Perils of false signals

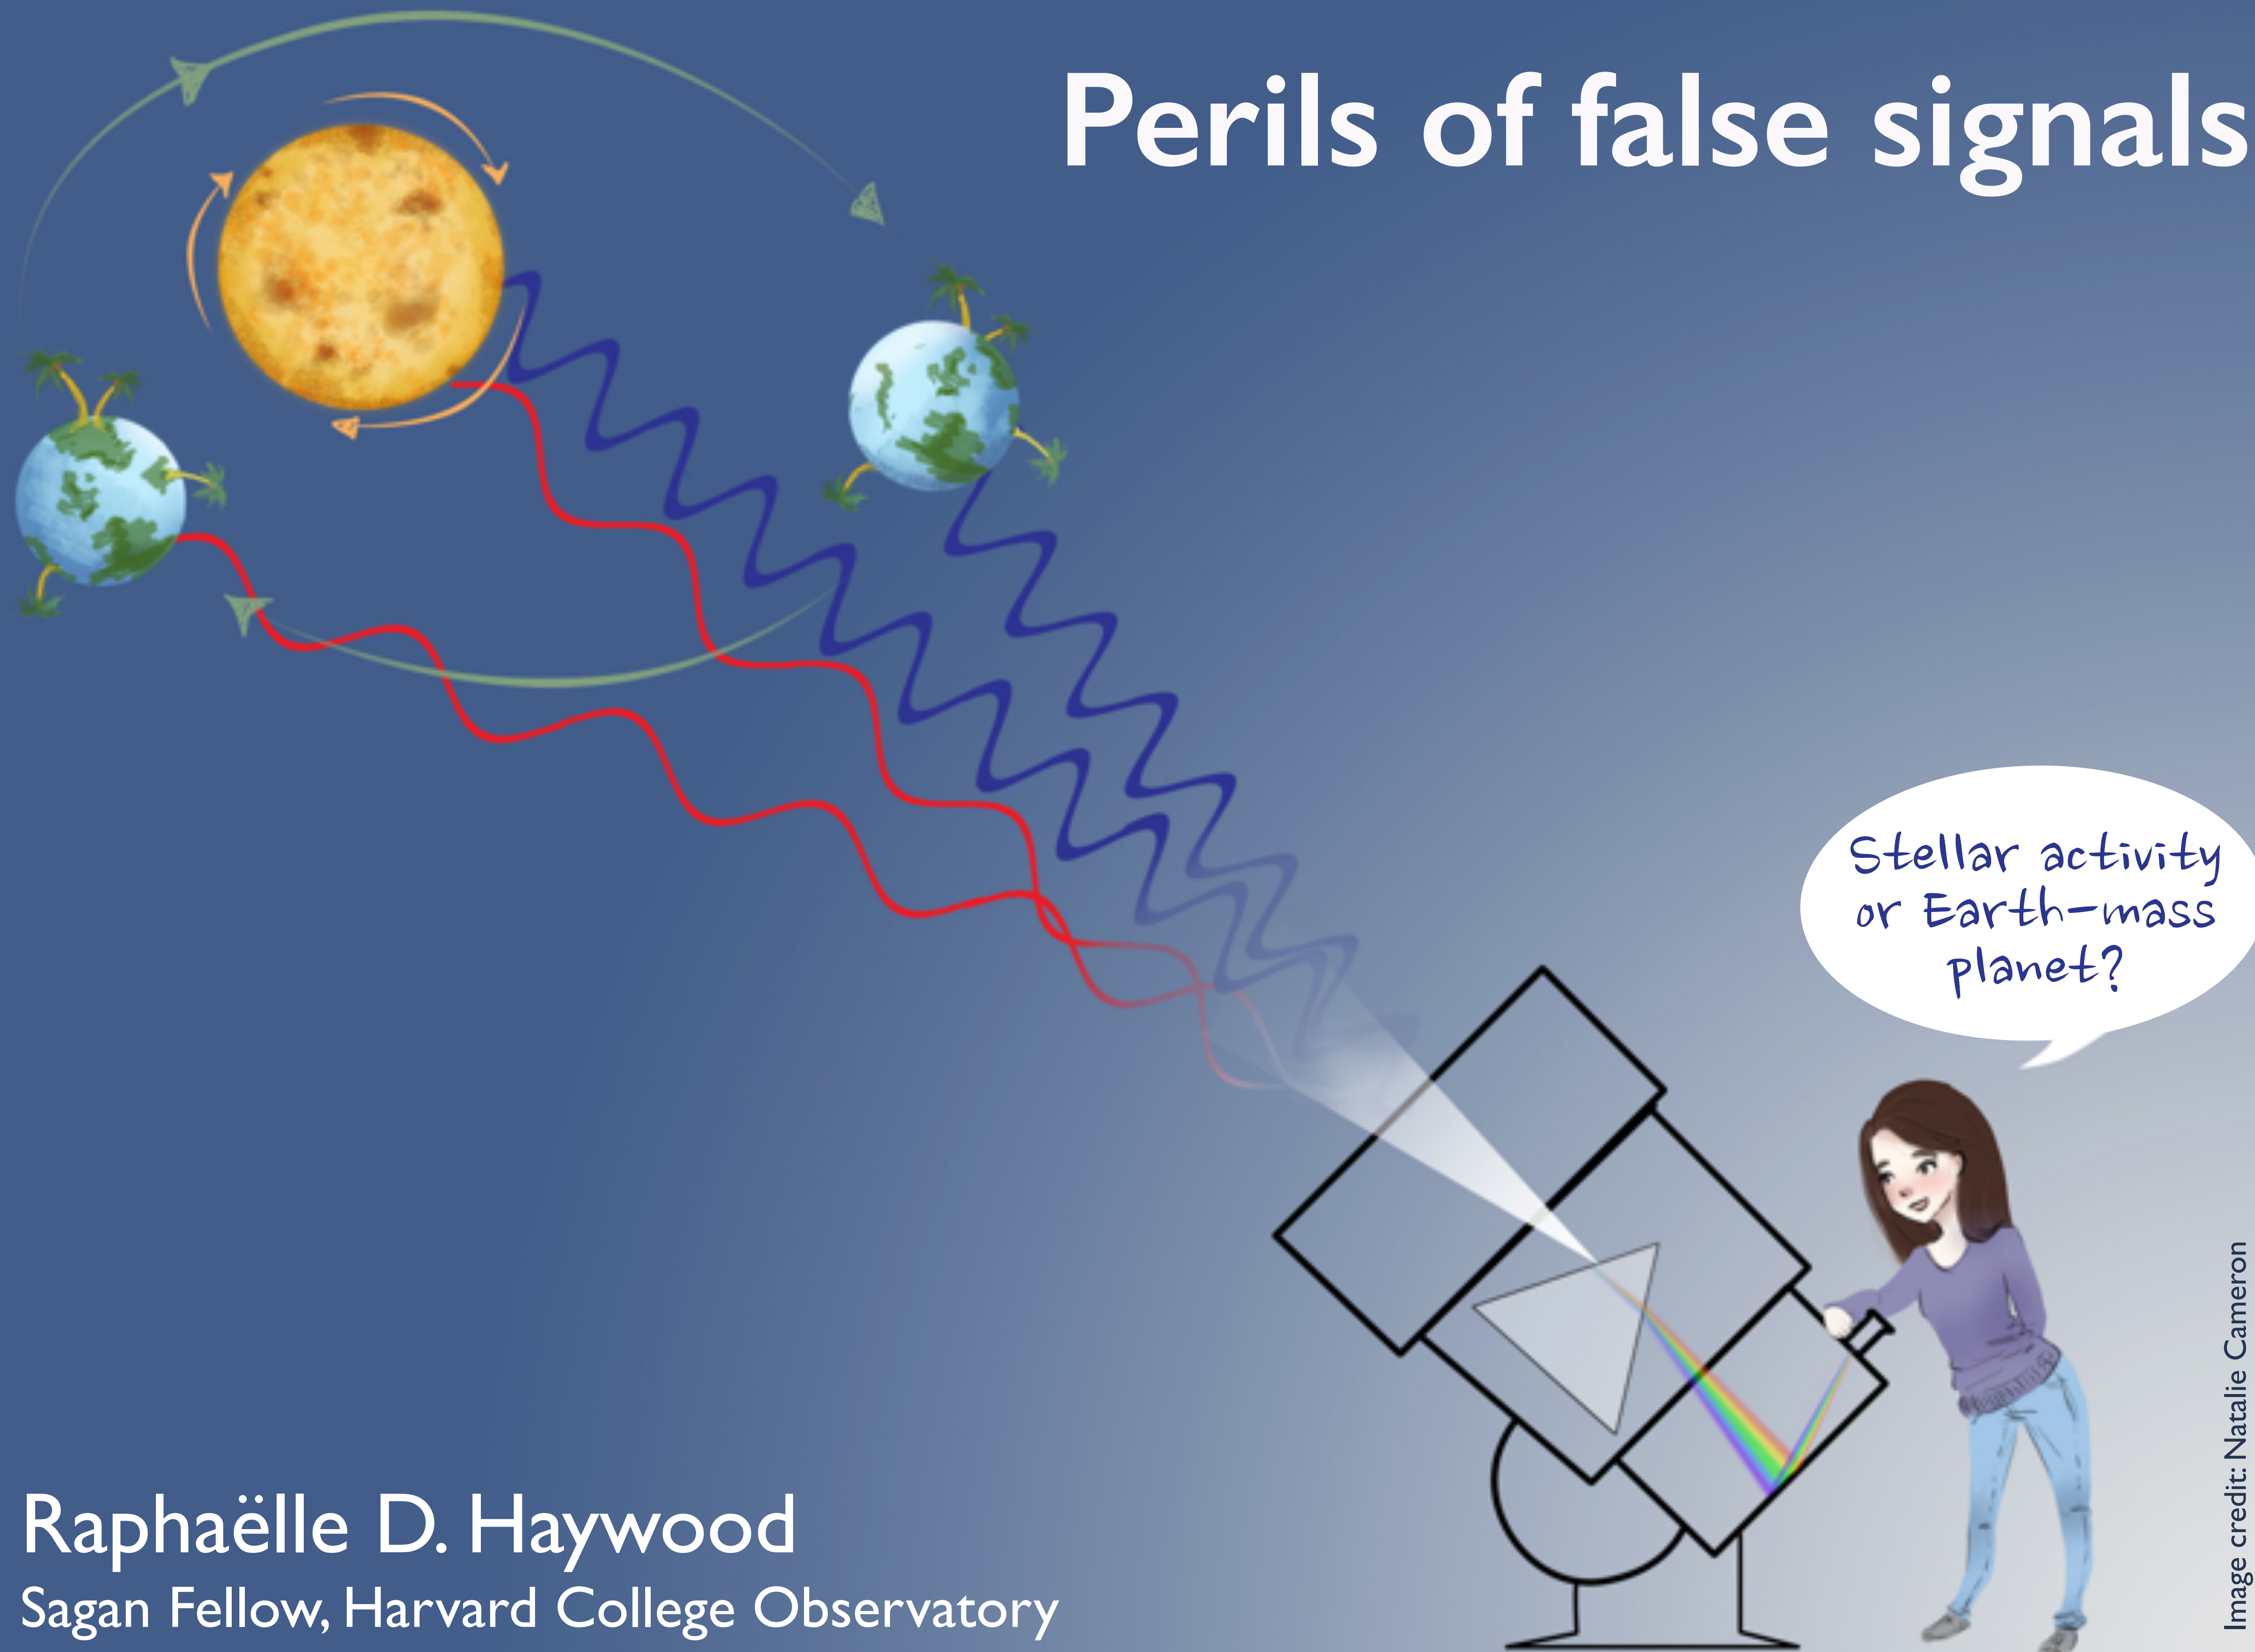

Raphaëlle D. Haywood
Sagan Fellow, Harvard College Observatory

Image credit: Natalie Cameron

Accounting for stellar activity (or any kind of unknown correlated systematics) in RV analyses

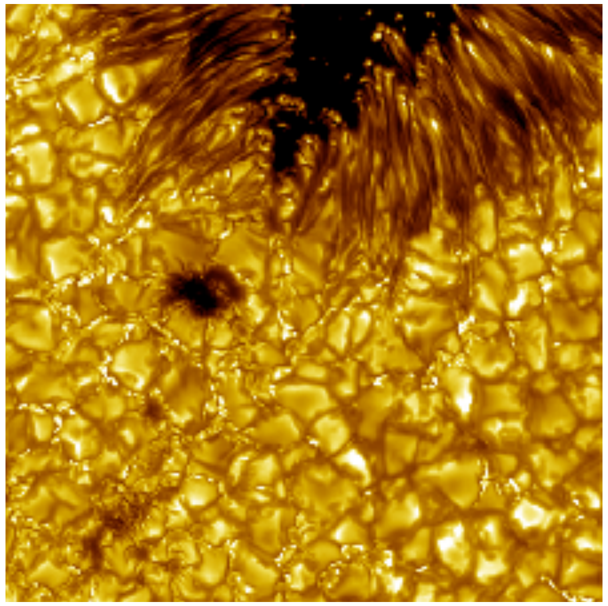

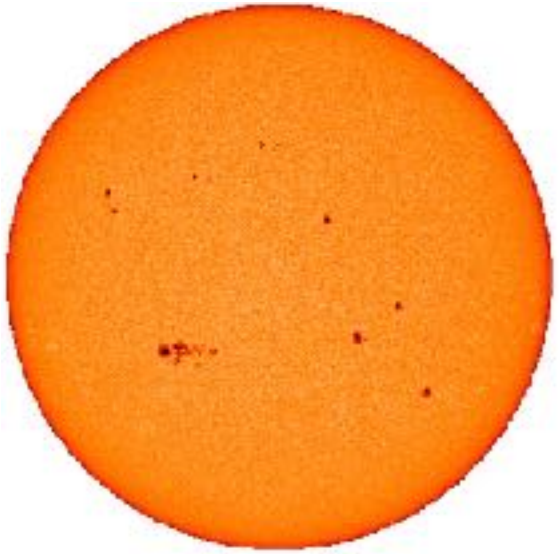

Granulation: Swedish Telescope, V. Henriques

Full Sun: SDO/HMI continuum

Accounting for stellar activity (or any kind of unknown correlated systematics) in RV analyses

- The main **goal** of RV follow-up is to determine accurate, precise planet masses, in the presence of stellar activity

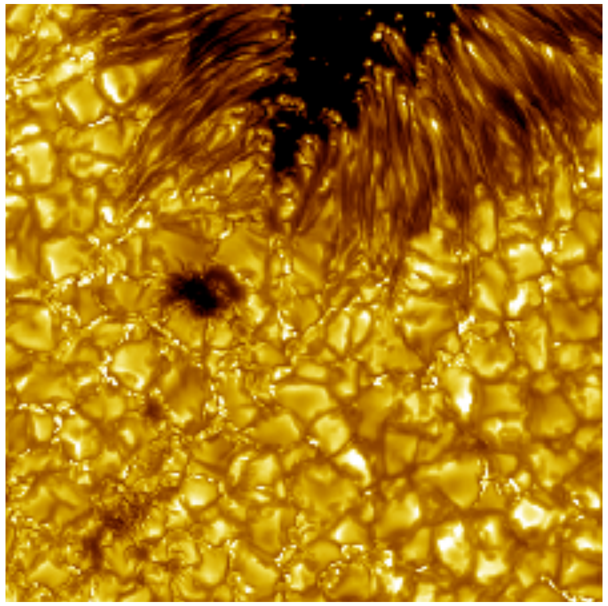

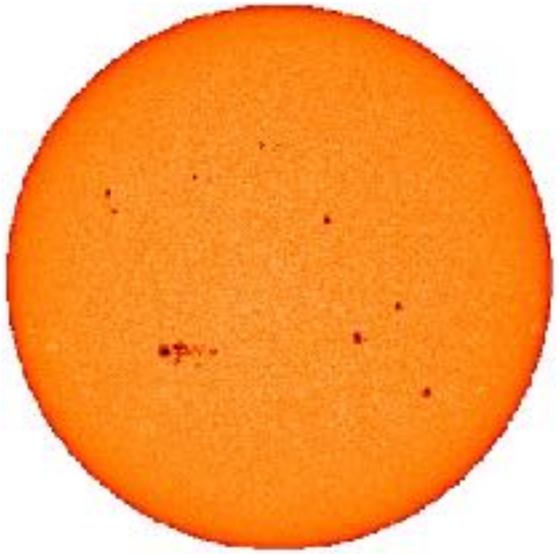

Granulation: Swedish Telescope, V. Henriques

Full Sun: SDO/HMI continuum

Accounting for stellar activity (or any kind of unknown correlated systematics) in RV analyses

