments used in the early aughts were developed far before anyone had detected a single exoplanet. That is, none were designed with exoplanetary science cases in mind. The Hubble Space Telescope (HST), for instance, can often only capture


¹⁰In special cases, with excellent instrumentation, we are *able* to spatially disentangle a planet from its host star. Doing so is easiest if the planet is large, hot, and far from its star. These systems turn out to be quite rare (Bowler 2016). While “direct imaging” is not currently the workhorse method of exo-atmospheric characterization, modern telescopes are pushing the direct imaging technique to smaller, cooler, and older planets (e.g., Matthews et al. 2024; Crotts et al. 2025; Lagrange et al. 2025).




(a) A schematic of the physical system. An observer measures an exoplanet's absorption and emission based on its visibility.



(b) The (noiseless) emergent exoplanetary spectra Doppler shift as the exoplanet moves through its orbit (which is shown in (a)).



(c) The result of cross-correlating a template spectrum against the spectra in (b), combining all spectral lines into a net signal strength as a function of phase (a “trace plot”), tracking the exoplanet's orbit.



(d) A K_p - V_{sys} map, summing cross-correlation values along an orbit for each tested K_p and V_{sys} . The map is maximized (brightest) at the K_p and V_{sys} that match the exoplanet's orbit (the sinusoid in (c)).

Figure 1.5: An outline of high-resolution cross-correlation spectroscopy.

portions of a transit at a time; its low-Earth orbit ensures that there are gaps in its data when it passes behind the Earth and the exoplanet is out of view. It took substantial investment to mitigate the systematics that manifest only when performing observations of exoplanet transits (e.g., [Zhou et al. 2017](#)).

1.3.2 Spectroscopy

Observing an exoplanet’s transit and/or its thermal emission at one wavelength generally does not provide much leverage on its atmosphere. But performing these measurements at *multiple* wavelengths — *spectroscopy* — can yield information-rich data products. Such information is due to a series of quantum mechanical effects: photons from the host star and / or the planetary atmosphere interact with gaseous, solid, or liquid particles in the planet’s atmosphere before traversing vast swathes of space and arriving at our telescopes. These interactions are highly wavelength-dependent.

Consider an atmosphere composed of pure neutral atomic hydrogen. Per the Bohr model, the single electron orbiting the single proton can only exist in a certain set of orbits. At higher-level orbits, the electron (and therefore the system) has more energy; at lower-level orbits, less. Due to this relationship, the energy levels of neutral atomic hydrogen are quantized ([Bohr 1923](#)).

Incident radiation can interact with this system. If a photon with energy matching the exact energy difference between two levels in the atomic system interacts with the atom, the photon is absorbed and the electron changes orbits. The same argument extends to molecules. Because the multiple atoms in the molecular system break spherical symmetry, molecules have rotational and vibrational energy states in addition to the electronic energy states. If incident photons match

the gaps between these energy levels, the photons are absorbed and the molecule is excited into different rotational, vibrational, or electronic states.

While the energy levels themselves are quantized, spectral lines have finite width. Various “line-broadening mechanisms” shift the photons that individual atoms or molecules can absorb, either by motions of the species (Doppler-shifting the radiation field) or by shifting the energy levels themselves. Thus, there is a range of photons that can be emitted or absorbed by these systems even for a single spectral line.

In transit, these quantum mechanical effects have a clear implication: host star light percolating through the planet’s atmosphere will be preferentially blocked at wavelengths (energies) matching the energy levels of specific atmospheric gases. Therein lie the “fingerprints” of atomic and molecular absorbers. The transit spectrum reveals the radii at which the planet becomes opaque, or “optically thick,” at multiple wavelengths. In other words, at each measured wavelength, the depth of the transit “dip” depends on how much starlight the planet blocks, and how much light is blocked depends on the planet’s gaseous and aerosol composition. Strong spectral lines indicate wavelengths with high spectral opacity; only a small amount of gas at the tenuous top of the atmosphere is required for the planet to become optically thick. Conversely, weak spectral lines indicate that the opacity is weak; they require the high density at smaller planetary radii to become optically thick. This “transmission spectroscopy” probes exoplanetary limbs (Fig. 1.4– 1.5a), typically at pressures ranging from 0.1 mbar to 1 μ bar.

Similar logic holds for planetary emission. An excited atom or molecule relaxes by emitting a photon at characteristic wavelengths, with the intensity set by the local temperature. The emission spectrum shows the temperature of the blackbody at which the atmosphere becomes optically thick at that wavelength. Strong absorption lines occur in spectral regions with high opacity and

indicate a decreasing thermal profile along the radial coordinate; the emerging radiation becomes optically thick in a cold region of the atmosphere. Conversely, strong emission lines indicate an increasing thermal profile (a stratosphere); the atmosphere is hotter as altitude increases. Emission spectroscopy probes deeper in the atmosphere than transmission spectroscopy, from roughly 1 bar to 1 mbar.¹¹

Gaseous detections from the ground and space, particularly in the past few years, have provided a set of expectations for typical absorbers. Many of these results are interpretable with equilibrium chemistry: in the infrared, H₂O and CO tend to dominate (e.g., [Snellen et al. 2010](#); [de Kok et al. 2013](#); [Line et al. 2021](#)), with CO₂ becoming observable at higher metallicity (e.g., [Rustamkulov et al. 2022](#); [Xue et al. 2024](#)) and OH at higher temperature (e.g., [Brogi et al. 2023](#)). However, disequilibrium chemistry must be invoked to explain the presence of SO₂, which appears when photochemistry is active and at high enough metallicity ([Tsai et al. 2023](#)), and the behavior of CH₄, which is present when the planet’s internal heat and vertical mixing are modest (e.g., [Bell et al. 2023](#); [Madhusudhan et al. 2023](#); [Welbanks et al. 2024](#)). The optical is rife with atoms and ions, with the broad lines of Na and K (e.g., [Seidel et al. 2021a](#); [Tabernero et al. 2021](#)) interspersed by forests of lines contributed by Fe I, Fe II, Mg I, Mn I, and more (e.g., [Kesseli et al. 2022](#); [Borsato et al. 2023](#)). These atomic lines are sometimes mixed with the molecular features of metal oxides, namely TiO and VO. The potential presence of these species has been of interest to the community for decades, as their strong optical opacity could be responsible for generating thermal inversions (e.g., [Fortney et al. 2008](#)). Due in part to the difficulty of precisely knowing the opacity

¹¹There is a third type of spectroscopy that I will not consider: reflected light spectroscopy. Part of the reason for this omission is that hot Jupiters incidentally have relatively low albedos, so not much starlight is reflected back to the observer; only a select few datasets exhibit evidence for such an effect ([Coulombe et al. 2025](#); [Pelletier et al. 2026](#)). Reflected light spectroscopy did play a large role in exoplanetology’s past (with many early attempts at exoplanet spectroscopy searching for it; e.g., [Charbonneau et al. 1999](#); [Collier Cameron et al. 1999](#)), and it will play a large role in exoplanetology’s future (characterizing Earthlike worlds; e.g., [Feinberg et al. 2024](#); [Tuchow et al. 2024](#)).

function of these species (e.g., [Hoeijmakers et al. 2015](#); [De Regt et al. 2022](#)), these species have only recently been robustly detected in exoplanet atmospheres ([Prinoth et al. 2022](#); [Pelletier et al. 2023](#)).

These strong spectral lines peek above continuum sources of opacity, such as clouds (which absorb largely independently of wavelength, aside from broad slopes and Christiansen features; [Christiansen 1884](#); [Wakeford & Sing 2015](#); [Grant et al. 2023a](#)). Additionally, interactions *between* species provide smooth sources of opacity, as they can occur with a wide range of possible energies. Notable collisional pairs include $\text{H}_2\text{--H}_2$ and $\text{H}_2\text{--He}$ for hydrogen–helium-dominated atmospheres. For hot atmospheres, the H^- interactions contribute, as well ([Lothringer et al. 2018](#)). Although these continuum sources seem like a nuisance (with their flat opacity washing away the wavelength-dependent spectral signature from gaseous absorbers), they actually provide a needed baseline for measuring such wavelength independence ([Welbanks & Madhusudhan 2019](#)) and are scientifically interesting in their own right ([Lothringer & Barman 2020](#)).

The planets that are most amenable to spectroscopy are hot Jupiters. We can motivate this favorability as follows. On a practical level, hot Jupiters are on short orbital periods, making scheduling their observations far easier than for slower-moving planets. For a more physics-based justification, we can derive a metric for the signal-to-noise ratio of a transit spectrum observation ([Kempton et al. 2018](#)):

$$\text{SNR}_{\text{transit}} \propto \frac{T_{\text{eq}} R_p^3}{M_P}, \quad (1.9)$$

where T_{eq} is the equilibrium temperature of the planet and M_P is its mass. Hot, inflated planets therefore maximize this quantity. We can derive a similar metric for emission spectroscopy:

$$SNR_{\text{emission}} \propto \frac{2hc^2}{\lambda^5} \frac{1}{e^{\frac{hc}{\lambda k_B T_{eq}}} - 1} R_p^2, \quad (1.10)$$

for Planck’s constant h , speed of light c , and observational wavelength λ . This function (the Planck spectrum scaled by the area of the planet) monotonically increases with temperature. Once again, large, hot planets are favored.

Altogether, remote-sensing spectroscopy provides leverage on the composition and thermal structure of exoplanet atmospheres. Unfortunately, the connections between spectra and atmospheric properties do not follow simple, closed-form expressions; realistic models of exoplanet spectra require numerical computation.

1.3.3 Statistics

At this point, we have physical models that we understand, and we have observational data that we understand. We would ideally like to compare our models with the data to demonstrate which families of models best describe what we observe. To do so, for better or for worse, we must invoke statistics.

The first metric that one might use is a simple χ^2 statistic:

$$\chi^2 = \sum_i \frac{(m_i - f_i)^2}{\sigma_i^2}, \quad (1.11)$$

for measured data f_i , model points m_i , and uncertainty on data σ_i at each wavelength i . This goodness-of-fit metric does indeed quantify how well a single model fits a dataset. But now we would like to test *other* models and see how well they fit the data. We could generate a grid of models — say, over a range of temperatures and a range of water abundances — and see which best




Figure 1.6: Three model transmission spectra calculated with `petitRADTRANS` (Mollière et al. 2019). Varying one parameter at a time (i.e., increasing the atmospheric temperature) leads to stark changes, but other parameters complicate the picture. Here, an increase in the metal content of the atmosphere could masquerade as a moderate increase in temperature. Degeneracies such as this motivate the use of Bayesian computation.

fits the data. We must select a planetary mass when we do so. We can then excitedly report to our collaborators that a planet has some amount of water and its atmosphere is at some temperature. But how sure are we about these “point estimates”? And how sensitive are our answers to the measurement uncertainty of the planet’s mass? As mentioned previously, exoplanet spectroscopy is a highly nonlinear and degenerate problem (Fig. 1.6). Can our dataset be nearly equivalently explained by an atmosphere that is twice as hot with half as much water?

The branch of statistics known as “Bayesian statistics” offers a natural framework for connecting our prior knowledge of the system to our models, our data, and our post-experiment uncertainty regarding the exoplanet’s properties. Bayesian computation relies on Bayes’ theorem (Bayes 1763):

$$p(M|D) = \frac{p(D|M)p(M)}{p(D)}. \quad (1.12)$$

In this expression, $p(M|D)$ is the probability of some combination of model inputs, M (such as a specific planet temperature, water abundance, and mass), given the data, D . This is our final desired quantity — the *posterior probability distribution* of model inputs, weighted by how likely they are given the dataset. Computing this quantity requires $p(D|M)$, the probability of the data given the model. This can be formally expressed as a likelihood function. A standard such function is proportional to the χ^2 function itself; χ^2 is an appropriate likelihood function when noise in the data is distributed according to a Gaussian distribution. $p(M)$ is the prior distribution — the probability of the set of model parameters before we see the data, perhaps based on prior measurements. This framework therefore allows us to weight model inputs by our *a priori* expectation that some combinations might not be plausible.

The “Bayesian evidence” in the denominator of Eq. 1.12 is often discarded, as it is identical for every tested parameter combination. Its full form integrates the likelihood function (describing how well any combination of inputs fits the data) over the full parameter space, weighted by the prior probability of each combination of parameters. In essence, it encapsulates how well this model *could possibly* explain the data, including every possible combination of inputs. This quantity is therefore useful for comparing the performance of competing model parameterizations.

We now have a statistical framework that suits us well. In principle, we can compute these quantities within a grid in parameter space. However, this approach is impractical. For instance, if our models reside in a 20-dimensional parameter space, and we would like to test 10 values along each axis, and each model takes 1 second to compute, this process would take more than

3 trillion years. This approach is clearly not ideal. So, we intelligently explore parameter space using statistical samplers. Common approaches include Markov Chain Monte Carlo ([Brooks 1998](#)) and nested sampling ([Skilling 2004](#)). Sampling algorithms render these Bayesian experiments tractable, with runtimes brought down to just weeks, or even days, especially after parallelizing computation over multiple computational units.

In exoplanet atmospheres, we term this Bayesian estimation scheme — inferring the properties of an exoplanet atmosphere via spectroscopy — a *retrieval*. This framework has a long lineage in Earth sciences and connections to information theory ([Shannon 1948](#); [Rodgers 1976](#); [Line et al. 2012](#)). So long as we trust the data, atmospheric forward model, and statistical sampler, retrievals are powerful tools that allow us to measure the physical properties of exoplanet atmospheres.

1.4 Spectroscopy, now at higher resolution

The methods presented thus far are only indirectly sensitive to the dynamics of atmospheres via chemistry and temperature. A more recent technique — ground-based, high-resolution cross-correlation spectroscopy (HRCCS; e.g., [Snellen 2025](#)) — allows direct measurement of exoplanetary winds and, in some cases, improved constraints on atmospheric chemistry.

The problem is the same as before: we seek to make a small, differential measurement to isolate the planetary signal from stellar and background sources. From the ground, though, it is possible to build large, vacuum-sealed spectrographs. This fact enables much higher spectral resolution, up to $R = \frac{\lambda}{\Delta\lambda} \sim 140000$ (e.g., [Pepe et al. 2010](#); [Yuk et al. 2010](#); [Follert et al. 2014](#); [Seifahrt et al. 2016](#)).

1.4.1 Isolating the planetary signal

We generally do not isolate the planetary signal in HRCCS relative to absolute flux levels. In fact, dividing the flux measured in-transit by the flux measured out-of-transit is rarely enough to make the planetary signal evident, due primarily to the difficulty of maintaining pointing stability and constant signal-to-noise with ground-based telescopes.¹² Rather, we separate the planetary signal from the stellar signal by virtue of the Doppler-shifting of the exoplanetary spectrum. When the exoplanet moves through its orbit toward and away from the observer, the absorption and emission lines blueshift and redshift accordingly.¹³ These motions greatly exceed the minimum velocities that ground-based, high-resolution spectrographs can resolve. In contrast, the stellar lines remain quasistationary. While the star certainly moves in its orbit (around the barycenter of the system), such motion is generally an order of magnitude slower than the exoplanetary signal.

To reach ground-based telescopes, photons from astrophysical sources must travel through our sky. This process preferentially absorbs photons at certain wavelengths, imprinting a *telluric spectrum* on ground-based data. Luckily, this spectrum can also be treated as quasistationary. While the Earth’s atmosphere does move (e.g., the turbulence experienced by airplanes), such motion is again slower than the multiple km s^{-1} of the exoplanetary signal.

In general, there are two approaches to removing such quasistationary signals: empirical methods (such as singular value decomposition, or SVD; [de Kok et al. 2013](#)) or physics-based methods (such as Molecfit; [Smette et al. 2015](#)). Both approaches aim to remove the dominant spectral components, ideally leaving only white (uncorrelated) noise and the exoplanetary signal.

¹²An exception is in the very high signal-to-noise regime, in which individual lines exceed the noise floor.

¹³This effect is related to how the siren on a fire truck sounds higher-pitched as it approaches and lower-pitched as it recedes; the movement of the source changes the distance between successive wave peaks.

In practice, it is difficult to assess when a majority of the telluric and stellar correlated residuals have been removed from a dataset. This fact, in combination with the recency of HRCCS retrieval methodology, the continuing emergence of best practices, and the computational intensity of such techniques, makes HRCCS challenging.

Even if the dominant sources of contamination have been removed, an individual planetary spectral line competes with the noise floor. To stack our signal, we resort to cross-correlation: detecting the exoplanet’s spectral signature in velocity space, essentially combining the contributions from all spectral features at once. Assuming that a spectrum that contains N_{lines} spectral features with approximately equal strength, the signal-to-noise of this measurement procedure scales as (e.g., [Birkby 2018](#)):

$$SNR_{CCF} \propto \sqrt{N_{lines}}. \quad (1.13)$$

Yet even this level of stacking signal is often not enough to extract the exoplanetary signal (except for with the highest-quality datasets). We must stack again, by combining the signal from multiple consecutive exposures. How would we do so? The planet is moving. Therefore, we could assume an orbit following Kepler’s laws and shift all exposures to that orbit’s reference frame. However, this process involves interpolating noisy data, which would incur undesirable error. Instead, we can assume a single Keplerian orbit, Doppler shift our *model* to the appropriate velocity at each exposure, cross-correlate, and sum across exposures.

The current best practice is quite close to this latter approach. The only modification is that instead of testing a single Keplerian orbit, a grid is tested in $K_p - V_{sys}$ space, where the planet’s line-of-sight radial velocity as a function of time is defined as:

$$RV(t) = K_p \sin(2\pi\phi(t)) + V_{sys}, \quad (1.14)$$

where ϕ is the planet’s orbital phase (ranged 0 to 1; Fig. 1.5c). In this expression, K_p is the planet’s radial velocity semiamplitude, and any velocity offset is captured with V_{sys} . This approach assumes a circular orbit for the planet. Such an assumption is well-motivated by the notion that at such close proximity to their host stars, hot Jupiters’ orbits rapidly circularize and their spins are rapidly driven to synchronous rotation by tidal effects (e.g., [Rasio & Ford 1996](#)). The $K_p - V_{sys}$ map is maximized near the exoplanet’s true orbit (Fig. 1.5d), which in most cases has already been determined by independent detection methods (such as radial velocity measurements of the *host star* as it orbits the system’s barycenter; [Lovis et al. 2010](#)).

Testing multiple orbits provides certainty in the detection of the exoplanet’s atmosphere; we often quantify HRCCS detection significance by comparing the distribution of CCF values in the $K_p - V_{sys}$ peak compared to out of the peak (e.g., [Kasper et al. 2021](#)). From a data perspective, testing multiple orbits also encompasses any uncertainty in the system parameters. For instance, K_p can be derived as:

$$K_p = \frac{2\pi a}{P} \sin(i), \quad (1.15)$$

where a is the planet’s semi-major axis, i is its inclination with respect to the observer ($\pi/2$ radians if the planet passes directly between the star and the observer), and P is the planet’s orbital period. Uncertainties in any of these quantities would percolate through to uncertainties in K_p . Furthermore, the system may physically be moving away from the Earth in its path through the Milky Way. There may also exist ephemeris uncertainties, which would manifest as velocity offsets. Fi-

nally, perhaps there are departures from a circular orbit ; an eccentricity of just 0.01 could induce a radial velocity shift of a few km/s ([Montalto et al. 2011](#)). Other sources of noise in velocity space include correlated stellar or telluric residuals, aliases between the template and other gases in the exoplanet atmosphere, and self-aliases.

Nowadays, we routinely characterize exoplanet atmospheres using this method, with dozens of detections spanning a wide range of planets ([Kempton & Knutson 2024](#); [Snellen 2025](#)).

1.4.2 Sensitivity to exo-atmospheric winds

After subtracting off the velocity corresponding to the exoplanetary orbit, it is common to find a residual velocity shift in the data. In high-quality data, this shift in velocity can be attributed to net Doppler shift in exoplanetary atmospheres. That is, motions in the planetary rest frame (due to planetary rotation and winds) manifest as shifting and broadening of spectral lines, which cause shifting and broadening of the cross-correlation function. **High-resolution spectroscopy, therefore, provides the only way to directly measure the winds in an exoplanet atmosphere.**

1.5 Popping the spherical cow: Limits of the 1D retrieval paradigm

High-resolution spectroscopy offers a concrete path to measuring quantities that are highly relevant to understanding exoplanets, and ultimately understanding the Solar System in the broader exoplanetary context. Realizing that potential requires a key component: comparing *complete* models to the data. If our model are inaccurate, then so are the inferences — often in subtle and difficult to diagnose ways.

Theory indeed suggests that hot Jupiters are far from homogeneous. This claim goes back

to our spin-synchronized picture — GCMs predict day–night gradients that can exceed thousands of Kelvin. These thermal gradients, coupled with rotation, can in turn set up wind gradients and chemical gradients. Our isothermal blob obviously no longer holds. The proverbial “spherical cow” of exoplanet spectroscopy — the seemingly necessary assumption of homogeneity — must be released.

Why not compare *these* models to data? We do, sometimes (including as in my own work). But they cannot be used for parameter estimation, which requires millions of model evaluations. This would cost roughly 1 billion CPU hours, with a wall clock time of roughly 100,000 years!

The limitations of 1D atmospheric modeling have been discussed at length for space-based spectroscopy ([Line & Parmentier 2016](#); [Rocchetto et al. 2016](#); [Blecic et al. 2017](#); [MacDonald & Madhusudhan 2017](#); [Caldas et al. 2019](#); [Lacy & Burrows 2020](#); [MacDonald et al. 2020](#); [Pluriel et al. 2020](#); [Taylor et al. 2020](#); [Espinoza & Jones 2021](#); [Pluriel et al. 2021](#); [Chubb & Min 2022](#); [MacDonald & Lewis 2022](#); [Nixon & Madhusudhan 2022](#); [Zingales et al. 2022](#); [Falco et al. 2024](#); [Schleich et al. 2024](#)). It stands to reason that the biases incurred by inaccurate modeling are even more conspicuous in HRCCS, because it provides access to the exoplanetary wind field. This additional effect can produce spectral line shapes that 1D models likely cannot match with simple shifts and broadening. Indeed, significant departures from the 1D picture — containing information about wind fields, thermal profiles, and planetary rotation — begin to crop up at higher spectral resolution ([Kempton & Rauscher 2012](#); [Showman et al. 2013a](#); [Kempton et al. 2014](#); [Flowers et al. 2019](#); [Beltz et al. 2020](#); [Wardenier et al. 2021](#)).

My work enters here. There are a number of claims that are worth investigating before diving into 3D retrievals. After all, we generally seek to obey Occam’s Razor: explain complex observed phenomena with simple physics first, then escalate the physics as needed. These claims are:

1. 3D effects impact data.
2. 3D effects outcompete typical noise sources.
3. 3D effects bias inferences made with 1D frameworks.
4. 3D inference is computationally plausible.

I will, over the course of this dissertation, present evidence for these statements.

Chapter 2: No umbrella needed: Confronting the hypothesis of iron rain on WASP-76 b with post-processed general circulation models

2.1 Introduction

Decades after the discovery of the first exoplanet around a normal star ([Mayor & Queloz 1995](#)), a harbinger of the population of hot Jupiters, observational, theoretical, and laboratory efforts have elucidated some of the key processes and features relevant to these planets (e.g., [Showman & Guillot 2002](#); [Cooper & Showman 2006](#); [Knutson et al. 2007b](#); [Fortney et al. 2008](#); [Crossfield et al. 2010](#); [Rauscher & Menou 2012](#); [Demory et al. 2013](#); [Dobbs-Dixon & Agol 2013](#); [Kataria et al. 2015](#); [Mendonça et al. 2016](#); [Fleury et al. 2019](#); [Winter et al. 2020](#)). In the past eight years, however, an even more extreme class of planet — the “ultra-hot Jupiter” ([Arcangeli et al. 2018](#); [Bell & Cowan 2018](#); [Parmentier et al. 2018](#)) — has emerged, once again requiring the construction of new theoretical frameworks. Featuring persistent dayside temperatures in excess of 2200 K, ultra-hot Jupiter atmospheres are thought to include thermal dissociation of molecules and thermal ionization of metals. The presence of gas-phase metals at low pressures within these atmospheres also renders ultra-hot Jupiters notably susceptible to temperature inversions by virtue of strong optical absorption of host star light (e.g., [Lothringer & Barman 2019](#)). In short, observations of ultra-hot Jupiters have revealed them to be sites for extreme physics, consequently requiring significant model updates to understand them in detail.

One particularly intriguing aspect of ultra-hot Jupiter atmospheres is their 3D structure. Such highly irradiated objects are predicted to have very strong day-night temperature gradients, and perhaps chemical gradients, too (e.g., [Showman et al. 2009](#); [Rauscher & Menou 2012](#); [Kataria et al. 2016](#)); these atmospheres’ short chemical timescales relative to dynamical timescales should

drive them to chemical equilibrium (Baeyens et al. 2021). These features make it challenging to justify treating such planets as 1D objects, and the departures from 1D behavior are especially apparent in certain current and future (Feng et al. 2020; MacDonald et al. 2020; Pluriel et al. 2020; Taylor et al. 2020) observations, including those at high spectral resolution, which are the subject of this chapter.

To account for the unique 3D properties of ultra-hot Jupiters, various updates to “traditional” hot Jupiter general circulation models (GCMs) have been required. For instance, recent models have incorporated the cooling and heating effects of hydrogen dissociation and recombination (Tan & Komacek 2019; Mansfield et al. 2020). Moreover, the Lorentz forces and Ohmic dissipation related to free electrons in weakly ionized atmospheres have been captured to varying degrees of sophistication in GCM frameworks, often as a Rayleigh drag force (e.g., Perna et al. 2010; Rauscher & Menou 2013; Komacek & Showman 2016) or by coupling to the magnetic induction equation (Batygin et al. 2013; Rogers & Showman 2014; Hindle et al. 2019; Hindle et al. 2021). Finally, the net thermal flux at the lower boundary can substantially impact the GCM-predicted thermal structure and circulation higher in the atmosphere (e.g., Rauscher & Menou 2012; Showman et al. 2015), making this parameter a useful tuning knob to explore the effect of differing interior entropies on global dynamics.

A prime example of the unique benefits of the intersection of the 3D nature of (ultra-)hot exoplanet atmospheres and high-resolution spectroscopy lies in the ultra-hot Jupiter WASP-76 b, a gas giant orbiting an F7 star (West et al. 2016) that is well-studied at lower resolution (Fu et al. 2017; Fisher & Heng 2018; Tsiaras et al. 2018; Edwards et al. 2020; von Essen et al. 2020; Fu et al. 2021). Using the high-resolution ($R \approx 138,000$) ESPRESSO spectrograph on the VLT (Pepe et al. 2010, 2013), the Ehrenreich et al. (2020) team was able to produce novel, high-SNR, phase-

resolved transmission spectra of this target across two separate transits. Curiously, an anomalous Doppler signature in the planet’s transmission spectrum was detected: between 0 and -5 km s^{-1} at ingress, but roughly -11 km s^{-1} by egress. These speeds far exceed the few km s^{-1} planet-frame velocities detected on other planets (Snellen et al. 2010; Brogi et al. 2016). The detection team attributes this variable and strong blueshift to an asymmetric distribution of atomic iron in the planet’s atmosphere, with a considerable amount of iron existing in the gas phase east of the substellar point but cooling and condensing out as it makes its way to the much colder nightside. This asymmetry would cause a progressively blueshifted signal as the gas-phase iron region rotates into view over the course of the planet’s transit. Ehrenreich et al. (2020) posit that their signal is composed of two independent Doppler components: solid-body rotation of $\pm 5.3 \text{ km s}^{-1}$, the tidally locked equatorial velocity of the planet, and a uniform day-night wind contributing an additional -5.3 km s^{-1} across both limbs, with gas phase iron only present on the evening limb of the planet (Figure 2.1). Hence, the approaching limb would exhibit a blueshift from both rotation and winds totaling -10.6 km s^{-1} , and the receding limb would produce no Doppler signal, as it would contain no gas-phase iron to absorb starlight.

The Ehrenreich et al. (2020) “toy model,” as they put it, is supported by a separate, archival analysis of HARPS data (Kesseli & Snellen 2021), which similarly finds increasing blueshift over the course of WASP-76 b’s transit. Additionally, the magnitude of the blueshift and speed of the day-night flow are confirmed by Tabernero et al. (2021) and Seidel et al. (2021b), who perform in-depth secondary analyses of the Ehrenreich et al. (2020) dataset. These works also produce rich datasets in their own right, with the Tabernero et al. (2021) study reporting the detection of multiple chemical species (Li I, Na I, Mg I, Ca II, Mn I, K I), their atmospheric heights, and their corresponding blueshifts.

The aforementioned iron chemistry gradient interpretation would be the first of its kind among exoplanet atmospheres. However, the necessity of a chemical gradient to explain the extant observations has been questioned by the forward models of [Wardenier et al. \(2021\)](#), who self-consistently calculate transmission spectra from a 3D model of this planet and find that either iron condensation or a significant temperature asymmetry could match the [Ehrenreich et al. \(2020\)](#) data. Furthermore, Fe clouds are not necessarily favored in hot Jupiter atmospheres; microphysics models indicate that the nucleation rate of Fe is low, causing Fe clouds to be sequestered deep in the atmosphere ([Gao et al. 2020](#); [Gao & Powell 2021](#)).

In this chapter, in addition to seeking to examine the iron rain hypothesis, we also aim to understand the physical, chemical, and dynamic processes at play in the atmosphere of WASP-76 b. We do so with a suite of GCM simulations, applying a spectroscopic post-processing scheme to evaluate the likelihood of several plausible physical scenarios in the face of the constraints provided by the [Ehrenreich et al. \(2020\)](#), [Tabernero et al. \(2021\)](#), and [Kesseli & Snellen \(2021\)](#) data. A companion work, [May & Komacek et al. \(2021\)](#), examines the dynamics of this planet with the aid of Spitzer phase curves; this chapter focuses on the detected blueshift signature and chemical species by post-processing GCM simulations to produce high-resolution transmission spectra. Together, these works aim to compare the results of photometric phase curves and high-resolution transmission spectroscopy to paint a coherent picture of WASP-76 b's 3D atmosphere.




Figure 2.1: A schematic diagram of the “toy model” introduced in [Ehrenreich et al. \(2020\)](#). WASP-76 b is pictured at three different phases (left to right: ingress, center of transit, and egress). Colder regions of WASP-76 b’s atmosphere, where Fe condenses and produces no absorption, are the primary regions probed at ingress. Over the course of its transit, the planet rotates and brings hotter regions, which contain gas-phase iron, into view for the observer. This geometry makes the strong combined velocity of the trailing limb ($\approx -5.3 \text{ km s}^{-1}$ from winds, $\approx -5.3 \text{ km s}^{-1}$ from rotation) observable as a Doppler blueshift by mid-transit. Atmosphere size not to scale.

2.2 Model Description

2.2.1 GCM

As described above, we aim to produce physically motivated, self-consistent comparisons to the [Ehrenreich et al. \(2020\)](#) data by spectrally post-processing the outputs of a GCM. We use the same GCM model introduced and described in detail in our companion work, which also includes the full GCM parameter assumptions and values ([May & Komacek et al. 2021](#), Table 3). Broadly, we use the MITgcm ([Adcroft et al. 2004](#)) to solve the primitive equations of meteorology on a cubed-sphere grid with a horizontal resolution of C48 and 70 vertical layers, extending evenly in log-pressure from 100 bars to 10 microbars. To connect the modeled atmospheric circulation to




Figure 2.2: 1D temperature-pressure (T-P) profiles of WASP-76 b as output by our 3D GCMs, with the condensation curves of Mg_2SiO_4 , Fe, and Al_2O_3 from [Mbarek & Kempton \(2016\)](#) and terminators overplotted. Though the GCMs’ outputs extend to 100 bars, our radiative transfer is limited to the (plotted) region at pressures less than 1 bar. Every 20th profile is plotted for ease of visualization, resulting in 902 samples of 18048 total latitude-longitude points. Regions near the substellar point (0° / 360°) tend to be hotter on average than the rest of the planet, and regions near the antistellar point (180°) tend to be cooler on average than the rest of the planet. Significant scatter at a given longitude is controlled by latitudinal variation. The GCMs with deeper RCBs tend to have cooler nightsides than those with shallower RCBs, though all sets have regions in which Fe, Mg_2SiO_4 , and Al_2O_3 condense (assuming equilibrium chemistry). Models with weaker drag (i.e., longer drag timescales) tend to produce much greater asymmetry between the east and west planetary limbs.




Figure 2.3: Maps of atmospheric temperature at the terminator of WASP-76 b for weak (1×10^7 s) and strong (1×10^3 s) drag timescales and three phases (ingress, center of transit, and egress). The planet core is not shown to scale, and we restrict our visualization to the region that is probed by high-resolution transmission spectroscopy ($\lesssim 1$ bar — our model domain slightly exceeds 1 bar in some regions, as our post-processed atmosphere is interpolated onto a grid evenly spaced in altitude). Black dashed lines correspond to the millibar isobar. As also shown in Figure 2.5, the strong drag case enforces more overall east-west symmetry in the atmosphere of WASP-76 b. The temperature inversion that can be seen in Figure 2.2 is also present here — as would be expected for an ultra-hot Jupiter with high-altitude optical absorbers.

heating and cooling within the atmosphere, the GCM is coupled with a two-stream, double-gray radiative transfer scheme with the TWOSTR (Kylling et al. 1995) package of DISORT (Stamnes et al. 1988). A crucial update warranted for this project is the inclusion of local cooling where hydrogen dissociates, relevant to the intense dayside stellar irradiation of ultra-hot Jupiters, and local heating at locations where atomic hydrogen recombines, relevant to the much cooler limbs and nightsides of ultra-hot Jupiters (Tan & Komacek 2019; Mansfield et al. 2020).

As described in our companion work (May & Komacek et al. 2021), our GCMs (Figure 2.2) incorporate a spatially constant Rayleigh drag force proportional to $\mathbf{v}_{\text{wind}}/\tau_{\text{drag}}$, for horizontal wind velocity $\mathbf{v}_{\text{wind}}$ and drag timescale τ_{drag} . Our grid of GCMs encompasses a parameter sweep of drag timescales and two endmembers of radiative-convective boundary (RCB) depth, with the aim of exploring a wide range of possible dynamical states and interior states of WASP-76 b. Our models are computed over a parameter sweep of drag timescales ($\tau_{\text{drag}} \in [1 \times 10^3 \text{s}, 1 \times 10^7 \text{s}]$, with five values sampled evenly in log space), motivated by the unknown strength of turbulence and magnetic forces associated with the ionized dayside. The longest drag timescales correspond to atmospheres that have similar temperature structures and wind speeds to those with no applied frictional drag. We model upward heat fluxes from the deep planetary interior by including a “surface” beneath our model domain that does not interact with the atmospheric flow aside from heating it. The upward fluxes for our limiting cases of RCB depth are 3543.75 W m^{-2} ($T_{\text{int}} = 500 \text{ K}$) and $4.474 \times 10^7 \text{ W m}^{-2}$ ($T_{\text{int}} = 5300 \text{ K}$) for the deep and shallow RCB cases, respectively.

2.2.2 Radiative transfer model

2.2.2.1 Overview

For our study, we make use of a line-of-sight ray-striking radiative transfer model that includes Doppler effects, inherited from [Kempton & Rauscher \(2012\)](#), [Kempton et al. \(2014\)](#), [Rauscher & Kempton \(2014\)](#), and [Flowers et al. \(2019\)](#). We begin with the GCM output, which we interpolate from a grid evenly spaced in log pressure onto a grid evenly spaced in altitude — thus ensuring that our simulated rays propagate in straight lines. We further only consider the GCM output at pressures generally greater than 1 bar (the cut is made in altitude, not pressure, so the bottom-most pressure value is weakly latitude/longitude dependent), as the atmosphere is fully optically thick in transmission geometry at pressures greater than this value (Figure 2.3). We convert between pressure and altitude assuming hydrostatic equilibrium, accounting for the non-uniform mean molecular weight between GCM grid cells due to the thermal dissociation of H_2 in hotter regions of the atmosphere. Our calculation assumes a nominal atmosphere in cooler regions that is 83.6% H_2 , 16% He, and 0.4% metal-bearing molecules (the latter having an average mean molecular weight $\mu = 12$) by number.

Accounting for the 3D geometry of our GCM structures, we calculate the slant optical depth,

$$\tau_\lambda = \int \kappa_\lambda ds, \quad (2.1)$$

along a given trajectory after determining the opacity κ_λ in each grid cell as a local function of temperature and pressure (see Section 2.2.2.2 for a discussion of our opacity sources), where ds is the line-of-sight path length through an individual grid cell. The wavelength at which κ_λ is

evaluated is adjusted depending on the line-of-sight velocity of the grid cell v_{LOS} , as per

$$\begin{aligned}\lambda &= \lambda_0 \left(1 - \frac{v_{\text{LOS}}}{c} \right) \\ v_{\text{LOS}} &= -u \sin(\phi) - v \cos(\phi) \sin(\theta) \\ &\quad - \Omega(R_{\text{P}} + z) \sin(\phi) \cos(\theta)\end{aligned}\tag{2.2}$$

for rest wavelength λ_0 , east-west velocity u , north-south velocity v , latitude θ , longitude ϕ , altitude z , speed of light c , and planetary angular rotation speed Ω . These Doppler shifts assume that the orbital velocity of the planet has already been entirely accounted for; we address potential consequences of this assumption with respect to WASP-76 b in Section 2.5.2.1. As described in [Kempton & Rauscher \(2012\)](#), higher-order line-shifting effects such as gravitational redshifting and microturbulent broadening are expected to be negligible.

After calculating the above τ_λ , we can calculate the intensity I_λ along each line of sight, assuming absorption-only radiative transfer:

$$I_\lambda = I_{\lambda,0} e^{-\tau_\lambda}\tag{2.3}$$

for stellar intensity incident upon a grid cell, $I_{\lambda,0}$. We assume that the stellar spectrum follows a blackbody distribution for $T = 6329$ K, the effective temperature of WASP-76 ([Ehrenreich et al. 2020](#)). The total flux transmitted through the atmosphere is arrived at by integrating I_λ over the sky-projected solid angle of each respective grid cell through which a beam emerges. The




Figure 2.4: Representative model spectra computed in this study. The bottommost, gray spectrum is computed without including Doppler shifts from planetary winds and rotation and includes Fe I and Fe II as the sole opacity sources. The middle, gold spectrum is the same as the bottom, but now including Doppler effects. The topmost, maroon spectrum includes Doppler effects and all opacity sources discussed in Section 2.2.2.2. The inset axis is centered near the Ca II H line (393.4 nm); the vertical dashed lines indicate the peak of a representative Fe feature for the gray and gold spectra, respectively. In the optical range of WASP-76 b, the inclusion of Doppler effects broadens and blueshifts strong lines, such as those depicted.

wavelength-dependent transit depth D_λ is then given by

$$D_\lambda = \frac{F_{out} - F_{in}}{F_{out}}, \quad (2.4)$$

where F_{out} is the out-of transit flux (i.e. the stellar flux) and F_{in} is the in-transit flux, accounting for both the blockage of stellar light by the optically thick core of the planet and the transmitted flux through the atmosphere.

Our spectral simulations are computed over the 379 – 789 nm range to coincide with the ESPRESSO data from [Ehrenreich et al. \(2020\)](#), with a total of 298,766 individual wavelength points corresponding to a resolution of $R \approx 4 \times 10^5$. After computation, our spectra are convolved to the native ESPRESSO resolution in the singleHR mode ($R \approx 138,000$; [Pepe et al. 2013](#)).

Because of very large file sizes and RAM requirements associated with the substantial number of wavelength points, we generate the spectra in discrete wavelength “chunks” on a high-performance computing cluster and then stitch the entire spectrum together once the full calculation is complete.

2.2.2.2 Opacity sources

Because WASP-76 b is an ultra-hot Jupiter observed at optical wavelengths, we compile a set of opacities that is appropriate for these conditions. The atomic species that we utilize are Fe I, Fe II, Mg I, Mn I, Na I, Ca II, Ti, Li I, and K I; and our molecular species are TiO, VO, OH, and H₂O. Rayleigh scattering from H₂, H, and He, as well as H⁻ bound-free and free-free absorption are included as continuum opacity sources. Finally, collision-induced opacity from collisional pairs of H₂-H₂, H₂-H, H₂-He, and H-He is considered. This list of opacity sources includes all of the species with reported detections in WASP-76 b optical spectra from [Ehrenreich et al. \(2020\)](#), [Tabernero et al. \(2021\)](#), and [Fu et al. \(2021\)](#). We do not include ZrO, which was searched for in WASP-76 b by [Tabernero et al. \(2021\)](#) but resulted in a non-detection. In modeling WASP-76 b, we generate two sets of transmission spectra, one with Fe (I and II) and continuum opacity sources only, and another with the full set of molecular, atomic, and continuum absorbers (see Figure 2.4 and Section 2.4.1.1).

All opacities are generated on a temperature-pressure grid spanning from 500 – 5000 K and $1 \times 10^{-6} - 1 \times 10^3$ bar, covering the range of the WASP-76 b GCM output. We initially calculate opacities on an evenly spaced wavenumber grid, from a wavenumber of 0 to 30,000 cm⁻¹ with a spacing of 0.01 cm⁻¹. These results are then downsampled with the k-table feature of HELIOS ([Malik et al. 2017](#)).

We use the newest available open-source line lists in constructing our molecular opacities: VO (VOMYT; McKemmish et al. 2016), TiO (TOTO; McKemmish et al. 2019), OH (MoLLIST; Tennyson et al. 2016), H₂O (POKAZATEL; Polyansky et al. 2018). Atomic opacities are drawn from the Kurucz database (Kurucz 1995). In selecting our opacities, we make use of the most abundant isotopologue for each species. Our atomic sources do not include pressure broadening — only Doppler (thermal) and natural broadening.¹ From these line lists, we calculate the relevant opacities with the GPU-enabled HELIOS-K code (Grimm & Heng 2015).

Opacities from individual species are combined by a weighting of their respective mixing ratios, under the assumption that the atmosphere of WASP-76 b resides in a state of chemical equilibrium. To determine the chemical equilibrium mixing ratios, we employ the FastChem model (Stock et al. 2018), which accounts for gas-phase chemistry only. In Section 2.2.2.4, we describe how we treat condensation for the purposes of our modeling.

2.2.2.3 Phase-dependent transmission spectra

As in Flowers et al. (2019), to compare our model transmission spectra directly against the Ehrenreich et al. (2020) results, we must calculate transmission spectra as a function of orbital phase throughout the duration of transit. To account for orbital phase dependencies, we apply the following procedure:

1. *Account for phase-dependent back-lighting of the planet (i.e. stellar limb-darkening effects).*

At different points of its transit, the planet will occult regions of its host star of varying

¹The Doppler effects from planetary rotation and winds dominate over the broadening introduced by pressure, thermal, and quantum effects here, because the signal for high-resolution spectra originates from high-altitude, low-pressure regions of planetary atmospheres. Furthermore, our science case is primarily concerned with the location of line centers, not the exact nature of line wings. Therefore, we anticipate that the exclusion of pressure broadening here does not significantly impact our final spectra.

brightness. Furthermore, at a fixed orbital phase, different regions of the planet’s limb will be back-lit by varying intensities of stellar light. Similarly to [Flowers et al. \(2019\)](#), we calculate the normalized stellar intensity at the center of each cell of the 2-D projected planetary grid produced by our GCM at each modeled orbital phase of the planet. We use the quadratic limb-darkening coefficients reported by [Ehrenreich et al. \(2020\)](#) to establish the stellar center-to-limb intensity profile, and we take into account the 89.623° orbital inclination of WASP-76 b ([Ehrenreich et al. 2020](#)) to determine where the planet resides on the stellar disk as a function of its orbital phase. We make the assumption of constant impact parameter b over the course of transit.² This procedure allows us to calculate a back-lighting factor f , which ranges from 0 to 1, effectively replacing the constant $I_{\lambda,0}$ from Equation 2.3 with a variable $I_{\lambda,0} \times f(\theta', z, \varphi)$, for a given orbital phase φ and 2-D projected polar angle θ' .

2. *Account for the decreasing of the continuum by interpolating a light curve produced by the `batman` code* ([Kreidberg 2015](#)). Step 1 ensures that less light is transmitted through the planet’s atmosphere than a uniform stellar disk would emit. Step 2 further enforces that the inner, optically thick core of the planet is simulated crossing a limb-darkened star, as opposed to a star of uniform brightness.
3. *Account for the planet’s rotation over the course of transit.* Because the planet is continually rotating as it travels across the face of its host star, we must transform the GCM coordinate system so that the correct observer-facing hemisphere is modeled at each instance during

²In reality, a planet on an inclined orbit will not have a constant b over the entire duration of transit; rather, the planet’s distance from the stellar equator will be decreased at ingress and egress, reaching its maximum at center of transit. Our tests reveal that the relative error induced by the constant b assumption is on the order of 4% in distance, which results in a change on the order of 1 m/s at the blueshift level (see Section 2.3.1). Hence, our b treatment is justified.

transit. For simplicity, we assume zero obliquity, which allows us to calculate the coordinate transform simply by assigning a linear offset to each planetary longitude; i.e., $\phi_{\text{rotated}} = \phi + \varphi$.

In choosing phases at which to evaluate our models, we more densely sample ingress and egress because the transmission signal quickly varies in strength and velocity during these times. The sampling that we employ is:

$$d\varphi = \begin{cases} 0.416^\circ, & -15.94^\circ \leq \varphi \leq -8.25^\circ \\ 1.39^\circ, & -8.25^\circ \leq \varphi \leq 6.0^\circ \\ 0.416^\circ, & 6.0^\circ \leq \varphi \leq 15.94^\circ \end{cases} \quad (2.5)$$

where $\varphi_{\text{max}} = 15.94^\circ$ represents the phase at which the planet’s trailing limb no longer occults the stellar disk. This procedure results in 50 total orbital phases being explicitly modeled with our radiative transfer post-processing code.

2.2.2.4 Condensation and clouds

The large blueshifts observed by [Ehrenreich et al. \(2020\)](#) have been attributed to iron condensation on the nightside and cooler “morning” limb of WASP-76 b. To confront this hypothesis within the framework of our transmission spectrum modeling, we account for clouds and condensation in the following ways.

To incorporate the effects of condensation out of the gas phase, we utilize the iron condensation curve for a solar composition mixture computed in [Mbarek & Kempton \(2016\)](#). At each cell in our GCM-produced atmosphere, we interpolate the condensation curve at the cell’s pressure; if the

temperature in the cell is lower than the interpolated condensation temperature, then we remove iron from the total opacity calculation at that location by setting its opacity contribution to zero. This approach maximizes the effect of Fe condensation, as in reality homogeneous nucleation of Fe is inefficient (Gao & Powell 2021). This step effectively treats the physics considered in the Ehrenreich et al. (2020) toy model, which only accounted for gas phase removal as the process responsible for the large net blueshifts in WASP-76 b’s transmission spectrum.

In our full-species modeling of WASP-76 b’s transmission spectrum, we allow other species to condense out of the atmosphere following a similar procedure. We condense Ca (I and II) following the Ca_2SiO_4 condensation curve, Ti and TiO following CaTiO_3 , K following KCl, Na following Na_2S , Mn following MnS, and Cr following Cr_2O_3 . Our prescription of complete removal implies full rainout from the gas phase once the species-limiting condensate forms.

Condensation not only causes gas-phase depletion but can also result in the formation of an optically thick cloud made up of liquid droplets or solid particles. Such a cloud layer would block transmission of stellar light through the atmosphere in the location where the cloud forms and could therefore also be the main driver of the “missing” iron absorption in the receding limb of WASP-76 b’s atmosphere — a scenario that has not yet been explored in the literature for this planet. To account for such cloud layers, our models further contain a post-processing implementation of gray, optically thick clouds. These clouds, which are self-consistent with our radiative transfer (but not with our GCMs), are placed such that they span up to several (1 – 10) scale heights above the condensation curve of iron, forsterite (Mg_2SiO_4), or aluminum oxide (Al_2O_3).

Fe is selected because it is the species that is putatively condensing in WASP-76 b’s atmosphere; Mg_2SiO_4 is modeled because per the microphysics model of Gao et al. (2020), silicate clouds are chemically favored over iron ones and appear to be a defining feature in the transmis-

sion spectra of giant planets with comparable irradiation to WASP-76 b; and Al_2O_3 is considered because it is the highest-temperature condensate predicted in hot Jupiter atmospheres (Mbarek & Kempton 2016).

Notably, our implemented clouds never fully reach the limiting-case thickness of 10 scale heights. Because our temperature-pressure profiles tend to cross condensation curves twice (Figure 2.2), clouds cannot exist in the uppermost regions of the atmosphere — the upper boundary of the cloud is (indirectly) set by the pressure at which the atmospheric inversion or isothermal region occurs. This effect can be seen in the bottom row of Figure 2.5, in which the “10 scale height” clouds do not extend much past the “1 scale height” clouds. Hence, while our clouds can exist up to 10 scale heights in thickness (provided they fulfill their criterion with respect to a given condensation curve), in practice they are limited by the temperature/pressure structure of the atmosphere.

Similarly to our condensation treatment, our cloud effect is modeled *post-hoc*, in that the radiative effects of cloud formation and the resultant changes to opacities are not accounted for natively in the GCM (see Section 2.5.3 for a discussion of this limitation). Furthermore, our simple cloud model does not account for the wavelength-dependent scattering and/or absorption from cloud particles — a simplification that is justified because cloud opacities are typically smoothly varying with wavelength, in contrast with the sharply wavelength-dependent atomic and molecular features that are primarily responsible for the high-resolution transmission spectrum signal in cross-correlation. Additionally, cloud particle sizes are represented by broad distributions, and no clouds in a planetary atmosphere are characterized by a single particle size (e.g., Pont et al. 2013).




Figure 2.5: Column-integrated surface density of Fe abundance (Σ_{Fe}) of WASP-76 b. Each figure column corresponds to an orbital phase (left to right: ingress, center of transit, and egress), and each row corresponds to a model run (top to bottom: strong drag, weak drag, strong drag with condensation, weak drag with condensation, weak drag with Fe clouds). All models shown have a deep RCB. One scale height clouds are drawn with a darker gray, and 10 scale height clouds are drawn with a lighter gray. (In practice, the latter do not achieve a full thickness of 10 scale heights because the upper atmosphere of WASP-76 b presents conditions incompatible with iron condensation — see text for more information.) The strong drag case is more symmetric, especially at center of transit, than the weak drag case.

2.2.3 Suite of models

The suite of post-processed spectral models explored in this work is laid out in Table 4.2. We use the following terminology to refer to our various model implementations. By “condensation on,” we mean that we are applying the rainout prescription described in Section 2.2.2.4. “Fe-only” models refer to those in which the only non-continuum opacity source is gas-phase iron. Our “full-species” models include the full set of opacities listed in Section 2.2.2.2.

Furthermore, when referring to cloud inclusion, “Fe clouds on” implies the placement of an optically thick cloud above the Fe condensation curve with a thickness equivalent to a specified number of scale heights. Similarly, “Mg₂SiO₄” or “Al₂O₃ clouds on” refer to the same, albeit above the Mg₂SiO₄ or Al₂O₃ condensation curves, respectively. The inclusion of clouds in our models is decoupled from our condensation procedure, allowing us to model these effects together or separately.

The majority of our work holds WASP-76 b on a circular orbit (i.e., $e = 0$). We investigate the effect of eccentric orbits in Section 2.5.2.1.

2.3 Analysis

2.3.1 Cross-correlation

With our model spectra in hand, our next step is to determine the information that can be recovered from these spectra. For high spectral resolution observations of exoplanet atmospheres, signal recovery has typically been accomplished via a cross-correlation analysis (e.g., Snellen et al. 2010; Birkby et al. 2013; Brogi et al. 2016). This procedure has an advantage over the fitting of

individual spectral lines (Section 2.3.2) in that it leverages a key aspect of high-resolution data — a forest of weak lines that together can produce a high-SNR detection of a chemical species. To mimic the data analysis process undertaken by Ehrenreich et al. (2020), thereby allowing us to compare our models directly to their results, we implement a cross-correlation procedure into our modeling routine.

Our cross-correlation function (CCF) c is computed by combining our data, x , with a mask or template, T , at a given velocity, v , as follows (e.g., Baranne et al. 1996; Pepe et al. 2002; Allart et al. 2017; Hoeijmakers et al. 2019):

$$c(v) = \sum_{i=0}^N x_i T_i(v) \quad (2.6)$$

where the mask or template is Doppler-shifted by velocity v and interpolated onto the wavelength grid of x for summing. This calculation produces a map of cross-correlation against velocity (Figure 2.6); the location of the peak of this distribution is reported as the net Doppler shift of x . We Doppler-shift T between -250 km s^{-1} and 250 km s^{-1} in steps of 1 km s^{-1} .

We first employ a model-on-model cross-correlation. To do so, we construct as our template a simplified, “Doppler-off” spectrum, which is a model otherwise equivalent to the simulated data but with no Doppler shifts applied in the radiative transfer calculation. This template is then convolved down to the resolution of ESPRESSO using a Gaussian kernel of appropriate width. We then perform the previously described cross-correlation procedure.

Once we perform our cross-correlation, we fit a Gaussian function to our CCF with the `scipy curve_fit` routine. The peak value of the fitted Gaussian is then reported as the velocity shift of a given spectrum.




Figure 2.6: Normalized cross-correlation functions of VO (gray), Fe I (black), and Ca II (gold). Note that while the CCFs of Fe I and Ca II have clearly defined peaks (indicated by vertical dashed lines), the VO CCF is very noisy, with no discernible peak value. Most of our computed CCFs are similarly bimodal in quality.

For CCFs with well-characterized peaks, we perform a second, narrower Gaussian fit in a window 6 km s^{-1} wide centered on the peak identified by the previous step. Doing so ensures that asymmetries in the CCF far from the peak do not influence the final determination of the peak value. Our fitting process allows us to report CCF peak values to precision higher than the 1 km s^{-1} sampling of our CCFs; limited testing of sampling CCFs at higher velocity resolution yielded results consistent with our chosen sampling. CCFs that are not well-characterized (e.g., VO; Figure 2.6) are not considered in further analysis.

2.3.2 Individual line-fitting

For certain strong lines, we perform individual fitting to determine Doppler shifts of those lines. This process is an important complement to the cross-correlation procedure described above, and it allows us to compare our forward models directly to the observational results of [Tabernero et al. \(2021\)](#), who focused on specific strong lines of a range of chemical species. Individual-line

Doppler shifts are also scientifically interesting because lines of different strengths probe different altitudes of the planet’s atmosphere, and hence different portions of the planet’s wind field and temperature structure (e.g., [Kempton & Rauscher 2012](#)).

To fit individual lines, we fit a Gaussian profile to a narrow spectral window centered on the line in our center-of-transit spectrum. The corresponding Doppler shift (v_{Doppler}) is calculated by comparing the wavelength at the maximum of the fit Gaussian ($\lambda_{\text{max,fit}}$) to the rest wavelength (λ_0):

$$v_{\text{Doppler}} = c \left(1 - \frac{\lambda_{\text{max,fit}}}{\lambda_0} \right) \quad (2.7)$$

For ease of comparison to the [Tabernero et al. \(2021\)](#) data, we fit the same lines examined by that study. Furthermore, we adopt their method for fitting weak lines: As an intermediate approach between cross-correlation of many lines and fitting a single line, [Tabernero et al. \(2021\)](#) combine multiple (maximum 5) weak lines in velocity space and fit them jointly to produce a stronger signal. As with that work, we apply this approach for Fe, Mn, and a subset of Mg lines.

This line-fitting method allows us to probe both the Doppler shifts of individual lines and their depths — which in turn establish the altitude at which these lines become optically thick.

2.4 Results

2.4.1 Cross-correlation

2.4.1.1 Fe-only

For our first set of results, we perform our analysis on our iron-only transmission spectra.




Figure 2.7: Representative planet-frame Doppler shifts as a function of orbital phase for our Fe-only models, as computed by our cross-correlations of forward-modeled spectra. Overplotted are data from the [Ehrenreich et al. \(2020\)](#) WASP-76 b observations and their corresponding 1σ error bars (corresponding to the top panel of that work’s Extended Data Figure 7). Our worst-fitting model consistent with our GCM has strong drag ($\tau_{\text{drag}} = 1e3$ s) and a shallow RCB (solid gray line). Conversely, our best-fitting model consistent with our GCM has weak drag ($\tau_{\text{drag}} = 1e7$ s) and a deep RCB (dashed teal line). Finally, our best-fitting model with 1 scale height clouds inclusive of condensation effects has weak drag, a deep RCB, and Al_2O_3 composition (dotted gray line).




Figure 2.8: Terminator line-of-sight velocity maps of our deep RCB WASP-76 b GCMs, including contributions from both winds and rotation, plotted in the style of Figure 2.3 (two drag timescales and three phases). Our (Fe) modeled clouds are over-plotted as a gray band in the bottom row, for the 10 scale height implementation. Our clouds, as opposed to our wind field, is not a terminator slice, but rather a depiction of every (projected) sightline that intersects a cloud in the full 3D geometry. As in maps depicting other quantities, the line-of-sight velocities are more symmetric for the high-drag case, while they are more asymmetric for the weak-drag case. Stronger blueshifting is further present at egress than ingress for the weak-drag case.

