From Cross-Correlations to Likelihoods:

A fully Bayesian approach to high-resolution retrievals

(see also Matteo's talk yesterday)

Dr. Neale Gibson
Trinity College Dublin, The University of Dublin
n.gibson@tcd.ie

Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Image Credit: ESO / Y. Beletsky

Today I'll focus on transmission spectroscopy

- composition
- atmospheric chemistry
- scattering properties
- temperature structure
- line shapes/shifts ➡ dynamics
- biomarkers?

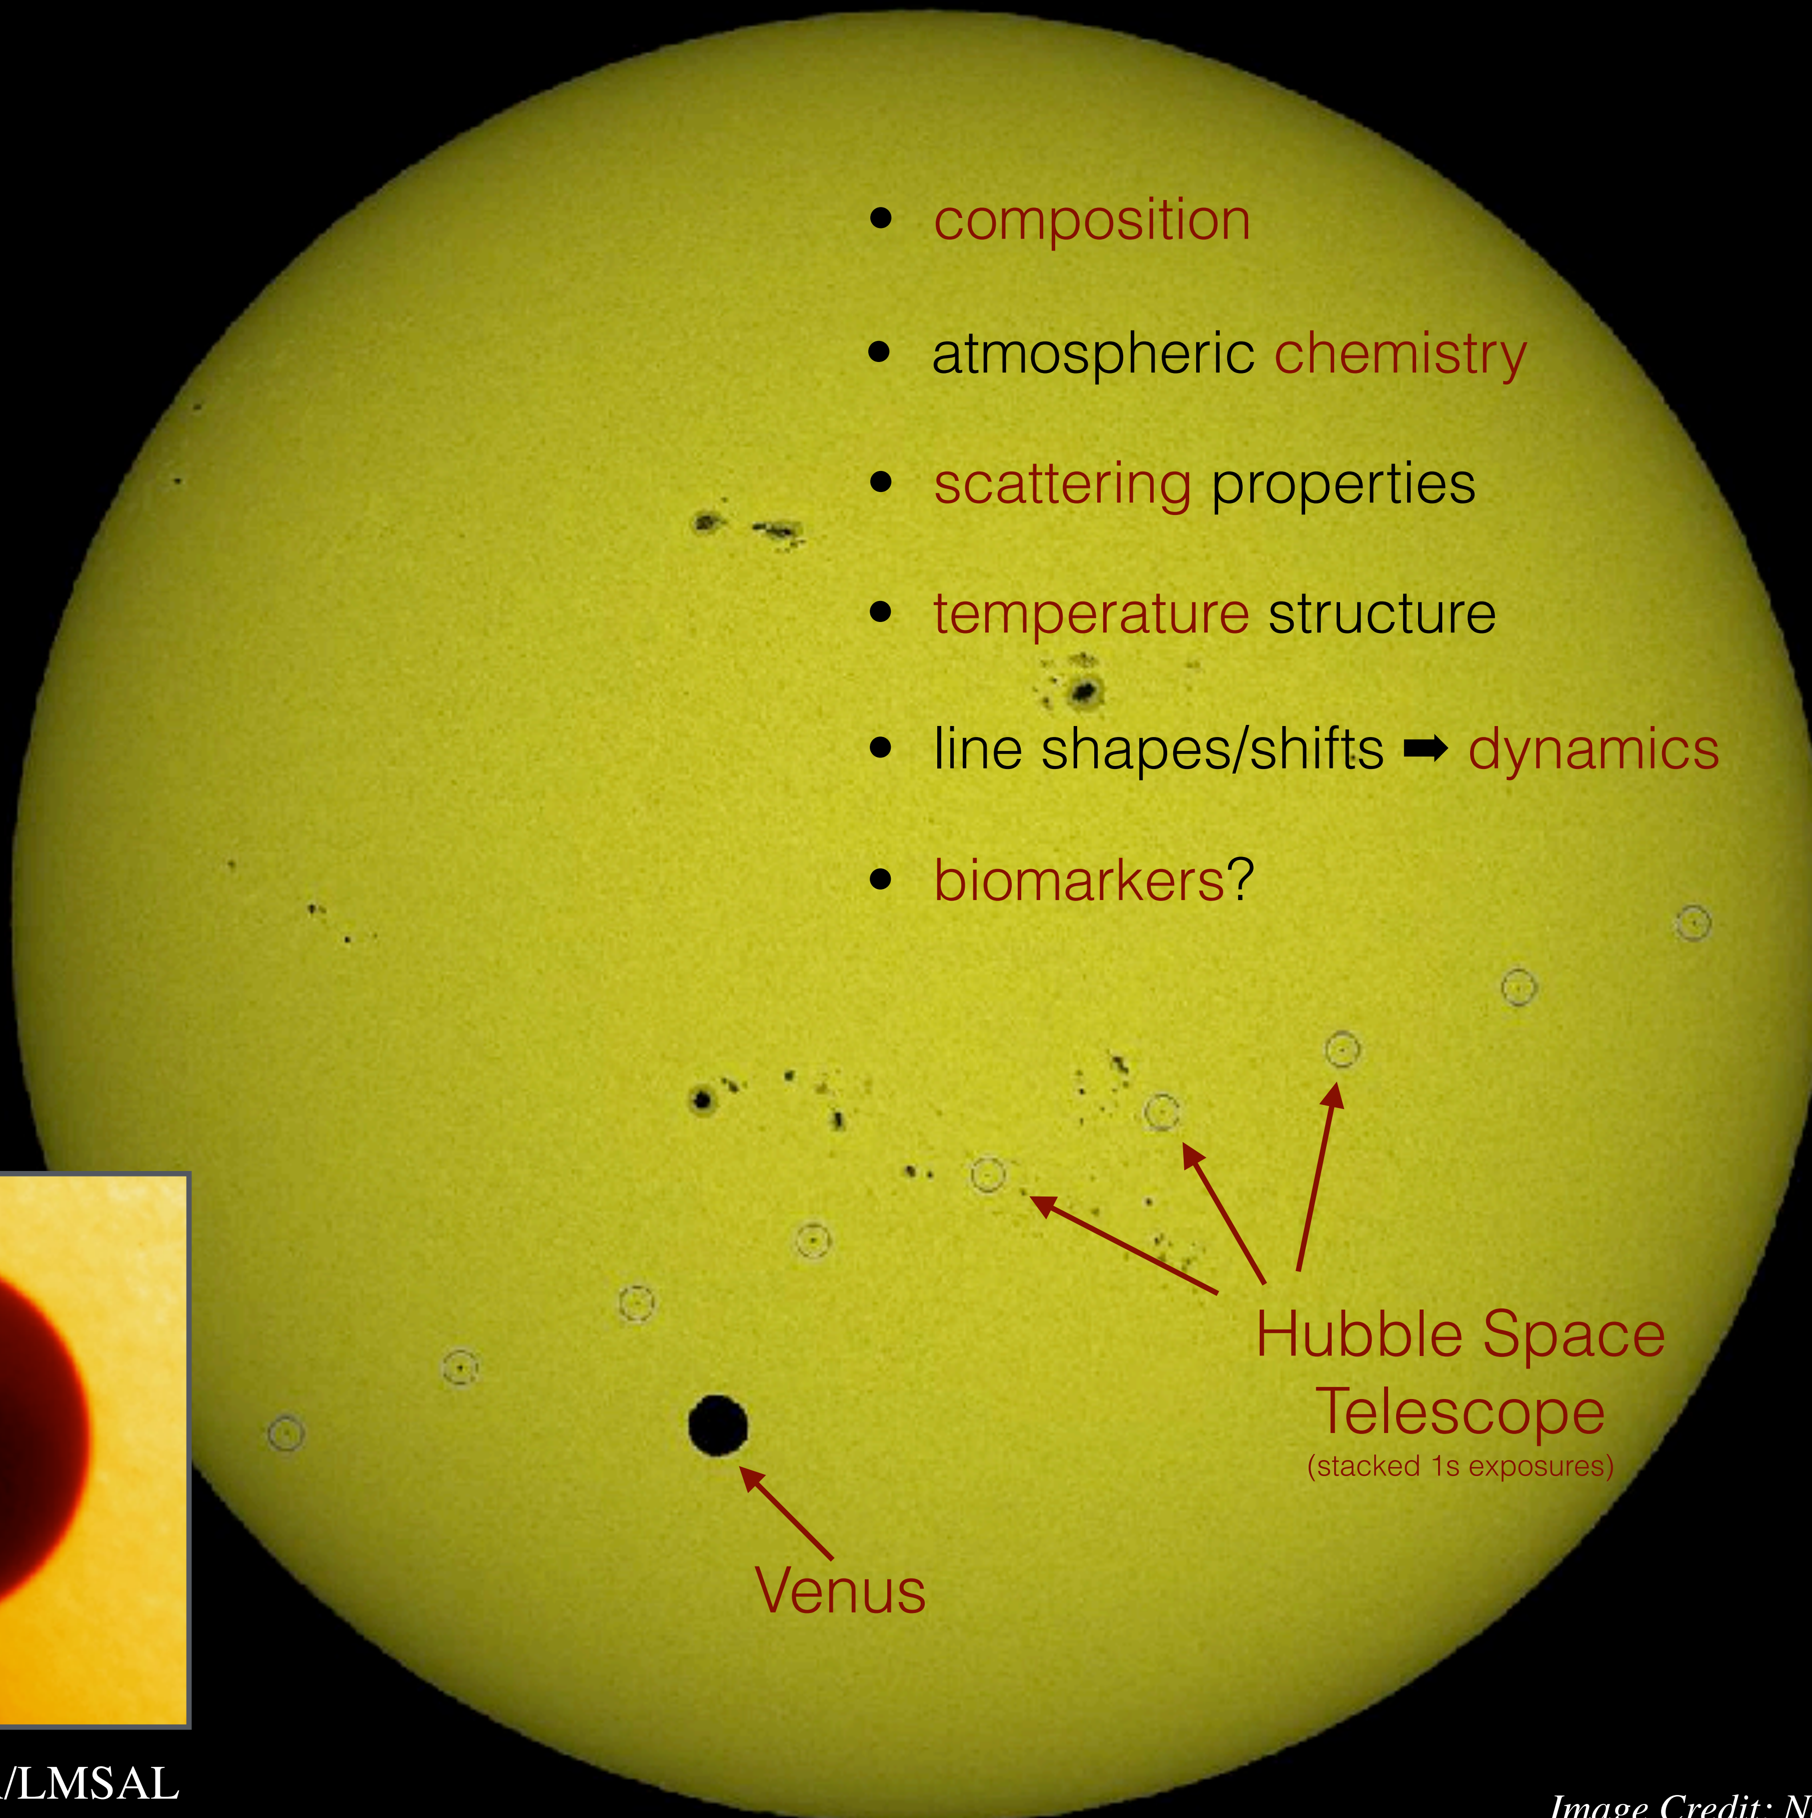

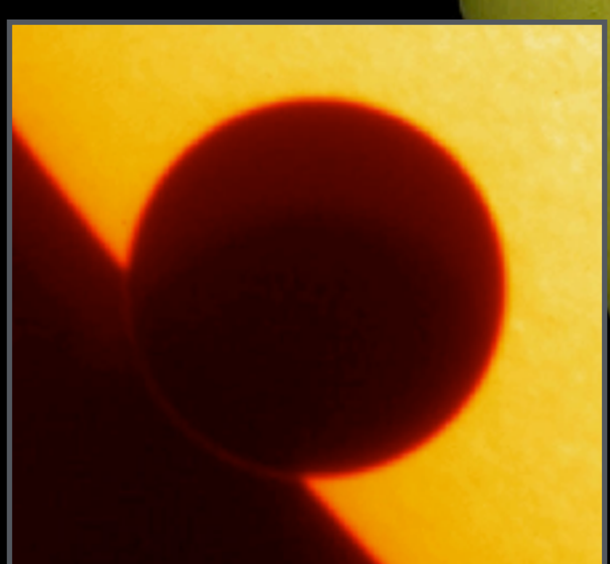

Image credit: NASA/LMSAL

Image Credit: Nasa.org/Thierry Lagault

‘Traditionally’ we use low-resolution time-series observations…

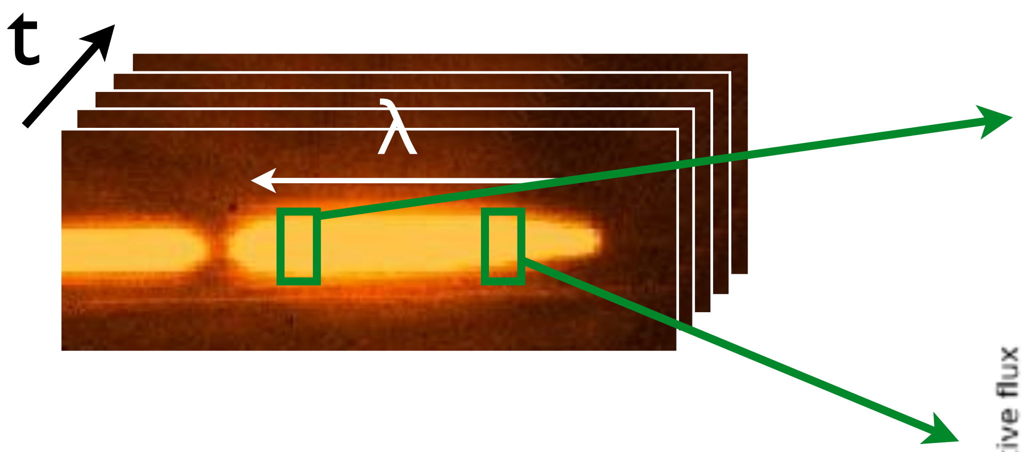

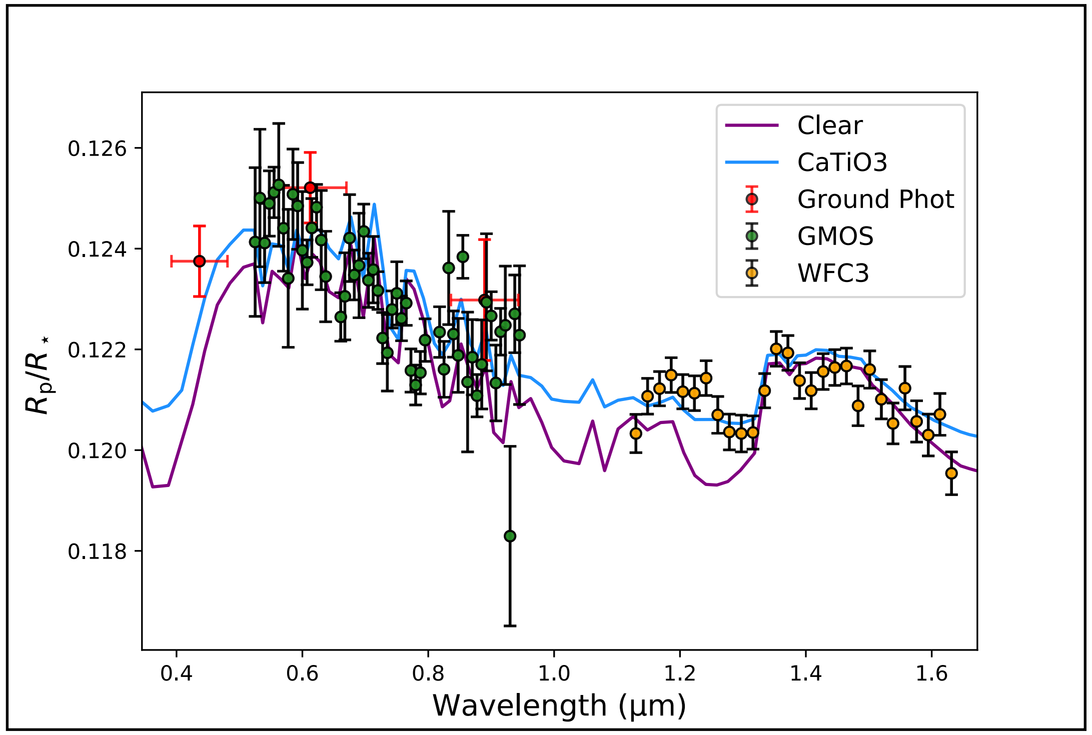

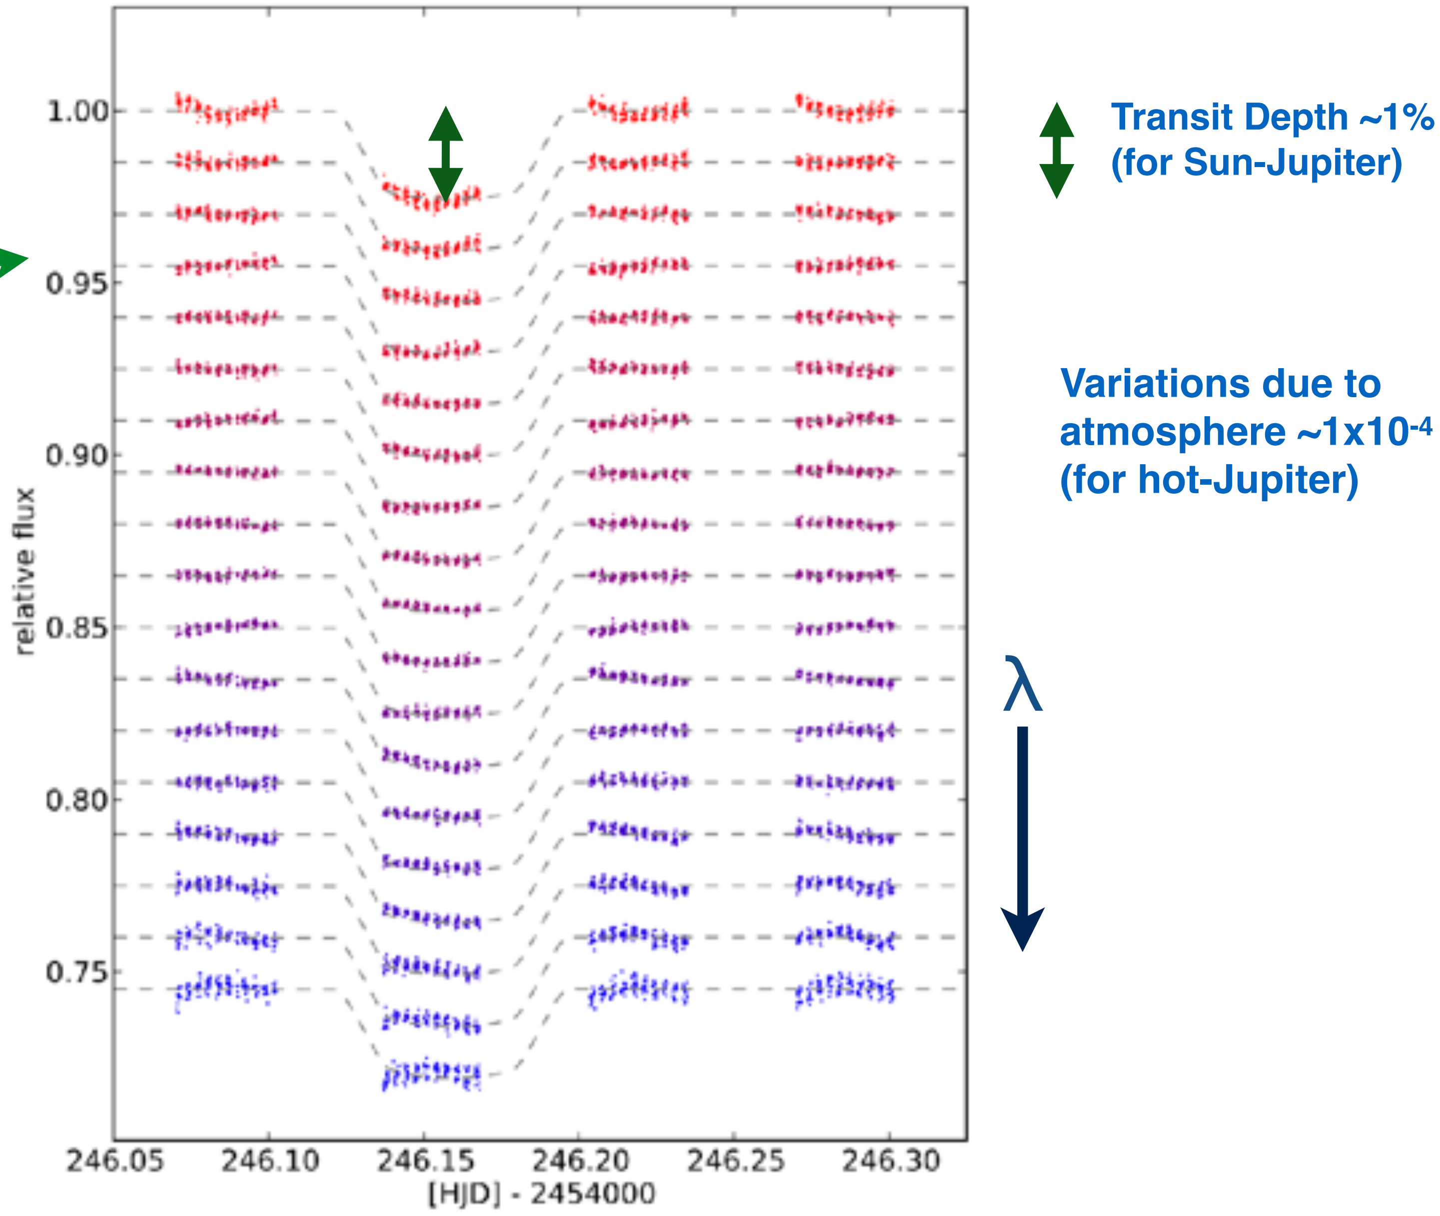

Gibson et al. (2011); Wilson et al. (2021)

…which can be extended to eclipses and phase curves.