- The main **goal** of RV follow-up is to determine accurate, precise planet masses, in the presence of stellar activity
- Often, we do not know enough about the physical processes responsible for activity-driven RV variations, so we cannot **model** them.

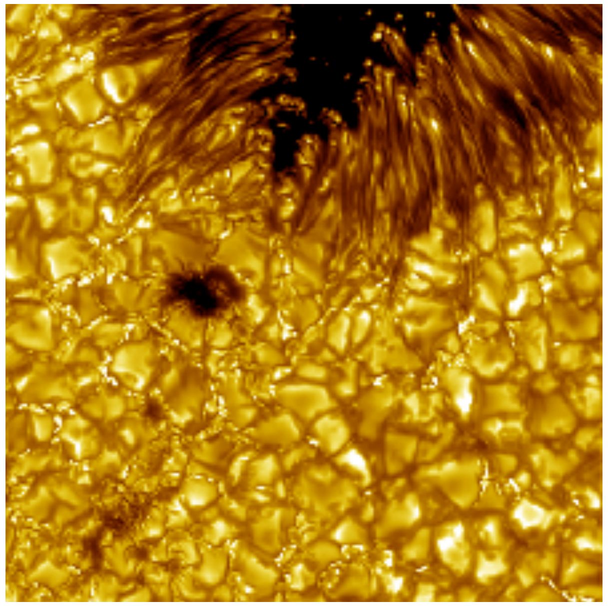

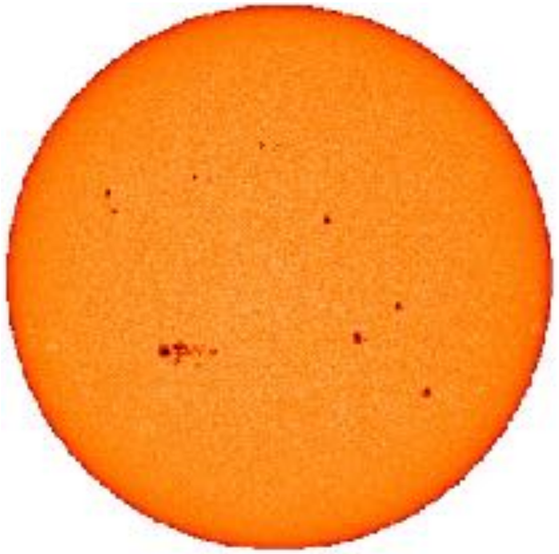

Granulation: Swedish Telescope, V. Henriques

Full Sun: SDO/HMI continuum

Accounting for stellar activity (or any kind of unknown correlated systematics) in RV analyses

- The main **goal** of RV follow-up is to determine accurate, precise planet masses, in the presence of stellar activity
- Often, we do not know enough about the physical processes responsible for activity-driven RV variations, so we cannot **model** them.
- But if we do not account for activity: our mass determinations can be biased, *i.e* inaccurate.