We find that controlling the drag timescale over the explored parameter range has the greatest first-order effect on Doppler shift over all phases (Figure 2.7), with the weak drag models providing a much better match to the [Ehrenreich et al. \(2020\)](#) data than the strong drag models. Decreasing the drag timescale in an atmosphere is equivalent to making drag processes more efficient. In the Helmholtz decomposition framework of [Hammond & Lewis \(2021\)](#), about half the day-night heat redistribution in a fiducial hot Jupiter atmosphere is driven by substellar-antistellar flow; shorter drag timescales would serve to slow the speed of this flow, reducing heat transport. The remainder of the day-night heat redistribution, per [Hammond & Lewis \(2021\)](#), is due to rotational flow — e.g., the superrotating equatorial jet. The speed of this jet would similarly suffer directly from a shorter drag timescale. Additionally, per [Showman et al. \(2013a\)](#), stronger drag also reduces the pumping of eddy angular momentum onto the dayside equator because drag damps the propagation of planet-scale wave patterns, further decreasing the jet strength. In sum, increased drag reduces longitudinal asymmetries and slows wind speeds, both of which reduce net Doppler shifts in transmission spectra.

As a result of the aforementioned effects, our high-drag models have more homogenized limbs, whereas low-drag models have limbs that are more heterogeneous, with a visible equatorial jet structure (Figures 2.3, 2.5, & 2.8). These features can be seen in our simulated Doppler shifts (Figure 2.7). The strong drag case has Doppler shifts that are roughly symmetric about the center of transit, with only a slight ($< 2 \text{ km s}^{-1}$) blueshift at mid-transit. Moreover, the maximum and minimum of the Doppler shift are both near $\pm 4\text{--}5 \text{ km s}^{-1}$, which is close to the equatorial rotational velocity of WASP-76 b ($\pm 5.3 \text{ km s}^{-1}$). These features are expected for a solely rotationally broadened profile — which would be predicted for the high-drag models, which have weaker winds. Additionally, a clear prediction from Figure 2.8 is that the strong drag case is redshifted at

ingress and blueshifted at egress — as seen in Figure 2.7.

Similarly, the Fe abundance map (Figure 2.5) combined with the wind map (Figure 2.8) explains the behavior of the weak drag case in Figure 2.7. Namely, the Fe abundance is asymmetrically distributed, preferentially on the eastern limb, where the strongest blueshifts for this model are. Therefore, the model produces Doppler shifts that are not evenly distributed around 0 km s^{-1} across phase and are more strongly blueshifted, which is more consistent with the Ehrenreich et al. (2020) observations.

Including iron condensation removes the gas-phase iron in the cooler regions of the atmosphere (Figure 2.5). As can be inferred from our temperature-pressure structure (Figure 2.2), condensation is primarily confined to the nightside in all our GCMs. Especially in our weak drag models, condensation is less prominent at the equator, whereas it is more prominent at the mid-latitudes and poles. The effect of condensation tends to be marginal, dependent on cloud placement, and not necessarily linear, reflective of the complexities in the 3D thermal structure of the atmosphere.

We further find that our deep RCB models are a better fit to the data, generating stronger net blueshifts. A shallow RCB is expected to create a more homogeneous atmosphere, because convective motions quickly mix entropy across large spatial scales, reducing day-night temperature contrasts on deep isobars relative to the deep RCB case. Atmospheric circulation in (ultra-)hot Jupiters is driven and supported by day-night pressure gradients (e.g., Showman & Guillot 2002), so the smaller day-night contrasts in the shallow RCB case cause winds to be slower. These slower winds in turn decrease east-west limb asymmetries as well as the strength of the blueshift signal (as shown in Figures 2.7 & 2.9).

We find that the presence of optically thick clouds, as opposed to gas-phase iron condensation, also strongly and consistently increases the measured blueshift (Figure 2.7). This is because

clouds will preferentially form on the cooler, receding (eastern) limb, hence blocking its contribution to the transmission spectrum; these clouds would lessen the redshift contribution of the receding limb, thereby boosting the measured blueshift.

As seen in Figure 2.2, Al_2O_3 condenses out at hotter temperatures than Fe. Hence, placement of clouds above the Al_2O_3 condensation curve as opposed to the Fe condensation curve results in a cloud deck situated in warmer regions of the atmosphere. The effect of this change at the blueshift level is highly nonlinear, as it depends somewhat strongly on the exact placement (and thickness) of the cloud. Notably, though, the optically thick clouds produce a “jump” in blueshift near a phase of -7° that the other models cannot (Figure 2.7) — a distinct feature of the [Ehrenreich et al. \(2020\)](#) data. While this jump does not extend to the exact observed magnitude, it appears that an Al_2O_3 cloud deck (or a similarly positioned cloud of arbitrary composition) provides the type of spatial inhomogeneity required to reproduce the [Ehrenreich et al. \(2020\)](#) data in broad strokes.

Interestingly, although the deep RCB / strong drag GCM would have much more condensation in its atmosphere (Figure 2.2), the deep RCB / weak drag GCM has much stronger blueshifts (Figure 2.9). It thus appears that the importance of the drag timescale and faster winds leading to more east-west asymmetry “wins out” over the existence of condensation alone.

Our net center-of-transit Doppler shifts are summarized in Figure 2.9 for all of the model permutations explored. In total, we find that the following effects produce increased blueshifts in WASP-76 b’s atmosphere when iron is considered alone: weak drag, deep RCBs, and optically thick clouds composed of Al_2O_3 , Mg_2SiO_4 , or Fe. Gas-phase iron condensation increases net blueshifts in some cases, but (contrary to the [Ehrenreich et al. 2020](#) interpretation) actually decreases net blueshifts in others.

By toggling our back-lighting effect, we determine that accurate treatment of limb-darkening




Figure 2.9: Center-of-transit Doppler shifts for all models explored in this work. From left to right, we show the effect of increasing the drag timescale of our GCM (i.e., decreasing the effect of drag within the atmosphere). Within a given row, we include gas-phase condensation for some models (“True” vs. “False”), as described in Section 2.2.2.4. For the weakest drag timescales, we also explore the effect of placing optically thick clouds above different species’ condensation curves (Mg_2SiO_4 , Fe, and Al_2O_3), as described in Section 2.2.2.4. Cells representing models that are not computed are filled in black. The upper sets of models presented here only includes iron as a spectral opacity source (“Fe-only”), whereas the lower sets of models include all gas opacity sources discussed in Section 2.2.2.2 (“Full Species”). For our full species models, we explore the effect of cross-correlating against different single-opacity template spectra (e.g., “Full species: Fe” convolves against a full species spectrum against an Fe template) and increasing our clouds’ scale heights (left to right: 1, 5, 10 scale heights). The strongest blueshift is detected for the all-species model / Fe cross-correlation mask with condensation on, and 10 scale height thick Al_2O_3 clouds. Aside from the top row, all models are computed with a deep RCB. For reference, the observed transit-averaged Doppler shift identified by Tabernero et al. (2021) via Fe cross-correlation is $-8.25 \pm 0.25 \text{ km s}^{-1}$ for the first analyzed transit (T1) and $-8.75 \pm 0.56 \text{ km s}^{-1}$ for the second analyzed transit (T2). Line analysis by Tabernero et al. (2021) in T1 yielded Doppler shifts of $4.1 \pm 5.1 \text{ km s}^{-1}$ and $-4.4 \pm 2.5 \text{ km s}^{-1}$ for the Ca II H&K lines, respectively. T2 analysis produced Doppler shifts of $1.0 \pm 3.0 \text{ km s}^{-1}$ and $-2.1 \pm 1.9 \text{ km s}^{-1}$ for these lines.

(see Section 2.2.2.3) has a median effect over the course of transit on the order of 0.1 km s^{-1} . Considering the mean errorbar size on the Ehrenreich et al. (2020) blueshift data (Figure 2.7), it appears that although including this effect improves physical consistency, it does not have strong (or even currently detectable) effects at the Doppler shift level.

2.4.1.2 Full species

For our next set of model runs, we include all of the opacity sources described in Section 2.2.2.2 and cross-correlate these “full species” models against Doppler-off templates that include one chemical species at a time. We are able to identify peaks in the CCFs at most in-transit phases for Fe, Ti, Mn II, Ca II, and Cr, indicating that these species are present at an observable level in our modeled spectra.

Similarly to the Fe-only models, we find that spectrally post-processed, unaltered GCMs cannot reproduce the Ehrenreich et al. (2020) piecewise blueshift trend without added modifications such as condensation or cloud formation (Figure 2.10, panels (a) and (b)). Most of the detected species tend to follow the trend of Fe, especially in the weak drag case. The clear exception is Ca II. It exhibits much weaker blueshifts than the other species in the weak drag case, and is mostly redshifted over the course of transit in the strong drag case. Interestingly, its Doppler shift remains mostly constant out to egress in the weak drag case, whereas it experiences a sharper blueshift in egress in the strong drag case. This trend implies that in the weak drag case, the Ca II signature is proportionally stronger in the approaching limb, so as the receding limb exits the stellar disk in egress, the Doppler signal remains largely unchanged until the approaching limb exits as well. In contrast, in the strong drag case, it appears that the Ca II signal is more evenly distributed on

both limbs, so when the receding limb begins to exit the stellar disk, the total contribution of Ca II signal shifts to the approaching limb, rapidly increasing the measured blueshift over egress due to the rotational component.

These characteristics can be physically motivated by considering the temperature structure of the atmosphere. The abundance of Ca II in equilibrium chemistry is strongly dependent on thermal ionization, and hence temperature. Any temperature asymmetries, then, will be reflected in asymmetries in the abundance profile of Ca II. Per Figure 2.3, our weak drag case is much hotter on the approaching limb, because the lack of strong drag is conducive to the formation of a superrotating equatorial jet that advects heat from the planet’s hotspot, causing a thermal offset in the direction of the approaching limb. Therefore, in the weak drag case, the approaching limb will be more abundant in Ca II than the receding limb. In the strong drag case, however, the lack of a jet and hotspot offset imposes limb homogeneity, giving each limb relatively equal weight in the Ca II signal. See Section 2.5.2.2 for a further discussion of Ca II and its potential non-hydrostatic behavior in observations.

Adding clouds to our models (Figure 2.10, panel (c)) has the same effect identified in the Fe-only models: stronger blueshifts and a trend that better matches the Ehrenreich et al. (2020) data for clouds with higher condensation temperatures.

Notably, the cloud thickness seems to strongly control the jump in blueshift prior to center of transit — the greater the maximum thickness of the cloud, the earlier in phase the jump occurs. The best-fitting model in this grid includes clouds up to 10 scale heights above the Al_2O_3 condensation curve (corresponding to a 0.7 mbar cloud top on the western limb and a 1.1 mbar cloud top on the eastern limb). A cloud top up to 10 scale heights above the Fe condensation curve performs nearly, but not quite as well. Even in the Al_2O_3 cloud case, though, the fit to the data is not perfect;

namely, the magnitude of our Doppler shift is a few km s^{-1} short of the observed data at egress. (Figure 2.10, panel (c)).

Contrary to the cloudless case, including gas-phase condensation when clouds are also included increases the measured blueshift (Figures 2.9 & 2.10). This effect is often very slight; in the 10 scale height, Al_2O_3 cloud case, including condensation increases blueshift by 6 m s^{-1} at center of transit.

The overall best fit to the Ehrenreich et al. (2020) data from our suite of models is obtained for the case of low drag ($\tau_{\text{drag}} = 10^7$), a deep RCB, a 10 scale height, optically thick Al_2O_3 cloud, and gas-phase iron condensation. Even this model, however, is not perfect. For instance, this model fails to match the egress blueshift magnitude of the Ehrenreich et al. (2020) data (Figure 2.7). We discuss these points in further detail in Section 2.5.

2.4.2 Line-fitting

Overall, we are able to detect and fit nearly all of the individual absorption lines explored by Tabernero et al. (2021). We additionally find that our line Doppler shifts are all consistent within 2σ with the Tabernero et al. (2021) data (Figure 2.11), and nearly all our line heights are consistent within 1σ for at least one of the Tabernero et al. (2021) transits (Figure 2.12).

As also seen in our cross-correlation analysis, the lines from our weak-drag models have stronger blueshifts than our strong drag models. Condensation has a negligible effect on the Doppler shifts of the non-condensed species and a small effect on the Doppler shifts of the condensed species. Introducing clouds has the expected impact of increasing blueshift across species (Figure 2.11). Yet some species (Ca II, Mn I) are more affected by the inclusion of clouds than




Figure 2.10: Phase-resolved CCFs of our spectra including all opacity sources. Subplot (a) is the strong drag case, subplot (b) is the weak drag case, subplot (c) showcases the weak drag case for varying cloud treatments (only convolving with an Fe template), and subplot (d) compares the condensation on/off cases. Most species with detectable CCFs follow the Doppler trend of Fe I, with the notable exception of Ca II. Note that the CCF of Ca II is too noisy to fit during ingress in subplot (c). As in the Fe-only models, introducing thick clouds to our full-species spectrum reproduces the sharp jump in Doppler shift noted by [Ehrenreich et al. \(2020\)](#) and [Kesseli & Snellen \(2021\)](#), and the condensation effect alters the fit for condensed species depending on whether clouds are included; non-condensed species are unaffected by condensation treatments. Our best-fitting model (up to 10 scale height Al_2O_3 clouds, weak drag, condensation on, deep RCB) is shown in black in panel (d).




Figure 2.11: Doppler shifts investigated by [Tabernero et al. \(2021\)](#) in comparison with the lines from the forward models in our study. Including 10 scale height, gray Fe clouds in our model (purple squares) blueshifts our lines to varying degrees across species. The weak drag models are slightly preferred over the strong drag case. Although endmembers of drag (teal circle for strong drag, teal diamond for weak drag) are plotted for clarity, intermediate drag seems to be slightly more consistent with the observed Doppler shift data overall.




Figure 2.12: Height of lines investigated by [Tabernero et al. \(2021\)](#) in comparison with the lines from representative forward models in our study. Our strong drag case is plotted with teal circles, whereas our weak drag case is plotted with teal diamonds. Most of our lines are consistent with the [Tabernero et al. \(2021\)](#) line heights — with the exception of the Ca II lines, which have observed heights consistent with an escaping atmosphere that extends beyond our model domain.

others (Na I, Li I). This is not surprising, as different species are expected to be abundant in different temperature-pressure regimes, and they would thus be blocked by our pressure-dependent cloud treatment to differing extents.

Broadly, as with the [Ehrenreich et al. \(2020\)](#) data, the [Tabernero et al. \(2021\)](#) Doppler shift data favor the weakest drag case over the strongest, though they cannot discriminate between different condensation cases. The line height data cannot favor one model over another, given their error bars. Even so, both data sets can be used to demonstrate an overall degree of consistency of our models with the available data.

All of the strong lines in our forward models fall within the altitude range of $1.078 R_P$ to $1.095 R_P$ (Figure [2.12](#)). For most species, this is commensurate with the [Tabernero et al. \(2021\)](#) line height data. Notably, however, our models significantly fail to reproduce the observed Ca II line strengths (see Section [2.5.1.2](#)).

2.5 Discussion

2.5.1 Comparison to other work

2.5.1.1 Comparison to the [Ehrenreich et al. \(2020\)](#) “Toy Model”

As shown in Figure [2.8](#), our 3D, physically self-consistent GCM output wind field is more complex than the substellar-antistellar and rotational wind field described by the [Ehrenreich et al. \(2020\)](#) toy model. For instance, blueshifts from winds are not uniform across the entire limb, and there even exists a redshifted annulus on the terminator (bottom row of Figure [2.8](#)), which

represents return flow from the nightside to the dayside.³ Furthermore, our GCMs are unable to reproduce the wind speeds required by the [Ehrenreich et al. \(2020\)](#) toy model. Our weak drag GCM has an eastern limb average line-of-sight velocity of -6.5 km s^{-1} ; this is in excess of the -5.3 km s^{-1} day-night flow in the toy model. When averaged with the western limb, though, the average line-of-sight velocity reduces to -2.3 km s^{-1} — significantly less than the toy model wind field.

A crucial component of the [Ehrenreich et al. \(2020\)](#) toy model is an asymmetry in the Fe abundance distribution in the atmosphere of WASP-76 b. Broadly, our iron abundance maps (Figure 2.5) indicate that, regardless of drag timescale, there does exist a projected asymmetry of iron on the eastern limb by the end of transit. The eastern limb is preferentially blueshifted (Figure 2.8), so more gas-phase iron on the eastern limb produces a larger measured blueshift. While this asymmetry is exacerbated by gas-phase condensation of Fe (top two rows vs. next two rows of Figure 2.5), our self-consistent, GCM-based forward models with Fe condensation alone do not readily match either the magnitude or the shape of the [Ehrenreich et al. \(2020\)](#) observed blueshifts, and we find that we must invoke additional physics (e.g., clouds) in order to achieve a better fit.

2.5.1.2 Comparison to [Tabernero et al. \(2021\)](#)

As shown in Figures 2.11 and 2.12, our forward models generally provide good agreement with the multiple species detected across the entire ESPRESSO wavelength range by [Tabernero](#)

³Specifically, we can confirm that this redshifted annulus represents the western flank of the Rossby gyres. [Tsai et al. \(2014\)](#) demonstrate that the shift of a planet’s hotspot east of its substellar point is not only a product of heat advection, but is also indicative of a phase shift of planet-scale standing wave (Matsuno-Gill) patterns. Our deep RCB models have stronger day-night contrasts than our shallow RCB models — hence, the former have stronger equatorial jets and larger phase shifts in their Matsuno-Gill patterns than the latter. The redshifted annulus present in the shallow RCB model (not plotted) can be reproduced in the deep RCB model by rotating the latter into negative phase, which effectively cancels out the additional phase shift.

et al. (2021). However, we bring up several noteworthy exceptions below.

Tabernero et al. (2021) report a detection of $H\alpha$ in one of their two analyzed transits. In our Doppler-off, single-species template models, though, we note that the hydrogen Balmer lines do not protrude above the continuum (comprised of collision-induced absorption, H^- opacity, and Rayleigh scattering). However, when performing a blind search in our full-species spectrum, we find that there does exist an absorption line within the $H\alpha$ spectral region. When compared to the $H\alpha$ rest wavelength, this line (which we identify as a blend of TiO and VO) displays a blueshift relative to rest $H\alpha$ that is 2σ consistent with the reported $H\alpha$ data. We propose that either $H\alpha$ was mistakenly identified by Tabernero et al. (2021) or that NLTE, non-hydrostatic, and/or non-equilibrium chemistry effects may be boosting the strength of the $H\alpha$ line compared to what is predicted by our GCM post-processing scheme.

Conversely, Tabernero et al. (2021) were not able to detect Cr or Ti in the WASP-76 b high-resolution transmission spectrum, whereas our chemical equilibrium forward models are readily able to identify these species in cross-correlation. This mismatch may point to unaccounted-for opacity sources that wash out the Cr and Ti features, increased continuum absorption that masks the lines of these species (perhaps excess H^- brought about by photoionization, for example), or non-solar abundances of Cr and Ti.

A notable result of the Tabernero et al. (2021) work is the height of their reported Ca II lines. The Ca II H line was detected at a height of $1.57 R_p$, and the Ca II K line was detected at a height of $1.78 R_p$ — both at high significance and at much higher altitudes than in our models. The error bars on these detections are large enough such that they are still consistent with our model’s line heights at 2σ . However, the consistency of the observed Ca II H&K line heights with one another supports their fidelity, as their similar line strength implies that they should become optically thick at roughly

the same altitude. Furthermore, the deep absorption of the Ca II triplet near 850 nm detected by two other instruments is supporting evidence for abundant high-altitude Ca II in WASP-76 b (Casasayas-Barris et al. 2021; Deibert et al. 2021). Finally, high-altitude Ca II H&K lines have been detected in other ultra-hot Jupiter atmospheres and have been attributed to hydrodynamic outflows and high-altitude photoionization (Yan et al. 2019; Borsa et al. 2021). Therefore, the discrepancy with respect to our models may be linked to the hydrostatic assumption of our GCM (see Section 2.5.2.2), which only extends up to $1.19 R_P$. Exploring the outflow hypothesis is beyond the scope of this chapter, and we encourage further non-hydrostatic investigations of this and similar planets.

2.5.1.3 Comparison to Wardenier et al. (2021)

Similarly to this work, Wardenier et al. (2021) aimed to post-process a 3D GCM with a radiative transfer code to explore the role of Fe condensation in WASP-76 b. Given that they use a different GCM (the nongray SPARC/MITgcm; Showman et al. 2009) and different radiative transfer scheme (Monte Carlo photon-tracking) from those used in this work, comparing basic results of both works serves to validate both approaches.

Our results are generally consistent with the 3D post-processing results of Wardenier et al. (2021). Both studies favor weak drag over strong drag models and resort to GCM modifications to reproduce the Doppler shift behavior reported by Ehrenreich et al. (2020) and Kesseli & Snellen (2021).

Other details of the modeling approach differ between our two works. For example, Wardenier et al. (2021) are able to use their non-gray GCM to explore the effect of optical opacities

on the predicted thermal structure. Our work explores a much larger spectral range than [Wardenier et al. \(2021\)](#), enabled by our more computationally efficient ray-striking radiative transfer approach. This method allows us to avoid potential biases incurred by only modeling a small portion of the Fe I band structure. (We find this to be up to a 1.5 km s^{-1} effect that can further alter the Doppler shift trend over phase, due to Fe lines over a narrow wavelength range only probing a commensurately narrow range of altitudes.) We further model additional opacity sources beyond those considered by [Wardenier et al. \(2021\)](#).

Notably, [Wardenier et al. \(2021\)](#) were unable to reproduce the post-ingress jump in Doppler shift exhibited by the [Ehrenreich et al. \(2020\)](#) and [Kesseli & Snellen \(2021\)](#) data without artificially restricting the longitude range of iron condensation within their model domain or removing TiO and VO opacity from their GCM opacities (thus altering the thermal structure of the planet’s atmosphere). We are able to produce the same jump in our simulations — effectively muting portions of the gas-phase Fe Doppler shift signal — but by using a more physically motivated model involving optically thick clouds (Figure 2.10). At no point do we actually change the GCM output (and hence the underlying physics); rather, we use *post-hoc*, painted-on effects to reproduce the Doppler shift observations.

2.5.1.4 Consistency with May & Komacek et al. 2021

By comparing the results presented in this chapter with those from our companion work on the Spitzer phase curve of WASP-76 b ([May & Komacek et al. 2021](#)), we can attempt to seek out a set of globally consistent models that describe the ensemble of high-resolution and broadband photometric data. We find that models with cold interiors (i.e., a deep RCB) are favored by both

sets of observations, as such a model allows for the cold nightside and large phase curve amplitude observed at $4.5\ \mu\text{m}$, constrained by the Spitzer study. However, the phase curve data prefer the strong drag ($\tau_{\text{drag}} \leq 10^4\ \text{s}$) GCMs, driven largely by the near-zero phase curve offset observed in both Spitzer wavelength channels. This result is in contrast with our own conclusions that the high-resolution data can only be explained by weak drag models, which allow for a significant east-west limb asymmetry.

The mixed success of our joint modeling efforts to consistently describe all observables — hot spot offsets, phase curve amplitudes, phase-resolved Fe blueshifts, and multi-species blueshifts — hints that our physical understanding of WASP-76 b is as of yet incomplete.

2.5.2 Alternative explanations for the anomalous blueshift

2.5.2.1 A small eccentricity

We have shown that our forward models can reproduce much of the observed Doppler shift trend once optically thick clouds are included (Figure 2.7, Figure 2.10, panel (c)). The full magnitude of the observed Doppler shift, however, is not fully accounted for by this approach. The orbit of WASP-76 b itself may provide the rest of the solution.

A Kozai-Lidov mechanism is often invoked in the literature to explain the migration of hot Jupiters to their current-day short-period orbit (for reviews, see [Dawson & Johnson 2018](#) and [Fortney et al. 2021](#)). Any Kozai-like process, though, would result in vestigial eccentricity, as the mechanism presupposes a high past eccentricity that then enables tidal interactions to reduce the semi-major axis. Other formation mechanisms also have the potential to drive large eccentricities for Jupiter-mass planets as well — for instance, migration through the disk into a magnetospheric




Figure 2.13: Similar to Figure 2.7, but now including the effect of a small eccentricity, both with (maroon lines) and without (teal lines) clouds. Circular orbit models are plotted with dashed lines; best-fit, eccentricity-inclusive models are plotted with solid lines; and MCMC posterior draws are plotted in faded lines. Our best-fit combination of eccentricity and longitude of periastron (solid teal) is consistent with observations (West et al. 2016; Fu et al. 2021). This work’s best explanation of the Ehrenreich et al. (2020) Doppler shift signature includes 10 scale height Al_2O_3 clouds, condensation, and a small eccentricity. The models shown here both include weak drag and a deep RCB.

cavity could do so (Debras et al. 2021). Finally, hot Jupiters subject to other potential formation channels (even such as in-situ formation; e.g., Batygin et al. 2016) could have initially low eccentricities excited by secular interactions (e.g., Wu & Lithwick 2011) or an external perturber (e.g., Zakamska & Tremaine 2004).

As demonstrated by Montalto et al. (2011), the effect of even a small ($e = 0.01$) unaccounted-for eccentricity can produce velocity shifts for hot Jupiters on the order of km s^{-1} . With respect to WASP-76 b, West et al. (2016) placed a 3σ upper bound on eccentricity at 0.05. Fu et al. (2021) additionally constrained the planet’s eccentricity as $0.016^{+0.018}_{-0.011}$. A nonzero — and, for our purposes, significant — eccentricity is therefore certainly within the realm of possibility given the extant data.

In the analyses of Ehrenreich et al. (2020) and Tabernero et al. (2021), the eccentricity of

the planet is held to 0. The authors’ rationale appeals to a short (on the order of tens of Myr) circularization timescale, the age of the star (~ 2 Gyr), a low stellar rotation rate (~ 0.03 d), and observational constraints favoring small eccentricities. If we allow for even a small eccentricity in our models (perhaps due to the chance timing of another body’s perturbation, or recent delivery of WASP-76 b to its current orbit), however, we find that we do not need to appeal to *ad hoc* removal of iron at specific longitudes (the favored explanation of [Wardenier et al. 2021](#)) to produce the observed blueshift signature. Because WASP-76 b subtends more than 30° of phase during its transit, including a small eccentricity can both increase the magnitude of our blueshift at every point in transit and produce a trend consistent with the [Ehrenreich et al. \(2020\)](#) results.

Using the `exoplanet` package ([Foreman-Mackey et al. 2021a](#)) to model planet-frame RV deviations from a circular orbit within a χ^2 fitting routine (see Appendix C for fitting details), we find that the data are best explained by an eccentricity of $e = 0.018$ and a longitude of periastron $\omega = 125^\circ$ for a cloudless atmosphere.⁴

But eccentricity alone cannot produce the data’s characteristic jump in Doppler shift post-ingress that our optically thick cloud models reproduced. Combining 10 scale height Al_2O_3 clouds and condensation with a small eccentricity produces this work’s best qualitative model fit (Figure 2.13, with $e = 0.017$ and $\omega = 117^\circ$). The quantitative fit also improves correspondingly: When including eccentricity, the χ^2 value for the clear case decreases from 401 to 121, and the χ^2 value for the cloudy case decreases from 329 to 112. The blueshift at the correct phases is increased by the eccentricity trend, thereby fully reproducing both the magnitude and time dependence of the [Ehrenreich et al. \(2020\)](#) data.

⁴As noted by [Showman et al. \(2013a\)](#), the quantity $e \cos(\omega)$, which can be derived from transit or secondary eclipse observations, is more constraining of an anomalous center-of-transit blueshift than e alone.

Both our cloudless and cloudy eccentricity-inclusive modeling results are consistent within 1σ with the ω derived from [Fu et al. \(2021\)](#) ($62^{+67}_{-82}^\circ$). Furthermore, both estimates are well within the 3σ limit derived by [West et al. \(2016\)](#).

While the above eccentricity argument is perhaps satisfying, it may not accurately represent reality. Increasing blueshift as a function of phase has also been observed in another planet, the ultra-hot Jupiter WASP-121b ([Bourrier et al. 2020a](#); [Borsa et al. 2021](#)). It is perhaps more likely that there exists an underlying physical mechanism common to both WASP-121b and WASP-76 b than their orbital configurations happening to be similar; the parameter estimation of WASP-121b at discovery constrains $e < 0.07$ at 3σ ([Delrez et al. 2016](#)), and later studies have held that planet to a circular orbit ([Evans et al. 2016, 2017](#); [Mikal-Evans et al. 2019](#); [Borsa et al. 2021](#)). Furthermore, [Bourrier et al. \(2020a\)](#) place a tighter constraint on eccentricity ($e < 0.0078$) at 3σ from TESS data. It is yet to be determined whether a post-processed GCM could account for the blueshift trend in WASP-121b without modifications such as added eccentricity (e.g., as in the weak drag, deep RCB case in Fig. 2.7). Hence, the question of whether cloudless models or the addition of optically thick clouds alone could provide a population-level explanation of phase-resolved phenomena warrants further investigation.

2.5.2.2 Non-hydrostatic effects

As explored in Section 2.5.1.2, the presence of non-hydrostatic effects in WASP-76 b’s atmosphere is supported by the behavior of the Ca II H&K lines observed in [Tabernero et al. \(2021\)](#). In turn, it is possible that a hydrodynamic outflow may be responsible for some portion of the anomalous blueshift signature.

By virtue of its close proximity to its host star, WASP-76 b should experience strong UV radiation flux; such irradiation can result in mass loss in its upper layers (e.g., [Erkaev et al. 2007](#); [Des Etangs et al. 2010](#); [Linsky et al. 2010](#); [Ehrenreich et al. 2015](#)). Indeed, [Seidel et al. \(2021b\)](#) recently found evidence for high-altitude vertical winds in WASP-76 b, which may deliver material to a hydrodynamic outflow (or serve as the base of the outflow itself) in the planet’s exosphere. As proposed for WASP-121b ([Bourrier et al. 2020a](#)), an anisotropic expansion of the atmosphere could lead to variable Doppler shifts over the course of transit.

Both our ray-striking radiative transfer and our GCM presuppose that WASP-76 b’s atmosphere is in local hydrostatic equilibrium. While our assumption may be valid strictly in the pressure regimes that we consider in our study — simulations of atmospheric loss generally begin in the nanobar regime (e.g., [Murray-Clay et al. 2009](#); [Bourrier et al. 2015, 2016](#)), whereas the upper boundary of our atmosphere is at a pressure of $\approx 11 \mu\text{bar}$ — enforcing the hydrostatic assumption prevents us from exploring this hypothesis. Fully non-hydrostatic models that include winds driven by UV radiation would be required to determine whether a physically plausible outflow could become optically thick enough at the correct phases to meaningfully contribute to the WASP-76 b Doppler shift trend.

2.5.3 Limitations of our model

As with any modeling study, it is challenging to incorporate all of the relevant physical and chemical processes in a self-consistent manner. It is worth briefly noting here several of the limitations of our modeling effort that may impact our ability to fully capture the physical state of WASP-76 b’s atmosphere.

Firstly, our GCM employs a double-gray radiative transfer scheme. While this scheme has benefits compared to more accurate, nongray approaches (e.g., speed, ability to diagnose underlying dynamics), the heating and cooling predicted by double-gray radiative transfer is inherently limited in spectral regions of strong wavelength dependence, given that these types of models use two opacities as opposed to precise, wavelength-dependent opacities ([Showman et al. 2009](#); [Kataria et al. 2015](#); [Lee et al. 2021](#)). For example, such a modeling approach is limited in its ability to generate the strong thermal inversions predicted by 1D and 3D modeling of ultra-hot Jupiter atmospheres (e.g., [Parmentier et al. 2018](#); [Lothringer & Barman 2019](#)). Additionally, while our post-processing scheme includes treatments of condensation and cloud formation, our GCM is cloud-free. This introduces a fundamental inconsistency between the two components of our modeling effort and misses the potential for radiative feedback from clouds, which can significantly impact the thermal structures of hot Jupiter atmospheres (e.g. [Lee et al. 2016](#); [Lines et al. 2018, 2019](#); [Roman & Rauscher 2019](#); [Roman et al. 2021](#)).

Secondly, as noted in our companion work ([May & Komacek et al. 2021](#)), one explanation for the inability of our models to fully reproduce the [Ehrenreich et al. \(2020\)](#) data in a self-consistent manner could be our GCM’s approximate treatment of drag. If the regions probed by high-resolution spectroscopy contain atmospheric flow that is sufficiently coupled to the magnetic field (i.e., there is enough ionization for plasma effects to become notable; [Perna et al. 2010](#)), collisions between electron, ion, and neutral populations can produce drag effects that depart from our Rayleigh approximation. In particular, the strength of the drag should depend strongly on the local temperature (via the amount of thermal ionization), the direction of the flow relative the magnetic field ([Rauscher & Menou 2013](#)), and the geometry of the global magnetic field ([Batygin & Stanley 2014](#)), and in some cases these interactions can actually serve to accelerate the flow of neutrals

(Koskinen et al. 2014; Rogers 2017; Hindle et al. 2021). Recently, Beltz et al. (2021) found that accounting for locally calculated magnetic drag in a double-gray GCM of WASP-76 b significantly altered the dayside atmospheric flow, driving flow toward the poles at low pressures and strongly reducing the extent of the equatorial jet. The remaining question — whether magnetic acceleration could account for the gap between our models and the Ehrenreich et al. (2020) data — is an excellent avenue for future work, as ultra-hot Jupiters have strong atmospheric thermal ionization and hence high electrical conductivity, and if they have significant dipolar magnetic fields, they will almost certainly experience strong Lorentz forces.

Finally, the accuracy of our models at the spectrum level is bounded by our treatment of chemistry: Both the GCM and the chemical equilibrium models computed with `FastChem` assume solar metallicity and the absence of any disequilibrium processes. These assumptions may impact the thermal structure of our GCM; enhanced (i.e., greater than solar) metallicity in GCMs has been shown to promote greater day-night temperature contrasts (e.g., Showman et al. 2009; Kataria et al. 2015) — which could in turn drive day-night flow and affect our blueshift conclusions or alter chemistry (Steinrueck et al. 2019; Drummond et al. 2020). Furthermore, our condensation treatment does not account for a “cold trap” scenario (e.g., Showman et al. 2009; Spiegel et al. 2009). Specifically, our temperature-pressure profiles (Figure 2.2) intersect some condensation curves at multiple pressures; our approach to modeling condensation can in some cases remove gas-phase abundance at a deep pressure but allow it at a higher pressure. This is not physically plausible, as a cold trap would make the lower-pressure regions thermodynamically inaccessible to a gas-phase species if higher-pressure regions condense it out.

From an observational perspective, Fu et al. (2021) indicate that the emission spectrum of WASP-76 b at low resolution is consistent with solar metallicity and equilibrium chemistry. This

result motivates our not accounting for disequilibrium, although the upper atmosphere on WASP-76 b’s dayside could experience significant photochemistry by virtue of its intense stellar insolation (e.g., [Line et al. 2010](#); [Moses et al. 2011](#); [Lavvas et al. 2014](#); [Molaverdikhani et al. 2019](#); [Shulyak et al. 2020](#)). Thus, while our modeling approach does not explore the entire relevant chemical parameter space, it does not run contrary to previous conclusions about the state of WASP-76 b’s atmosphere.

2.6 Conclusions

To examine the hypothesis of iron rain on WASP-76 b and delve into the planet’s 3D thermal, chemical, and wind structures, we post-process a grid of GCMs with a ray-striking radiative transfer code at high spectral resolution. We further apply a variety of post-processing schemes to account for condensation and/or rainout of iron and other species. Our major findings are as follows:

1. Reproducing the large blueshifts observed by [Ehrenreich et al. \(2020\)](#), [Kesseli & Snellen \(2021\)](#), and [Tabernero et al. \(2021\)](#) in the context of our GCMs almost certainly requires that WASP-76 b’s atmosphere is in a weak-drag state and has a deep RCB (i.e., a low internal temperature). These two conditions are necessary to drive the strong wind speeds and significant east-west limb asymmetries that are required to produce large blueshifts ($\gtrsim 5 \text{ km s}^{-1}$) during transit.
2. With our cloud-free models, we are unable to reproduce the magnitude and time dependence of the anomalous iron blueshift signature reported in [Ehrenreich et al. \(2020\)](#) in a way that is consistent with the thermal and wind structure predicted by our GCM output. This con-

clusion is robust even when accounting for gas-phase iron condensation. We are therefore unable to reproduce the behavior of the [Ehrenreich et al. \(2020\)](#) toy model when applying our self-consistent GCM-based modeling approach.

3. The maximum (egress) magnitude of the observed blueshift *can* be reproduced if we allow for a small and previously unaccounted-for eccentricity of $e \approx 0.017$, which is within the bounds allowed by previous observational constraints ([Fu et al. 2021](#)).
4. Including an optically thick cloud up to 10 scale heights thick in our post-processing scheme allows us to reproduce the sharp jump in Doppler shift prior to mid-transit that is evident in the [Ehrenreich et al. \(2020\)](#) and [Kesseli & Snellen \(2021\)](#) data. While the cloud composition does not strongly affect the Doppler signature, (gray) Al_2O_3 clouds provide the best fit, with Fe clouds and Mg_2SiO_4 clouds fitting marginally worse. Similarly, while including gas-phase condensation along with thick clouds provides the best fit, the fit improvement by doing so is marginal. Combining these clouds with a small eccentricity accurately reproduces both the magnitude and the time-dependent trend of the [Ehrenreich et al. \(2020\)](#) Doppler shift observations. Notably, the cloud does not actually reach 10 scale heights at most locations in the atmosphere, as the atmospheric temperature-pressure structure can become too hot for our modeled clouds at high altitudes.
5. Our model spectra line heights and Doppler shifts are broadly consistent with the [Tabernero et al. \(2021\)](#) analysis of the WASP-76 b ESPRESSO data for a variety of chemical species. We find that individual line Doppler shifts can differentiate between drag timescale scenarios, whereas the altitudes at which these lines form cannot. Phase-resolved, cross-correlation-derived Doppler shifts are uniquely suited for differentiating between drag timescale scenar-

ios and condensation/cloud scenarios, as they probe many lines (and hence altitudes) over multiple geometries.

6. Ca II is distinct among the studied chemical species. Its line heights as observed by [Tabernero et al. \(2021\)](#) are likely explained by non-hydrostatic effects (i.e., atmospheric escape). Its modeled Doppler shift is distinctly redder than all other investigated species, and its phase-resolved Doppler shift seems to be strongly dependent on the assumed planetary temperature structure.

Overall, we note broad qualitative agreement between our GCM outputs (specifically those with weak drag and deep RCBs) and the existing high-resolution data for WASP-76 b, as presented by [Ehrenreich et al. \(2020\)](#), [Kesseli & Snellen \(2021\)](#), and [Tabernero et al. \(2021\)](#). However, quantitatively reproducing both the observed magnitude and time dependence of the Doppler shift data is more elusive. Specifically, our models call into question that the large blueshifts of iron can be simply explained by condensation of this species out of the gas phase. Our work instead indicates that WASP-76 b’s anomalous blueshift during transit cannot solely be interpreted as iron condensation, insofar as the thermal and wind structures output by our GCM accurately predict the physical state of the planet. Instead, we have invoked a combination of condensation, optically thick clouds, and a small but non-zero eccentricity to best explain the [Ehrenreich et al. \(2020\)](#) data in particular. Other aspects of the WASP-76 b high-resolution data (i.e., the results presented by [Tabernero et al. 2021](#)) are well-explained by a larger subset of our forward models. We point out that our challenges in reproducing the observed Doppler shift signatures from [Ehrenreich et al. \(2020\)](#) do not call into question the fidelity of the data themselves, which have been independently verified across two instruments, three studies, and six nights of observation.

WASP-76 b remains a well-studied and interesting ultra-hot Jupiter, and we have shown in this work that the high-resolution data are highly diagnostic of 3D processes in the planet's atmosphere. Ultimately, additional data for this planet and for other ultra-hot Jupiters with similar properties, along with improved 3D and 1D modeling approaches, will shed light on the relative importance of clouds, condensation, and atmospheric circulation in shaping the properties of this intriguing class of exoplanet.

Chapter 3: Diagnosing Limb Asymmetries in Hot and Ultrahot Jupiters with High-resolution Transmission Spectroscopy

3.1 Introduction

Recent years have seen the intrinsic 3-dimensionality of these objects be uniquely constrained with transit HRCCS results. Observational studies such as [Louden & Wheatley \(2015\)](#), [Ehrenreich et al. \(2020\)](#), and [Kesseli et al. \(2022\)](#) have isolated signals from the morning and evening limbs of planetary atmospheres, unveiling Doppler shifts of multiple chemical species at multiple points in transit—and hence over multiple longitudinal slices. Such studies have revealed asymmetries in the probed Doppler velocity field (i.e., changes in the Doppler shift of the CCF maximum as a function of orbital phase), which are often attributed to physical asymmetries in the atmosphere.

However, as reviewed in Section [3.2](#), an asymmetric signal in HRCCS can arise from a combination of different classes of mechanisms: 1) chemistry, 2) clouds, 3) dynamics, 4) orbital properties, and 5) thermal structure. Disentangling these effects is not a straightforward process. This may especially be the case if transmission spectra must be stacked together to achieve a higher signal-to-noise ratio (SNR), thereby smearing phase information.

In this chapter, we aim to explore the general question of asymmetry in exoplanet atmospheres, with particular focus on its manifestations in high-resolution transmission spectroscopy. Section [3.2](#) examines what drives asymmetry in exoplanet atmospheres; here we define a metric that quantifies limb-to-limb asymmetry. In Section [3.3](#), we elaborate on diagnostics of specific mechanisms that may drive such asymmetries. This section additionally emphasizes how these diagnostics may be used to support or falsify compelling “toy models” motivated by the drivers

described in Section 3.2. Finally, we summarize our results in Section 3.4.

3.2 Selected drivers

There exist a number of potential drivers of asymmetry in high-resolution transmission spectroscopy. But what are the relative strengths of these drivers?

Previous works have considered the effects of condensation, longitude-dependent winds, and orbital eccentricity in producing such asymmetries (Wardenier et al. 2021; Savel et al. 2022). Table 3.1 includes these and a number of other potential drivers of asymmetry (along with potential diagnostics; Section 3.3). While many drivers are listed in Table 3.1, we consider in this work the relative strengths of two potentially first-order effects: the “scale height effect” and differences in equilibrium chemistry abundance across the limbs of the planet. Being both temperature-dependent effects, the distinction between the two is particularly subtle from an observational perspective, and hence interesting from a theoretical perspective.

The scale height effect is due to the larger scale height in hotter regions (e.g., Miller-Ricci et al. 2008), such that they are “puffed up” and cover more solid angle on the sky. These hotter regions therefore contribute more to the observed net Doppler signal in HRCCS. The scale height effect is seen in Kempton & Rauscher (2012) as a slight, increasing blueshift over transit and as slight ingress/egress differences. The effect there is not as dramatic as in planets with larger east–west limb asymmetry, such as WASP-76 b (West et al. 2016; Wardenier et al. 2021; Savel et al. 2022).

With respect to equilibrium chemistry: because of the strong day–night contrasts in (ultra)hot Jupiter atmospheres, there exist strong spatial variations in temperature. The day–night contrasts

result in east–west contrasts because the equatorial jet advects hot gas ahead of the substellar point to the evening limb and relatively cold gas from the antistellar point to the morning limb. Furthermore, as the planet rotates on its spin axis during transit, the hotter side of the planet progressively rotates into view, exacerbating these differences at egress. Ignoring all disequilibrium processes and scale height differences, there should therefore exist strong spatial variations in gas-phase atmospheric composition; at a given bulk composition, equilibrium chemistry implies variations in chemistry solely as a function of temperature and pressure. It is expected that asymmetries in transmission could hence vary as a function of temperature due to differences in chemistry alone.

Chemical gradients are invoked to explain a number of observational datasets (e.g., [Ehrenreich et al. 2020](#); [Kesseli & Snellen 2021](#)). However, other temperature-dependent effects, such as the scale height effect, may instead be driving observed asymmetries. With this distinction in mind, it is prudent to consider the difference in strength between these two effects and whether one considerably outweighs the other.

3.2.1 Asymmetry metric

To quantify the asymmetry of chemical abundance in a planetary atmosphere, we construct a west–east asymmetry metric, A_{WE} :

$$A_{WE} = \frac{1}{C} \int_{\text{west}} \log_{10} \left(\int n_{\alpha}(T, P) dl \right) d\Omega - \frac{1}{C} \int_{\text{east}} \log_{10} \left(\int n_{\alpha}(T, P) dl \right) d\Omega, \quad (3.1)$$

where, for species α , n is the number density in an atmospheric cell, $d\Omega$ is the solid angle subtended

by a given sky-projected radius–latitude cell, and there are C total cells per limb. By equilibrium chemistry, n is a function solely of temperature T and pressure P within a given cell in the modeled 3D atmosphere. For each 2D sky-projected radius–latitude cell, dl is integrated along the line of sight through the planet’s 3D modeled atmosphere. This metric takes into account regions of the planet outside the terminator (which impacts transmission spectra even at low resolution; e.g., [Caldas et al. 2019](#), [Wardenier et al. 2022](#)) by ray-striking through a 3D atmosphere.

A_{WE} essentially reduces to the difference in mean (log) abundance between the two limbs. The sign of this quantity encodes information about the asymmetry, as well: positive A_{WE} implies that the western limb is more abundant in a species, whereas negative A_{WE} implies that the eastern limb is more abundant in a species.

3.2.2 Model atmospheres

As of yet, we have remained agnostic to the model that generates the temperature–pressure structure and defines the grid cells for an A_{WE} calculation. Some of the most complex and physics-rich descriptions of 3D exoplanet temperature–pressure structures are given by general circulation models (GCMs; e.g., [Showman et al. 2009](#)). In this study, however, we seek to gain intuition for the basic scaling of asymmetry with planetary temperature (which drives the scale height and equilibrium chemistry gradients), and the added physical complexity of GCMs could add “noise” to this “signal”—it would be difficult to isolate the effect of increasing planetary temperature alone. Furthermore, we here consider unphysical situations in order to determine the magnitude of the resulting difference with the correct physics. Finally, GCMs are very computationally expensive to run and have a number of free parameters to tune, and we here aim to explore a nontrivial grid

of models over a representative range of parameter space.

We opt for a simple, parameterized approach instead of pursuing a full GCM description of our atmospheres for this specific experiment. Our model atmospheres have two parameters: a normalized east–west contrast $\tilde{\Delta}T = (T_{\text{east}} - T_{\text{west}})/T_{\text{eq}}$ and an equilibrium temperature T_{eq} . A normalized east–west contrast is a natural choice over an absolute east–west contrast for this work; namely, it prevents negative temperatures at low T_{eq} , and it has physical meaning motivated by dynamical theory (e.g., [Tan & Komacek 2019](#)). In these models, the choice of $\tilde{\Delta}T$ also uniquely enforces the east-west temperature differences. The limb-to-limb difference cannot exceed the day–night difference; based on the GCMs of [Tan & Komacek \(2019\)](#) and a set of phase curve observations ([Parmentier & Crossfield 2018](#)), we do not expect a day–night contrast to exceed 0.6, so we hold our east–west contrast below this value.

Hence, we here sweep our parameterized atmospheric models in $\tilde{\Delta}T$ from 0.1–0.6, in addition to sweeping in T_{eq} from 1000 K – 4000 K. Each atmosphere is characterized by two isothermal temperature–pressure profiles. Defining

$$\begin{aligned} T_{\text{east}} &= T_{\text{eq}} + \Delta T/2 \\ T_{\text{west}} &= T_{\text{eq}} - \Delta T/2 \end{aligned} \tag{3.2}$$

and noting that $\Delta T = T_{\text{east}} - T_{\text{west}}$, it therefore follows that

$$\begin{aligned}
T_{\text{east}} &= \frac{T_{\text{eq}}}{1 - \tilde{\Delta}T/2} \\
T_{\text{west}} &= \frac{T_{\text{eq}}}{1 + \tilde{\Delta}T/2}.
\end{aligned}
\tag{3.3}$$

With the substellar longitude at 0° , all cells with a longitude $\phi < 180^\circ$ —the warmer evening limb—are assigned temperature T_{east} . Conversely, all cells with a longitude $\phi > 180^\circ$ —the cooler morning limb—are assigned temperature T_{west} . Pressures in the atmosphere run as low as $1 \mu\text{bar}$, as one of the benefits of HRCCS is that it can probe low pressures such as these (e.g., [Kempton et al. 2014](#); [Gandhi et al. 2020](#); [Hood et al. 2020](#)). The bottom of the atmosphere is set at 0.5 bar; our previous 3D forward models run in [Savel et al. \(2022\)](#) across the optical and near-infrared indicate that for our test case of WASP-76 b ([West et al. 2016](#)), this region is the deepest that can be probed given the expected continuum opacity. The parameterized modeled atmospheres in this study have no set wind fields, as in our models (motivated by and assuming chemical equilibrium), winds do not control A_{WE} —only the chemical abundance of a given cell does.

We calculate A_{WE} to assess the relative strength of the scale height and equilibrium chemistry effects. To infer the strength of the scale height effect, we construct pairs of model atmospheres. In each pair, one atmosphere is constructed self-consistently: pressure falls off per hydrostatic equilibrium, with the scale height set by the temperature on either limb. For the models here, we hold composition constant across both limbs, thereby holding μ constant at 2.36 (appropriate for a solar-composition gas dominated by molecular H_2 ; [Kempton & Rauscher 2012](#)). See Section 3.2.4 for a discussion of this caveat. The other atmosphere in the pair is constructed on the same pressure grid as the western limb at all longitudes. That is, the eastern limb is not simulated




Figure 3.1: Asymmetry (as defined in Equation 3.1) of all chemical species considered in this study in our parameterized atmospheres at an equilibrium temperature of 2200 K. These models do not self-consistently inflate the hotter limb of the parameterized model (i.e., they do not observe the “scale height effect”). The shading of each species represents the normalized temperature difference, $\tilde{\Delta}T$, across the two limbs of our parameterized atmospheres; the lightest boxes have $\tilde{\Delta}T = 0.1$, whereas the darkest have $\tilde{\Delta}T = 0.6$. For illustrative purposes, we color in green tick marks for species with detections noted in Guillot et al. (2022) (and including the recent CO_2 detection; Ahrer et al. 2022). We also draw a vertical line denoting 0 asymmetry. Without taking the scale height effect into account, positive ions form much more predominantly on the warmer limb (i.e., have negative asymmetry) than other species and reach the greatest asymmetry values.




Figure 3.3: Asymmetry (per Equation 3.1) for Sr II, Al, H₂O, and CO in our parameterized atmospheres. Our grid sweeps over equilibrium equilibrium temperature and normalized temperature difference across limbs, and includes models that observe the scale height effect (circles) and do not (squares). We find that species with strong temperature-dependent abundances (e.g., Sr II) are less dominated by the scale height effect than species with weaker temperature-dependent abundances.

as inflated compared to the western limb—removing the scale height effect from the projected model atmosphere in transmission.

3.2.3 Equilibrium chemistry

To calculate the number densities of our species in each modeled atmospheric cell (n_α), we construct a grid in temperature–pressure space using the `FastChem` equilibrium chemistry code (Stock et al. 2018) and interpolate the grid based on local atmospheric cell temperature and pressure. We initialize the code with solar abundances from Lodders (2003). Our chemistry code does not explicitly include any condensation or cloud-formation processes.

Even disregarding questions of species detectability in HRCCS data, it is worth considering that not all species with `FastChem` thermochemical data have freely available opacity data. With this constraint in mind, we restrict our A_{WE} molecule calculations to molecules with opacity data available on ExoMol,¹ a popular opacity database for exoplanet atmosphere modeling.

3.2.4 Asymmetry metric: application

We calculate A_{WE} for our grid of parameterized atmospheres. Disregarding the scale height effect, we find that positive ions tend to form preferentially on the hotter limb of our models at an equilibrium temperature of 2200 K (Figure 3.1). This is expected, as thermal ionization should increase the abundance of positive ions at higher temperatures. Furthermore, larger east–west temperature asymmetries lead to larger abundance asymmetries.

Including the scale height effect increases the asymmetry for neutral atoms and molecules, as can be seen by comparing the right-hand sides of Figures 3.1–3.2. Furthermore, there is more

¹<https://www.exomol.com/>

homogeneity across the A_{WE} values across positive ions, negative ions, neutral atoms, and neutral molecules (Figure 3.2). In particular, while higher $\tilde{\Delta}T$ still implies higher absolute asymmetry in neutral species, the scale height effect makes it such that the warmer limb almost always has higher projected asymmetry.

It is therefore clear that the scale height effect strongly tamps down genuine variation in species abundance due to equilibrium chemistry. However, the fact that inter-species variation in asymmetry remains implies that this variation in abundance is not completely washed out by the scale height effect; if the scale height effect truly and fully dominated, all species would have the same A_{WE} value.

When considering individual species more closely, we find that certain species are particularly differentially affected by the scale height effect. For example, Figure 3.3 shows that there is a stark difference in whether the scale height effect is included for Fe. However, this is not as much the case for, e.g., Sr II. The meaning behind this result is evident in the equilibrium abundance calculations of Fe and Sr II: Fe is less sensitive to temperature variations than Sr II. This result is expected, as the onset of Sr II is determined by the temperature at which Sr I can be effectively ionized. This is generally the case for positive ions—the temperature effect on chemical abundance wins out over the scale height effect, as seen by the left-hand sides of Figures 3.1–3.2. Physically, this behavior is because the Saha equation is more strongly dependent on temperature than most chemical equilibrium reaction rates.

The results of this experiment indicate that the most temperature-sensitive species are strongly influenced by both abundance changes and scale height differences. Conversely, to isolate the scale height effect, it would be therefore useful to consider a species with very weakly temperature-dependent abundance; in this case, if a strong asymmetry were detected, it could be attributed to a

scale height effect (or other non-equilibrium chemistry or physics). We explore this idea further in Section 3.3.

Note that this approach, aside from its simplified temperature–pressure structure, does not account for a variety of physics. Namely, it does not include the effects of hydrogen dissociation and recombination that occurs in the ultra-hot Jupiter regime (Tan & Komacek 2019). Inclusion of this physics would serve to decrease the mean molecular weight in the atmosphere, increasing the scale height for the hotter, eastern limb, thereby amplifying the observed asymmetry. Additionally, at the lower-temperature end, we did not include the effects of certain species being sequestered into clouds (e.g., silicate clouds). We will model the Doppler shift impact of optically thick clouds in Section 3.3.1.2. Finally, our approach does not include disequilibrium effects (e.g., vertical / horizontal mixing) that may alter asymmetries. Therefore, the results shown here motivate asymmetries due to equilibrium chemistry alone, which we expect to be a first-order driver of asymmetry; disequilibrium chemistry is not expected to be significant in the ultrahot Jupiter regime (e.g., Tsai et al. 2021).

We further did not include the effect of temperature- and pressure-dependent opacities. At the spectrum level, a temperature asymmetry would be exaggerated by the fact that, e.g., Fe absorbs more on the hotter limb than the colder limb because its opacity increases with temperature. This would mean that the detected net Doppler shift is even more strongly weighted to the hotter limb.

Despite these limitations in our modeling, the trends listed above should hold to first order and provide intuition about the relative strengths of two potential drivers of asymmetry in exoplanet atmospheres. Broadly, it holds that the scale height effect appears to dominate in general, but relative differences in abundances of species as a function of temperature still matter. Given the limitations of simple models, we will move on to more self-consistent atmospheric modeling in

the following sections.

3.3 Selected diagnostics

3.3.1 Diagnostics for specific mechanisms

Per Section 3.2, even differentiating between two drivers of asymmetry in exoplanet atmospheres is nontrivial. Drivers can compete to varying degrees to produce a similar result: an asymmetric trend in net Doppler shifts in HRCCS.

However, by exploiting nuances in the HRCCS Doppler shift signal and by independent means, it may be possible to disentangle even drivers that produce similar effects. Table 3.1 lists example drivers of asymmetries in HRCCS and how they might be diagnosed. The associated works listed in the table may not directly propose these diagnostics, but at minimum they provide foundational material for them.

Of course, exhibiting a single diagnostic does not mean that a given physical mechanism is in play. Other mechanisms could surely be present, and uniquely constraining a single mechanism as dominant would require ruling out the others, as well. For instance, both day–night winds and morning limb condensation could result in a net blueshifted CCF. But if, for example, a nightside temperature were derived from a phase curve that was far too hot for any known condensate to form, then day–night winds would be much preferred to condensation as a physical solution. Together, collections of diagnostics are hence able to test the dominance of individual mechanisms.

In the following sections, we explore a few tests for specific physical mechanisms of asymmetry: using CO as a baseline molecule to identify the scale height effect and tracking the blueshifts

of multiple species to identify the presence of clouds. We furthermore evaluate the effectiveness of diagnostics that may be used to evaluate a number of different mechanisms: averaging HRCCS data into two phase bins and using finely phase-resolved HRCCS data. We additionally show how these diagnostics can further motivate or rule out “toy models” that at first may appear convincing.

3.3.1.1 CO as a baseline molecule

We have demonstrated (Section 3.2.4) that species with strongly temperature-dependent abundances are the least susceptible to the scale height effect. Conversely, observing a species with very *weak* temperature-dependent abundance could indicate whether the scale height effect is in play.

Consider CO. In Figure 3.3, its A_{WE} values are clustered around 0 without the scale height effect, with relatively weak dependence on $\tilde{\Delta}T$. However, CO’s A_{WE} values are strongly negative when the scale height effect is included. We propose using CO as a tracer of the scale height (and other chemistry-unrelated) effects.