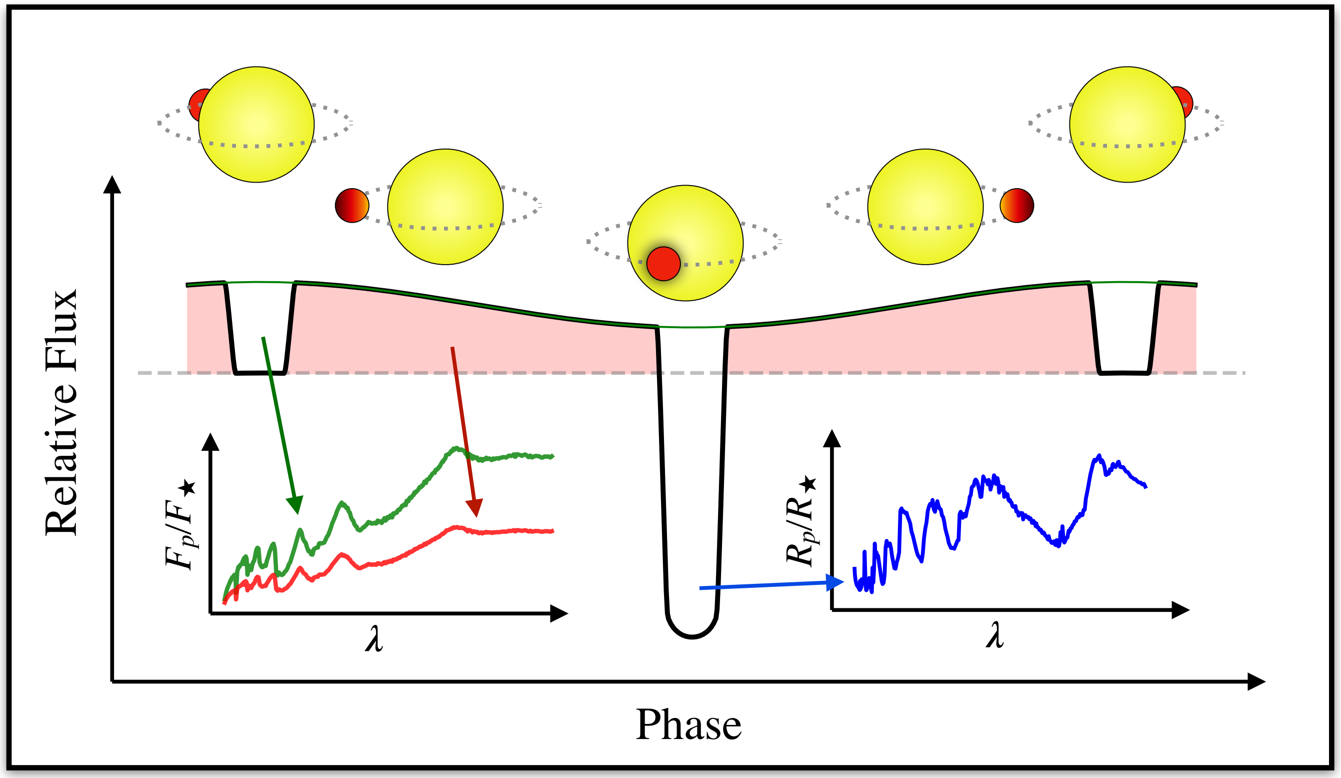

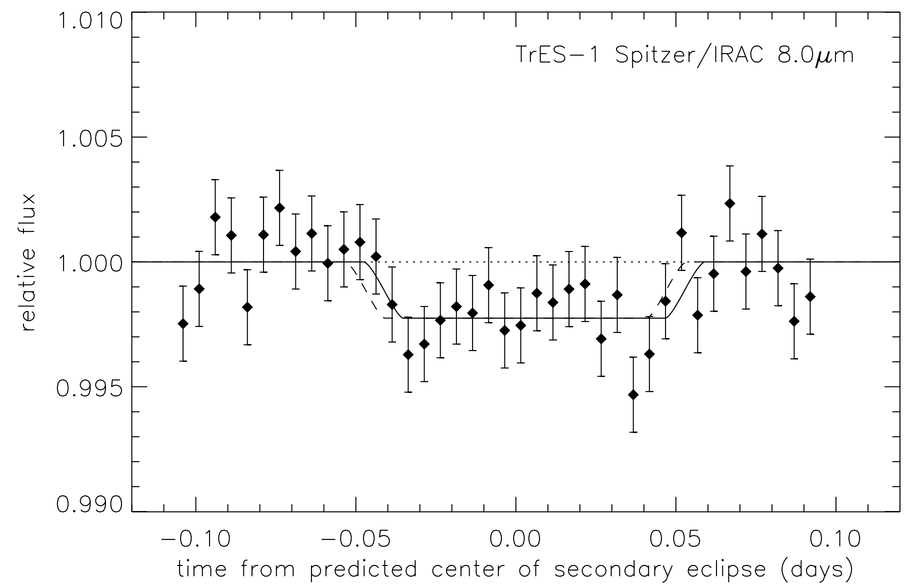

Charbonneau et al (2005); Deming et al. (2005)

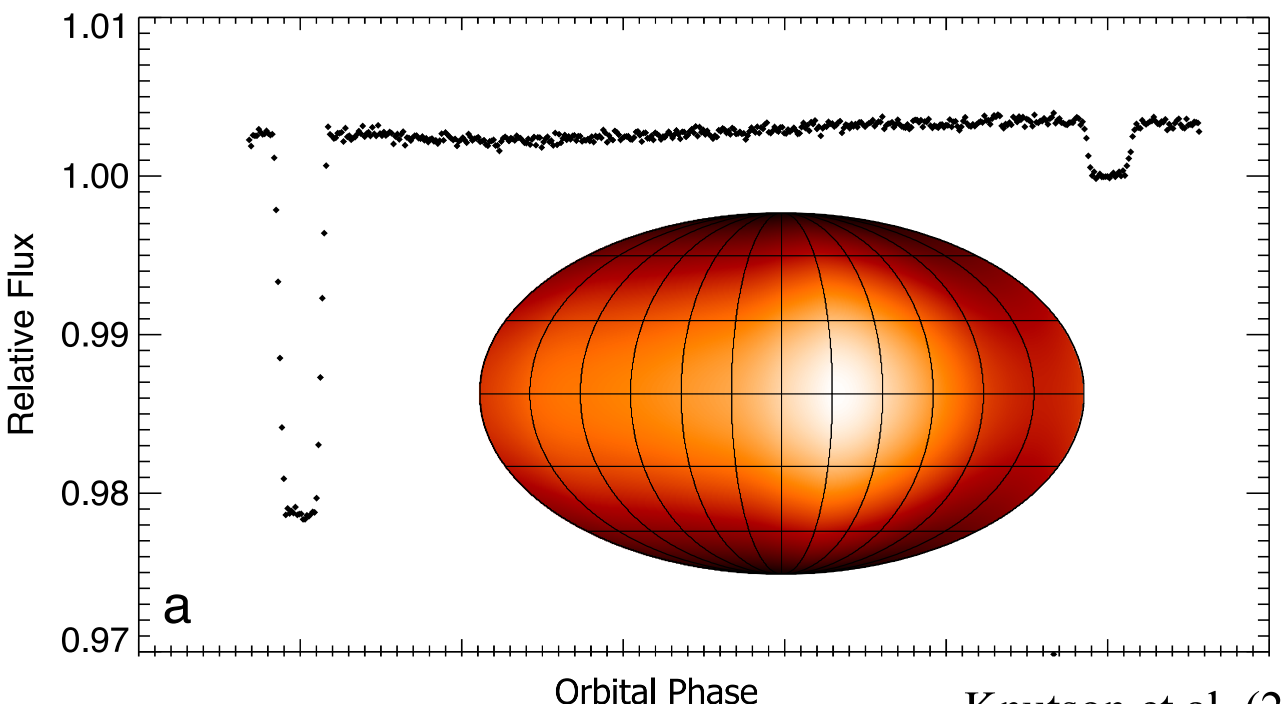

Knutson et al. (2007)

With low-res observations, atmospheric retrieval is decoupled from the light curve analysis

Wilson et al (2021); Gibson et al. (2020)

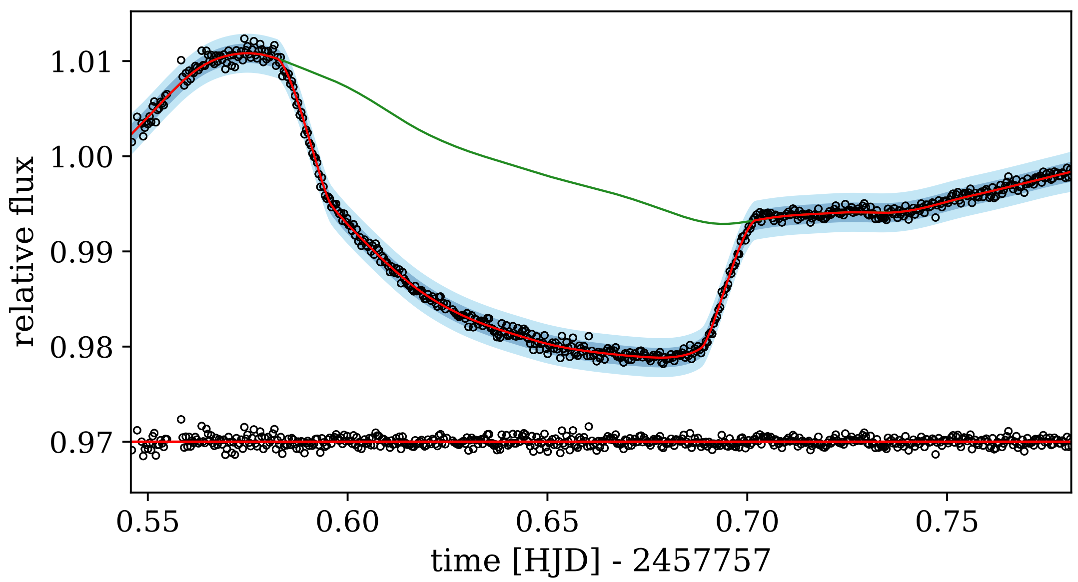

"Observers"

1) Extract light curves from raw data

2) Fit light curves and <u>*infer noise properties*</u>

"Modellers"

3) Model exoplanet spectrum to infer atmospheric properties

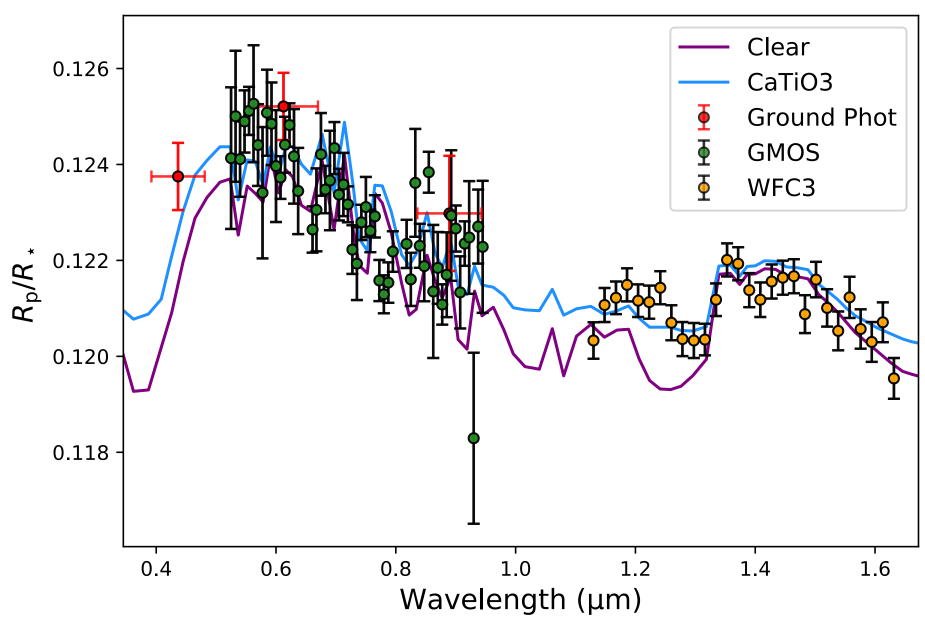

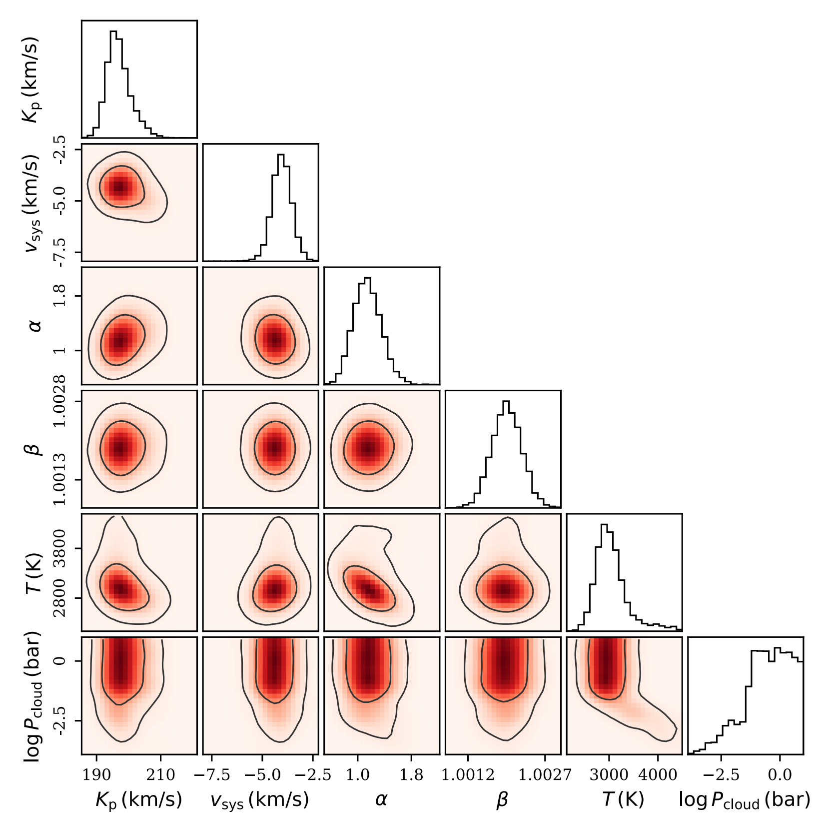

- A major advantage that we can decouple retrievals from low-level data analysis

High-res observations are much trickier

Pre-processed spectra

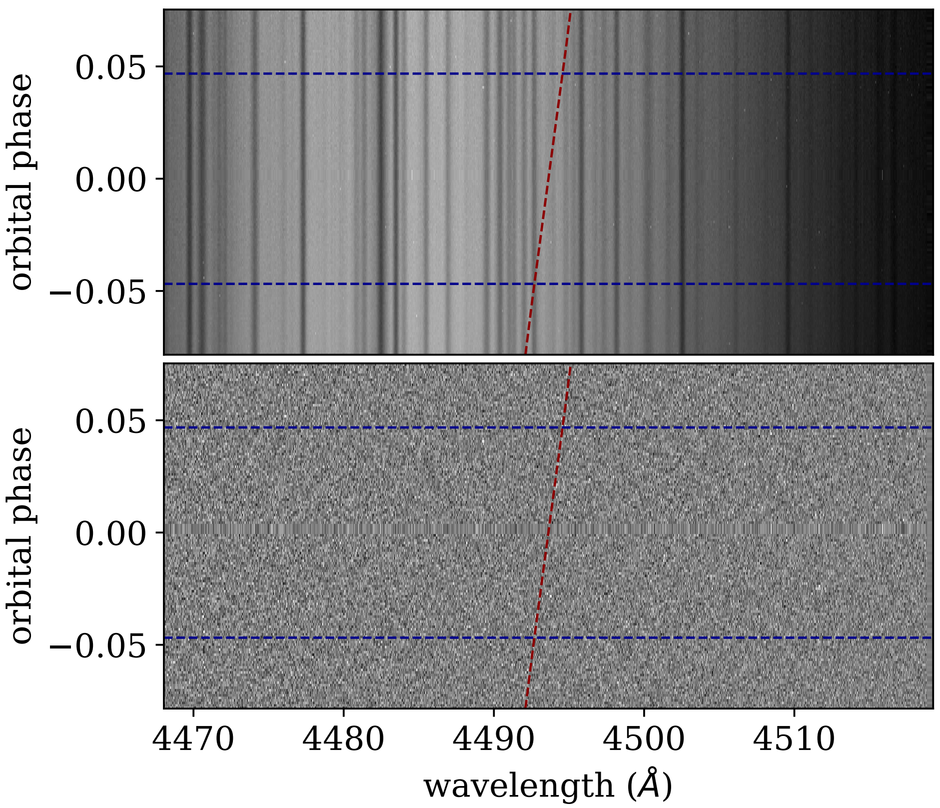

Directly from spectra to atmospheric modelling!

Must infer noise properties

Posterior distribution

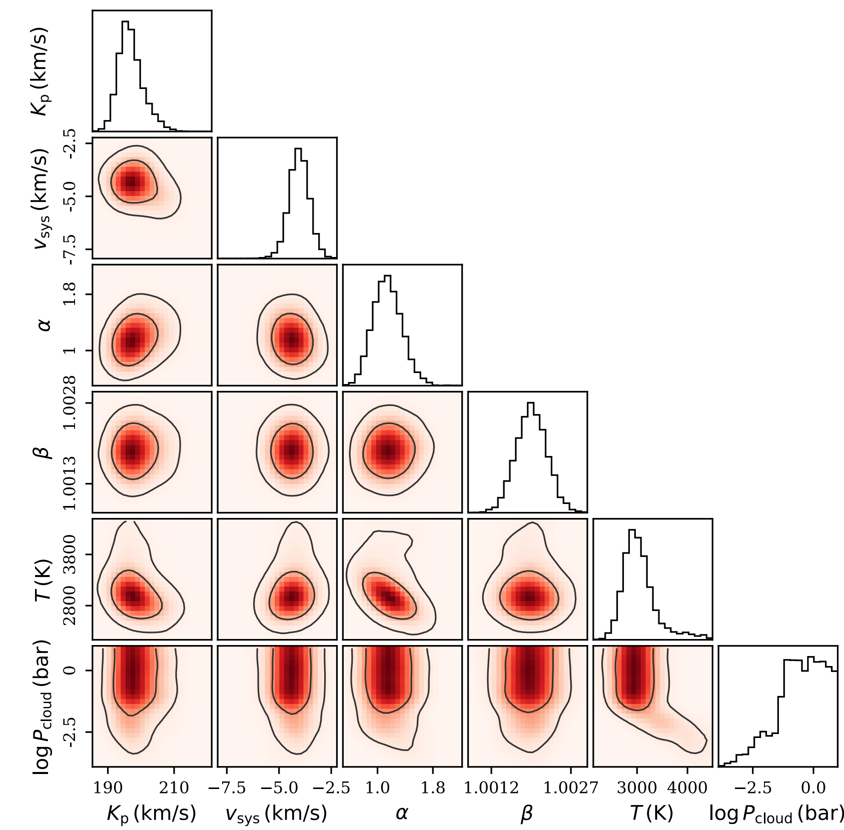

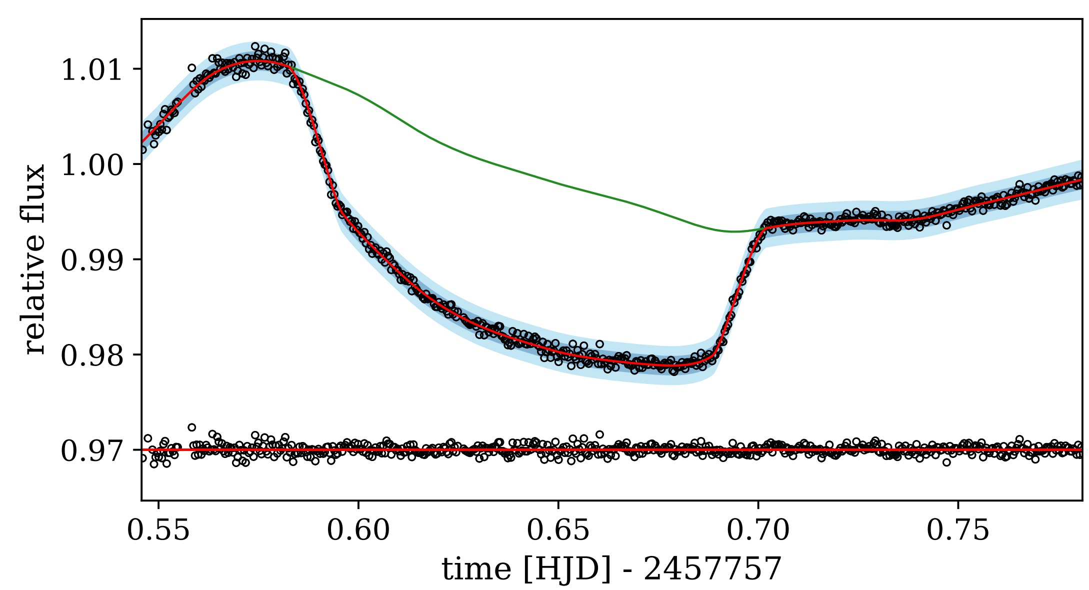

Talk Outline

- A fully Bayesian approach to high-resolution retrievals
- How can we deal with noise?
- What extra information can we get?
- Demonstration on UVES data of WASP-121b