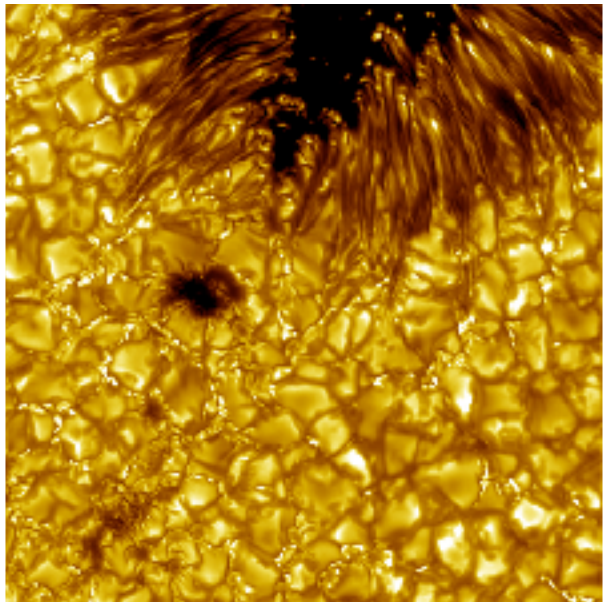

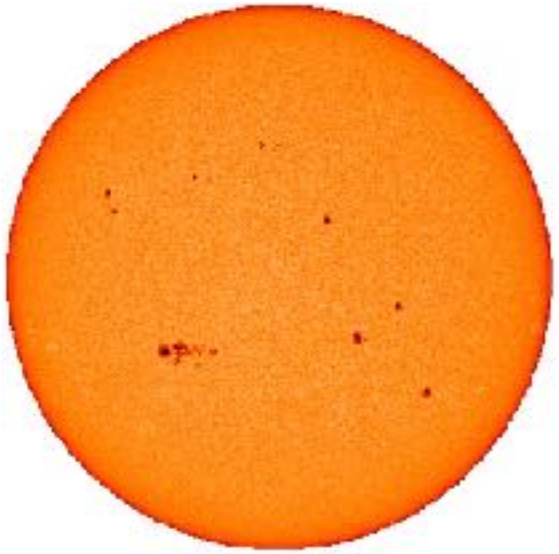

Granulation: Swedish Telescope, V. Henriques

Full Sun: SDO/HMI continuum

Accounting for stellar activity (or any kind of unknown correlated systematics) in RV analyses

- The main **goal** of RV follow-up is to determine accurate, precise planet masses, in the presence of stellar activity
- Often, we do not know enough about the physical processes responsible for activity-driven RV variations, so we cannot **model** them.
- But if we do not account for activity: our mass determinations can be biased, *i.e* inaccurate.
- Solution: we account for the uncertainty induced by stellar activity by treating activity as **noise**.

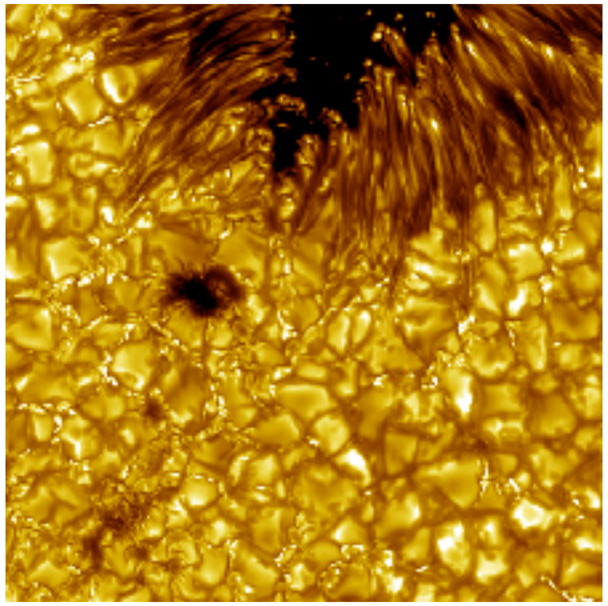

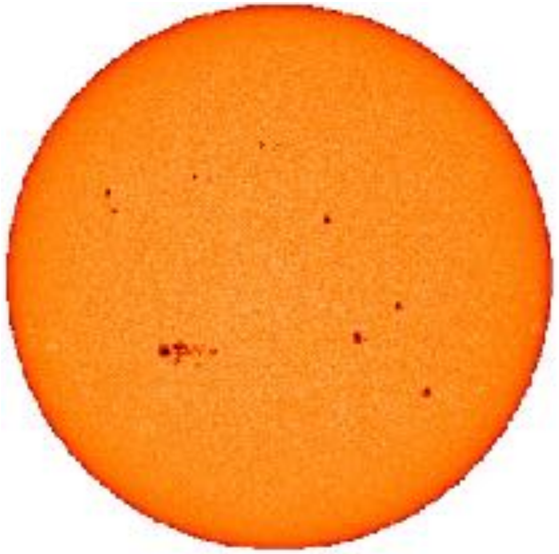

Granulation: Swedish Telescope, V. Henriques

Full Sun: SDO/HMI continuum

Accounting for stellar activity (or any kind of unknown correlated systematics) in RV analyses

- The main **goal** of RV follow-up is to determine accurate, precise planet masses, in the presence of stellar activity
- Often, we do not know enough about the physical processes responsible for activity-driven RV variations, so we cannot **model** them.
- But if we do not account for activity: our mass determinations can be biased, *i.e* inaccurate.
- Solution: we account for the uncertainty induced by stellar activity by treating activity as **noise**.
- Stars are rotating and their surfaces are constantly evolving: we must treat their activity as **correlated noise**.

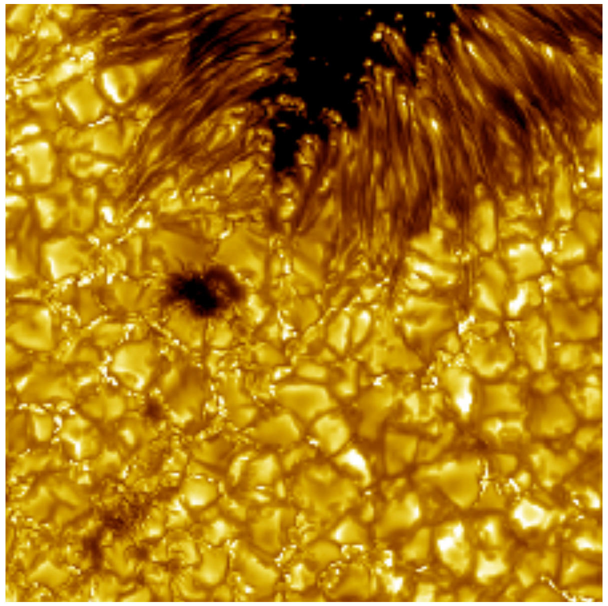

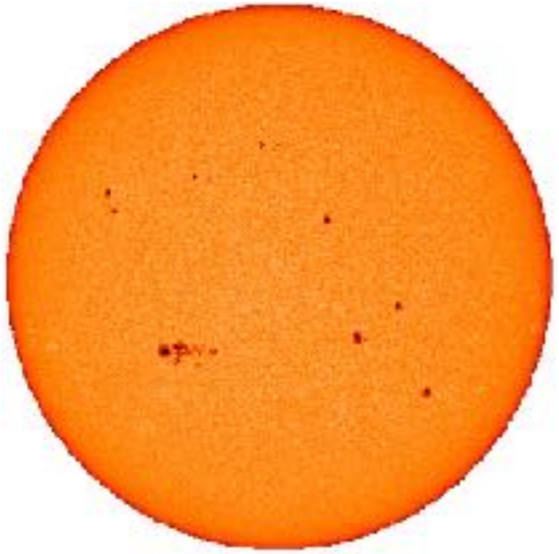

Granulation: Swedish Telescope, V. Henriques

Full Sun: SDO/HMI continuum

Outline

- Accuracy and precision
- Signal and noise
- Uncorrelated and correlated noise
- (A brief intro to) Gaussian process regression
- Astrophysically motivated GPs to account for stellar activity in RV observations
- Conclusions
- References for further reading/coding