As shown in Figure 3.4, the abundance of CO is relatively stable between 1000 K and 3500 K. [Beltz et al. \(2022\)](#) note that this stability holds over the temperature–pressure range of the observable atmosphere of the ultra-hot Jupiter WASP-76 b. Indeed, this feature remains true over the general temperature–pressure range of ultra-hot Jupiters. For illustrative purposes, we calculate the 1D temperature–pressure profiles of a hot Jupiter (WASP-39b; [Faedi et al. 2011](#)) and an ultra-hot Jupiter (WASP-18b; [Hellier et al. 2009](#)). These profiles, calculated with the HELIOS 1D radiative-convective model (with full heat redistribution), indicate that the observable atmosphere for these planets is largely within a region of near-constant CO mixing ratio. The stability of CO




Figure 3.4: Volume mixing ratio of CO as a function of pressure and temperature as calculated by `FastChem`. Overplotted are the onset of ultra-hot Jupiters (as defined by their dayside temperature; [Parmentier et al. 2018](#)), the CO/CH₄ equivalency curve from [Visscher \(2012\)](#) as a function of pressure, the Fe condensation curve from [Mbarek & Kempton \(2016\)](#), and 1D temperature–pressure profiles for a hot Jupiter (WASP-39b) and an ultra-hot Jupiter (WASP-18b) as computed with `HELIOS` ([Malik et al. 2017](#)). Both the condensation curve and the equivalency curve are computed at solar metallicity. Considering the regime of ultra-hot Jupiter atmospheres, CO is a relatively stable chemical species.

is due to three factors: its strong chemical bonding, its lack of participation in gas-phase chemical reactions, and its lack of condensation.

Since the strong triple bond of CO makes it difficult to thermally dissociate, CO remains stable at temperatures that would dissociate molecules with weaker bonds, such as H₂O (Parmentier et al. 2018), which has two single bonds. Beyond roughly 3500 K, even the triple bond becomes susceptible to thermal dissociation; hence, the few exoplanets with significant portions of their atmosphere hotter than this temperature (e.g., KELT-9b, with $T_{\text{eq}} \approx 4050$ K; Gaudi et al. 2017) would likely exhibit spatial variation in CO abundance. Most ultra-hot Jupiters, though, should fall shy of this regime. Furthermore, the high photoionization threshold of CO (relative to, e.g., H₂O; Heays et al. 2017) means that it is not commonly photodissociated (Van Dishoeck & Black 1988). Even when it is photodissociated, recycling pathways exist in hot Jupiters that can replenish CO abundance, keeping it near equilibrium abundance even inclusive of photochemistry (Moses et al. 2011). Hence, the assumption of non-dissociation of CO is reasonably justified across much of the ultra-hot Jupiter population.

Additionally, CO does not commonly participate in thermochemical reactions and is the dominant carbon carrier in our temperature–pressure range of interest. While at lower temperatures the dominant carbon carrier becomes CH₄, the ultra-hot Jupiter regime is squarely beyond the CO/CH₄ equivalency curve (Figure 3.4; Visscher 2012). Therefore, even aside from thermal dissociation, CO should not participate in gas-phase thermochemistry that would alter its abundance.

Finally, CO does not form any high-temperature condensates expected in ultra-hot Jupiter atmospheres. The condensation temperature of CO (≈ 80 K at 1 bar; Lide 2006; Fray & Schmitt 2009) is far below the temperature–pressure range of ultra-hot Jupiters. This quality makes CO a

less complicated tracer of, e.g., atmospheric dynamics than species that do condense in this region of parameter space, such as Fe, Mg, or Mn (Mbarek & Kempton 2016). Therefore, while the calculations of Figure 3.4 do not include gas-phase condensation, the resultant spatial constancy of CO should still be robust even when condensation is considered. CO is thus a more straightforward molecule to model than other, condensing species, as it does not participate in the complex microphysics of condensation and cloud formation (see, e.g., Gao et al. 2021).

Beyond its spatial uniformity, there are further observational reasons that CO is an appealing species to target. Namely, CO has very strong spectroscopic bands placed across the infrared wavelength range (e.g., Li et al. 2015) that do not overlap with other strong absorbers and are relatively well understood (Li et al. 2015). Additionally, the high cosmic abundance of C and O (Lodders 2003) means that, unlike many of the species in the previous section, CO is readily detectable (and has become a standard detection in HRCCS; Snellen et al. 2010; de Kok et al. 2013; Rodler et al. 2013; Brogi et al. 2014, 2016; Flowers et al. 2019; Giacobbe et al. 2021; Line et al. 2021; Pelletier et al. 2021; Zhang et al. 2021; Guilluy et al. 2022; van Sluijs et al. 2023a).

Given its stability and observational advantage, we propose that CO can be used as a faithful tracer of the atmosphere—whether it is inflated in some regions, what its wind profile is, whether regions are blocked by clouds, etc. In turn, CO may then be leveraged to better motivate sources of asymmetry that affect other species. While other species with low A_{WE} in Figure 3.1 (e.g., He, Fe, MgH, Rb II) would also appear to be good candidates for baseline species, these species are either largely spectroscopically inactive, have variable abundance over broader temperature–pressure ranges, or can condense. A caveat to the above is that while CO is a faithful longitudinal tracer, it is not an unbiased *radial* tracer (as seen in Figure 3.4). As with all chemical species, CO has its own balance between deep and strong lines that depends on the waveband considered (see,

e.g., Section 3.3.3.1). Therefore, the net CO Doppler signal does not uniformly weight the wind profile across all altitudes. Again, this is a bias inherent to all chemical species.

3.3.1.2 A decreasing blueshift test for clouds

As noted in Table 3.1, clouds may introduce strong asymmetry into HRCCS data. Savel et al. (2022) demonstrated that gray, optically thick clouds produce stronger blueshifts in the Doppler shift signal of WASP-76 b than the blueshifts in clear models, also changing the trend of Doppler shift with phase. But, again as shown in Table 3.1, these changes at the Doppler shift level are not sufficient to uniquely identify clouds as the driver of an observed asymmetry. Combinations of observable quantities that would uniquely identify clouds as the source of observed HRCCS asymmetry are therefore necessary.

To devise such a test, we investigate in this work a limiting-case cloudy model. As in Savel et al. (2022), we construct gray, optically thick, post-processed clouds in our 3D ray-striking code. We here make another assumption, though: that the clouds are confined to the cooler, morning limb, as opposed to having a distribution dictated by a specific species' condensation curve. This distribution is based on planetary longitude (between longitudes of 180° and 360°). This approach is motivated by the results of Roman et al. (2021), who found that a subset of cloudy GCMs exhibited a cloud distribution strongly favoring the western limb.² Our approach benefits from providing limiting-case intuition for how cloudiness affects Doppler shift signals while avoiding the complex questions of how clouds form and which species contribute the most opacity (Gao & Powell 2021; Gao et al. 2021).

²These GCMs produced clouds on a temperature–pressure basis, and did not model clouds as tracers. Therefore, they do not capture potential disequilibrium cloud transport (e.g., as done in Komacek et al. 2022), which may alter the degree of patchiness within the cloud deck.

Briefly, our modeling methodology is as follows:

1. Double-gray, two-stream GCM. GCMs such as this one solve the primitive equations of meteorology, which are a reduced form of the Navier-Stokes equations solved on a spherical, rotating sphere with a set of simplifying assumptions.³ The output of these models is temperature, pressure, and wind velocity as a function of latitude, longitude, and altitude. We use the GCM that was shown to best fit the [Ehrenreich et al. \(2020\)](#) WASP-76 b data in [Savel et al. \(2022\)](#).
2. Equilibrium chemistry with `FastChem`. As in Section 3.2.3, we interpolate a model grid of chemistry to determine local abundances of a number of chemical species as determined by temperature and pressure conditions of the GCM output.
3. Ray-striking radiative transfer. Using a code modified from [Kempton & Rauscher \(2012\)](#) (as detailed in [Savel et al. 2022](#)), we compute the high-resolution absorption spectrum of our planetary atmosphere by calculating the net absorption of stellar light along lines of sight through our GCM output. This absorption is calculated inclusive of net motions along the lines of sight from atmospheric winds and rotation, inducing Doppler shifts relative to that of a static atmosphere’s spectrum. Limb-darkening is calculated with a quadratic limb-darkening law in the observable planetary atmosphere and with the `batman` code ([Kreidberg 2015](#)) for the portion of the star blocked by the optically thick planetary interior.

Given its increasing utility as a benchmark planet for HRCCS studies (e.g., [Ehrenreich et al. 2020](#); [Kesseli & Snellen 2021](#); [Landman et al. 2021](#); [Seidel et al. 2021a](#); [Wardenier et al. 2021](#);

³These assumptions are 1) local hydrostatic equilibrium, such that vertical motions are caused purely by the convergence and divergence of horizontal flow, 2) the “traditional approximation,” which removes the vertical coordinate from the Coriolis effect, and 3) a thin atmosphere.




Figure 3.5: Atmospheric Doppler shifts, which should remain in the HRCCS signal after the orbital motion is subtracted, as a function of orbital phase for our forward models. Shown are representative species that span Doppler shifts and are noted as potentially observable by [Kesseli et al. \(2022\)](#): Fe, Sr II, and Sc. Cloud-free models are represented with solid lines, whereas models with fully cloudy morning limbs are represented with dashed lines. The first half of transit (RV_1) and second half of transit (RV_2) Doppler shifts for Fe from [Kesseli et al. \(2022\)](#) are overplotted as horizontal lines, with width determined by observational errors. Our cloudy models are much more strongly blueshifted than their cloud-free counterparts, become less blueshifted over transit, and do not have significant CCF peaks at early phases.

[Kesseli et al. 2022](#); [Sánchez-López et al. 2022](#)), we model the ultra-hot Jupiter WASP-76 b ([West et al. 2016](#)). We calculate 25 spectra inclusive of Doppler effects equally spaced in phase from the beginning to end of transit. For our cross-correlation template, $\mathcal{T}$, we use a model that does not include Doppler effects.

We then cross-correlate our template against our calculated spectrum, y :

$$c(v) = \sum_{i=0}^N y_i(\lambda) \mathcal{T}_i(v, \lambda), \quad (3.4)$$

where the mask or template is Doppler-shifted by velocity v and interpolated onto the wavelength grid, λ , of y for summing. Our CCF is computed on a grid of velocities from -250 km s^{-1} and 250 km s^{-1} with a step of 1 km s^{-1} . The final net planet-frame Doppler shift is calculated by fitting

a Gaussian to the peak of the CCF.

The results of our experiment are shown in Figure 3.5. When we allow clouds to extend over the entire morning limb, note that all species become less blueshifted over time. Because the limb that is rotating away from the observer (the “receding limb”) is entirely blocked off by clouds, there is no wavelength-dependent absorption for that limb. Therefore, the contribution of redshifting from solid-body rotation on the receding limb is not present—the only Doppler shift contributions are from evening limb rotation and evening limb winds, which are generally in the same direction as the rotation. Hence, there are much stronger blueshifts at earlier phases than in the clear models.

However, at later phases, the non-cloudy regions of the atmosphere rotate into the receding limb, thereby contributing some rotational redshift to the net Doppler shift signal.⁴ At the earliest phases, the cloudy models do not have enough wavelength-dependent absorption to produce a significant CCF peak. Notably, all species follow this trend, as the blocking of clouds as modeled here is wavelength-independent and altitude-independent. This behavior is shown in Figure 3.5 for Fe, Sc, and Sr II—all species identified in Kesseli et al. (2022) as having high potential observability for ultra-hot Jupiters.

From Figure 3.5, it is also apparent that the cloud-driven trend of decreasing blueshift in phase is not matched by the observations of Kesseli et al. (2022). As found in Savel et al. (2022) in comparison to the Ehrenreich et al. (2020) data, while the absolute magnitude of the cloudy model’s Doppler shift better match the data than the clear model’s, the cloudy model trend over Doppler shift is not matched by the data. In sum, this limiting-case model of opaque, morning

⁴The degree of rotation during transit varies as a function of semimajor axis and host star radius, and hence from planet to planet. While we only model WASP-76 b, other planets also have large (e.g., compared to the angles probed by transmission spectroscopy) rotations during transit (Wardenier et al. 2022).

limb clouds does not appear to be a first-order effect driving existing observational trends. This does not necessarily mean that clouds are not the driving factor behind limb asymmetries; it may simply be that a more physically motivated model for partial cloud coverage of the limb could fit the available data better.

Also of note in Figure 3.5 is that the egress signatures of the clear and cloudy models are quite distinct. Near a phase of roughly 14 degrees, the clear model produces a sharp change in Doppler shift for all species as the leading (rotationally redshifted) limb begins to leave the stellar disk. This sharply blueshifting behavior continues to the end of egress, until the last sliver of the trailing (rotationally blueshifted) limb has left the stellar disk as well. In the cloudy case, however, the leading limb leaving the stellar disk has no effect, as it is fully cloudy. While this effect is evident in these models, it may be less evident in observations, which naturally cannot finely sample ingress and egress phases.

3.3.2 Phase bins

We have thus far examined drivers of asymmetry and potential diagnostics of specific mechanisms. Next, we will evaluate a few HRCCS data types to determine how robust they are and their potential ability to constrain a number of different physical mechanisms that give rise to HRCCS asymmetry.

The first of these data types is HRCCS Doppler shifts that are binned in phase. A substantial fraction of HRCCS studies present detections and Doppler shifts integrated over the entirety of transit (e.g., [Giacobbe et al. 2021](#)). This approach maximizes detection SNR, which may be necessary for a given set of observations (e.g., because of a low-resolution spectrograph, small

telescope aperture, faint star, low species abundance, low number of absorption lines, or weak intrinsic absorption line strengths). While it is possible to reveal aspects of limb asymmetry with this approach, especially when comparing detections of multiple species to one another, phase-resolving the transit (and observing isolated ingresses and egresses when possible) will certainly give a more direct probe of east–west asymmetries. Binning HRCCS data in phase across transit may provide a desirable balance between revealing asymmetry and maintaining high SNR.

We seek to address this question by phase-binning modeled Doppler shifts to examine its biases with respect to the underlying model. We follow this experiment with a comparison to the phase-binned observations of [Kesseli et al. \(2022\)](#).

3.3.2.1 Theoretical phase binning

We average our phase-resolved calculations into two bins: the first and second half of transit. Once our CCFs are calculated we average them in phase to effectively reduce our data to two single bins: the first half of transit and the second half of transit. We make versions of these two half-transit bins that include or exclude the ingress and egress phases (when the planet is only partially occulting the star).

Motivated by recent detections in the near-infrared ([Landman et al. 2021](#); [Sánchez-López et al. 2022](#)), we search for absorption from various molecules⁵ in our models, focusing on the CARMENES ([Quirrenbach et al. 2014](#)) wavelength range and resolution for direct comparison against observational results using that instrument. Of these molecules, we find that OH, H₂O, and CO produce significant absorption over the modeled wavelength range, with OH and H₂O

⁵We use the MoLLIST ([Brooke et al. 2016](#)), POKAZATEL ([Polyansky et al. 2018](#)), and Li2015 ([Li et al. 2015](#)) line lists for OH, H₂O, and CO, respectively.




Figure 3.6: Single-species (OH, CO, H₂O, HCN) 3D forward-modeled spectra of WASP-76 b. These spectra are simulated over the CARMENES waveband and resolution. Doppler effects are not included in these spectra, which are modeled at center of transit. H₂O is the dominant absorber in this bandpass, followed by OH. HCN exhibits no spectral features above the continuum for WASP-76 b in this bandpass.

producing the strongest features (Figure 3.6). We find that HCN does not produce any noticeable absorption under the assumption of chemical equilibrium and solar composition, implying either more exotic chemistry for WASP-76 b’s atmosphere (i.e., photochemistry or non-solar abundances; Moses et al. 2012), or that the detection of HCN in this atmosphere (Sánchez-López et al. 2022) was spurious (perhaps due to the nature of the HCN opacity function; Zhang et al. 2020). We furthermore find a moderate ($\approx 4 \text{ km s}^{-1}$) increase in blueshift for our modeled H₂O. While this increase in blueshift is commensurate with the increase in blueshift described for H₂O in Sánchez-López et al. (2022), we are once again unable to match the high reported velocities (here -14.3 km s^{-1}) with our self-consistent forward models.

Figure 3.7 shows the results of this experiment. As the error for each phase bin, we take the average error of phase bins from Kesseli et al. (2022) (1.55 km s^{-1}). We define the two phase bins as inconsistent if the peak of their respective CCFs are inconsistent at 2σ .

We find that excluding ingress and egress phases can strongly reduce the difference in derived Doppler shift between phase bins. Furthermore, we find that, as expected from Kempton &




Figure 3.7: CCFs of individual species averaged over two phase bins. Each column corresponds to different species (OH, CO, H₂O), and each row corresponds to different bin selection: without including ingress and egress, including the full transit, and only including ingress and egress. Central bars between the CCFs are colored blue if the difference between the CCFs is greater than optimal Doppler shift errors (1.55 km/s, in black; [Kesseli et al. 2022](#)); otherwise, they are colored red. In our models, CO only displays detectable CCF differences when only including ingress and egress. The SNR in each plot refers to the difference between the two phase bins' CCF peaks relative to the optimal Doppler shift errors.

[Rauscher \(2012\)](#), differences between bins are maximized when just considering the ingress and egress phases.⁶

While higher-order drivers of asymmetry are clearly not detectable with phase bins (e.g., at what longitude condensation may begin to play a role; [Wardenier et al. 2021](#)), certain drivers of asymmetry are accessible with this method. For example, ignoring for now the exact details of error budgets, all species in Figure 3.7 clearly blueshift over the course of transit. This provides potential evidence for, among other things, a spatially varying wind field, condensation, optically thick clouds, or a scale height effect. Furthermore, per the results of Section 3.3.1.1 the detection of CO’s blueshifting indicates that something besides equilibrium chemistry is driving at least some of the asymmetry in the atmosphere. These underlying models are cloud-free, so these results imply sensitivity to, e.g., the scale height effect.

3.3.2.2 Comparison to [Kesseli et al. \(2022\)](#)

With our models calculated, we can now explore the ability of phase-resolved spectra to confront toy models by comparing the models to observations. A prime observational work that made use of phase binning is [Kesseli et al. \(2022\)](#); there, the authors search for asymmetries in two phase bins for a wide variety of species, motivated by the strength of those species’ opacity functions in the data’s wavelength range.

To consider a toy model: based on previous studies ([Ehrenreich et al. 2020](#); [Tabernero et al. 2021](#); [Savel et al. 2022](#)), it appears that Ca II does not follow the Fe-like Doppler shift trend first observed by [Ehrenreich et al. \(2020\)](#). Rather, it appears that Ca II, with its strong opacity and

⁶We would expect that binning with fewer spectra (just including the ingress phases) would increase the associated error on Doppler shift at each bin. However, the point of this exercise is to illustrate the magnitude of ingress/egress Doppler shift discrepancy; observational strategies such as stacking multiple transits could reduce errors in practice and make these differences discernible.




Figure 3.8: The net Doppler shifts of [Kesseli et al. \(2022\)](#) (error bars) as compared to this work’s models (crosses). The first phase bin is drawn thinner than the second phase bin; observed phase bins are connected by a dotted line for visibility’s sake. The species are ordered and colored by total observed detection SNR. Rows without crosses correspond to species that we could not recover via cross-correlation in our models. Our models are able to explain some species (e.g., Na), fail to explain others (e.g., Cr) and fail to detect yet others (e.g., K).

resultant deep lines, may be probing a non-hydrostatic region of the atmosphere (Casasayas-Barris et al. 2021; Deibert et al. 2021; Tabernero et al. 2021). This region of the atmosphere cannot be captured by the models of this work and Savel et al. (2022).

Without a model of atmospheric escape, it seems difficult to elevate the above picture beyond “toy model” status. However, by phase-resolving multiple species, a clearer picture can emerge.

For our comparison with Kesseli et al. (2022), we use the same line lists as in that study: the National Institute of Standards and Technology (NIST; Kramida et al. 2019) line lists. It is crucial to use the same line lists for comparisons of HRCCS studies—different line list databases can contain vastly discrepant numbers of line transition, which greatly affects the resultant opacity function (see, for instance, Figure 11 of Grimm et al. 2021).

The results of our comparison with the species detected in Kesseli et al. (2022) are shown in Figure 3.8. As in Savel et al. (2022), these baseline models—no clouds, no condensation, no orbital eccentricity—cannot fully explain the Doppler shifts of Fe observed in WASP-76 b. However, the comparison across multiple different species provides further constraints. Figure 3.8 shows that Fe, V, Cr, Ca II, and Sr II are strongly discrepant from our models for at least one half of transit, whereas Na, Mg, Mn, and Ni are reasonably well described by our models for both the first and second half of transit. Furthermore, Fe, V, and Cr all have stronger blueshifts in the second phase bin than in our models. The similar level of disagreement between Fe, V, and Cr implies that they share a common driver of asymmetry. This result in turn implies that whatever driver affects them affects the regions in which these species form similarly — be it clouds, condensation, etc.

To bridge the toy models presented in Section 3.2.3 to our Kesseli et al. (2022) comparison, we compute a set of high-resolution spectra exactly as above, but with the same altitude grid at all latitudes and longitudes in an effort to effectively turn off the scale height effect while maintaining




Figure 3.9: Modeled phase-resolved Doppler shifts for select NIR-absorbing species, with representative error bars (Ehrenreich et al. 2020) drawn on. We find that OH and H₂O have distinct Doppler signatures from CO; however, OH and H₂O have Doppler shifts that are indistinguishable from one another with current best-case error bars (e.g., Ehrenreich et al. 2020). Considering CO as a “baseline species” here allows one to better understand how H₂O and OH may change through the atmosphere.

chemical limb inhomogeneities. Post-processing this (self-inconsistent) model yields less than half the Doppler shift asymmetry as compared to our self-consistent models. This experiment confirms the intuition that the scale height effect is a first-order asymmetry effect.

Finally, we consider the Ca II toy model previously described. Certain lightweight and/or ionized species may be entrained in an outflow, as indicated by some previous observations (e.g., Tabernero et al. 2021) of very deep absorption lines in transmission that must extend very high up in altitude. The differential behavior of the Ca II and Sr II Doppler shifts lends more credence to this hypothesis.

In sum, by taking advantage of phase-binned spectra, it is possible to better identify drivers of HRCCS asymmetry. Additionally, our predictions in Figure 3.8 indicate that most species should

have roughly the same Doppler shift patterns. In stark contrast, observations reveal much larger variations in velocity across different species. While some interpretation may be due to spurious detections, physics that is not included in our model (e.g., outflows, condensation) may be playing a driving role.

3.3.3 Full phase-resolved spectra

Currently, the most information-rich diagnostic available to probe asymmetry in HRCCS is phase-resolved cross-correlation functions (e.g., [Ehrenreich et al. 2020](#); [Borsa et al. 2021](#))—that is, net Doppler shifts associated with the absorption spectrum evaluated over multiple points in transit. With these data, one should be able to directly constrain longitudinally dependent drivers of asymmetry, providing the best chance of disentangling the physical mechanisms outlined in Section [3.2](#). But how far can we push these data?

3.3.3.1 Example: probing physics in the NIR

To explore this question, we take as an example a three-species (OH, H₂O, and CO) near-infrared (NIR) dataset over a CARMENES-like waveband as in Section [3.3.2](#). Figure [3.9](#) shows the Doppler shifts of these species as a function of phase, produced for single species at a time as in Section [3.3.2](#), but without any averaging.

Without considering any data, a compelling toy model would be as follows: H₂O is thermally dissociated on the hotter, approaching limb, so it preferentially exists on the receding limb. OH is a product of H₂O photodissociation, so it forms preferentially on the approaching limb. CO is constant everywhere; therefore, CO should not experience much of a trend in Doppler shift, OH




Figure 3.10: Results of an investigation into anomalous Ca II blueshift between different model runs. In panel (a), it can be seen that forward models that include absorption due to Sc opacity yield a larger Ca II blueshift than models that lack Sc (Fe Doppler shift is included for comparison). Panel (b) illustrates the cause of this anomalous blueshift: a Sc line overlapping one line in the optical Ca II doublet. These results imply that overlapping line profiles can subtly contaminate calculated Doppler shifts.

should be more blueshifted than CO, and H₂O should be more redshifted than CO.

We shall see, however, that additional, complicating physics is revealed by fully phase-resolved spectra. For our models, the relevant underlying physics is as follows:

1. **Altitude-dependent winds:** H₂O lines are more strongly blueshifted than CO lines at all phases because the H₂O line cores over the wavelength range of the CARMENES band-pass more predominantly form at higher altitudes. At high altitudes, the atmospheric flow switches from dominantly rotational (via an eastward equatorial jet) to dominantly divergent (via day–night winds) (Hammond & Lewis 2021). This result is the opposite of what would be expected from the above-described toy model, revealing the shortcomings of simple models and how they can sometimes mislead us.
2. **Equilibrium chemistry:** H₂O and CO are less blueshifted than OH because OH preferentially forms on the approaching, blueshifted limb of the planet. OH being more blueshifted than the other molecules is in agreement with the predictions of the toy model.

3. **Equilibrium chemistry:** The relationship between H_2O and OH changes as a function of phase because the ratio $\text{OH}/\text{H}_2\text{O}$ increases as a function of temperature, and hotter regions of the planet rotate into view over transit. This finding is also qualitatively in agreement with the toy model. However, per Figure 3.9, this effect is unfortunately not likely to be observable given the error bars in current data sets.

Now the question remains: Can we observe in real data the trends matching these model explanations? As a simple experiment, we can apply error bars representative of the best observing nights on the best instrument with the most observable chemical species (roughly 2 km/s, as drawn as vertical error bars in Figure 3.8; Ehrenreich et al. 2020) and determine whether these trends are still detectable. With our errorbars now applied to our simulated data, only the first explanation—that H_2O forms at higher altitudes than CO—can fully be addressed, assuming that Doppler shifts for both species can be obtained. The second explanation can only be partially addressed—we can still determine that CO is less blueshifted than OH.

3.3.3.2 Warning: blending of Doppler shifts

The disentangling of physics in Section 3.3.3.1 rests on a fundamental assumption: that each cross-correlation template directly tracks only a single species. Indeed, one of the promises of HRCCS is the ability to uniquely constrain individual species’ abundance; with individual line profiles resolved, different species should be readily identifiable from one another in cross-correlation space (e.g., Brogi & Line 2019). Furthermore, our noiseless models should be even less susceptible to degeneracies between different species’ spectral manifestations.

Panel (a) of Figure 3.10 seems to contradict the notion of complete line profile independence

across species. For models run in [Savel et al. \(2022\)](#), Sc was excluded. Motivated by the search for atoms in [Kesseli et al. \(2022\)](#), however, we included Sc in this work’s models. Surprisingly, we found a subsequent significant difference in the Doppler shifts recovered from our cross-correlation analysis in our Sc-inclusive models.

Panel (b) of Figure [3.10](#) reveals the source of the discrepancy. In the optical, Ca II opacity is dominated by a doublet; one of the lines in this doublet partially overlaps with a strong, narrow Sc line. When both species combined in a forward model, the Sc line produces absorption just blueward of this Ca II line’s core; hence, the cross-correlation of the Ca II template yields a spurious blueshift. There did exist other modeling differences between the two spectra (e.g., the [Savel et al. \(2022\)](#) models included TiO and VO), but none of these differences strongly impacted the Doppler shift of Ca II.

Because Ca II in the optical has only two strong lines, it is particularly susceptible to this type of error. All it takes is one slight overlap with another species near a Ca II doublet core, and the Ca II Doppler signal can be significantly biased. Species with forests of lines (e.g., Fe in the optical) should hence be more robust to chance overlaps with other species’ lines.

To guard against this error for species with few lines, we recommend cross-correlating templates against one another to get a first-order sense for the extent of species overlap in Doppler space. Furthermore, we recommend performing these analyses on HRCCS with combined species models, as opposed to single-species models. This approach could involve a retrieval framework ([Brogi & Line 2019](#); [Gandhi et al. 2019](#); [Gibson et al. 2020](#)), which couples a statistical sampler to an atmospheric forward model to determine the exoplanet spectrum that best fits the data, inclusive of multiple chemical species at once.

3.4 Conclusions

The past few years have yielded asymmetric Doppler signals from exoplanet atmospheres as a function of phase. Compelling “toy models” notwithstanding, a number of physical processes can drive these asymmetries, and it can be difficult to uniquely constrain the cause of an asymmetry.

In this study, we determine that if an asymmetry is observed:

1. It may be due to a scale height difference across the atmosphere, not a chemistry difference across the atmosphere. Comparing a signal of a species in HRCCS to a baseline species that is guaranteed to be chemically stable over the atmosphere can better motivate whether the asymmetry could be due to chemistry. CO is an excellent baseline species for ultra-hot Jupiters, as it is stable over these planets’ expected temperature–pressure space, has many spectral lines in the near-infrared accessible to ground-based spectrographs, and has been detected in numerous studies.
2. The asymmetry can be highly informative even if it is binned in phase, especially if multiple species are considered. For instance, much larger Doppler shifts (both blue and red) of certain species relative to the predictions of hydrostatic GCMs can be used as evidence for outflowing material.
3. The asymmetry may be boosted by including (and perhaps only considering) ingress and egress phases. Ingress and egress spectra are the the gold standard for asymmetric signals so long as the signal to noise is high enough.
4. The asymmetry may be influenced by line confusion between species, even at high resolution. Species with very few lines (e.g., a single doublet) in the observed waveband are es-

pecially susceptible to contamination by other species in cross-correlation analysis, and they should be carefully checked against theoretical models for possible contaminating opacity sources.

5. If all species exhibit a similar asymmetry—especially if they all become less blueshifted over the course of transit—the asymmetry may be due to a large-scale effect, such as clouds blanketing the cooler limb.
6. Per our comparison of near-infrared absorbers in the CARMENES waveband, the toy model predictions of the H₂O Doppler shift relative to CO was inaccurate, as it did not include information about the vertical coordinate. With H₂O lines on average probing higher in the atmosphere than CO in this waveband, they probed a different part of the flow, departing from expectations of the toy model.

By aiming to systematically understand even just a few drivers of asymmetry, this work has made it clear that HRCCS—already arguably abstract given its general inability to produce visible planetary spectra—has yet more nuance to uncover. As data quality continues to increase, it will become increasingly necessary to understand the relationships between higher-order physical effects.

Table 3.1: Example drivers of phase asymmetries of ultra-hot Jupiters

Mechanism	Proposed diagnostic	Selected work(s)
Escape	H Lyman- α transit duration Very deep, e.g., Ca II lines Strong vertical winds Very broadened, e.g., Na I lines Excess He absorption (10830Å) Velocities of ions	Owen et al. (2023) Fossati et al. (2013) Seidel et al. (2021a) Hoeijmakers et al. (2020) Oklopčić & Hirata (2018) , Spake et al. (2018) This work
H variation	Blueshifting CCF w/ phase Strongly varying CO velocity	Kempton & Rauscher (2012) Wardenier et al. (2021) , this work
H ₂ dissoc.	Small phase curve amplitude Muted spectrum from H ⁻	Mansfield et al. (2020) Arcangeli et al. (2018)
Weak drag	Blueshifted CCF Large phase curve offset	Wardenier et al. (2021) , Savel et al. (2022) May & Komacek et al. (2021)
Cold interior	Blueshifted CCF Large phase curve amplitude	Savel et al. (2022) May & Komacek et al. (2021)
Fast jet	$\text{FWHM}_{\text{CCF}} > \text{FWHM}_{\text{solid-body}}$ Phase curve offset	Brogi et al. (2016) Knutson et al. (2007b)
Day–night wind	Blueshifted CCF	Snellen et al. (2010)
Eq. chemistry	Limb abundance asymmetry τ_{chem} vs. $\tau_{\text{dynamical}}$	This work Showman et al. (2013a)
Photochemistry	Diseq. abundances of gases	Tsai et al. (2017)
Condensation	Model with GCM Blueshifting CCF w/ phase	Wardenier et al. (2021) , Savel et al. (2022) Ehrenreich et al. (2020)
Eccentricity	All lines at similar velocities Independent orbital constraints	Montalto et al. (2011) , Savel et al. (2022) Montalto et al. (2011)
Clouds	All gases blueshift less w/ phase Blueshifted CCF Lines only present in HRCCS CCFs of water bands	This work Savel et al. (2022) Kempton et al. (2014) , Hood et al. (2020) Pino et al. (2018)
Lorentz forces	Reduced hotspot offset Large phase curve amplitude Westward hotspot offset	Beltz et al. (2021) Beltz et al. (2021) Hindle et al. (2021)
Variable winds	Varying velocity w/ phase Strong vs. weak lines' velocities	Kempton & Rauscher (2012) Kempton & Rauscher (2012)
Tidal strain	Blueshifting CCF over phase Light curve fitting	Future work Akinsanmi et al. (2019)
$\kappa(T)$	Blueshifting CCF over phase	Wardenier et al. (2021)

Chapter 4: A New Lever on Exoplanetary B Fields: Measuring Heavy Ion Velocities

4.1 Introduction

For the foreseeable future, exoplanets can only be investigated by remote sensing. Of the exoplanetary properties accessible by remote sensing, magnetic fields are one of the most fundamental yet least known. While they influence many aspects of planetary science—such as heat redistribution (e.g., [Menou 2012](#)), habitability (e.g., [Meadows & Barnes 2018](#)), evolution (e.g., [Driscoll & Barnes 2015](#)), and interior structure (e.g., [Zarka et al. 2015](#))—we simply do not know their strengths (e.g., [Brain et al. 2024](#)).

Theory does not provide much guidance on B field strengths for hot gas giants. Depending on the dynamo scaling law chosen, these planets’ B field strengths could be on the order of Jupiter’s (~ 4 G at the equator; [Sánchez-Lavega 2004](#); [Stevens 2005](#); [Zaghoo & Collins 2018](#)), an order of magnitude smaller (e.g., [Christensen 2010](#)), or nearly two orders of magnitude larger ([Yadav & Thorngren 2017](#)).

Observations have *also* not provided strong magnetic field constraints. Extensive efforts have been made to connect planetary magnetic field strength to, e.g., the circulation of neutral gas ([Perna et al. 2010](#); [Menou 2012](#); [Rauscher & Menou 2012](#); [Batygin et al. 2013](#); [Rauscher & Menou 2013](#);

Table 4.1: Ba II detections in exoplanet atmospheres

Reference	Planet	Ba II S/N	$u_{\text{Ba II}}$ (km/s)	Fe I S/N	$u_{\text{Fe I}}$ (km/s)
Silva et al. (2022)¹	WASP-76b	7.4	-4.00 ± 0.58	22.2	-3.92 ± 0.34
Silva et al. (2022)	WASP-121b	10.6	-1.9 ± 0.5	16.3	-2.6 ± 0.3
Borsato et al. (2023)²	KELT-9b	4.5	-19.0	15.3	-19.0
Pelletier et al. (2023)³	WASP-76b	4.95	-8.80 ± 0.59	12.94	-6.39 ± 0.29
Prinoth et al. (2023)	WASP-189b	9.12	-0.97 ± 1.14	22.7	-4.31 ± 0.52

Rogers & Showman 2014; Rogers & Komacek 2014; Rogers & McElwaine 2017; Beltz et al. 2021; Knierim et al. 2022) and their velocities measured with high-resolution spectroscopy (Kempton & Rauscher 2012; Beltz et al. 2022, 2023), UV transit timing asymmetries from bow shocks (e.g., Vidotto et al. 2011; Kislyakova et al. 2014), radio auroral emission from the planet (e.g., Zarka 2007; Narang et al. 2024; Shiohira et al. 2024) and star (e.g., Pineda & Villadsen 2023), and helium line velocity measurements (Schreyer et al. 2024) and spectropolarimetry (Oklopčić et al. 2020). Despite these efforts, success has been isolated (Ben-Jaffel et al. 2022), indirect (Cauley et al. 2019), or difficult to reconcile across repeated observations (e.g., Lally & Vanderburg 2022; Turner et al. 2023).

Between theory and observations, B field strengths are unknown to orders of magnitude uncertainty. There is clear need for novel techniques to provide more leverage on this problem. In this chapter, we present a proof of concept: using measured heavy ion–neutral velocity differences to infer exoplanet magnetic field strengths. Unlike neutrals, ions directly “feel” the magnetic field. And unlike light ions, heavy ions require more energy input to be entrained in an escaping flow, making them more robust tracers of field lines. Thus far, I have presented HRCCS as an excellent technique for diagnosing the winds of exoplanet atmospheres; here, I will demonstrate a physical interpretation allowed by such techniques. HRCCS is particularly well suited for the proposed B field measurement because it has yielded numerous heavy ion detections and velocity measurements in just the past few years (Table 4.1; Silva et al. 2022; Borsato et al. 2023; Pelletier et al. 2023; Prinoth et al. 2023).

We present our basic magnetohydrodynamic (MHD) arguments in Section 4.2. We then briefly demonstrate why heavy ions are less likely to be lofted into an atmosphere outflow (Section 4.3). Section 4.4 lays out our primary proof of concept, calculating the ion–neutral velocity

difference for a representative ultra-hot Jupiter atmosphere. Finally, we summarize and provide thoughts on further work in Section 4.5.

4.2 Magnetohydrodynamics

Below, we set up a physical system similar to MHD approaches that calculate “magnetic drag.” Such models consider the drag that neutral particles experience in a partially ionized atmosphere as they collide with charged particles, the latter of which experience Lorentz forces from field lines (e.g., [Zhu et al. 2005](#); [Perna et al. 2010](#); [Rauscher & Menou 2013](#); [Beltz et al. 2021](#)). We here focus not on the motion of neutrals as arrested by ions, but rather the relative motions of the *ions themselves*.⁴ If the flow structures are not tightly coupled, then strong ion–neutral velocity differences, or “slip velocities” (e.g., [Hopkins et al. 2024](#)), would correspond to strong planetary magnetic fields.

By combining the electron and ion momentum equations, [Koskinen et al. \(2014\)](#) find

$$(\mathbf{u}_p - \mathbf{u}_n) \approx \frac{\mathbf{j} \times \mathbf{B}}{n_p m_i \nu_{in}} \quad (4.1)$$

for plasma velocity $\mathbf{u}_p$, neutral velocity $\mathbf{u}_n$, current $\mathbf{j}$, magnetic field $\mathbf{B}$, plasma number density n_p , ion mass m_i , and collision frequency ν_{in} .

Evaluating $\mathbf{j}$ is non-trivial. We can make the simplifying assumption of “resistive MHD,” equivalent to assuming that the conductivity tensor can be expressed as a scalar. Then the current is (e.g., [Perna et al. 2010](#); [Koskinen et al. 2014](#))

⁴We emphasize that an ion–neutral force difference is the mechanism that produces magnetic drag. Therefore, a decoupling of the neutral and charged gas velocity structures is consistent with previous exoplanet MHD studies. What remains to be seen is how tightly these flow structures are coupled.

$$j = \sigma u_n B, \quad (4.2)$$

for conductivity σ , considering only the component of the current that is entirely orthogonal to the (assumed deep-seated, dipolar, aligned with the rotation of the planet) B field. Therefore, the difference in velocities between ions and neutrals can be written as

$$|\mathbf{u}_p - \mathbf{u}_n| \approx \frac{\sigma u_n B^2}{n_p m_i \nu_{in}}. \quad (4.3)$$

We now calculate the ion–neutral collision frequency as (analogously to, e.g., [Helling et al. 2021](#))

$$\nu_{in} = \sigma_{\text{coll}} n_{\text{gas}} v_{\text{rel}}, \quad (4.4)$$

where σ_{coll} is the cross-section for ion–neutral collisions, n_{gas} is the gas density, and v_{rel} is the relative velocity between the neutrals and ions.

In the high-temperature limit, the thermal velocity is expected to be greater than the bulk flow difference $u_p - u_n$.⁵ We hence consider our relative velocity to be the thermal velocity of hydrogen, which is greater than the heavy ions we consider in a given collision.⁶

⁵This approximation breaks down when the slip speed exceeds the thermal speed, at B fields of roughly 100 G. For superthermal slip speeds, bulk flow will contribute to a higher collision rate, and the system may be driven toward *subthermal* slip velocities ([Hillier 2024](#); [Hopkins et al. 2024](#)).

⁶We remark for this calculation, as does [Helling et al. \(2021\)](#), that hot enough atmospheres are not composed of pure molecular hydrogen—an appreciable fraction of H_2 thermally dissociates into atomic H ([Bell & Cowan 2018](#); [Komacek & Tan 2018](#); [Tan & Komacek 2019](#); [Mansfield et al. 2020](#); [Roth et al. 2021](#)). The primary impact of dissociation would be to decrease the mass of the dominant neutral particle, thereby increasing the thermal velocity, increasing the collision frequency, and therefore decreasing the slip velocity for a given B field.

$$v_{\text{rel}} \approx \sqrt{\frac{k_B T}{m_{H_2}}}. \quad (4.5)$$

We also assume that a heavy ion is much larger than a neutral atom in an average collision:

$$\sigma_{\text{coll}} \approx \pi r_i^2. \quad (4.6)$$

Taken together, the slip velocity can therefore be written as

$$|\mathbf{u}_p - \mathbf{u}_n| \approx \frac{\sigma u_n B^2}{n_p m_i n_{\text{gas}} \pi r_i^2} \sqrt{\frac{m_{H_2}}{k_B T}}. \quad (4.7)$$

We also note that, from [Perna et al. \(2010\)](#), the conductivity can be written simply as

$$\sigma = \frac{n_e e^2}{m_e n_n \mathcal{M}_{en}}, \quad (4.8)$$

where n_e is the electron density, e is the electric charge, m_e is the electron mass, n_n is the neutral gas number density, and $\mathcal{M}_{en}$ is the electron–neutral momentum transfer coupling term (in units of speed). [Draine et al. \(1983\)](#) provide the following estimate for $\mathcal{M}_{en}$:

$$\mathcal{M}_{en} = 10^{-15} \sqrt{\frac{128 k_B T}{9 \pi m_e}}. \quad (4.9)$$

Hence

$$\sigma = 10^{15} \frac{n_e e^2}{m_e n_n} \sqrt{\frac{9 \pi m_e}{128 k_B T}}. \quad (4.10)$$

The full expression for the ion–neutral velocity difference is slightly unwieldy, but it is worth

examining.

Altogether, noting that $n_e = n_p$ in the [Koskinen et al. \(2014\)](#) formalism, we have:

$$|\mathbf{u}_p - \mathbf{u}_n| \approx 10^{15} \frac{u_n e^2 B^2}{m_i n_n n_{\text{gas}} r_i^2 k_B T} \sqrt{\frac{9m_{H_2}}{128\pi m_e}}. \quad (4.11)$$

We now identify a few noteworthy scalings. The larger and more massive the ion, the stronger the coupling is to the neutral flow. Hence, some of the heaviest ions detected in exoplanet atmospheres to date (such as Tb II; [Borsato et al. 2023](#)) may not be ideal candidates for tracing B fields. We argue, however, that the lightest ions are not preferable for tracing B fields because they may be entrained in hydrodynamic outflows (explored further in Section 4.3), so a moderately massive ion is ideal.

We also highlight the strong dependence on gas density in Equation 4.11. The lower the density, the less strongly the ion and neutral flows are coupled. This scaling is sensible, because lower density leads to fewer collisions. In hydrostatic atmospheres, gas density falls off exponentially with altitude. So, the best candidate ions for trace B fields have strong opacity, causing the gas to become optically thick at low pressures, in turn allowing spectroscopy to probe low-pressure regions. The tradeoff is that if too low pressures are probed, one may be observing the upper atmosphere and therefore the non-hydrostatic outflow (which begins, to order of magnitude, at 1 nanobar for a standard hot Jupiter and scales linearly with surface gravity; [Murray-Clay et al. 2009](#)).

Also note the dependence on the neutral gas number density, n_n . The higher this quantity is (i.e., the lower the electron number density, n_e), the weaker the relationship between B field and ion–neutral velocity difference. Therefore, a more ionized atmosphere will better enable ions to

trace B fields.

Finally, note the strong dependence on B . Here is where the strength of this proposed technique lies: regardless of assumptions about hydrogen dissociation, or the exact free electron population, the ion–neutral velocity difference is highly dependent on the magnetic field. So, while precisely constraining B fields will require more theoretical work, a high velocity difference can reasonably imply a strong B field, and a low velocity difference can imply a weak B field.

For very strong B fields, this slip velocity can become very large. For such high velocities, we require an additional Lorentz factor:

$$|\gamma_p \mathbf{u}_p - \mathbf{u}_n| \approx 10^{15} \frac{u_n e^2 B^2}{m_i n_n n_{\text{gas}} r_i^2 k_B T} \sqrt{\frac{9m_{H_2}}{128\pi m_e}}, \quad (4.12)$$

where

$$\gamma_p = \frac{1}{\sqrt{1 - \frac{u_p^2}{c^2}}}. \quad (4.13)$$

Equation 4.12 implies that, for atmospheres with very high B fields and ionization fractions, the (relativistic) kinetic energy of the ions in the atmosphere may exceed the binding energy of a given atmospheric shell. This effect would of course violate hydrostatic equilibrium, resulting in strong vertical velocities. Strong vertical velocities have been inferred via resolved alkali lines at high resolution (e.g., [Seidel et al. 2020](#)), though such signatures are likely connected to hydrodynamic outflow driven by instellation. If ion speeds are not regulated by, e.g., current instabilities (e.g., [Norman & Smith 1978](#); [Hopkins et al. 2024](#)), relativistic plasma velocity would also be associated with inverse-Compton scattering and synchrotron radiation from electrons rapidly accelerating along (assumed curved) B field lines. Presently, there is no evidence of such radiation.




Figure 4.1: Calculated properties of the model atmosphere described in Table 4.2. Left: the limb temperature–pressure structure calculated with HELIOS (Malik et al. 2017, 2019). Regions where the atmosphere is convective are highlighted in orange. Pressures generally probed in high-resolution transmission spectroscopy (e.g., Kempton et al. 2014; Gandhi et al. 2020; Hood et al. 2020; Wardenier et al. 2023) are filled in green. Middle: free electron volume mixing ratio calculated by FastChem equilibrium (in black; Kitzmann et al. 2024) and VULCAN disequilibrium (shaded blue; Tsai et al. 2021) chemistry. We run VULCAN over a wide range of parameters—hence, there is a width to this region. Right: ion–neutral velocity difference (Equation 4.11), calculated under equilibrium and disequilibrium chemistry assuming two B field strengths. The gold line indicates the approximate minimum detectable velocity difference via high-resolution spectroscopy (e.g., Ehrenreich et al. 2020). The width of the disequilibrium chemistry grid is due almost entirely to the changing input stellar flux.

4.3 Light ions are not preferable probes of B fields

As presented, the results of Section 4.2 hold for any ion. However, not all ions would be *effective* B field probes—an ideal one remains in the hydrostatic interior. Even if ions roughly follow field lines in the upper atmosphere (Owen 2019), being exposed to the stellar wind introduces a host of other complicating ion processes (e.g., Gronoff et al. 2020). Furthermore, ion velocity structures would be easier to interpret without having to consider a hydrodynamic outflow (e.g., Yelle 2004; Murray-Clay et al. 2009). We maintain, therefore, that it is simpler to interpret ions following field lines if the ions are not entrained in these outflows.

Table 4.2: Parameters of the 1D ultra-hot Jupiter model presented in Section 4.4

Parameter	Value
g (cm s ⁻²)	900
a (AU)	0.025
R (R_{Jup})	1.8
T_{int} (K) ⁷	560
T_{star} (K)	6500
R_{star} ($R_{\odot}$)	1.5
K_{zz} (cm ² s ⁻¹)	[10 ⁸ , 10 ⁹ , 10 ¹⁰ , 10 ¹¹ , 10 ¹² , 10 ¹³]
Stellar XUV spectrum	F5V star, ⁸ HD 189733 ⁹

We argue that the *heavier* an ion is, the more likely it is to not be entrained in outflows. Atmospheric escape is of course a very complex process, with 3D numerical simulations revealing detailed geometric structure (e.g., [Bourrier & Des Etangs 2013](#); [Tripathi et al. 2015](#); [Debrecht et al. 2019](#)) and the stellar environment ultimately controlling many observables (e.g., [Owen et al. 2023](#)). We present below a simple motivating argument in the planet frame.

To briefly demonstrate why a heavy ion would be less likely to participate in an atmospheric outflow, we consider the case of XUV (1–120 nm)-driven hydrodynamic (i.e., photoevaporative) escape. In this regime, escaping hydrogen can drag heavier elements along ([Hunten et al. 1987](#)). [Schaefer et al. \(2016\)](#) define the “critical XUV flux” required for this hydrogen drag to occur as

$$F_{\text{XUV}}^{\text{crit}} = \frac{4b_{12}\mathcal{V}_1^2}{\epsilon k_B T R_P} ((\mu_2/\mu_1) - 1) X_1, \quad (4.14)$$

where b_{12} is the binary diffusion coefficient ([Mason & Marrero 1970](#); [Zahnle & Kasting 1986](#)) between the primary escaping species (1) and the dragged species (2), $\mathcal{V}$ is the potential energy of a single particle of the primary escaping species, ϵ is an efficiency factor that is usually between 0.15 and 0.3 ([Watson et al. 1981](#); [Kasting & Pollack 1983](#); [Chassefière 1996](#); [Tian 2009](#); [Schaefer et al. 2016](#)), R_P is the planetary radius, μ_i is the mean molecular weight of species i , and X_1 is the

molar concentration of the primary escaping species.

As an example, we take the ratio of the critical fluxes required to drag Ba II and Ca II. Assuming that the two have similar binary diffusion coefficients with respect to the primary escaping species,

$$\frac{F_{\text{XUV, Ba II}}^{\text{crit}}}{F_{\text{XUV, Ca II}}^{\text{crit}}} = \frac{(\mu_{\text{Ba II}}/\mu_1) - 1}{(\mu_{\text{Ca II}}/\mu_1) - 1}. \quad (4.15)$$

If these ions are dragged by atomic H, this ratio evaluates to ~ 3.5 . That is, a factor of ~ 3.5 higher XUV flux is required for Ba II escape vs. Ca II escape. This factor is marginally greater than the factor of ~ 3 in both the variation of XUV flux across intermediate-mass main-sequence stars (e.g., [Fossati et al. 2018](#)) and the precision attainable with XUV reconstructions from observations (e.g., [Schaefer et al. 2016](#); [Youngblood et al. 2016](#)). Notably, Ca II has been observed with absorption so strong as to be detected by fitting individual line profiles (as opposed to via cross-correlation), with transit depths on the scale of the planetary Roche lobe (e.g., [Deibert et al. 2021](#); [Tabernero et al. 2021](#); [Silva et al. 2022](#)), whereas Ba II has not. Because the transit depths in the Ba II line cores are smaller than the planetary Roche lobe, the atmospheric layers that Ba II spectral lines probe appear to be within the hydrostatic atmosphere, whereas Ca II line cores almost certainly probe outflowing regions.

The previous conclusion hinges on a mean molecular weight argument. Any other heavy ion will fare similarly—it will remain in the hydrostatic interior, where it will more straightforwardly follow field lines, under a broader range of conditions than a light ion.

4.4 Proof of concept: Barium ion in an ultra-hot Jupiter

Sections 4.2 and 4.3 motivate using a moderately heavy, high-opacity, low-abundance ion to probe exoplanetary B fields. We now flesh out our argument with a case study: Ba II in an ultra-hot Jupiter atmosphere. These atmospheres are hot enough to be appreciably partially (thermally) ionized (e.g., [Arcangeli et al. 2018](#); [Parmentier et al. 2018](#)), making them excellent test cases for our ion-focused framework.

We begin by calculating the limb temperature–pressure (TP) structure of a generic ultra-hot Jupiter atmosphere (Table 4.2) with the HELIOS radiative-convective equilibrium code ([Malik et al. 2017, 2019](#)). We perform our calculation at solar metallicity and carbon-to-oxygen ratio (C/O), including gaseous metals and H^- opacity, at a zenith angle of 83° . This angle is loosely where we expect the actinic flux (i.e., the flux capable of contributing to photochemistry) to reach an optical depth of 1 ([Tsai et al. 2023](#)). The result of this calculation is shown in Figure 4.1. The TP profile is convective at depth and exhibits a temperature inversion, typical for ultra-hot Jupiters (e.g., [Haynes et al. 2015](#); [Evans et al. 2017](#); [Sheppard et al. 2017](#); [Lothringer et al. 2018](#); [Coulombe et al. 2023](#)). In our model, the temperature inversion is predominantly driven by a combination of TiO, VO (deeper than 1 mbar), Fe I (shallower than 1 mbar), Fe II (near 1 microbar), and alkalis (with contribution toward the top of the atmosphere¹⁰).

We next consider equilibrium chemistry by interpolating FastChem ([Kitzmann et al. 2024](#)) tables along our model TP profile, using fits to equilibrium constants including ions up to Uranium.

This step produces the free electron, neutral gas, and Ba II number densities as a function of

¹⁰This result is similar to [Mollière et al. \(2015\)](#), who found that alkalis can cause inversions in hot Jupiter atmospheres without TiO and VO when the C/O ratio is roughly 1. In their case, the alkalis produce inversions because the main molecular cooling agents are depleted through carbon-rich chemistry. In our hotter atmosphere, however, alkalis contribute to inversions at solar C/O due to depletion of these cooling agents via (thermal) dissociation.

pressure. Under equilibrium chemistry, K is the primary electron donor between roughly 10 bars and 1 mbar (Figure 4.2). Between ~ 1 mbar and 1 microbar, Na, Al, Ni, and Mg together yield 2/3 of the total free electron population. Finally, below this pressure, the bulk of the electron donation is from Fe, Si, and Mg. The quantities calculated in this step, along with a given B field strength, allow us to calculate $u_i - u_n$ from Equation 4.11.

Our model predicts that Ba II–neutral velocity differences well exceed 1 km s^{-1} for reasonable B field values (Figure 4.1).¹¹ This method is evidently sensitive to B field strengths to a measurable degree.

We note that this method requires relatively strong constraints on the pressure that both the ion and neutral spectral lines probe. Uncertainties in the contribution function by a dex in pressure would correspond to uncertainties in the estimated B field of more than an order of magnitude. Even so, an order of magnitude constraint on an exoplanet B field would still be highly useful. The calculation also assumes that both the ion and neutral probe the same pressure level. This approximation holds so long as the measured velocity of the neutral corresponds to the neutral velocity at the measured *pressure* of the ion, or vice versa. For instance, the calculation is valid if Ba II is measured at 1 microbar and Fe I at 10 microbars, and the Fe I flow is approximately constant with pressure below 10 microbars. Simulations would be able to assess how reasonable this scenario might be.

Thus far, we have only considered thermal ionization in our model atmosphere. Highly irradiated gas giants may additionally be subject to photoionization at low pressures, which pushes the free electron number density out of thermal equilibrium (e.g., Yelle 2004; Cho et al. 2008;

¹¹Interestingly, our model also predicts that ions and neutrals are very well coupled deep in the atmosphere. This result essentially affirms the magnetization regimes laid out in Koskinen et al. (2014).




Figure 4.2: The primary electron donors in our chemical calculations. The thick solid lines represent equilibrium chemistry. Thin solid lines represent VULCAN disequilibrium chemistry with the F5V stellar spectrum, and thin dotted lines represent disequilibrium chemistry with the HD 189733 stellar spectrum.

Koskinen et al. 2010; Helling et al. 2021). To assess the impact of this uncertainty, we run a grid of VULCAN (Tsai et al. 2021) photochemical simulations based on our HELIOS temperature structure. We use a chemical reaction network that includes neutrals *and* ions for this calculation, as previously shown in Lee et al. (2020) and Coulombe et al. (2023).¹² Varying the unknown eddy diffusion coefficient, K_{zz} , by orders of magnitude, in addition to toggling molecular diffusion and changing the input stellar spectra (Table 4.2, Figure 4.3), consistently points to a single result: uncertainty in the XUV spectrum is the primary driver of uncertainty in our calculation. The strong scaling of $u_i - u_n$ with the planetary B field still allows order of magnitude estimates of the latter from the former, but quantitative estimates of the planetary B field will require estimates of the stellar flux at wavelengths shorter than ~ 100 nm.¹³

¹²The official release of `ion-vulcan` is in preparation.

¹³Our chemical network does not currently include the ionization of neutral barium (e.g., Griesmann et al. 1992) or of Ba II. Therefore, we cannot assess the effect of photoionization on the Ba II population itself. However, the intrinsic number density of this ion is so low that the contribution of Ba II to the total gas number density (the only way the abundance of Ba II features in Equation 4.11) is very small. The n_i assumed here under equilibrium chemistry could increase by 6 orders of magnitude before affecting $u_i - u_n$ at the percent level—making our results very robust to this source of uncertainty.




Figure 4.3: The stellar UV spectra input to VULCAN.

When including photochemistry, K is still the main ion donor at ~ 1 bar, but a combination of K, Na, and Ca dominates ion donation between 1 bar and almost 10 microbars. Below this pressure, H II is the major ion donor. This upper atmosphere behavior is the largest discrepancy with respect to equilibrium chemistry ion donation. Presumably, this effect arises because H_2 is photodissociated into atomic H (e.g., [Yelle 2004](#); [Muñoz 2007](#)), which is then readily (photo)ionized. Note that this departure only occurs for the photochemical model with more XUV flux. In the other model, we revert to the equilibrium expectation that Mg and Fe predominantly donate ions in the upper atmosphere. Just as with the free electron population, the more XUV flux is considered, the greater the departure from equilibrium chemistry expectations. Broadly, photochemistry does not appear to have a very significant impact on our results. This robustness is likely due to the high temperatures of our ultra-hot Jupiter causing fast chemical reaction rates and a tendency toward chemical equilibrium.