Also see talk by Brogi

Overview of the Cross-Correlation Technique

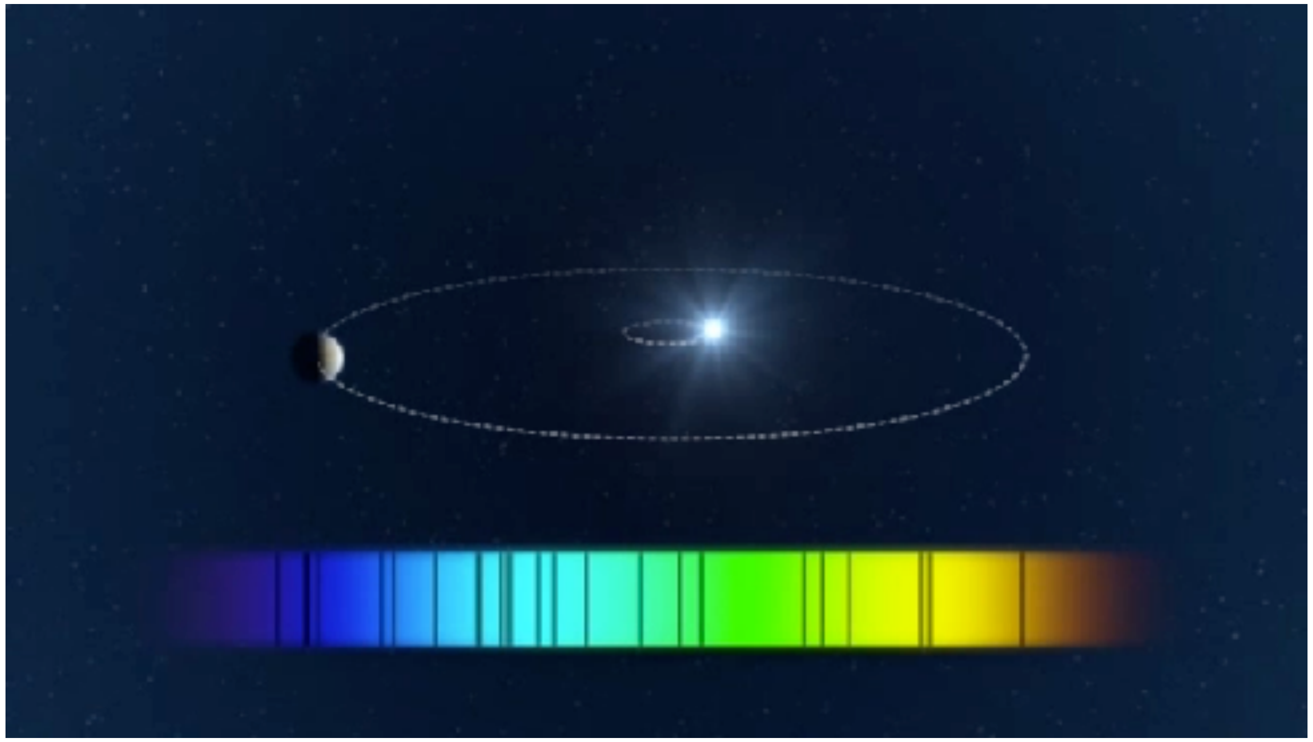

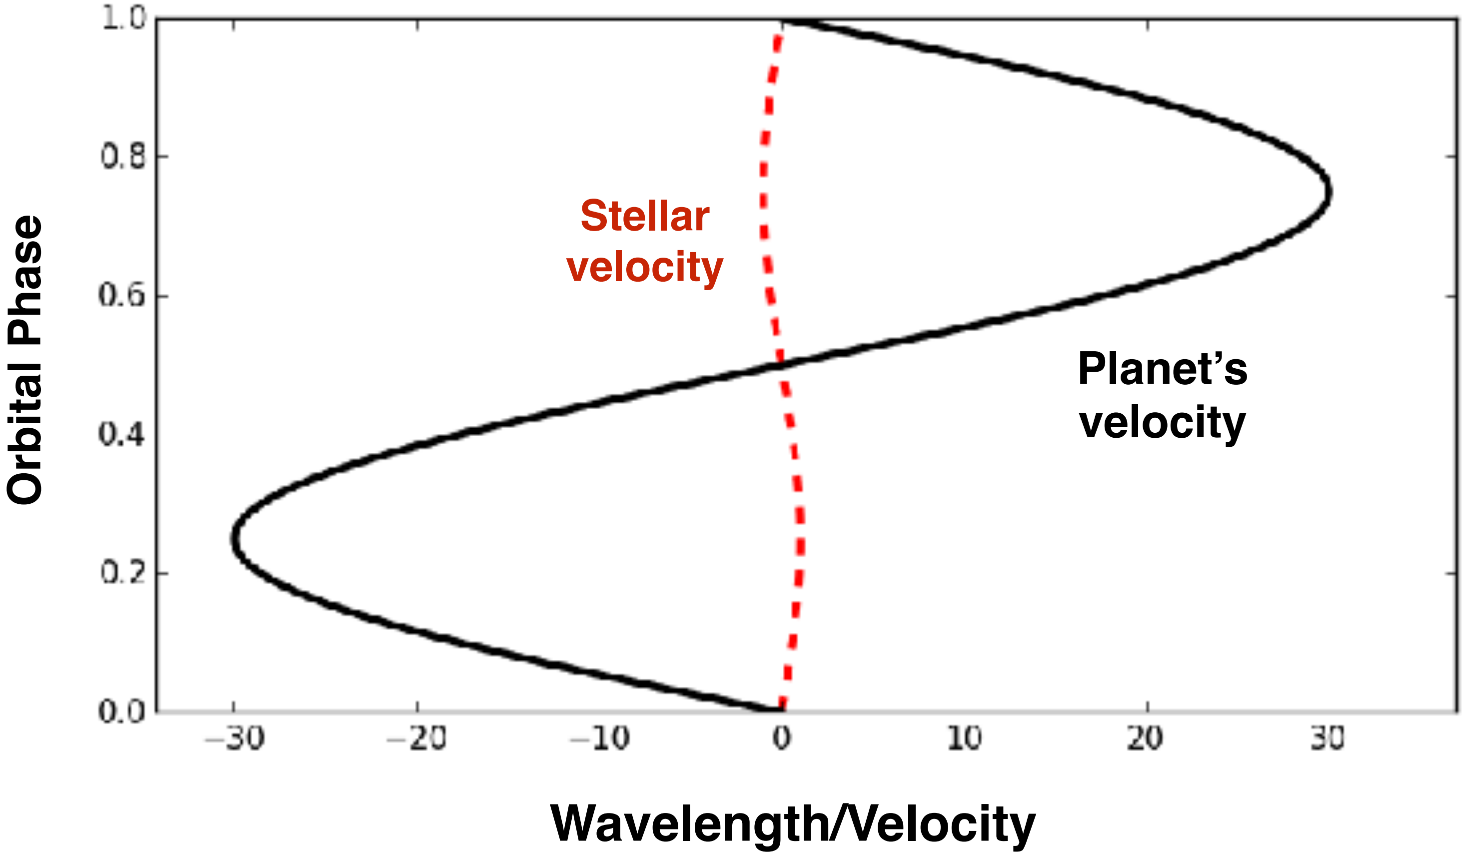

Neale Gibson, Trinity College Dublin

Doppler-resolved Transmission Spectroscopy

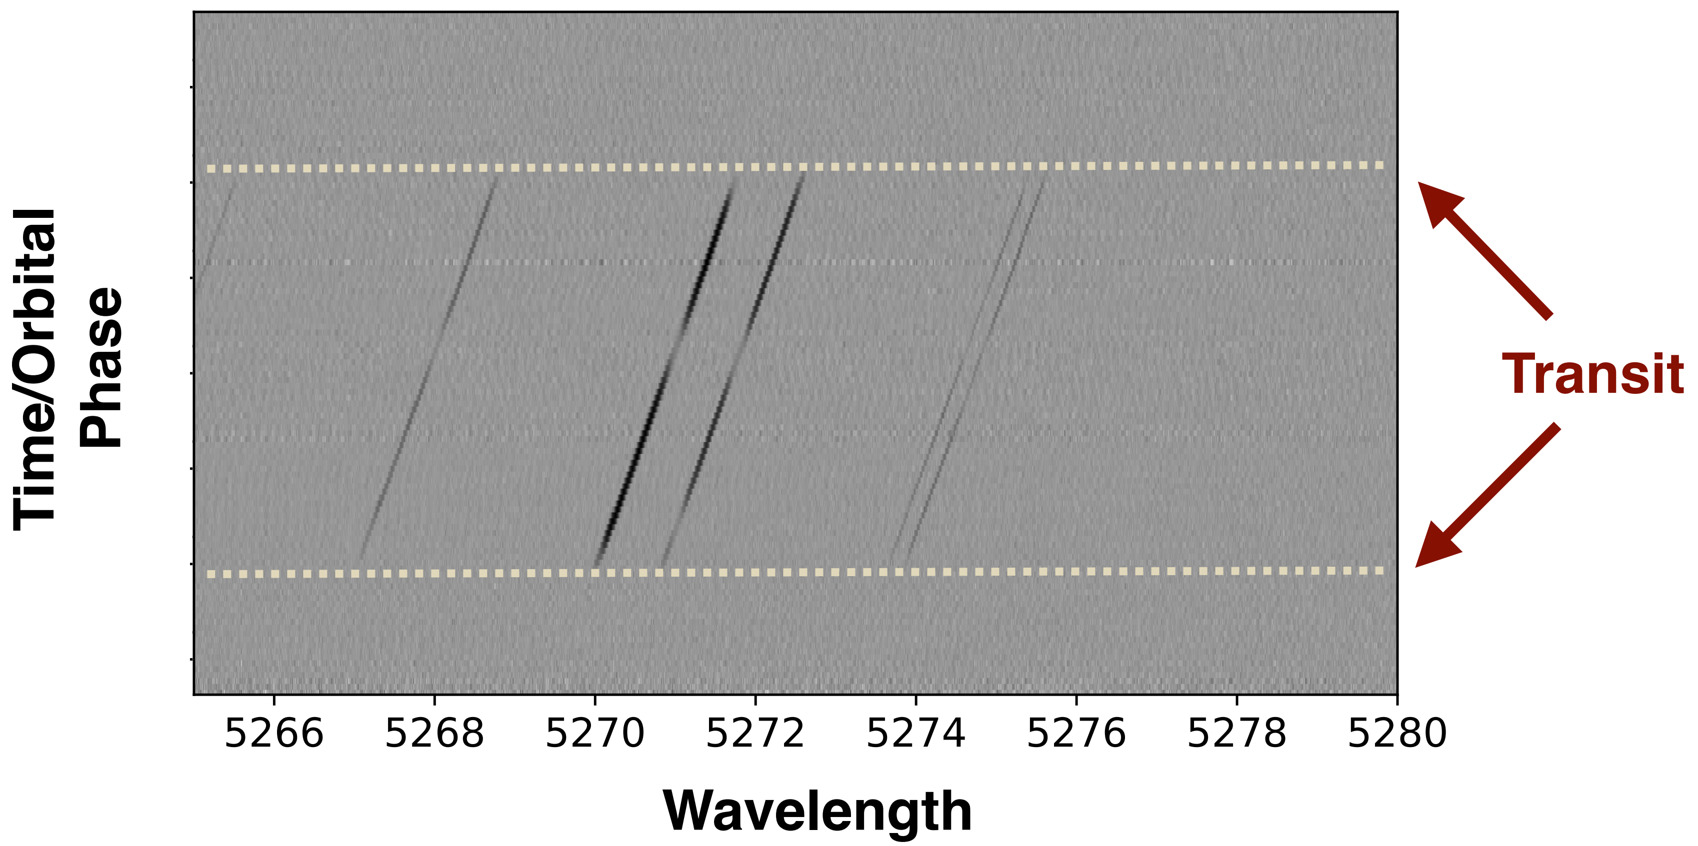

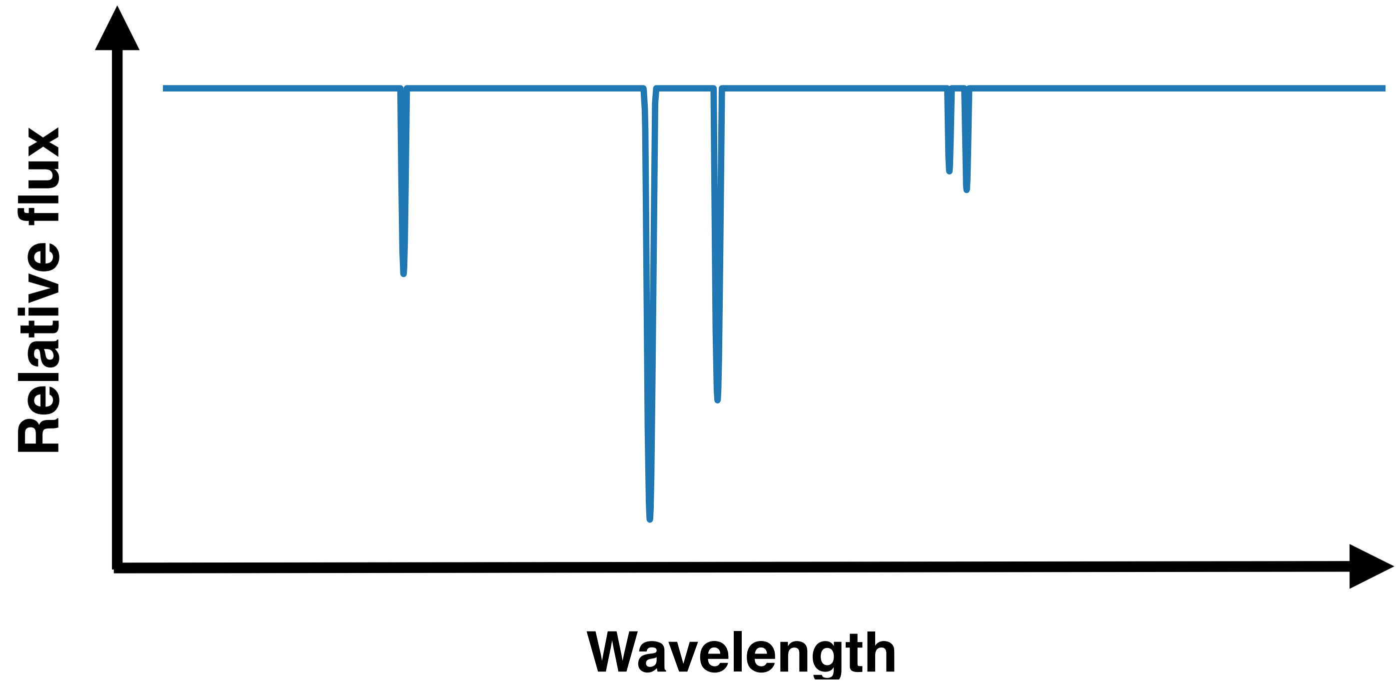

- Generally speaking we don't see individual lines (although see talk by J. Seidel) - we need to also sum over the spectral lines

Step 1: Pre-processing the data

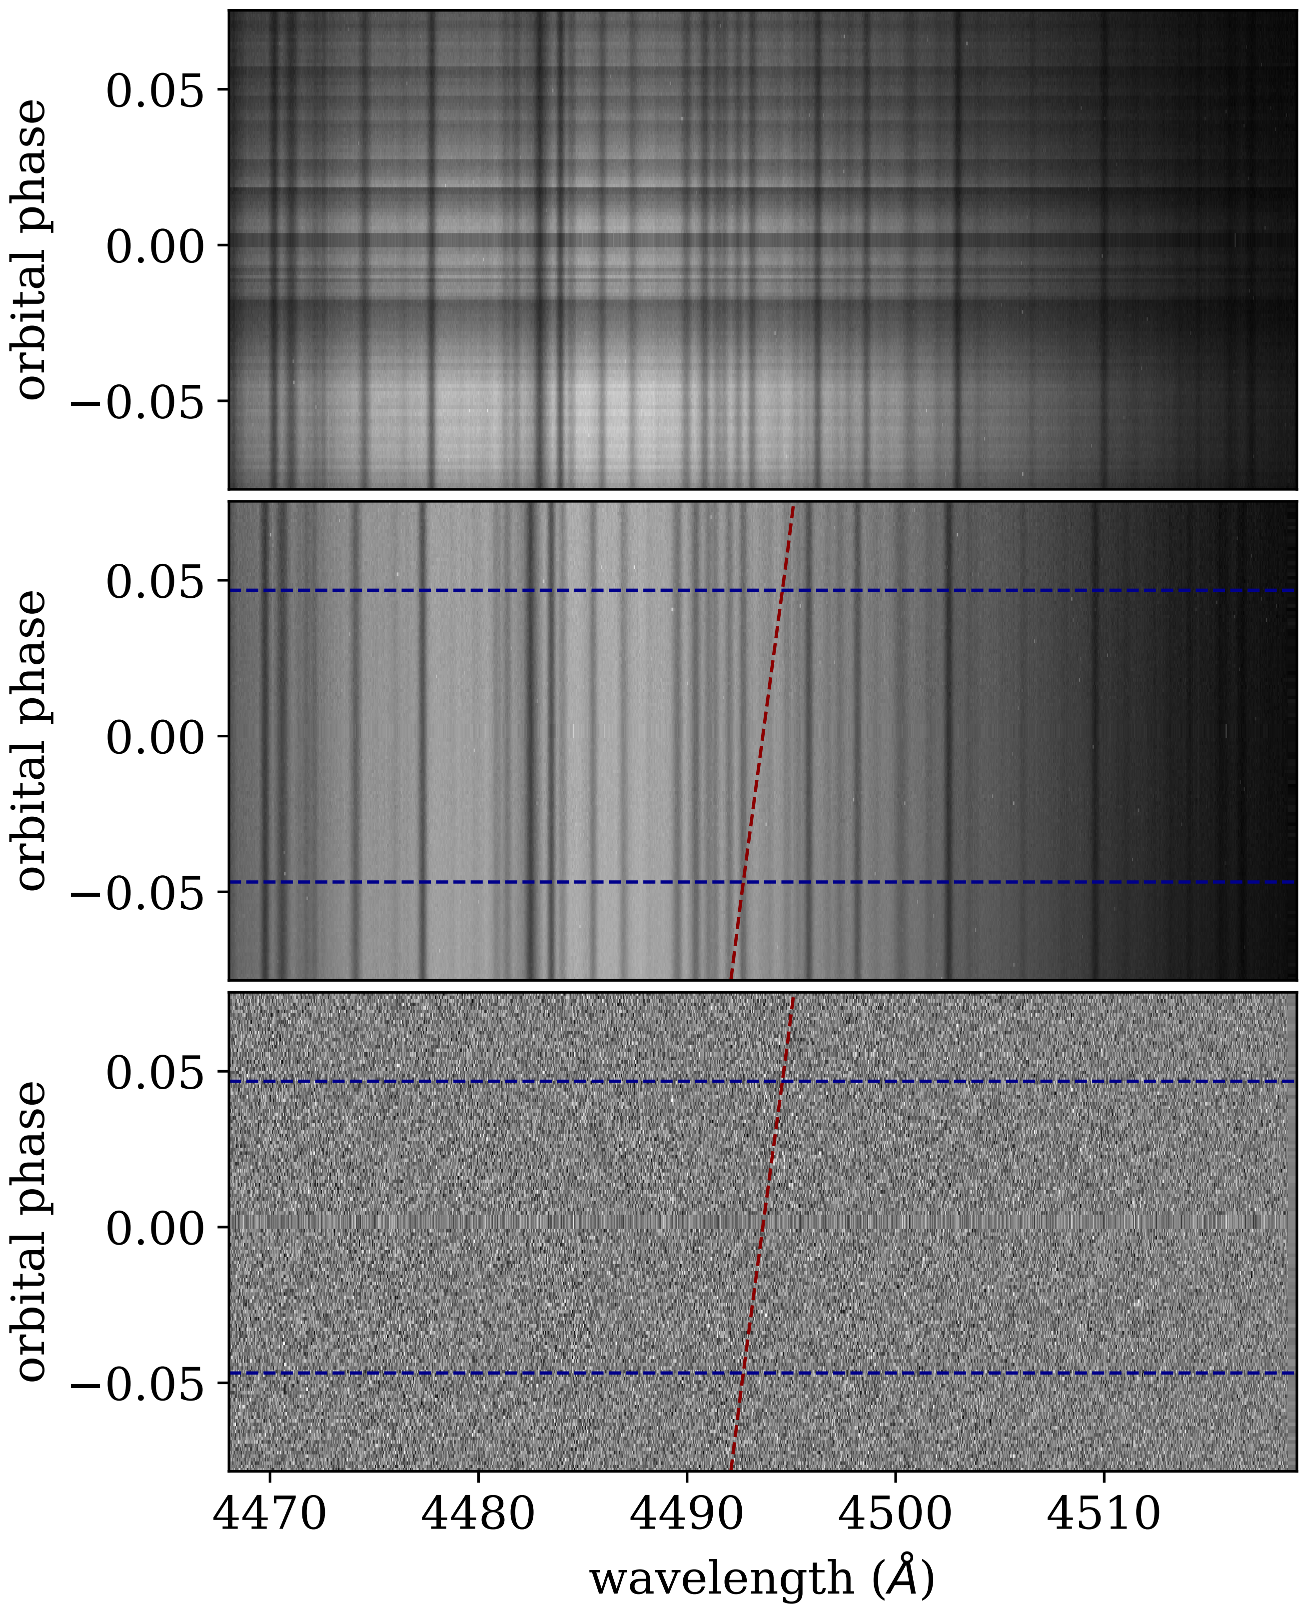

Goal is to normalise spectra and remove stellar and telluric lines

1) Align spectra and normalise

2) Divide through by average spectrum

3) Apply a de-trending algorithm
(e.g. SysRem or PCA)

This process also modifies the underlying exoplanet signal (see Matteo's talk!)

Step 2: Cross-Correlation with atmospheric template

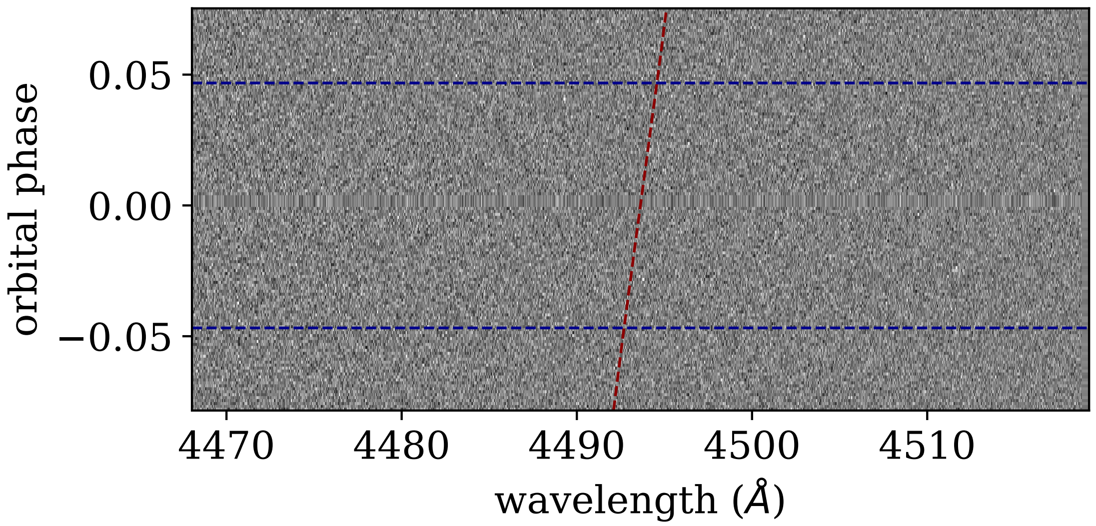

Cross-correlate data with atmospheric template

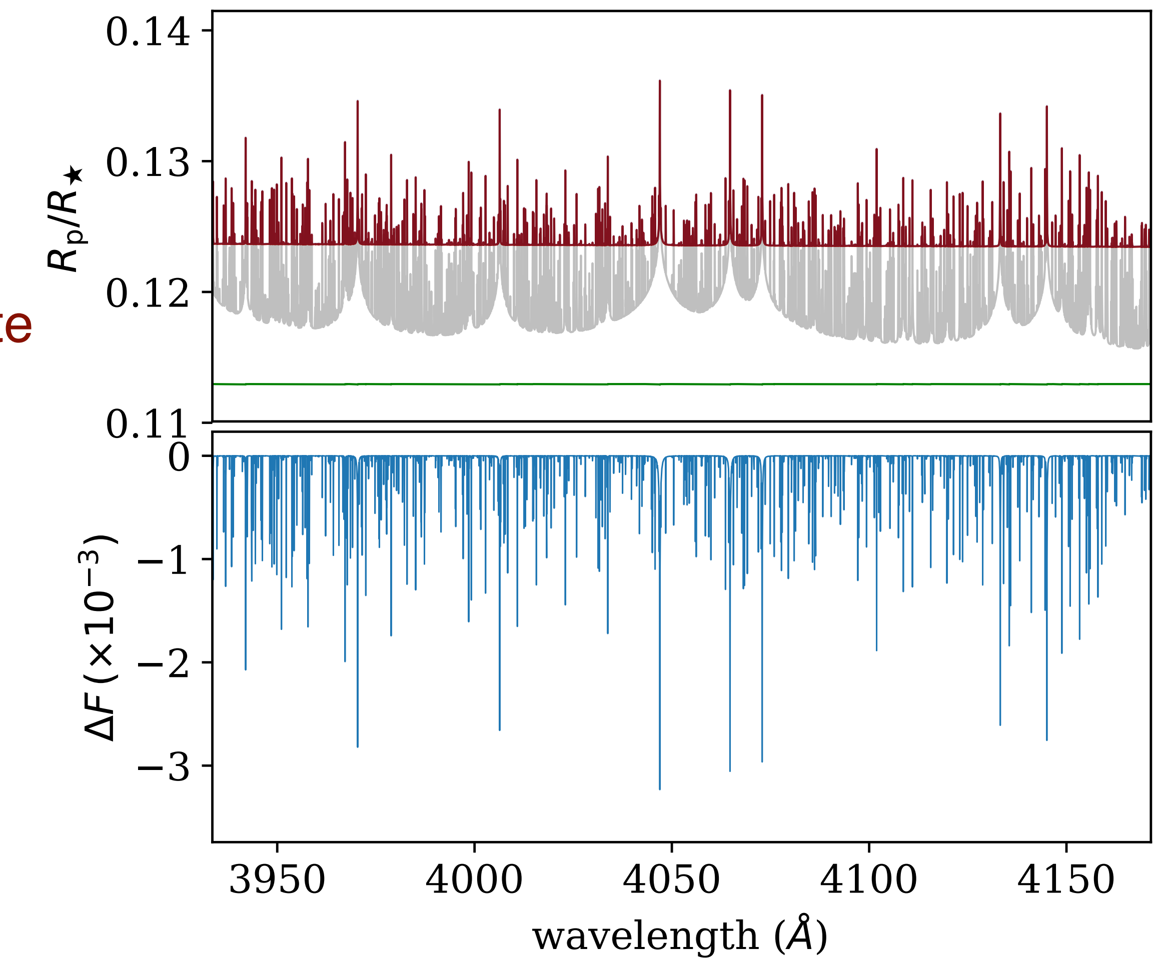

CCF: time vs systematic velocity

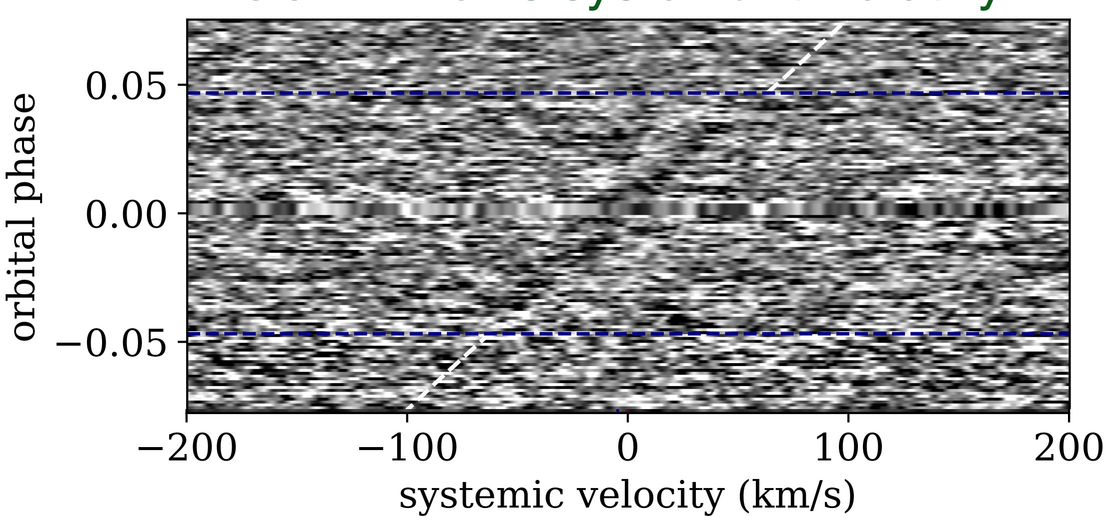

1) Generate model atmosphere

2) Multiple template by data and integrate

3) Do for a range of velocity 'lags' ($v_{\rm sys}$)

Optimal way to compute CCF:

$$\mathrm{CCF}(v_{\rm sys}) = \sum_i \frac{f_i\, m_i(v_{\rm sys})}{\sigma_i^2}$$

Step 3: Summation over time/velocity

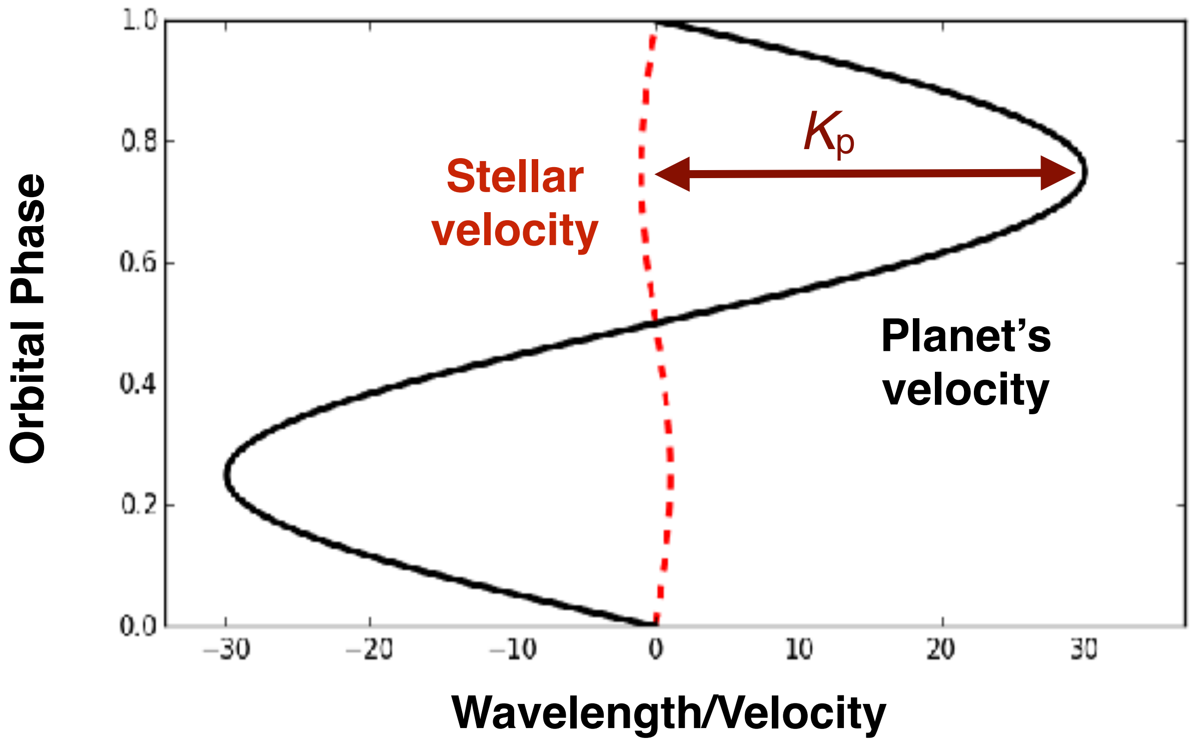

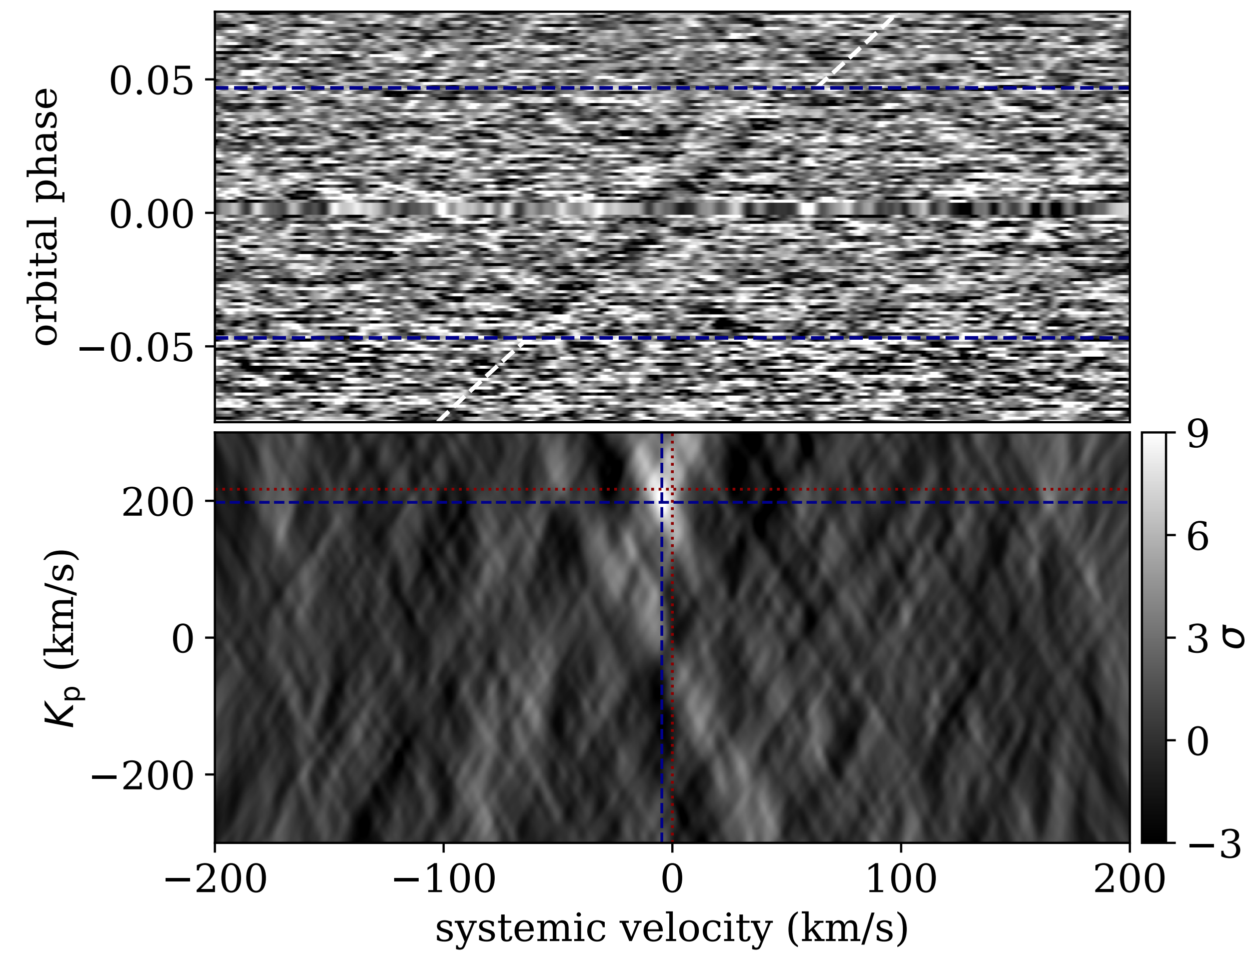

1) Weight signal by transit model

2) Sum CCF over (Keplerian) velocity of the planet

3) Repeat for range of values of K_p to get K_p-v_{sys} 'map'