Accuracy and precision of a parameter estimate

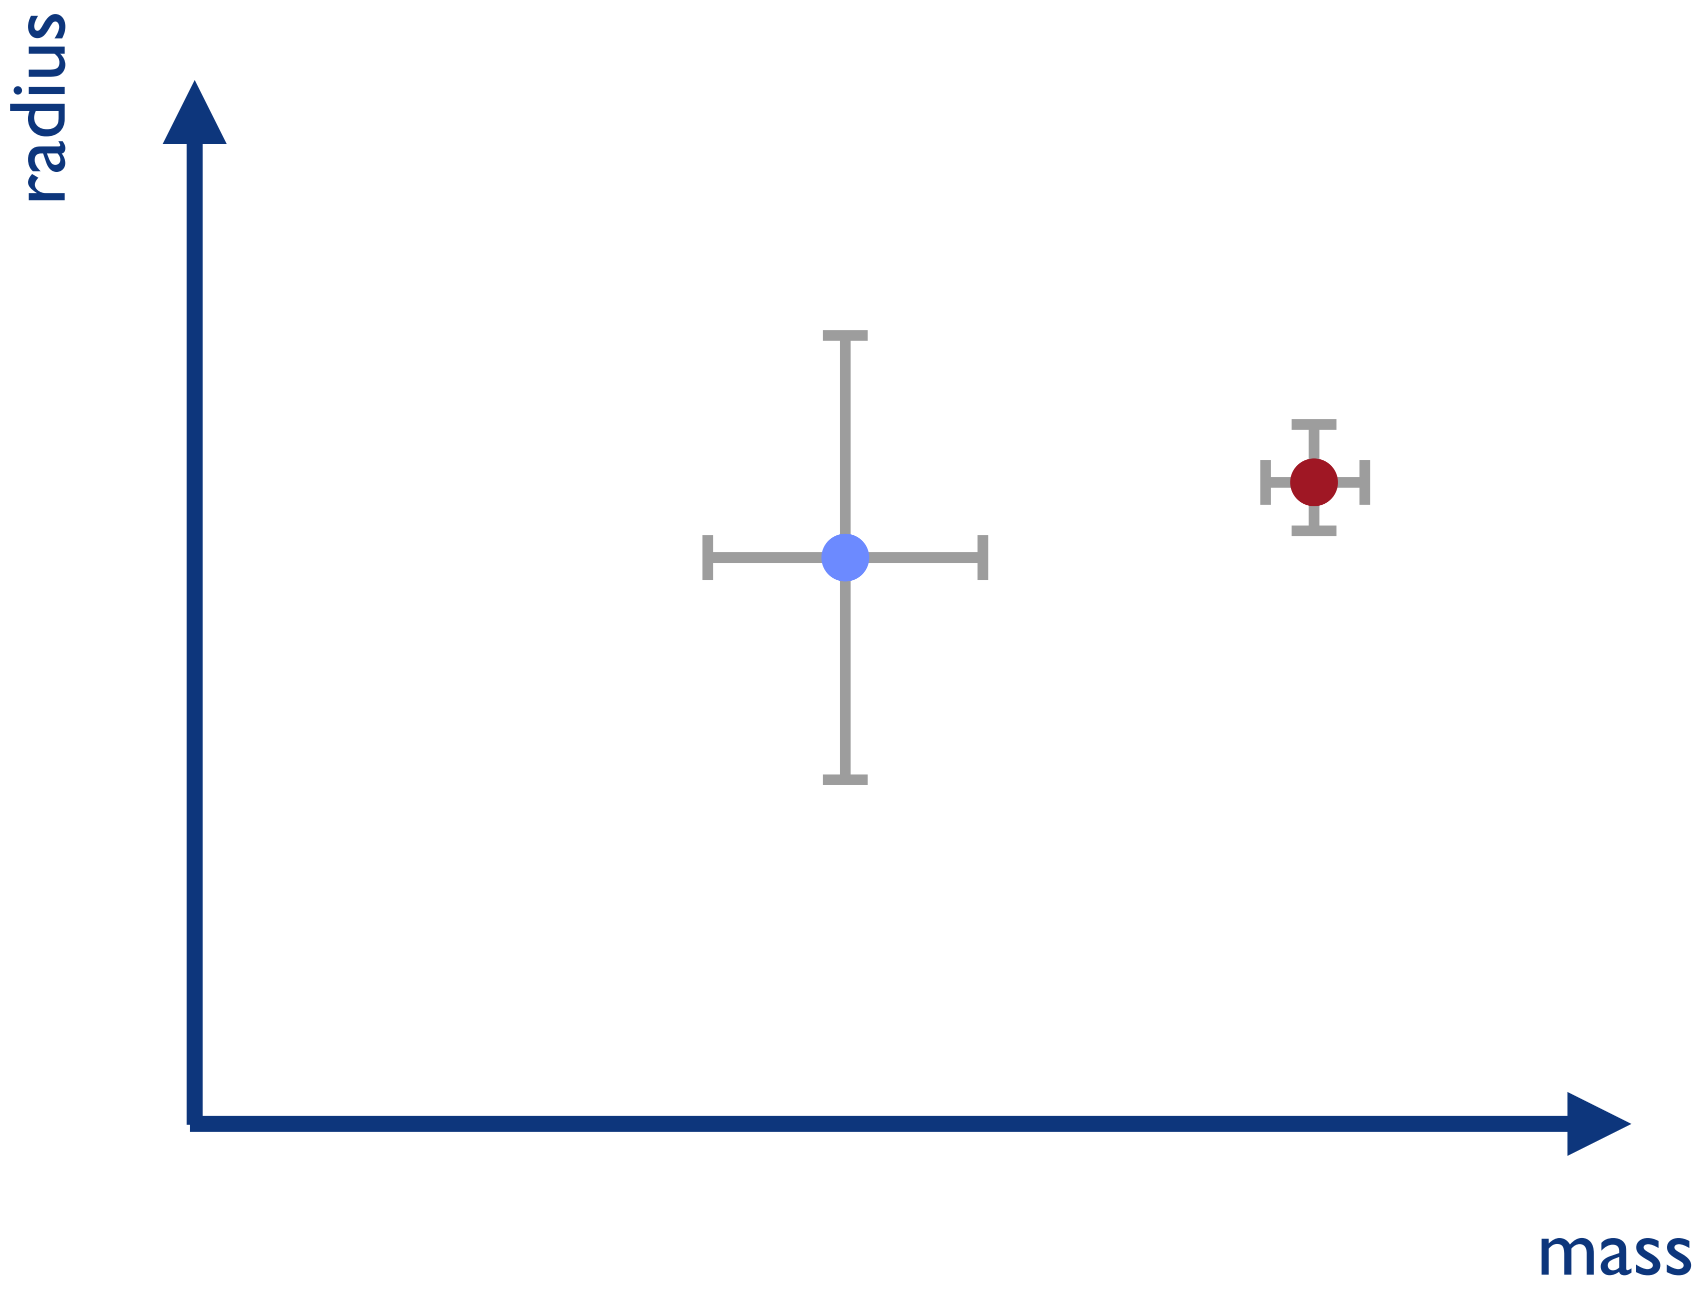

Refer back to Tom Loredo's talk yesterday

Accuracy and precision of a parameter estimate

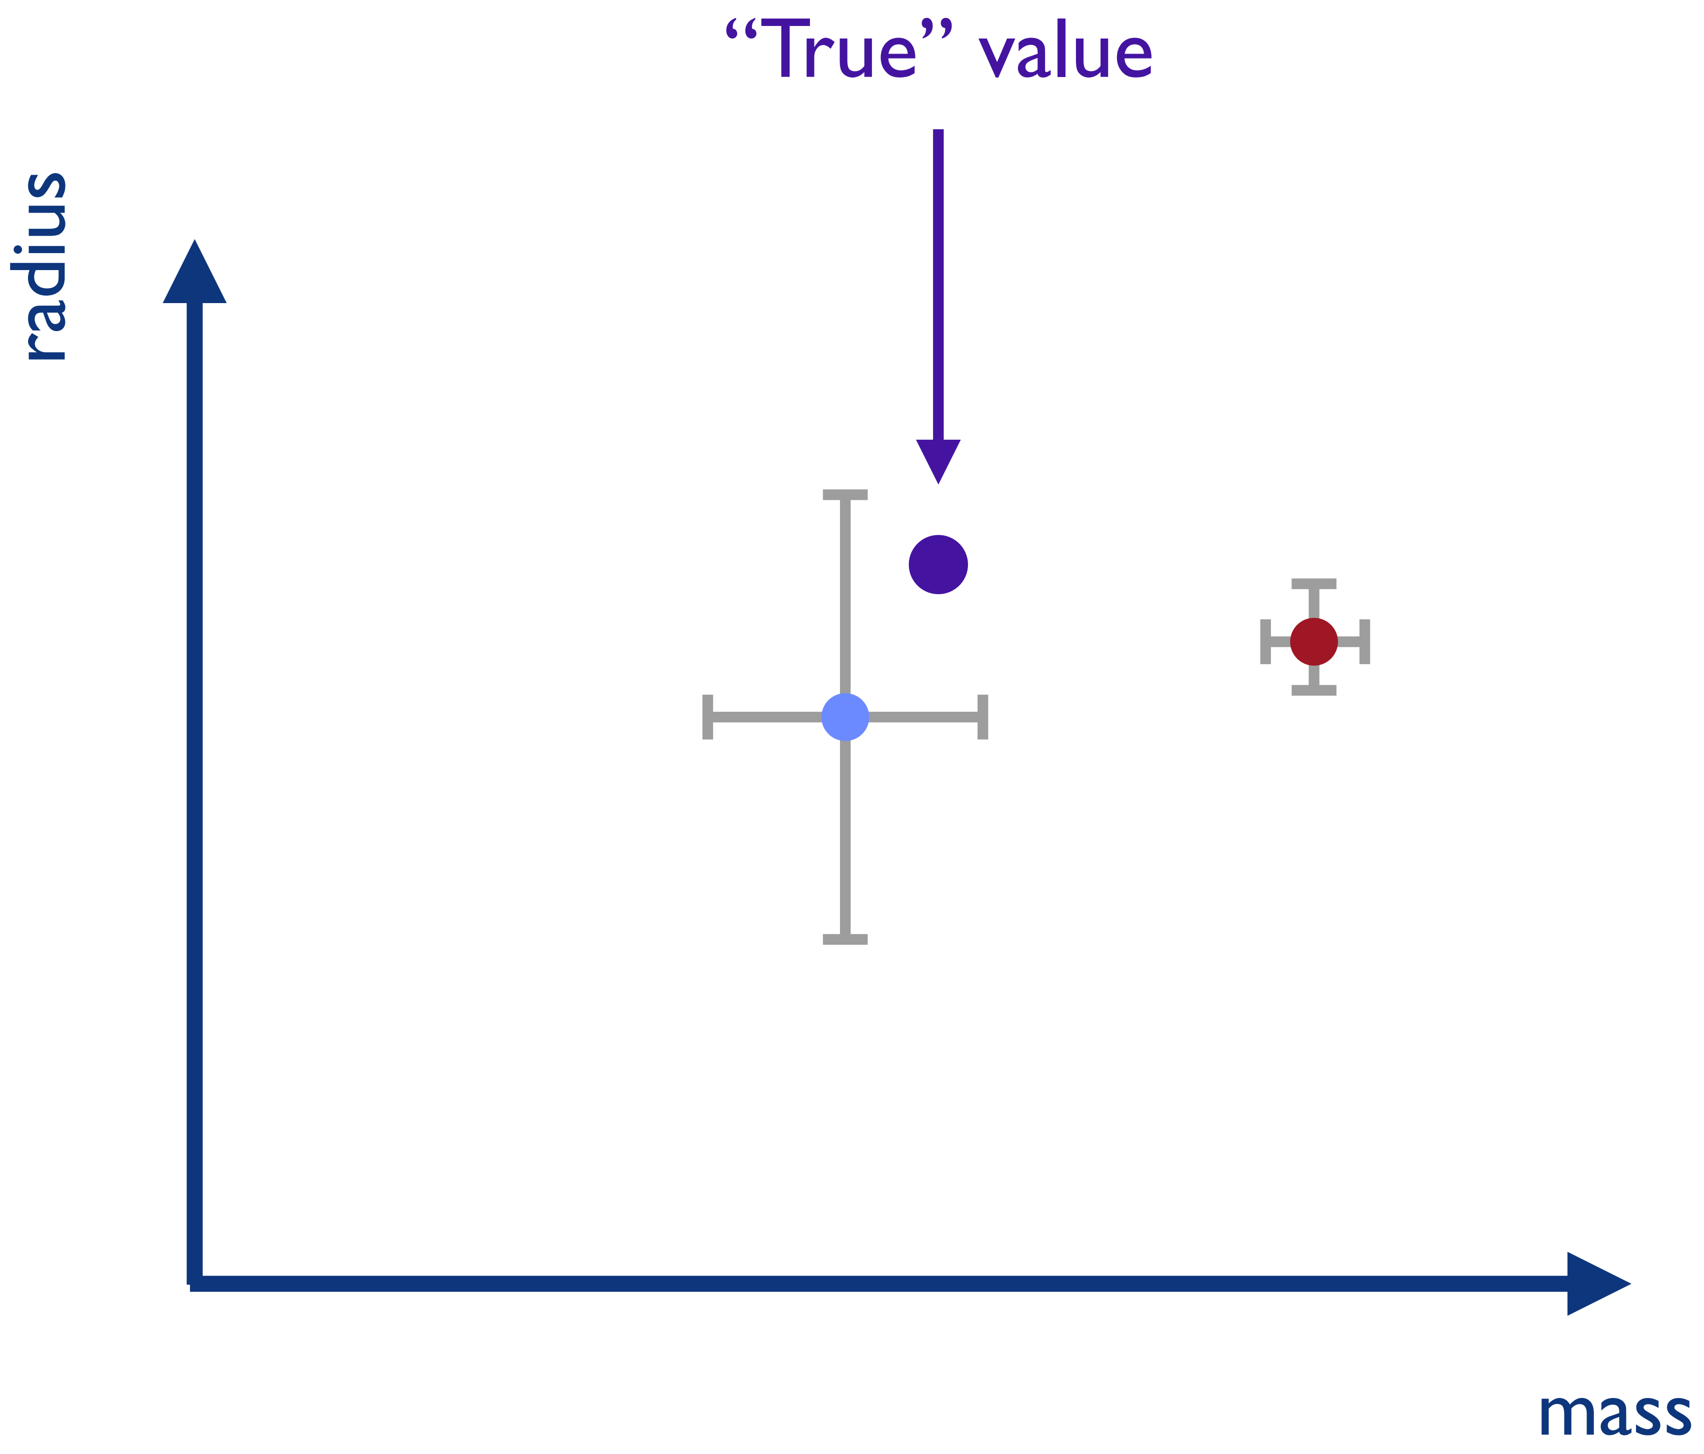

Refer back to Tom Loredo’s talk yesterday

Accuracy and precision of a parameter estimate

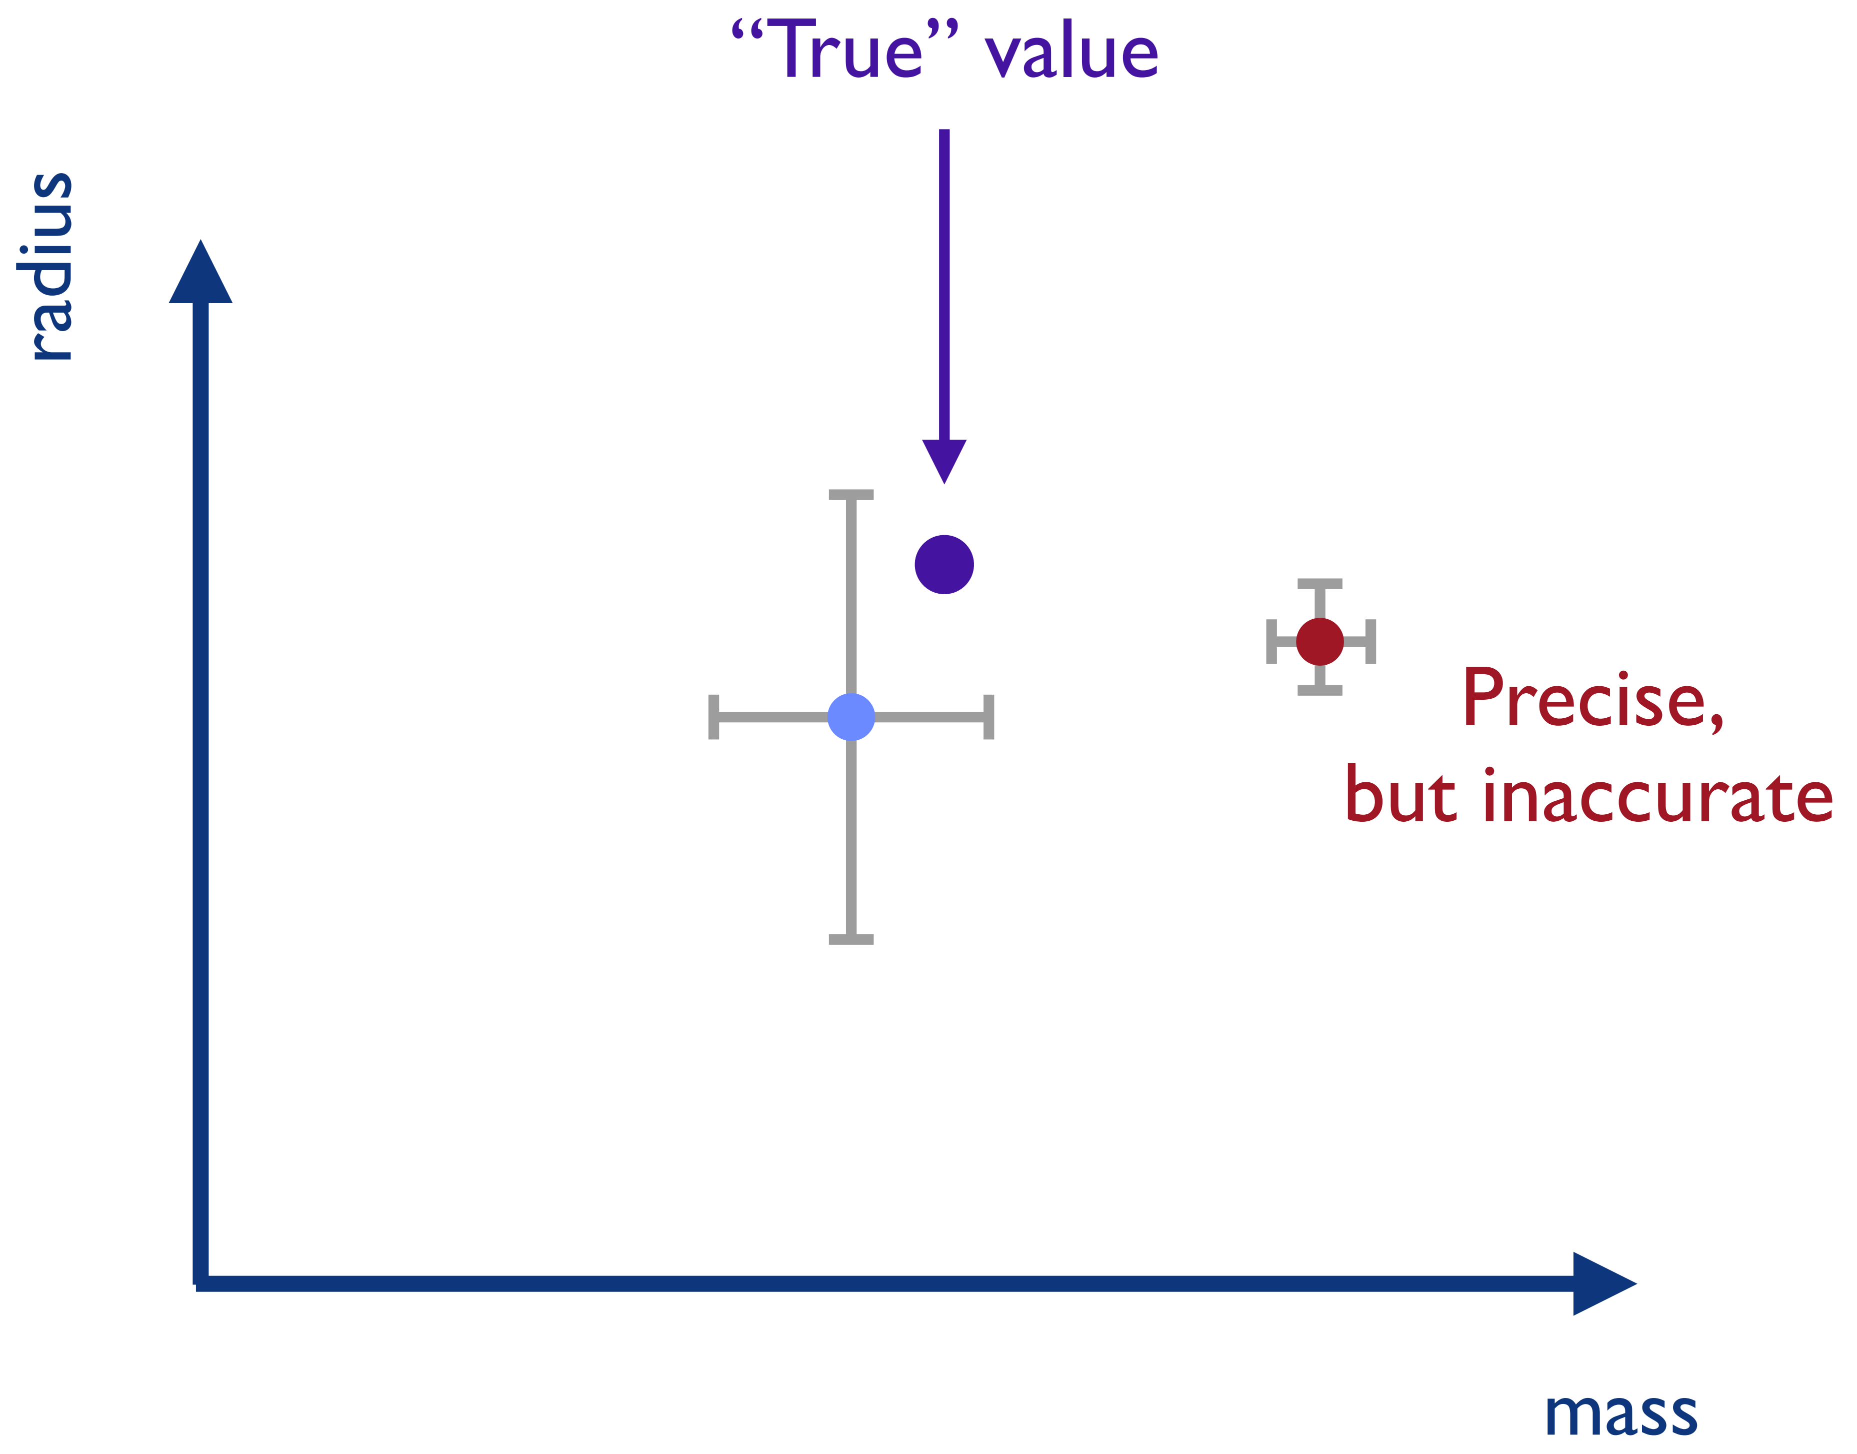

Refer back to Tom Loredo's talk yesterday

Accuracy and precision of a parameter estimate

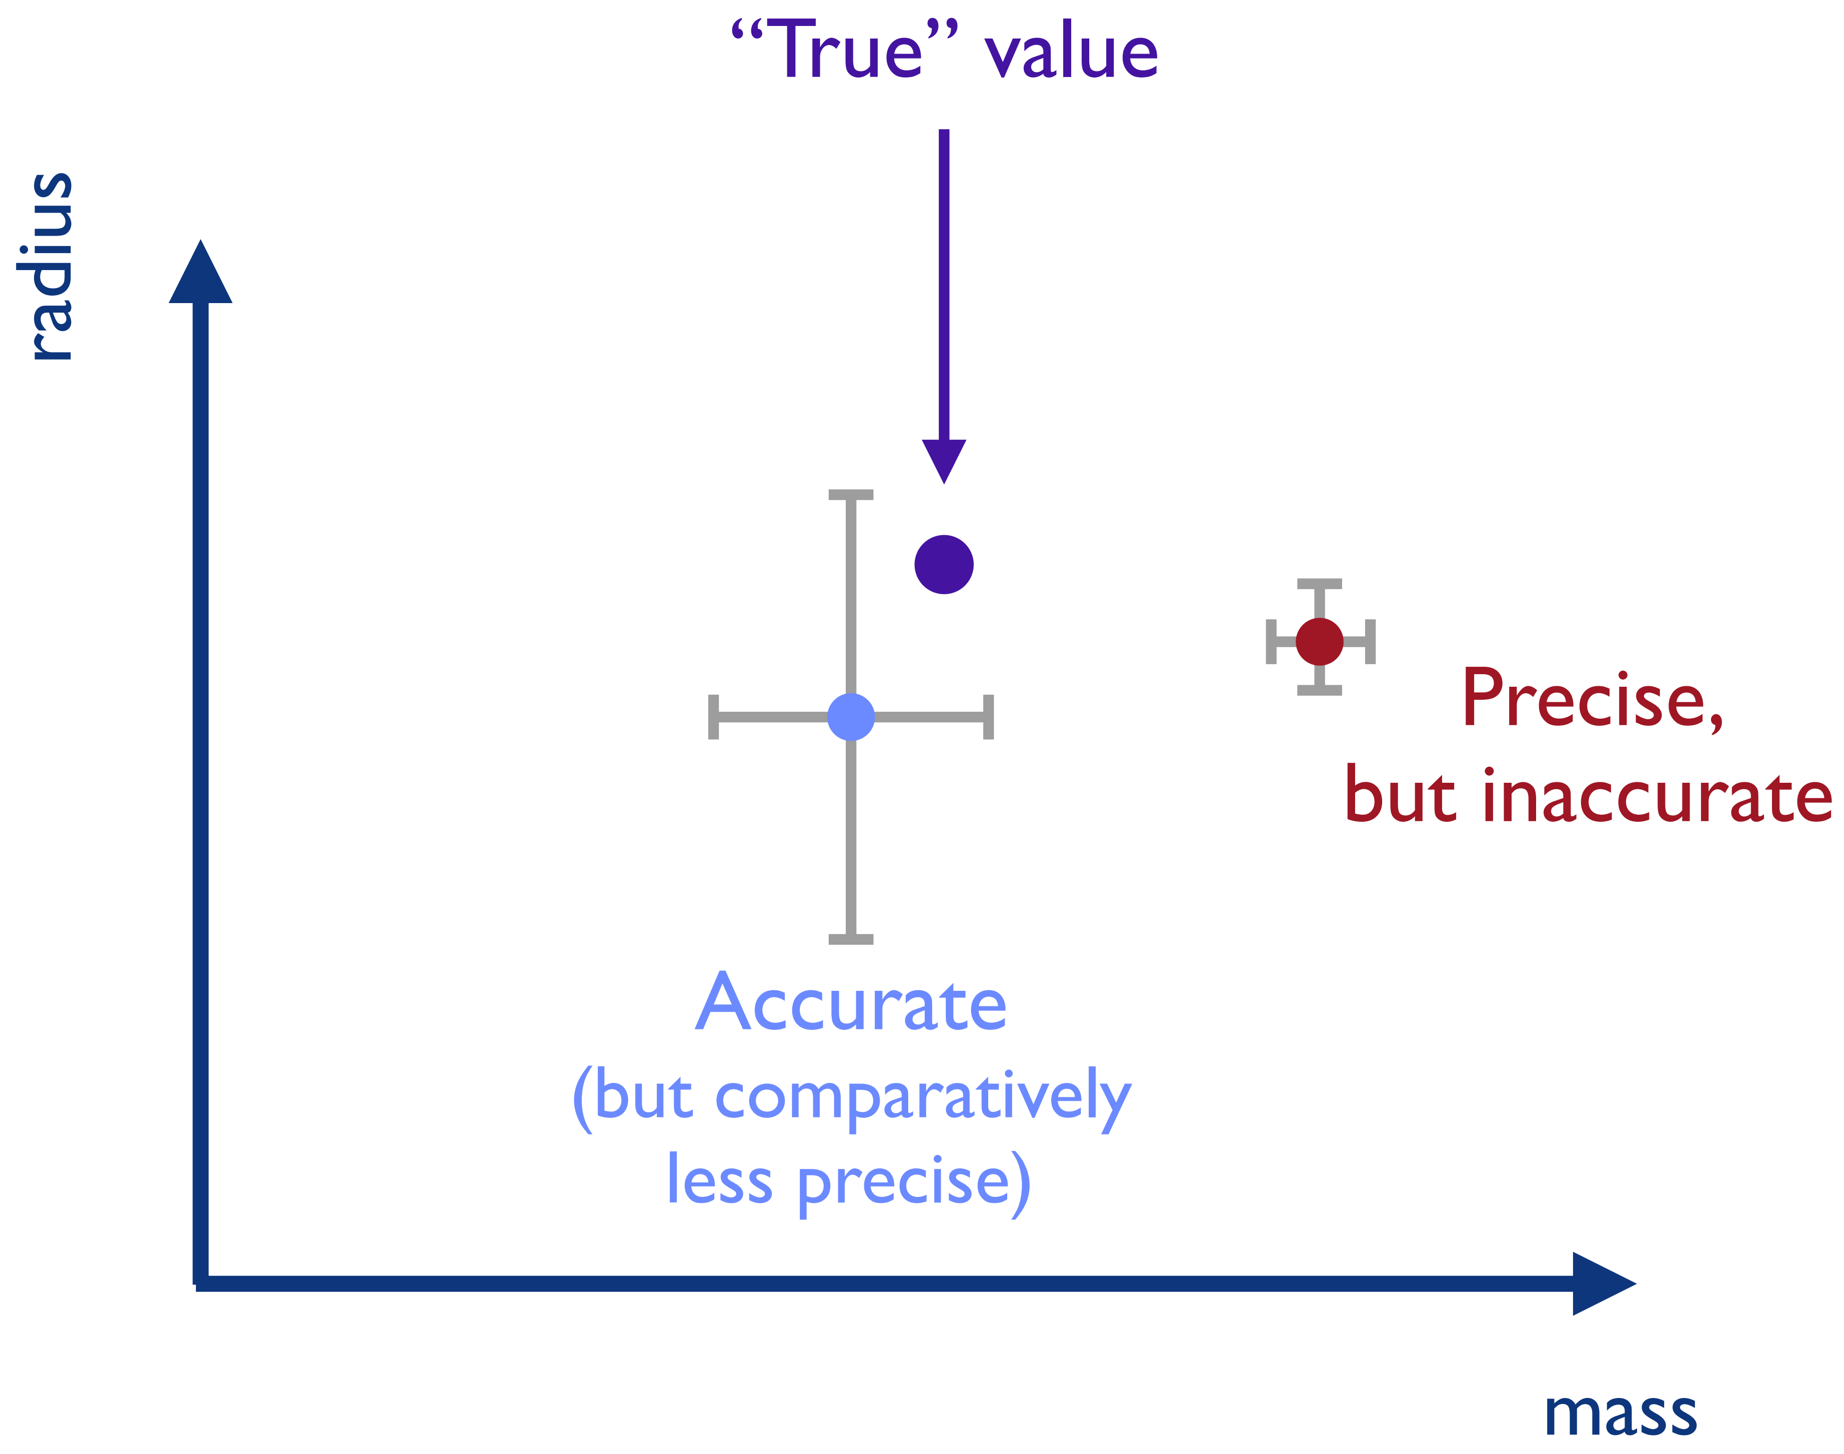

Refer back to Tom Loredo's talk yesterday

Accuracy and precision of a parameter estimate

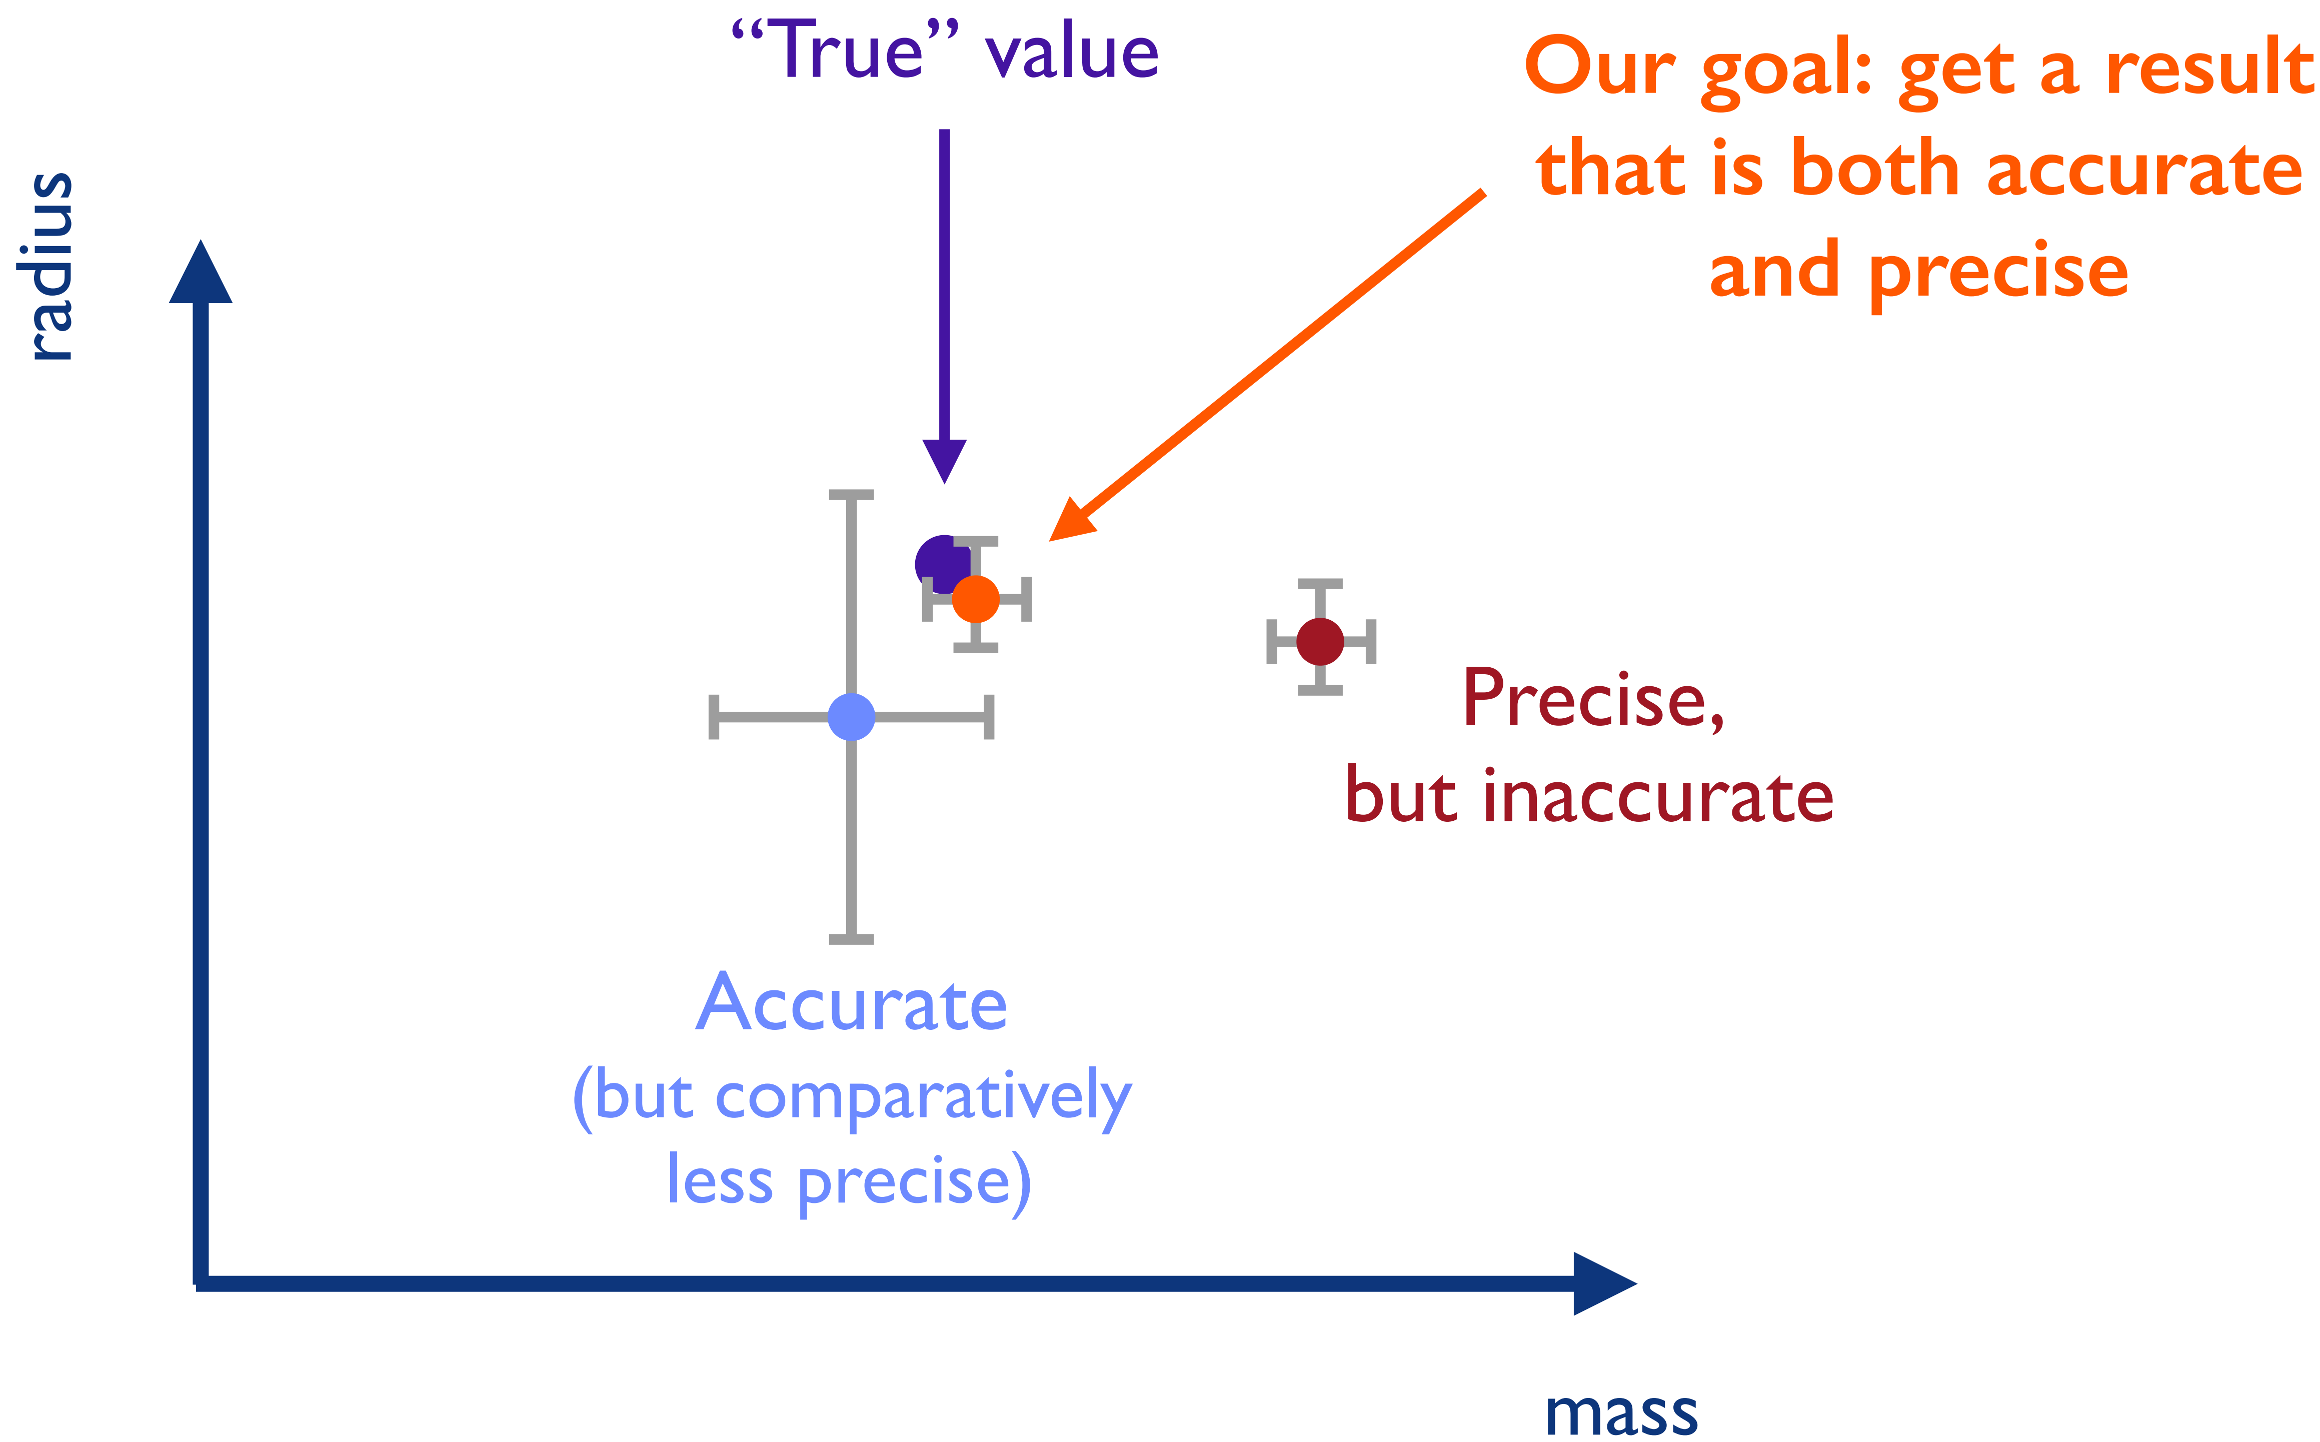

Refer back to Tom Loredo's talk yesterday

Signal and noise

Refer back to Tom Loredo's talk yesterday