To round out our proof of concept, we calculate the transmission spectrum of this atmosphere under chemical equilibrium. We do so with the transmission version of the 1D CHIMERA code (e.g., [Line et al. 2013](#)). We calculate our spectrum at a resolution of $R = 300,000$ and subsequently




Figure 4.4: An $R = 145,000$ CHIMERA optical transmission spectrum of an ultra-hot Jupiter under chemical equilibrium. A spectrum with Ba II as the only spectrally active trace species is overplotted (gold), indicating the location and strength of Ba II absorption lines in our model atmosphere. Ba II lines in this model probe pressures near a microbar. We generate an analogous spectrum for neutral iron, as well (black).

convolving it to $R = 145,000$, similar to the resolution of the ESPRESSO spectrograph (Pepe et al. 2021). We include Ba II, Fe, Mg, Ca II, Na, K, and VO as opacity sources.¹⁴ We also include Rayleigh scattering from H and He. We calculate the vertically varying equilibrium abundances of these gases according to the HELIOS-provided TP profile. We do not explicitly include H^- . H^- opacity is roughly constant over this wavelength range, so we simply model it as a continuum opacity (e.g., Pelletier et al. 2023). We set this continuum rather arbitrarily at 10 mbar, emulating the effect of a large free electron population.

We show the resultant spectrum in Figure 4.4. It is clear that multiple Ba II lines peek out above the continuum, with a few of the strongest reaching a slant optical depth of unity at pressures less than 1 microbar. Referring back to Figure 4.1, slip velocities evidently are expected to be measurable for a wide range of B field strengths within this pressure range. Fe appears to be a suitable neutral against which to compare; it has a number of strong lines that probe similar pressure levels.

¹⁴All species' opacity data are drawn from Kurucz (2011). The exception is VO, which is drawn from McKemmish et al. (2016).

4.5 Conclusions

Magnetic fields in exoplanets are largely unconstrained by both theory and observation. Motivated by the need for more leverage on this fundamental quantity, we demonstrate that tracking heavy ion–neutral velocity differences is useful for estimating planetary magnetic fields. Because ions are deflected by magnetic fields and neutrals are not, strong differences between the two populations’ velocities imply strong magnetic fields. A recommended recipe for applying this technique is as follows:

1. Constrain the XUV flux of the host star (such as extrapolating measurements from Hubble COS observations; e.g., [France et al. 2016](#)).
2. Assess whether the XUV flux is strong enough to appreciably drag Ba II in escaping flow (e.g., Equation [4.14](#)).
3. Measure the velocities of a suitable heavy ion (such as Ba II) and a suitable heavy neutral (such as Fe) with high-resolution spectroscopy. Ideally, this step is performed in a retrieval context, thus constraining the thermal structure of the atmosphere along with other potential aliases and/or degeneracies (e.g., [Madhusudhan & Seager 2009](#); [Brogi & Line 2019](#); [Gandhi et al. 2019](#); [Gibson et al. 2020](#); [Pelletier et al. 2023](#)).
4. Calculate the electron density at the pressure level probed by the two gas species, ideally accounting for photoionization.
5. With the available information, estimate the B field strength (Equation [4.11](#)).

From an observational perspective, this approach is ready to be applied immediately. For a

substantial range of B field strengths, the predicted velocity differences here are large enough to observe (Table 4.1). Specifically, ultra-hot Jupiters orbiting bright stars appear to be particularly amenable to this technique, as their atmospheres are highly irradiated enough to support a substantial steady-state ion population and large enough to appreciably occult their host stars—in addition to exhibiting relatively low photon noise. Even for these planets, discrepancies have been seen between high-resolution measurements made with different instruments (e.g., Table 4.1) or with different epochs on the same planet, which may be astrophysical or may be revealing some key differences in data processing procedures. Such issues may be resolved with a larger sample size: broader observational campaigns and archival efforts to measure heavy ion–neutral velocity differences could build up a population of such detections, with ramifications for understanding hot Jupiter interior structure and dynamo theory in an interesting forcing regime (large energy flux under slow rotation; e.g., Christensen 2010). In particular, it would be valuable to assess the inferred B field strength as a function of equilibrium temperature and rotation, as both of these planetary properties could impact dynamo field generation (e.g., Kao et al. 2016; Shulyak et al. 2017; Yadav & Thorngren 2017; Brain et al. 2024).

Our proposed method benefits from the availability of relevant observations, as measured slip velocity precision (a few tenths to 1 km/s; Table 4.1) appears amenable to order-of-magnitude B field constraints (Figure 4.1). Even so, ours is a simple toy model, so it currently neglects a number of confounding factors. Further theoretical work is needed to enable quantitative estimates of (potentially very diverse) exoplanetary magnetic fields. Avenues for future work include accounting for:

- neutral flow feedback processes and microphysical regulating effects, which would likely act

against superthermal ion flow and require a reformulation of the induction equation assumed in this work (e.g., [Norman & Smith 1978](#); [Hillier 2024](#); [Hopkins et al. 2024](#))

- more explicit treatment of multi-fluid MHD (e.g., [Wardle & Ng 1999](#); [Hopkins et al. 2024](#)), as Ba II is not the only ion species—and, in fact, is not the primary electron donor
- B field geometric complexity, such as departures from a dipole in the deep-seated field (Jupiter has substantial quadrupole and octopole B field contributions; e.g., [Acuna & Ness 1976](#)) or induced magnetic fields (e.g., [Rogers & Showman 2014](#); [Rogers & McElwaine 2017](#); [Soriano-Guerrero et al. 2023](#))
- tensor conductivity (e.g., [Koskinen et al. 2014](#))
- the prospect of using multiple heavy ion–neutral pairs to calibrate B field estimates for a single planet
- more physical photochemistry, such as ionic ambipolar diffusion and the photoionization of Ba
- whether XUV reconstructions are sufficient for reducing the uncertainty in the free electron population (e.g., [Youngblood et al. 2017](#); [Melbourne et al. 2020](#); [Teal et al. 2022](#))
- 3D ion–neutral velocity differences, especially as projected onto high-resolution measurements of line-of-sight velocities (e.g., [Kempton & Rauscher 2012](#); [Showman et al. 2013a](#); [Flowers et al. 2019](#); [Beltz et al. 2020](#); [Harada et al. 2021](#); [Malsky et al. 2021](#); [Wardenier et al. 2021](#); [Beltz et al. 2023](#); [Savel et al. 2023](#); [Wardenier et al. 2023](#)). At the poles, dipolar field lines are parallel to lines of sight, whereas they are perpendicular at the equator. It is not obvious how the latitude dependence of our effect averages over, e.g., the terminator

Of these considerations, one of the most crucial to address is the neutral flow feedback processes. Our model is clearly not yet useful in the strong B field limit, where ion velocity must be limited by interactions with the bulk flow, because the ion velocities otherwise become unreasonably large.

Of similar importance are 3D effects. On one hand, 3D considerations could open up further avenues for investigation, with the potential for inferring not only exoplanetary B field strength, but also their *orientation*, via phase-resolved ion–neutral velocity differences. On the other hand, “3D-ness” introduces complexities: velocity differences between gases can be caused by a variety of spatial variations (e.g., [Ehrenreich et al. 2020](#); [Wardenier et al. 2021](#); [Savel et al. 2022](#); [Beltz et al. 2023](#); [Savel et al. 2023](#); [Wardenier et al. 2023](#)). This concern further motivates including self-consistent 3D forward models in high-resolution retrievals. If a 3D model can “model out” the slip velocity contribution due to, e.g., presence of Ba II on one limb and not the other, any residual slip velocity could be attributed to magnetic effects. This obstacle additionally presents an opportunity for using complementary methods for constraining magnetic field strength (e.g., measuring neutral flow velocity; e.g., [Beltz et al. 2021](#)). Agreement across multiple different B field measurement techniques would greatly increase confidence in a given constraint.

Looking forward: this technique is in principle applicable to any gas giant that is hot enough for sufficient ionization and is amenable to characterization with high-resolution spectroscopy. Especially once the extremely large telescopes (ELTs) are on-sky (e.g., [Gilmozzi & Spyromilio 2007](#); [Johns et al. 2012](#)), observers will be able to directly resolve a large number of exoplanetary spectral lines (e.g., [Palle et al. 2023](#)), providing greater sensitivity to individual species’ velocity profiles through line profiles (e.g., [Seidel et al. 2021a](#)). These telescopes will additionally be able to target fainter host stars with high-resolution spectroscopy, thus expanding our planetary

population even further. Moving beyond individual planets will potentially allow us to marginalize over unknown physics. Successfully doing so would reveal, at a basic level, how magnetic fields are generated in planetary bodies.

Chapter 5: cortecs: A Python package for compressing opacities

5.1 Introduction

Previous chapters have indicated that HRCCS is uniquely sensitive to the dynamics of exoplanet atmospheres. A cost of additional sensitivity is computational expense: high-resolution spectra require high-resolution models to interpret, and high-resolution models require solving the radiative transfer equation over tens of thousands of wavelengths (e.g., [Beltz et al., 2023](#); [Gandhi et al., 2023](#); [Line et al., 2021](#); [Maguire et al., 2023](#); [Prinoth et al., 2023](#); [Savel et al., 2022](#); [Wardenier et al., 2023](#)). This problem will become even more pressing as retrievals become more complex — i.e., as they seek to interpret the 3D structure of exoplanet atmospheres. How can we ensure that these science cases will be tractable?

To decrease computational runtime, some codes have parallelized the problem on GPUs (e.g., [Lee et al., 2022](#); [Line et al., 2021](#)). However, GPUs cannot in general hold large amounts of data in their video random-access memory (VRAM) (e.g., [Ito et al., 2017](#)); only the cutting-edge, most expensive GPUs are equipped with VRAM in excess of 30 GB (such as the NVIDIA A100 or H100). While VRAM available in consumer hardware has increased in recent years, RAM and VRAM management is therefore a clear concern when producing high-resolution spectra.

How do we decrease the RAM footprint of these calculations? By far the largest contributor to the RAM footprint, at least as measured on disk, is the opacity data. For instance, the opacity data for a single gas species across the wavelength range of the Immersion GRating INfrared Spectrometer spectrograph (IGRINS, [Mace et al., 2018](#)) takes up 2.5 GB of non-volatile memory (i.e., the file size is 2.5 GB) at `float64` precision and at a resolving power of 400,000 (as used in [Line et al. \(2021\)](#); with 39 temperature points and 18 pressure points, using, e.g., the Polyansky

et al. (2018) water opacity tables). In many cases, not all wavelengths need to be loaded, e.g. if the user is down-sampling the resolution of their opacity function. Even so, it stands to reason that decreasing the amount of RAM/VRAM consumed by opacity data would strongly decrease the total amount of RAM/VRAM consumed by the radiative transfer calculation.

One solution is to isolate redundancy: While the wavelength dependence of opacity is sharp for many gases, the temperature and pressure dependencies are generally smooth and similar across wavelengths (e.g., Barber et al., 2014; Coles et al., 2019; Polyansky et al., 2018). This feature implies that the opacity data should be compressible without significant loss of accuracy at the spectrum level.

While our benchmark case (see Benchmark) demonstrates the applicability of `cortecs` to high-resolution opacity functions of molecular gases, the package is general and can be applied to any opacity data that has pressure and temperature dependence, such as the opacity of neutral atoms or ions. Our benchmark only shows, however, that the amounts of error from our compression technique is reasonable in the spectra of exoplanet atmospheres at pressures greater than a microbar for a single composition. This caveat is important to note for a few reasons:

1. Based on error propagation, the error in the opacity function may be magnified in the spectrum based on the number of cells that are traced during radiative transfer. The number of spatial cells used to simulate exoplanet atmospheres (in our case, 100) is small enough that the `cortecs` error is not large at the spectrum level.
2. Exoplanet atmospheres are often modeled in hydrostatic equilibrium at pressures greater than a microbar (e.g., Barstow et al., 2020; Showman et al., 2020). When modeling atmospheres in hydrostatic equilibrium, the final spectrum essentially maps to the altitude at which the gas




Figure 5.1: The results of a simple toy model demonstrating a `cortecs` failure mode. The toy model is a single column, with the number of grid cells within the column varied (along the x-axis). For each number of cells, we calculate the difference in the optical depth τ as calculated using `cortecs`-compressed opacities (using the polynomial compression method) and the second-to-last wavelength point of native `PLATON` methane opacity.

becomes optically thick. If `cortecs`-compressed opacities were used to model an optically thin gas over large path lengths, however, then smaller opacities would be more important. `cortecs` tends to perform worse at modeling opacity functions that jump from very low to very high opacities, so it may not perform optimally for these optically thin scenarios.

3. The program may perform poorly for opacity functions with sharp temperature–pressure dependence (e.g., the Lyman series transitions of hydrogen; [Kurucz, 2017](#)). That is, the data may require so many parameters to be fit that the compression is no longer worthwhile.

The effects of points 1 and 2 are shown in [Figure 5.1](#). The column simulated varies in temperature from 100 K to 500 K at a constant pressure of 0.1 Pa and a constant density of 0.003 g/cm³. The absolute value of the optical depth is far from the true value at this wavelength. Over short distances, this error is infinitesimal. Over much larger distances, the error accumulates, leading to large differences in $e^{-\tau}$ (the transmission as yielded by Beer’s Law). Hence, in this limiting case, the poor fit at very low opacity levels is not acceptable because of the large physical distances simulated.

5.2 Methods

`cortecs` seeks to compress redundant information by representing opacity data not as the opacity itself but as fits to the opacity as a function of temperature and pressure. We generally refer to this process as “compression” as opposed to “fitting” to emphasize that we do not seek to construct physically motivated, predictive, or comprehensible models of the opacity function. Rather, we simply seek representations of the opacity function that consume less RAM/VRAM. The compression methods we use are “lossy” — the original opacity data cannot be exactly recovered with our methods. We find that the loss of accuracy is tolerable for at least the hot Jupiter emission spectroscopy application (see Benchmark below).

We provide three methods of increasing complexity (and flexibility) for compressing and decompressing opacity: polynomial-fitting, principal components analysis (PCA, e.g., [Jolliffe & Cadima, 2016](#)) and neural networks (e.g., [Alzubaidi et al., 2021](#)). The default neural network architecture is a fully connected neural network; the user can specify the desired hyperparameters, such as number of layers, neurons per layer, and activation function. Alternatively, any `keras` ([Chollet & others, 2015](#)) model can be passed to the fitter. Each compression method is paired with a decompression method for evaluating opacity as a function of temperature, pressure, and wavelength. These decompression methods are tailored for GPUs and are accelerated with the JAX code transformation framework ([Bradbury et al., 2018](#)). An example of this reconstruction is shown in [Figure 5.2](#).




Figure 5.2: Top panel: The original opacity function of CO (Rothman et al., 2010) (solid lines) and its `cortecs` reconstruction (transparent lines) over a large wavelength range and at multiple temperatures and pressures. Bottom panel: the absolute residuals between the opacity function and its `cortecs` reconstruction. Note that opacities less than 10^{-60} are not generally relevant for the benchmark presented here; an opacity of $\sigma_\lambda = 10^{-60}$ would require a column nearly 10^{35} m long to become optically thick at a pressure of 1 bar and temperature of 1000 K.

5.3 Workflow

A typical workflow with `cortecs` involves the following steps:

1. Acquiring: Download opacity data from a source such as the ExoMol database ([Tennyson et al., 2016](#)) or the HITRAN database ([Gordon et al., 2017](#)).
2. Fitting: Compress the opacity data with `cortecs`’s `fit` methods.
3. Saving: Save the compressed opacity data (the fitted parameters) to disk.
4. Loading: Load the compressed opacity data from disk in whatever program you’re applying the data—e.g., within your radiative transfer code.
5. Decompressing: Evaluate the opacity with `cortecs`’s `eval` methods.

The accuracy of these fits may or may not be suitable for a given application. It is important to test that the error incurred using `cortecs` does not impact the results of your application—for instance, by using the `cortecs.fit.metrics.calc_metrics` function to calculate the error incurred by the compression and by calculating spectra with and without using `cortecs`-compressed opacities. We provide an example of such a benchmarking exercise below.

5.4 Benchmark: High-resolution retrieval of WASP-77Ab

As a proof of concept, we perform a parameter inference exercise (a “retrieval,” [Madhusudan & Seager, 2009](#)) on the high-resolution thermal emission spectrum of the fiducial hot Jupiter WASP-77Ab ([August et al., 2023](#); [Line et al., 2021](#); [Mansfield et al., 2022](#)) as observed at IGRINS. The retrieval pairs `PyMultiNest` ([Buchner et al., 2014](#)) sampling with the `CHIMERA` radiative transfer code ([Line et al., 2013](#)), with opacity from H_2O ([Polyansky et al., 2018](#)), CO ([Rothman](#)

et al., 2010), CH₄ (Hargreaves et al., 2020), NH₃ (Coles et al., 2019), HCN (Barber et al., 2014), H₂S (Azzam et al., 2016), and H₂ – H₂ collision-induced absorption (Karman et al., 2019). The non-compressed retrieval uses the data and retrieval framework from (Line et al., 2021), run in an upcoming article (Savel et al. 2024, submitted). For this experiment, we use the PCA-based compression scheme implemented in `cortecs`, preserving 3 principal components and their corresponding weights as a function for each wavelength as a lossy compression of the original opacity data.

Using `cortecs`, we compress the input opacity files by a factor of 13. These opacity data (as described in the previous paragraph) were originally stored as 2.0 GB .h5 files containing 39 temperature points, 18 pressure points, and 373,260 wavelength points. The compressed opacity data are stored as 154 MB files of PCA coefficients and 1.1 KB files of PCA vectors (which are reused for each wavelength point). These on-disk memory quotes are relatively faithful to the in-memory RAM footprint of the data when stored as `numpy` arrays (2.1 GB for the uncompressed data and 160 MB for the compressed data). Reading in the original files takes 1.1 ± 0.1 seconds, while reading in the compressed files takes 24.4 ± 0.3 ms. Accessing/evaluating a single opacity value takes 174.0 ± 0.5 ns for the uncompressed data and 789 ± 5 ns for the compressed data. All of these timing experiments are performed on a 2021 MacBook Pro with an Apple M1 Pro chip and 16 GB of RAM.

Importantly, we find that our compressed-opacity retrieval yields posterior distributions (as plotted by the `corner` package, Foreman-Mackey, 2016) and Bayesian evidences that are consistent with those from the benchmark retrieval using uncompressed opacity (Figure 5.3) within a comparable runtime. The two posterior distributions exhibit slightly different substructure, which we attribute to the compressed results requiring 10% more samples to converge (about 5 hours




Figure 5.3: The posterior distributions for our baseline WASP-77Ab retrieval (teal) and our retrieval using opacities compressed by `cortecs` (gold).

of extra runtime on a roughly 2 day-long calculation) and residual differences between the compressed and uncompressed opacities. The results from this exercise indicate that our compression/decompression scheme is accurate enough to be used in at least some high-resolution retrievals.

Method	Compression factor	MAD	Compression time (s)	$t_{\text{decompress}}$ (s)
PCA	13	0.30	2.6×10^1	2.3×10^2

Method	Compression factor	MAD	Compression time (s)	$t_{\text{decompress}}$ (s)
Polynomials	44	0.24	7.8×10^2	3.6×10^3
Neural network	9	2.6	1.4×10^7	3.6×10^4

Comparison of compression methods used for the HITEMP CO line list ([Rothman et al., 2010](#)) over the IGRINS wavelength range at a resolving power of 250,000, cumulative for all data points. Note that the neural network compression performance and timings are only assessed at a single wavelength point and extrapolated over the full wavelength range.

Chapter 6: Peering into the Black Box: Forward Modeling of the Uncertainty Budget of High-resolution Spectroscopy of Exoplanet Atmospheres

6.1 Introduction

Previous chapters have discussed many of the benefits of HRCCS. This technique is unique in its ability to directly constrain exoplanetary winds (e.g., [Snellen et al. 2010](#); [Kempton & Rauscher 2012](#); [Showman et al. 2013b](#)), and by resolving tens of thousands of individual spectral lines, it has enabled precise constraints on chemical composition (e.g., [Line et al. 2021](#); [Gandhi et al. 2022](#); [Maguire et al. 2022](#); [Boucher et al. 2023](#)). With clear evidence of three-dimensional physics and chemistry ([Flowers et al. 2019](#); [Beltz et al. 2020](#); [Ehrenreich et al. 2020](#); [Savel et al. 2021](#); [Wardenier et al. 2021](#); [Gandhi et al. 2022](#); [Beltz et al. 2023](#); [Nortmann et al. 2024](#)) and the potential to unveil dozens of chemical species in the optical and near-infrared (e.g., [Carleo et al. 2022](#); [Kesseli et al. 2022](#); [Prinoth et al. 2022](#)), HRCCS datasets promise a strong complement to their lower-resolution counterparts, one that will endure even in the era of JWST atmospheric studies ([Gardner et al. 2006](#); [Greene et al. 2016](#); [Brogi & Line 2019](#)).

Complicating the high scientific potential of HRCCS are the subtleties in its analysis. Unlike in low-resolution studies, exoplanet spectra in general are not detected in each observed spectral channel with HRCCS, but are rather buried in the noise and not visible “by eye.” Therefore, the presence of the planetary signal in HRCCS must be inferred from cross-correlation functions (e.g., [Snellen et al. 2010](#)). This feature of HRCCS is accompanied by “black box”-like data processing. While the steps to this analysis are fundamentally deterministic and *can* be understood with forward modeling, the complexity of the underlying physical processes (namely tellurics) in turn requires complex statistical treatments.

Further complicating matters is the axis of researcher choice. Previous studies have sought to shed light on the impact of these processing steps, comprehensively exploring the sensitivity of HRCCS results to analysis choices (de Kok et al. 2013; Cabot et al. 2019; Langeveld et al. 2021; Gully-Santiago & Morley 2022; Meech et al. 2022; Brogi et al. 2023; Cheverall et al. 2023). Furthermore, studies often “inject” additional signal into data to check the responsivity of their dataset to true signal (e.g., Gibson et al. 2022). Even so, it is difficult to exactly isolate in which circumstances the data, the analysis techniques, or some combination drives the behavior of the results. Taken together, these aspects have produced an opaque uncertainty budget for HRCCS.

In this chapter, we seek to provide intuition for the less intuitive aspects of HRCCS data processing and analysis to better understand the HRCCS uncertainty budget. We do so with a forward-modeling approach, simulating HRCCS emission spectroscopy datasets and analyzing them as one would analyze real data. Our focus is threefold: on variable tellurics and their removal, on wavelength solution stability, and on the level of photon noise. All of these uncertainty sources are unexplored in HRCCS from the modeling perspective, and they are hypothesized to explain outstanding questions in HRCCS (e.g., the sub-unity scaling of the planetary spectrum found by van Sluijs et al. 2023a). For this work, our model aims to mimic the HRCCS dataset of WASP-77A b from Line et al. (2021) observed with the IGRINS spectrograph (Yuk et al. 2010; Park et al. 2014; Lee et al. 2017; Mace et al. 2018). We focus on this dataset because of its high quality, the availability of many telluric standard stars at this observatory, and the validation of this dataset’s analysis by recent JWST results (August et al. 2023). Though our work focuses on simulating emission spectra, our results have direct implications for transmission spectroscopy studies as well, and we will comment on them in detail. Section 6.2 describes the methods of this work, including our HRCCS data simulation pipeline—which we name `scope`—and our reduction of

computational burden. Section 6.3 contains the results of our forward-modeling experiments. Section 6.4 discusses our results, including a list of known caveats. Finally, we conclude this work in Section 6.5.

6.2 Methods

To better understand the impacts of HRCCS data-processing steps and choices, we simulate HRCCS observations with the `scope` pipeline (Fig. 6.1). After detailing the steps involved in simulating these datasets, we then describe the signal extraction procedures used. Afterward, we discuss how we motivate the data-driven tellurics forward model specifically, and finally how we assess spectrograph instability to motivate the instabilities injected in our datasets.

6.2.1 Baseline observation simulation pipeline and data processing

Our forward-modeling approach is similar in spirit to the methods described in Brogi & Line (2019). To create an HRCCS dataset, `scope`:

1. Calculates time stamps of simulated observations
2. Creates a data cube
3. Includes and Doppler shifts planetary spectrum
4. Includes and Doppler shifts stellar spectrum
5. Includes telluric absorption
6. Includes the blaze function




Figure 6.1: Schematic of `scope` pipeline. This example uses a planetary emission spectrum calculated with the CHIMERA (Line et al. 2013) atmospheric model as input. This simulation is meant to reconstruct the IGRINS dataset used to identify the emission of WASP-77A b by Line et al. (2021). The slight discontinuity in the planetary emission tracks is due to a gap in exposures corresponding to these observations. The SNR of this simulation matches the SNR of the Line et al. 2021 observations (SNR=12.8); 9 other noise realizations span SNRs from 11.0 to 19.0.

7. Includes photon noise

8. Includes wavelength solution instability (ignored in the baseline case)

6.2.1.1 Integration time calculator

Setting exposure times in HRCCS requires trading off two considerations. Exposing longer decreases the photon noise; however, exposing shorter means the planetary lines have less time to Doppler shift a planetary pixel, or “smear” (see Boldt-Christmas et al. 2023 for an investigation of this tradeoff).

For this study, we use the exposure times (70 s) and time sequences from Line et al. (2021).

scope can also calculate the exposure time before a planet’s lines cross a pixel. Because IGRINS has 3.3 pixels per resolution element with a resolving power per element of roughly 45,000 (e.g., [Mace et al. 2018](#)), the velocity scale of an individual resolution element is $\approx 6.7 \text{ km s}^{-1}$, and the velocity scale of a single pixel is roughly $\approx 2.0 \text{ km s}^{-1}$. With this value determined, we calculate the maximum integration time before this “crossing,” t_{int} , as

$$t_{\text{int}} = \frac{v_{\text{pixel}}}{\frac{dv_{\text{LOS,planet}}}{dt}}, \quad (6.1)$$

for the planetary change in line-of-sight velocity dv_{planet}/dt and the velocity scale of pixel v_{pixel} . We use the `exoplanet` ([Foreman-Mackey et al. 2021b](#)) package to numerically evaluate the line-of-sight acceleration of the modeled planet at each point in its orbit, and we take the maximum time as the constant exposure time for our set of observations. The maximum integration time we calculate for WASP-77A b over the simulated phase range (260 s) is much greater than the actual adopted exposure time (70 s).

It should be noted that overheads should also be considered when setting exposure times. While they do not necessarily impact Doppler smearing, they do factor into the total amount of photons collected during the night and therefore the statistical significance of any recovered signals. The readout time for the WASP-77A b observations with IGRINS in [Line et al. \(2021\)](#) is 28 s per frame, which is a significant fraction of the exposure time.

Setting the exposure time sets the time sequence of our simulated observations, which in turn controls the number of observations in our data cube (Section [6.2.1.2](#)), the Doppler shifting of the planetary (Section [6.2.1.3](#) and stellar (Section [6.2.1.4](#)) spectra, the telluric variation (Section [6.2.1.5](#)), and the signal-to-noise ratio (SNR) that can be achieved for our observations

(Section 6.2.1.7).

6.2.1.2 Data cube

In HRCCS datasets, observed data are customarily packaged in data cubes as, e.g., matrices of shape $N_{\text{exposures}} \times N_{\text{orders}} \times N_{\text{pixels}}$ (e.g., [Line et al. 2021](#), though note that $N_{\text{exposures}}$ and N_{orders} are often swapped in the literature; e.g., [Brogi et al. 2023](#)). To construct our data cube, we match the number of exposures (79) in the [Line et al. \(2021\)](#) dataset. We then extract the order and pixel grid (44 orders by 1848 pixels) from the WASP-77A b dataset. As in [Line et al. \(2021\)](#), we drop 10 orders that are severely impacted by tellurics in the IGRINS waveband.

6.2.1.3 Planet spectrum

In this study, we model planetary emission with the CHIMERA code ([Line & Yung 2013](#); [Line et al. 2013, 2014, 2015, 2016, 2017](#); [Brogi & Line 2019](#); [Line et al. 2021](#)), using the best-fitting atmospheric and velocity parameters from [Line et al. \(2021\)](#), calculated at a native resolution of 250,000 prior to convolution. Our temperature–pressure profile is calculated using the [Madhusudan & Seager \(2009\)](#) prescription.

To include a planetary spectrum (scaled by the square of the planetary radius) in our simulated observations, we first convolve the spectrum down to the instrumental resolution. We also convolve the spectrum with a rotational kernel to include the effect of planetary rotational broadening (e.g., [Gray 2021](#)). Finally, we calculate the planetary orbital radial velocity at each time step in our data cube and Doppler shift the planetary spectrum to that velocity. The planetary spectrum must be calculated to regions well outside the instrumental wavelength range, so that Doppler

shifting and convolving the spectrum does not introduce spurious edge effects.

For a subset of our simulations, we alter both the planetary and stellar (see 6.2.1.4) signals to include the observer’s movement around the Solar System’s barycenter. Doing so introduces another net time-varying velocity component to the measured signal.

6.2.1.4 Stellar spectrum

Incorporating a stellar spectrum involves similar convolving and Doppler shifting as the planetary spectrum (Section 6.2.1.3), the key difference being that the stellar spectrum is Doppler shifted to its radial velocity around its system barycenter, as opposed to the planetary orbital motion. Both the planet and star can be shifted by the system’s overall radial velocity, which we set to the 1.6845 km s^{-1} systemic velocity of WASP-77A b (Maxted et al. 2012).

For the simulations presented in this work, we exclusively use PHOENIX stellar models (Husser et al. 2013) scaled by the square of the stellar radius. For our emission benchmark, we use the same PHOENIX stellar model corresponding to the known parameters of WASP-77A ($T_{\text{eff}} = 5500 \text{ K}$; $\log g = 4.33$; $\text{Fe/H} = 0.00$; Maxted et al. 2012). In the IGRINS waveband (roughly 1.45 to 2.45 microns), these models have a median native resolution of 580,000; like the planetary models, they are then convolved down to instrumental resolution.

6.2.1.5 Tellurics

To include telluric absorption in our simulation, we first determine the airmass at each timestamp of our data cube. For the simulations presented in this work, we use the airmasses from the Line et al. (2021) observations, which range from 1.5 to 3.1. We then incorporate tellurics by

multiplying the data cube by telluric spectrum at each timestamp, interpolating to the wavelength grid of each order. See Section 6.2.2.1 for this work’s calculation of tellurics.

6.2.1.6 Blaze function

The final spectral component of `scope` is the instrumental blaze function (the wavelength-dependent instrumental response to uniform illumination). Given that this work focuses on reproducing IGRINS observations, we provide a fit specifically to the IGRINS blaze function based on standard star observations.

We access the IGRINS standard star data by scraping the Raw & Reduced IGRINS Spectral Archive (RRISA).¹ We gather observations from 28 nights and search the headers of those observations for the targets that were used as telluric standards. This process yields a total of 93 observations of 4 A0V standard stars.

With our standard star observations acquired, we next determine the blaze function from our observations. In sum, we do so by extracting the IGRINS blaze function on an order-by-order basis with an iterative spline-fitting procedure. To remove spectral features due to absorption lines, we first fit a given order with a low-order spline, then divide the data by the fit. We iteratively remove any flux that is 1.5σ below the mean of this fit-divided data until this process converges.² We repeat this process until no other points are removed from the dataset. We then fit a tenth-order polynomial to this absorption-cleaned standard star spectrum, which has the flexibility to capture the behavior of the blaze function.

The initial data products from which we fit the blaze function are normalized to unity, so they

¹<https://igrinscontact.github.io/>

²We find that 1.5σ is adequate to drop out the main absorption features while preserving the continuum

cannot capture differences in throughput between orders. To capture order-dependent throughput variations, we fit a piecewise second-order polynomial to the throughput variations across orders from the [Line et al. \(2021\)](#) dataset.

6.2.1.7 Noise

With all spectral components included, we add shot noise per pixel as an optional, final step. We initially determine the signal-to-noise ratio (SNR) that can be achieved with our data cube’s time steps using the exposure time calculator for IGRINS ([Le et al. 2015](#)). This approach assumes a seeing of $0.6''$ and requires picking a stellar magnitude. For this work, we fix $\text{SNR} = 250$ for one set of simulations (to match the [Line et al. \(2021\)](#)) and $\text{SNR} = 60$ to investigate a low-SNR limiting case. The lower-SNR case effectively simulates observations of a fainter star with the same integration time.

To capture the throughput of the instrument, we scale our simulated flux cube such that the 90th percentile of the flux cube matches the 90th percentile of the [Line et al. \(2021\)](#) counts on a per-order basis—that is, each exposure is assumed to have the same throughput. Finally, we simulate a Poisson draw for each pixel to generate our noisy data.

6.2.1.8 PCA

With our HRCCS data cube simulated, the next step to simulating the HRCCS analysis is to perform signal extraction. Principal components analysis (PCA) is a method for abstracting low-dimensional trends from high-dimensional datasets (e.g., [Jolliffe & Cadima 2016](#)). In the context of HRCCS, it is used to remove the dominant spectral components—the blaze function,

the stellar spectrum, the tellurics—and preserve shot noise and the planetary spectrum (e.g., [Line et al. 2021](#)). Prior to performing PCA (which we implement with single value decomposition, or SVD), we take standard data-cleaning steps to maintain the statistical validity of the operation, such as sending NaNs to 0, mean-centering the input dataset, and dividing it by its standard deviation. We decompose the data matrix using the `numpy.linalg` package, setting the first few singular values to zero. The matrix is then reconstructed without those first few singular components, effectively removing the stationary components in the data cube.

6.2.1.9 The cross-correlation function

With the simulated HRCCS dataset ostensibly cleaned of non-planetary spectral components, `scope` is now able to compare exoplanet atmospheric models to the simulated dataset. This comparison is fundamentally performed via the cross-correlation function (e.g., [Brogi & Line 2019](#))

$$\text{CCF}(t, s) = \frac{\sum^n F(\lambda, t)m(\lambda - s)}{n\sqrt{v_F v_m}}, \quad (6.2)$$

where m is the template model that is cross-correlated against the data, v_F is the variance of the data, v_m is the variance of the model, s is the amount by which the template model has been shifted, and there are n total data points. The data are here mean-subtracted prior to cross-correlation. Because this is a forward-modeling exercise, there exists a template model that exactly matches our input dataset—which would not be the case for a real planetary atmosphere. Despite the presence of this artificial template match, the forward-modeling exercise is still useful because the data is subjected to the various HRCCS data-processing steps that can sculpt the signal.

The above equation describes the CCF at a single Doppler shift. In general, cross-correlation

is performed along an axis—i.e., the resultant CCF is a vector with values corresponding to each Doppler shift of the template model, maximized at the Doppler shift where the model best matches the data. This process allows model comparison to a single observation. HRCCS observations, however, generally involve multiple exposures taken at different orbital phases. The planet therefore has a different line-of-sight velocity in each exposure. HRCCS analyses encapsulate this effect by computing “ K_P – V_{sys} maps,” for planetary orbital velocity K_P and systemic velocity V_{sys} . We follow the standard literature (e.g., [Brogi et al. 2012](#); [Brogi & Line 2019](#)) to compute our K_P – V_{sys} maps:

1. For a given pair of K_P and V_{sys} , calculate the line-of-sight planetary velocity as a function of orbital phase.
2. Doppler shift the template model to the expected velocity at each exposure for this orbital solution.
3. Scale the template with a tunable scale factor (default set to 1) and by the data after removing the contributions from the data’s first N_{PCA} components
4. Perform PCA on the scaled template model
5. Cross correlate each exposure in the data with the corresponding Doppler-shifted template model.
6. Sum the cross-correlation values for this solution across all exposures to arrive at the value for this K_P – V_{sys} pair.

Once the map is completed (as a grid spanning $\pm 100 \text{ km s}^{-1}$ in K_P and $\pm 100 \text{ km s}^{-1}$ in V_{sys} around the expected values), we compute detection significance. There exist multiple ways to

determine significance from CCF maps, such as iterative sigma clipping (Kasper et al. 2021) and the Welsh’s T test (Brogi et al. 2013). In this work, we calculate detection significance simply by subtracting the map by its mean and by dividing by its standard deviation far from the expected peak in the map (with K_P spanning 93.06 km s^{-1} to 172.06 km s^{-1} and V_{sys} spanning -100 km s^{-1} to -11 km s^{-1}), in order to match the approach of Line et al. (2021).

6.2.1.10 Retrievals

The CCF map does not provide quantitative estimates on the properties of a planet’s atmosphere. To calculate these, we must turn to Bayesian inference, which in turn requires a mapping to a likelihood function (Brogi & Line 2019; Gandhi et al. 2019; Gibson et al. 2020). On top of the forward model and statistical sampler required for low-resolution retrievals (e.g., Madhusudhan 2019), high-resolution retrievals require Doppler-shifting of the template model at a given K_P – V_{sys} pair and model processing (i.e., performing PCA on models in the way that they were performed on data; Brogi & Line 2019).

In this work, we perform HRCCS retrievals using the CCF–likelihood mapping derived in Brogi & Line (2019). To avoid mismatch between the CCF template and the planetary model within the simulated dataset, we also use CHIMERA as our forward model.

For our sampling algorithm, we significantly depart from previous works. Past retrievals have used Markov Chain Monte Carlo (MCMC; specifically, `emcee`; Gibson et al. 2020; Foreman-Mackey et al. 2021b; Maguire et al. 2022) or nested sampling (specifically, `PyMultiNest`; Buchner et al. 2014; Brogi & Line 2019; Line et al. 2021) to sample posterior distributions. However, vanilla MCMC can have difficulty in exploring posterior distributions that are strongly degenerate

or multimodal (e.g., [Madhusudhan 2018](#)), and nested sampling algorithms often take long periods of time to converge. In this work, we use the `flowMC` ([Wong et al. 2022](#), Appendix D.1.2) package to accelerate the posterior distribution estimation. This method pairs a local, gradient-based sampler (to effectively characterize local minima) with an adaptive, non-local sampler (to effectively find global minima). Switching to this sampler, in addition to other optimizations (e.g., re-implementing the retrieval in JAX; Appendix D.1.1; [Bradbury et al. 2018](#)) decreased runtime from 8 days on 24 CPUs and 1 GPU (with `PyMultiNest`) to 15 hours on 1 CPU and 1 GPU, or only 5 hours when restricted to local sampling.³

6.2.2 Perturbations to baseline pipeline

Beyond the baseline modeling of Section 6.2.1, we seek in this work to explore higher-order observation-related effects that may appear in HRCCS data. We focus in particular on two such effects: time-varying tellurics and moderate wavelength instability.

6.2.2.1 Modeling tellurics

Absorption by the Earth’s atmosphere is traditionally modeled out of HRCCS datasets based on synthetic modeling of the Earth’s atmosphere (e.g., `Molecfit`; [Kausch et al. 2015](#); [Smette et al. 2015](#)) and/or data-driven approaches (e.g., PCA; [Giacobbe et al. 2021](#)). It is unclear how well the synthetic models capture potential time dependence of observed telluric absorption at telescopes, especially with respect to “microtellurics” (e.g., [Cunha et al. 2014](#); [Wang et al. 2022](#)). It is similarly unclear to what extent the finer details of tellurics are fully removed by existing

³For these performance benchmarks, we used an Nvidia A100-80GB GPU, and we used CPUs on a 4-core Intel Ice Lake node with 1 TB of RAM.




Figure 6.2: The eigenweights of our PCA decomposition for IGRINS standard star observations for a single order (order 5 in the H band). In this case, the first two components clearly covary with airmass, and multiple components exhibit subtle sub-structure as as function of date.

data-driven methods. These questions can be answered with a forward-modeling approach: by reconstructing telluric absorption using real observations, we can assess its impact on HRCCS data with controlled “experiments.”

In this work, we construct a data-driven forward model of tellurics by performing PCA on observed telluric spectra. Then, using these PCA components, we build a forward model to simulate the telluric spectrum as a function of airmass and date. The aim of this simplified model is not to accurately predict the telluric spectrum at future times, but only to yield a reasonably true-to-life representation of the way that the telluric spectrum varies in a subset of IGRINS data, so that we can simulate a realistic representation of telluric variability.

To construct a reasonable estimate of telluric absorption as observed by IGRINS, we begin

again with our set of A0V standard star observations (as scaled by Vega models⁴) from the IGRINS Raw and Reduced Spectral Archive. Some nights contain multiple (up to 6) observations; the median difference between exposures is 36 days. We then mean-scale each dataset and perform PCA on the cube of spectra (with dimensions of 93 spectra by 2048 pixels), keeping the first four eigenvectors and eigenweights. These first four components contain a majority (more than 75%) of our dataset’s variance. Capping the components at four risks losing some higher-order telluric variability signals that may still impact the planet detection, but given the quality of telluric standard star data available we didn’t want to risk imprinting photon noise by going to higher numbers of components, so this should be regarded as an incomplete test of telluric variability’s effects.

From each observation, we extract the spectrum, the airmass, and the timestamp. Humidity and precipitable water vapor (PWV) are known to correlate with singular vectors in HRCCS data analysis; PWV in particular can vary on the order of minutes, potentially driving strong time-variability in HRCCS signals ([Chiavassa & Brogi 2019](#); [Smith et al. 2024a](#)). Rapid changes in PWV are not incorporated in our tellurics model and may cause more complex variability than we account for. In light of this, we also test the impact of the ambient humidity recorded with each observation.

These eigenvectors and corresponding eigenweights describe the primary variation of the telluric absorption at each subsequent exposure. We assume that the primary time-dependence of this absorption is driven by airmass, followed by variations in time. We assess this hypothesis with a linear regression for these data on airmass and time. We leave out 10% of the initial dataset (12

⁴Upon inspection near the Balmer lines, we find that this Vega model correction imprints limited residuals, resulting in a relatively pure telluric spectrum.




Figure 6.3: Simulated telluric spectra, zoomed in on a narrow spectral region to display the effect of changing airmass (top) and date (bottom). Airmass is clearly the first-order effect, but some lines are particularly time-variable.

of 120 spectra) to assess the quality of our regression.

We find that the telluric absorption observed by IGRINS exhibits both a trend in airmass and a trend in time. To select the best representation of eigenweights with airmass and time, we predict the left-out 10% of spectra and seek to minimize the χ^2 of those predictions. Of the models explored to characterize our eigenweights, the best-performing model scales the eigenweights linearly in airmass and quadratic in time. We do not find that including humidity markedly increases the performance of our model (potentially because of our poor time sampling), so we continue with a model that assumes only airmass and time.

The results of our tellurics forward modeling experiment are shown in Figs. 6.2–6.3. The first two eigenweights in particular exhibit linear trends in airmass. The corresponding shifts in the contributions of the first two eigenvectors with airmass produce deepening absorption lines at

higher airmass.

While Fig. 6.2 shows light clustering in eigenweights as a function of time, it is not immediately evident from this figure how the full spectral model changes with time. The spectrum-level changes with time are more evident in Fig. 6.3, with two classes of spectral regions emerging: static lines and lines that vary on the scale of months. We verify that the two strongest lines with time variability in the plotted wavelength region are in regions of strong water opacity (Polyansky et al. 2018).

As an extension of our tellurics model, we additionally include a forward model that lengthens the modeled timescale (one night) by a factor of 100. This approach effectively increases the variability of the tellurics and approximates a multi-night observing campaign.

Our data-driven approach broadly matches the line positions and depths of simulations with ATRAN (Lord 1992; Fig. 6.4), a radiative transfer telluric modeling code that takes as input the altitude, latitude, water vapor overburden, and zenith angle of observations. Differences in the absolute continuum level are likely due to our approach of rescaling the telluric standard stars prior to performing PCA. Because the continuum level of the flux is lost by HRCCS data-processing (e.g., Brogi & Line 2019), these differences in the continuum between telluric approaches are not expected to impact inference.

6.2.2.2 Telluric removal: airmass-detrending and Molecfit

Beyond our baseline PCA procedure for telluric-removal, we explore alternative telluric-removal methods to understand the sensitivity of our results to such approaches.

Detrending datasets by airmass has a long history in HRCCS (e.g., Snellen et al. 2010). We




Figure 6.4: A comparison of our data-driven telluric reconstruction against ATRAN (Lord 1992) radiative transfer simulations over one IGRINS order.

implement it in this work by fitting a second-order polynomial to our dataset to each pixel as a function of time, motivated by the expected change in line depth as the telescope points through longer columns of air. More robust approaches in the literature generally follow this step with higher-order regression on the residuals of the fit and filtering to remove additional substructure (e.g., Brogi et al. 2012). While our limited approach requires less tuning and is adequate to understand changes in significance relative to changes in simulated data, the approach is likely not reliable enough to be applied on real datasets because we do not perform additional filtering steps.

To further assess the impact of PCA on detection significance, we performed tests where we apply matrix standardization prior to the PCA, multi-linear regression afterward, and dividing out the PCA fit instead of subtracting (e.g., Brogi et al. 2023). Doing so emphasizes relative errors (therefore overweighting, e.g., tellurics).

Finally, we test correction of telluric absorption with `Molecfit`. In contrast with the more empirical approach of airmass-detrending, `Molecfit` models the absorption of the Earth’s atmosphere by fitting radiative transfer-derived spectra to (subsets of) data. Based on the absorption

noted in [Smette et al. \(2015\)](#), we fit our `Molecfit` runs with respect to H_2O , CH_4 , and CO_2 . When available, we pin input values (e.g., with relative humidity) to values recorded by the [Line et al. \(2021\)](#) IGRINS observations. Otherwise, we resort to `Molecfit` defaults (e.g., with ambient temperature).

6.2.2.3 Stability of wavelength solution

Unstabilized spectrographs, such as IGRINS, are susceptible to changes in the wavelength solution over the course of a night. For instance, [Brogi et al. \(2023\)](#) report wavelength solution variability of $\pm 0.46 \text{ km s}^{-1}$ over the course of their IGRINS observations. As described in Section [6.2.1](#), `scope` can incorporate inaccuracies in the wavelength solution to determine their effect on inference from HRCCS data. We implement this effect by incorporating an additional Doppler shift to our data cube, unrelated to astrophysical sources.

To motivate the functional form and degree of wavelength instability in our model, we assess the wavelength instability of the [Line et al. \(2021\)](#) dataset. This instability found in the dataset will then be included in our HRCCS forward models. We seek in this dataset a simple shift and drift (i.e., a changing shift in time) of the wavelength solution in time as recorded by IGRINS. To assess this variation, we try two approaches: fitting for the position of a single line (confirmed non-stellar spectral feature that occurs across all observations), and for the position of all telluric lines simultaneously by cross-correlating against an ATRAN telluric model ([Lord 1992](#)). We chose this pair of approaches to explore limiting cases (i.e., considering shifts of one line vs. all lines at once).

Our line-fitting experiment indicates that, over the course of a single night, IGRINS data

exhibit significant shifts in wavelength solution. Our cross-correlation drifts are on the order of 0.3 km s^{-1} , whereas the single-line drifts are on the order of 0.1 km s^{-1} . These shifts are on the order of a tenth of a pixel in resolution, and therefore may play a role in subsequent cross-correlation analysis. The similarity between these two drift estimates over time indicates that we have identified a real effect in these data. Notably, this level of precision is on the order of the wavelength solution precision reported by [Brogi et al. \(2012\)](#) on CRIRES ([Ulrich Käufel 2008](#)), using a wavelength solution approach similar to [Line et al. \(2021\)](#). When injecting wavelength instability into the pipeline model, we adopt a drift of 0.3 km s^{-1} .

6.3 Results

6.3.1 K_P-V_{sys} maps

We begin by presenting the K_P-V_{sys} forward-modeling maps, as these data products are the simplest scientifically meaningful units that `scope` produces. These maps indicate the significance of the planetary detection and can additionally reveal spurious signals.

6.3.1.1 Varying the number of PCA components

In this section, we report how the planetary spectrum detection significance changes in the presence of different levels of noise to test its sensitivity to PCA treatment.

[Fig. 6.5](#) shows the effect of increasing the number of PCA components on the CCF detection significance from 0 to 4. Of first note is that cross-correlating the planetary spectrum against a pure planetary data cube—an unphysical test case with no star, tellurics, or blaze function—yields a very high-significance detection even when the photon noise level is relatively high (i.e., in the




Figure 6.5: The effect of removing different numbers of PCA components from HRCCS data on the CCF detection significance in our baseline results. This is calculated in noiseless simulations (top left), $\text{SNR} = 250$ simulations (top right), and $\text{SNR} = 60$ simulations (bottom left) with 10 different noise realizations. Including more spectral components (e.g., the stellar spectrum) severely decreases the planet’s detection significance, but a small number of PCA components can recover the planet detection about the detection threshold. These significance values are compared to a commonly accepted HRCCS detection threshold ($\text{SNR}=4.5$; e.g., [Cabot et al. 2019](#)).




Figure 6.6: The spectral components learned by four PCA components from a simulated HRCCS dataset that only includes the planetary signal and instrumental blaze function. The PCA algorithm is able to approximate the diagonal trace of the planetary atmosphere with a superposition of a few strong stationary components. When other spectral components are present (including photon noise), PCA is not able to approximate the planetary signal, as it is most sensitive to the time-invariant signals with greater variance across wavelengths.

SNR=60 simulation).

Once more spectral components are included, the detection significance of the planetary atmosphere is severely diminished, as those components have features that dominate over planetary ones. In all cases, at least one PCA component is required to bring the detection significance above the detection threshold (commonly held to be about SNR=4.5; Cabot et al. 2019).

When modeling just the planet and star combination without noise, the detection significance begins to diminish when more than 1 PCA component is used. This reduction of the planetary detection significance is possible even though PCA cannot capture the moving planetary signal as a single component. As we show in Fig. 6.6, with enough freedom and without many other spectral components, multiple components can together approximate the planetary signal as a sum of stationary components. This pattern, however, is not present in a realistically noisy case (SNR=250).

While removing a low number of PCA components does not severely impact planetary detection, we do see reduced detection significance for a very high numbers of removed components. In particular, we can produce spurious non-detections when many components are removed; that




Figure 6.7: A set of spurious HCN signals in our baseline simulations. The maxima of these K_P-V_{sys} are denoted with a white “x.” While HCN is not included in this set of models, its signal can be found in cross-correlation when using ATRAN tellurics (top left), data-driven tellurics (top right), with low photon noise (top row), with high photon noise (bottom row), and in different noise realizations (bottom left).

is, as the PCA algorithm is provided more flexibility, it can create artificial correlated structures (peaks and troughs) in the K_P-V_{sys} maps. As a limiting example, when we remove the maximum possible number of PCA components (79), the planet signal entirely disappears. Despite these dangers, the experiments presented here indicate that using a reasonably low number of PCA components on realistically noisy data does not meaningfully lessen the detection significance.

6.3.1.2 Instrumental and telluric variability

We next tested two effects that could potentially interfere with atmospheric characterization with HRCCS: variability in the instrumental wavelength solution and variability in the tellurics.

By forward-modeling the wavelength instability that we recover from IGRINS observations, we find that such instability introduces essentially no decrease in detection significance in the low-noise case and a 0.5 SNR decrease in detection significance in the high-noise case. Similarly, our time-variability exercise indicates that the tellurics as we model them do not vary enough within a given night to impact HRCCS observations. The smooth changing of our line depths is simply not rapid enough to induce strong changes in K_P - V_{sys} maps. As a limiting-case investigation, our longer-timescale telluric variability yields a decrease in SNR of 0.8.

6.3.1.3 Barycentric motion

In the high-SNR regime (photon SNR=250), we find that including barycentric motion reduces the CCF SNR by about 1 on average. In the noiseless regime, there is negligible difference between the two; in the low-SNR regime (photon SNR=60), there is also negligible difference between the two.

6.3.1.4 Spurious signals

Our results are able to produce lower-significance signals related to species that are not included in the forward model. We focus in particular on HCN, motivated by the idea that the gas's intrinsic line strength and periodicity in wavelength might make it more susceptible to spurious signals than other species (Zhang et al. 2020). To test this, we generate a dataset that includes all best-fitting values of species from Line et al. (2021) aside from HCN. Despite its exclusion in the forward model, `scope` measures a spurious HCN signal at a significance of SNR=3.5 in the low-noise case, albeit at the wrong K_P value (Fig. 6.7). Note that this peak is matched by a trough

of roughly the same magnitude in this simulation.

We also note spurious HCN signals when using different sets of tellurics (trading the ATRAN radiative transfer tellurics for our data-driven tellurics forward model), different noise levels (in both the low and high photon noise case), and in different noise realizations. The structure of the K_P-V_{sys} varies when the tellurics model and the noise instance change, and the strength of the spurious signal changes with all model changes. The spurious signal is robust to telluric treatment and noise instance. Additionally, the K_P-V_{sys} structure in a spurious signal scenario is altered by both the noise instance and tellurics simulation.

6.3.2 Likelihood sampling

We have shown (Section 6.3.1.1) that the planetary detection significance is severely weakened when other spectral components are introduced to a simulated HRCCS dataset.

As a next step, we move beyond the K_P-V_{sys} maps by optimizing the likelihood function over those two velocity terms and a linear scale factor, a . The scale factor (e.g., [Pino et al. 2020](#); [Brogi et al. 2023](#); [van Sluijs et al. 2023a](#)) stretches the planetary spectrum, which may accommodate associated stretching from data processing steps such as PCA. This likelihood sampling approach is analogous to a full atmospheric retrieval, except that all planetary atmospheric parameters are held constant—only the velocity terms and scale factor are fit here. This task is relatively computationally intensive; we perform our fitting with the `emcee` ([Foreman-Mackey et al. 2021b](#)) package, parallelizing across 25 CPUs with MPI to avoid consuming GPU resources.

We find that including four PCA components does not impact the retrieved scale factor (Fig. 6.8). Our result is consistent at both noise levels considered, with the width of the poste-

riors simply increasing by a factor roughly equal to the noise increase (roughly 3).

6.3.3 Retrievals

As the final step in our suite of experiments, we perform full atmospheric retrievals on our simulated datasets to assess whether any biases are introduced.

6.3.3.1 Baseline comparison to [Line et al. \(2021\)](#)

Given the complexity of the algorithms used, we firstly benchmark our simulations against the [Line et al. \(2021\)](#) observational results to validate `scope`.

Broadly, we are able to reproduce the basic results of [Brogi & Line \(2019\)](#): the retrieval framework adequately recovers the input parameters. Furthermore, our modeled parameter constraints are about as precise as the [Line et al. \(2021\)](#) constraints. For instance, our constraints on water ($-4.08^{+0.16}_{-0.13}$ dex) and V_{sys} ($0.09^{+0.81}_{-0.86}$ km s $^{-1}$) approximately match the corresponding water ($-3.93^{+0.10}_{-0.09}$ dex) and V_{sys} ($-7.71^{+5.27}_{-5.41}$ km s $^{-1}$) constraints from [Line et al. \(2021\)](#) in both value and uncertainty (Fig. 6.10 with an 8 km s $^{-1}$ offset applied in our simulations).

6.3.3.2 Comparison to likelihood sampling approach

Our baseline retrieval shows a departure from the results in Section 6.3.2: while the scale factor is still consistent with 0, there is an asymmetry in the scale factor posterior (Fig. 6.9). This asymmetry coincides with the degeneracy between this parameter and two of the temperature–pressure (TP) profile parameters that are themselves asymmetric.




Figure 6.8: Results of an MCMC fit to the velocity parameters (in km s^{-1}) and scale factor for simulated HRCCS data at two photon noise levels. This model used four PCA components to remove the stellar, telluric, and instrumental features in the simulated data. Using PCA in this experiment did not bias the scale factor toward smaller values, indicating that using PCA alone does not scale the planetary signal in HRCCS data.




Figure 6.9: Gas and velocity constraints from our PyMultiNest high-resolution atmospheric retrieval. Including strongly time-dependent tellurics (maroon) yields distinct posteriors from our time-dependent telluric (teal) model.

6.3.3.3 Sensitivity to variable tellurics

Our previous experiment in Section 6.3.1.2 showed that including time-dependent tellurics decreases the significance of a planetary detection. How does such an effect impact retrieved quantities?

When using our longer-timescale telluric variability, the abundances of our retrieved emission spectra are correctly recovered, but other quantities are strongly biased. The K_P is biased to $-3.87^{+1.18}_{-1.14}$ km s⁻¹ from $2.44^{+1.41}_{-1.35}$ km s⁻¹ (Fig. 6.9). The first log pressure point is similarly biased to $-1.76^{+0.43}_{-0.36}$ from $1.06^{+0.95}_{-0.86}$. Finally, gas parameters that were previously bound—water and carbon dioxide—become non-Gaussian and/or unbound.




Figure 6.10: Gas constraints from our baseline retrieval on simulated data of WASP-77A b. We are able to recover gas constraints about as well as [Line et al. \(2021\)](#), with bounded constraints on H_2O , CO, and the CO isotope ratio in the lower-noise case ($\text{SNR}=250$). The higher-noise case ($\text{SNR}=60$) substantially changes the likelihood surface for the gas constraints shown, with less of an impact on velocity and temperature–pressure parameters.

6.3.3.4 Sensitivity to noise levels

When we increase the noise in our modeled HRCCS datasets, it appears that the gas constraints (Fig. 6.10) are more affected by increased photon noise than the velocity constraints are (Fig. 6.8). While the posteriors on the velocity parameters simply become broader with increased noise, the posteriors of the gas abundances become significantly warped—and in some cases, they become unbounded when they were previously bounded (i.e., with H_2O and CO) and become confidently incorrect (i.e., with HCN).