(This accounts for unknown K_p and also to evaluate systematics)

Example of UVES/WASP-121b

- Transit observation of WASP-121b: an “ultra-hot” Jupiter (>2,500K)

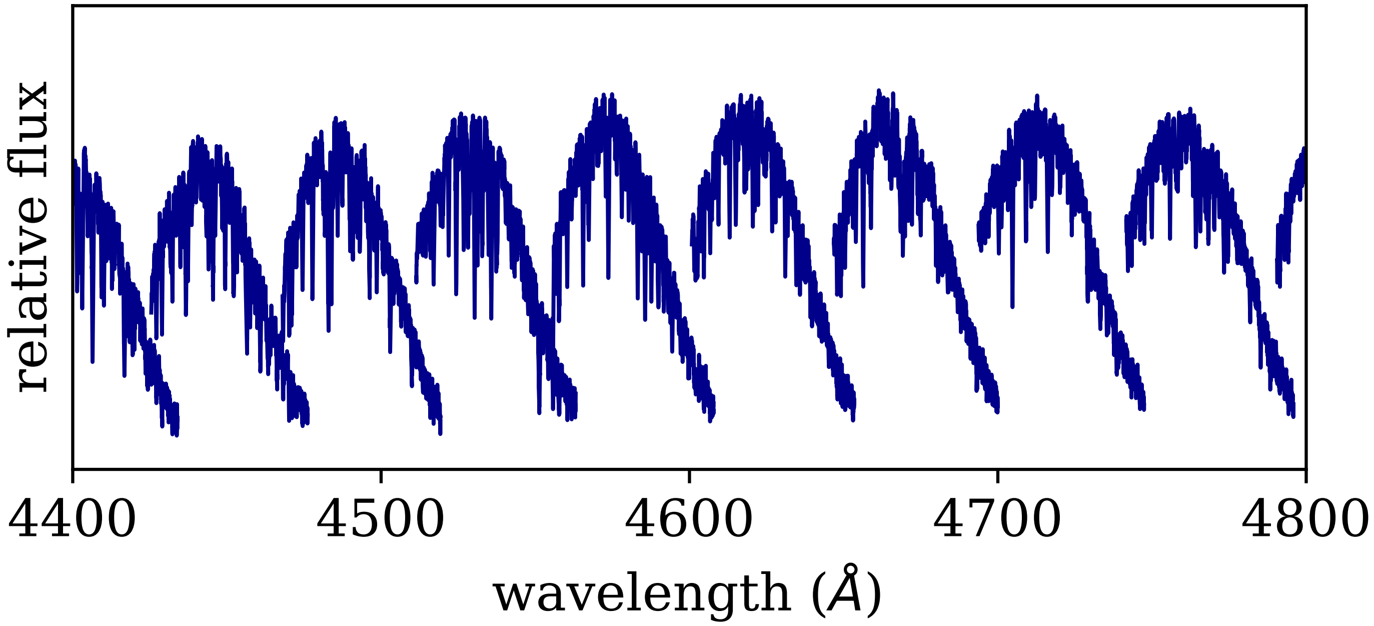

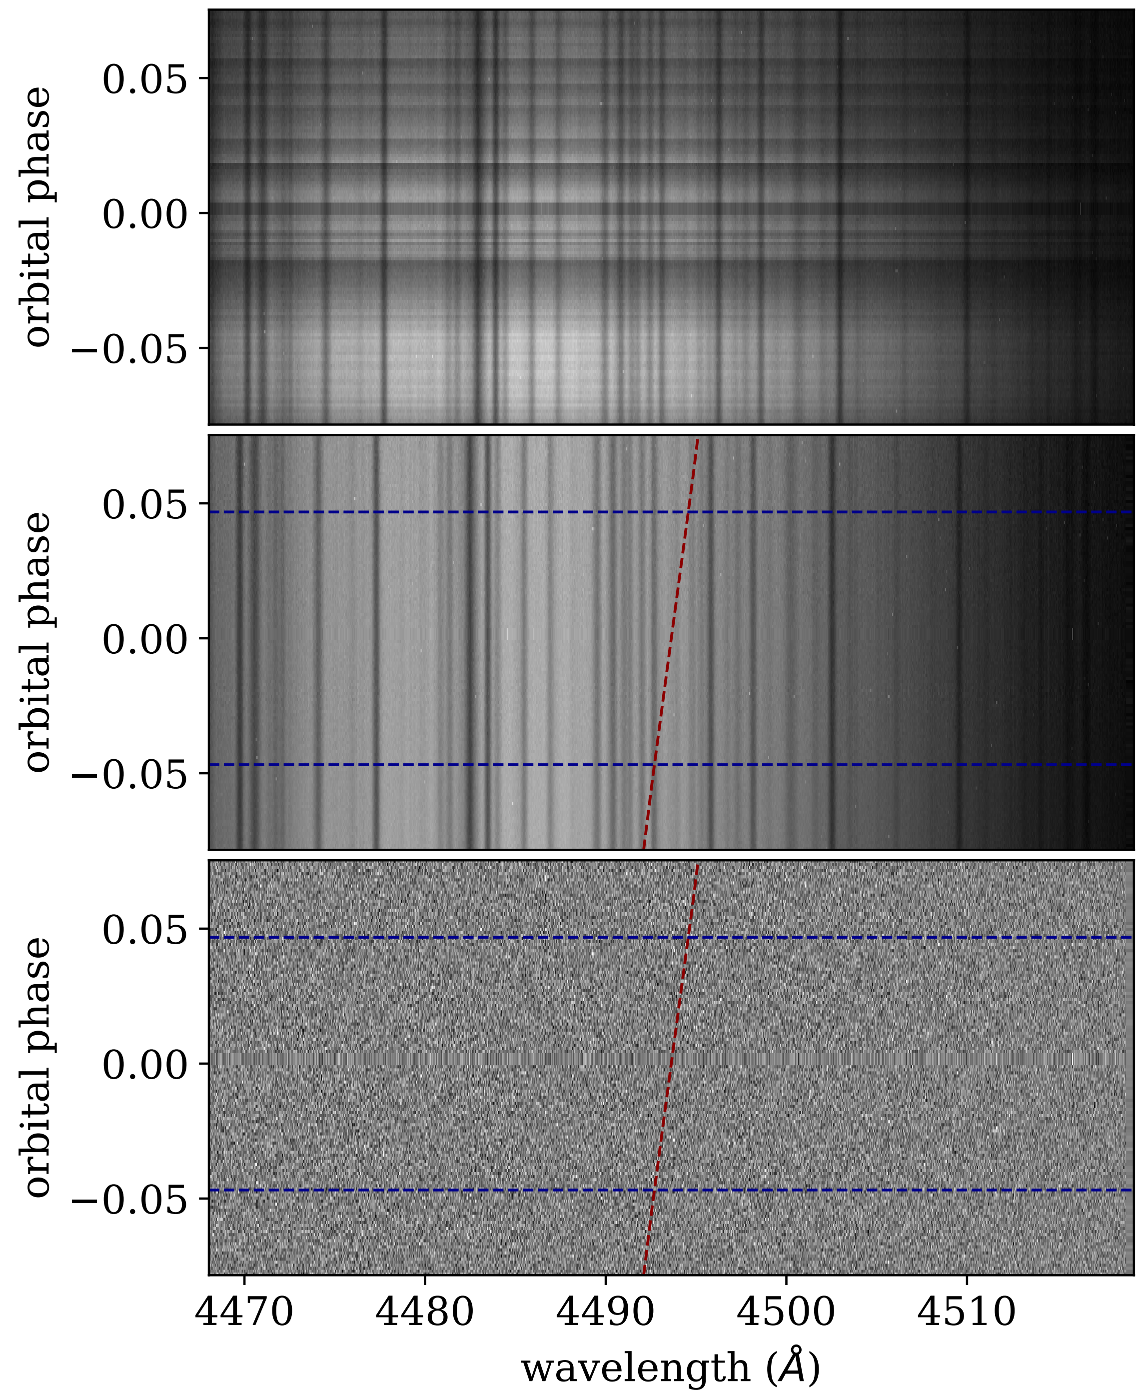

- R~100,000; $\lambda = 3{,}700 - 8{,}800\,\text{Å}$
- 137 spectra with 85 spectral orders

  = >35 Million Data points!
- Work with the extracted orders, not merged spectra

Example of UVES/WASP-121b

- WASP-121b: T_{eq}>2,500K

@ low-resolution: H_2O, TiO?, VO? + inversion

- Low-resolution spectra showed excess absorption from unknown species

Low-resolution spectra from HST (Evans et al. 2018)

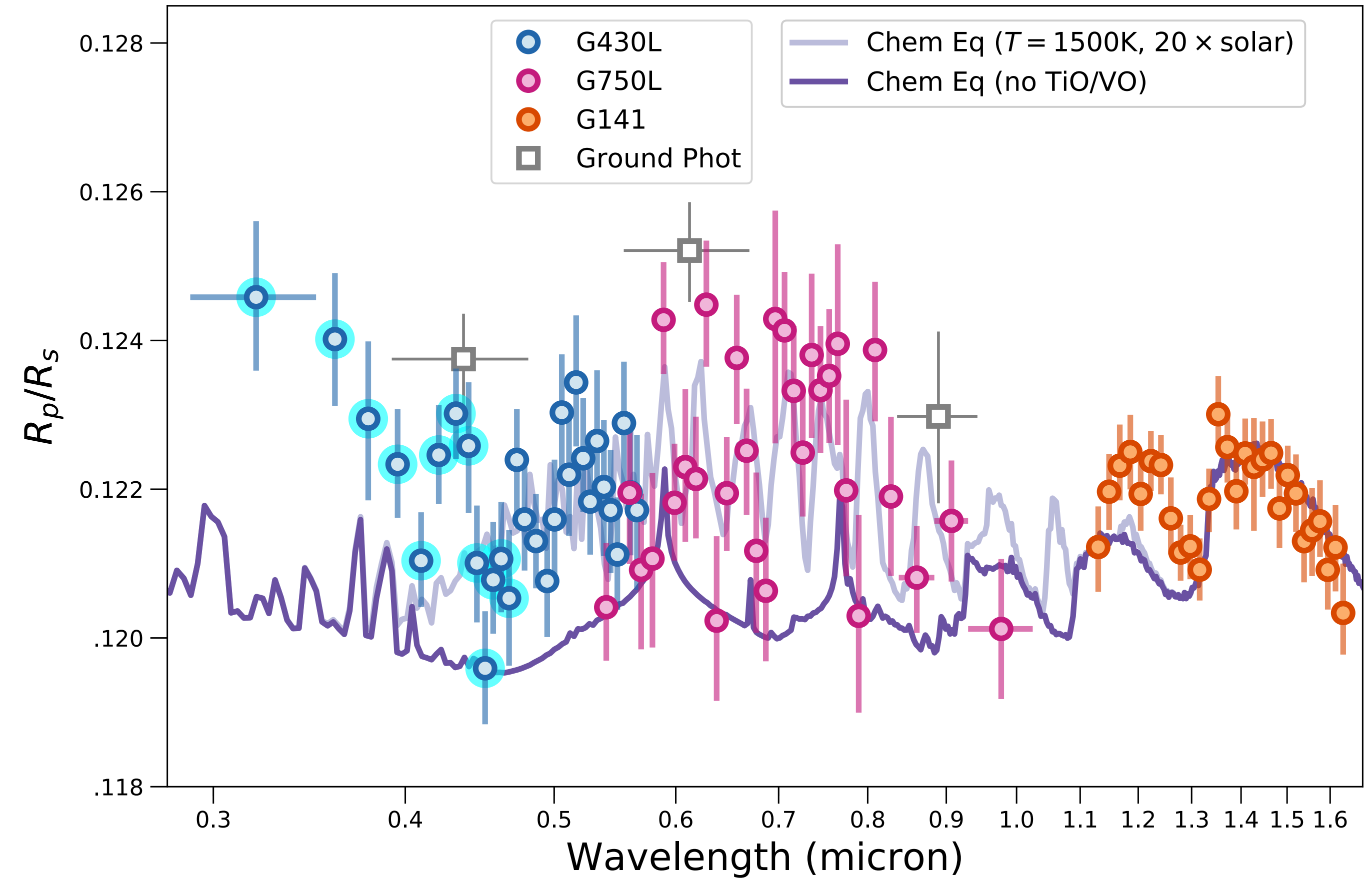

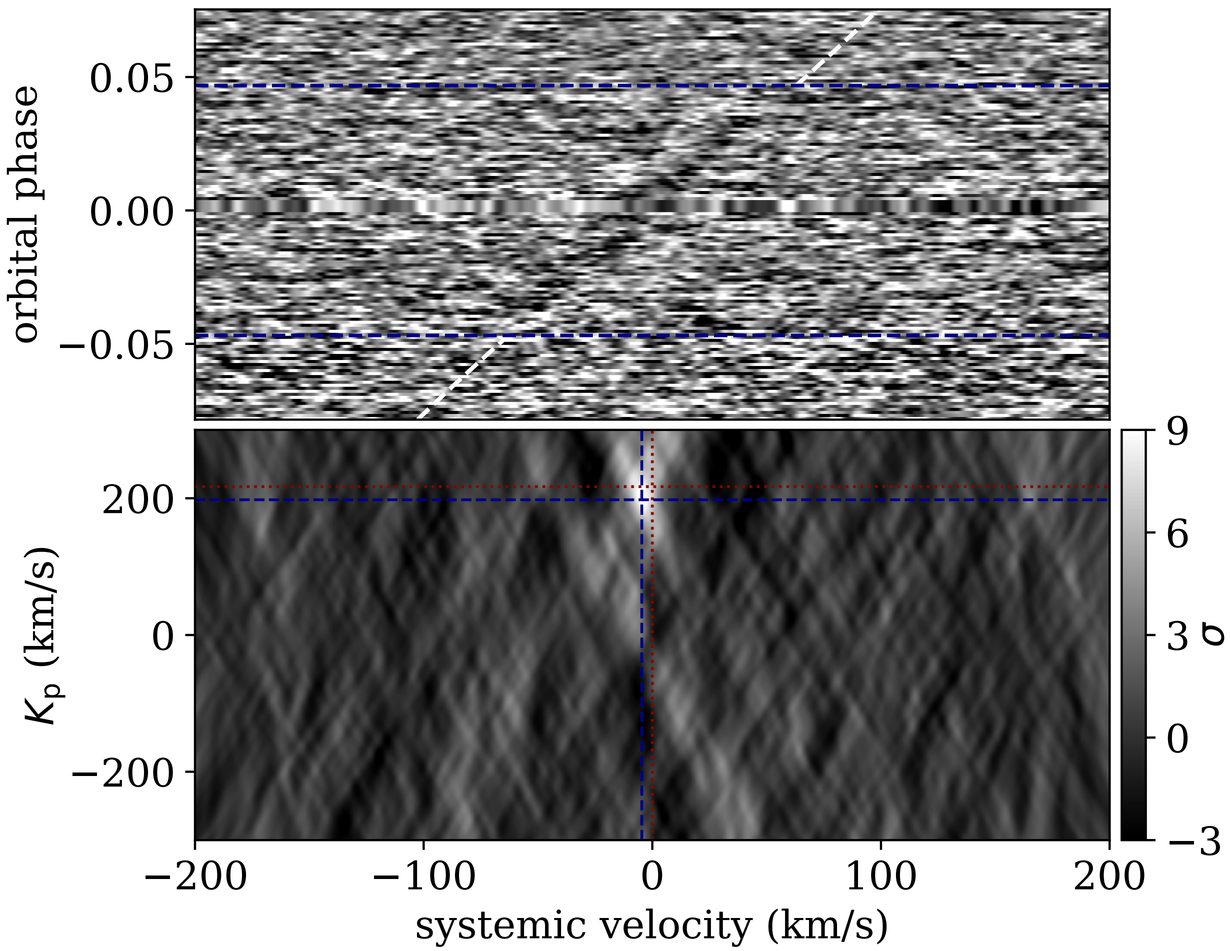

Fe signal in WASP-121b using UVES (Gibson et al. 2020)

- *@ high-res: Clearly see the signal separated from the star in velocity-space*

- Gives an *unambiguous* detection: absorber is probably (mostly) Fe!

High-res shows no sign of TiO, VO, + many other metals detected. See Merritt et al. (2020; 2021), Bourrier et al. (2020); Hoeijmakers et al. (2020); and more…

A spectral inventory of species in exoplanet atmospheres

Species detected from transmission spectroscopy:
Na, K, Fe, Li , Mg, Mn, O, C, Ca, H, He, Sc, Si, V, Ti…
H_2O, CO, CO_2, C_2H_2, CH_4, HCN, TiO, AlO, VO
+ ions (of Fe, Mg, Ca, Ti,…)
(updated from Madhusudhan 2019 review; probably out of date)

- Recent explosion in detections largely driven by metals in ultra-hot Jupiters at optical wavelengths (e.g. Hoeijmakers et al. 2018; Lothringer et al. 2018)

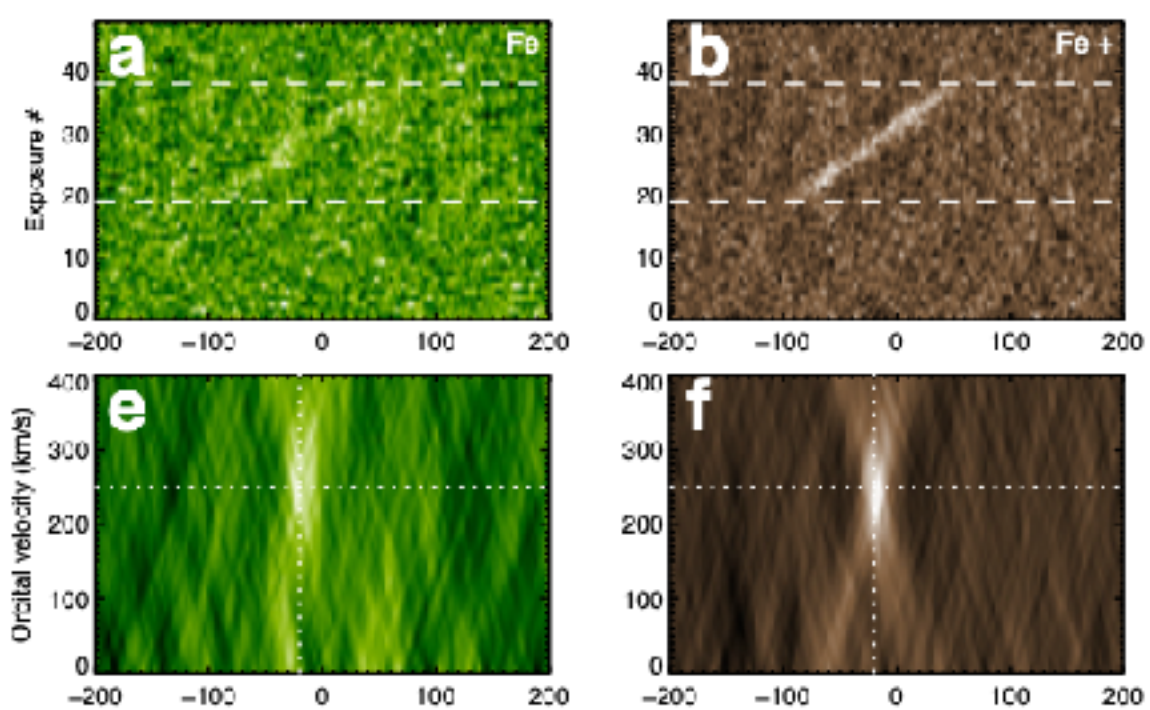

(see also Matteo's results on molecular species in the IR)

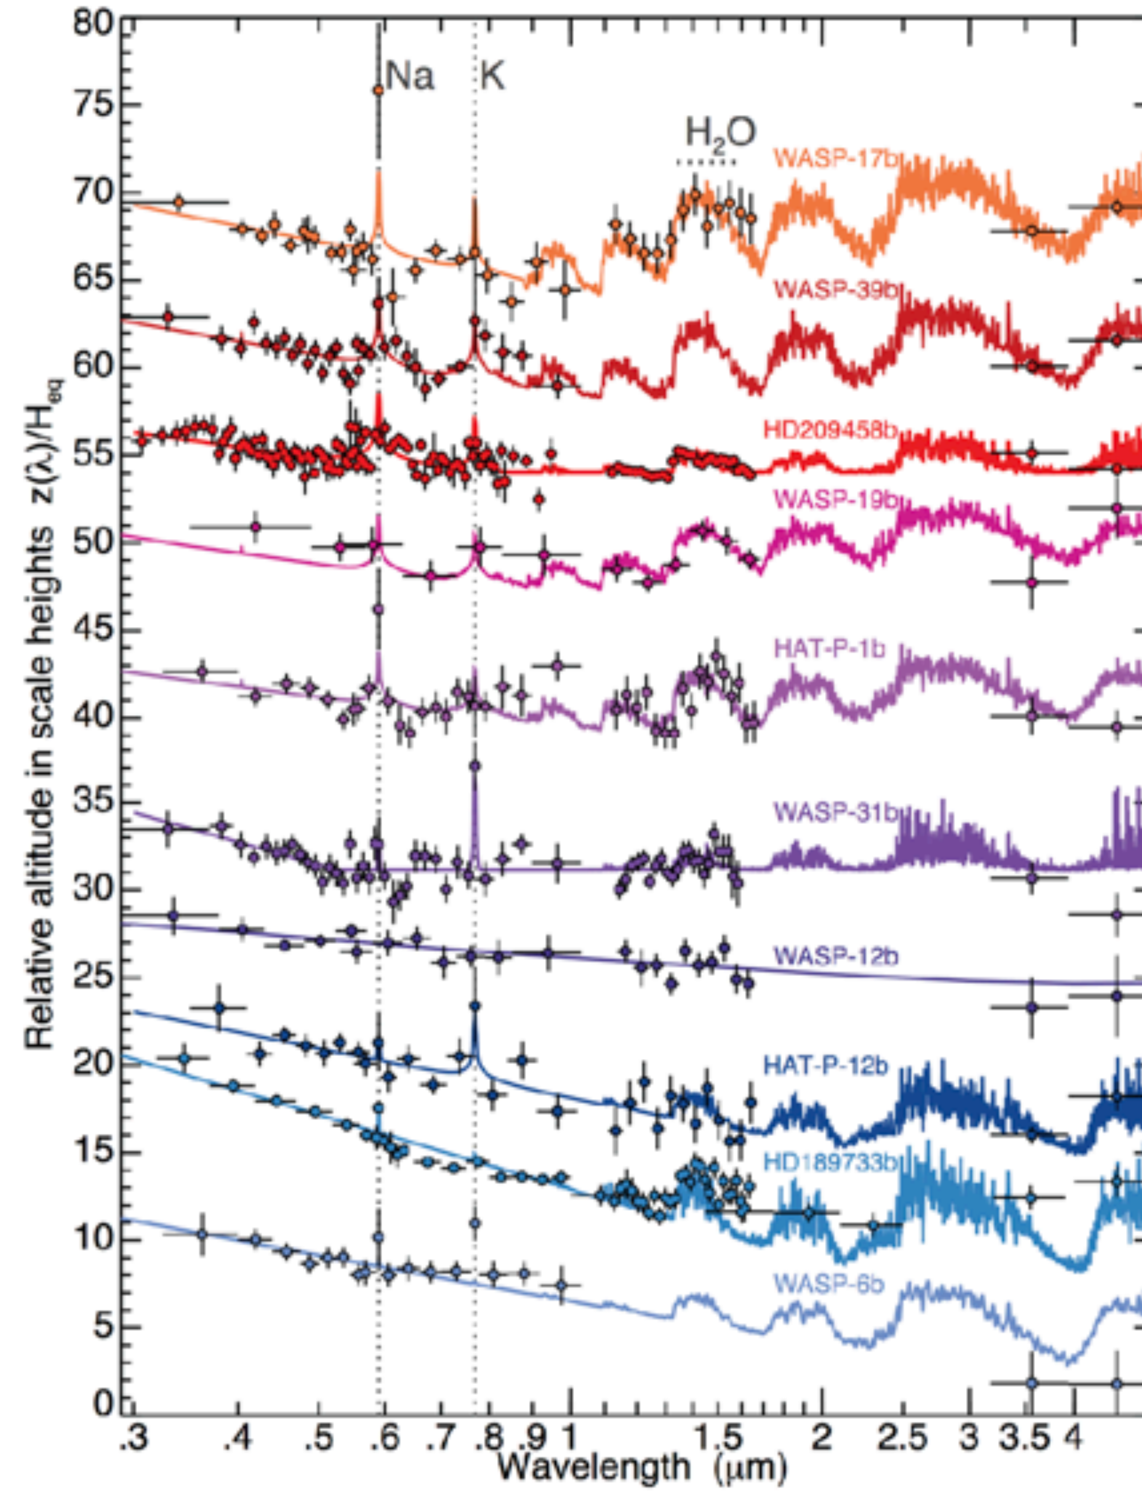

Sing et al. (2015)

Why do we need a likelihood approach?