6.4 Discussion

6.4.1 Confirmation of HRCCS principles

As evidenced by our K_P – V_{sys} maps, likelihood sampling, and retrieval results, our simulation reassuringly confirms the basic premise of previous HRCCS proofs-of-concepts. The planetary signal can be recovered in cross-correlation despite its intrinsic spectral weakness; some stellar and telluric removal is required to do so; and the velocity of the planet is imprinted on the data.

The robustness of our results persists even in the presence of mild wavelength instability. This behavior may be expected; the instability measured occurs on the sub-pixel level, so the net effect is to smear out the planetary signal at the sub-pixel level.

6.4.2 Noise and time-varying tellurics can strongly affect HRCCS inference

While our baseline, more optimistic models reaffirm HRCCS analysis techniques, we identify two potential complicating factors: noise and time-varying tellurics.

6.4.2.1 On noise

Extrapolating from the K_P – V_{sys} maps and likelihood sampling experiments, it would appear that increasing the amount of photon noise simply broadens the uncertainty on all inferred parameters. However, a full retrieval reveals that the high photon noise case gives poor constraints on gas abundances. Additionally, when photon noise is increased, the upper bound on HCN becomes less informative. In light of the retrieval results, the standard for what constitutes an “acceptable amount of noise” therefore appears to be higher for retrievals than for the other HRCCS data products.

One reason for the differing gas and velocity posteriors under noise could be that the velocity parameters are sensitive to the *position* of spectral lines, which photon noise does not change, whereas the gas abundances are sensitive to relative line ratios and shapes (e.g., [Brogi & Line 2019](#); [Gandhi et al. 2023b](#)), which are perturbed by increased photon noise. Additionally, *all* the planetary lines contain the planet’s bulk velocity information (aside from differences due to atmospheric winds and rotation), whereas inferences from gas abundances of additional species are restricted to subsets of the planetary lines. Therefore, while low-SNR targets may be detected in cross-correlation and may have accurately measured velocities, measuring gas abundances in this regime proves more challenging.

6.4.2.2 On time-varying tellurics

Our K_P – V_{sys} maps demonstrated negligible differences for our standard time-varying tellurics, and these differences were still modest when using the longer-timescale telluric variability. Similar experiments with airmass-detrending and `Molecfit` yield stronger deviations (from

SNR=10 to SNR=2 for the former).

The exact reason for these differences is unclear. We modeled our time-dependent tellurics using PCA from the start. Therefore, the PCA-based telluric-removal method may be able to identify and remove this substructure more effectively because of the shared PCA method. Alternatively, the other two approaches as we implement them may be fundamentally less well suited for removing telluric absorption. *This experiment indicates that realistic tellurics likely do not impact IGRINS observations in previously unaccounted-for ways—but that higher telluric variability can decrease detection significance.*

In terms of ramifications, the majority of HRCCS emission studies analyze data from a single night, and they therefore would not be strongly affected by mildly varying tellurics. Additionally, multi-epoch HRCCS observations that aim to keep the planetary spectral lines constant and remove all other sources as variable (e.g., [Lockwood et al. 2014](#); [Buzard et al. 2020](#); [Finnerty et al. 2021](#)) would in principle be unaffected, as they can treat tellurics as a time-variable process to be removed. However, combinations of multiple nights of data with strongly shifting planetary spectral features may be susceptible to additional error from even mildly varying tellurics over a long enough baseline. This issue could potentially be addressed by performing telluric removal for a given night at a time, then summing the CCF values after the fact (e.g., [Kasper et al. 2021](#)).

As described in Section 6.3, we observe a strong bias in temperature–pressure parameters under longer-timescale telluric variability. The reason for this result may be that the retrieval algorithm seeks to compensate for improperly removed telluric signals by mutating the planetary spectrum to account for those spectral components’ residual spectral features. This framing also accounts for the K_P bias—the presence of spectral features from the star and tellurics at a different velocity from the planet significantly biases the planetary V_{sys} . *This experiment demonstrates that*

the velocity parameter and gas abundances can be biased by tellurics that are not fully removed; the same intuition should apply to transmission spectroscopy.

While we are able to assess the effect of time-dependent tellurics on the scale of days, a caveat is that our observations do not have a short enough cadence to capture minute- or hour-scale telluric variability (Section 6.4.5).

6.4.3 Effects of PCA on the planetary signal

When all realistic spectral components are included, increasing the number of PCA components does not *decrease* the planetary significance for up to four PCA components. While the noise-free, pure-planet case can be learned by PCA (Fig. 6.6), simply adding a competing noise term disrupts the coherence of the planetary signal enough so that the planetary signal does not leak into the higher-variance PCA terms.

In our telluric-dominated near-infrared regime, when we perform matrix standardization and multi-linear regression, our signal sizes decrease, with a different number of PCA components now preferred. We anticipate that this approach would not strongly affect retrievals given the reprocessing of the model.

Furthermore, when holding atmospheric parameters constant, the scale factor is fit to be Gaussian around 1 with four PCA components. While the atmospheric signal is certainly stretched by the PCA operation, our result implies it is not stretched in a way that cannot be corrected by model reprocessing in this regime. *This experiment indicates that PCA alone does not reduce CCF significance when using a low number of principal components.*

However, the limiting-case PCA application (with 79 components removed) shows that re-

moving too many PCA components at some point becomes detrimental to the planetary signal. This result indicates that as demonstrated by, e.g., [Cheverall et al. \(2023\)](#), there exists an *optimal* number of PCA components to remove for a given dataset. While removing some PCA components is beneficial, removing too many clearly prohibits detecting the planet. This point is expected from the perspective of the definition of PCA—together, all the PCA components explain all the variance in a dataset, so at some point much of the planetary signal will be encompassed by some of the PCA components. In our set of simulations, the optimal number of PCA components appears to be 4. This number, however, is specific only to our simulations as they are modeled, and real data with more complexity (e.g., additional sources of systematic uncertainty) may require a larger number of components to be removed. Again, this finding motivates the search for metrics that characterize the optimal number of components to remove (e.g., [Cheverall et al. 2023](#)).

6.4.4 On spurious results

Our modeling results are able to produce spurious signals for HCN at a different K_P than input (Section 6.3.1.4). A significant signal, even when none is injected, appears across different specific tellurics models and noise realizations, but its location can vary. Because of this, it seems that spurious signals are a result of a combination of genuine line confusion (i.e., the cross-correlation template picking up on telluric signal) and specific noise instances.

Unfortunately, preventing this spurious signal would not be as simple as limiting the retrieval to the expected K_P . Given the known strong correlation between these two parameters, strongly limiting one without biasing the other may be difficult. Furthermore, the expected K_P for a given observation may simply be incompletely known.

Instead of restricting K_P to a single value, one potential way to address off- K_P spurious signals would be to enforce a wide uniform prior on K_P centered on the expected value, thereby penalizing strongly differing K_P . The width of this prior could be determined based on the maximum deviation expected from rotation alone (Wardenier et al. 2023), with an additional allowance for measurement uncertainty in K_P . This approach can be readily incorporated into retrieval frameworks, and it may be possible to self-consistently include them in K_P - V_{sys} map searches for individual gases.

It may be that our spurious signals are driven by poorly corrected tellurics. If so, masking out these pixels (e.g., Boucher et al. 2023; Kanumalla et al. 2024) may decrease the impact of this noise source on detection significance. As a test, we mask out the noisiest pixels in our cross-correlation step. We determine which pixels to mask via the inverse survival function of the normal distribution, Brogi et al. (2023). Doing so allows adaptive thresholding based on the noise characteristics of the dataset. We verify that even when we mask out these selected pixels, we can generate spurious signals of HCN up to a significance of 4.4 in our baseline set of simulations (that is, without the strongly varying tellurics).

It may be that HCN is more likely than other species to generate spurious signals, because of its few, strong, wavelength-periodic lines (Zhang et al. 2020). Robustly testing this hypothesis would be instructive for interpreting HRCCS inferences and is a promising avenue for future work—especially in the context of the potential tension between HRCCS and JWST inferences of HCN in the atmosphere of the planet HD 209458 b (Hawker et al. 2018; Giacobbe et al. 2021; Xue et al. 2024).

6.4.5 Caveats

While our work incorporates many of the expected first-order astrophysical and instrumental sources in HRCCS datasets, we do not include an exhaustive set. Our limitations can be grouped into five camps:

- *limited data reduction and processing steps.* There are steps that are often required to turn a raw data product into science-ready HRCCS data that we do not implement here. For instance, we do not perform bad pixel correction, throughput correction, masking, or high-pass filtering (e.g., [Brogi et al. 2023](#); [van Sluijs et al. 2023a](#)). We choose not to implement these processes here because they are not assumed to be primary drivers of error in HRCCS data analysis and they are more instrument-specific.
- *telluric observation cadence.* We are only able to assess the longer-term variability in tellurics, as the observations of telluric standard stars are not on a short enough cadence to capture, e.g., minute-scale variability. This gap motivates dedicated observing programs for measuring short-timescale telluric variability.
- *astrophysical simplicity.* We have furthermore not considered effects such as the multidimensionality of the planetary atmosphere, which strongly sculpts HRCCS data ([Kempton & Rauscher 2012](#); [Showman et al. 2013b](#); [Zhang et al. 2017](#); [Flowers et al. 2019](#); [Beltz et al. 2020](#); [Savel et al. 2021](#); [Wardenier et al. 2021](#)), and higher-order effects including condensation, non-local thermodynamic equilibrium, and stellar variability. The broader question of model uncertainty, though certainly intertwined with HRCCS datasets, is out of scope for this work.

- *generalizability.* This work models one planet observed by one instrument. The dataset that we compare our simulations to, furthermore, provided a relatively high SNR detection and clear inference. Hence, the simulations here may be more optimistic than they should be in a general case for other planets observed with other instruments. Additionally, while the trends in significance as a function of, e.g., number of PCA components may generalize to other datasets, the exact values will of course not.
- *model-on-model simulations.* Our approach inherently assumes an underlying match between the simulated dataset and the forward spectral model. This approach yields unrealistically high detection significances because in real observations, the planetary spectrum is likely not a perfect fit for any model. Furthermore, our simulations include no correlated noise, and other simulation work indicates that red noise and more complex throughput calculations can disrupt the expected scaling of SNR with $\sqrt{N_{\text{lines}}}$ (e.g., [Birkby 2018](#); [Lopez-Morales et al. 2019](#); [Currie et al. 2023](#)).
- *separation of telluric lines and planetary lines.* Based on the planetary orbit, the emission spectra that we simulate are well separated from the telluric spectral features. This separation would not be the case for transmission spectra, in which the radial planetary velocity intersects with the stellar velocity, and the planet’s velocity may change comparatively less over the course of observations. Therefore, simulations of HRCCS transmission spectra may yield different results from those shown here, motivating future work.

6.5 Conclusions

The complex data representation of high-resolution cross-correlation spectroscopy can make it difficult to understand the uncertainty budget of an inference. In this work, we aim to clarify some of more uncertain aspects of HRCCS data processing and analysis. With `scope`, a newly developed forward model for HRCCS datasets, we find that:

- We are able to reduce the computational burden of retrieval algorithms by an order of magnitude through implementing numerical speed-ups with JAX and more efficient statistical sampling with `flowMC`.
- Under our baseline approach, standard HRCCS data processing algorithms work as expected. In our modeling, using a reasonable number (4) of PCA components does not destructively impact the planet signal.
- These results are also robust against mild tellurics as we infer that IGRINS observes them within a single night. This is additionally the case for wavelength instabilities at the level evident in IGRINS data after standard corrections have been performed.

Pushing our forward model beyond the baseline cases reveals some HRCCS failure modes.

- While the planetary velocity parameters are robust against increased photon noise, gas constraints deteriorate strongly into non-Gaussian posteriors.
- When strongly time-varying tellurics are improperly removed, gas abundances are biased. The velocity parameters are impacted in this case as well, yielding discrepant Doppler shifts. This bias may arise for multi-epoch observations if each epoch is not cleaned separately.

- We search for and are able to produce spurious HCN signals at an incorrect planetary K_P at up to SNR=4.4.

The full likelihood surface of these types of observations remains fundamentally underexplored. Further investigations of the dependence of HRCCS on observational setup, analysis techniques, and more complex physics will be required to completely flesh out the strengths and limitations of this powerful technique. Doing so will allow more involved planning and interpreting of HRCCS datasets. By open-sourcing `scope`, we hope to provide the community with a tool that will enable these investigations.

Chapter 7: Precise Determination of the Metallicity and C/O of WASP-39b From a Single JWST Instrument Mode with Phase-Resolved Cross Correlation Retrievals

7.1 Introduction

JWST ([Gardner et al. 2006](#)) is already delivering on one of its promises: to provide impactful, information-rich measurements of transiting exoplanet atmospheres (e.g., [Ahrer et al. 2022](#); [Feinstein et al. 2023](#); [Grant et al. 2023a](#); [Bell et al. 2024](#)). In high signal-to-noise JWST transmission spectra, detections of spectral features arising from atmospheric gases are incontrovertible ([Coulombe et al. 2023](#); [Program 2023](#)), revealing exoplanetary processes in unprecedented detail ([Tsai et al. 2023](#); [Challener et al. 2024](#)). Underpinning this success is JWST’s capacity to detect a wide variety of molecules, many of which were not accessible with previous facilities, as well as its ability to robustly constrain both atmospheric metallicities ($[M/H]$) and carbon-to-oxygen ratios (C/O).

Measurements of metallicity and C/O with JWST have been highly anticipated because of their purported role as signposts of planet formation ([Öberg et al. 2011](#); [Madhusudhan 2012](#)). Planets’ assembly in relation to snowlines in protoplanetary disks should have imprints on their composition, especially via abundance ratios (e.g., [Piso et al. 2016](#); [Hobbs et al. 2022](#)). Accurately constraining both $[M/H]$ and C/O requires a full spectroscopic accounting of the key carbon and oxygen reservoirs in exoplanet atmospheres (typically H_2O , CO_2 , CO , and CH_4 for hot and warm jovians). Thus, the most informative JWST datasets typically combine spectra from multiple observing modes to compile the full molecular inventory. The primary advantage of the multi-instrument approach is that broader wavelength coverage can include contributions from multiple

gaseous absorbers, some with multiple absorption bands observed, which aids in breaking degeneracies between molecular abundances and aerosol properties (e.g., [Benneke & Seager 2013](#)). Both simulations ([Batalha & Line 2017](#)) and observations ([Heinke et al. 2026](#)) have shown that the combination of the NIRISS/SOSS and NIRSpec/G395H observing modes, spanning $\sim 0.5\text{--}5\ \mu\text{m}$, provides the best constraints on [M/H] and C/O. Single-mode observations, lacking such spectral range, can yield constraints on atmospheric chemistry that are otherwise surprising and difficult to reconcile with expectations from planet formation theory ([Bean et al. 2023](#); [Meech et al. 2025](#); [Claringbold et al. 2026](#)).

In this Letter, we challenge the idea that multiple observing modes are required to produce robust measurements of [M/H] and C/O by leveraging novel analysis techniques. At the same time, we are able to address another outstanding concern with JWST data analysis related to the detectability of the CO molecule in exoplanet atmospheres, which we illustrate through the example of WASP-39 b ([Faedi et al. 2011](#)). CO was detected at high significance (7σ) in WASP-39 b’s NIRSpec/PRISM transmission spectrum ([Rustamkulov et al. 2022](#)), in line with predictions from equilibrium chemistry (e.g., [Stock et al. 2018](#)). Yet it was not found in the higher spectral resolution NIRSpec/G395H data of the same planet ([Alderson et al. 2023](#)), presenting a puzzle.

The solution is found by considering the behavior of the fundamental CO band at $4.6\ \mu\text{m}$. In typical binned low-resolution JWST data, the band appears wide and weak, despite the intrinsic strength of the individual ro-vibration lines that comprise the feature. In two follow-up efforts, [Grant et al. \(2023b\)](#) and [Esparza-Borges et al. \(2023\)](#) both showed that CO can be detected in the same G395H data that produced the [Alderson et al. \(2023\)](#) non-detection, but at *native spectral resolution*. These detections were both enabled by specialized analysis techniques tailored to detect the high-resolution peaks in the CO opacity function. When unbinned, the band structure

of CO provides additional spectral correlation that can be distinguished from noise. The mystery is therefore resolved — CO is indeed present in WASP-39 b’s atmosphere, but standard analysis techniques making use of *binned* G395H data do not provide enough sensitivity to the molecule. The implicit tradeoff is that while binned G395H data could not reach the S/N of PRISM to constrain the CO pseudocontinuum, the lower-S/N, higher-resolution unbinned G395H data were more sensitive to the strong and regular wavelength dependence of CO.

The difficulty of detecting CO is not limited to WASP-39 b. Few papers have claimed detections of CO in hot Jupiters using JWST data (e.g., [Fu et al. 2024](#); [Gapp et al. 2025](#); [Meech et al. 2025](#)). Even fewer have robust, bounded constraints on the abundance of CO. This paucity is despite the fact that CO is expected to be the dominant carbon carrier in hot Jupiter atmospheres ([Moses et al. 2011](#)).

We offer a way around this lack of sensitivity, taking particular interest in the [Esparza-Borges et al. \(2023\)](#) approach: cross-correlation. The authors detected CO at 6σ with this technique, with the significance increasing to 7.5σ when including minor CO isotopologues. Follow-up work demonstrated that the cross-correlation approach can also be successfully applied to other molecules, detecting an inventory of molecules consistent with standard analyses (i.e., H₂O and CO₂; [Esparza-Borges et al. 2025](#)). These results imply that existing JWST datasets genuinely contain additional signal that traditional techniques are missing.

The [Esparza-Borges et al. \(2023\)](#) method, while innovative, is limited to molecular *detection*; it does not measure the abundances of gases or quantify their uncertainties. Doing so would maximize quantitative constraints on atmospheric chemistry — and through chemistry, potentially planet formation (e.g., [Öberg et al. 2011](#); [Mollière et al. 2022](#); [Feinstein et al. 2025](#)). Fortunately, the community of ground-based exoplanet atmosphere observers has spent the last decade address-

ing a similar problem of extracting the signatures of unseen absorbers in high-resolution spectra and quantifying their abundances (e.g., [Brogi & Line 2019](#); [Gibson et al. 2020](#)). By mapping the cross-correlation to a likelihood function, models and data can be compared in a Bayesian framework, connecting cross-correlation techniques to the well-established discipline of atmospheric retrievals ([Madhusudhan & Seager 2009](#); [Benneke & Seager 2012](#)). This approach has the added benefit of explicitly revealing degeneracies between molecules and aliases that could otherwise hide in cross-correlation.

Another point of note is that most analyses of JWST transmission spectra to-date, including the cross-correlation studies of WASP-39 b by [Esparza-Borges et al. 2023](#), have made use of the *time-integrated* transit. That is, the analyses are performed by first fitting the full U-shaped transit to extract a single value — the transit depth — at each wavelength. Subsequent interpretation then occurs on this time-integrated spectrum, which is typically referred to as “the” transmission spectrum. Ground-based high-resolution studies of exoplanet atmospheres typically follow a different approach, in which the *phase-resolved* data are maintained throughout the analysis and are fitted directly (e.g., [Gandhi et al. 2023b](#); [Prinoth et al. 2023](#); [Boldt-Christmas et al. 2025](#)). The reason for this is that the cross-correlation procedure detects exoplanets’ signals in velocity space. As hot Jupiters such as WASP-39 b transit, their projected orbital velocities shift by $\sim 10 \text{ km s}^{-1}$, which is a substantial fraction of the G395H pixel scale along the spectral axis (74 km s^{-1} at native pixel resolution). It is plausible, then, that (effectively) integrating over the time axis may smear out signal that would otherwise be accessible in phase-resolved spectra with cross-correlation-based methods. Recent phase curve analyses ([Sing et al. 2024](#); [de Regt et al. 2026](#); [Ouyang et al. 2026](#); [Snellen et al. 2026](#)) imply that this notion may hold for JWST data, as well.

This Letter presents the first phase-resolved cross-correlation retrieval on JWST transmission

spectra. We apply a modified version of techniques previously performed only on ground-based exoplanet spectra to analyze the NIRSpec/G395H spectrum of WASP-39 b. By focusing on the G395H instrument mode, we endeavor to accurately and precisely measure both $[M/H]$ and C/O for this benchmark hot Jupiter. This aim is possible because all of the major C- and O-bearing molecules provide considerable absorption across the G395 bandpass, and the cross-correlation approach has improved ability to resolve degenerate solutions through its sensitivity to the many individual ro-vibration lines that comprise each molecular band. This is in comparison to traditional low-resolution techniques, which for example only fit for a *partial* water absorption band in NIRSpec/G395 measurements, leading to the possibility of inaccurate abundance inference due to oversimplified cloud parameterizations or detector edge effects.

The remainder of this Letter is structured as follows. Section 7.2 describes our methodology, with particular focus on the different steps we take with respect to “traditional” JWST retrievals. In Section 7.3, we apply our techniques on archival NIRSpec/G395H observations of WASP-39 b (Alderson et al. 2023). We discuss the robustness of these results and their implications in Section 7.4. Finally, we conclude and remark on future directions in Section 7.5.

7.2 Methods

Our analysis (depicted in Fig. 7.1), as is standard for atmospheric retrievals, has four components: spectral data, a forward atmospheric model, a likelihood function, and a sampling algorithm. For completeness, we review these components as they are implemented in this study, with special attention to our departures from standard analyses of JWST transmission spectra.

7.2.1 The data

We elect to perform our analysis on the WASP-39 b G395H spectrum (JWST ERS 1366, PI Batalha) previously published in [Alderson et al. \(2023\)](#). The NIRSpec/G395H observing mode has the highest spectral resolution of the modes commonly used for transiting exoplanet science, and this is the only JWST transmission spectrum with published cross-correlation detections to date ([Esparza-Borges et al. 2023, 2025](#)). We begin by re-reducing the WASP-39 b G395H data, as there have been changes to the JWST data calibration files since the original ERS data reductions. We use the `Eureka!` pipeline ([Bell et al. 2022](#)), which serves as a wrapper for the `jwst` pipeline ([Bushouse et al. 2023](#)) in Stages 1 and 2, with the addition of an extra step to perform group-level background subtraction prior to ramp fitting. This step removes $1/f$ noise due to detector readout and increases the accuracy of ramp fitting.¹

Corrections to JWST’s barycentric motion are applied during the data reduction process when setting the common wavelength solution for our spectroscopic lightcurves in Stage 2. We do not account for systemic velocity or the orbital motion of the planet in the data reduction stage. Therefore, the velocity offset that we expect is simply the system’s velocity in the barycentric frame: $-58.4421 \text{ km s}^{-1}$ ([Mancini et al. 2018](#)). We concatenate the data from both detectors (NRS1 and NRS2), resulting in a single dataset spanning roughly 2.8–5.2 microns at a (non-constant) resolution of $R \approx 2700$.

Stage 3 of `Eureka!` performs optimal spectrum extraction ([Horne 1986](#)). Stage 4 generates light curves, and as in [Esparza-Borges et al. \(2023\)](#), we opt for native *pixel* binning rather than native (spectral) resolution element binning, extracting every pixel with an assigned wavelength

¹`Eureka!` `.ecf` files used to generate these data will be made available upon publication.

value. The native pixel binning is expected to retain the greatest information content and the greatest sensitivity to the intrinsic CO band structure. This approach comes at the cost of a potential increase in pixel-to-pixel correlations, which we discuss in Section 7.4.2.1. Stage 5 of `Eureka!` performs light curve fitting. We adopt the orbital parameters and limb darkening recommendations from [Carter et al. \(2024\)](#). For each channel, we model the transit with `batman` ([Kreidberg 2015](#)), fitting only for transit depth in each channel. We include a linear ramp in time and a jump step function to model the change in flux after the mirror tilt event [Alderson et al. \(2023\)](#).

As previously described, our nominal atmospheric retrievals are not performed on a typical transmission spectrum. Rather, we perform our primary atmospheric fits at the spectroscopic lightcurve level, dividing out the temporal ramp and step function in each channel (with the timing of the step function set by a fit to the white light lightcurve). We package the lightcurves into arrays of shape $(N_{\text{exposure}}, N_{\text{pixel}})$. We do not perform any binning along the time axis, retaining 160 exposures. Separately, to directly compare our approach to traditional methods, we also perform secondary analyses on a more standard data product: a transmission spectrum binned to $R = 100$ from the native-resolution fits discussed above.

7.2.2 The forward model

Performing inference requires a mapping from the latent parameter space (i.e., the exoplanet atmosphere’s properties) to the data space (i.e., the spectrum). To this end, we generate model transmission spectra with the `CHIMERA` code ([Line & Yung 2013](#)). Our spectra include gas opacity from H_2O ([Polyansky et al. 2018](#)), CO ([Li et al. 2015](#)), CO_2 ([Yurchenko et al. 2020](#)), and SO_2 ([Underwood et al. 2016](#); [Tóbiás et al. 2018](#)). These gas species warrant inclusion because they




Figure 7.1: Schematic comparing our retrievals (left) and standard JWST retrieval methodology (right). The retrievals that we present in this work begin with a different data product and fit for additional parameters with a different likelihood function.

have been firmly detected in the atmosphere of WASP-39 b (e.g., [Rustamkulov et al. 2022](#)) and are expected based on forward modeling. We do not include methane in the list of species in our forward model because it has not been detected in this planet’s atmosphere, and it is not anticipated to be a dominant carbon carrier (see discussion in Section 7.3.2). This parameterization is appropriate for a “free retrieval” — we do not enforce that the volume mixing ratios of the considered gases follow predictions from equilibrium chemistry. We also incorporate opacity from Rayleigh scattering, an optically thick gray cloud deck (at a fit pressure level P_{cloud}), and collision-induced absorption from $\text{H}_2 - \text{H}_2$ and $\text{H}_2 - \text{He}$ pairs.

Our forward models for this cross-correlation study differ from traditional JWST retrieval forward models in a few key ways:

1. We calculate our atmospheric models at very high spectral resolution ($R \simeq 10^6$). We then convolve this spectrum with the approximate instrumental line-spread function and bin to data resolution (using the `spectres` package to ensure flux conservation; [Carnall 2017](#)).
2. To ensure that we fully resolve line cores at such high spectral resolution, we extend our model domain to very low pressures (10^{-12} bars).
3. We Doppler-shift our forward model to match a velocity offset at each orbital phase.
4. We limb-darken our spectra as a function of orbital phase. Following [Gandhi et al. \(2022\)](#), we perform this darkening for the morning and evening limbs separately because the morning and evening limbs occult different regions on the stellar disk. That is, the atmosphere is broken into two regions, each backlit with its own limb-darkened stellar profile.
5. We fit for the broadening of WASP-39 b’s spectral lines due to solid-body rotation by broad-

ening our model with a convolutional kernel. This kernel accounts for the contributions of the approaching and receding limbs to the blue and red wings of the spectral lines, respectively (Maguire et al. 2022).

We introduce another parameter, K_p , which is the change in planetary radial velocity (RV) over the course of the transit (assuming a circular orbit):

$$RV = K_p \sin(2\pi\phi) + V_{\text{sys}}, \quad (7.1)$$

for planetary orbital phase ϕ ranging from 0 (mid-transit) to 1. While we expect V_{sys} to coincide with the velocity of the WASP-39 system ($-58.4421 \text{ km s}^{-1}$; Mancini et al. 2018), we allow V_{sys} to vary as a free parameter. For each sample in the retrieval, we compare our (Doppler-shifted) atmospheric model to the flux at each exposure, as is typically done for ground-based retrievals (e.g., Gibson et al. 2020; Line et al. 2021).

In a departure from standard practices in cross-correlation analyses both from the ground and space (e.g., Brogi & Line 2019; Gibson et al. 2020; Esparza-Borges et al. 2023), we preserve the shape of the spectral continuum in our cross-correlation retrievals. Given the wavelength coverage and moderate spectral resolution of this dataset, typical continuum “removal” procedures are likely to also remove key spectral features (such as the blue slope of the water feature in the G395 bandpass). Essentially, the continuum is difficult to define given the broadness and ubiquity of features, so fitting and removing it would subtly remove atmospheric signal. Early tests on these data supported this intuition, revealing obvious biases in retrieved temperature (pushing toward the upper edge of the prior) and chemistry (very strong departures from expected abundance ratios). We therefore conclude that preserving the spectral continuum is necessary for obtaining accurate

compositional inferences.

7.2.3 The likelihood function

The comparison between model and data requires a likelihood function to assess how well the models match the data. For our likelihood function, we again depart from standard JWST retrievals, which follow the standard Gaussian likelihood function:

$$\mathcal{L} = \sum_{i=1}^N \left(\frac{y_i - f(x_i)}{\sigma_i} \right)^2, \quad (7.2)$$

for spectral data y_i , forward model f computed on each wavelength point x_i , and formal spectral uncertainties σ_i . We instead employ the [Brogi & Line \(2019\)](#) likelihood:

$$\ln \mathcal{L} = -\frac{N}{2} \left(\ln s_d s_m + \ln \left[\frac{s_d}{s_m} + \frac{s_m}{s_d} - 2C \right] \right), \quad (7.3)$$

where N is the number of data points, s_m and s_d are the empirically computed variance of the model and data, respectively, and C is the correlation coefficient. We mean-subtract both the model and the data prior to cross-correlation.² This approach “nulls” the standard Gaussian likelihood function with respect to the (assumed Gaussian-distributed) uncertainty, σ , substituting σ for its maximum-likelihood estimator. The mapping essentially fits for the noise by performing error inflation across the spectral axis for each phase. It is otherwise mathematically equivalent to the standard Gaussian likelihood function in Eq. 7.2 (see, e.g., [Gibson et al. 2020](#)).

²We note that mean-subtracting the spectrum before cross-correlation likely limits our ability to retrieve information regarding cloud decks.

7.2.4 The sampling

The likelihood function can be thought of as a “cost function” to optimize, given the constraints of the prior distribution. Calculating the likelihood function in full (e.g., on a grid) is intractable given the runtime of a single model evaluation and the high dimensionality of the retrieval problem. Retrievals therefore couple forward models to samplers that aim to efficiently explore parameter space.

While CHIMERA is generally coupled to `PyMultiNest` (Buchner et al. 2014) to perform the statistical sampling for relevant retrievals, we choose instead to couple CHIMERA to the `nautilus` sampling package (Lange 2023). We do so because the statistical literature indicates the `MultiNest` algorithm (Feroz et al. 2009) provides underestimated posterior widths on the order of 10% and biased Bayesian evidence on the order of $\ln(Z)$ of 1 unless extremely conservative settings are used (Albert 2020; Nelson et al. 2020; Ih & Kempton 2021; Miller et al. 2021; Lemos et al. 2023; Dittmann 2024; Group et al. 2025). `nautilus` has been used in the context of atmospheric retrievals (Gebhard et al. 2023); we benchmark its performance, confirm `MultiNest`’s biases, and demonstrate `nautilus`’ substantial increase in efficiency over `MultiNest` (per live point) in Appendix E.1.

We present our prior distribution ranges in Table 7.1. For the gaseous absorbers, abundances are each drawn from a log-uniform distribution between -12 and 0 dex. We fit for the abundance of the $^{13}\text{C}^{16}\text{O}$ isotopologue *relative* to the abundance of $^{12}\text{C}^{16}\text{O}$, as we expect our retrieval to be sensitive to the *ratio* of these species’ sharply wavelength-dependent opacity structures (e.g., Gandhi et al. 2023b). We place the additional constraint that the minor $^{13}\text{C}^{16}\text{O}$ abundance must lie within 5 dex of $^{12}\text{C}^{16}\text{O}$, generously encompassing the carbon isotope fractionation from




Figure 7.2: Our native pixel resolution cross-correlation constraints on gases in WASP-39 b using G395H data vs. those acquired with standard methods ($R = 100$, phase-integrated). Our method is sensitive to previously hidden gases' abundance. Specifically, we obtain a bounded abundance for CO. Vertical dashed lines mark the 16th and 84th percentiles for each distribution.

Parameter	Prior
<i>Gaseous absorbers</i>	
$\log_{10} {}^{12}\text{C}^{16}\text{O}$ [dex]	$\mathcal{U}(-12, 0)$
$\log_{10} \text{H}_2\text{O}$ [dex]	$\mathcal{U}(-12, 0)$
$\log_{10} \text{CO}_2$ [dex]	$\mathcal{U}(-12, 0)$
$\log_{10} \text{SO}_2$ [dex]	$\mathcal{U}(-12, 0)$
$\log_{10} ({}^{13}\text{C}^{16}\text{O}/{}^{12}\text{C}^{16}\text{O})$ [dex]	$\mathcal{U}(-5, 0)$
<i>Thermal structure</i>	
T [K]	$\mathcal{U}(200, 5000)$
<i>Velocity</i>	
V_{rot} [km/s]	$\mathcal{U}(0.1, 10)$
V_{sys} [km/s]	$\mathcal{U}(-73.4421, -43.4421)$
K_p [km/s]	$\mathcal{U}(-300, 300)$
<i>Miscellaneous</i>	
xR_p	$\mathcal{U}(0.5, 2.0)$

Table 7.1: Prior distributions for our retrieval parameters. $\mathcal{U}(a, b)$ denotes a uniform distribution between a and b . Samples in which the gas abundances sum to greater than unity are rejected.

the Solar System ($\log_{10} {}^{12}\text{C}^{16}\text{O} - \log_{10} {}^{13}\text{C}^{16}\text{O} \approx 1.9$; Clayton & Nittler 2004), the interstellar medium ($\log_{10} {}^{12}\text{C}^{16}\text{O} - \log_{10} {}^{13}\text{C}^{16}\text{O} \approx 1.8$; Milam et al. 2005), and planetary-mass companions ($\log_{10} {}^{12}\text{C}^{16}\text{O} - \log_{10} {}^{13}\text{C}^{16}\text{O} \approx 1.5$; Line et al. 2021). We allow the isothermal atmospheric temperature to widely bracket the equilibrium temperature of WASP-39 b (1170 K; Mancini et al. 2018). We also fit a scaling factor to the planetary radius, xR_p , fixing the reference pressure to 1 bar.

As for the velocity parameters, our priors allow a very wide range of planetary rotational velocities (0.1 to 10 km s^{-1}) that bracket the expected velocity under synchronous rotation ($\approx 1.6 \text{ km s}^{-1}$). We allow systemic velocities (V_{sys}) within 15 km s^{-1} of the reported systemic velocity ($\approx -58 \text{ km s}^{-1}$; Mancini et al. 2018). This restricted prior width avoids self-aliasing of the CO cross-correlation function (e.g., Esparza-Borges et al. 2023) while remaining more than three orders of magnitude wider than the uncertainty distribution of the system’s measured veloc-

ity (Mancini et al. 2018). We freely vary K_p over a wide range (-300 – 300 km s^{-1}) covering the expected value under a circular orbit ($\approx 129 \text{ km s}^{-1}$) (Mancini et al. 2018); we do not expect this value to be well constrained given how slowly this planet moves and the relative low resolution of the observing mode.

We perform evidence testing to assess the significance of our detections of individual species. That is, we perform additional retrievals without individual molecules, then compare the Bayesian evidence of each retrieval to the full “reference” model to assess the significance of each gas’s inclusion in our model (e.g., Benneke & Seager 2013; Welbanks et al. 2026).

7.3 Results

We now discuss the results of our phase-resolved, cross-correlation retrievals on native pixel resolution data. We address the gas abundance and velocity parameters first, then present the inferred distributions of C/O and atmospheric metallicity. Our full retrieval corner plots can be found in Appendix E.2.

7.3.1 Gas abundances and velocity parameters

Our retrieved gas abundances are presented in Fig. 7.2. As with previous analyses of this dataset (Alderson et al. 2023), we find that the abundance of both the major and minor CO isotopologues are essentially unconstrained using traditional analysis techniques (gray contours). Our $R = 100$ transmission spectrum retrieval does not retrieve a bounded CO abundance, and the gas is preferred with only $\Delta \ln(Z) = 0.2$.

In contrast, we obtain a bounded CO abundance ($-2.1^{+0.2}_{-0.3}$; purple contours) with the cross-

correlation likelihood function applied to the native pixel resolution, phase-resolved spectrum. This retrieved abundance is consistent with the much less informative constraint provided by traditional retrieval techniques and from other works in the literature that combine spectra from multiple JWST modes (see Section 7.4.1). Notably, our retrieved CO abundance lies near the upper bound associated with the non-detection using traditional methods. The presence of CO is strongly preferred in our cross-correlation analysis, with $\Delta \ln(Z) = 23$.

Our cross-correlation retrieval also finds bounded constraints on the gas-phase abundances of H_2O (-2.5 ± 0.3), CO_2 (-3.4 ± 0.3), and SO_2 (-5.6 ± 0.2). Tests with a non-isothermal thermal profile (Guillot 2010, not shown) also yielded consistent results. These constraints agree well with expectations from the literature (Constantinou et al. 2023; Kawahara et al. 2025). They are also consistent with our $R = 100$ retrievals within $\approx 2\sigma$ for each gas except SO_2 . For the latter, the discrepancy is intriguing, but it does not impact our results on metallicity and C/O due to its low intrinsic abundance. This is the first cross correlation detection of SO_2 in the literature, and its interpretation merits future work beyond the scope of this Letter. Overall, for all gases other than CO, our cross-correlation retrievals achieve comparable precision to the $R = 100$ retrievals.

For water in particular, we retrieve a slightly lower abundance than with traditional methods, despite using the same opacities in both cases to fit the same underlying data set. In this case, we believe the cross-correlation retrieval to be more accurate because we probe many individual water spectral lines with this technique; whereas the standard approach only fits a single *partial* water band and is therefore susceptible to detector edge systematics as well as degeneracies with aerosol parameters.

Beyond H_2O , given that all of our retrieved abundances, with the exception of CO, are lower with the cross correlation retrieval, we also hypothesize that at lower resolution CO is degenerate

with the abundances of the rest of the gaseous species. If this is the case, the other gases must increase in abundance to fit the CO opacity function in the low-resolution retrieval, whereas this degeneracy is resolved at native pixel resolution. However, additional forward modeling and mock retrieval work beyond the scope of this Letter would be required to confirm this hypothesis.

Our velocity parameters are not tightly constrained but are consistent with their expected values. Notably, our V_{sys} measurement is strongly bimodal (see Fig. E.3). We find that this bimodality is strongly correlated with the CO isotopologue abundance. The bluer mode is likely the “correct” mode for two reasons. First, it is consistent with the known systemic velocity of WASP-39 within 1σ . Second, this bluer mode correlates with a plausible value for the minor CO isotopologue’s relative abundance (consistent with the interstellar medium). The redder velocity mode prefers a wide range of very low minor isotopologue abundances, implying that this bimodality is due to a spectral correlation between the major and minor CO isotopologue. Interestingly, this indicates that the G395H data do have some sensitivity to the minor CO isotopologue, in line with the findings of [Esparza-Borges et al. \(2023\)](#). Our K_p constraint is unbound, with a spread of $\approx 100 \text{ km s}^{-1}$.

The velocity constraint on V_{sys} also appears to be driven by the presence of CO. When we remove the opacity of both CO isotopologues from our retrieval, the constraint on V_{sys} becomes unbounded. Furthermore, when we fit a separate V_{sys} for the NRS1 and NRS2 detectors, the V_{sys} for NRS1 pushes to the red edge of the prior, while the V_{sys} for NRS2 (which contains the bulk of the CO opacity in the G395 bandpass) converges to a similar result as our retrieval with a single V_{sys} for both detectors. This result implies that the velocity constraint in our retrieval is driven by the comb-like structure of the CO opacity.

7.3.2 C/O and metallicity

We use our freely retrieved molecular abundances to construct estimates of the volatile-proxy metallicity (i.e., $[(\text{C}+\text{O})/\text{H}]$) and C/O ratio of the gas-phase envelope (Fig. 7.3). The metallicity is constrained to be $15_{-6}^{+9} \times \text{solar}$, and the C/O is constrained to be $0.68_{-0.14}^{+0.10}$. This result is broadly consistent with other estimates from the literature based on JWST data for WASP-39 b, as we discuss in greater detail in Section 7.4.1.

Compared to the cross-correlation retrieval, our retrieval using traditional techniques obtains a consistent (but somewhat higher) metallicity ($30_{-18}^{+18} \times \text{solar}$). However, because it does not constrain CO, the latter finds a much lower C/O ($0.04_{-0.01}^{+0.08}$). These distributions are shown in gray contours in Fig. 7.3. We can see from Fig. 7.2 that the low inferred C/O is presumably an artifact of the non-detection of the CO molecule at low spectral resolution, combined with somewhat higher retrieved H₂O and SO₂ abundances. In particular, the CO non-detection pulls down the median of the CO posterior distribution by 2 orders of magnitude relative to the bounded abundance measured by the cross-correlation retrieval. Because CO is the dominant carbon reservoir in WASP-39 b, its non-detection has an outsized impact on the inferred C/O. This effect can also be seen in the yellow contours in Fig. 7.3, where we report the metallicity and C/O from our cross-correlation retrieval, but neglecting the contribution from CO gas, resulting in a comparably low C/O.

We note that CH₄ is not included in our retrieval setup because previous studies have shown no evidence for this molecule in WASP-39 b (e.g., [Rustamkulov et al. 2022](#); [Alderson et al. 2023](#)), despite JWST’s strong leverage to detect CH₄ at high confidence. Furthermore, to alter our C/O or metallicity constraints by more than 1% would require CH₄ to be present at a mixing ratio of > -3 dex. Such high abundances of methane are strongly disfavored in other retrievals that have

been performed on JWST data (e.g., [Kawahara et al. 2025](#)).

7.4 Discussion

7.4.1 Comparison to other work

Firstly, our results agree with [Esparza-Borges et al. \(2023\)](#), [Grant et al. \(2023b\)](#), and [Esparza-Borges et al. \(2025\)](#): at higher spectral resolution, NIRSpec/G395H provides substantial evidence for the presence of CO. Our results build on this statement by providing a quantitative *measurement* of the CO abundance using these same archival data. Our retrievals do not reveal a bounded constraint on the minor carbon isotope due to a bimodal correlation between this parameter and V_{sys} . In the velocity mode corresponding to the systemic velocity of WASP-39, the minor CO isotopologue has an abundance $-1.6^{+0.8}_{-1.3}$ relative to the major isotopologue, indicating sensitivity to this species. Even so, Bayesian odds ratios do not indicate that we have formally detected it ($\Delta \ln Z < 1$).

Our analysis is also broadly consistent with other studies that have applied retrievals to binned JWST data of WASP-39 b, both in the G395H bandpass ([Niraula et al. 2023](#); [Wilkinson et al. 2024](#); [Kawahara et al. 2025](#)) and in other instrument modes, such as NIRSpec PRISM ([Constantinou et al. 2023](#); [Khorshid et al. 2024](#); [Roy-Perez et al. 2025](#)), NIRISS/SOSS ([Fisher et al. 2024](#); [Changeat et al. 2025](#)), PRISM and SOSS combined ([Constantinou & Madhusudhan 2024](#)), MIRI LRS ([Powell et al. 2024](#)), NIRCам, SOSS, PRISM, and G395H separately ([Lueber et al. 2024](#)), and NIRISS, NIRCам, NIRSpec, and MIRI together ([Ma et al. 2025](#)). Specifically, these studies and our own agree in inferring an elevated metallicity and low-to-moderate C/O (both relative to solar) for WASP-39 b. Our retrieved C/O using cross-correlation techniques trends to-




Figure 7.3: The (volatile-proxy) envelope metallicity and C/O ratio of WASP-39 b, as inferred by our native pixel resolution, cross-correlation, phase-resolved free retrieval (purple). With a single NIRSpec/G395H transit, we are able to simultaneously constrain both chemical parameters. As points of comparison, we also show the posterior distributions obtained without including the contribution of CO abundance (yellow) and using standard JWST retrieval techniques (gray). 1σ and 2σ contours are outlined in black for all three posterior distributions. The orange star indicates the solar values. The ticks on the marginals mark the 16th and 84th quantiles of each distribution.

ward the highest values reported in the literature, and we justify that this inference has improved accuracy based on the bounded abundance measurement of CO.

To more directly compare our retrieval results to an array of traditional retrievals, we turn to the outputs of the JWST Transiting Exoplanet Community Early Release Science Program (JTEC-ERS) Model Synthesis effort (Welbanks et al., in prep; [Welbanks et al. 2023](#)). This study aims to establish best practices for atmospheric retrievals in the JWST era by benchmarking a number of retrieval codes run by different research groups against a single dataset: the panchromatic JWST spectrum of WASP-39 b spanning 0.5–5.5 μm . Our results are consistent within 2σ with each molecules’ free retrieval average from the Model Synthesis effort, and our results are consistent with the Model Synthesis metallicity within 0.3σ .³

7.4.2 What drives the signal quality?

7.4.2.1 Why does native pixel resolution perform well?

Our results confirm that higher spectral resolution provides more leverage on the CO band structure. In the limit of very low spectral resolution, the CO opacity “comb” becomes smeared out; in the limit of very high spectral resolution, these features are very clearly differentiable from a flat line and the opacity structure of other competing gaseous species. Therefore, while the standard retrieval yields substantial probability mass at moderately high CO mixing ratios, it finds that the CO pseudocontinuum is not necessary to explain the data (hence the formal non-detection).

The cost of extracting native pixel resolution spectra is the enhanced pixel-to-pixel correlations in these data. The cross-correlation likelihood presented in Eq. 7.3 was developed specifically

³L. Welbanks, private communication.

for this regime, in which the observed flux may not be well described with well-known Gaussian uncertainties. Additional tests that we performed seem to align with this intuition: traditional retrievals performed at native pixel resolution using the Gaussian likelihood function did not yield bounded CO constraints. The question of the ideal resolution for traditional retrievals is an intriguing one that merits exploration in future work.

7.4.2.2 Why does phase-resolving perform well?

As discussed, WASP-39 b subtends a non-negligible range of orbital phase during its transit. Its transmission spectrum is therefore Doppler-shifted as a function of time. This effect should be especially pronounced for species such as CO, given how narrow its spectral features are at native pixel resolution. Accounting for this Doppler-shifting by comparing atmospheric models directly to lightcurves — and appropriately shifting the atmospheric models themselves — is a straightforward solution. This approach is routinely employed in high-resolution ground-based exoplanet studies (e.g., [Brogi & Line 2019](#)), and more recently to a handful of JWST phase curves and secondary eclipses ([Sing et al. 2024](#); [Ouyang et al. 2026](#); [Snellen et al. 2026](#)). We have shown here that it extends well to JWST NIRSpec/G395H transits, even in the limit of slower-moving exoplanets.

7.5 Conclusions

This Letter provides the proof of concept that an exoplanet’s atmospheric C/O ratio and metallicity can be measured robustly with the JWST NIRSpec/G395H instrument mode, alone. This result is made possible by applying techniques that had previously been exclusively employed

by the ground-based observing community — i.e., phase-resolved cross-correlation retrievals on native pixel resolution spectra. In the case of WASP-39 b’s NIRSpec/G395H data, our method yields additional sensitivity to the presence of CO compared to standard techniques. We confidently detect the gas and produce a bounded abundance measurement. We simultaneously measure bounded abundances for all of the key carbon- and oxygen-bearing species present in WASP-39 b’s atmosphere.

Our results firmly demonstrate that traditional approaches that bin JWST data in wavelength and in time can dilute information about exoplanet atmospheric composition. We also show that the non-detection of CO, which is a common occurrence in NIRSpec/G395 datasets, can heavily bias inferred C/O ratios (as well as metallicities to a lesser extent) by “hiding” the dominant carbon-carrying molecule in hot Jupiter atmospheres. Applying the high-resolution phase-resolved retrieval techniques from this Letter to additional G395H observations of transiting exoplanets is likely to be fruitful, providing a pathway to maximally make use of legacy datasets.

Chapter 8: 1D on 3D: The impact of inhomogeneous atmospheric properties on high-resolution transmission spectroscopy retrievals

8.1 Introduction

How does gas move and change in synchronously rotating gas giants? At pressures substantially less than 1 bar—where spectroscopy probes—previous chapters have shown that we cannot know the answer without considering atmospheres as truly three-dimensional (3D) bodies. Indeed, the processes that set gas abundance, temperature, velocity, and cloud distributions in this regime are fundamentally 3D (e.g., [Showman & Polvani 2011](#); [Komacek & Showman 2016](#); [Zhang & Showman 2018](#); [Hammond & Lewis 2021](#)), so these quantities are predicted to exhibit substantial spatial variations (e.g., [Kataria et al. 2016](#); [Mendonça et al. 2018](#); [Beltz et al. 2021](#); [Parmentier et al. 2021](#); [Roman et al. 2021](#); [Steinrueck et al. 2021](#); [Komacek et al. 2022](#); [Helling et al. 2023](#); [Zamyatina et al. 2023](#)).

Inferences of exoplanet atmospheres, however, rarely take such 3D variations into account. Indeed, retrievals generally model atmospheres as varying only with planetary altitude, if they contain any spatial variation. Such parsimony is generally by design: because simpler models are often preferred by data (e.g., [Madhusudhan & Seager 2009](#); [Feng et al. 2016](#); [Welbanks & Madhusudhan 2022](#)), using 1D models tends to limit overfitting, lower computational burden, and reduce the dimensionality of statistical sampling (e.g., [Barstow & Heng 2020](#)).

Even as HRCCS retrievals have grown in promise ([Gibson et al. 2020](#); [Line et al. 2021](#); [Gandhi et al. 2023b](#); [Pelletier et al. 2023](#)), significant challenges to their underlying 1D assumptions have made the need for more complex models more pressing. Observations using this tech-

nique indicate the presence of spatial variations in exoplanet atmospheric chemistry, temperature, and gas velocity across a number of planets (Flowers et al. 2019; Beltz et al. 2020; Ehrenreich et al. 2020; Borsa et al. 2021; Kesseli & Snellen 2021; Gandhi et al. 2022; Kesseli et al. 2022; Brogi et al. 2023; Pelletier et al. 2023; van Sluijs et al. 2023b; Hood et al. 2024; Maguire et al. 2024; Wardenier et al. 2024; Nortmann et al. 2025). Forward modeling exercises have supported these observations, demonstrating that using 3D models distinctly alters the signal strengths and spectral line shapes of gas species detected with HRCCS (Kempton & Rauscher 2012; Kempton et al. 2014; Harada et al. 2021; Wardenier et al. 2021; Beltz et al. 2022; Savel et al. 2022; Beltz et al. 2023; Savel et al. 2023; Wardenier et al. 2023; Blain et al. 2024; Wardenier et al. 2024). Two conclusions can be drawn from these studies. First, the influence of 3D effects on data is quantifiable. Second, neglecting these effects can alter the significance of — and even bias — atmospheric inferences. Taken together, it is clear that the standard 1D approach used in retrievals requires significant enhancements to accurately and maximally extract information from exoplanet atmospheres with HRCCS.

We would like to include 3D effects into inference tools, both to mitigate biases and to understand exoplanet atmospheres more deeply. These 3D effects could include variable chemistry, such as condensation or dissociation fronts, which would manifest in spectral line ratios. Thermal variations, such as limb-to-limb temperature differences, could be evident in the spectral line contrast relative to the continuum. And atmospheric dynamics, such as transitions from equatorial jets to day–night flow, could be disentangled by their influence on line broadening. In principle, spatial variation in these effects would be evident as spectral changes as a function of orbital phase, especially for the hottest planets (Wardenier et al. 2022).

Before altering retrieval frameworks to account for such 3D effects, however, we seek to

understand *a priori* how and to what extent these 3D effects may influence inferences. Namely, we seek to answer the questions: *what types of spatial variation most impact HRCCS retrievals, and how do they manifest?* This topic has been rigorously explored in the context of low-resolution spectroscopy (Line & Parmentier 2016; Rocchetto et al. 2016; Blecic et al. 2017; MacDonald & Madhusudhan 2017; Caldas et al. 2019; Lacy & Burrows 2020; MacDonald et al. 2020; Pluriel et al. 2020; Taylor et al. 2020; Espinoza & Jones 2021; Pluriel et al. 2021; Chubb & Min 2022; MacDonald & Lewis 2022; Nixon & Madhusudhan 2022; Zingales et al. 2022; Falco et al. 2024; Schleich et al. 2024). Less attention has been paid to so-called “3D biases” in HRCCS retrievals, in large part because of the computational complexity and relative novelty of this technique. Given the resolved line profiles and imprints of atmospheric dynamics on HRCCS data, though, it stands to reason that 3D biases may in fact be *more* pronounced in HRCCS retrievals than low-resolution retrievals, motivating careful consideration of this problem.

In this work, we assess the impact of 3D structure on 1D HRCCS retrievals by performing 1D retrievals on a wide array of simulated high-resolution spectra generated from 3D model atmospheres. In doing so, we aim to methodically reveal how a 1D HRCCS retrieval responds to spatial variation in a controlled simulation environment. The results of this study will enable the future development of improved HRCCS retrievals that natively account for 3D effects and directly recover 3D information. We describe our simulation pipeline in Section 8.2. Section 8.3 lays out the results of our experiments. We discuss implications, applications, and caveats of our work in Section 8.4. Finally, we summarize our findings and remark on future directions in Section 8.5.

8.2 Methods

Our simulation pipeline (Fig. 8.1) involves six steps: 1) 3D atmospheres, 2) 3D forward-modeled spectra, 3) HRCCS data simulations, 4) 1D forward-modeled spectra, 5) cross-correlation, and 6) retrieval. These are laid out sequentially in the following sub-sections.

8.2.1 3D atmospheres

We begin with a prediction of an exoplanet’s underlying 3D atmospheric structure (i.e., temperature, pressure, and velocity as a function of latitude, longitude, and altitude). We consider two classes of model atmospheres: general circulation models (GCMs) and simplified atmospheres informed by GCM calculations. The latter are used to inform how different aspects of 3D atmospheric structure and dynamics impact high-resolution spectra and HRCCS retrievals.

8.2.1.1 General circulation models

Our GCMs follow the formalism of [Beltz et al. \(2021\)](#). We describe this framework below.

The models calculate the temperature and velocity structure of an atmosphere as a function of latitude, longitude, and pressure. These quantities are computed via spectral methods with a horizontal resolution of T31 (corresponding to approximately 4° increments at the equator) and either 45 or 65 vertical layers that are evenly spaced in log pressure. The vertical dimension extends from 100 bar to 10 μ bar.

The dynamical core of the GCM solves the primitive equations of meteorology — a simplified form of the Navier-Stokes equations with a long history of Solar System and exoplanetary application (e.g., [Showman et al. 2020](#)). The top three pressure levels of the GCMs (“sponge lay-




Figure 8.1: Schematic of our simulation pipeline. We begin with a 3D model of the atmosphere (either a GCM output or a simplified atmospheric structure), then post-process it with 3D ray-striking radiative transfer to yield a transmission spectrum. We then simulate an observation of this spectrum using the `scope` package (Savel et al. 2025). Next, we generate a 1D forward model and cross-correlate it against our simulated dataset. Finally, we perform parameter inference on the simulated dataset with a full HRCCS retrieval.




Figure 8.2: The zonal-mean zonal wind plots of the two GCMs used in this work. The hot Jupiter (left) has a faster and more spatially extended jet than the ultra-hot Jupiter (right), in large part because the latter is damped via coupling to an intrinsic magnetic field (Showman & Polvani 2011; Beltz et al. 2021).

ers”) feature an additional scale-free vertical drag term to minimize the unphysical growth of (e.g., gravity or acoustic) waves in the simulation (e.g., Forget et al. 1999; Klemp et al. 2008; Mayne et al. 2014; Wills & Schneider 2016; Christie et al. 2024). The atmospheric circulation is coupled to the radiation field through heating and cooling. Radiative transfer is implemented with the double-gray approximation as in Rauscher & Menou (2012).

The GCMs used in this study approximate drag in a “kinematic magnetohydrodynamics” approach, incorporating the resistive magnetohydrodynamics drag timescale derived in Perna et al. (2012). The drag term assumes a deep-seated, dipolar planetary magnetic field that is aligned with the poles; the drag works against east–west flow, converting kinetic energy into localized Ohmic heating. This drag is evaluated locally, which is crucial when there exist large differences in conductivity across the planet—for instance, between the permanent day and night sides of ultra-hot Jupiters.