Cross-correlation does not allow us to compare different atmospheric models

- To obtain quantitative constraints, we need to fit models to the data
  - How much Fe in the atmosphere?
  - What is the temperature? etc…
- Ideally would perform retrieval similar to low-resolution datasets

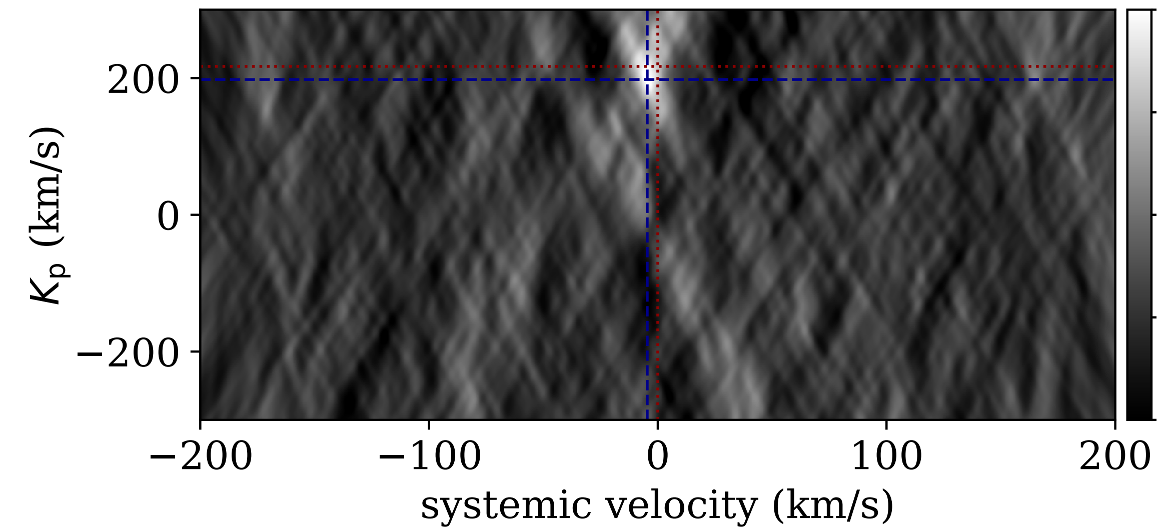

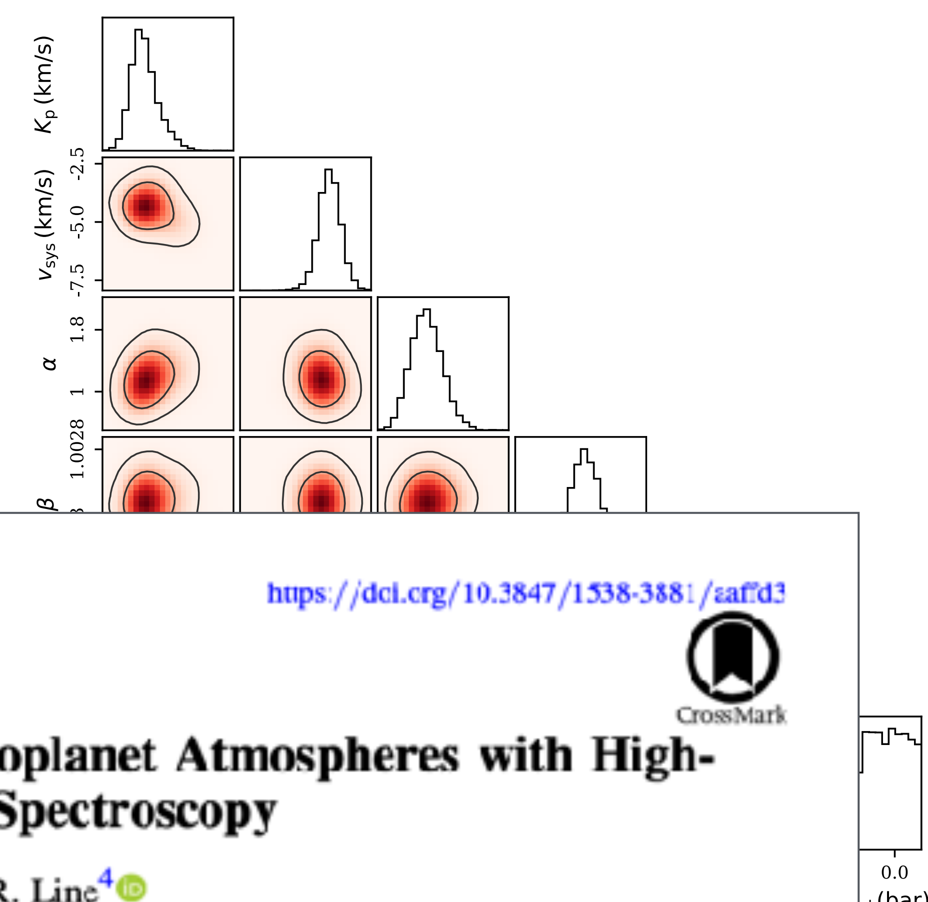

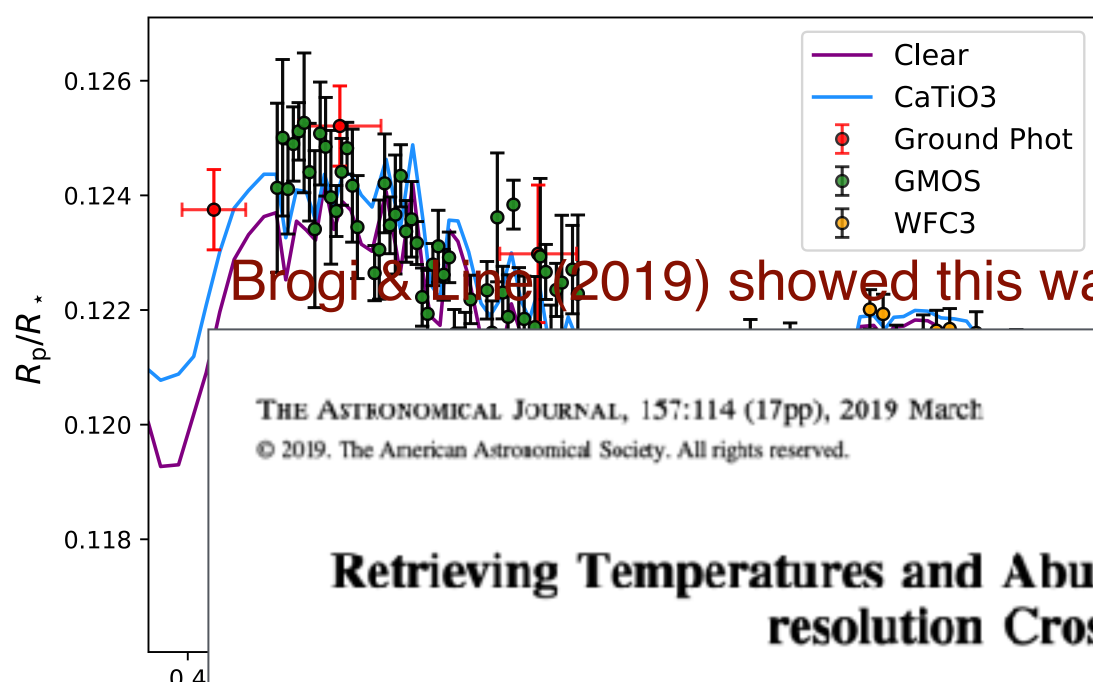

Brogi & Line (2019) showed this was possible!

THE ASTRONOMICAL JOURNAL, 157:114 (17pp), 2019 March
© 2019. The American Astronomical Society. All rights reserved.
https://doi.org/10.3847/1538-3881/aaffd3

CrossMark

Retrieving Temperatures and Abundances of Exoplanet Atmospheres with High-resolution Cross-correlation Spectroscopy

Matteo Brogi[1,2,3] and Michael R. Line[4]
[1] Department of Physics, University of Warwick, Coventry CV4 7AL, UK; m.brogi@warwick.ac.uk
[2] INAF—Osservatorio Astrofisico di Torino, Via Osservatorio 20, I-10025, Pino Torinese, Italy
[3] Centre for Exoplanets and Habitability, University of Warwick, Gibbet Hill Road, Coventry CV4 7AL, UK
[4] School of Earth and Space Exploration, Arizona State University, Tempe, AZ 85287, USA
Received 2018 November 5; revised 2019 January 4; accepted 2019 January 16; published 2019 February 13

Can we use the same methods as transit light curve fitting?

- Similar to how we analyse transit light curves - can we infer noise properties from the data?

  (we have many years experience of this from dealing with systematics)

- 'Standard' likelihood is the product of independent Gaussians:

$$\mathcal{L}(\theta) = \prod_i \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(-\frac{1}{2}\frac{(f_i - m_i)^2}{\sigma_i^2}\right)$$

$$\mathcal{L}(\theta) \propto \exp\left(-\frac{1}{2}\chi^2\right)$$

- Should we fix noise, or fit for it? Can use both! (e.g. Gibson et al. 2013):

$$\mathcal{L}(\theta) = \prod_i \frac{1}{\sqrt{2\pi(\beta\sigma_i)^2}} \exp\left(-\frac{1}{2}\frac{(f_i - \alpha m_i)^2}{(\beta\sigma_i)^2}\right)$$

- Gibson et al. (2013)

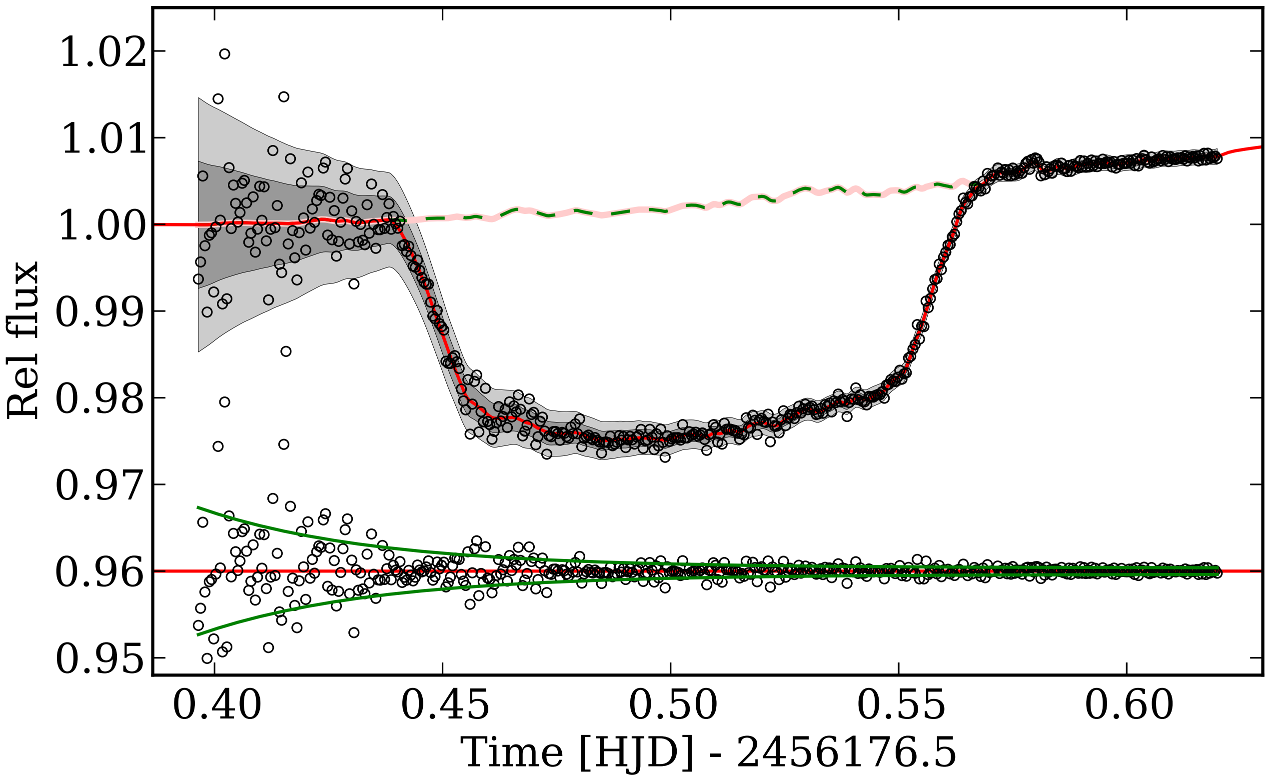

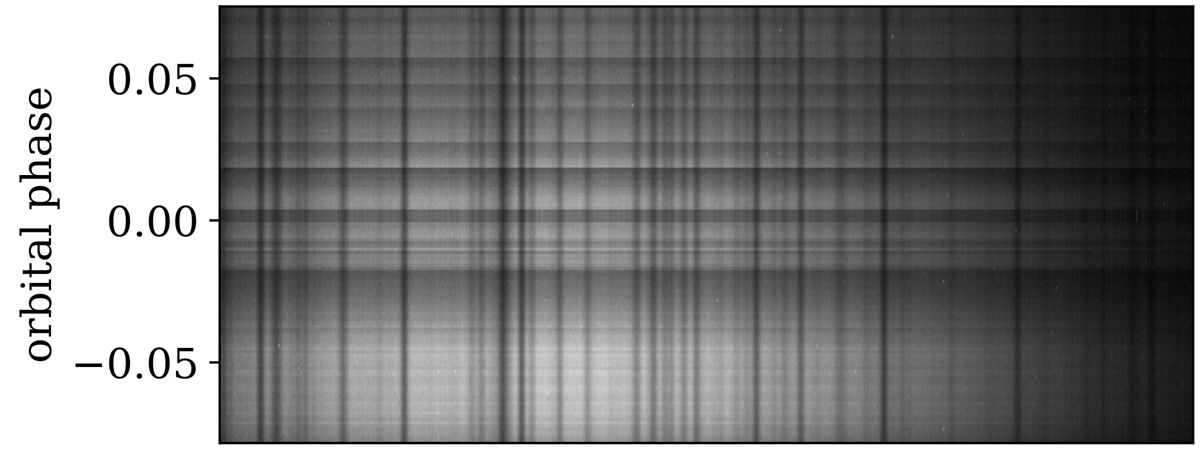

From Cross-Correlation ‘Maps’ to Likelihood ‘Maps’

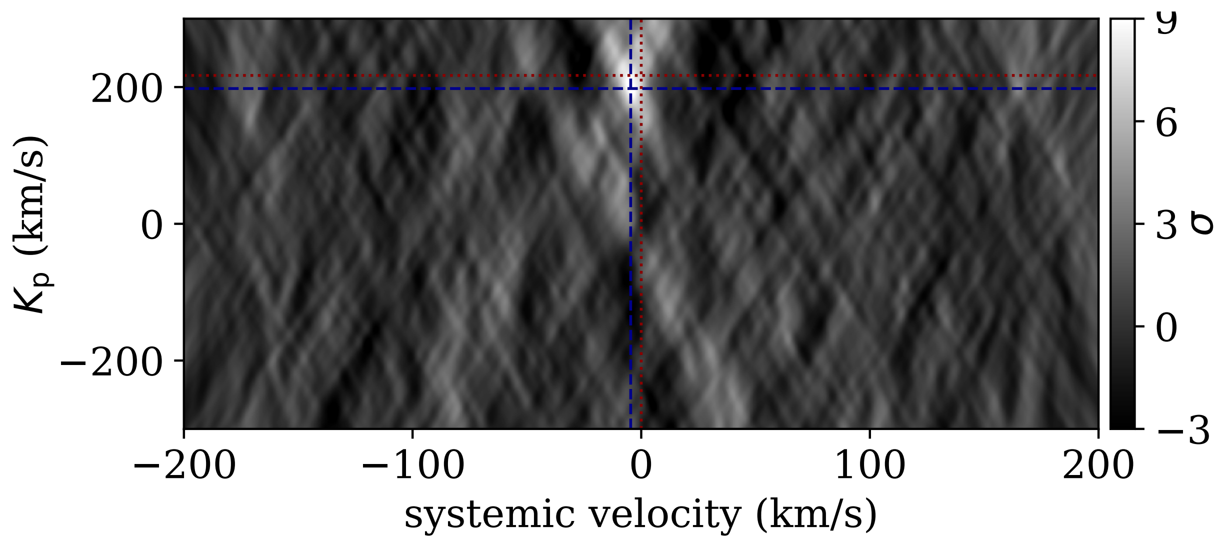

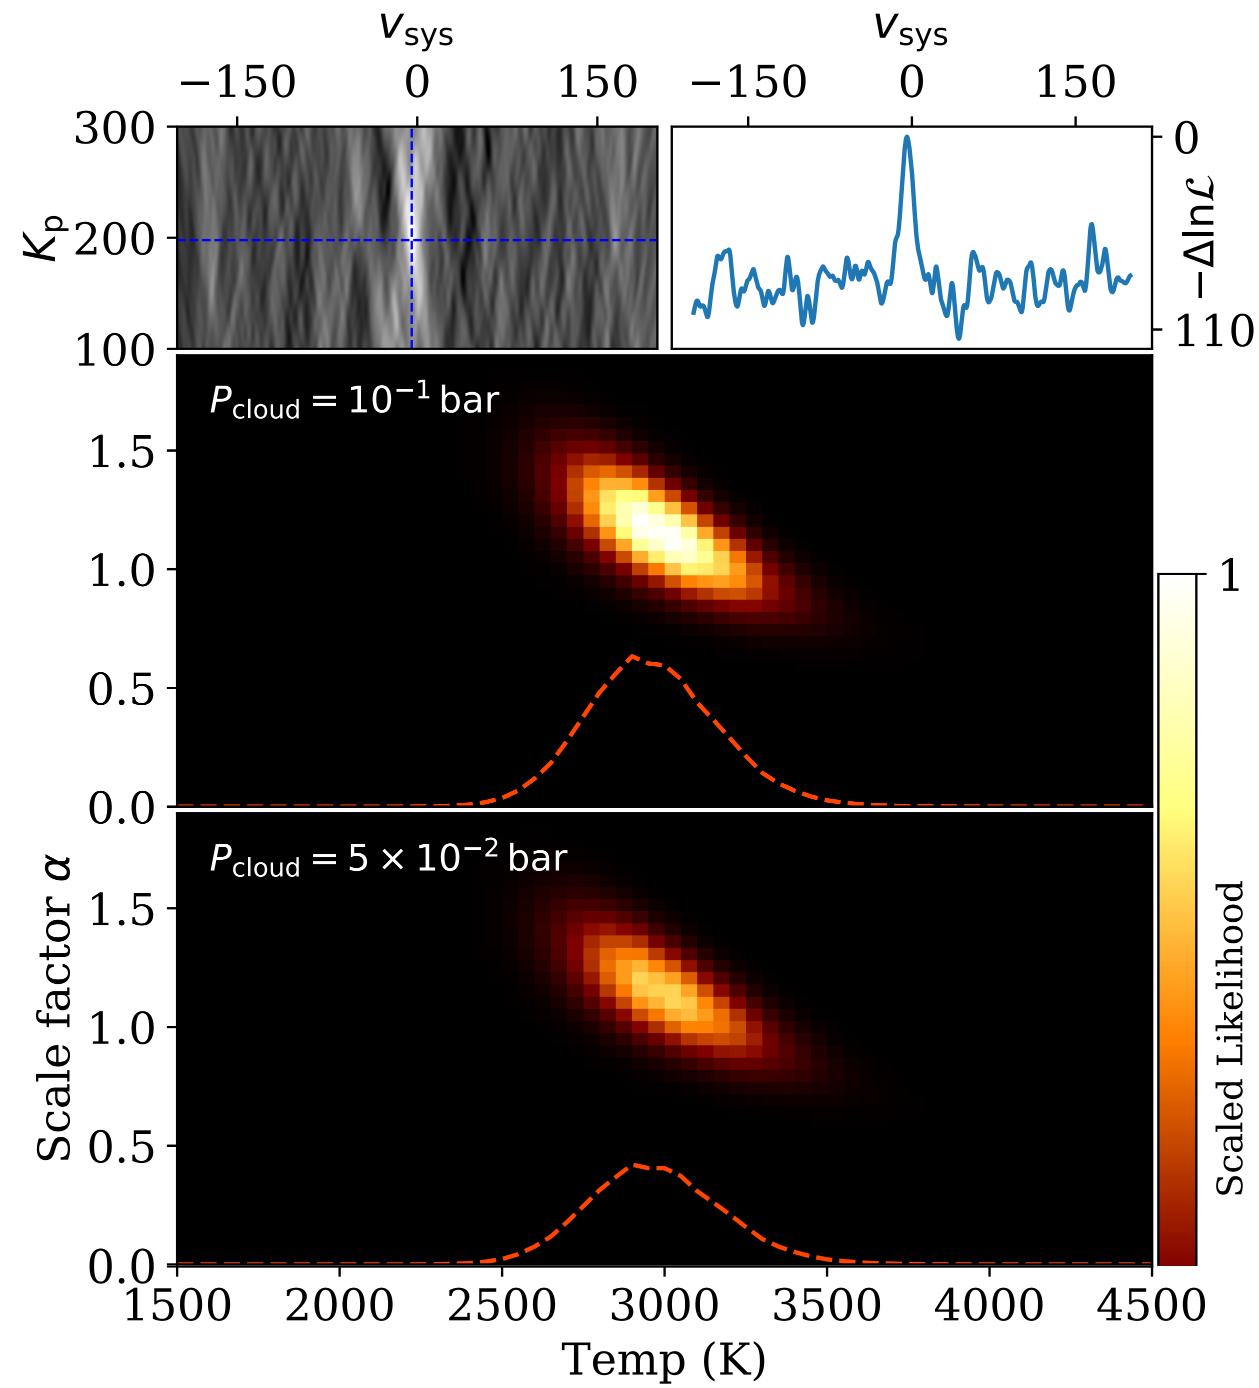

Use log L in practice:

$$\log \mathcal{L}(\theta) = -\frac{1}{2}\chi^2 - N \log \beta + \mathrm{Cnst}$$

$$\chi^2 = \sum_i^N \left(\frac{f_i - \alpha m_i}{\beta \sigma_i} \right)^2$$

$$\chi^2 = \frac{1}{\beta^2} \left[\sum_i \frac{f_i^2}{\sigma_i^2} + \alpha^2 \sum_i \frac{m_i^2}{\sigma_i^2} - 2\alpha \sum_i \frac{f_i\, m_i}{\sigma_i^2} \right]$$

Cnst

Model dependent

$$\mathrm{CCF}\,(v_\mathrm{sys}) = \sum_i \frac{f_i m_i(v_\mathrm{sys})}{\sigma_i^2}$$

e.g. Lockwood et al. (2014); Brogi & Line (2019)

How do we switch to full Bayesian retrievals?

Cannot use a ‘grid search’ for complex modelling

Now we forget about cross-correlation entirely!