We choose to study two planets that bracket stark changes in expected thermal structure,

Table 8.1: Planetary Parameters

Parameter	Hot Jupiter	Ultra-hot Jupiter
R_p (m)	1×10^8	1.3×10^8
T_{eq} (K)	1500	2160
T_{int}^1 (K)	500	500
$R_\star$ ($R_\odot$)	1.155	1.16
b	0.499	0.14
i ($^\circ$)	86.78	88.0
P (days)	3.3	1.81
M_p (M_{Jup})	0.26	0.92
K_p (km/s)	148.48	198.36
a (AU)	0.045	0.033
g (m/s^2)	8	6.8
μ_{fixed}^2	2.303	2.031
t_{dur} (hours)	3	3.8

composition, and atmospheric dynamics: an ultra-hot Jupiter modeled after WASP-76b ($T_{eq} = 2160$ K; [West et al. 2016](#)) and a “classic” hot Jupiter modeled after HD 209458b ($T_{eq} = 1500$ K; [Charbonneau et al. 2000](#)). Hereafter, we will refer to the hotter model simply as our “ultra-hot Jupiter model,” and the cooler planet as our “hot Jupiter model”. The hotter planet will have greater day–night and radial thermal gradients than the cooler analog due to its enhanced dayside heating ([Komacek & Showman 2016](#); [Lothringer et al. 2018](#); [Tan & Komacek 2019](#)). In the hottest regions of the hotter planet, molecules (most notably water and hydrogen) can thermally dissociate, producing sharp chemical gradients ([Arcangeli et al. 2018](#); [Bell & Cowan 2018](#); [Komacek & Tan 2018](#); [Parmentier et al. 2018](#)). Temperature differences between the two planets also influence their dynamics: radiative-hydrodynamic theory predicts that stronger radiative forcing of the hotter planet will drive faster winds, both in the lower-pressure day–night flow and in the deeper equatorial jet ([Komacek & Showman 2016](#); [Roth et al. 2024](#)).³ These winds, however, may be weakened

³For *much* hotter planets, this trend is not expected to continue. The heat distribution effect of H_2 dissociation and recombination can *decrease* the day–night contrast at higher temperatures ([Tan & Komacek 2019](#)), reducing the speed of the equatorial jet in the hottest ultra-hot Jupiters ([Showman & Polvani 2011](#)).

by the previously described magnetohydrodynamic effects in the hotter atmosphere (Perna et al. 2012; Rauscher & Menou 2013; Beltz et al. 2021; Christie et al. 2025), as we find to be the case in our models (see Fig. 8.2).

For our ultra-hot Jupiter, we use a GCM simulated with a 3 G dipolar magnetic field (published in Beltz et al. 2021). The magnetic field of this planet (and exoplanets in general) is largely unconstrained (e.g., Brain et al. 2024); a B field of at least 3 G provided the best match to Doppler shift measurements of WASP-76b (Ehrenreich et al. 2020; Beltz et al. 2023). For our hot Jupiter, we opt for a 0 G model (i.e., no magnetic drag included). The basic planetary properties of our simulations are listed in Table 8.1, and the zonal-mean flow patterns for the GCMs are shown in Fig. 8.2. For detailed numerical setup of the GCMs, see Beltz et al. (2020) and Beltz et al. (2021).

Our tests indicate that some of the strongest spectral lines in the optical are saturated when the model domain is confined to pressures greater than 10 microbars. To alleviate this numerical issue, we perform an isothermal and barotropic extension of the GCM output to a model domain top of 1 microbar. That is, we perform a constant extension of the atmospheric temperature and wind fields in each atmospheric column. This assumption is reasonable for our study because transmission spectroscopy is not particularly sensitive to temperature gradients (e.g., Miller-Ricci et al. 2008). The constant barotropic extension of the wind field is self-consistent with the isothermal extension, as it relies on the day–night temperature gradients remaining constant with pressure. This approach essentially limits vertical wind shear, which is not a focus of this work.

Table 8.2: Summary of experiments performed

Quantity	Model Type	Baseline	Comments
Temperature	Radial variations Latitudinal variations Longitudinal variations Day–night contrast East–west asymmetry 3D thermal structure	Isothermal	Magnitude of thermal variations drawn from general circulation models (see Section 8.2.1).
Chemistry	Eq. chemistry Varying H ₂ O Morning H ₂ O dissociation	Const. chemistry	Const. from eq. chemistry at T_{eq} .
Dynamics	Include winds Include rotation Include winds and rotation	Static atmosphere	Winds interpolated from general circulation models
Misc.	Rotation-free Const. transit depth		

8.2.1.2 Idealized atmospheres

As mentioned previously, we supplement our GCMs with a set of idealized atmospheres. The properties of these atmospheres are chosen to introduce individual aspects of 3D variation (e.g., rotation) or to vary along individual spatial axes (e.g., longitude). Analyzing this set of atmospheres in addition to the full GCMs isolates which types of spatial differences most significantly impact 1D HRCCS retrievals, whereas the GCMs demonstrate the combined effects of multiple forms of spatial heterogeneity.

We present our grid of simplified atmospheres in Table 8.2. The atmospheres vary from one another in the setup of their 3D temperature, chemistry, and dynamics. We additionally perform several miscellaneous experiments to test specific 3D effects that we wish to quantify. To retain a degree of consistency with the GCMs, each set of experiments begins with a full GCM output, then isolates a specific spatially-varying quantity. For instance, for the “radial variations ultra-hot Jupiter” model, we start with the substellar temperature – pressure (TP) profile from our ultra-

hot Jupiter GCM. We then apply this profile around the whole planet at each latitude–longitude grid point, thus suppressing any 2D temperature variations. This procedure allows us to isolate physically motivated variations in the quantities of interest, while maintaining the same spatial grid between models.

The temperature experiments are our most extensive, given that temperature differences between limbs are expected to be a leading-order effect in transmission spectrum asymmetries (Wardenier et al. 2021; Savel et al. 2023; Roth et al. 2024; Arnold et al. 2025). This set of experiments assumes constant chemistry and mean molecular weight (held to μ_{fixed}) set by the equilibrium temperature of the planet. We perform our chemistry calculations using the `FastChem` (Stock et al. 2022) equilibrium chemistry code at solar metallicity and C/O ratio. We use `FastChem`’s ion network including atoms up to uranium.

To understand the impact of 2D variations in thermal structure, we introduce the following simplified atmospheric prescriptions. Our latitudinal variations model allows the GCM variations in azimuth. The longitudinal variations model follows similarly, preserving the equatorial longitudinal variation at all latitudes and holding each radial T–P profile to be isothermal. Our day–night temperature gradient model allows only variation in a single spatial parameter from the substellar to antistellar point. Finally, our east–west temperature gradient model follows similarly, with a longitudinal offset applied.

To generate a model with a pure day–night temperature contrast, we employ some of the formalism of Nixon & Madhusudhan (2022). We take $\hat{\theta}$ to be the angle measured from the substellar point:

$$\cos(\hat{\theta}) = \frac{s}{r}, \quad (8.1)$$

where s is the location of the ray such that $s = 0$ at the terminator and r is the distance between a ray and the center of the planet. Our “day–night contrast” model linearly varies temperature in $\hat{\theta}$ between the substellar and antistellar temperatures of the underlying GCM. We generate our “east–west asymmetry” model similarly, but with the entire model rotated 30° eastward (i.e., a purely longitudinal shift).

We approach our chemistry-varying simulations similarly, assuming an otherwise isothermal atmosphere and no winds for the underlying model. The first of these experiments enforces equilibrium chemistry throughout the atmosphere as set by the local temperature field and pressure, again as interpolated from `FastChem` tables (Stock et al. 2022). For a second experiment, we fix the volume mixing ratios of all gas species to the equilibrium chemistry prediction at the planet’s equilibrium temperature at 1 mbar. The exception is water, which we allow to vary according to the GCM’s temperature and pressure fields. For our hot Jupiter model, the water abundance varies from -3.42 dex to -4.86 dex. For the ultra-hot Jupiter, the abundance varies from -3.42 dex to -11.81 dex—due to water’s thermal dissociation in high-temperature environments. We do so to isolate the effect of a single gas species’ variation — in this case, H_2O , because its abundance is expected to have significant 3D variation in the ultra-hot planet. Finally, we approximate a limiting case of the thermal dissociation of water in ultra-hot Jupiters (e.g., Parmentier et al. 2018) by removing water entirely from the evening limb. Even though the hot Jupiter model is cooler, we include a “cool limb water” model for this planet, too, as a test of limiting-case chemical heterogeneity. For this set of experiments, we locally vary the atmospheric mean molecular weight

consistently with the chemistry, as doing so may strengthen the spectral features of dayside chemical species (e.g., [Pluriel et al. 2020](#)) For our hot Jupiter, μ varies between 1.4 and 2.37. For our ultra-hot Jupiter, μ falls to as low as 1.3.

We additionally wish to understand the effects of 3D atmospheric *dynamics* on 1D retrievals. There are two main such effects: winds and planetary rotation. As we will describe in Section 8.2.1.1, we include these effects in our 3D post-processed spectra by Doppler shifting the gas opacity in each grid cell according to the local line-of-sight velocity. In our experiments, winds predicted by the GCM are simply included or neglected at the spectral post-processing stage. All our models include some rotational broadening prescription, as HRCCS generally broaden their 1D spectra to approximate the effect of planetary rotation (e.g., [Debras et al. 2024](#); [Kanumalla et al. 2024](#); [Maguire et al. 2024](#)). For a subset of our models, we include self-consistent rotation, which locally calculates the projected line-of-sight velocity from planetary rotation in every grid cell. This calculation captures latitudinal, longitudinal, and radial variations in the projected rotational velocity. All other models are subjected to a post-hoc 1D rotational broadening kernel to capture the effect of uniform, solid-body rotation (see Appendix F.1 for more information and a derivation; [Gandhi et al. 2022](#)). That is, we treat no atmosphere as truly static, though only some models include the GCM’s rotational and wind fields.

HRCCS observations motivate further, miscellaneous tests. For instance, H_2O and CO are often found together and retrieved at the same velocity, even in datasets where they are known to have velocity differences (e.g., [Brogi et al. 2023](#); [Mansfield et al. 2024](#)). To assess the impact of this choice on inferences, we include a subset of models that offset CO absorption by a constant velocity. Separately, the impact of spatially varying water abundance may be exacerbated by self-consistent rotation. This effect could perhaps result in the curious depletion of water noted in

some HRCCS studies (e.g., [Pelletier et al. 2021](#)). So, we include a model that features both water dissociation and 3D rotation, combining the impact of 3D chemistry and 3D atmospheric dynamics. These tests aim to explore potential degeneracies between abundance and velocity parameters and explicitly stress test the flexibility of the 1D retrievals.

8.2.2 3D modeled spectra

With our atmospheric structures simulated, we next seek to calculate these models' transmission spectra. We do so with the 3D ray-striking radiative transfer code developed in [Kempton & Rauscher \(2012\)](#) and extended in [Savel et al. \(2022\)](#). Briefly: we compute net extinction from atmospheric gas along lines of sight through 3D atmospheres. The extinction is calculated as the sum of absorption and Rayleigh scattering, the latter of which is assumed to be purely out of the beam (justified by, e.g., [Seager & Sasselov 2000](#); [Hubbard et al. 2001](#); [MacDonald & Lewis 2022](#)).

When applicable (i.e., for the models that include 3D atmospheric dynamics), we account for local motions in each atmospheric grid cell (including winds and planetary solid-body rotation) by Doppler shifting the gas absorption as a function of wavelength based on the local line-of-sight velocity. This 3D broadening is described as ([Kempton & Rauscher 2012](#))

$$\lambda = \lambda_0 \left(1 - \frac{v_{\text{LOS}}}{c} \right)$$

$$v_{\text{LOS}} = -u \sin(\phi) - v \cos(\phi) \sin(\theta) - \Omega(R_{\text{P}} + z) \sin(\phi) \cos(\theta), \quad (8.2)$$

where λ_0 is the rest wavelength, v_{LOS} is the line-of-sight velocity in a grid cell, c is the speed of

light, u is the east–west velocity, v is the north–south velocity, ϕ is longitude, θ is latitude, z is altitude (above 1 bar in pressure), and Ω is the planetary angular rotation speed. Note the longitudinal dependence of the u term, the longitudinal and latitudinal dependencies of the v term, and the longitudinal, latitudinal, and radial dependencies of the Ω term. This calculation is therefore distinct from the post-hoc 1D broadening kernel described above, which only takes into account the latitudinal dependence of Ω . The wind terms are also not captured by a single velocity shift, because even a constant velocity in the planetary frame is modulated by the ϕ and θ dependencies in the *projected* velocity. Hence, constant u and v produce additional spectral line broadening that is not captured in standard broadening kernels.

We include the effect of stellar limb-darkening on the transit depth (within the white-light radius of the planet, R_P only) using the `batman` package (Kreidberg 2015). The effect of limb-darkening on the atmospheric transmission is also calculated by reducing the amount of light incident along a given ray.

We perform this calculation for 63 orbital phases. This number of phases is similar to the number of in-transit exposures attainable for hot Jupiters orbiting bright stars (e.g., Ehrenreich et al. 2020; Mansfield et al. 2024). Our phase list includes a perfect center-of-transit exposure for numerical convenience. With the exception of our “rotation-free” model, we rotate the atmosphere over the course of the transit under the assumption of synchronized rotation. We assume no “smearing” of the planetary spectrum occurs during a given exposure — i.e., that planetary motion is negligible over the course of a single exposure.

Our calculations require model atmospheres that are on a fixed-altitude grid (as opposed to pressure). To transform the GCM outputs onto such a grid, we interpolate our model atmospheres onto new spatial grids under the assumption of hydrostatic equilibrium (as is consistent with our

GCMs’ assumptions). When constructing the altitude grid, we set the scale height as

$$H = \frac{k_B T}{\mu g}, \quad (8.3)$$

where k_B is Boltzmann’s constant, T is temperature, μ is mean molecular weight, and g is gravity. We consider each relevant quantity locally in each cell. For instance, spatially varying T impacts H in the non-isothermal models, and spatially varying μ impacts H in the models with spatially varying chemistry. We vary g relative to 1 bar radii quoted in Table 8.1.

We perform our ray-striking simulations in the near-infrared, spanning 1–2.6 microns at a resolution of 250,000. This specific wavelength range is motivated by the desire to compare our spectra to IGRINS observations, which have yielded several hints of 3D effects (e.g., [Brogi et al. 2023](#); [Mansfield et al. 2024](#)). We list these opacity sources in Appendix F.

8.2.3 Simulated data pipeline

At this point, we have calculated noiseless high-resolution transmission spectra. To perform meaningful retrievals on these spectra, we must simulate HRCCS data cubes that approximate the statistical properties of real data.

We simulate HRCCS observations with the `scope` package ([Savel et al. 2025](#)). We now briefly describe the steps taken, especially noting differences from our previous implementation.

1. Doppler shift the forward model to the orbital velocity of the planet at the simulated phases.
2. Broaden the model with a 1D convolutional rotation kernel if 3D broadening (Section 8.2.1.1) is not applied.

3. Convolve the planet down to instrumental resolution with a Gaussian kernel. We pick the IGRINS instrument (Yuk et al. 2010; Park et al. 2014; Mace et al. 2018) as a fiducial example.
4. Account for the stellar flux by selecting an appropriate PHOENIX stellar model spectrum, $F_\star$, and multiplying it by $(1 - F_{p,\lambda})$, where $F_{p,\lambda}$ is the exoplanetary transmittance.
5. Interpolate $F_\star(1 - F_{p,\lambda})$ to the instrumental pixel grid.
6. Apply an instrumental blaze function. We adopt the blaze function from observations of WASP-77Ab (Line et al. 2021) as our fiducial example.
7. Account for telluric absorption by multiplying $F_\star(1 - F_{p,\lambda})$ by F_{tell} , a telluric spectrum. Our F_{tell} is a simple data-driven tellurics forward model trained on archival IGRINS spectra of standard stars (Savel et al. 2022).
8. Add a representative amount of photon noise to our data cube, corresponding to a median photon SNR per pixel of roughly 350 (high, but not unreasonably so compared to observations; Line et al. 2021).

After we generate our synthetic data cube, we normalize it and perform singular vector decomposition (SVD, also referred to as PCA in the literature; e.g., de Kok et al. 2013) on the simulated data. This is a standard data processing step in HRCCS, used to remove the static stellar and telluric contamination while preserving the Doppler-shifting planetary signal (e.g., de Kok et al. 2013; van Sluijs et al. 2023b). This step yields two data products: a residual matrix $\mathcal{R}$ (containing noise and the planetary signal) and an SVD matrix $\mathcal{P}$. We consistently remove 4 principal components from our simulated datasets. This number of components is arbitrary but

is consistent with the range that is standard in observational practice (e.g., [Smith et al. 2024a](#); [Wardenier et al. 2024](#); [Panwar et al. 2025](#)).

This process necessarily introduces randomness into our calculations in Step 7. To test the effect of shot noise on our final results, at this stage we produce two sets of simulations with different noise instances.

8.2.4 1D forward model

A forward model is required to make inferences (either by cross-correlation analysis or formal retrieval) from our synthetic datasets. We make use of the CHIMERA 1D transmission spectrum code ([Line et al. 2013](#); [Kreidberg et al. 2014](#)) as our forward model. We use the GPU-optimized, HRCCS version of CHIMERA (e.g., [Line et al. 2021](#)), which computes spectra over wide wavelength ranges in parallel. Our version of CHIMERA assumes an isothermal TP profile. Gas abundances are left “free” (i.e., not self-consistent; we only enforce that the volume mixing ratios sum to 1 and that they affect the mean molecular weight) and are held constant throughout the modeled atmosphere. We use the same opacities as for the 3D forward models.

In these 1D models, gravity varies as a function of altitude as set by distance from the fit reference pressure. We fix our planets’ radii to their corresponding white-light radius (Table 8.1). We implement Rayleigh scattering in our transmission spectrum using the CRC polarizabilities for these species, in accordance with our 3D radiative transfer code. We use the same opacities for the 1D forward models as the 3D forward models.

Prior to performing any cross correlation or retrieval analyses, we successfully benchmarked our 3D forward model against CHIMERA to ensure that any biases in inferred atmospheric prop-

erties are due to true differences between modeled physics and chemistry (i.e., 1D vs. 3D differences), as opposed to numerical and implementation differences between the codes themselves (see Appendix F).

8.2.5 Cross-correlation

Before moving directly to retrievals, we will pause and consider cross-correlation “ K_P – V_{sys} maps”, which have become a standard analysis tool in high-resolution spectroscopic characterization of exoplanet atmospheres. The K_P – V_{sys} map plots the degree to which the planet signal is recovered along two axes in velocity space:

$$RV = K_P \sin(2\pi\phi) + V_{\text{sys}}, \quad (8.4)$$

where RV is the planet’s measured radial velocity, ϕ is its orbital phase (ranging 0–1), K_P is the amplitude of the planetary orbital motion,⁴ and V_{sys} is the systemic velocity of the system with respect to the Solar System’s barycenter. We assume perfectly circular orbits in this work.

These maps are standard data products and are often used to justify running retrievals in the first place—i.e., gases that are detected in K_P – V_{sys} maps are then included in subsequent retrievals (e.g., [Pelletier et al. 2023](#)). Furthermore, a wide range of studies focus entirely on K_P – V_{sys} maps themselves. In addition to providing evidence for the presence of gas species, the structure of these maps can indicate atmospheric dynamics and 3D atmospheric structure via velocity offsets (e.g., [Snellen et al. 2010](#); [Flowers et al. 2019](#); [Asnodkar et al. 2022](#); [Prinoth et al. 2022](#)) and the shape of the central “blob” ([Wardenier et al. 2021, 2023](#)).

⁴Not to be confused with the stellar reflex motion.

To generate our K_P – V_{sys} maps, we cross-correlate our synthetic data with their closest 1D, isothermal counterparts. For these cross-correlation templates, we calculate the abundance of each gas at 1 mbar and the equilibrium temperature of each planet. These abundances are fixed throughout the vertical domain of the template models. Once the models are calculated, we construct K_P – V_{sys} maps by cross-correlating $\mathcal{R}$ with the 1D template $\mathcal{T}$:

$$\text{CCF}(\mathcal{R}, \mathcal{T}, s) = \frac{\sum^n \mathcal{R}(\lambda) \mathcal{T}(\lambda - s)}{n \sigma_{\mathcal{R}} \sigma_{\mathcal{T}}}, \quad (8.5)$$

for cross-correlation function $\text{CCF}(\mathcal{R}, \mathcal{T}, s)$, number of data points n , standard deviations $\sigma_{\mathcal{R}}$ and $\sigma_{\mathcal{T}}$, and lag s . The cross-correlation function is calculated along all lags s , which are then converted to velocities by treating shifts along the wavelength axis as Doppler shifts.

Note that $\mathcal{T}$ is multiplied with the SVD matrix $\mathcal{P}$, and SVD is performed on their product before cross-correlating with $\mathcal{R}$. This “model reprocessing” step ensures that the planetary signal we compare to the data has been subjected to the same warping as the data during the SVD process and is a standard step in HRCCS analyses (e.g., [Gibson et al. 2020](#); [Line et al. 2021](#); [Brogi et al. 2023](#)).

8.2.6 High-resolution retrievals

The final rung in our simulation pipeline is the HRCCS retrieval. Our approach uses the [Line et al. \(2021\)](#) HRCCS retrieval pipeline, with the main distinction being that we use the transmission spectrum version of CHIMERA (as opposed to thermal emission geometry). The code combines the 1D forward model described in Section 8.2.4 with the PyMultiNest nested sampler ([Buchner](#)




Figure 8.3: Simulated transmission spectra corresponding to a subset of the model grid described in Table 8.2. These spectra are calculated at a resolution of $R = 10^6$ for visualization purposes.

2016) to globally sample the posterior distribution. This approach has the benefit of being able to capture high-dimensional, multi-modal posterior distributions. Because we are primarily interested in parameter inference and not model comparison, we set the evidence tolerance to 0.5 and the efficiency to 0.8, with 500 live points. We do not perform importance sampling for these retrievals.

We assess model goodness-of-fit using the Brogi & Line (2019) likelihood:

$$\log(L) = -\frac{n}{2} \left[\log(s_f s_g) + \log \left[\frac{s_f}{s_g} + \frac{s_g}{s_f} - 2\text{CCF}(\mathcal{R}, \mathcal{T}, s) \right] \right], \quad (8.6)$$

for model variance s_f and data variance s_g . For a given model at a given $K_P - V_{\text{sys}}$, the likelihood function is computed at all phases at once and summed. The sampler then seeks to optimize the likelihood function with respect to $K_P - V_{\text{sys}}$, in addition to the gas abundance inventory, broadening kernel width, reference pressure, and temperature structure. For a subset of our retrievals,

we also include a scale factor, a , fit in log space. This parameter stretches or shrinks the planetary spectrum.

8.3 Results

To build intuition before diving into retrievals, we begin our results section with the output from our 3D forward modeled spectra and K_P – V_{sys} maps.

8.3.1 3D forward models

We compute a subset of our 3D forward-modeled spectra at very high spectral resolution ($R = 10^6$) to illustrate the ways in which physics intrinsically impacts spectral line profiles. These effects broadly follow discussions in the literature (Brown 2001; Fortney et al. 2010; Kempton & Rauscher 2012; Showman et al. 2013a; Kempton et al. 2014; Rauscher & Kempton 2014; Wardenier et al. 2021; Savel et al. 2022; Beltz et al. 2023); we show them (Fig. 8.3) to make direct connections to our model grid.

Together, atmospheric dynamics both shift and broaden line profiles (right panel of Fig. 8.3). Rotation of the planet alone smears out lines, decreasing their depths, with both blueshifted (approaching limb) and redshifted (receding limb) components emerging. Winds also introduce blueshifted and redshifted components to the line profile, but these Doppler shifts do not follow solid-body rotation, so the line profile can readily depart from purely symmetric broadening. Notably, winds also do not necessarily induce a simple wavelength shift: for instance, an equatorial jet can broaden line profiles by producing blueshifting for the approaching limb and redshifting for the receding limb. Day–night flow primarily blueshifts line profiles, but it can *also* induce line broad-

ening if there is asymmetry in the day–night flow strength. Including the effects of both winds and planetary rotation yields a complex line structure with depths and line positions substantially departing from a “dynamics-free” model.

Note that stronger lines experience more relative broadening from rotation and more shifting due to winds. The first effect is explained by strong lines becoming optically thick at higher altitudes (where solid-body rotation is faster; e.g., Eq. 4 of [Kempton & Rauscher 2012](#)) than weak lines. Note that this effect cannot be captured with standard convolutional kernels, which assume that all velocity signatures originate from the same planetary radius. Winds can more strongly impact these strong lines, as well. For example, our hot Jupiter’s wind profiles tend to have faster winds toward the top of the atmosphere (Fig. 8.2); details of radiative forcing and sponge layers aside, a simple momentum balance argument suggests that it is easier to move the lower-density parcels at high altitudes than higher-density, deeper parcels.

While atmospheric dynamics similarly affect all species’ line profiles (under the assumption of spatial chemical homogeneity), including spatially varying chemistry introduces effects that depend on which species a particular line originates from. For instance, in chemical equilibrium, the higher CH_4 abundance introduces a forest of CH_4 lines that create a pseudo-continuum in the spectral region shown in Fig. 8.3. Analogously, depleting H_2O from the warm limb decreases the strength of the strong H_2O lines. That is, varying chemistry can cause certain lines to appear or disappear altogether, or to change strength relative to one another.⁵ Line *positions*, however, are largely unchanged by different assumptions of chemistry.

Finally, the primary effect of spatially varying temperature is to alter line strengths. Line

⁵The strong effect of chemistry (let alone spatially varying chemistry) on line ratios, coupled with the sensitivity of HRCCS to line ratios (e.g., [Gibson et al. 2022](#); [Gandhi et al. 2023b](#)), illustrates why HRCCS is such an excellent technique for constraining chemical compositions.

strengths change because changes in temperature directly alter the atmospheric scale height, which is proportional to line strength in transmission spectroscopy (e.g., [Brown 2001](#); [Miller-Ricci et al. 2008](#)). Note that not all lines are impacted identically, though—once again, the location in the atmosphere in which the line becomes optically thick determines the change in line profile. In this case, the location at which the spectrum becomes optically thick controls how much the line is stretched by the altered temperature field. For example, a very strong line would not be as impacted by radially varying temperature if the atmosphere is isothermal at pressures less than 1 bar. Additional structure is imparted because of the temperature dependence of different lines’ underlying opacity functions; hotter gas imparts a greater degree of thermal broadening to the line profiles.

In summary: in transmission, atmospheric dynamics alters spectral line profiles through both shifting and broadening, spatially varying chemistry changes line strengths (in an absolute and relative sense), and inhomogeneous temperature structures primarily impact line strength while subtly impacting broadening and line ratios as well.

8.3.2 K_P – V_{sys} maps

The results from our cross-correlation against 1D models are shown in Fig. 8.4. For both planetary cases, the “isotherm” model represents the baseline cross-correlation of our 1D CHIMERA template against a fully isothermal 3D atmosphere post-processed by our 3D code. Chemistry is held constant throughout the atmosphere (with mixtures determined by equilibrium chemistry at the planet’s equilibrium temperature at 1 mbar in pressure), with the only atmospheric dynamics imparted by the post-hoc 1D convolutional rotation kernel. In essence, it represents cross-

HD 209458b




Figure 8.4: K_P - V_{sys} maps from our grid of models. The text in the top left of each panel refers to K_P offsets from the known orbit of the planet and proper motion of its host star, the top right analogous offsets in V_{sys} , and bottom left signal-to-noise ratios. For each planet, we consider three families of spatial variation: temperature, chemistry, and atmospheric dynamics (see Table 8.2). We additionally include cross-correlation against a fully 3D model — i.e., the output from a GCM.

correlating two identical 1D models. All other models in the grid represent increasing amount of heterogeneity in the underlying model, while the 1D cross-correlation template is kept the same. Our grid therefore stepwise introduces dimensionality, building up from static, homogeneous atmospheres to fully 3D atmospheres as predicted by our two GCMs.

Overall, the spatially homogeneous models are associated with the strongest detection significances, with the more heterogeneous models yielding lower S/N. This case is of course a highly idealized case with overpredicted S/N, because real atmospheres have 3D structure (which we have removed for this step in the experiment). This result is partially by design, as the 1D template is selected to match the heterogeneous model. Even so, it is still instructive to quantify how each addition of 3D complexity reduces the detection significance, and how the shape and location of the cross-correlation peak changes. All of our maps feature net detections in excess of $S/N = 5$.

For most maps, the K_P - V_{sys} structure is as expected — an essentially vertical stripe that flares out at greater values of ΔK_P . This feature is due to the nature of the velocity constraints pre- and post-center of transit, where there is a zero-crossing in planetary radial velocity (Brogi et al. 2016; Wardenier et al. 2021). This flaring is larger for the ultra-hot Jupiter, which exhibits greater shifts in radial velocity over the course of its transit. Our simulations exhibit some offsets in K_P and V_{sys} . The distributions of these maps’ peaks are generally wider in K_P than in V_{sys} . Most of our models feature strong model anticorrelations at V_{sys} values immediately outside the peak, whereas these strong anticorrelations are not as present in the K_P direction.

A notable departure from this symmetry is the “cool limb water” chemistry model, which features a slope in K_P - V_{sys} . This pattern is indicative of the strong chemical heterogeneities in our spatially varying underlying atmosphere. The intuition is similar to the post-eclipse emission phase: the spectral signal is dominated by one hemisphere of the planet that is rotating away from

the observer. The slope is in this

An exception to the trend of decreased signal is the ultra-hot Jupiter isothermal case, which yields *lower* significance than the other, more thermally heterogeneous models. While at first perhaps unintuitive, this effect is likely due to the main absorber in IGRINS waveband being water. The isothermal temperature chosen for the model is indeed hot enough for the partial (thermal) dissociation of water, resulting in a lower detection significance than in the other models, which include regions with cooler temperatures.

Now we examine models with spatially varying temperature: as heterogeneity is introduced, the K_P-V_{sys} maps simply decrease their maximum signal strength. These models do not develop much in the way of strong velocity offsets, and the overall K_P-V_{sys} map structure remains similar. This behavior is due to changing temperature (in the absence of all dynamics) not changing the average line *positions*, but creating a greater mismatch in the line *profile* between the underlying simulation and the 1D cross-correlation template.

The dynamics family of models also shows subtle asymmetries and decreases in signal strength. Including rotation does induce a substantial Doppler *shift* and broadens the central K_P-V_{sys} peak, because a wider range of line-of-sight velocities is introduced at the limb. Note also that the faster-rotating ultra-hot Jupiter demonstrates a larger decrease in detection significance (relative to the isothermal model) than does the hot Jupiter. This is because the now-broad line profile is poorly matched by the 1D model, as its convolutional kernel cannot effectively reproduce solid-body rotation over a wide range of pressures. Including winds more strongly decreases the CCF signal (due to further smearing of the line profile from both rotational and divergent velocity structures) and creates a blueshift on the order of a few kilometers per second. Including both effects quite strongly decreases the K_P-V_{sys} significance because of the combination of both effects’

line profile distortions. The more highly irradiated ultra-hot Jupiter exhibits a stronger blueshift than does the hot Jupiter, owing to its faster day–night winds driven by greater instellation.

The fully 3D model, with its most complex line profile (from combining various lines of sight with varying chemistry, temperature, and atmospheric dynamics) has the lowest significance and has the most smeared-out K_P – V_{sys} map. For both planets, these models show offsets in both K_P and V_{sys} .

8.3.3 Retrievals

We now present the main results of this work: the impact of spatially variable atmospheres on 1D HRCCS *retrievals*.

For the 1D isothermal atmosphere, we correctly retrieve the input atmosphere (and zero velocity offset, scale factor of unity, etc.). We retrieve a bounded constraint on the H_2O abundance and upper bounds for the other gas abundances.

We next present selected results from varying each quantity (temperature, chemistry, and dynamics) in turn. We report results for both retrievals per model (i.e., we report the retrieval for each noise instance).

8.3.3.1 Temperature

Our retrieved isotherms are shown with respect to the input TP profiles in Fig. 8.5. Overall, the temperatures themselves appear to be more or less accurately retrieved. Spatially varying temperature alone did not significantly alter any of the retrieved velocity, cloud top pressure, or reference pressure parameters. The simple stretching of the spectrum by variable temperature

cannot be fit by the retrieval via line shifts (velocity), continuum changes (cloud top pressure and reference pressure), or changes to line (pressure) broadening (reference pressure).

Note, however, that the isothermal temperature is *overconstrained* for the spatially varying models. For the hot Jupiter model, both retrievals underestimate the spread in temperatures in the terminator region by nearly a factor of 2. This effect is even more exaggerated for the ultra-hot Jupiter thermal retrieval: the retrieved temperature is mildly biased toward hotter temperatures (by just over 1σ), and the spread in temperature is underestimated by a factor of more than 4.

For the hot Jupiter model, our simplest thermal departure from an isotherm — a radially varying TP profile — did not bias the retrieved abundances. Nor did varying temperature longitudinally or including the full GCM temperature field. The spatial variations in temperature (and how they interact with the opacity functions’ temperature dependence) is not drastic enough in this case to bias gas abundances.

However, including this spatially varying temperature did *widen* the gas abundance posteriors. For the first retrieval seed, varying temperature radially broadened the posterior from $\text{H}_2\text{O} = -3.19^{+0.5}_{-0.28}$ to $\text{H}_2\text{O} = -3.19^{+0.8}_{-0.28}$. In the second retrieval, the posterior broadened from $\text{H}_2\text{O} = -3.24^{+0.39}_{-0.26}$ to $\text{H}_2\text{O} = -3.25^{+0.77}_{-0.47}$ — roughly doubling the uncertainty of the inference.

This was not the case for all thermal variations, however. In the most extreme case — the first retrieval varying the full 3D structure of the ultra-hot Jupiter GCM — H_2O abundance was biased 3.3σ higher ($-2.13^{+0.18}_{-0.36}$ dex, with input -3.3 dex). This bias was significantly reduced in the second noise instance (down to 1σ). This result is due to the far greater thermal variation in the ultra-hot Jupiter GCM than the hot Jupiter; this greater variation is less easily captured by the relatively inflexible 1D forward model.

Taken together, spatially varying temperature fields can be incorrectly retrieved by isothermal




Figure 8.5: The retrieved and input temperature–pressure profiles. Retrieved isotherms (mean and 1σ distributions) are shown in vertical bars. The input temperature–pressure profiles are color-coded by morning and evening limb where applicable. For the GCM-input temperature structure, only limb profiles are shown.

1D retrievals, but they are often overconstrained due to the simplified parameterization. That is, the retrieval tends to recover the average temperature profile accurately in many instances, but not the intrinsic temperature spread. We found a number of instances in which abundance inferences are biased, corresponding to setups in which the temperature is most strongly inhomogeneous.

8.3.3.2 Chemistry

Our varying chemistry models exhibit departures from the homogeneous-chemistry model in a number of retrieved parameters. Specifically, this a posteriori mode is revealed with overly high gas abundances, rotational velocities, and cloud-deck amplitudes. This three-way degeneracy is a result of a competition to alter the line contrast.

The degree of bias depends on the degree of chemical variation. In the equilibrium chemistry model, the biased mode is only retrieved in one of the simulated retrievals (Fig. 8.6). In the model with water on only the cool limb, both retrievals are biased. Furthermore, the “cool limb water” retrieval revealed a tight and biased abundance for H_2S , which was not bounded in the isothermal model. That is, the sharper gradient in chemistry (water only on the cool limb) gave the most robust bias of the 1D model, which struggled to capture such a gradient in chemistry.

Both varying chemistry cases yield less biased posterior distributions when a scale factor is fit within the retrieval. This effect is due to the scale factor providing additional flexibility to the 1D retrieval, as the “dilution factor” does in JWST emission retrievals (Taylor et al. 2020).

In summary, heterogeneous chemistry can bias retrieved gas abundance, cloud-top pressure, and rotational velocity. A scale factor may mitigate this bias.




Figure 8.6: Select retrieved parameters from the HD 209458b chemical equilibrium model. Input values are indicated with solid black lines. We show the biases induced in the retrieved water, rotational velocity, and cloud-top pressure parameters. This experiment is repeated for two different noise instances (Retrievals 1 & 2).

8.3.3.3 Atmospheric dynamics

Fig. 8.7 shows the impact of our 3D rotation and winds on the retrieved velocity parameters. Across all seeds, and across both planetary models, the 1D broadening kernel, K_P , and V_{sys} parameters fail to capture the line broadening from winds and rotation. This failure is most evident in the behavior of the 1D broadening kernel. While the K_P and V_{sys} parameters are offset to reasonable degrees (on the order of the flow speed), the broadening kernel spins far faster than the input rotation speed. Whenever rotation is included, the kernel pushes against the fast-rotating edge of the prior. Our models indicate, therefore, that 3D velocity fields can impart significant biases on a 1D rotational kernel.

8.3.3.4 Misc.

Of our miscellaneous models’ results, the most stark is of the “full” model: performing our 1D retrieval on the full 3D output of GCMs. We find that that some biases do not interact (e.g., we see a combination of the biases presented above), while some work to cancel one another out. Noise appears to impact the bias of water abundance, whereas TP biases and offset velocity parameters are less subject to noise.

For our hot Jupiter model, the full model provided a combination of the biases presented—i.e., it had an overconstrained TP profile, broader water constraints, offset K_P and V_{sys} , and a very high rotational velocity.

The interaction of some effects *removes* some biases. For instance, the cloud-top pressure is no longer overconstrained, as it was in the chemical equilibrium model. Even so, for one retrieval, the cloud-top pressure posterior distribution does have a mode at 0.05 bar, and the constraint is

degenerate with the water abundance (where it is not with the isothermal retrieval).

The water abundance constraint is more complicated. For one retrieval, this full model yields an accurate water abundance constraint, while the other is nearly 4σ underabundant.

8.4 Discussion

We begin this section with a review of what parameters are biased in our simulations, and which are not. We then move on to discuss implications for retrievals, connections to the literature, implications for the Extremely Large Telescopes (ELTs), and caveats.

8.4.1 What are the 3D biases in high-resolution retrievals?

Our simulations indicate that 1D HRCCS retrievals can yield improper posterior distributions when applied to 3D atmospheres of varying complexity. We discuss this effect for each spatially varying quantity in turn.

Firstly, spatially varying temperature fields can be overconstrained by isothermal retrievals. This effect appears to be exacerbated in the presence of stronger spatial variation (e.g., in our ultra-hot Jupiter model compared to our hot Jupiter model). This issue has been reported in tests on real high-resolution transmission spectra (e.g., [Pelletier et al. 2023](#)).

Furthermore, variable temperature fields can bias inferred gas abundances, especially when the TP profile is not very flexible. This bias can be attributed to the changing line ratios shown in Fig. 8.3. With different lines now being stretched differentially (based on where in the atmosphere they become optically thick), an inflexible isotherm cannot stretch all model lines to match the data’s lines at the input abundance. So, the retrieval is biased away from the ground-truth

abundance.

Spatially varying chemistry can also induce biases in retrieved gas abundances. Once again, the Fig. 8.3 line profiles provide intuition. The equilibrium chemistry model has shallower line contrasts than the constant-chemistry model, as increased methane absorption along some lines of sight introduces a pseudo-continuum. HRCCS data processing removes the continuum, so it is difficult for the retrieval to uniquely identify methane as the culprit. In turn, the retrieval attempts to reduce the line contrast by a combination of raising the gray continuum (lowering cloud top pressure), broadening the line profile (increasing rotation), and decreasing the atmospheric scale height (increasing the mean molecular weight by greatly increasing the water abundance). Similar behavior emerges when depleting water in the warm limb—this model reduces the net water line contrast, which the retrieval then attempts to fit by other means (increasing rotation and cloud-top altitude). That is, the 1D retrieval struggles to fit decreased line contrasts due to spatially varying chemistry.

With respect to velocity parameters: our main result is that the 1D convolution kernel that we use cannot fully capture line broadening due to 3D winds and rotation. This effect is likely because of the fact that even “simple” solid-body rotation varies as a function of altitude, and high-resolution transmission spectroscopy probes a wide range of pressures (e.g., [Pelletier et al. 2023](#)), so a wide range of rotational broadening can be found. Winds can also introduce complex broadening to lines, as seen in Fig. 8.3. This effect has also been noted by [Flowers et al. \(2019\)](#).

We emphasize the three-way degeneracy that emerges in our HRCCS simulations as the line profiles become increasingly complex. As in the varying chemistry case, the broadening parameter, cloud top pressure, and atmospheric abundance can all increase or decrease spectral line contrast. This behavior reflects a unique consideration when porting code and intuition from space-based

retrieval studies.

8.4.2 What is *safe* from 3D biases in high-resolution retrievals?

Despite the arguably grim picture painted above, there are a number of situations in which it may be appropriate to apply 1D models to HRCCS data.

The first of these situations is the classic hot Jupiter population in a mildly noisy data regime. If the planet is too hot, rotation may be too fast, winds could be too fast, and the day–night temperature gradient could be substantial (e.g., [Komacek & Showman 2016](#)). At the ultra-hot Jupiter boundary, rotation may be slower, but there will be the greatest east–west temperature gradient ([Roth et al. 2024](#)). If the planet is too cold, the atmosphere may be more homogeneous (due to efficient heat redistribution), but the velocity change may be too low for HRCCS studies. Furthermore, at the coldest end, the atmospheric scale heights may be too small for adequate signal in transmission geometry, and synchronous rotation may be broken (i.e., the planet may be a fast rotator). So there may perhaps exist a “sweet spot” in the middle of the instellation axis for planets that move quickly enough, are hot enough, and yet are not “too 3D.”

A further consideration is “dynamically simple” ultra-hot Jupiters. The most egregious velocity biases emerged when both rotation and winds were included. Ultra-hot Jupiters by definition rotate (relatively, for the tidally locked population) quickly. The 1D models would be most appropriate, then, for a population in which the winds are “turned off,” limiting the degree of complex line broadening. This constraint can be met in the very strong drag limit, which allows only very sluggish winds.

While it may seem difficult to *a priori* know that a given ultra-hot Jupiter is dynamically

simple with near-infrared HRCCS, one could extrapolate from other measurements. For instance, JWST may indicate a very small hotspot offset for a planet in the near-infrared, and optical transmission spectra may yield very low Doppler shifts for a wide range of metallic species (Wardenier et al. 2021; Savel et al. 2022; Beltz et al. 2023). With some prior assessment, then, some ultra-hot Jupiters may still yield fruitful inferences with 1D kernels even in the high-SNR regime.

Finally, it is worth noting that using an isothermal model does not lead to biased abundances in comparison to a vertically varying TP profile. This result aligns well with what Pelletier et al. (2023) found when testing TP parameterizations against their data. In their study, while more restrictive, simpler TP profiles overconstrained the *retrieved TP profile*, they did not introduce biases in, e.g., the retrieved abundances. For those investigators who are interested primarily in inferring the composition of planetary atmospheres (as opposed to thermal structure, for instance), our results and those of Pelletier et al. (2023) indicate that an isothermal retrieval may be adequate—so long as there are very substantial planet-scale thermal variations.

The Gandhi et al. (2022) results also indicate that *chemical gradients* may not induce strong biases in HRCCS data. This takeaway is evident from the fact that their two-component models did not retrieve, e.g., an Fe abundance that was strongly discrepant from the 1D retrieval in their dataset. It may be, of course, that the specific planet they were studying (WASP-76b) simply does not have strong Fe chemical asymmetries across the terminator.

8.4.3 Are there any indications that a 1D retrieval is biased?

Perhaps. A scale factor biased away from unity, for instance, may imply that there is some spatial variation in chemistry that could be accounted for. Low scale factors have been recovered in

some HRCCS observations (e.g., [van Sluijs et al. 2023b](#)). Unfortunately, a non-unity scale factor may instead be due to observational considerations, e.g. highly variable tellurics ([Savel et al. 2025](#)).

Another clue of a biased 1D retrieval could be a retrieved rotational velocity that is much faster than synchronous rotation. This diagnostic is tricky, however, because it could be evidence of genuine atmospheric superrotation—which would be highly interesting to characterize.

This problem—that 1D retrievals may be biased to unknown degrees by 3D effects—is related to the idea of “stealth bias” in gravitational wave inference. Stealth bias refers to the idea that attempting to match data with an incorrect or incomplete template can induce subtle biases into posterior distributions that otherwise appear reasonable ([Cutler & Vallisneri 2007](#); [Vallisneri & Yunes 2013](#); [Vitale & Del Pozzo 2014](#)). A prevailing philosophy to address this problem is to simply fit data with more complex models, even if the additional physics in such models is low-amplitude. We recommend doing the same for HRCCS, especially in the high SNR regime.

8.4.4 How should we update retrievals?

Ultimately, we seek to extract 3D effects from data. The existence and relative strength of biases can help us motivate different types of retrieval updates. While we plan to investigate this topic more thoroughly in a future work, some candidates for retrieval upgrades include:

- *temperature*: radial variations are already captured in many retrieval codes, with, e.g., the [Madhusudhan & Seager \(2009\)](#) TP profile or the flexible parameterization of [Pelletier et al. \(2021\)](#). Ultra-hot Jupiters in particular seem ripe candidates for introducing day–night gradients (e.g., [Nixon & Madhusudhan 2022](#)).

- *chemistry*: even if the scale factor can mitigate chemical biases, it does not help us understand much of *what* is inhomogeneous. The inability of our retrievals to match smaller line contrasts motivates introducing two-component models (i.e., calculating two spectra that represent different aspects of the net observed spectrum). A simple approach could involve linear mixing, as is done for patchy clouds in space-based retrievals (e.g., [Line & Parmentier 2016](#)). A related approach could be motivated by [Wardenier et al. \(2021\)](#), who note that the (area-weighted) cross-correlation function is a linear operator—meaning that the cross-correlation functions from two separate limbs can simply be added together. We also suggest running retrievals that allow different gases to be separated in velocity space, which could be a proxy for spatial variation (e.g., [Brogi et al. 2023](#)).
- *dynamics*: a more flexible treatment is needed here, perhaps in the form of even more complex convolutional kernels. Such kernels should at minimum account for the varying projected velocity of solid-body rotation as a function of altitude. Ideally, these kernels would also allow for spatially varying wind profiles, as opposed to pure day–night flow. Note that there may exist degeneracies between these parameters and spatially inhomogeneous temperature profiles and abundance profiles.

Introducing these upgrades—implementation aside—will introduce challenges. These retrievals will be more computationally intensive, and they will include more parameters. The former can be alleviated by taking steps to decrease the number of samples required for a converged retrieval, e.g. by exploring the cutting edge of Bayesian sampling (e.g., [Lange 2023](#)). To prevent the overfitting associated with the an increased number of parameters, of course, rigorous statistical tests (e.g., Bayes factors, WAIC) must be required that naturally incorporate Occam’s razor.

These changes will be difficult to incorporate, but we posit that the underlying 3D information will make such efforts worthwhile.

8.4.5 How do these “3D biases” compare to those from low-resolution spectroscopy?

Some of the 3D biases discussed in this work are close analogs to those found for JWST datasets. For instance, as in our work, strong day–night temperature gradient can bias gas abundances in JWST retrievals (Caldas et al. 2019; Lacy & Burrows 2020; Pluriel et al. 2020; Nixon & Madhusudhan 2022). Furthermore, as in Caldas et al. (2019), the exact degree of bias depends on the exact simulation.

Interestingly, we find that the TP profiles in isothermal HRCCS retrievals are biased less by a dayside “shielding” effect than is discussed in the JWST literature. The 3D thermal structure ultra-hot Jupiter had only a mild bias toward hotter temperatures (consistent across both spectral noise instances). We would perhaps find a stronger temperature bias if we systematically explored day–night contrasts as, e.g., Caldas et al. (2019) do. It is also possible that our models are less biased toward the dayside because HRCCS probes higher altitudes than JWST does, and the geometry favors the terminator at higher altitudes.

Our result for the scale factor has connections to JWST retrievals, as well. Taylor et al. (2020) found that multiplying a single 1D emission spectrum with a “dilution factor” could account for the inhomogeneities of an underlying two-component atmosphere. We find a similar result for *chemistry*, in that retrieving on a chemically inhomogeneous atmosphere with a scale factor can help prevent biases. A notable distinction in HRCCS is that observational considerations can *also*

stretch or shrink the signal (including data processing steps or variable tellurics; e.g., [Savel et al. 2025](#)), beyond the “missing physics” that this approach would also marginalize over for JWST data.

Our study also, of course, has the extra dimension of velocity fields, which exacerbates some of the 3D complexity of the problem. In particular, the broadening kernel seems to participate in a degeneracy building upon the well-known cloud-top pressure and abundance degeneracy (e.g., [Welbanks & Madhusudhan 2019](#)). In the purely isothermal case this degeneracy matters less, but in the regime where line profiles are not fit perfectly by the 1D (i.e., where the underlying atmosphere is fundamentally 3D but a 1D model is being used), this discrepancy becomes more prominent.

8.4.6 How sensitive are K_P – V_{sys} maps to 3D effects?

Our work essentially recovers the intuition provided by [Flowers et al. \(2019\)](#) and [Beltz et al. \(2020\)](#): *that 1D models cannot match intrinsically 3D atmospheres in high-resolution spectroscopy*. While this point may at first seem obvious — a mismatch between a cross-correlation template and the signal should decrease the detection significance — it is not given that *physically plausible* departures from homogeneity exceed the HRCCS noise floor. The experiments in [Flowers et al. \(2019\)](#), [Beltz et al. \(2020\)](#), [van Sluijs et al. \(2025\)](#) and this work firmly demonstrate that spatial heterogeneity alters 1D inference at the detectable level.

Beyond the decreased detection significance, not much in the K_P – V_{sys} map structures alone provides insight into the sources and strengths of “3D-ness.” Some of our models exhibit velocity offsets, but it is not evident whether they are due to changing chemistry, dynamics, or temperature. That is, the alterations to K_P – V_{sys} maps generated with 1D cross-correlation templates do not

appear to be unique and can even be mimicked by noise, so it is difficult to diagnose 3D effects using them alone. This degeneracy reinforces the idea that running full 3D models to compare to data is more informative than using toy models when the aim is disentangling 3D effects in HRCCS (Wardenier et al. 2021; Savel et al. 2022; Beltz et al. 2023; Wardenier et al. 2024).

On the more optimistic side, this result implies that, in the high signal-to-noise regime, K_P – V_{sys} maps may not require highly computationally intensive or highly tuned models to answer first-order questions about exoplanet atmospheres, e.g., whether water is present. That is, in our models (absent correlated noise, systematics, etc.), a simple K_P – V_{sys} maps provides robust enough evidence of the presence of a species—as long as offsets are not interpreted to be solely due to dynamics, for instance. Note that in the lower signal-to-noise regime, though, 3D effects may be sufficient to drop the K_P – V_{sys} detection below the formal detection threshold.

One interesting departure from this picture is our “cool limb water” case. In this model, we recover a noticeable slant in the K_P – V_{sys} map, and we produce strong K_P offsets. The latter are especially indicated in Wardenier et al. (2023) for inhomogeneous atmospheres. Our simulations agree with the analytic arguments of Wardenier et al. (2023)—that K_P offsets within a reasonable range can be generated from chemical inhomogeneities alone.

In terms of changing signal strength in this controlled experiment, it seems as though changing line ratios (via spatially varying chemistry) and line strengths (via spatially varying temperature) impacts signal strength comparatively less. What matters most is changing line shape and position—the effects of atmospheric dynamics. These shifts, importantly, cannot be matched by a single shift of a 1D spectrum in velocity, because individual lines are shifted differentially. Again, however, this is not an obviously *unique* change in significance, and inferring the details of atmospheric dynamics from K_P – V_{sys} alone (beyond potentially attributing some V_{sys} offset to winds)




Figure 8.7: Select retrieved parameters from the inhomogeneous winds and rotation models. Input values are shown with a solid black line. We show the biases induced in the K_P , V_{sys} , and rotational velocity parameters.

remains nontrivial and warrants further study.

8.4.7 What are the implications for the Extremely Large Telescopes?

The Extremely Large Telescopes (ELTs; e.g., [Gilmozzi & Spyromilio 2008](#); [Johns et al. 2012](#)) will present both opportunities and challenges to HRCCS by producing even higher-SNR spectra of exoplanet atmospheres. While it may seem as though these data will immediately translate to tighter constraints on low-SNR quantities of interest (e.g., carbon isotope ratios; [Line et al. 2021](#)), such high SNR may come with a cost. In particular, it stands to reason that the previously described line profile discrepancies and their associated retrieval biases will be placed into even starker relief in this noise regime. Our high-SNR simulations indicate that 1D models will face significant difficulties fitting 3D data. With the ELTs’ first light dates approaching in less than half a decade, *it is evident that we will need to upgrade our models to accurately fit our data.*

Other astrophysical fields have struggled to fit very new, high signal-to-noise data, with model uncertainty dominating the error budget (e.g., in the brown dwarf community; [Gonzales et al. 2022](#)). Statistically, these difficulties can be circumvented by introducing tolerance parameters (i.e., rescaling the *measured* spectral uncertainties to implicitly account for *model* uncertainty; e.g., [Line et al. 2015](#); [Gandhi et al. 2023a](#)) that may even explicitly incorporate model uncertainty ([Piette & Madhusudhan 2020](#)). In the case of exoplanet atmospheres, we have demonstrated that a significant portion of the upcoming model uncertainty budget will be due to *known* “missing 3D physics.” We therefore will soon have an opportunity to learn about 3D effects—provided that we prepare in advance.

8.4.8 Caveats

- We include equilibrium chemistry as our “most inhomogeneous” chemistry model, but we would like to emphasize that equilibrium chemistry is not necessarily the *most physical* chemistry prescription. GCMs that include chemical kinetics modules often find spatially homogenized chemistry when allowing for chemical transport, especially for planets at moderate instellations (e.g., [Zamyatina et al. 2023](#)).
- Our GCMs do not include two effects that are known to control heat transport in ultra-hot Jupiters: H_2 dissociation / recombination ([Bell & Cowan 2018](#); [Komacek & Tan 2018](#); [Tan & Komacek 2019](#); [Mansfield et al. 2020](#)) and cloud feedback ([Parmentier et al. 2021](#); [Roman et al. 2021](#); [Komacek et al. 2022](#)). So, the heat transport for our ultra-hot model misses some physical components—though the model remains a complex, physically motivated examples of spatially variable atmospheres (and we note that no GCM yet includes all physics relevant to hot Jupiters).
- Our GCMs employ double-gray radiative transfer. See [Lee et al. \(2021\)](#) and [Malsky et al. \(2024\)](#) for in-depth comparisons of these methods. TP profiles calculated under the double-gray approximation necessarily become isothermal in the upper atmosphere. This behavior (along with our isothermal extension) is likely unproblematic when the end goal is to calculate transmission spectra, which are not particularly sensitive to TP gradients.
- We have limited our study to the hydrostatic regions of exoplanet atmospheres. There exists considerable evidence that HRCCS observations are sensitive to atmospheric outflows (e.g., [Seidel et al. 2021a](#)). We do not consider the biasing impact of outflow velocity fields, nor do

we consider the optically thin abundance field that can extend up to and beyond the planetary Roche lobe. Our results are therefore valid when considering 3D biases purely within the deeper, hydrostatic regions of atmospheres.

- We only consider one instrument (IGRINS). This approach limits our results to a single wavelength range and resolution. It is certainly possible that 3D biases are further exacerbated at higher resolution (wherein more lines probing different pressures are resolved in finer detail) and / or in different wavelengths (the optical, for instance, in which the intrinsic atomic line profiles are narrower).

8.5 Conclusions

HRCCS has emerged as a promising technique to explore the composition and dynamics of exoplanet atmospheres. Yet both theory and observations indicate that our current 1-dimensional approach to HRCCS inference may be limited. To assess the ways in which spatial variation can bias 1D high-resolution atmospheric retrievals, we perform a systematic 1D retrieval investigation of different types of “3D-ness” in self-consistent and parameterized atmospheres.

In our simulated retrievals, we find that:

- isothermal retrievals on 3D atmospheres overconstrain the temperature, often leading to abundance biases. Unlike low-resolution spectroscopy, the temperature is not necessarily biased toward dayside conditions, perhaps because HRCCS probes lower pressures.
- spatial variations in chemistry can lead to overconfident and biased gas abundance constraints. A scale factor can eliminate these biases in the context of our simulations.

- often-used convolutional kernels cannot capture line-broadening effects due to rotation and exoplanetary winds.
- there exists a three-way degeneracy between broadening kernels, gas abundance, and cloud-deck pressure in the presence of noisy data and complex line profiles.

It is likely true that all the hot Jupiters orbiting bright, nearby stars have been discovered. In the coming years, the frontier of HRCCS may therefore naturally shift toward deeper studies of already characterized targets. The 3D effects discussed here and elsewhere in the literature may therefore prove to be one of the most fruitful avenues for understanding the nature and origins of these enigmatic objects.

Chapter 9: Conclusions and future work

9.1 Summary

The work contained in this dissertation reveals and seeks to better extract the information in spectroscopy of hot Jupiters, with the ultimate aim of revealing these exoplanets' 3D nature. In Chapter 2, I challenged a prevailing interpretation of the archetypal asymmetric exoplanet, WASP-76 b. I spectrally post-processed a suite of general circulation models to forward model the planet's high-resolution optical spectrum. I then cross-correlated these 3D models with spectral templates to generate synthetic Doppler shifts. Our self-consistent models could match neither the shape nor the magnitude of the Doppler shifts reported by [Ehrenreich et al. \(2020\)](#), even when we modeled iron condensation in the atmosphere — the hypothesis favored by the observational team. We instead flesh out two compelling alternatives for the asymmetry in the Doppler shifts: a deck of optically thick clouds and a small, unaccounted-for orbital eccentricity. This work demonstrated that high-resolution spectroscopy can clearly rule out climate models (such as strong global fluid drag) It also began to highlight the degeneracies of this problem and the difficulties posed for climate models.

Extending the work from Chapter 2, Chapter 3 examined the population of hot Jupiters and the range of possible mechanisms that may drive their spatial asymmetries. The main result from this work is that the chemical stability of carbon monoxide makes it an excellent tool for disentangling possible drivers of asymmetry. Because the molecule has relatively similar abundance over a wide range of temperatures and pressures, measured asymmetries in the CO signal are not likely due to chemical gradients. Rather, they imply gradients in the thermal and/or velocity fields.