1) Produce a ‘2D’ forward model as a function of K_p + v_{sys} (for each order)

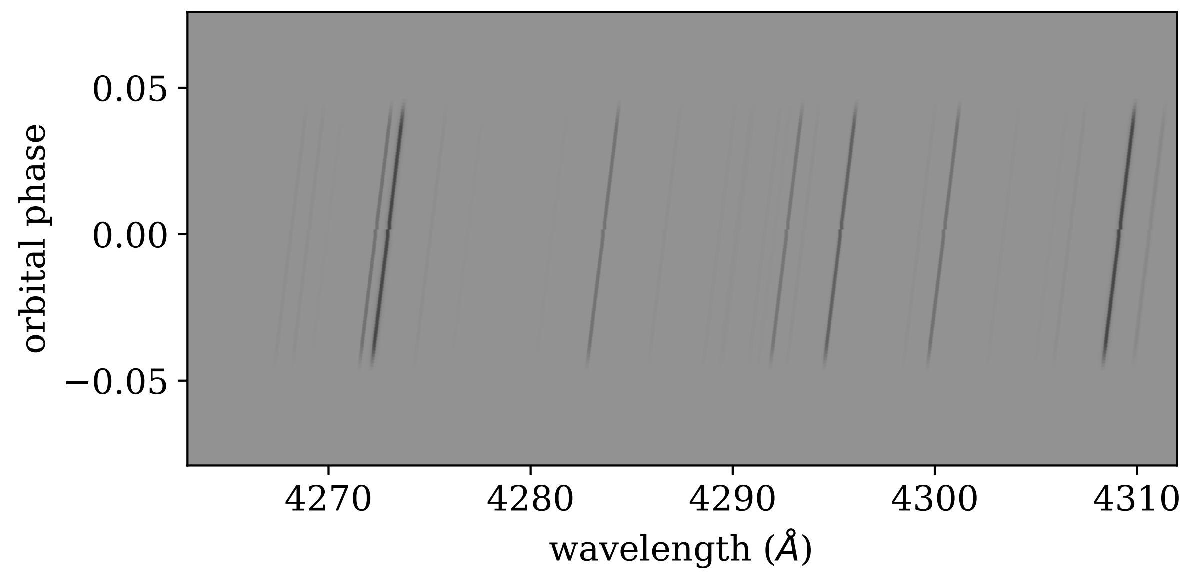

2) Use standard likelihood:

$$\log \mathcal{L}(\theta) = -\frac{1}{2}\chi^2 - N \log \beta$$

3) Define prior + posterior:

$$\log\ p(\theta|\mathcal{D}) = \log\ \mathcal{L}(\mathcal{D}|\theta) + \log\ \pi(\theta)$$

4) Feed posterior to MCMC (or favourite algorithm)

5) Understand your atmosphere (!?)

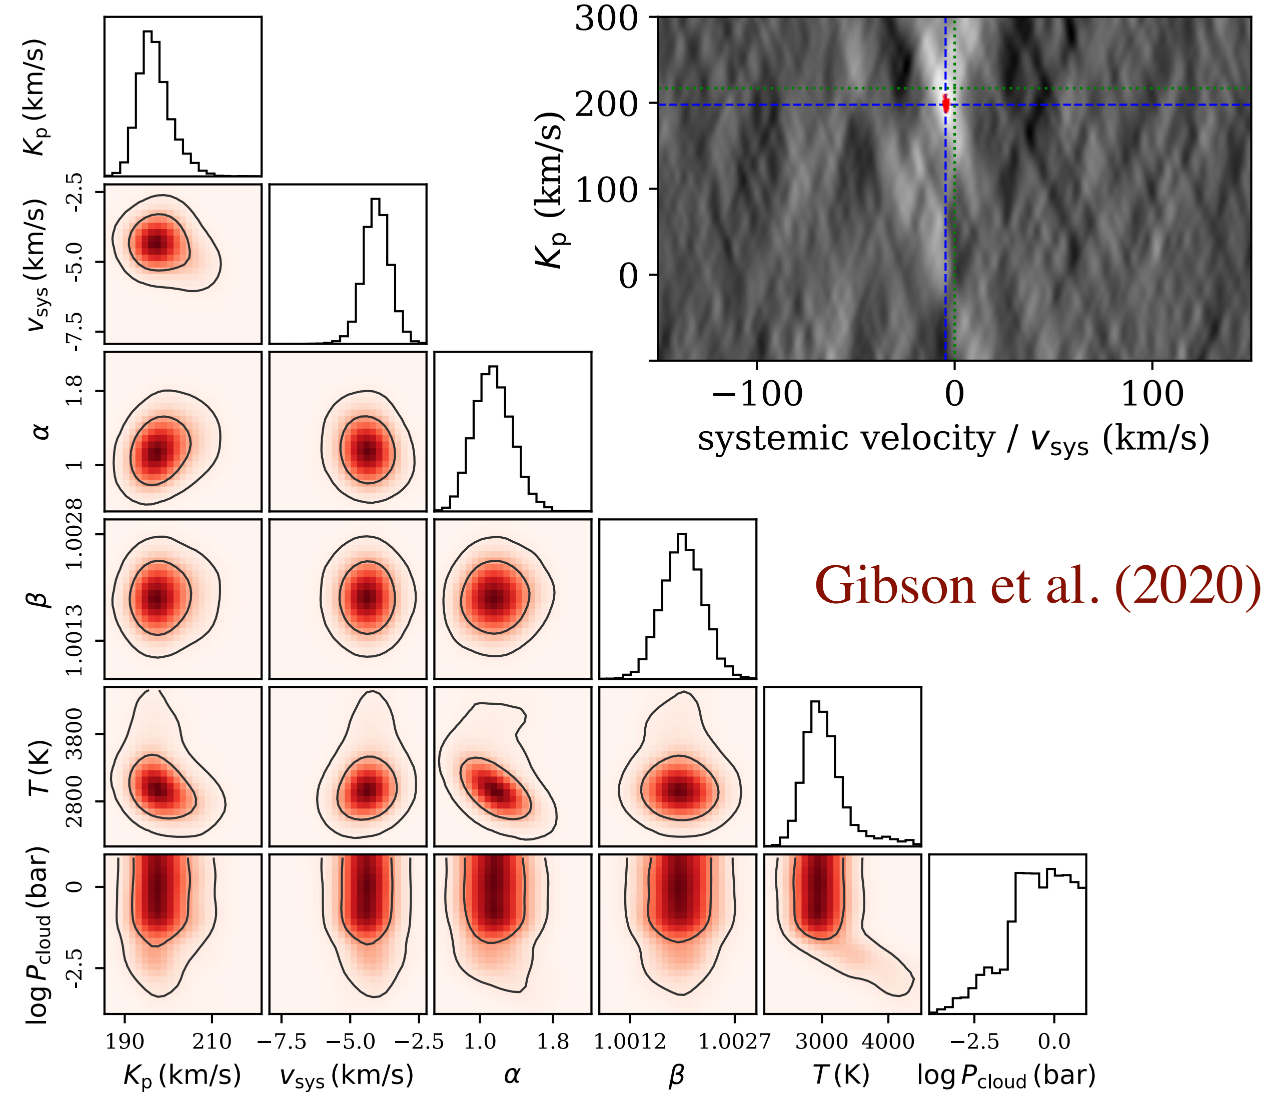

Gibson et al. (2020)

Comparison with Brogi & Line (2019)

Gibson et al. (2020):

$$\mathcal{L}(\theta) = \prod_i \frac{1}{\sqrt{2\pi(\beta\sigma_i)^2}} \exp\left(-\frac{1}{2}\frac{(f_i - \alpha m_i)^2}{(\beta\sigma_i)^2}\right)$$

We use this likelihood directly

These are equivalent in practice (for high-res)

Hybrid approach 'Modified' Brogi & Line (Gibson et al. 2020):

$$\ln \mathcal{L}(\theta) = -\frac{N}{2} \ln\left[\frac{1}{N}\left(\sum_i \frac{(f_i - \alpha m_i)^2}{\sigma_i^2}\right)\right]$$

Brogi & Line (2019):

$$\mathcal{L}(\theta) = \prod_i \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{1}{2}\frac{(f_i - m_i)^2}{\sigma^2}\right)$$

(intuitively equivalent to setting reduced χ^2 to 1 by scaling σ)

$$\ln \mathcal{L}(\theta) = -\frac{N}{2} \ln\left[\frac{1}{N}\left(\sum_i (f_i - m_i)^2\right)\right]$$

The Forward Model Atmosphere

- Gibson et al. (2020) used a semi-analytic model; isothermal, no pressure broadening (Heng & Kitzmann 2017)

$$R(\lambda) = R_0 + H\left[\gamma + \ln\left(\frac{P_0}{mg}\sqrt{\frac{2\pi R_0}{H}}\right)\right] + H\ln\sum_j \chi_j\sigma_j(\lambda)$$

(Really useful if you want to play about with detections/retrievals without a detailed model)

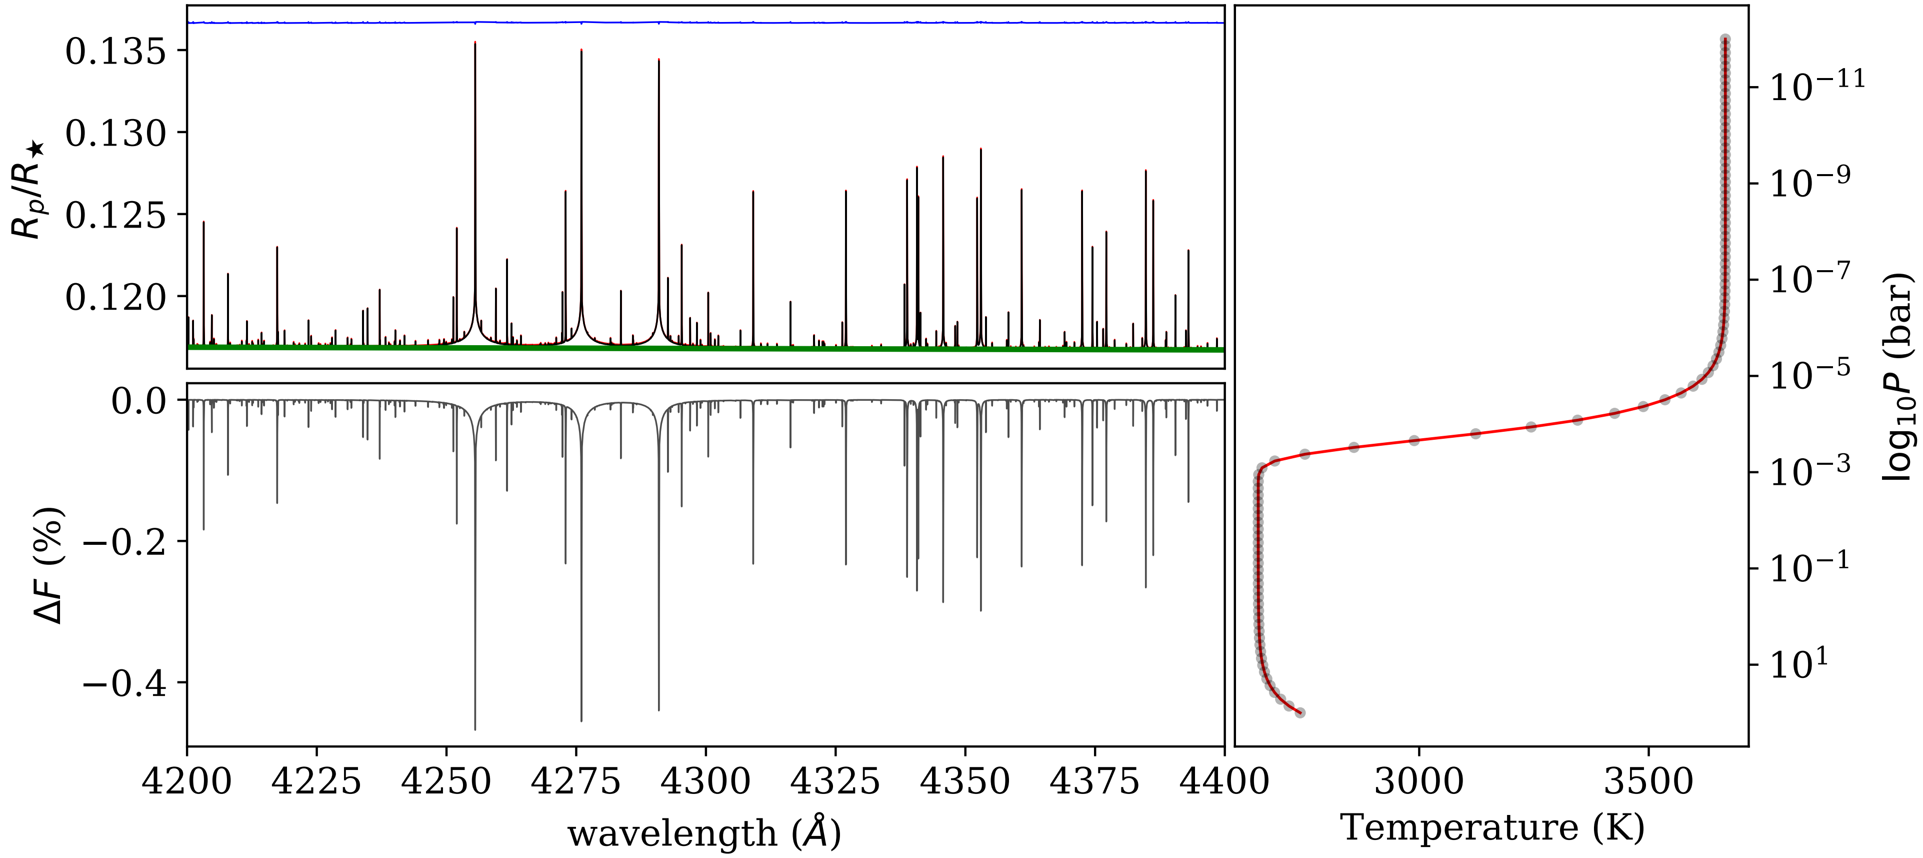

- Now using a more sophisticated model with T-P profile + 100 layers. Validated against petitRADTRANS (Molliere et al. 2019)

Models typically computed at R~200,000-300,000, then convolved with a Gaussian kernel (our model is nothing fancy, just designed to be simple and fast!)

How do we account for SysRem messing around with our data?

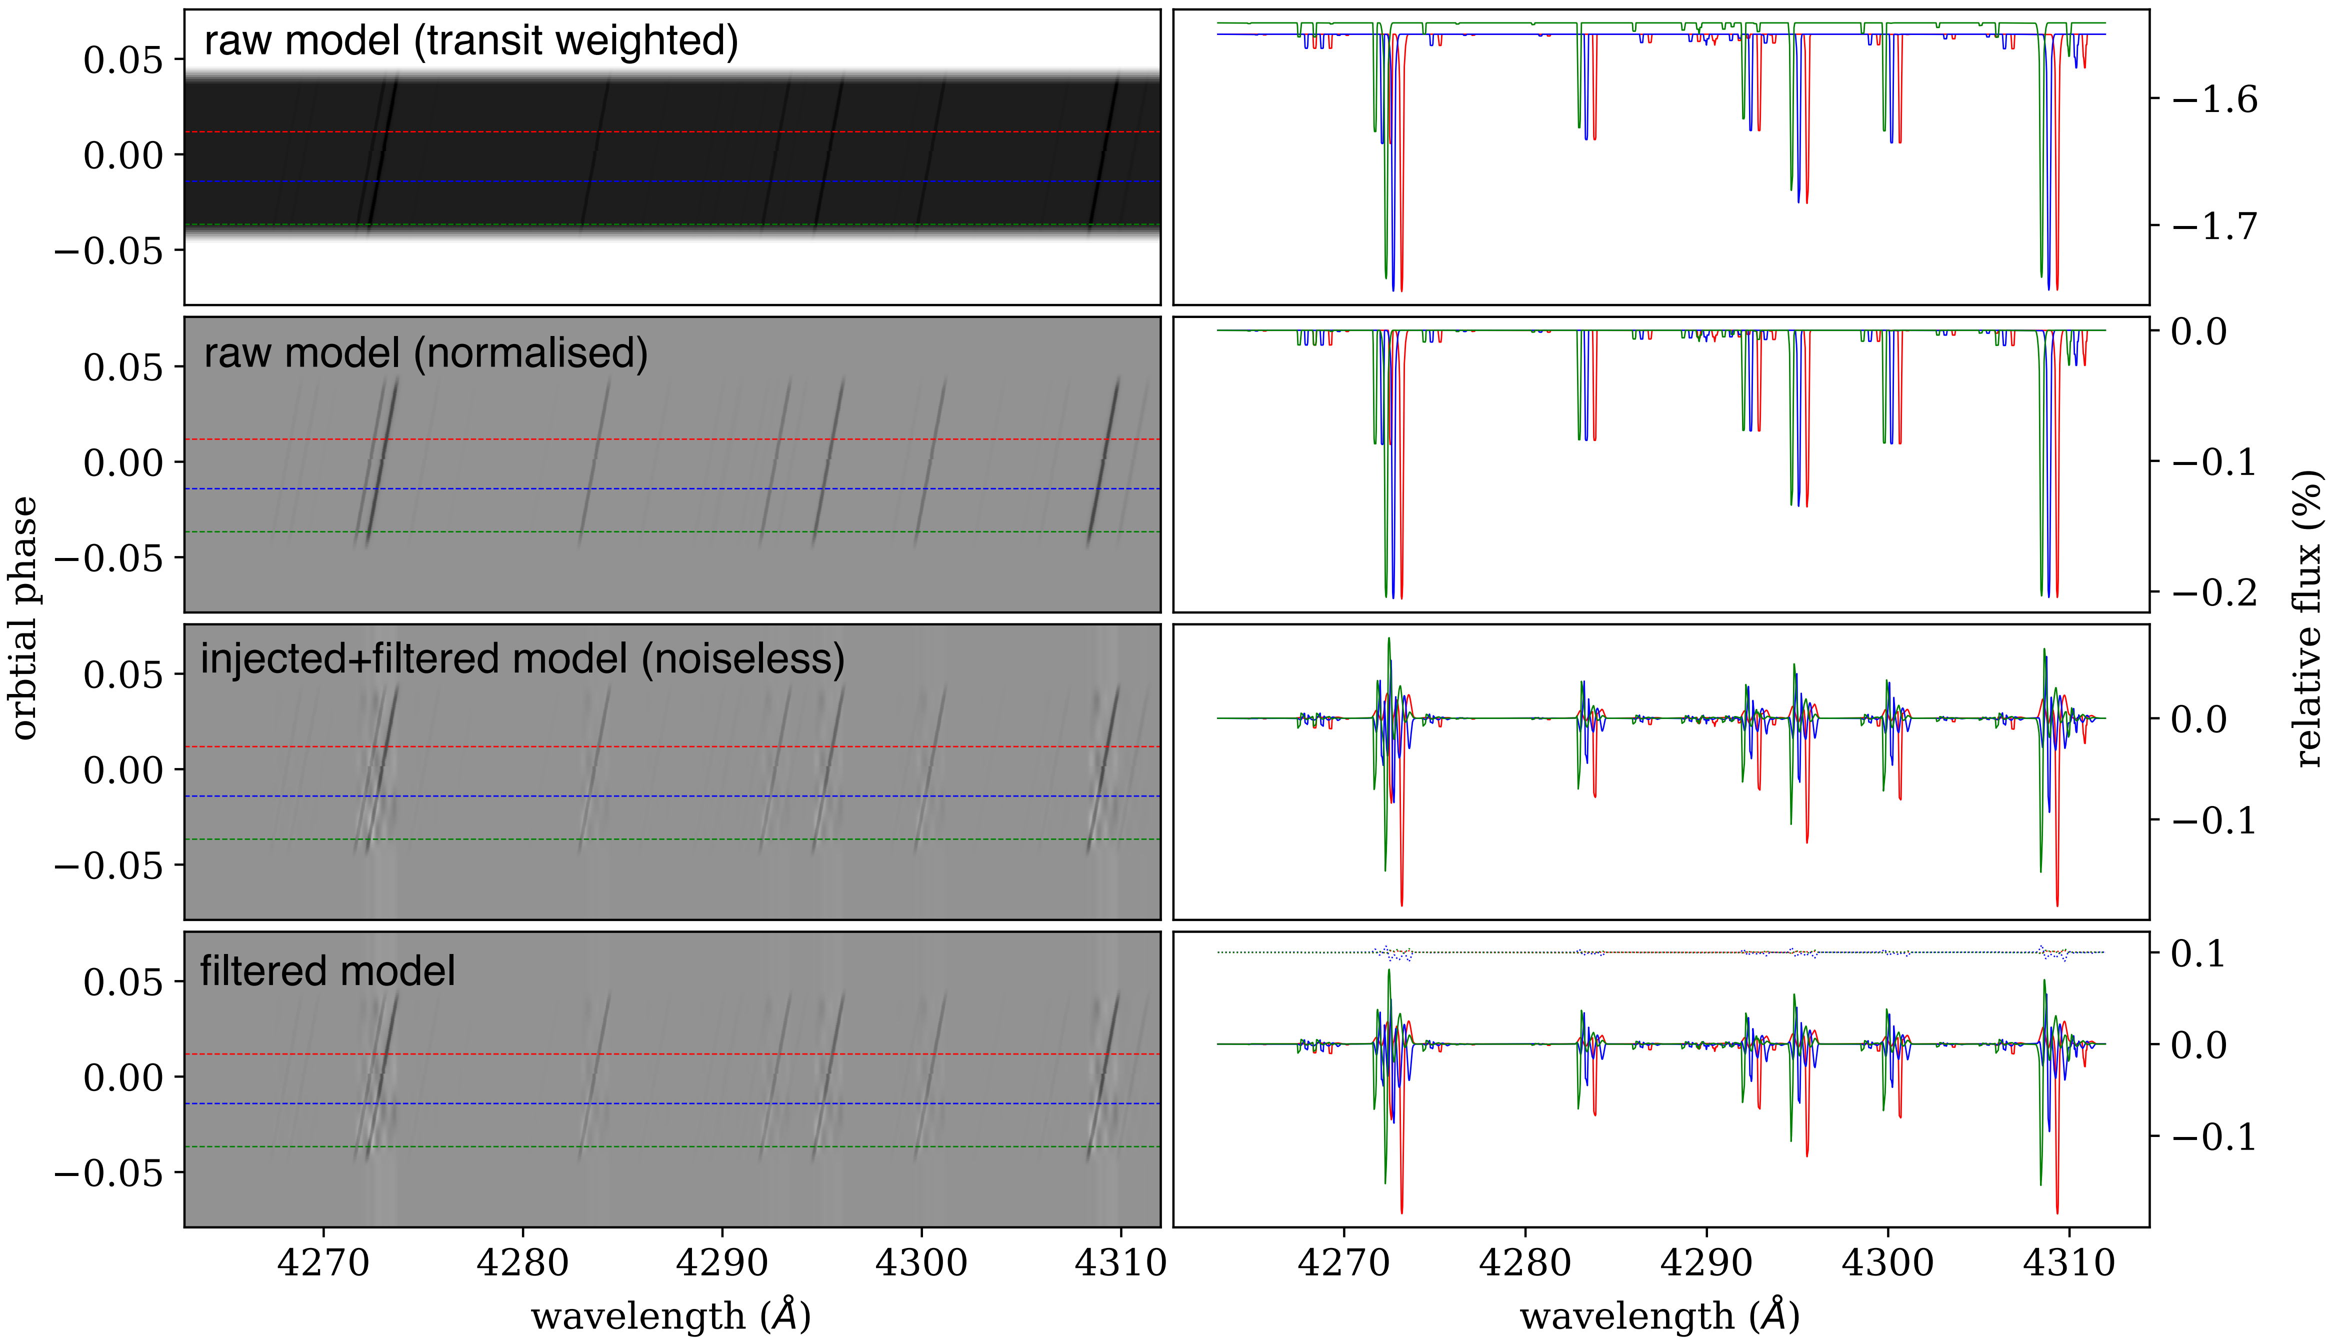

1) We run SysRem on the data

2) Perform filtering on the model f*or every likelihood calculation* to match the data-preprocessing

(This is one of the key differences between high and low res retrievals!)