Chapter 4 illustrated that magnetic fields may *also* drive measurable differences between

gas velocities in exoplanet atmospheres. Despite controlling an array of planetary properties — from atmospheric escape to habitability — the strengths of exoplanet magnetic fields are unknown to orders of magnitude. I demonstrated that the so-called “magnetic drag” picture implies that, in certain regimes, ions and neutrals should have different velocities — a potential diagnostic of magnetic field strength. This approach opens up a new line of interpreting HRCCS observations and provides additional context for performing 3D inference based on wind measurements alone.

Chapter 5 showcased a tool that reduces the RAM of gas opacity data (a limiting factor for HRCCS modeling) by up to an order of magnitude. The key necessary observation was that we need not save the full opacity structure to disk; rather, we can save polynomial fits to the opacity function and re-evaluate the function on the fly. By doing so, I enabled more complex inference on cheaper hardware. Given the increased computational complexity of multi-dimensional inference, this would be particularly useful for practically scaling up these next-generation retrievals.

In Chapter 6, I built a framework for connecting my models more directly to observations by including realistic sources of noise. To do so, I built an end-to-end simulator of high-resolution spectroscopy, including losses and correlated noise from the star, telluric absorption, and the instrumental blaze function. I then used this tool to demonstrate potential failure modes of high-resolution observations. In particular, this work showed that highly variable tellurics can bias inferences made with this ground-based technique, as the data-cleaning procedure assumes a quasi-stationary telluric signal.

I next connected HRCCS inference methods to JWST via cross-correlation retrievals of JWST high-resolution modes in Chapter 7. This procedure yields more confident and more precise estimates of CO abundance. Most importantly, this method can measure a reliable C/O ratio with a single instrument mode; the conventional wisdom is that broader wavelength coverage is required.

I look to scale this method to the vast trove of relevant JWST datasets.

With mock 1D retrievals on a hierarchy of 3D atmospheres (from isothermal models to GCMs), Chapter 8 demonstrated the current limitations of HRCCS. For instance, treatments of planetary rotation are too inflexible, and retrieved chemistry is biased by 3D thermal structure.

Together, these chapters establish that 3D inference of hot Jupiter atmospheres is both necessary and tractable.

9.2 Future directions

I believe that my work and the work of the field will trend in the following directions in the coming years.

9.2.1 The 3D processes of exoplanet atmospheres with high-resolution spectroscopy

A particularly exciting direction for extending this work would be developing, applying, and scaling **3D forward models** to use in HRCCS retrievals. My work has demonstrated that these models can be developed and should be applied; the natural next step is to do so. Given the surprises that exoplanets and their atmospheres have had in store since their first discovery, I am particularly in favor of creating flexible parameterizations of the quantities we seek to extract — for instance, 3D spherical harmonics to parameterize the temperature field. This approach would let the data “speak for themselves.” The results of these retrievals of course can subsequently be compared in detail to grids of, e.g., general circulation models. My intuition is that convolutional kernel-based approaches, while rapid to evaluate, are not flexible enough to capture the full com-

plexity of the velocity field. In particular, they lack the ability to parameterize *radial* differences in the velocity field, which are perhaps one of the most physically interesting types of velocity differences.

Such a 3D model, once developed (and perhaps accelerated with frameworks such as JAX; [Bradbury et al. 2018](#)), could be applied to the now-vast archive of HRCCS datasets. Beyond studying individual objects, I would therefore hope to also learn about how these quantities of interest — e.g., the spatial distribution of temperature and wind — scale with fundamental planetary properties. Pure hydrodynamic simulations expect that wind speeds should tick upward with increasing instellation ([Komacek & Showman 2016](#)). This scaling may not hold if flow is instead regulated by magnetic effects, which are also expected to increase with instellation ([Beltz et al. 2021](#)). No significant trend of wind speed vs. instellation has yet been found in HRCCS observations ([Gandhi et al. 2023b](#)). This non-detection may be due to limitations in the data, and/or it may be due to limitations in the models. Future population-level work with more complete models on better data may yield further insights.

From the GCM perspective, much of the first-order physics seems to be worked out quite nicely at this time. Yet challenges remain. For example, the field now has quite advanced models of magnetohydrodynamics and its effect on circulation ([Beltz et al. 2021](#); [Christie et al. 2025](#)). However, none of these models couple realistic radiative transfer to the full magnetic induction equation, primarily due to the numerical difficulties of capturing the vast day–night conductivity gradient. The question of whether magnetic fields can truly drive large-scale atmospheric variability therefore remains out of reach. Furthermore, there is not yet one model that can capture all such effects — in particular, non-gray radiative transfer, hydrogen dissociation / recombination, gas-phase chemistry, magnetic fields, and cloud microphysics ([Kataria et al. 2013](#); [Komacek &](#)

[Tan 2018](#); [Beltz et al. 2021](#); [Zamyatina et al. 2023](#); [Lee & Ohno 2025](#)). This gap presents a very difficult (and interesting!) modeling challenge that I anticipate will occupy the field for a number of years.

One particular piece of physics that I think is missing from interpreting wind observations with HRCCS is upper atmosphere dynamics. All the models that I have described thus far are in hydrostatic equilibrium, from the retrievals' forward models to the GCMs. Planetary bodies need not — and, in fact, do not — obey this assumption. At low pressures, appreciable amounts of gas can be lost from planetary atmospheres at appreciable velocities ([Murray-Clay et al. 2009](#); [Owen 2019](#)). Such effects have been observed in exoplanets for decades now, both spectroscopically ([Vidal-Madjar et al. 2003](#)) and indirectly via trends in exoplanet demographics ([Fulton et al. 2017](#)). The common spectral tracers of atmospheric escape are hydrogen and helium, which are abundant, low mass, and can easily absorb high-energy photons. Other gaseous species, however, are likely dragged along in these flows. Evidence for this effect includes fast inferred radial outflows in the Na line ([Seidel et al. 2021a](#)) and very strong Ca absorption ([Deibert et al. 2021](#)). Less direct evidence that points to model incompleteness includes data that exceed GCM-predicted velocities ([Savel et al. 2021](#); [Wardenier et al. 2021](#); [Beltz et al. 2023](#)). Considering the radiative and dynamical impacts of upper atmospheres on emerging HRCCS spectra would require uniting subfields that are traditionally distinct. Doing so may resolve a few key extant observational problems.

9.2.2 High-resolution spectroscopy

HRCCS is rapidly moving into an intriguing regime. This technique is currently best suited for large, hot planets on fast-moving orbits around bright host stars. These features naturally make deep investigations of single hot Jupiters obvious targets. But large-scale surveys struggle as a result — there are only so many of these planets around bright stars, and the most observable targets already boast comprehensive atlases of atomic and molecular detections. So what is next?

The frontier of slower-moving exoplanets may be the challenge of choice. These are less-obvious targets, as traditionally the differential measurement in HRCCS is the movement of the planetary spectral signal across the detector. While it appears like low-velocity planets may simply only be accessible with JWST, I believe that there is real potential for HRCCS to improve on this front. One lesson lies in WASP-127 b as observed with CRIRES+. This dataset initially seemed to contain no useful signal, until it was realized that nodding the telescope has yielded slightly different wavelength solutions for subsequent exposures. Separately, excellent progress is being made in directly modeling tellurics (e.g., [Leet et al. 2019](#)) and utilizing out-of-transit spectra ([Cheverall & Madhusudhan 2024](#)) to maximally remove quasi-stationary correlated noise from these datasets. I believe that with extra data processing capabilities, we will be able to extract some of these slower-moving planets’ signals. Doing so will open up new science cases beyond the hot Jupiter population — such as exploring the atmospheres of the curious Neptune Ridge population of planets ([Castro-González et al. 2024](#)).

The ELTs, of course, will drastically change this picture. With much larger apertures, they will open up a much wider volume of space for detailed atmospheric characterization. This will be highly beneficial for executing true population-level hot Jupiter science ([Palle et al. 2025](#)). The

field may also experience an interesting problem: our highest signal-to-noise targets may have *too* high-quality data. While optimistically, we will be able to constrain properties such as atmospheric vertical wind shear ([Kempton & Rauscher 2012](#)), equatorial jet width, and more, we may run into problems that the brown dwarf community faces with fitting such pristine datasets (e.g., [Piette & Madhusudhan 2020](#)).

The ELTs will enable characterization of planets with signals smaller than the typical hot Jupiter. Photoevaporation ensures that there do not exist many hot Neptunes and hot sub-Neptunes — as manifests in the so-called Neptune Desert ([Lundkvist et al. 2016](#); [Mazeh et al. 2016](#)). Some of the most observationally interesting targets in the ELT era may end up being lava ocean planets ([Léger et al. 2011](#)). They move fast, they are very hot, and they would open up a wide discovery space: what are the interiors of these rocky planets like? Mass–radius inference based on interior modeling is notoriously degenerate. Can observations of atmospheres exhaled by magma oceans help? How efficient is ongoing atmospheric escape for terrestrial-mass exoplanet atmospheres? Can a partially collapsed atmosphere be contained? Are these objects variable, perhaps due to clouds? Lava ocean planets represent another opportunity to test our Solar System theories, this time in the rocky planet regime and using physics and chemistry learned from the hot Jupiter population.

9.2.3 Inferring the properties of exoplanet atmospheres

As our data quality improves, our retrieval methodology may have to adjust accordingly. The “reference model problem” — that claimed detections are sensitive to the models used to compare against one another — is particularly relevant ([Welbanks et al. 2025](#)). I am eager to explore how

to efficiently test for the presence of numerous potential gases while remaining in the Bayesian framework that has proven to be so useful for this field.

I am also interested in extending my work on JWST cross-correlation retrievals. This approach brings us “closer to the data,” and I believe it warrants a reanalysis of existing (high-resolution) JWST exoplanet spectroscopy. I would particularly like to bring this modeling even closer to the raw data, for instance by simultaneously extracting all lightcurves and fitting the atmospheric models.

Finally, I am curious to see if the field begins to move away from retrievals as the “gold standard” as data quality continues to improve. For instance, the brown dwarf community has introduced methods for identifying cloud species without relying on retrievals ([Mollière et al. 2025](#)), and the JWST limb asymmetries field is performing retrieval-free comparative studies ([Fu et al. 2025](#)). It may be worthwhile to ask ourselves: **when do the statistics actually serve us, and when are we overcomplicating matters with models we know are incomplete?** This question goes back to what I explored in Chapter 2 — we know that the GCMs do not perfectly match the data, and yet we could still loosely constrain the atmospheric climate state.

9.2.4 The study of exoplanet atmospheres

We certainly understand exoplanet formation better after thirty years of study. Incorporating clues from planetary architecture will perhaps be the next step in making concrete claims tying current-day exoplanet atmospheres to their formation conditions (e.g., [Kirk et al. 2024](#)). This approach will not only give exoplanet atmospheres more *leverage* — i.e., can answer its questions more effectively — but it will also help make atmospheric spectroscopy a more sustainable line of

inquiry in the long term by connecting this field to adjacent ones. I also look forward to seeing more observations of refractory-to-volatile ratios, which may prove to be a more reliable tracer of planet formation than the C/O ratio (e.g., [Lothringer et al. 2021](#); [Smith et al. 2024b](#); [Pelletier et al. 2025](#)).

As previously mentioned, a “small planet frontier” is approaching as our spectroscopic data become more precise. This will naturally give rise to one of the key motivations for studying atmospheres that I have neglected until now: finding life beyond Earth. We may still be decades out from robustly detecting “biosignatures” (if we ever do), as it appears that the Habitable Worlds Observatory (HWO) will be required to make these assertions ([Currie et al. 2023](#); [Hardegree-Ullman et al. 2023, 2025](#); [Seager et al. 2025](#)).

There will still be work to do in the intervening years. Planetary systems are interconnected. Take the Earth, for instance: the biosphere interacts with geology and atmospheric chemistry. An alien observer with their own HWO would not just have to measure our atmosphere; they would have to interpret their measurements in the context of the full planetary system to rule out false positives ([Schwieterman et al. 2018](#)). I am skeptical that this field will ever reach a slam-dunk “chirp” like the gravitational waves community ([Abbott et al. 2016](#)), but I do believe we will be able to say something statistical. The field is developing frameworks for interpreting and communicating the potential nature of biosignatures ([Catling et al. 2018](#)). That is, we will not be able to answer the question: “Where is life?” But we may be able to answer: “In our local slice of the Universe, are we alone?”

Appendix A: Software used in this dissertation

- `astropy` ([Astropy Collaboration et al. 2013, 2018, 2022](#))
- `batman` ([Kreidberg 2015](#))
- CHIMERA ([Line et al. 2013](#))
- `corner` ([Foreman-Mackey et al. 2016](#))
- `cortecs` ([Savel et al. 2024a](#))
- `emcee` ([Foreman-Mackey et al. 2013](#))
- `exoplanet` ([Foreman-Mackey et al. 2021c](#))
- `flowMC` ([Wong et al. 2022](#))
- FASTCHEM ([Stock et al. 2018, 2022](#); [Kitzmann et al. 2024](#))
- GitHub Copilot ([Chen et al. 2021](#))
- HELIOS ([Malik et al. 2017, 2019](#))
- HELIOS-K ([Grimm & Heng 2015](#))
- IPython ([Perez & Granger 2007](#))
- JAX ([Bradbury et al. 2018](#))
- Jupyter ([Granger & Pérez 2021](#))
- `matplotlib` ([Hunter 2007](#))
- MultiNest ([Feroz et al. 2009](#))

- `nautilus` ([Lange 2023](#))
- `numpy` ([Harris et al. 2020](#))
- `Numba` ([Lam et al. 2015](#))
- `pandas` ([McKinney 2010](#); [pandas development team 2020](#))
- `PLATON` ([Zhang et al. 2019, 2020, 2025](#))
- `PyMultiNest` ([Buchner et al. 2014](#))
- `scipy` ([Virtanen et al. 2020](#))
- `scope` ([Savel et al. 2025](#))
- `spectres` ([Carnall 2017](#))
- `tqdm` ([da Costa-Luis 2019](#))
- `VULCAN` ([Tsai et al. 2017](#))

Appendix B: Facilities used in this dissertation

- zaratan, the UMD supercomputer
- JWST
- IGRINS

Appendix C: Appendix from Chapter 2

Our procedure to fit our eccentricity-inclusive model to the [Ehrenreich et al. \(2020\)](#) blueshift data is as follows:

1. Initialize an `exoplanet KeplerianOrbit` object with stellar and orbital parameters from [Ehrenreich et al. \(2020\)](#). Additionally, compile a `Theano` ([Theano Development Team 2016](#)) function associated with the `KeplerianOrbit` object that takes eccentricity and longitude of periastron as inputs.
2. Write a likelihood function for our two orbital parameters (eccentricity and longitude of periastron) and a parameter characterizing the degree of error bar underestimation ($\log(f)$), assuming Gaussian errors on the data.
3. Identify the global minimum of the negative log likelihood using the `minimize` routine from `SciPy`. We find a minimum at $e = 0.019$, $\omega = 123.1^\circ$, $\log(f) = -1.84$ for our cloudless model and a minimum at $e = 0.018$, $\omega = 115.6^\circ$, $\log(f) = -1.89$ for our optically thick cloud model.
4. Initialize 32 `emcee` walkers in a small Gaussian ball around the global minimum identified in the previous step. Repeat this step separate runs with two different priors: a uniform prior over the relevant parameter ranges and a prior using the constraints provided by [Fu et al.](#)

(2021) (see Table C.1).

5. Run the `emcee` walkers for 20,000 steps.
6. Visually identify whether the walkers effectively explore the parameter space via trace plots.
7. Check that the sampler’s acceptance fraction is between 10% and 90%. All of our chains have acceptance fractions near 62%.
8. Discard the first few multiples of the autocorrelation time as “burn-in.” We generally discard 1,000 samples, which is more than sufficient for this criterion.
9. Thin the samples by a factor of half the autocorrelation time (≈ 20) to generate reasonably independent samples for the posterior distribution.
10. Assess convergence using the [Geweke \(1992\)](#) criterion.

We find that our results are relatively insensitive to our choice of prior (of the two explored; Figure C.1) and pass our diagnostic tests.

As a separate note, both our cloudless and cloudy MCMC runs converge on similar median eccentricities. The eccentricities in both runs meet the $e > 2.45\sigma_e$ metric for statistically significant nonzero eccentricity defined by [Lucy & Sweeney \(1971\)](#), which accounts for there being zero phase to explore at exactly zero eccentricity.




Figure C.1: Posterior corner plot corresponding to our 10 scale height Al_2O_3 cloud MCMC run. Both priors result in similar posteriors, implying that the posterior is heavily influenced by the data. Quantities are reported for the asymmetric normal prior (drawn from the [Fu et al. 2021](#) constraints).

Table C.1: MCMC values

Model	Param	Priors	Fit value	l_{auto}	Indep. samples
No clouds, weak drag, deep RCB	e	$\mathcal{U}(0, 1),$ $\mathcal{P}(0.016, 0.011, 0.018)^{\dagger}$	$0.018^{+0.004}_{-0.004}$	46, 46	433, 434
—	ω	$\mathcal{U}(0^\circ, 180^\circ),$ $\mathcal{P}(62, 82, 67)$	$125^{+9.87^\circ}_{-6.87^\circ}$	47, 48	417, 409
—	$\log(f)$	$\mathcal{U}(-10, 1)$	$-1.79^{+0.20}_{-0.20}$	41, 43	485, 459
Al_2O_3 clouds, weak drag, deep RCB, cond. on	e	$\mathcal{U}(0, 1),$ $\mathcal{P}(0.016, 0.011, 0.018)$	$0.017^{+0.004}_{-0.004}$	44, 46	450, 432
—	ω	$\mathcal{U}(0^\circ, 180^\circ),$ $\mathcal{P}(62, 82, 67)$	$117^{+8.43^\circ}_{-5.92^\circ}$	48, 43	412, 457
—	$\log(f)$	$\mathcal{U}(-10, 1)$	$-1.84^{+0.20}_{-0.20}$	39, 42	504, 467

Appendix D: Appendix from Chapter 6

D.1 Optimization

The simulations required for HRCCS retrievals are computationally intensive. At its core, the issue is that comparing simulations to datasets from modern echelle spectrographs with large spectral grasps requires solving the radiative transfer equation at hundreds of thousands of wavelength points. Previous studies combated this the runtime issue by sampling a representative wavelength range within the instrumental band (e.g., [Kempton & Rauscher 2012](#); [Wardenier et al. 2021](#)), with simple parallellization over multiple CPUs (e.g., [Savel et al. 2021](#)), or by parallelizing on GPUs ([Line et al. 2021](#); [Lee et al. 2022](#)). The parallelization advances in particular have taken advantage of the fact that the radiative transfer problem can cast as embarassingly parallel over wavelength (or wavelength packet, for Monte Carlo approaches).

To make the HRCCS retrievals in this work computationally feasible enough for repeated use and testing, we undertook an effort to decrease the retrieval code’s computational runtime. Across all optimizations, we decreased the runtime from 8 days, 24 CPUs, and 1 GPU to 15 hours, 1 CPU, and 1 GPU.

D.1.1 Optimization with JAX

JAX ([Bradbury et al. 2018](#)) is a code transformation framework in Python. It provides a number of features that enable faster statistical sampling, including automatic differentiation (e.g., [Margossian 2019](#)), just-in-time (JIT) compilation, and GPU acceleration. Because of the code team’s main focus being machine learning on accelerators such as GPUs, JAX does not trivially allow common control flows such as if statements. This feature adds nontrivial development cost to port codes to the framework. Previous studies have used JAX for retrievals of exoplanet atmospheres ([Kawahara et al. 2022](#)); to our knowledge, this work marks the first time that JAX has been used in HRCCS retrievals. Using JAX alone successfully decreased the retrieval runtime, but it also allowed gradient-based sampling approaches that reduced the number of samples required for retrieval convergence (see Appendix [D.1.2](#)).

Other speed-ups were realized in the process of porting the code to JAX. For instance, we adjusted the code to use linear interpolation of Doppler-shifted spectra via `numpy` instead of a `scipy` spline fit; this step was necessary for the JAX migration because JAX does not support implicit tracing of external libraries.

Previously, the likelihood evaluation of a single atmospheric model was originally CPU parallelized across exposures (i.e., across planetary orbital phases); however, the JAX optimization is fast enough that it can be done without resorting to explicit parallelization with, e.g., `joblib`. The underlying radiative transfer calculations on the GPU also incurred a 3x speedup, but this is marginal compared to the total evaluation time of a single likelihood call.

One caveat to the JAX speed-up of this code is the compilation time, which can run in excess of an hour. This cost is negligible, however, in the context of runtimes that generally exceed

days. Recent JAX development indicates that compiled code can be cached and reused, so this compilation cost will likely not be a major concern in the future.

Another unique aspect of JAX is that it may copy the memory of global objects. To prevent this behavior while allowing functions to access memory that is only read in once, we partial the opacities into the function definitions themselves.

D.1.2 Using `flowMC`

`flowMC` (Wong et al. 2023b) is a JAX-based statistical sampling package that couples a local (often gradient-based) sampler with a non-local sampler. Because this work (to our knowledge) marks the first usage of exoplanet atmospheric retrievals with `flowMC` and a normalizing flow non-local sampler, we detail in this section a set of tests that we conduct to verify the utility of these techniques for our science case.

`flowMC` essentially adds a non-local proposition to a Markov chain Monte Carlo (Gelman et al. 2004) proposal step in conjunction with the standard local step. While local-stepping is known to converge, reaching a converged state can heavily consume compute resources in the presence of nontrivially shaped posterior distributions (Gelman & Rubin 1992; Tierney 1994; Neal 2003). This non-local proposition is designed to learn the complex structure in the posterior geometry; this learning is achieved on the fly with a normalizing flow. Normalizing flows are a machine learning technique—a type of generative model that can be sampled (e.g., Kobyzev et al. 2020; Papamakarios et al. 2021). This model can also be *used* to compute the likelihood of its samples. They operate by performing a change of variables to map from the current sampler location to another posterior mode—the target distribution. This process improves the acceptance of MCMC

proposals as the sampling continues (e.g., [Gabri  et al. 2022](#)). When combined with a gradient-based sampler, this sampling scheme is well-suited for high-dimensional, multimodel sampling problems (e.g., [Betancourt 2017](#); [Mangoubi et al. 2018](#); [Wong et al. 2023a](#)).

The critical intuition is that, even a normalizing flow that is not trained perfectly will allow walkers to “teleport” from one mode to another, with the “teleportation” probability weighted by the target mode’s probability mass. Hence, normalizing flows are able to quickly construct the relative mass of modes in a posterior, at the cost of alternating between sampling the posterior and training the normalizing flow on new samples.

First, we test that our normalizing flow can approximate the topology of a converged `PyMultiNest` run. For this study, we make use of the RQSpline normalizing flow architecture. With this check successfully verified (Fig. [D.1](#), Moving on to retrievals: We find that running `flowMC` produces posterior distributions that match the `PyMultiNest` posteriors. Given that our parameters can span different orders of magnitude (e.g., the water abundance vs. temperature at base of the atmosphere), we tune the step size (via the “mass matrix”) of our local MALA sampler ([Grenander & Miller 1994](#)) to be inversely proportional to the gradient of the likelihood at the initialization point. While this approach is not optimal, it does well enough to ensure that that correlations between parameters of different magnitudes are well explored. Furthermore, we rescale our top of atmosphere temperature so that it is on the same scale as our abundances (on the order of 10), as opposed to its true range (on the order of 1000).




Figure D.1: The learned normalizing flow topology as compared with `PyMultiNest` equally weighted samples from a baseline retrieval. The normalizing flow is able to learn the topology of the posterior distribution. Note that this exercise does not constitute actual parameter inference—rather, we are demonstrating the flexibility of the global sampling scheme to approximate the final global posterior distribution.

Appendix E: Appendix from Chapter 7

E.1 Benchmarking nested sampling

As mentioned previously, in this work we couple CHIMERA to `nautilus` because it provides benefits over the more standard `PyMultiNest` sampler. The primary functional difference between `nautilus` and `PyMultiNest` is the way in which the two codes define iso-likelihood surfaces as they sweep through parameter space. `PyMultiNest` defines ellipsoidal bounds and inflates the bounds by some ϵ ; `nautilus` learns the iso-likelihood surface by training an ensemble of neural networks.

Due to the relative novelty of applying `nautilus` in the context of atmospheric retrievals, we seek to benchmark its performance in simulated test cases. We test the performance of the `nautilus` package and `MultiNest` along three axes: posterior accuracy, evidence accuracy, and runtime / efficiency.¹ To keep this problem tractable, we do not perform these tests on the full, nominal JWST retrievals presented in this manuscript. Rather, we test the performance of `nautilus` by coupling it to the `PLATON` forward model (Zhang et al. 2019, 2020, 2025). Our benchmark case is the transmission spectrum of HD 209458b as observed with Hubble/STIS, Hubble/WFC3, and Spitzer (Knutson et al. 2007a; Deming et al. 2013; Evans et al. 2015). We fit this with an 8-parameter model: stellar radius, planet mass, planet radius, planetary (isothermal) temperature,

¹All code for these tests was run on a 2021 MacBook Pro with a 16 GB Apple M1 Pro chip.




Figure E.1: The posterior distributions output by `nautilus` and `MultiNest` for our PLATON benchmarking exercise.

a log scattering slope factor, a log planet metallicity, a log (gray) cloud top pressure, and a scalar multiplier to the observed flux uncertainties.

Our baseline calculations include 500 live points for `MultiNest` and a sampling efficiency of 0.3, per the code’s documented suggestions for calculating Bayesian evidence. For `nautilus`, we use the developer’s recommended setting of 2000 live points. We additionally set the `n_networks` parameter to 16. We do so because the cost of training more neural networks to learn the iso-likelihood ellipsoids is negligible compared to the cost of our likelihood calculation (about 1 s per forward model computation).




Figure E.2: The result of benchmarking the `nautilus` code against `pymultinest`. The estimated natural log of the Bayesian evidence is shown as a function of normalized runtime ($\tau_{\text{norm}} = 2000\tau_{\text{runtime}}/n_{\text{livepoints}}$). The nested sampling (NS) evidence and nested importance sampling (NIS) evidence as calculated by `MultiNest` are both shown.

E.1.1 Posterior precision

Fig. E.1 shows the posterior distributions calculated by both `nautilus` and `MultiNest`. We find that the `MultiNest`-derived posterior distribution widths tend to be underestimated on the order of 10% compared to the `nautilus` runs. This underestimation is particularly evident by eye in the 2D credible regions, in which differences in the joint posterior distributions (including correlations) become apparent. There does not appear to be a correlation between the simplicity of a marginalized distribution (e.g., how similar it is to a 1-dimensional Gaussian) and how well `MultiNest` matches the `nautilus` output. This underestimation is consistent with the findings of, e.g., [Ih & Kempton \(2021\)](#) and [Lemos et al. \(2023\)](#).

E.1.2 Evidence accuracy

We find that the standard evidence as reported by `MultiNest` is inaccurate and imprecise (Fig. E.2), as also noted by [Dittmann \(2024\)](#), [Lemos et al. \(2023\)](#), and [Albert \(2020\)](#), likely due to

the way in which `MultiNest` draws ellipsoidal bounds. While `MultiNest`'s importance nested sampling evidence is indeed more accurate and more precise than the standard reported evidence (Feroz et al. 2013), we find that its uncertainty is underestimated by a factor of approximately 1.5 with respect to the `nautilus` output. This effect is either a systematic offset or a hint of greater variance between runs.

Notably, our benchmarking test may be in the regime in which `MultiNest` comparatively performs slightly *better* than the novel retrievals presented in this Letter. Albert (2020) found that `MultiNest`'s evidence bias rapidly grows with the number of dimensions past 10 dimensions. The test problem in this appendix is 8-dimensional; the full retrievals we present in this manuscript reach 11 dimensions. We expect that any evidence biases would therefore be elevated in the context of our scientific problem. Given that the bias is expected to be dimension- and problem-dependent (Albert 2020), this evidence bias will likely result in Bayes factor biases, as well.

E.1.3 Runtime / efficiency

Despite having four times as many live points as our `MultiNest` runs, `nautilus` only took roughly two times as long to converge on average (Fig. E.2). Scaled by the number of live points, then, `nautilus` is far more efficient — as found in general test cases by Lange (2023).

E.2 Full retrieval outputs

Our full retrieval results will be made available on Zenodo upon publication. Full corner plots for our cross correlation retrieval and traditional retrieval are shown in Figs. E.3 and E.4, respectively.




Figure E.3: Posterior distribution of our cross-correlation retrievals, with gases and cloud deck pressure fit in $\log_{10}$. We label the relative log abundance of the minor CO isotopologue as “ISO.”




Figure E.4: Posterior distribution of our traditional retrievals.

Appendix F: Appendix from Chapter 8

F.1 A derivation of a transmission spectrum broadening kernel

Various kernels that are used to broaden 1D transmission spectra under the assumption of solid-body rotation have been shown in the literature ([Brogi et al. 2016](#); [Gandhi et al. 2022](#); [Maguire et al. 2022](#); [Boucher et al. 2023](#); [Klein et al. 2024](#)). We review a simple derivation of this kernel below from a distribution perspective, which we apply to most of our forward models.

At the approaching limb equator, the velocity relative to the observer is

$$v_{z,\max} = \frac{2\pi R_P}{P} \quad (\text{F.1})$$

for a planetary radius R_P and a rotational period P . At the receding limb equator, the line-of-sight velocity is

$$v_{z,\min} = -\frac{2\pi R_P}{P} \quad (\text{F.2})$$

At a given latitude θ measured from the equator, the distance swept out in a single rotation

is $R \cos(\theta)$, so the latitude-dependent line-of-sight velocity is

$$v_z(\theta) = \frac{2\pi R_P \cos(\theta)}{P} \quad (\text{F.3})$$

The probability distribution of θ is uniform between $-\pi$ and π ; there is equal probability of selecting any latitude on the projected sphere in our simplified, homogeneous model. The cosine of a uniform random variable, however, is not uniform. Assume a transformation f to a variable y from a variable x . If x has a distribution X , then the distribution of y is¹

$$Y(y) = X(f^{-1}(y)) \left| \frac{d}{dy} f^{-1}(y) \right| \quad (\text{F.4})$$

Here, we have $v_z = f(\theta) = \frac{2\pi R_P \cos(\theta)}{P}$, the inverse of which is $f^{-1}(v_z) = \arccos(\frac{P}{2\pi R_P} v_z)$.

Hence, the distribution of v_z is

$$Y(v_z) = X\left(\arccos\left(\frac{P}{2\pi R_P} v_z\right)\right) \left| -\frac{P}{2\pi R_P} \frac{1}{\sqrt{1 - \frac{P^2 v_z^2}{4\pi^2 R_P^2}}} \right| \quad (\text{F.5})$$

Because θ is uniform over its bounds, this distribution reduces to

¹<https://www.math.arizona.edu/~jwatkins/f-transform.pdf>

$$Y(v_z) = \frac{1}{\sqrt{\frac{4\pi^2 R_P^2}{P^2} - v_z^2}} \quad (\text{F.6})$$

The probability of selecting a particular velocity between v_z and $v_z + dv_z$ can be written as

$$P(v_z \in (v_z, v_z + dv_z)) \propto \frac{1}{\sqrt{\frac{4\pi^2 R_P^2}{P^2} - v_z^2}} dv_z \quad (\text{F.7})$$

This solution is shown to match well to a computed transform of a uniform distribution in Fig. F.1.

Normalized appropriately, this kernel can be calculated and convolved with a 1D exoplanet transmission spectrum to approximate the effect of solid-body rotation on the transmission spectrum. Note, of course, that this kernel is simpler than what others have derived (e.g., [Klein et al. 2024](#)), and it assumes homogeneous gas abundances, homogeneous temperature, and no winds.

F.2 Benchmarking against CHIMERA

We perform rigorous tests to assess how our code performs against the 1D CHIMERA model. This step is important because the former is used to generate the synthetic data, and the latter is the retrieval’s forward model. We need to first show that the two agree and produce consistent results for 1D atmospheres.

The benchmarking experiments of [Barstow et al. \(2020\)](#), [\(Niraula et al. 2022\)](#), and [Villanueva et al. \(2024\)](#) provide a starting point for this work, and like them we find that fixing 1) the precision




Figure F.1: A cosine transform of a uniform distribution, with the analytic solution from Eq. F.7 overplotted.

of constants, 2) the opacity functions, and 3) the altitude grids used by both models brings the codes into the expected agreement. Beyond this, addressing subtle effects (such as truncating the atmosphere at different altitudes) account for the majority of the remainder of the error.

Over representative portions of the hot Jupiter model spectrum, we achieve less than 5 ppm median absolute difference between CHIMERA and the ray-striking RT code for an isothermal atmosphere (Fig. F.2). Our implementation of an “isothermal atmosphere” in the 3D code is wrapped around the entire planet—that is, the entire atmosphere at every latitude–longitude–altitude point is set to be the same temperature.




Figure F.2: Benchmarking CHIMERA and our 3D forward model.

F.3 Opacity sources

Our gas opacities are summarized in Table F.1. We include nearly all gas species in the near-infrared (H_2O , CO , CH_4 , H_2S , NH_3 and HCN) that were included in [Line et al. \(2021\)](#); this list corresponds to all species that would be expected to contribute significant opacity over the modeled wavelength range based on chemical equilibrium calculations. We exclude the minor CO isotopologue, as there is no clear equilibrium chemistry expectation for its abundance.

Our infrared opacities are calculated on a grid from 1–2.63 microns at a resolution of 500,000 (later downsampled to a resolution of 250,000). The temperature range spans 250 K to 3000 K in 39 steps, with the steps increasing from 5 K to 200 K over the range. Pressures range from 300 bar to 1 μbar in 18 steps, which are roughly evenly (log-)spaced.

We swap between the CHIMERA and Exo-Transmit opacity formats using the `cortecs` ([Savel et al. 2024a](#)) code.

Table F.1: Opacity Sources and Citations

Infrared (1.4-2.6 μm , R=125,000)	
Opacity Source	Citation
CH_4	Hargreaves et al. (2020)
CO	Li et al. (2015)
HCN	Barber et al. (2014)
H_2O	Polyansky et al. (2018)
H_2S	Azzam et al. (2016)
NH_3	Coles et al. (2019)

Bibliography

- Bayes, T. 1763, Philosophical transactions, 370
- Gibbs, J. W. 1878, American journal of science, 3, 441
- Christiansen, C. 1884, Annalen der Physik, 259, 298
- Bohr, N. 1923, Nature, 112, 29
- Shannon, C. E. 1948, The Bell system technical journal, 27, 379
- Schwarzschild, M. 1958, Structure and evolution of the stars.
- Safronov, V. 1960, Annales d'Astrophysique, Vol. 23, p. 979, 23, 979
- Toomre, A. 1964, Astrophysical Journal, vol. 139, p. 1217-1238 (1964)., 139, 1217
- Mason, E., & Marrero, T. 1970, in Advances in atomic and molecular physics, Vol. 6 (Elsevier), 155-232
- Lucy, L., & Sweeney, M. 1971, The Astronomical Journal, 76, 544
- Rhines, P. B. 1975, Journal of Fluid Mechanics, 69, 417
- Acuna, M. H., & Ness, N. F. 1976, Journal of Geophysical Research, 81, 2917
- Rodgers, C. D. 1976, Reviews of Geophysics, 14, 609
- Prinn, R. G., & Barshay, S. S. 1977, Science, 198, 1031
- Weidenschilling, S. 1977, Monthly Notices of the Royal Astronomical Society, 180, 57
- Norman, C., & Smith, R. 1978, Astronomy and Astrophysics, vol. 68, no. 1-2, Aug. 1978, p. 145-155. Research supported by the Miller Foundation and CECAM., 68, 145
- Mizuno, H. 1980, Progress of theoretical physics, 64, 544
- Watson, A. J., Donahue, T. M., & Walker, J. C. 1981, Icarus, 48, 150
- Draine, B. T., Roberge, W. G., & Dalgarno, A. 1983, Astrophysical Journal, Part 1, vol. 264, Jan. 15, 1983, p. 485-507., 264, 485
- Kasting, J. F., & Pollack, J. B. 1983, Icarus, 53, 479
- Yung, Y. L., Allen, M., & Pinto, J. P. 1984, Astrophysical Journal Supplement Series (ISSN 0067-0049), vol. 55, July 1984, p. 465-506., 55, 465
- Bodenheimer, P., & Pollack, J. B. 1986, Icarus, 67, 391

- Horne, K. 1986, PASP, 98, 609, doi: [10.1086/131801](https://doi.org/10.1086/131801)
- Zahnle, K. J., & Kasting, J. F. 1986, *Icarus*, 68, 462
- Hunten, D. M., Pepin, R. O., & Walker, J. C. 1987, *Icarus*, 69, 532
- Stamnes, K., Tsay, S.-C., Wiscombe, W., & Jayaweera, K. 1988, *Applied optics*, 27, 2502
- Van Dishoeck, E. F., & Black, J. H. 1988, *The Astrophysical Journal*, 334, 771
- Gautier, D., & Owen, T. 1989, in *Origin and evolution of planetary and satellite atmospheres*, 487–512
- Wetherill, G. W. 1990, IN: *Annual review of earth and planetary sciences*. Volume 18 (A91-15219 04-46). Palo Alto, CA, Annual Reviews, Inc., 1990, p. 205-256., 18, 205
- McKay, C. P., Pollack, J. B., & Courtin, R. 1991, *Science*, 253, 1118
- Gelman, A., & Rubin, D. B. 1992, *Statistical science*, 7, 457
- Geweke, J. 1992, *Bayesian statistics*, 4, 641
- Griesmann, U., Esser, B., & Baig, M. A. 1992, *Journal of Physics B: Atomic, Molecular and Optical Physics*, 25, 3475
- Lord, S. D. 1992, *A new software tool for computing Earth's atmospheric transmission of near-and far-infrared radiation*, Vol. 103957 (Ames Research Center)
- Wolszczan, A., & Frail, D. A. 1992, *Nature*, 355, 145
- Grenander, U., & Miller, M. I. 1994, *Journal of the Royal Statistical Society: Series B (Methodological)*, 56, 549
- Tierney, L. 1994, *the Annals of Statistics*, 1701
- Kurucz, R. 1995, *Atomic Line List*
- Kyilling, A., Stamnes, K., & Tsay, S.-C. 1995, *Journal of Atmospheric Chemistry*, 21, 115
- Mayor, M., & Queloz, D. 1995, *Nature*, 378, 355
- Baranne, A., Queloz, D., Mayor, M., et al. 1996, *Astronomy and Astrophysics Supplement Series*, 119, 373
- Chassefière, E. 1996, *Icarus*, 124, 537
- Cho, J. Y., & Polvani, L. M. 1996, *Science*, 273, 335
- Lin, D. N., Bodenheimer, P., & Richardson, D. C. 1996, *Nature*, 380, 606
- Pollack, J. B., Hubickyj, O., Bodenheimer, P., et al. 1996, *icarus*, 124, 62

- Rasio, F. A., & Ford, E. B. 1996, *Science*, 274, 954
- Burrows, A., Marley, M., Hubbard, W., et al. 1997, *The Astrophysical Journal*, 491, 856
- Butler, R. P., Marcy, G. W., Williams, E., Hauser, H., & Shirts, P. 1997, *The Astrophysical Journal Letters*, 474, L115
- Brooks, S. 1998, *Journal of the royal statistical society: series D (the Statistician)*, 47, 69
- Butler, R. P., Marcy, G. W., Vogt, S. S., & Apps, K. 1998, *Publications of the Astronomical Society of the Pacific*, 110, 1389
- Seager, S., & Sasselov, D. D. 1998, *The Astrophysical Journal Letters*, 502, L157
- Charbonneau, D., Noyes, R. W., Korzennik, S. G., et al. 1999, *The Astrophysical Journal Letters*, 522, L145
- Collier Cameron, A., Horne, K., Penny, A., & James, D. 1999, *Nature*, 402, 751
- Fischer, D. A., Marcy, G. W., Butler, R. P., Vogt, S. S., & Apps, K. 1999, *Publications of the Astronomical Society of the Pacific*, 111, 50
- Forget, F., Hourdin, F., Fournier, R., et al. 1999, *Journal of Geophysical Research: Planets*, 104, 24155
- Wardle, M., & Ng, C. 1999, *Monthly Notices of the Royal Astronomical Society*, 303, 239
- Bodenheimer, P., Hubickyj, O., & Lissauer, J. J. 2000, *Icarus*, 143, 2
- Charbonneau, D., Brown, T. M., Latham, D. W., & Mayor, M. 2000, *The Astrophysical Journal*, 529, L45
- Mazeh, T., Naef, D., Torres, G., et al. 2000, *The Astrophysical Journal Letters*, 532, L55
- Seager, S., & Sasselov, D. D. 2000, *The Astrophysical Journal*, 537, 916
- Boss, A. P. 2001, *The Astrophysical Journal*, 563, 367
- Brown, T. M. 2001, *The Astrophysical Journal*, 553, 1006
- Hubbard, W., Fortney, J., Lunine, J., et al. 2001, *The Astrophysical Journal*, 560, 413
- Pepe, F., Mayor, M., Galland, F., et al. 2002, *Astronomy & Astrophysics*, 388, 632
- Showman, A., & Guillot, T. 2002, *Astronomy & Astrophysics*, 385, 166
- Hubeny, I., Burrows, A., & Sudarsky, D. 2003, *The Astrophysical Journal*, 594, 1011
- Lodders, K. 2003, *The Astrophysical Journal*, 591, 1220
- Neal, R. M. 2003, *The annals of statistics*, 31, 705

- Vidal-Madjar, A., Des Etangs, A. L., Désert, J.-M., et al. 2003, *Nature*, 422, 143
- Adcroft, A., Campin, J.-M., Hill, C., & Marshall, J. 2004, *Monthly Weather Review*, 132, 2845
- Canup, R. M. 2004, *Icarus*, 168, 433
- Clayton, D. D., & Nittler, L. R. 2004, *Annu. Rev. Astron. Astrophys.*, 42, 39
- Gelman, A., Carlin, J. B., Stern, H. S., & Rubin, D. B. 2004, *CRC Texts in Statistical Science*, 136
- Sánchez-Lavega, A. 2004, *The Astrophysical Journal*, 609, L87
- Skilling, J. 2004, *Bayesian inference and maximum entropy methods in science and engineering*, 735, 395
- Yelle, R. V. 2004, *Icarus*, 170, 167
- Zakamska, N. L., & Tremaine, S. 2004, *The Astronomical Journal*, 128, 869
- Ida, S., & Lin, D. 2005, *The Astrophysical Journal*, 626, 1045
- Milam, S., Savage, C., Brewster, M., Ziurys, L., & Wyckoff, S. 2005, *The Astrophysical Journal*, 634, 1126
- Stevens, I. R. 2005, *Monthly Notices of the Royal Astronomical Society*, 356, 1053
- Zhu, X., Talaat, E., Baker, J., & Yee, J.-H. 2005, in *Annales Geophysicae*, Vol. 23, Copernicus Publications Göttingen, Germany, 3313–3322
- Cooper, C. S., & Showman, A. P. 2006, *The Astrophysical Journal*, 649, 1048
- Gardner, J. P., Mather, J. C., Clampin, M., et al. 2006, *Space Science Reviews*, 123, 485
- Lide, D. R. 2006, *CRC Handbook of Chemistry and Physics*, ; Lide, DR, Ed, CRC Press: Boca Raton, Fla
- Rafikov, R. R. 2006, *The Astrophysical Journal*, 648, 666
- Chabrier, G., & Baraffe, I. 2007, *The Astrophysical Journal Letters*, 661, L81
- Cutler, C., & Vallisneri, M. 2007, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 76, 104018
- Erkaev, N., Kulikov, Y. N., Lammer, H., et al. 2007, *Astronomy & Astrophysics*, 472, 329
- Gilmozzi, R., & Spyromilio, J. 2007, *The Messenger*, 127, 3
- Hunter, J. D. 2007, *Computing in science & engineering*, 9, 90
- Knutson, H. A., Charbonneau, D., Noyes, R. W., Brown, T. M., & Gilliland, R. L. 2007a, *The Astrophysical Journal*, 655, 564

- Knutson, H. A., Charbonneau, D., Allen, L. E., et al. 2007b, *Nature*, 447, 183
- Muñoz, A. G. 2007, *Planetary and Space Science*, 55, 1426
- Perez, F., & Granger, B. E. 2007, *Computing in Science and Engineering*, 9, 21, doi: [10.1109/MCSE.2007.53](https://doi.org/10.1109/MCSE.2007.53)
- Zarka, P. 2007, *Planetary and Space Science*, 55, 598
- Cho, J. Y., Menou, K., Hansen, B. M., & Seager, S. 2008, *The Astrophysical Journal*, 675, 817
- Fortney, J. J., Lodders, K., Marley, M. S., & Freedman, R. S. 2008, *The Astrophysical Journal*, 678, 1419
- Gilmozzi, R., & Spyromilio, J. 2008, in *Ground-based and Airborne Telescopes II*, Vol. 7012, SPIE, 494–503
- Klemp, J., Dudhia, J., & Hassiotis, A. 2008, *Monthly Weather Review*, 136, 3987
- Lian, Y., & Showman, A. P. 2008, *Icarus*, 194, 597
- Miller-Ricci, E., Seager, S., & Sasselov, D. 2008, *The Astrophysical Journal*, 690, 1056
- Ulrich Käufl, H. 2008, in *Precision Spectroscopy in Astrophysics: Proceedings of the ESO/Lisbon/Aveiro Conference held in Aveiro, Portugal, 11–15 September 2006*, Springer, 227–230
- Feroz, F., Hobson, M., & Bridges, M. 2009, *Monthly Notices of the Royal Astronomical Society*, 398, 1601
- Fray, N., & Schmitt, B. 2009, *Planetary and Space Science*, 57, 2053
- Hellier, C., Anderson, D., Cameron, A. C., et al. 2009, *Nature*, 460, 1098
- Madhusudhan, N., & Seager, S. 2009, *The Astrophysical Journal*, 707, 24
- Murray-Clay, R. A., Chiang, E. I., & Murray, N. 2009, *ApJ*, 693, 23, doi: [10.1088/0004-637X/693/1/23](https://doi.org/10.1088/0004-637X/693/1/23)
- Murray-Clay, R. A., Chiang, E. I., & Murray, N. 2009, *The Astrophysical Journal*, 693, 23
- Showman, A. P., Fortney, J. J., Lian, Y., et al. 2009, *The Astrophysical Journal*, 699, 564
- Spiegel, D. S., Silverio, K., & Burrows, A. 2009, *The Astrophysical Journal*, 699, 1487
- Tian, F. 2009, *The Astrophysical Journal*, 703, 905
- Zahnle, K., Marley, M., Freedman, R. S., Lodders, K., & Fortney, J. 2009, *The Astrophysical Journal Letters*, 701, L20
- Batygin, K., & Stevenson, D. J. 2010, *The Astrophysical Journal Letters*, 714, L238

- Christensen, U. R. 2010, *Space science reviews*, 152, 565
- Crossfield, I. J., Hansen, B. M., Harrington, J., et al. 2010, *The Astrophysical Journal*, 723, 1436
- Davies, J. H., & Davies, D. R. 2010, *Solid earth*, 1, 5
- Des Etangs, A. L., Ehrenreich, D., Vidal-Madjar, A., et al. 2010, *Astronomy & Astrophysics*, 514, A72
- Fortney, J., Shabram, M., Showman, A., et al. 2010, *The Astrophysical Journal*, 709, 1396
- Guillot, T. 2010, *Astronomy & Astrophysics*, 520, A27
- Koskinen, T. T., Cho, J. Y., Achilleos, N., & Aylward, A. D. 2010, *The Astrophysical Journal*, 722, 178
- Leconte, J., Chabrier, G., Baraffe, I., & Levrard, B. 2010, *Astronomy & Astrophysics*, 516, A64
- Line, M. R., Liang, M.-C., & Yung, Y. L. 2010, *The Astrophysical Journal*, 717, 496
- Linsky, J. L., Yang, H., France, K., et al. 2010, *The Astrophysical Journal*, 717, 1291
- Lovis, C., Fischer, D., et al. 2010, *Exoplanets*, 27
- McKinney, W. 2010, in *Proceedings of the 9th Python in Science Conference*, ed. Stéfan van der Walt & Jarrod Millman, 56 – 61, doi: [10.25080/Majora-92bf1922-00a](https://doi.org/10.25080/Majora-92bf1922-00a)
- Pepe, F. A., Cristiani, S., Lopez, R. R., et al. 2010, in *Ground-based and Airborne Instrumentation for Astronomy III*, Vol. 7735, International Society for Optics and Photonics, 77350F
- Perna, R., Menou, K., & Rauscher, E. 2010, *The Astrophysical Journal*, 719, 1421
- Snellen, I. A., De Kok, R. J., De Mooij, E. J., & Albrecht, S. 2010, *Nature*, 465, 1049
- Yuk, I.-S., Jaffe, D. T., Barnes, S., et al. 2010, in *Ground-based and Airborne Instrumentation for Astronomy III*, Vol. 7735, SPIE, 653–662
- Yuk, I.-S., Jaffe, D. T., Barnes, S., et al. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7735, *Ground-based and Airborne Instrumentation for Astronomy III*, ed. I. S. McLean, S. K. Ramsay, & H. Takami, 77351M, doi: [10.1117/12.856864](https://doi.org/10.1117/12.856864)
- Faedi, F., Barros, S. C., Anderson, D. R., et al. 2011, *Astronomy & Astrophysics*, 531, A40
- Kurucz, R. L. 2011, *Canadian Journal of Physics*, 89, 417
- Léger, A., Grasset, O., Fegley, B., et al. 2011, *Icarus*, 213, 1
- Montalto, M., Santos, N., Boisse, I., et al. 2011, *Astronomy & Astrophysics*, 528, L17
- Moses, J. I., Visscher, C., Fortney, J. J., et al. 2011, *The Astrophysical Journal*, 737, 15

- Öberg, K. I., Murray-Clay, R., & Bergin, E. A. 2011, *The Astrophysical Journal Letters*, 743, L16
- Showman, A. P., & Polvani, L. M. 2011, *The Astrophysical Journal*, 738, 71
- Thies, I., Kroupa, P., Goodwin, S. P., Stamatellos, D., & Whitworth, A. P. 2011, *Monthly Notices of the Royal Astronomical Society*, 417, 1817
- Vidotto, A., Jardine, M., & Helling, C. 2011, *Monthly Notices of the Royal Astronomical Society: Letters*, 411, L46
- Walsh, K. J., Morbidelli, A., Raymond, S. N., O'Brien, D. P., & Mandell, A. M. 2011, *Nature*, 475, 206
- Wu, Y., & Lithwick, Y. 2011, *The Astrophysical Journal*, 735, 109
- Benneke, B., & Seager, S. 2012, *The Astrophysical Journal*, 753, 100
- Brogi, M., Snellen, I. A., De Kok, R. J., et al. 2012, *Nature*, 486, 502
- Hu, R., Seager, S., & Bains, W. 2012, *The Astrophysical Journal*, 761, 166
- Johns, M., McCarthy, P., Raybould, K., et al. 2012, *Ground-based and Airborne Telescopes IV*, 8444, 526
- Kempton, E. M.-R., & Rauscher, E. 2012, *The Astrophysical Journal*, 751, 117
- Line, M. R., Zhang, X., Vasisht, G., et al. 2012, *The Astrophysical Journal*, 749, 93
- Madhusudhan, N. 2012, *The Astrophysical Journal*, 758, 36
- Maxted, P., Anderson, D., Cameron, A. C., et al. 2012, *Publications of the Astronomical Society of the Pacific*, 125, 48
- Menou, K. 2012, *The Astrophysical Journal*, 745, 138
- Miller-Ricci Kempton, E., Zahnle, K., & Fortney, J. J. 2012, *The Astrophysical Journal*, 745, 3
- Moses, J., Madhusudhan, N., Visscher, C., & Freedman, R. 2012, *The Astrophysical Journal*, 763, 25
- Perna, R., Heng, K., & Pont, F. 2012, *The Astrophysical Journal*, 751, 59
- Rauscher, E., & Menou, K. 2012, *The Astrophysical Journal*, 750, 96
- Visscher, C. 2012, *The Astrophysical Journal*, 757, 5
- Astropy Collaboration, Robitaille, T. P., Tollerud, E. J., et al. 2013, *A&A*, 558, A33, doi: [10.1051/0004-6361/201322068](https://doi.org/10.1051/0004-6361/201322068)
- Batygin, K., Stanley, S., & Stevenson, D. J. 2013, *The Astrophysical Journal*, 776, 53
- Benneke, B., & Seager, S. 2013, *ApJ*, 778, 153, doi: [10.1088/0004-637X/778/2/153](https://doi.org/10.1088/0004-637X/778/2/153)

- Birkby, J., de Kok, R., Brogi, M., et al. 2013, *Monthly Notices of the Royal Astronomical Society: Letters*, 436, L35
- Bourrier, V., & Des Etangs, A. L. 2013, *Astronomy & Astrophysics*, 557, A124
- Brogi, M., Snellen, I., De Kok, R., et al. 2013, *The Astrophysical Journal*, 767, 27
- de Kok, R. J., Brogi, M., Snellen, I. A., et al. 2013, *Astronomy & Astrophysics*, 554, A82
- Deming, D., Wilkins, A., McCullough, P., et al. 2013, *The Astrophysical Journal*, 774, 95
- Demory, B.-O., De Wit, J., Lewis, N., et al. 2013, *The Astrophysical Journal Letters*, 776, L25
- Dobbs-Dixon, I., & Agol, E. 2013, *Monthly Notices of the Royal Astronomical Society*, 435, 3159
- Feroz, F., Hobson, M. P., Cameron, E., & Pettitt, A. N. 2013, arXiv preprint arXiv:1306.2144
- Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, *Publications of the Astronomical Society of the Pacific*, 125, 306
- Fossati, L., Ayres, T., Haswell, C., et al. 2013, *The Astrophysical Journal Letters*, 766, L20
- Husser, T.-O., Wende-von Berg, S., Dreizler, S., et al. 2013, *Astronomy & Astrophysics*, 553, A6
- Kataria, T., Showman, A. P., Lewis, N. K., et al. 2013, *The Astrophysical Journal*, 767, 76
- Line, M. R., & Yung, Y. L. 2013, *The Astrophysical Journal*, 779, 3
- Line, M. R., Wolf, A. S., Zhang, X., et al. 2013, *The Astrophysical Journal*, 775, 137
- Morley, C. V., Fortney, J. J., Kempton, E. M.-R., et al. 2013, *The Astrophysical Journal*, 775, 33
- Pepe, F., Cristiani, S., Rebolo, R., et al. 2013, *The Messenger*, 153, 6
- Perez-Becker, D., & Showman, A. P. 2013, *The Astrophysical Journal*, 776, 134
- Pont, F., Sing, D. K., Gibson, N., et al. 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 2917
- Rauscher, E., & Menou, K. 2013, *The Astrophysical Journal*, 764, 103
- Rodler, F., Kürster, M., & Barnes, J. R. 2013, *Monthly Notices of the Royal Astronomical Society*, 432, 1980
- Rugheimer, S., Kaltenegger, L., Zsom, A., Segura, A., & Sasselov, D. 2013, *Astrobiology*, 13, 251
- Showman, A. P., Fortney, J. J., Lewis, N. K., & Shabram, M. 2013a, *The Astrophysical Journal*, 762, 24
- . 2013b, *The Astrophysical Journal*, 762, 24

- Vallisneri, M., & Yunes, N. 2013, *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 87, 102002
- Barber, R., Strange, J., Hill, C., et al. 2014, *Monthly Notices of the Royal Astronomical Society*, 437, 1828
- Batygín, K., & Stanley, S. 2014, *The Astrophysical Journal*, 794, 10
- Brogi, M., De Kok, R., Birkby, J., Schwarz, H., & Snellen, I. 2014, *Astronomy & Astrophysics*, 565, A124
- Buchner, J., Georgakakis, A., Nandra, K., et al. 2014, *Astronomy & Astrophysics*, 564, A125
- Cunha, D., Santos, N. C., Figueira, P., et al. 2014, *A&A*, 568, A35, doi: [10.1051/0004-6361/201423723](https://doi.org/10.1051/0004-6361/201423723)
- Follert, R., Dorn, R., Oliva, E., et al. 2014, in *Ground-based and Airborne Instrumentation for Astronomy V*, Vol. 9147, SPIE, 476–485
- Kempton, E. M.-R., Perna, R., & Heng, K. 2014, *The Astrophysical Journal*, 795, 24
- Kislyakova, K. G., Holmström, M., Lammer, H., Odert, P., & Khodachenko, M. L. 2014, *Science*, 346, 981
- Koskinen, T., Yelle, R., Lavvas, P., & Cho, J. Y. 2014, *The Astrophysical Journal*, 796, 16
- Kreidberg, L., Bean, J. L., Désert, J.-M., et al. 2014, *The Astrophysical Journal Letters*, 793, L27
- Lavvas, P., Koskinen, T., & Yelle, R. V. 2014, *The Astrophysical Journal*, 796, 15
- Line, M. R., Knutson, H., Wolf, A. S., & Yung, Y. L. 2014, *The Astrophysical Journal*, 783, 70
- Lockwood, A. C., Johnson, J. A., Bender, C. F., et al. 2014, *The Astrophysical Journal Letters*, 783, L29
- Mayne, N., Baraffe, I., Acreman, D. M., et al. 2014, *Geoscientific Model Development*, 7, 3059
- O’Brien, D. P., Walsh, K. J., Morbidelli, A., Raymond, S. N., & Mandell, A. M. 2014, *Icarus*, 239, 74
- Park, C., Jaffe, D. T., Yuk, I.-S., et al. 2014, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 9147, *Ground-based and Airborne Instrumentation for Astronomy V*, ed. S. K. Ramsay, I. S. McLean, & H. Takami, 91471D, doi: [10.1117/12.2056431](https://doi.org/10.1117/12.2056431)
- Park, C., Jaffe, D. T., Yuk, I.-S., et al. 2014, in *Ground-based and Airborne Instrumentation for Astronomy V*, Vol. 9147, SPIE, 510–521
- Quirrenbach, A., Amado, P., Caballero, J., et al. 2014, in *Ground-based and airborne instrumentation for astronomy V*, Vol. 9147, SPIE, 531–542
- Rauscher, E., & Kempton, E. M. 2014, *The Astrophysical Journal*, 790, 79