Our model is validated with injection tests into *real* data

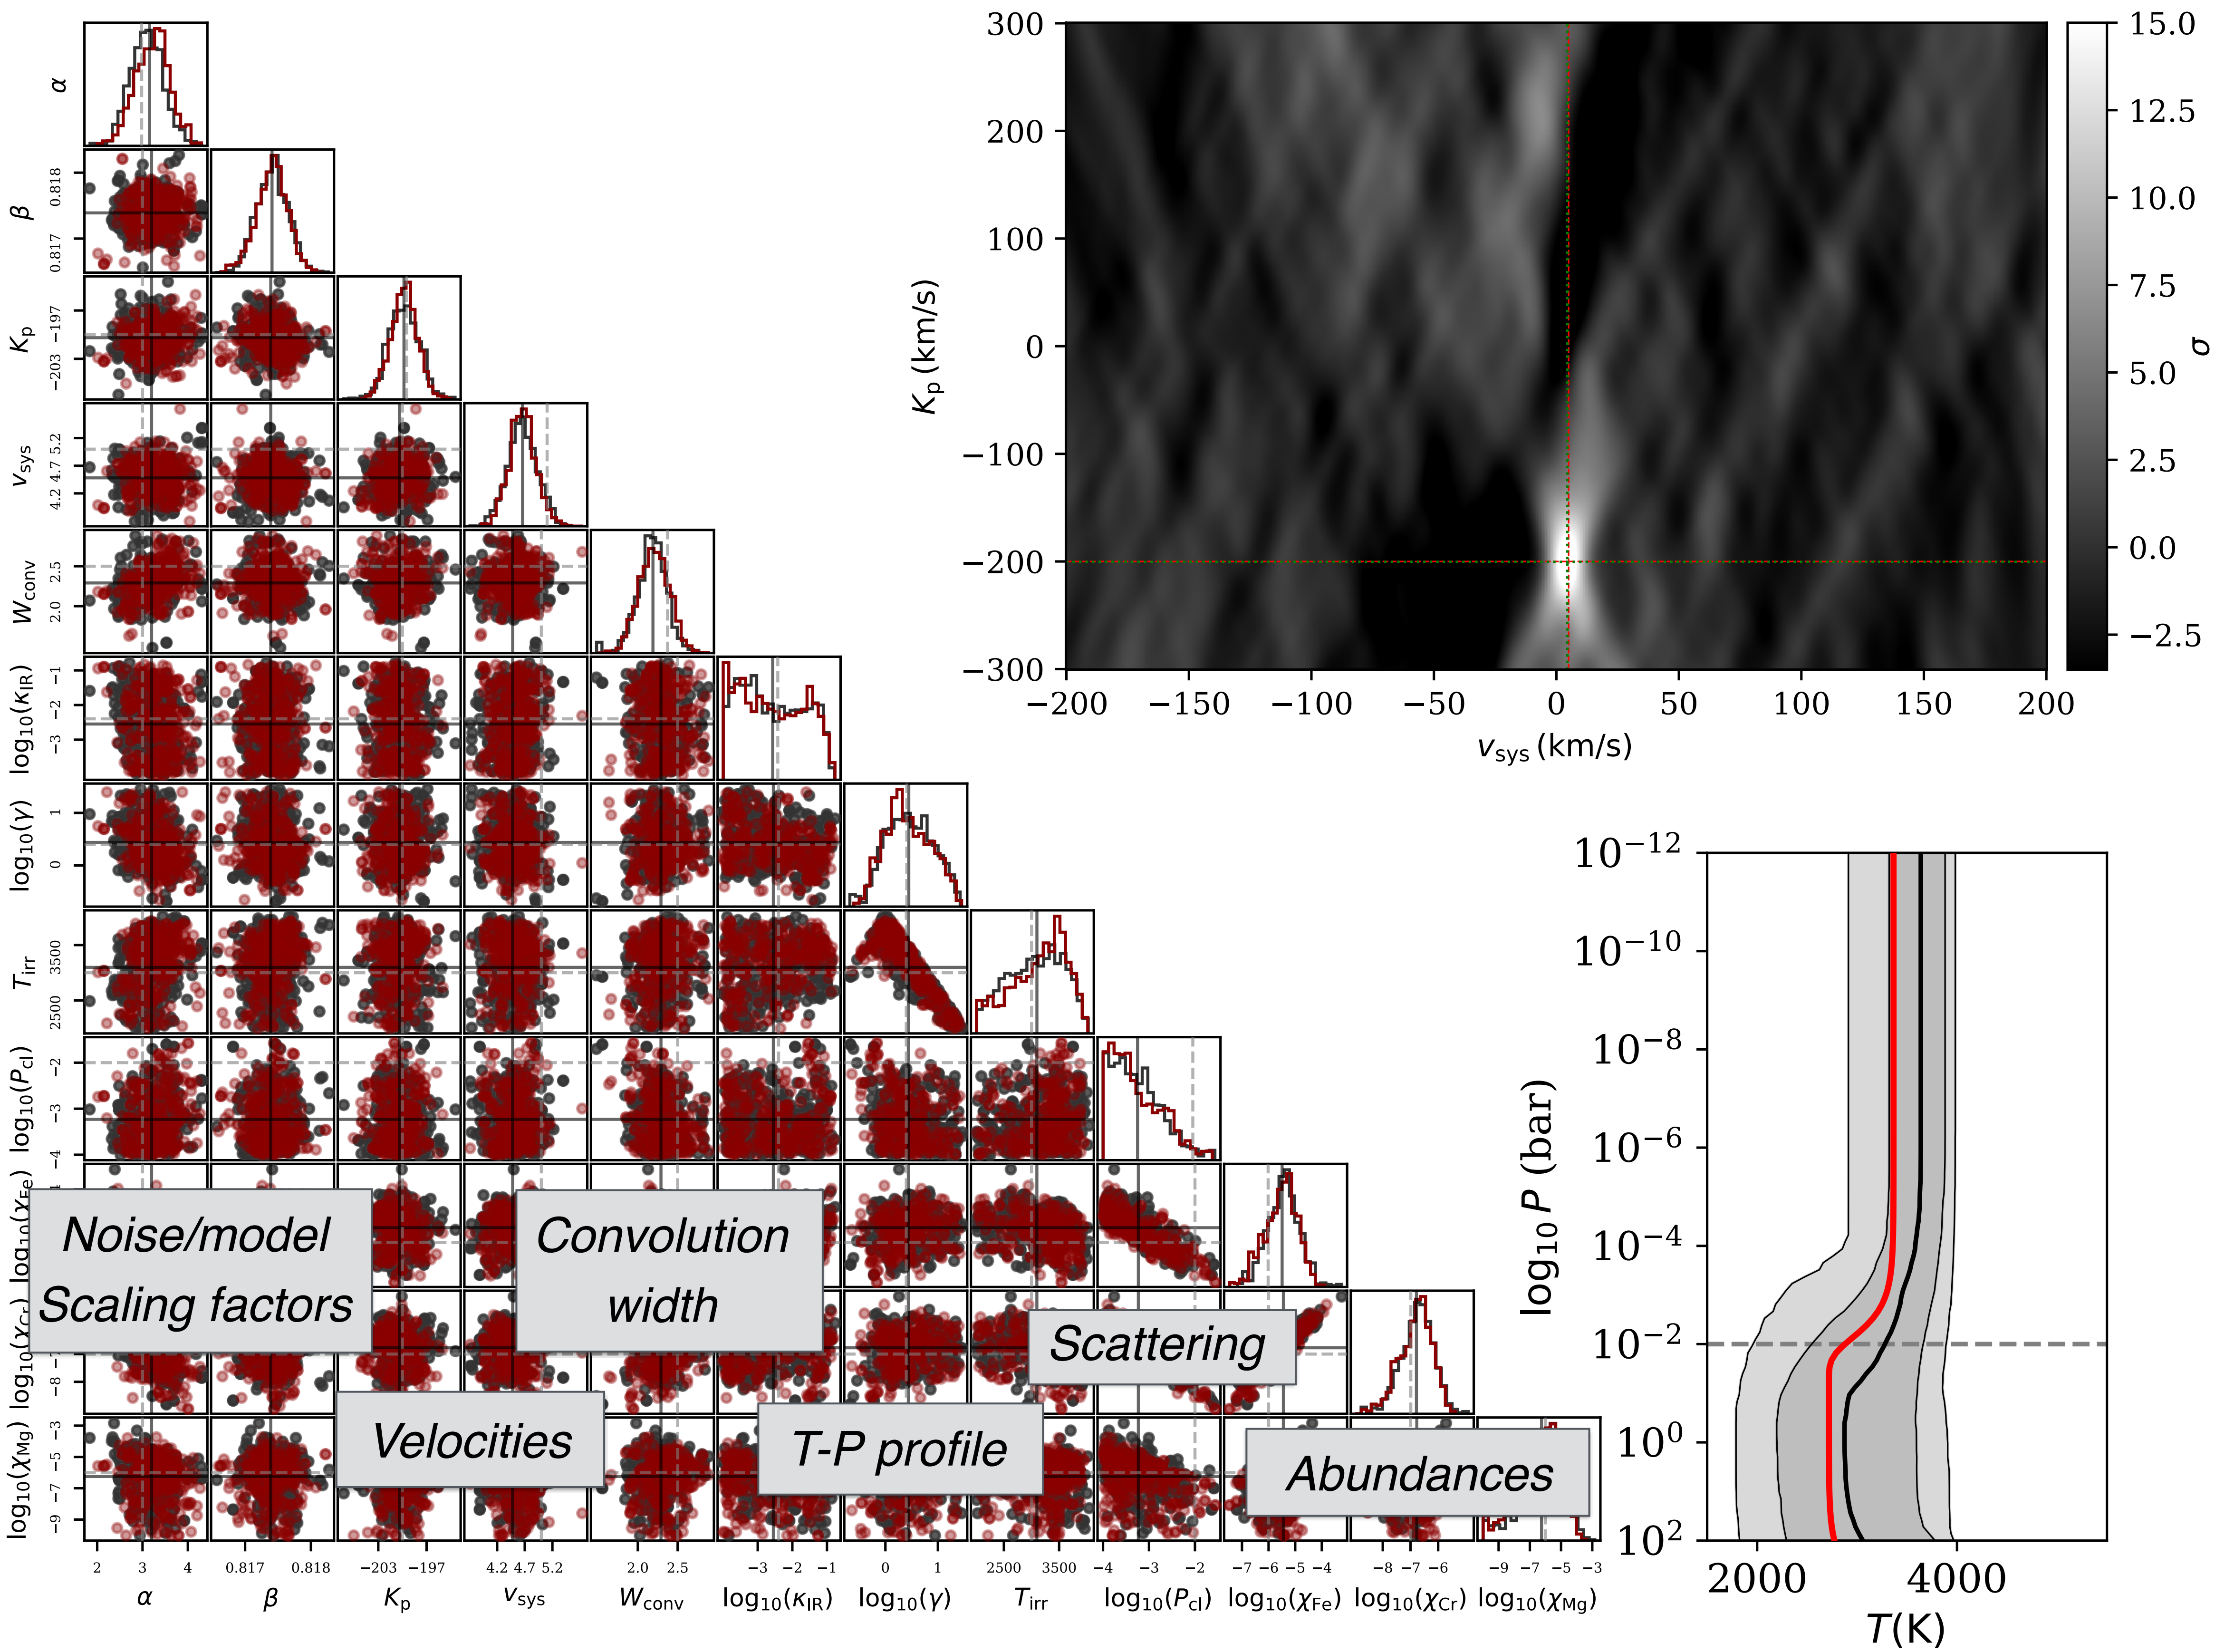

What can we do with a full retrieval?

Improve on measurements from CCF:

- Robust constraints on measured velocities
- Better stats on detection significances (?)
- Constrain width of convolution kernel (wind speeds and/or rotation)

In addition:

- Constrain the T-P profile
- Determine atmospheric abundances
- Marginalise over cloud/scattering properties
- Marginalise over the noise properties

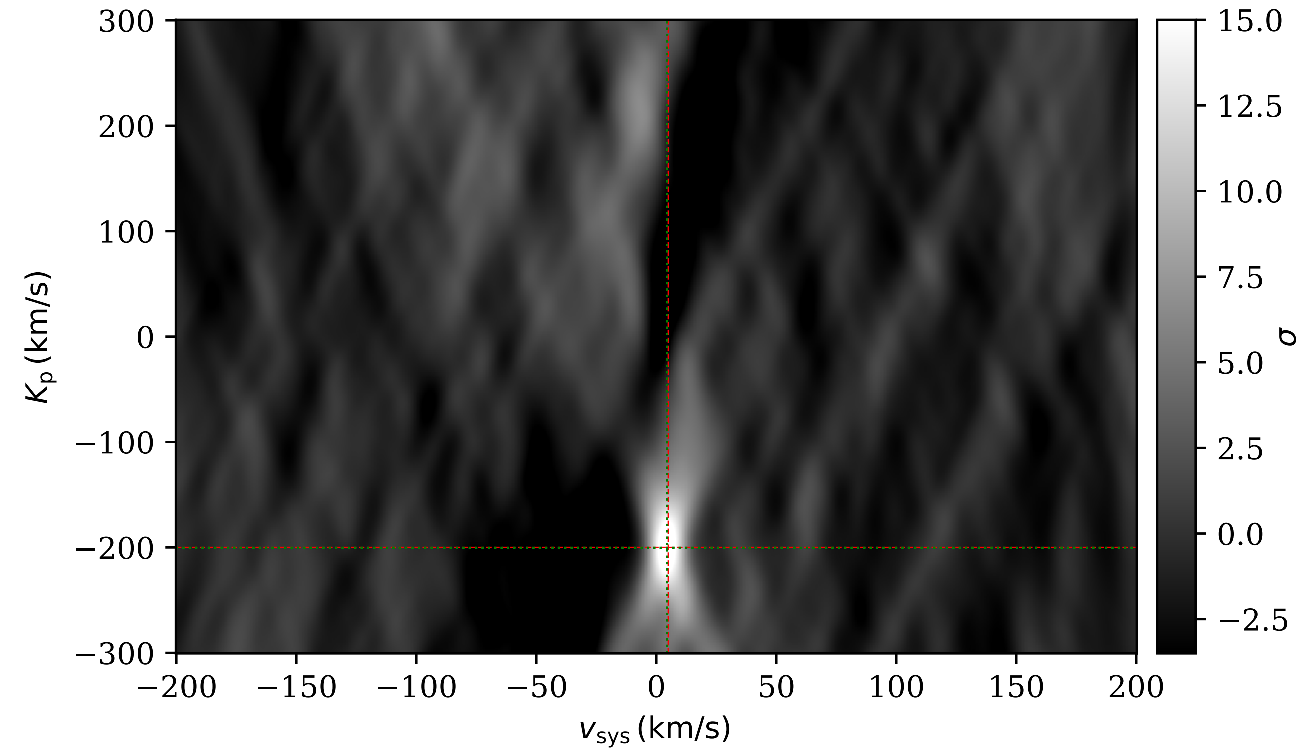

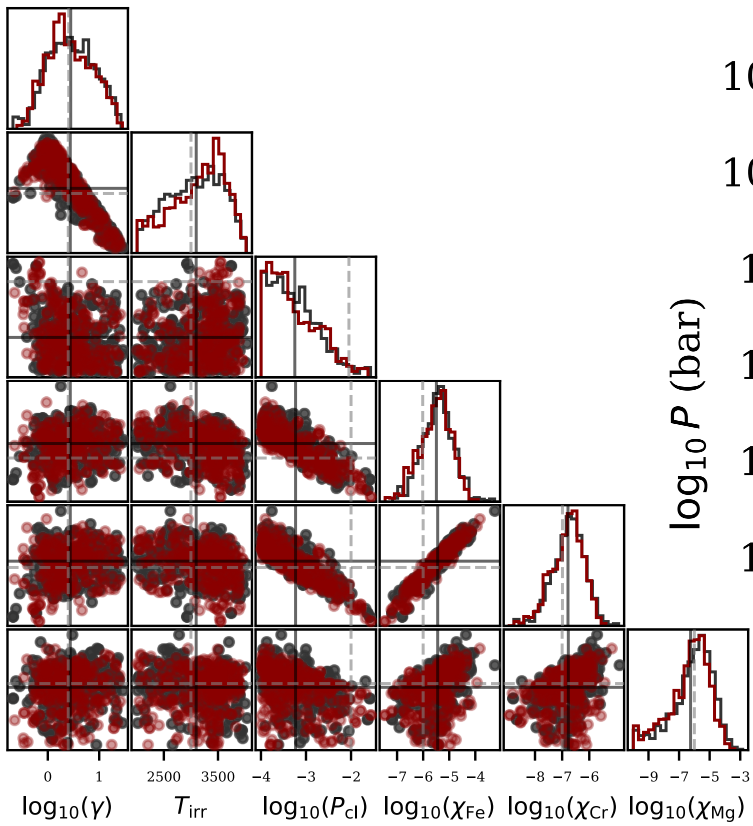

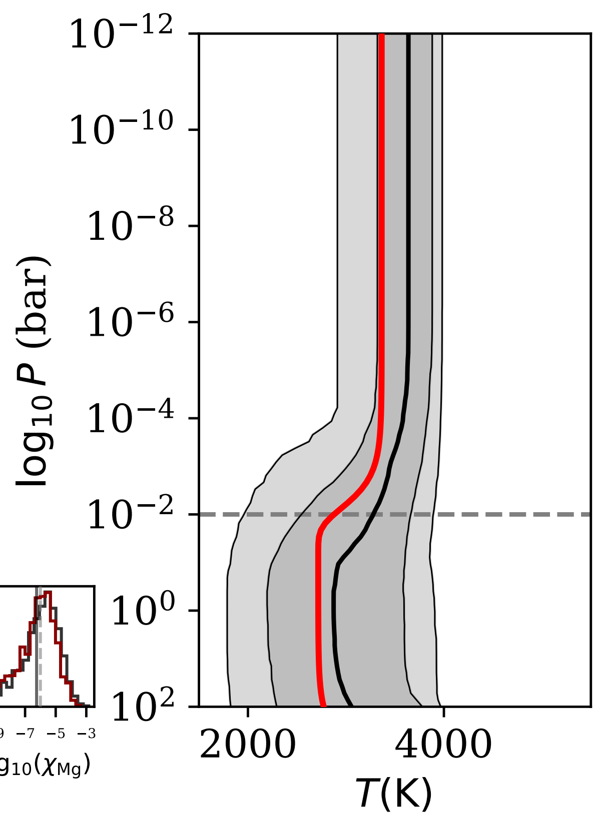

Retrieval example of WASP-121b UVES data

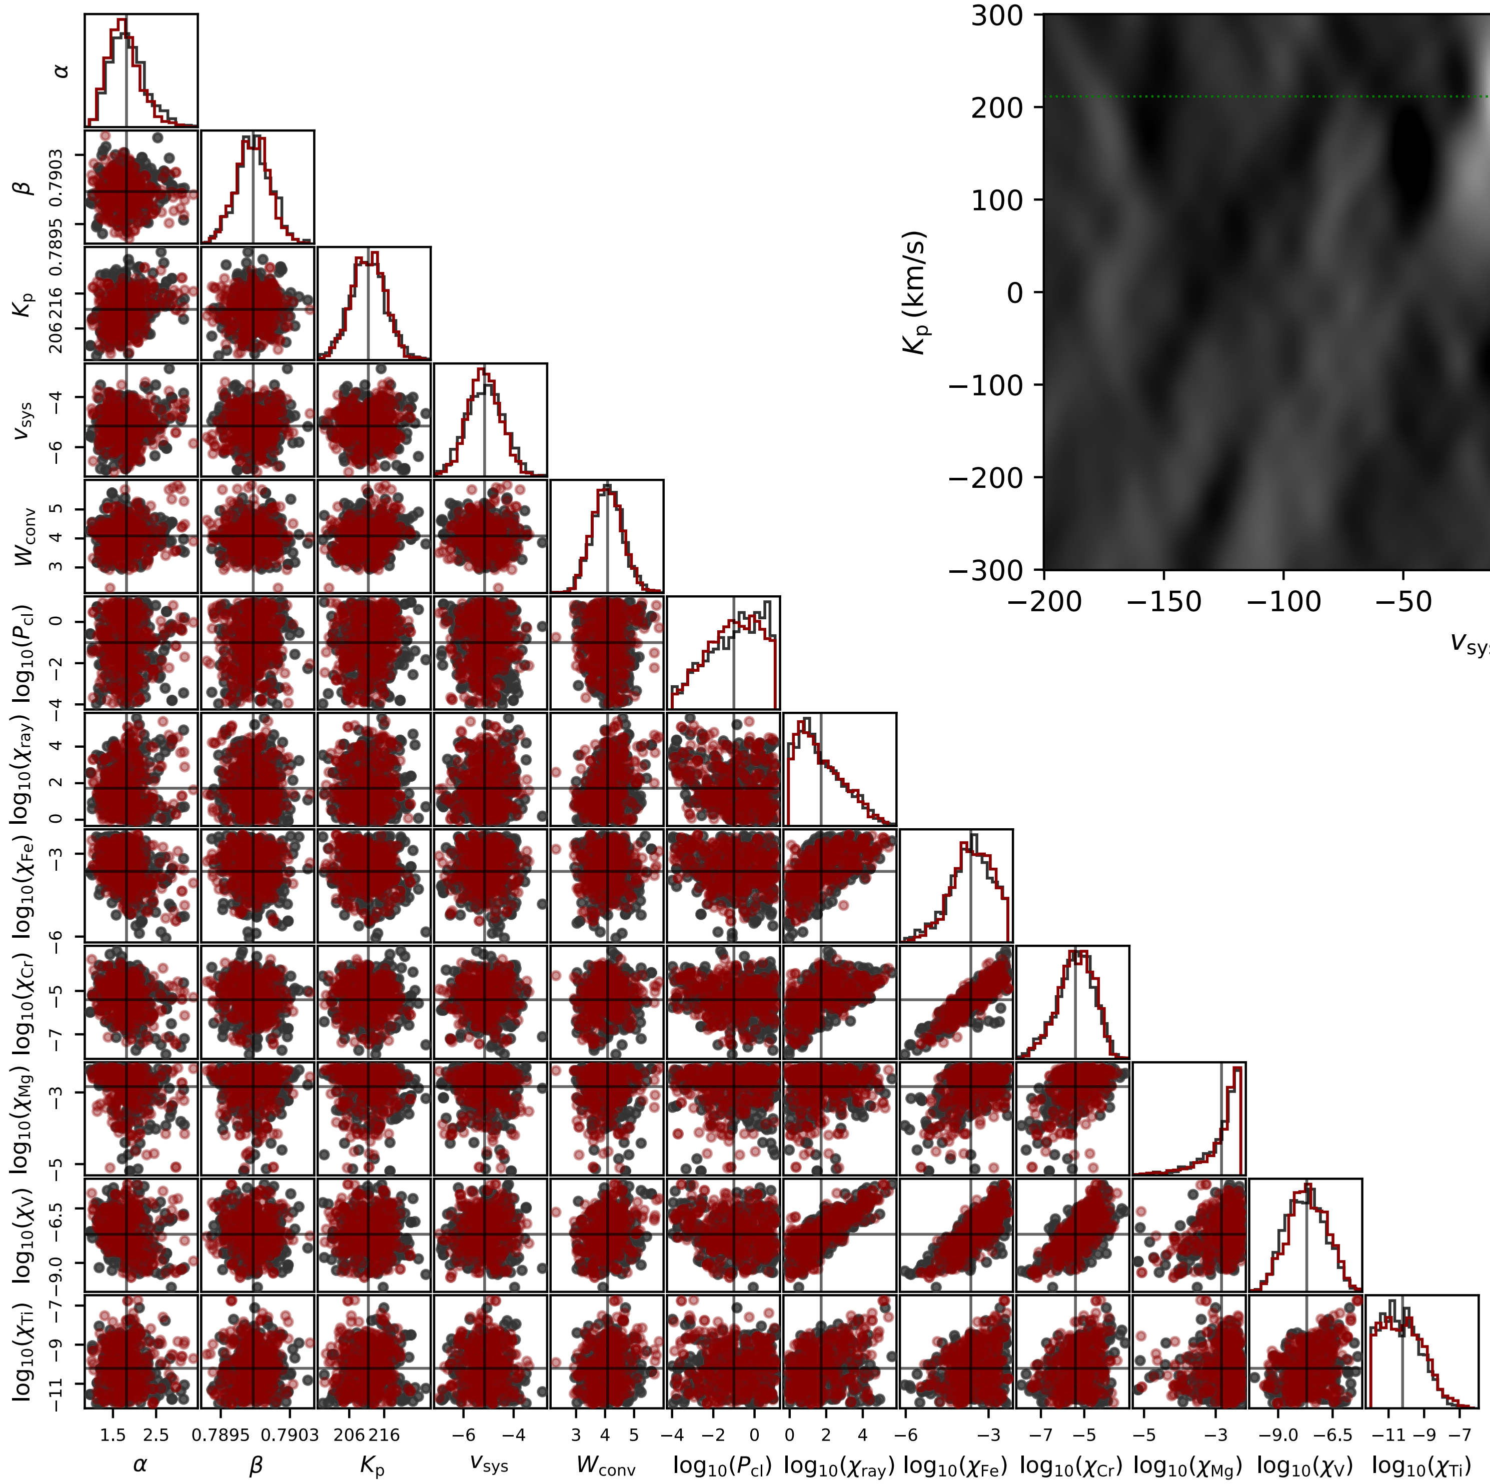

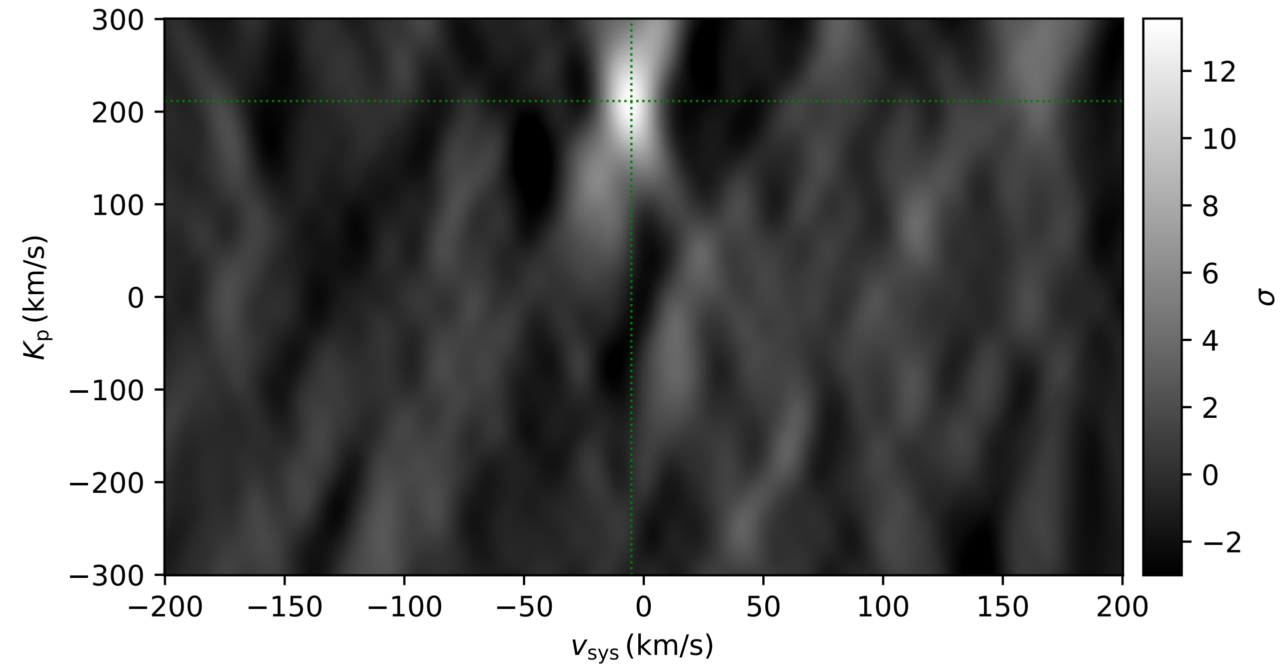

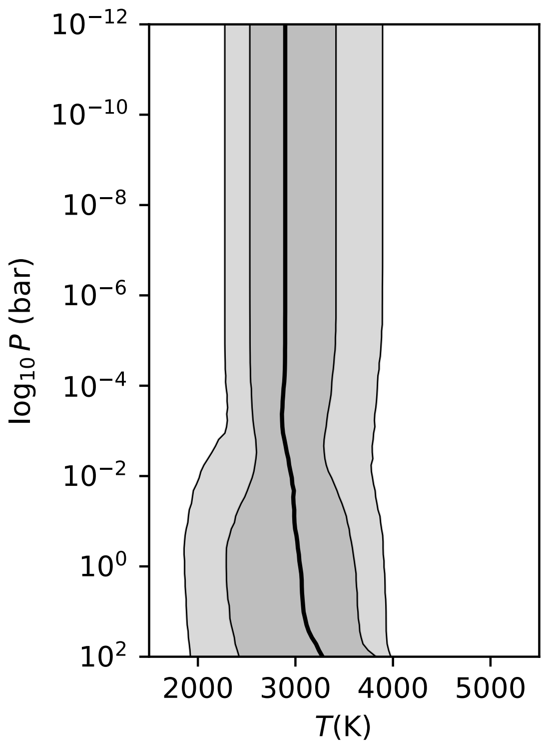

(Full retrieval ~2h on laptop, using 29 spectral orders, ~10,000,000 data points)

Retrieval example of WASP-121b UVES data

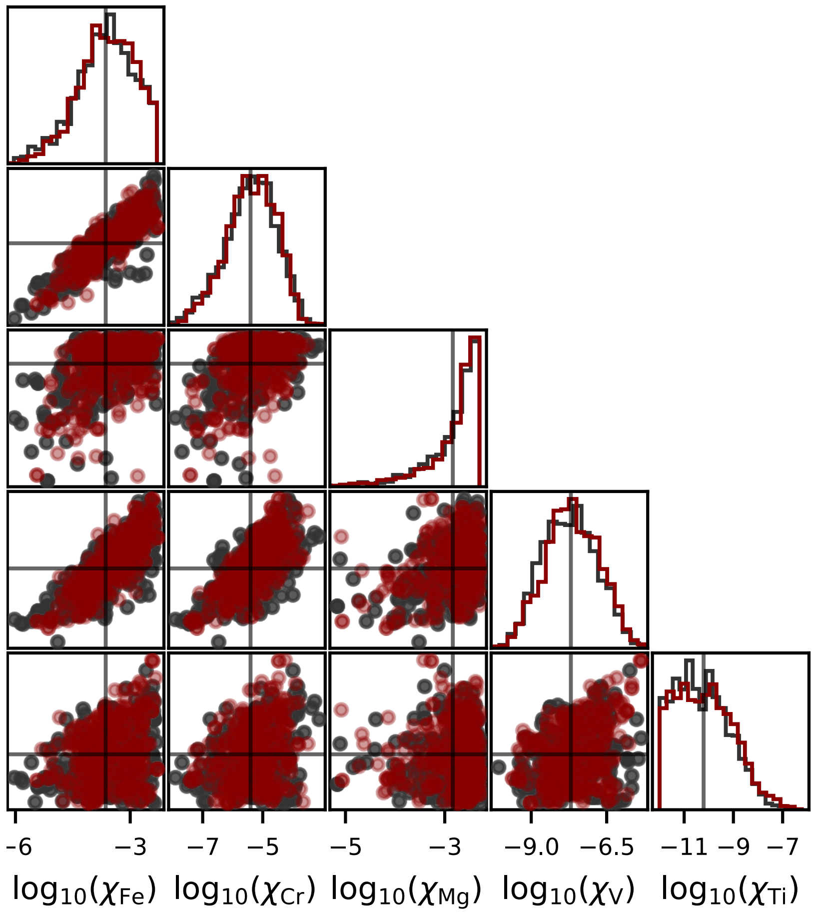

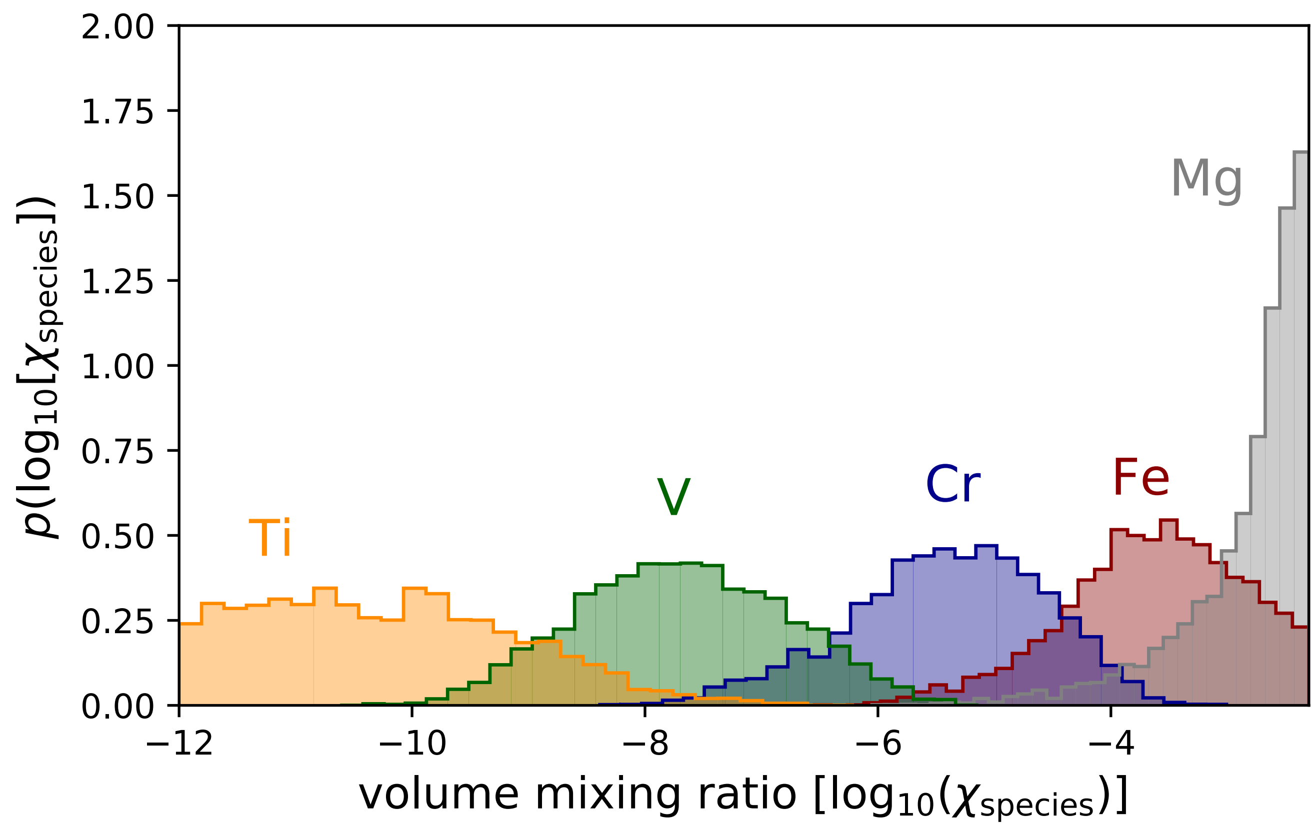

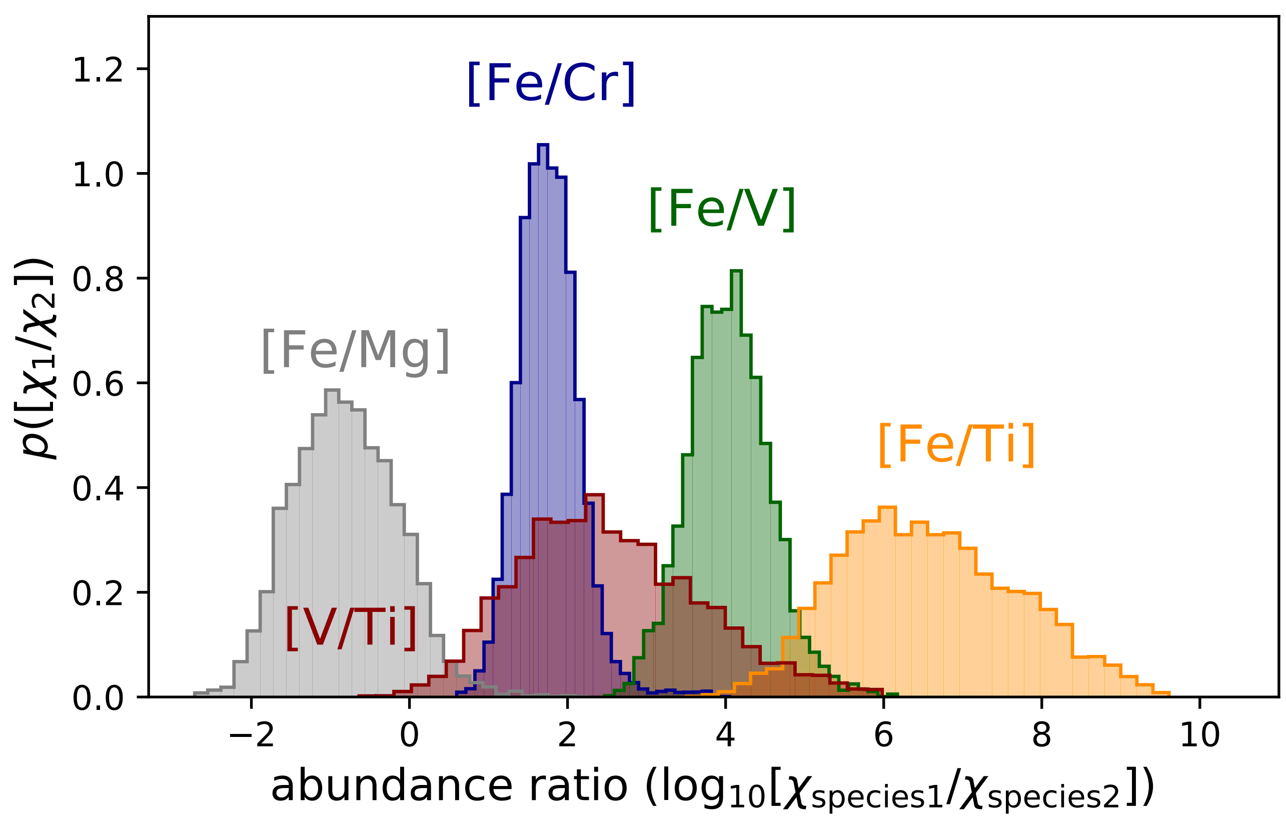

- Abundance constraints from transmission spectra are poor
- However, they are strongly correlated
- Therefore *relative* abundances are much easier to constrain: e.g. [Fe/Cr] = 1.77 ± 0.39
- Mg unphysically high? What does this mean?

What else can we do?

Inference is only as good as the model!
(+ the statistical model)

- WASP-121b cannot be explained by a hydrostatic atmosphere. Independent scale factor for each species?
- Can constrain velocity offset + broadening kernel. Do we need velocity shifts between species? What about different broadening?
- Time dependence of the signals?
- Combine datasets from different epochs instruments (including low-res)

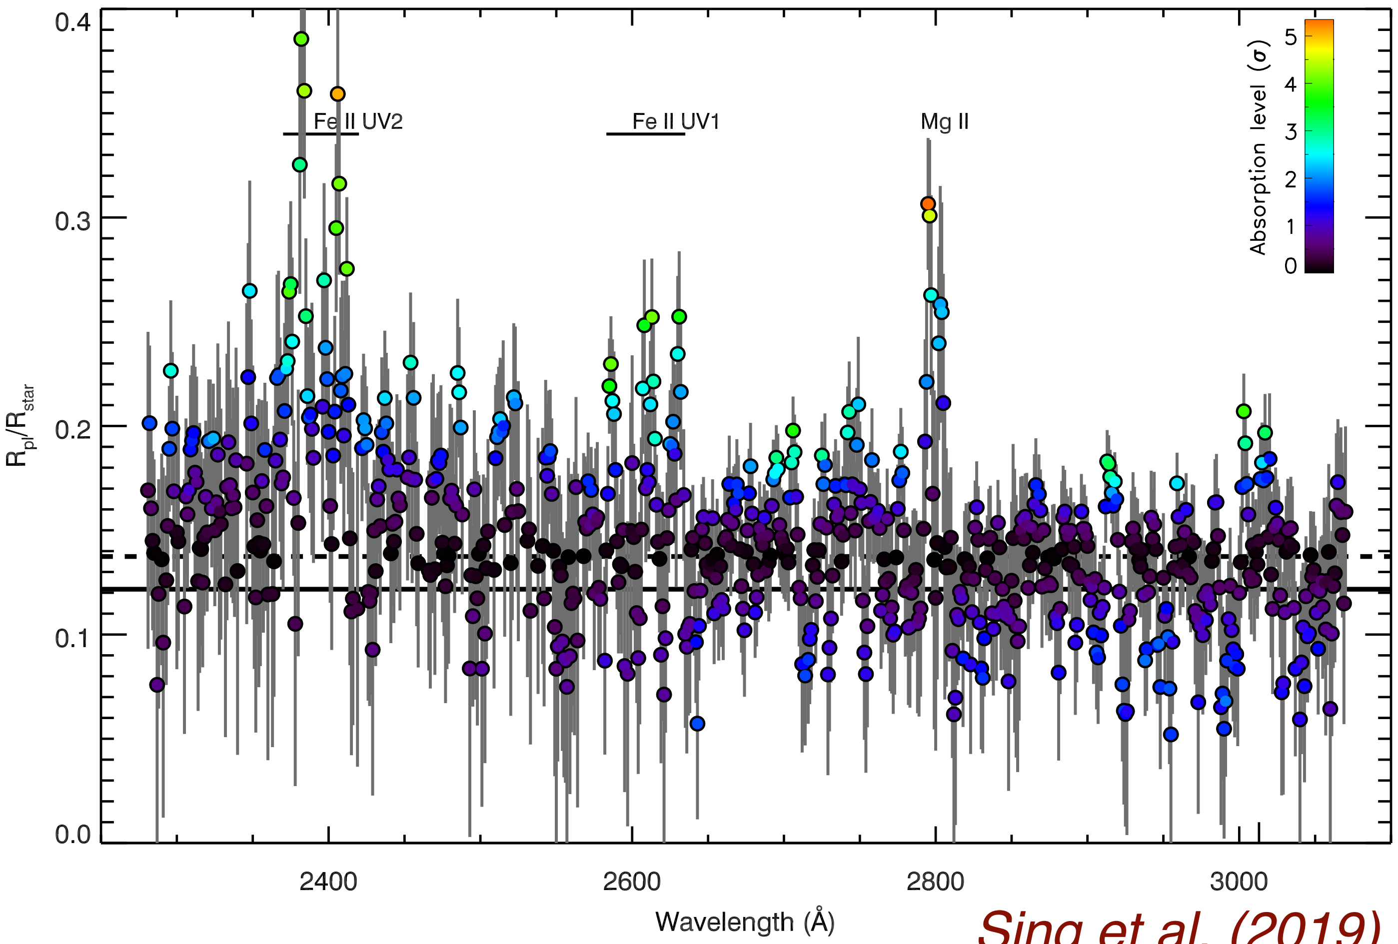

Sing et al. (2019)

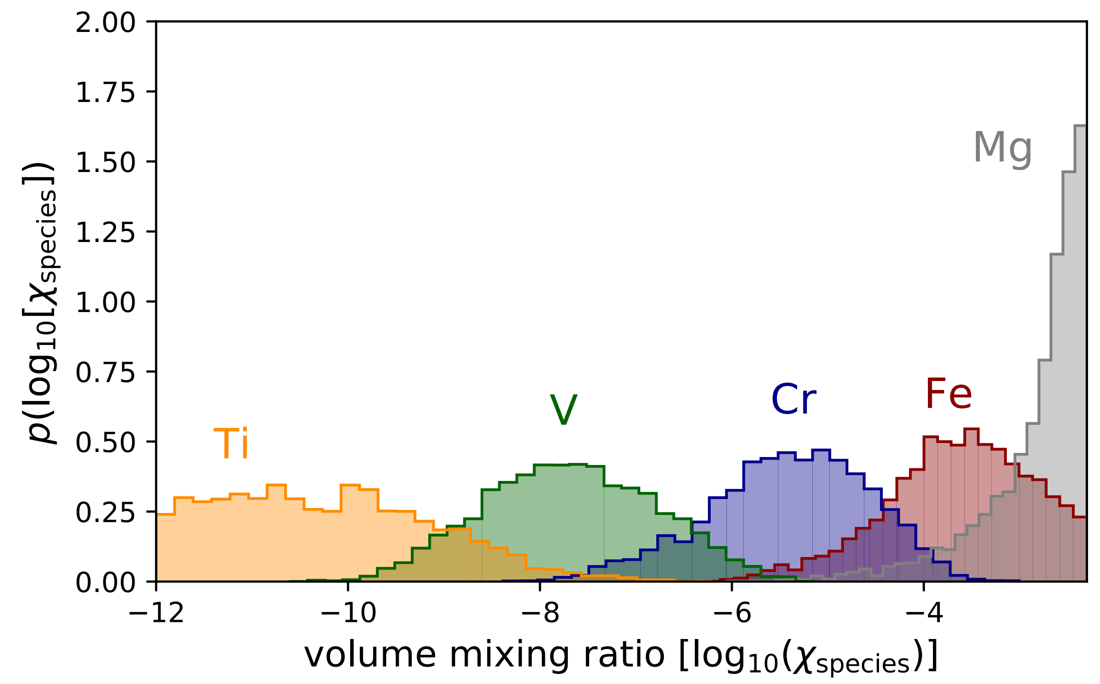

Conclusions

- High-resolution cross-correlation spectroscopy enables unambiguous detection of species
- However, cannot compare model templates
- The likelihood approach opens up a huge new parameter space to explore…
  - Careful in defining noise properties of the data
  - Account for SysRem (or equivalent) filtering
  - Huge number of datapoint means *everything* in the likelihood calculation must be optimised
- …(relative) abundances, T-P profiles, velocity shifts, broadening, time-dependence, aerosols, +much more with better models!
- Really important once we lose HST, with no ground-based low-res instrument to replace it

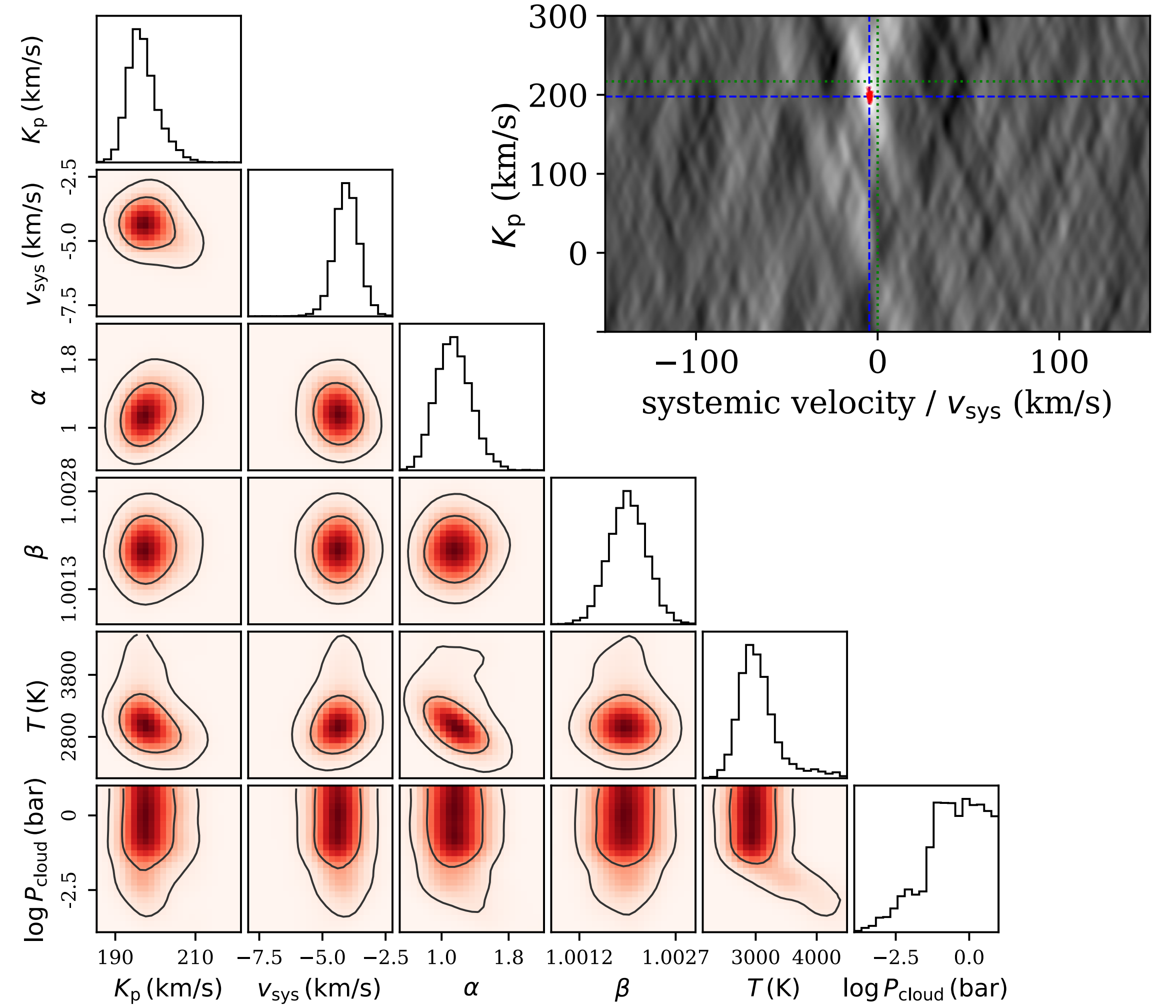

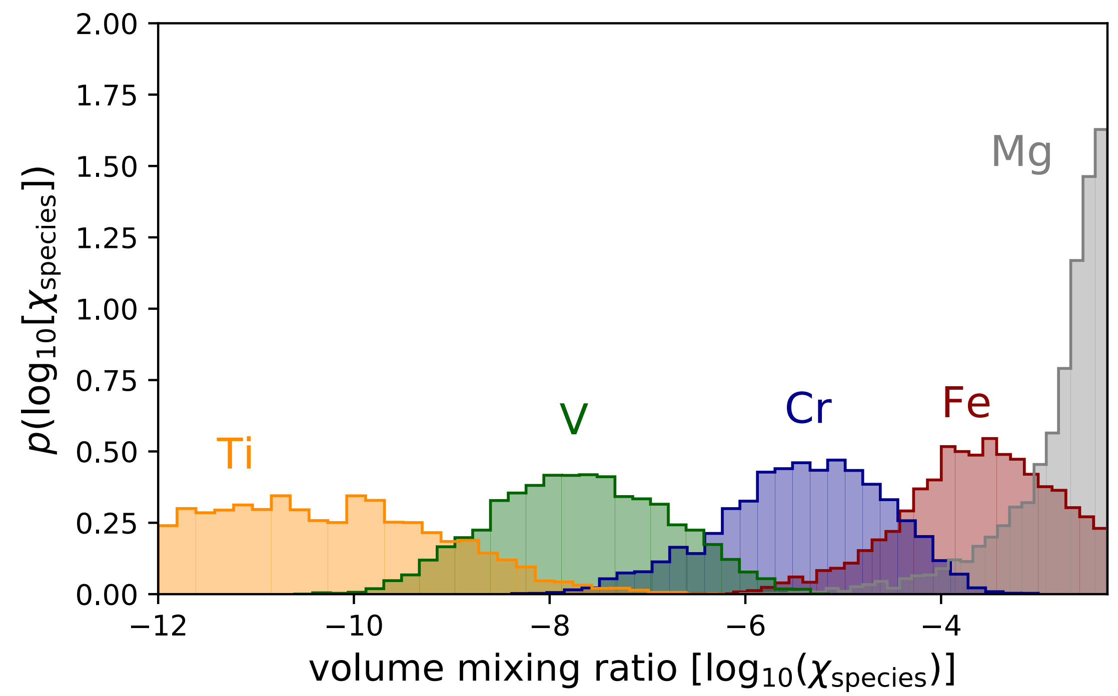

If you want to learn more:

See Brogi & Line (2019); Gibson et al. (2020, *+in prep*)

Neale Gibson, Trinity College Dublin