- Rogers, T., & Showman, A. 2014, *The Astrophysical Journal Letters*, 782, L4
- Rogers, T. M., & Komacek, T. D. 2014, *The Astrophysical Journal*, 794, 132
- Tsai, S.-M., Dobbs-Dixon, I., & Gu, P.-G. 2014, *The Astrophysical Journal*, 793, 141
- Vitale, S., & Del Pozzo, W. 2014, *Physical Review D*, 89, 022002
- Bourrier, V., Ehrenreich, D., & des Etangs, A. L. 2015, *Astronomy & Astrophysics*, 582, A65
- Driscoll, P. E., & Barnes, R. 2015, *Astrobiology*, 15, 739
- Ehrenreich, D., Bourrier, V., Wheatley, P. J., et al. 2015, *Nature*, 522, 459
- Evans, T. M., Aigrain, S., Gibson, N., et al. 2015, *Monthly Notices of the Royal Astronomical Society*, 451, 680
- Grimm, S. L., & Heng, K. 2015, *ApJ*, 808, 182, doi: [10.1088/0004-637X/808/2/182](https://doi.org/10.1088/0004-637X/808/2/182)
- Haynes, K., Mandell, A. M., Madhusudhan, N., Deming, D., & Knutson, H. 2015, *The Astrophysical Journal*, 806, 146
- Hoeijmakers, H., De Kok, R., Snellen, I., et al. 2015, *Astronomy & Astrophysics*, 575, A20
- Hospital, A., Goñi, J. R., Orozco, M., & Gelpí, J. L. 2015, *Advances and Applications in Bioinformatics and Chemistry*, 37
- Kataria, T., Showman, A. P., Fortney, J. J., et al. 2015, *The Astrophysical Journal*, 801, 86
- Kausch, W., Noll, S., Smette, A., et al. 2015, *Astronomy & Astrophysics*, 576, A78
- Kreidberg, L. 2015, *Publications of the Astronomical Society of the Pacific*, 127, 1161
- Lam, S. K., Pitrou, A., & Seibert, S. 2015, in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, 1–6
- Le, H. A. N., Pak, S., Jaffe, D. T., et al. 2015, *Advances in Space Research*, 55, 2509
- Li, G., Gordon, I. E., Rothman, L. S., et al. 2015, *The Astrophysical Journal Supplement Series*, 216, 15
- Line, M. R., Teske, J., Burningham, B., Fortney, J. J., & Marley, M. S. 2015, *The Astrophysical Journal*, 807, 183
- Louden, T., & Wheatley, P. J. 2015, *The Astrophysical Journal Letters*, 814, L24
- Mollière, P., van Boekel, R., Dullemond, C., Henning, T., & Mordasini, C. 2015, *The Astrophysical Journal*, 813, 47
- Ribas, Á., Bouy, H., & Merin, B. 2015, *Astronomy & Astrophysics*, 576, A52
- Showman, A. P., Lewis, N. K., & Fortney, J. J. 2015, *The Astrophysical Journal*, 801, 95

- Smette, A., Sana, H., Noll, S., et al. 2015, *Astronomy & Astrophysics*, 576, A77
- Tripathi, A., Kratter, K. M., Murray-Clay, R. A., & Krumholz, M. R. 2015, *The Astrophysical Journal*, 808, 173
- Wakeford, H. R., & Sing, D. K. 2015, *Astronomy & Astrophysics*, 573, A122
- Zarka, P., Lazio, J., & Hallinan, G. 2015, *Proceedings of Advancing Astrophysics with the Square Kilometre Array (AASKA14)*. 9-13 June, 2014
- Abbott, B., Jawahar, S., Lockerbie, N., & Tokmakov, K. 2016, *PHYSICAL REVIEW D Phys Rev D*, 93, 122003
- Azzam, A. A., Tennyson, J., Yurchenko, S. N., & Naumenko, O. V. 2016, *Monthly Notices of the Royal Astronomical Society*, 460, 4063
- Batygin, K., Bodenheimer, P. H., & Laughlin, G. P. 2016, *The Astrophysical Journal*, 829, 114
- Boley, A. C., Contreras, A. G., & Gladman, B. 2016, *The Astrophysical Journal Letters*, 817, L17
- Bourrier, V., Des Etangs, A. L., Ehrenreich, D., Tanaka, Y., & Vidotto, A. A. 2016, *Astronomy & Astrophysics*, 591, A121
- Bowler, B. P. 2016, *Publications of the Astronomical Society of the Pacific*, 128, 102001
- Brogi, M., De Kok, R., Albrecht, S., et al. 2016, *The Astrophysical Journal*, 817, 106
- Brooke, J. S., Bernath, P. F., Western, C. M., et al. 2016, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 168, 142
- Buchner, J. 2016, *Statistics and Computing*, 26, 383, doi: [10.1007/s11222-014-9512-y](https://doi.org/10.1007/s11222-014-9512-y)
- Delrez, L., Santerne, A., Almenara, J. M., et al. 2016, *MNRAS*, 458, 4025, doi: [10.1093/mnras/stw522](https://doi.org/10.1093/mnras/stw522)
- Evans, T. M., Sing, D. K., Wakeford, H. R., et al. 2016, *The Astrophysical Journal Letters*, 822, L4
- Feng, Y. K., Line, M. R., Fortney, J. J., et al. 2016, *The Astrophysical Journal*, 829, 52
- Foreman-Mackey, D., et al. 2016, *J. Open Source Softw.*, 1, 24
- France, K., Loyd, R. P., Youngblood, A., et al. 2016, *The Astrophysical Journal*, 820, 89
- Greene, T. P., Line, M. R., Montero, C., et al. 2016, *The Astrophysical Journal*, 817, 17
- Jolliffe, I. T., & Cadima, J. 2016, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374, 20150202
- Kao, M. M., Hallinan, G., Pineda, J. S., et al. 2016, *The Astrophysical Journal*, 818, 24

- Kataria, T., Sing, D. K., Lewis, N. K., et al. 2016, *ApJ*, 821, 9, doi: [10.3847/0004-637X/821/1/9](https://doi.org/10.3847/0004-637X/821/1/9)
- Kataria, T., Sing, D. K., Lewis, N. K., et al. 2016, *The Astrophysical Journal*, 821, 9
- Komacek, T. D., & Showman, A. P. 2016, *The Astrophysical Journal*, 821, 16
- Lee, G., Dobbs-Dixon, I., Helling, C., Bognar, K., & Woitke, P. 2016, *Astronomy & Astrophysics*, 594, A48
- Line, M. R., & Parmentier, V. 2016, *The Astrophysical Journal*, 820, 78
- Line, M. R., Stevenson, K. B., Bean, J., et al. 2016, *The Astronomical Journal*, 152, 203
- Lundkvist, M., Kjeldsen, H., Albrecht, S., et al. 2016, *Nature Communications*, 7, 11201
- Mazeh, T., Holczer, T., & Faigler, S. 2016, *Astronomy & Astrophysics*, 589, A75
- Mbarek, R., & Kempton, E. M.-R. 2016, *The Astrophysical Journal*, 827, 121
- McKemmish, L. K., Yurchenko, S. N., & Tennyson, J. 2016, *Monthly Notices of the Royal Astronomical Society*, 463, 771
- Mendonça, J. M., Grimm, S. L., Grosheintz, L., & Heng, K. 2016, *The Astrophysical Journal*, 829, 115
- Morbidelli, A., & Raymond, S. N. 2016, *Journal of Geophysical Research: Planets*, 121, 1962
- Mordasini, C., van Boekel, R., Mollière, P., Henning, T., & Benneke, B. 2016, *The Astrophysical Journal*, 832, 41
- Piso, A.-M. A., Pegues, J., & Öberg, K. I. 2016, *The Astrophysical Journal*, 833, 203
- Rocchetto, M., Waldmann, I., Venot, O., Lagage, P.-O., & Tinetti, G. 2016, *The Astrophysical Journal*, 833, 120
- Schaefer, L., Wordsworth, R. D., Berta-Thompson, Z., & Sasselov, D. 2016, *The Astrophysical Journal*, 829, 63
- Seifahrt, A., Bean, J. L., Stürmer, J., et al. 2016, in *Ground-based and Airborne Instrumentation for Astronomy VI*, Vol. 9908, SPIE, 332–345
- Tennyson, J., Yurchenko, S. N., Al-Refaie, A. F., et al. 2016, *Journal of Molecular Spectroscopy*, 327, 73
- Theano Development Team. 2016, arXiv e-prints, abs/1605.02688. <http://arxiv.org/abs/1605.02688>
- Underwood, D. S., Tennyson, J., Yurchenko, S. N., et al. 2016, *Monthly Notices of the Royal Astronomical Society*, 459, 3890

West, R., Hellier, C., Almenara, J.-M., et al. 2016, *Astronomy & Astrophysics*, 585, A126

Wills, R. C., & Schneider, T. 2016, *Journal of Climate*, 29, 3161

Youngblood, A., France, K., Loyd, R. P., et al. 2016, *The Astrophysical Journal*, 824, 101

Allart, R., Lovis, C., Pino, L., et al. 2017, *Astronomy & Astrophysics*, 606, A144

Batalha, N. E., & Line, M. R. 2017, *The Astronomical Journal*, 153, 151

Betancourt, M. 2017, arXiv preprint arXiv:1701.02434

Blecic, J., Dobbs-Dixon, I., & Greene, T. 2017, *The Astrophysical Journal*, 848, 127

Carnall, A. 2017, arXiv preprint arXiv:1705.05165

Evans, T. M., Sing, D. K., Kataria, T., et al. 2017, *Nature*, 548, 58

Fu, G., Deming, D., Knutson, H., et al. 2017, *The Astrophysical Journal Letters*, 847, L22

Fulton, B. J., Petigura, E. A., Howard, A. W., et al. 2017, *The Astronomical Journal*, 154, 109

Gaudi, B. S., Stassun, K. G., Collins, K. A., et al. 2017, *Nature*, 546, 514

Heays, A., Bosman, AD, v., & Van Dishoeck, E. 2017, *Astronomy & Astrophysics*, 602, A105

Komacek, T. D., Showman, A. P., & Tan, X. 2017, *The Astrophysical Journal*, 835, 198

Lee, J.-J., Gullikson, K., & Kaplan, K. 2017, *Igrins/Plp 2.2.0*, Zenodo, Zenodo, doi: [10.5281/zenodo.845059](https://doi.org/10.5281/zenodo.845059)

Line, M. R., Marley, M. S., Liu, M. C., et al. 2017, *The Astrophysical Journal*, 848, 83

MacDonald, R. J., & Madhusudhan, N. 2017, *Monthly Notices of the Royal Astronomical Society*, 469, 1979

Malik, M., Grosheintz, L., Mendonça, J. M., et al. 2017, *The astronomical journal*, 153, 56

Rogers, T. 2017, *Nature Astronomy*, 1, 1

Rogers, T., & McElwaine, J. 2017, *The Astrophysical Journal Letters*, 841, L26

Sheppard, K. B., Mandell, A. M., Tamburo, P., et al. 2017, *The Astrophysical Journal Letters*, 850, L32

Shulyak, D., Reiners, A., Engeln, A., et al. 2017, *Nature Astronomy*, 1, 0184

Tsai, S.-M., Lyons, J. R., Grosheintz, L., et al. 2017, *The Astrophysical Journal Supplement Series*, 228, 20

Yadav, R. K., & Thorngren, D. P. 2017, *The Astrophysical Journal Letters*, 849, L12

Youngblood, A., France, K., Loyd, R. P., et al. 2017, *The Astrophysical Journal*, 843, 31

- Zhang, J., Kempton, E. M.-R., & Rauscher, E. 2017, *The Astrophysical Journal*, 851, 84
- Zhou, Y., Apai, D., Lew, B. W., & Schneider, G. 2017, *The Astronomical Journal*, 153, 243
- Arcangeli, J., Désert, J.-M., Line, M. R., et al. 2018, *The Astrophysical Journal Letters*, 855, L30
- Astropy Collaboration, Price-Whelan, A. M., Sipőcz, B. M., et al. 2018, *AJ*, 156, 123, doi: [10.3847/1538-3881/aabc4f](https://doi.org/10.3847/1538-3881/aabc4f)
- Bell, T. J., & Cowan, N. B. 2018, *The Astrophysical Journal Letters*, 857, L20
- Birkby, J. L. 2018, arXiv preprint arXiv:1806.04617
- Bradbury, J., Frostig, R., Hawkins, P., et al. 2018, JAX: composable transformations of Python+NumPy programs, 0.3.13. [http://github.com/google/jax](https://github.com/google/jax)
- Brasser, R., Dauphas, N., & Mojzsis, S. 2018, *Geophysical Research Letters*, 45, 5908
- Catling, D. C., Krissansen-Totton, J., Kiang, N. Y., et al. 2018, *Astrobiology*, 18, 709
- Dawson, R. I., & Johnson, J. A. 2018, *Annual Review of Astronomy and Astrophysics*, 56, 175
- Fisher, C., & Heng, K. 2018, *MNRAS*, 481, 4698, doi: [10.1093/mnras/sty2550](https://doi.org/10.1093/mnras/sty2550)
- Fossati, L., Koskinen, T., Lothringer, J., et al. 2018, *The Astrophysical Journal Letters*, 868, L30
- Hawker, G. A., Madhusudhan, N., Cabot, S. H., & Gandhi, S. 2018, *The Astrophysical Journal Letters*, 863, L11
- Kempton, E. M.-R., Bean, J. L., Louie, D. R., et al. 2018, *Publications of the Astronomical Society of the Pacific*, 130, 114401
- Komacek, T. D., & Tan, X. 2018, *Research Notes of the AAS*, 2, 36
- Lines, S., Mayne, N., Boutle, I. A., et al. 2018, *Astronomy & Astrophysics*, 615, A97
- Lothringer, J. D., Barman, T., & Koskinen, T. 2018, *The Astrophysical Journal*, 866, 27
- Mace, G., Sokal, K., Lee, J.-J., et al. 2018, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 10702, *Ground-based and Airborne Instrumentation for Astronomy VII*, ed. C. J. Evans, L. Simard, & H. Takami, 107020Q, doi: [10.1117/12.2312345](https://doi.org/10.1117/12.2312345)
- Madhusudhan, N. 2018, arXiv preprint arXiv:1808.04824
- Mancini, L., Esposito, M., Covino, E., et al. 2018, *Astronomy & Astrophysics*, 613, A41
- Mangoubi, O., Pillai, N. S., & Smith, A. 2018, arXiv preprint arXiv:1808.03230
- Meadows, V. S., & Barnes, R. K. 2018, *Handbook of exoplanets*, 57
- Mendonça, J. M., Tsai, S.-m., Malik, M., Grimm, S. L., & Heng, K. 2018, *The Astrophysical Journal*, 869, 107

- Oklopčić, A., & Hirata, C. M. 2018, *The Astrophysical Journal Letters*, 855, L11
- Parmentier, V., & Crossfield, I. J. 2018, *Handbook of exoplanets*, 116
- Parmentier, V., Line, M. R., Bean, J. L., et al. 2018, *Astronomy & Astrophysics*, 617, A110
- Pino, L., Ehrenreich, D., Allart, R., et al. 2018, *Astronomy & Astrophysics*, 619, A3
- Polyansky, O. L., Kyuberis, A. A., Zobov, N. F., et al. 2018, *Monthly Notices of the Royal Astronomical Society*, 480, 2597
- Schwieterman, E. W., Kiang, N. Y., Parenteau, M. N., et al. 2018, *Astrobiology*, 18, 663
- Spake, J. J., Sing, D. K., Evans, T. M., et al. 2018, *Nature*, 557, 68
- Stock, J. W., Kitzmann, D., Patzer, A. B. C., & Sedlmayr, E. 2018, *Monthly Notices of the Royal Astronomical Society*, 479, 865
- Tóbiás, R., Furtenbacher, T., Császár, A. G., et al. 2018, *Journal of Quantitative Spectroscopy and Radiative Transfer*, 208, 152
- Tsiaras, A., Waldmann, I., Zingales, T., et al. 2018, *The Astronomical Journal*, 155, 156
- Zaghoo, M., & Collins, G. 2018, *The Astrophysical Journal*, 862, 19
- Zhang, X., & Showman, A. P. 2018, *The Astrophysical Journal*, 866, 2
- Akinsanmi, B., Barros, S., Santos, N., et al. 2019, *Astronomy & Astrophysics*, 621, A117
- Brogi, M., & Line, M. R. 2019, *The Astronomical Journal*, 157, 114
- Cabot, S. H., Madhusudhan, N., Hawker, G. A., & Gandhi, S. 2019, *Monthly Notices of the Royal Astronomical Society*, 482, 4422
- Caldas, A., Leconte, J., Selsis, F., et al. 2019, *Astronomy & Astrophysics*, 623, A161
- Cauley, P. W., Shkolnik, E. L., Llama, J., & Lanza, A. F. 2019, *Nature Astronomy*, 3, 1128
- Chiavassa, A., & Brogi, M. 2019, *Astronomy & Astrophysics*, 631, A100
- Coles, P. A., Yurchenko, S. N., & Tennyson, J. 2019, *Monthly Notices of the Royal Astronomical Society*, 490, 4638
- da Costa-Luis, C. 2019, *Journal of Open Source Software*, 4, 1277
- Debrecht, A., Carroll-Nellenback, J., Frank, A., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 483, 1481
- Fleury, B., Gudipati, M. S., Henderson, B. L., & Swain, M. 2019, *The Astrophysical Journal*, 871, 158

- Flowers, E., Brogi, M., Rauscher, E., Kempton, E. M.-R., & Chiavassa, A. 2019, *The Astronomical Journal*, 157, 209
- Gandhi, S., Madhusudhan, N., Hawker, G., & Piette, A. 2019, *The Astronomical Journal*, 158, 228
- Hindle, A., Bushby, P., & Rogers, T. 2019, *The Astrophysical Journal Letters*, 872, L27
- Hoeijmakers, H. J., Ehrenreich, D., Kitzmann, D., et al. 2019, *Astronomy & Astrophysics*, 627, A165
- Kramida, A., Olsen, K., & Ralchenko, Y. 2019, National Institute of Standards and Technology, US Department of Commerce
- Leet, C., Fischer, D. A., & Valenti, J. A. 2019, *The Astronomical Journal*, 157, 187
- Lines, S., Mayne, N., Manners, J., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 1332
- Lopez-Morales, M., Ben-Ami, S., Gonzalez-Abad, G., et al. 2019, *The Astronomical Journal*, 158, 24
- Lothringer, J. D., & Barman, T. 2019, *The Astrophysical Journal*, 876, 69
- Madhusudhan, N. 2019, arXiv preprint arXiv:1904.03190
- Malik, M., Kitzmann, D., Mendonça, J. M., et al. 2019, *The Astronomical Journal*, 157, 170
- Margossian, C. C. 2019, *Wiley interdisciplinary reviews: data mining and knowledge discovery*, 9, e1305
- McKemmish, L. K., Masseron, T., Hoeijmakers, H. J., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 2836
- Mikal-Evans, T., Sing, D. K., Goyal, J. M., et al. 2019, *Monthly Notices of the Royal Astronomical Society*, 488, 2222
- Molaverdikhani, K., Henning, T., & Mollière, P. 2019, *The Astrophysical Journal*, 883, 194
- Mollière, P., Wardenier, J., Van Boekel, R., et al. 2019, *Astronomy & Astrophysics*, 627, A67
- Owen, J. E. 2019, *Annual Review of Earth and Planetary Sciences*, 47, 67
- Roman, M., & Rauscher, E. 2019, *The Astrophysical Journal*, 872, 1
- Sainsbury-Martinez, F., Wang, P., Fromang, S., et al. 2019, *Astronomy & Astrophysics*, 632, A114
- Steinrueck, M. E., Parmentier, V., Showman, A. P., Lothringer, J. D., & Lupu, R. E. 2019, *The Astrophysical Journal*, 880, 14
- Tan, X., & Komacek, T. D. 2019, *The Astrophysical Journal*, 886, 26
- Thorngren, D., Gao, P., & Fortney, J. 2019a, *The Astrophysical Journal Letters*, 884, L6

- Thorngren, D., Gao, P., & Fortney, J. J. 2019b, *The Astrophysical Journal Letters*, 884, L6
- Welbanks, L., & Madhusudhan, N. 2019, *The Astronomical Journal*, 157, 206
- Yan, F., Casasayas-Barris, N., Molaverdikhani, K., et al. 2019, *Astronomy & Astrophysics*, 632, A69
- Zhang, M., Chachan, Y., Kempton, E. M.-R., & Knutson, H. A. 2019, *Publications of the Astronomical Society of the Pacific*, 131, 034501
- Albert, J. G. 2020, arXiv preprint arXiv:2012.15286
- Barstow, J. K., Changeat, Q., Garland, R., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 4884
- Barstow, J. K., & Heng, K. 2020, *Space science reviews*, 216, 1
- Beltz, H., Rauscher, E., Brogi, M., & Kempton, E. M.-R. 2020, *The Astronomical Journal*, 161, 1
- Bourrier, V., Ehrenreich, D., Lendl, M., et al. 2020a, *Astronomy & Astrophysics*, 635, A205
- Bourrier, V., Wheatley, P., Lecavelier des Etangs, A., et al. 2020b, *Monthly Notices of the Royal Astronomical Society*, 493, 559
- Buzard, C., Finnerty, L., Piskorz, D., et al. 2020, *The Astronomical Journal*, 160, 1
- Drummond, B., Hébrard, E., Mayne, N. J., et al. 2020, *Astronomy & Astrophysics*, 636, A68
- Edwards, B., Changeat, Q., Baeyens, R., et al. 2020, *AJ*, 160, 8, doi: [10.3847/1538-3881/ab9225](https://doi.org/10.3847/1538-3881/ab9225)
- Ehrenreich, D., Lovis, C., Allart, R., et al. 2020, *Nature*, 580, 597
- Feng, Y. K., Line, M. R., & Fortney, J. J. 2020, *The Astronomical Journal*, 160, 137
- Gandhi, S., Brogi, M., & Webb, R. K. 2020, *Monthly Notices of the Royal Astronomical Society*, 498, 194
- Gao, P., Thorngren, D. P., Lee, E. K. H., et al. 2020, *Nature Astronomy*, 4, 951, doi: [10.1038/s41550-020-1114-3](https://doi.org/10.1038/s41550-020-1114-3)
- Gibson, N. P., Merritt, S., Nugroho, S. K., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 2215
- Gronoff, G., Arras, P., Baraka, S., et al. 2020, *Journal of Geophysical Research: Space Physics*, 125, e2019JA027639
- Hargreaves, R. J., Gordon, I. E., Rey, M., et al. 2020, *The Astrophysical Journal Supplement Series*, 247, 55
- Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Nature*, 585, 357–362, doi: [10.1038/s41586-020-2649-2](https://doi.org/10.1038/s41586-020-2649-2)

- Hoeijmakers, H. J., Seidel, J. V., Pino, L., et al. 2020, *Astronomy & Astrophysics*, 641, A123
- Hood, C. E., Fortney, J. J., Line, M. R., et al. 2020, *The Astronomical Journal*, 160, 198
- Kobyzev, I., Prince, S. J., & Brubaker, M. A. 2020, *IEEE transactions on pattern analysis and machine intelligence*, 43, 3964
- Lacy, B. I., & Burrows, A. 2020, *The Astrophysical Journal*, 905, 131
- Lee, E., Casewell, S. L., Chubb, K. L., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 496, 4674
- Lothringer, J. D., & Barman, T. S. 2020, *The Astronomical Journal*, 159, 289
- MacDonald, R. J., Goyal, J. M., & Lewis, N. K. 2020, *The Astrophysical Journal Letters*, 893, L43
- MacDonald, R. J., Goyal, J. M., & Lewis, N. K. 2020, *ApJL*, 893, L43, doi: [10.3847/2041-8213/ab8238](https://doi.org/10.3847/2041-8213/ab8238)
- Mansfield, M., Bean, J. L., Stevenson, K. B., et al. 2020, *The Astrophysical Journal Letters*, 888, L15
- Melbourne, K., Youngblood, A., France, K., et al. 2020, *The Astronomical Journal*, 160, 269
- Miles, B. E., Skemer, A. J., Morley, C. V., et al. 2020, *The Astronomical Journal*, 160, 63
- Nelson, B. E., Ford, E. B., Buchner, J., et al. 2020, *The Astronomical Journal*, 159, 73
- Oklopčić, A., Silva, M., Montero-Camacho, P., & Hirata, C. M. 2020, *The Astrophysical Journal*, 890, 88
- pandas development team, T. 2020, pandas-dev/pandas: Pandas, latest, Zenodo, doi: [10.5281/zenodo.3509134](https://doi.org/10.5281/zenodo.3509134)
- Piette, A. A., & Madhusudhan, N. 2020, *Monthly Notices of the Royal Astronomical Society*, 497, 5136
- Pino, L., Désert, J.-M., Brogi, M., et al. 2020, *The Astrophysical Journal Letters*, 894, L27
- Pluriel, W., Zingales, T., Leconte, J., & Parmentier, V. 2020, *Astronomy & Astrophysics*, 636, A66
- Pluriel, W., Zingales, T., Leconte, J., & Parmentier, V. 2020, *A&A*, 636, A66, doi: [10.1051/0004-6361/202037678](https://doi.org/10.1051/0004-6361/202037678)
- Seidel, J., Ehrenreich, D., Pino, L., et al. 2020, *Astronomy & Astrophysics*, 633, A86
- Showman, A. P., Tan, X., & Parmentier, V. 2020, *Space Science Reviews*, 216, 1
- Shulyak, D., Lara, L. M., Rengel, M., & Nèmec, N.-E. 2020, *Astronomy & Astrophysics*, 639, A48

- Taylor, J., Parmentier, V., Irwin, P. G., et al. 2020, *Monthly Notices of the Royal Astronomical Society*, 493, 4342
- Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *Nature methods*, 17, 261
- von Essen, C., Mallonn, M., Hermansen, S., et al. 2020, *A&A*, 637, A76, doi: [10.1051/0004-6361/201937169](https://doi.org/10.1051/0004-6361/201937169)
- Winter, A. J., Kruijssen, J. D., Longmore, S. N., & Chevance, M. 2020, *Nature*, 586, 528
- Yurchenko, S., Mellor, T. M., Freedman, R. S., & Tennyson, J. 2020, *Monthly Notices of the Royal Astronomical Society*, 496, 5282
- Zhang, M., Chachan, Y., Kempton, E. M.-R., Knutson, H. A., et al. 2020, *The Astrophysical Journal*, 899, 27
- Baeyens, R., Decin, L., Carone, L., et al. 2021, *MNRAS*, 505, 5603, doi: [10.1093/mnras/stab1310](https://doi.org/10.1093/mnras/stab1310)
- Beltz, H., Rauscher, E., Roman, M. T., & Guilliat, A. 2021, *The Astronomical Journal*, 163, 35
- Borsa, F., Allart, R., Casasayas-Barris, N., et al. 2021, *Astronomy & Astrophysics*, 645, A24
- Casasayas-Barris, N., Orell-Miquel, J., Stangret, M., et al. 2021, arXiv preprint arXiv:2109.00059
- Chen, M., Tworek, J., Jun, H., et al. 2021, arXiv preprint arXiv:2107.03374
- Debras, F., Baruteau, C., & Donati, J.-F. 2021, *Monthly Notices of the Royal Astronomical Society*, 500, 1621
- Deibert, E. K., de Mooij, E. J., Jayawardhana, R., et al. 2021, *The Astrophysical Journal Letters*, 919, L15
- Espinoza, N., & Jones, K. 2021, *The Astronomical Journal*, 162, 165
- Finnerty, L., Buzard, C., Pelletier, S., et al. 2021, *The Astronomical Journal*, 161, 104
- Foreman-Mackey, D., Savel, A., Luger, R., et al. 2021a, *exoplanet-dev/exoplanet* v0.4.5, doi: [10.5281/zenodo.1998447](https://doi.org/10.5281/zenodo.1998447)
- Foreman-Mackey, D., Luger, R., Agol, E., et al. 2021b, *exoplanet: Gradient-based probabilistic inference for exoplanet data & other astronomical time series*, 0.5.1, Zenodo, doi: [10.5281/zenodo.7191939](https://doi.org/10.5281/zenodo.7191939)
- . 2021c, arXiv preprint arXiv:2105.01994
- Fortney, J. J., Dawson, R. I., & Komacek, T. D. 2021, *Journal of Geophysical Research: Planets*, 126, e2020JE006629
- Fu, G., Deming, D., Lothringer, J., et al. 2021, *AJ*, 162, 108, doi: [10.3847/1538-3881/ac1200](https://doi.org/10.3847/1538-3881/ac1200)

- Gao, P., & Powell, D. 2021, *The Astrophysical Journal Letters*, 918, L7
- Gao, P., Wakeford, H. R., Moran, S. E., & Parmentier, V. 2021, *Aerosols in exoplanet atmospheres*, Wiley Online Library
- Giacobbe, P., Brogi, M., Gandhi, S., et al. 2021, *Nature*, 592, 205
- Granger, B. E., & Pérez, F. 2021, *Computing in Science & Engineering*, 23, 7
- Gray, D. F. 2021, *The observation and analysis of stellar photospheres* (Cambridge university press)
- Grimm, S. L., Malik, M., Kitzmann, D., et al. 2021, *The Astrophysical Journal Supplement Series*, 253, 30
- Hammond, M., & Lewis, N. T. 2021, *Proceedings of the National Academy of Sciences*, 118, e2022705118
- Harada, C. K., Kempton, E. M.-R., Rauscher, E., et al. 2021, *The Astrophysical Journal*, 909, 85
- Helling, C., Woters, M., Samra, D., Molaverdikhani, K., & Iro, N. 2021, *Astronomy & Astrophysics*, 648, A80
- Hindle, A., Bushby, P., & Rogers, T. 2021, *The Astrophysical Journal Letters*, 916, L8
- Hindle, A. W., Bushby, P. J., & Rogers, T. M. 2021, *ApJ*, 922, 176, doi: [10.3847/1538-4357/ac0e2e](https://doi.org/10.3847/1538-4357/ac0e2e)
- Ih, J., & Kempton, E. M.-R. 2021, *The Astronomical Journal*, 162, 237
- Kasper, D., Bean, J. L., Line, M. R., et al. 2021, *The Astrophysical Journal Letters*, 921, L18
- Kesseli, A. Y., & Snellen, I. 2021, *The Astrophysical Journal Letters*, 908, L17
- Kuffmeier, M., Dullemond, C., Reissl, S., & Goicovic, F. 2021, *Astronomy & Astrophysics*, 656, A161
- Landman, R., Sánchez-López, A., Mollière, P., et al. 2021, *A&A*, 656, A119, doi: [10.1051/0004-6361/202141696](https://doi.org/10.1051/0004-6361/202141696)
- Langeveld, A. B., Madhusudhan, N., Cabot, S. H., & Hodgkin, S. T. 2021, *Monthly Notices of the Royal Astronomical Society*, 502, 4392
- Lee, E. K., Parmentier, V., Hammond, M., et al. 2021, arXiv preprint arXiv:2106.11664
- Line, M. R., Brogi, M., Bean, J. L., et al. 2021, *Nature*, 598, 580
- Lothringer, J. D., Rustamkulov, Z., Sing, D. K., et al. 2021, *The Astrophysical Journal*, 914, 12
- Malsky, I., Rauscher, E., Kempton, E. M.-R., et al. 2021, *The Astrophysical Journal*, 923, 62
- May, E. M., Komacek, T. D., Stevenson, K. B., et al. 2021, *The Astronomical Journal*, 162, 158

- Miller, M. C., Lamb, F., Dittmann, A., et al. 2021, *The Astrophysical Journal Letters*, 918, L28
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., & Lakshminarayanan, B. 2021, *The Journal of Machine Learning Research*, 22, 2617
- Parmentier, V., Showman, A. P., & Fortney, J. J. 2021, *Monthly Notices of the Royal Astronomical Society*, 501, 78
- Pelletier, S., Benneke, B., Darveau-Bernier, A., et al. 2021, *The Astronomical Journal*, 162, 73
- Pepe, F., Cristiani, S., Rebolo, R., et al. 2021, *Astronomy & Astrophysics*, 645, A96
- Pluriel, W., Leconte, J., Parmentier, V., et al. 2021, arXiv preprint arXiv:2110.09080
- Roman, M. T., Kempton, E. M.-R., Rauscher, E., et al. 2021, *The Astrophysical Journal*, 908, 101
- Roth, A., Drummond, B., Hébrard, E., et al. 2021, *Monthly Notices of the Royal Astronomical Society*
- Savel, A. B., Kempton, E. M.-R., Malik, M., et al. 2021, arXiv preprint arXiv:2109.00163
- Seidel, J., Ehrenreich, D., Allart, R., et al. 2021a, *Astronomy & Astrophysics*, 653, A73
- . 2021b, *Astronomy & Astrophysics*, 653, A73
- Steinrueck, M. E., Showman, A. P., Lavvas, P., et al. 2021, *Monthly Notices of the Royal Astronomical Society*, 504, 2783
- Tabernero, H., Osorio, M. Z., Allart, R., et al. 2021, *Astronomy & Astrophysics*, 646, A158
- Tsai, S.-M., Malik, M., Kitzmann, D., et al. 2021, *The Astrophysical Journal*, 923, 264
- Wardenier, J. P., Parmentier, V., Lee, E. K., Line, M. R., & Gharib-Nezhad, E. 2021, *Monthly Notices of the Royal Astronomical Society*, 506, 1258
- Wardenier, J. P., Parmentier, V., Lee, E. K. H., Line, M. R., & Gharib-Nezhad, E. 2021, *MNRAS*, 506, 1258, doi: [10.1093/mnras/stab1797](https://doi.org/10.1093/mnras/stab1797)
- Zhang, Y., Snellen, I. A. G., & Mollière, P. 2021, *A&A*, 656, A76, doi: [10.1051/0004-6361/202141502](https://doi.org/10.1051/0004-6361/202141502)
- Ahrer, E.-M., Alderson, L., Batalha, N. M., et al. 2022, arXiv preprint arXiv:2208.11692
- Asnodkar, A. P., Wang, J., Eastman, J. D., et al. 2022, *The Astronomical Journal*, 163, 155
- Astropy Collaboration, Price-Whelan, A. M., Lim, P. L., et al. 2022, *ApJ*, 935, 167, doi: [10.3847/1538-4357/ac7c74](https://doi.org/10.3847/1538-4357/ac7c74)
- Bell, T. J., Ahrer, E.-M., Brande, J., et al. 2022, arXiv preprint arXiv:2207.03585
- Beltz, H., Rauscher, E., Kempton, E. M. R., et al. 2022, *AJ*, 164, 140, doi: [10.3847/1538-3881/ac897b](https://doi.org/10.3847/1538-3881/ac897b)

- Beltz, H., Rauscher, E., Kempton, E. M.-R., et al. 2022, *The Astronomical Journal*, 164, 140
- Ben-Jaffel, L., Ballester, G. E., Muñoz, A. G., et al. 2022, *Nature Astronomy*, 6, 141
- Carleo, I., Giacobbe, P., Guilluy, G., et al. 2022, *The Astronomical Journal*, 164, 101
- Chubb, K. L., & Min, M. 2022, *Astronomy & Astrophysics*, 665, A2
- De Regt, S., Kesseli, A., Snellen, I., Merritt, S., & Chubb, K. 2022, *Astronomy & Astrophysics*, 661, A109
- Gabrié, M., Rotskoff, G. M., & Vanden-Eijnden, E. 2022, *Proceedings of the National Academy of Sciences*, 119, e2109420119
- Gandhi, S., Kesseli, A., Snellen, I., et al. 2022, *Monthly Notices of the Royal Astronomical Society*
- Gibson, N. P., Nugroho, S. K., Lothringer, J., Maguire, C., & Sing, D. K. 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 4618
- Gonzales, E. C., Burningham, B., Faherty, J. K., et al. 2022, *The Astrophysical Journal*, 938, 56
- Guillot, T., Fletcher, L. N., Helled, R., et al. 2022, arXiv preprint arXiv:2205.04100
- Guilluy, G., Giacobbe, P., Carleo, I., et al. 2022, *Astronomy & Astrophysics*, 665, A104
- Gully-Santiago, M., & Morley, C. V. 2022, *The Astrophysical Journal*, 941, 200
- Hobbs, R., Shorttle, O., & Madhusudhan, N. 2022, *Monthly Notices of the Royal Astronomical Society*, 516, 1032
- Kawahara, H., Kawashima, Y., Masuda, K., et al. 2022, *The Astrophysical Journal Supplement Series*, 258, 31
- Kesseli, A. Y., Snellen, I., Casasayas-Barris, N., Mollière, P., & Sánchez-López, A. 2022, *The Astronomical Journal*, 163, 107
- Knierim, H., Batygin, K., & Bitsch, B. 2022, *Astronomy & Astrophysics*, 658, L7
- Komacek, T. D., Tan, X., Gao, P., & Lee, E. K. 2022, arXiv preprint arXiv:2205.07834
- Lally, M., & Vanderburg, A. 2022, *The Astronomical Journal*, 163, 181
- Lee, E. K., Wardenier, J. P., Prinoth, B., et al. 2022, *The Astrophysical Journal*, 929, 180
- MacDonald, R. J., & Lewis, N. K. 2022, *The Astrophysical Journal*, 929, 20
- Maguire, C., Gibson, N. P., Nugroho, S. K., et al. 2022, *Monthly Notices of the Royal Astronomical Society*
- Meech, A., Aigrain, S., Brogi, M., & Birkby, J. L. 2022, *Monthly Notices of the Royal Astronomical Society*, 512, 2604

- Mollière, P., Molyarova, T., Bitsch, B., et al. 2022, *The Astrophysical Journal*, 934, 74
- Niraula, P., de Wit, J., Gordon, I. E., et al. 2022, *Nature Astronomy*, 6, 1287
- Nixon, M. C., & Madhusudhan, N. 2022, *The Astrophysical Journal*, 935, 73
- Prinoth, B., Hoeijmakers, H. J., Kitzmann, D., et al. 2022, *Nature Astronomy*, 6, 449
- Raymond, S. N., & Morbidelli, A. 2022, *Demographics of exoplanetary systems: Lecture notes of the 3rd advanced school on exoplanetary science*, 3
- Rustamkulov, Z., Sing, D. K., Liu, R., & Wang, A. 2022, *The Astrophysical Journal Letters*, 928, L7
- Rustamkulov, Z., Sing, D. K., Mukherjee, S., et al. 2022, arXiv e-prints, arXiv:2211.10487, doi: [10.48550/arXiv.2211.10487](https://doi.org/10.48550/arXiv.2211.10487)
- Sánchez-López, A., Landman, R., Mollière, P., et al. 2022, *Astronomy & Astrophysics*, 661, A78
- Savel, A. B., Kempton, E. M.-R., Malik, M., et al. 2022, *The Astrophysical Journal*, 926, 85
- Silva, T. A., Demangeon, O., Santos, N., et al. 2022, *Astronomy & Astrophysics*, 666, L10
- Stock, J. W., Kitzmann, D., & Patzer, A. B. C. 2022, *Monthly Notices of the Royal Astronomical Society*, 517, 4070
- Teal, D., Kempton, E. M.-R., Bastelberger, S., Youngblood, A., & Arney, G. 2022, *The Astrophysical Journal*, 927, 90
- Wang, S. X., Latouf, N., Plavchan, P., et al. 2022, *The Astronomical Journal*, 164, 211
- Wardenier, J. P., Parmentier, V., & Lee, E. K. 2022, *Monthly Notices of the Royal Astronomical Society*, 510, 620
- Welbanks, L., & Madhusudhan, N. 2022, *The Astrophysical Journal*, 933, 79
- Wong, K. W. K., Gabríel, M., & Foreman-Mackey, D. 2022, arXiv e-prints, arXiv:2211.06397. <https://arxiv.org/abs/2211.06397>
- Zingales, T., Falco, A., Pluriel, W., & Leconte, J. 2022, *Astronomy & Astrophysics*, 667, A13
- Alderson, L., Wakeford, H. R., Alam, M. K., et al. 2023, *Nature*, 1
- August, P. C., Bean, J. L., Zhang, M., et al. 2023, arXiv e-prints, arXiv:2305.07753, doi: [10.48550/arXiv.2305.07753](https://doi.org/10.48550/arXiv.2305.07753)
- Bean, J. L., Xue, Q., August, P. C., et al. 2023, *Nature*, 618, 43
- Bell, T. J., Welbanks, L., Schlawin, E., et al. 2023, *Nature*, 623, 709
- Beltz, H., Rauscher, E., Kempton, E. M.-R., Malsky, I., & Savel, A. B. 2023, *The Astronomical Journal*, 165, 257

- Boldt-Christmas, L., Lesjak, F., Wehrhahn, A., et al. 2023, arXiv e-prints, arXiv:2312.08320, doi: [10.48550/arXiv.2312.08320](https://doi.org/10.48550/arXiv.2312.08320)
- Borsato, N., Hoeijmakers, H., Prinoth, B., et al. 2023, *Astronomy & Astrophysics*, 673, A158
- Boucher, A., Lafrenière, D., Pelletier, S., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, stad1247
- Brogi, M., Emeka-Okafor, V., Line, M. R., et al. 2023, *The Astronomical Journal*, 165, 91
- Bushouse, H., Eisenhamer, J., Dencheva, N., et al. 2023, JWST Calibration Pipeline, 1.11.4, Zenodo, doi: [10.5281/zenodo.8247246](https://doi.org/10.5281/zenodo.8247246)
- Chachan, Y., Knutson, H. A., Lothringer, J., & Blake, G. A. 2023, *The Astrophysical Journal*, 943, 112
- Cheverall, C. J., Madhusudhan, N., & Holmberg, M. 2023, arXiv preprint arXiv:2303.01496
- Constantinou, S., Madhusudhan, N., & Gandhi, S. 2023, *The Astrophysical Journal Letters*, 943, L10
- Coulombe, L.-P., Benneke, B., Challener, R., et al. 2023, *Nature*, 620, 292
- Currie, M. H., Meadows, V. S., & Rasmussen, K. C. 2023, *The Planetary Science Journal*, 4, 83
- Esparza-Borges, E., López-Morales, M., Redai, J. I. A., et al. 2023, *The Astrophysical Journal Letters*, 955, L19
- Feinstein, A. D., Radica, M., Welbanks, L., et al. 2023, *Nature*, 1
- Gandhi, S., de Regt, S., Snellen, I., et al. 2023a, *The Astrophysical Journal Letters*, 957, L36
- Gandhi, S., Kesseli, A., Zhang, Y., et al. 2023b, *The Astronomical Journal*, 165, 242
- Gebhard, T. D., Wildberger, J., Dax, M., et al. 2023, arXiv preprint arXiv:2312.08295
- Grant, D., Lewis, N. K., Wakeford, H. R., et al. 2023a, *The Astrophysical Journal Letters*, 956, L29
- Grant, D., Lothringer, J. D., Wakeford, H. R., et al. 2023b, *The Astrophysical Journal Letters*, 949, L15
- Gupta, A., Miotello, A., Manara, C. F., et al. 2023, *Astronomy & Astrophysics*, 670, L8
- Hardegree-Ullman, K. K., Apai, D., Bergsten, G. J., Pascucci, I., & López-Morales, M. 2023, *The Astronomical Journal*, 165, 267
- Helling, C., Samra, D., Lewis, D., et al. 2023, *Astronomy & Astrophysics*, 671, A122
- Lange, J. U. 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 3181

- Lemos, P., Weaverdyck, N., Rollins, R. P., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 521, 1184
- Madhusudhan, N., Sarkar, S., Constantinou, S., et al. 2023, *The Astrophysical Journal Letters*, 956, L13
- Niraula, P., de Wit, J., Gordon, I. E., Hargreaves, R. J., & Sousa-Silva, C. 2023, *The Astrophysical Journal Letters*, 950, L17
- Owen, J. E., Murray-Clay, R. A., Schreyer, E., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 518, 4357
- Palle, E., Biazzo, K., Bolmont, E., et al. 2023, arXiv preprint arXiv:2311.17075
- Pelletier, S., Benneke, B., Ali-Dib, M., et al. 2023, *Nature*, 1
- Pineda, J. S., & Villadsen, J. 2023, *Nature Astronomy*, 7, 569
- Prinoth, B., Hoeijmakers, H. J., Pelletier, S., et al. 2023, *Astronomy & Astrophysics*, 678, A182
- Program, T. T. E. C. E. R. S. 2023, *Nature*, 614, 649
- Savel, A. B., Kempton, E. M.-R., Rauscher, E., et al. 2023, *The Astrophysical Journal*, 944, 99
- Soriano-Guerrero, C., Viganò, D., Perna, R., Akgün, T., & Palenzuela, C. 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 626
- Tsai, S.-M., Lee, E. K., Powell, D., et al. 2023, *Nature*, 617, 483
- Turner, J. D., Zarka, P., Griessmeier, J.-M., et al. 2023, arXiv preprint arXiv:2310.05363
- van Sluijs, L., Birkby, J. L., Lothringer, J., et al. 2023a, *Monthly Notices of the Royal Astronomical Society*, 522, 2145
- . 2023b, *Monthly Notices of the Royal Astronomical Society*, 522, 2145
- Wardenier, J. P., Parmentier, V., Line, M. R., & Lee, E. K. 2023, *Monthly Notices of the Royal Astronomical Society*, 525, 4942
- Welbanks, L., Line, M., & JWST Transiting Exoplanet Community ERS Team. 2023, in *American Astronomical Society Meeting Abstracts*, Vol. 241, American Astronomical Society Meeting Abstracts #241, 124.07
- Wong, K. W., Isi, M., & Edwards, T. D. 2023a, arXiv preprint arXiv:2302.05333
- Wong, K. W. k., Gabrié, M., & Foreman-Mackey, D. 2023b, *J. Open Source Softw.*, 8, 5021, doi: [10.21105/joss.05021](https://doi.org/10.21105/joss.05021)
- Zamyatina, M., Hébrard, E., Drummond, B., et al. 2023, *Monthly Notices of the Royal Astronomical Society*, 519, 3129

- Bell, T. J., Crouzet, N., Cubillos, P. E., et al. 2024, *Nature Astronomy*, 1
- Blain, D., Sánchez-López, A., & Mollière, P. 2024, *The Astronomical Journal*, 167, 179
- Brain, D. A., Kao, M. M., & O'Rourke, J. G. 2024, arXiv preprint arXiv:2404.15429
- Carter, A., May, E., Espinoza, N., et al. 2024, *Nature Astronomy*, 8, 1008
- Castro-González, A., Bourrier, V., Lillo-Box, J., et al. 2024, *Astronomy & Astrophysics*, 689, A250
- Challener, R. C., Rustamkulov, Z., Lee, E. K., et al. 2024, arXiv preprint arXiv:2406.10207
- Cheverall, C. J., & Madhusudhan, N. 2024, *The Astronomical Journal*, 167, 272
- Christie, D., Mayne, N., Zamyatina, M., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 532, 3001
- Constantinou, S., & Madhusudhan, N. 2024, *Monthly Notices of the Royal Astronomical Society*, 530, 3252
- Debras, F., Klein, B., Donati, J.-F., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 566
- Dittmann, A. J. 2024, arXiv preprint arXiv:2404.16928
- Falco, A., Leconte, J., Mechineau, A., & Pluriel, W. 2024, *Astronomy & Astrophysics*, 685, A125
- Feinberg, L., Ziemer, J., Ansdell, M., et al. 2024, in *Space Telescopes and Instrumentation 2024: Optical, Infrared, and Millimeter Wave*, Vol. 13092, SPIE, 548–559
- Fisher, C., Taylor, J., Parmentier, V., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 535, 27
- Fu, G., Welbanks, L., Deming, D., et al. 2024, *Nature*, 632, 752
- Hillier, A. S. 2024, *Philosophical Transactions of the Royal Society A*, 382, 20230229
- Hood, T., Debras, F., Moutou, C., et al. 2024, *Astronomy & Astrophysics*, 687, A119
- Hopkins, P. F., Squire, J., Skolidis, R., & Soliman, N. H. 2024, arXiv e-prints, arXiv:2405.06026.
<https://arxiv.org/abs/2405.06026>
- Kanumalla, K., Line, M. R., Mansfield, M. W., et al. 2024, arXiv preprint arXiv:2406.14072
- Kempton, E. M.-R., & Knutson, H. A. 2024, *Reviews in Mineralogy and Geochemistry*, 90, 411
- Khorshid, N., Min, M., Polman, J., & Waters, L. 2024, *Astronomy & Astrophysics*, 685, A64
- Kirk, J., Ahrer, E.-M., Penzlin, A. B., et al. 2024, *RAS Techniques and Instruments*, 3, 691

- Kitzmann, D., Stock, J. W., & Patzer, A. B. C. 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 7263
- Klein, B., Debras, F., Donati, J.-F., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 544
- Lueber, A., Novais, A., Fisher, C., & Heng, K. 2024, *Astronomy & Astrophysics*, 687, A110
- Maguire, C., Gibson, N. P., Nugroho, S. K., et al. 2024, *Astronomy & Astrophysics*, 687, A49
- Malsky, I., Rauscher, E., Roman, M. T., et al. 2024, *The Astrophysical Journal*, 961, 66
- Mansfield, M. W., Line, M. R., Wardenier, J. P., et al. 2024, *The Astronomical Journal*, 168, 14
- Matthews, E. C., Carter, A. L., Pathak, P., et al. 2024, *Nature*, 633, 789
- Narang, M., Puravankara, M., Chandra, C. I., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 529, 1161
- Nortmann, L., Lesjak, F., Yan, F., et al. 2024, arXiv preprint arXiv:2404.12363
- Powell, D., Feinstein, A. D., Lee, E. K., et al. 2024, *Nature*, 626, 979
- Roth, A., Parmentier, V., & Hammond, M. 2024, *Monthly Notices of the Royal Astronomical Society*, 531, 1056
- Savel, A. B., Bedell, M., & Kempton, E. M.-r. 2024a, *Journal of Open Source Software*, 9, 6104
- Savel, A. B., Beltz, H., Komacek, T. D., Tsai, S.-M., & Kempton, E. M.-R. 2024b, *The Astrophysical Journal Letters*, 969, L27
- Schleich, S., Saikia, S. B., Changeat, Q., et al. 2024, *Astronomy & Astrophysics*, 690, A336
- Schreyer, E., Owen, J. E., Spake, J. J., Bahroloom, Z., & Di Giampasquale, S. 2024, *Monthly Notices of the Royal Astronomical Society*, 527, 5117
- Shiohira, Y., Fujii, Y., Kita, H., et al. 2024, *Monthly Notices of the Royal Astronomical Society*, 528, 2136
- Sing, D. K., Evans-Soma, T. M., Rustamkulov, Z., et al. 2024, *The Astronomical Journal*, 168, 231
- Smith, P. C., Line, M. R., Bean, J. L., et al. 2024a, *The Astronomical Journal*, 167, 110
- Smith, P. C., Sanchez, J. A., Line, M. R., et al. 2024b, *The Astronomical Journal*, 168, 293
- Thorngren, D. P. 2024, arXiv preprint arXiv:2405.05307
- Tuchow, N. W., Stark, C. C., & Mamajek, E. 2024, *The Astronomical Journal*, 167, 139
- Villanueva, G. L., Fauchez, T. J., Kofman, V., et al. 2024, *The Planetary Science Journal*, 5, 64

Wardenier, J. P., Parmentier, V., Line, M. R., et al. 2024, Publications of the Astronomical Society of the Pacific, 136, 084403

Welbanks, L., Bell, T. J., Beatty, T. G., et al. 2024, Nature, 630, 836

Wilkinson, C., Charnay, B., Mazevet, S., et al. 2024, Astronomy & Astrophysics, 692, A113

Xue, Q., Bean, J. L., Zhang, M., et al. 2024, The Astrophysical Journal Letters, 963, L5

Arnold, K. E., Savel, A. B., Kempton, E. M.-R., et al. 2025, The Astrophysical Journal, 986, 187

Boldt-Christmas, L., Rains, A. D., Piskunov, N., et al. 2025, arXiv preprint arXiv:2508.18964

Changeat, Q., Bardet, D., Chubb, K., et al. 2025, Astronomy & Astrophysics, 699, A219

Christie, D. A., Evans-Soma, T. M., Mayne, N. J., & Kohary, K. 2025, Monthly Notices of the Royal Astronomical Society, 541, 2773

Coulombe, L.-P., Radica, M., Benneke, B., et al. 2025, Nature Astronomy, 9, 512

Crotts, K. A., Carter, A. L., Lawson, K., et al. 2025, The Astrophysical Journal Letters, 987, L41

Esparza-Borges, E., López-Morales, M., Pallé, E., et al. 2025, Monthly Notices of the Royal Astronomical Society, 543, 3456

Feinstein, A. D., Booth, R. A., Bergner, J. B., et al. 2025, arXiv preprint arXiv:2506.00669

Fu, G., Mukherjee, S., Stevenson, K. B., et al. 2025, The Astrophysical Journal Letters, 989, L17

Gapp, C., Evans-Soma, T. M., Barstow, J. K., et al. 2025, The Astronomical Journal, 169, 341

Group, T. D. H. D. S., Albert, J., Balázs, C., et al. 2025, Computer Physics Communications, 109756

Hardegree-Ullman, K. K., Apai, D., Haffert, S. Y., et al. 2025, The Astronomical Journal, 169, 171

Kawahara, H., Kawashima, Y., Tada, S., et al. 2025, The Astrophysical Journal, 985, 263

Lagrange, A.-M., Wilkinson, C., Mâlin, M., et al. 2025, Nature, 642, 905

Lee, E. K., & Ohno, K. 2025, Astronomy & Astrophysics, 695, A111

Ma, S., Saba, A., Al-Refaie, A. F., et al. 2025, arXiv preprint arXiv:2504.07823

Meech, A., Claringbold, A. B., Ahrer, E.-M., et al. 2025, Monthly Notices of the Royal Astronomical Society, 539, 1381

Mollière, P., Kühnle, H., Matthews, E. C., et al. 2025, Astronomy & Astrophysics, 703, A79

Nortmann, L., Lesjak, F., Yan, F., et al. 2025, Astronomy & Astrophysics, 693, A213

Palle, E., Biazzo, K., Bolmont, E., et al. 2025, Experimental Astronomy, 59, 29

- Panwar, V., Brogi, M., Kanumalla, K., et al. 2025, Monthly Notices of the Royal Astronomical Society, 541, 3561
- Pelletier, S., Benneke, B., Chachan, Y., et al. 2025, The Astronomical Journal, 169, 10
- Roy-Perez, J., Pérez-Hoyos, S., Barrado-Izagirre, N., & Chen-Chen, H. 2025, Astronomy & Astrophysics, 694, A249
- Savel, A. B., Bedell, M., Kempton, E. M.-R., et al. 2025, The Astronomical Journal, 169, 135
- Seager, S., Welbanks, L., Ellerbroek, L., Bains, W., & Petkowski, J. J. 2025, Proceedings of the National Academy of Sciences, 122, e2416188122
- Snellen, I. A. 2025, Annual Review of Astronomy and Astrophysics, 63
- van Sluijs, L., Beltz, H., Malsky, I., et al. 2025, The Astrophysical Journal, 990, 106
- Welbanks, L., Nixon, M. C., McGill, P., et al. 2025, Nature Astronomy, 1
- Wogan, N. F., Batalha, N. E., Zahnle, K., et al. 2025, The Planetary Science Journal, 6, 256
- Zhang, M., Paragas, K., Bean, J. L., et al. 2025, The Astronomical Journal, 169, 38
- Boss, A. P. 2026, arXiv e-prints, arXiv:2604.05178. <https://arxiv.org/abs/2604.05178>
- Claringbold, A. B., Fisher, C. E., Kirk, J., et al. 2026, Monthly Notices of the Royal Astronomical Society, stag143
- de Regt, S., Whiteford, N., Miles, B. E., et al. 2026, arXiv e-prints, arXiv:2607.00952. <https://arxiv.org/abs/2607.00952>
- Heinke, L., Min, M., Bouwman, J., et al. 2026, arXiv preprint arXiv:2604.01219
- Li, J., Bergin, E. A., Hirschmann, M. M., et al. 2026, The Astrophysical Journal Letters, 997, L29
- Ouyang, Q., Yan, F., Liu, S., et al. 2026, arXiv preprint arXiv:2605.30871
- Pelletier, S., Coulombe, L.-P., Splinter, J., et al. 2026, Astronomy & Astrophysics, 706, A2
- Snellen, I., Miguel, Y., Janssen, L., et al. 2026, arXiv e-prints, arXiv:2606.11866. <https://arxiv.org/abs/2606.11866>
- Welbanks, L., Nixon, M. C., McGill, P., et al. 2026, Nature Astronomy, 10, 234