Signal and noise

Any value obtained by an observation (eg. RVs) contains two components:

Refer back to Tom Loredo's talk yesterday

Signal and noise

Any value obtained by an observation (eg. RVs) contains two components:

1. **Signal:** carries the information of interest (eg. the planet's reflex motion)

Refer back to Tom Loredo's talk yesterday

Signal and noise

Any value obtained by an observation (eg. RVs) contains two components:

1. **Signal:** carries the information of interest (eg. the planet's reflex motion)

2. **Noise:** unknown processes that we do not understand (eg. stellar activity). Can either be random (uncorrelated), or non-random (correlated).

Refer back to Tom Loredo's talk yesterday

Signal and noise

Any value obtained by an observation (eg. RVs) contains two components:

1. **Signal:** carries the information of interest (eg. the planet's reflex motion) → We account for **signals** in our **model** (eg. a keplerian)

2. **Noise:** unknown processes that we do not understand (eg. stellar activity). Can either be random (uncorrelated), or non-random (correlated).

Refer back to Tom Loredo's talk yesterday

Signal and noise

Any value obtained by an observation (eg. RVs) contains two components:

1. **Signal:** carries the information of interest (eg. the planet's reflex motion)

   → We account for **signals** in our **model** (eg. a keplerian)

2. **Noise:** unknown processes that we do not understand (eg. stellar activity). Can either be random (uncorrelated), or non-random (correlated).

   → We account for **noise** in our **goodness of fit** (the chi square or the likelihood)

Refer back to Tom Loredo's talk yesterday

Chi square (χ^2) and likelihood ($\mathcal{L}$)

Refer back to Tom Loredo's talk yesterday

Chi square (χ^2) and likelihood ($\mathcal{L}$)

No noise:

$$\chi^2 = \frac{(\text{data-model})^2}{\text{errorbars}^2}$$

Refer back to Tom Loredo's talk yesterday

Chi square (χ^2) and likelihood ($\mathcal{L}$)

No noise:

$$\chi^2 = \frac{(\text{data-model})^2}{\text{errorbars}^2}$$

Uncorrelated noise
(Gaussian-distributed, "white" noise, "jitter"):

$$\chi^2 = \frac{(\text{data-model})^2}{\sqrt{(\text{errorbars}^2 + \text{jitterterm}^2)}}$$

Refer back to Tom Loredo's talk yesterday

Chi square (χ^2) and likelihood ($\mathcal{L}$)

No noise:

$$\chi^2 = \frac{(\text{data-model})^2}{\text{errorbars}^2}$$

Uncorrelated noise
(Gaussian-distributed, "white" noise, "jitter"):

$$\chi^2 = \frac{(\text{data-model})^2}{\sqrt{(\text{errorbars}^2 + \text{jitterterm}^2)}}$$

$$\left.\right\} \mathcal{L} = \exp(-\chi^2/2)$$

Refer back to Tom Loredo's talk yesterday

Chi square (χ^2) and likelihood ($\mathcal{L}$)

No noise:

$$\chi^2 = \frac{(\text{data-model})^2}{\text{errorbars}^2}$$

Uncorrelated noise
(Gaussian-distributed, "white" noise, "jitter"):

$$\chi^2 = \frac{(\text{data-model})^2}{\sqrt{(\text{errorbars}^2 + \text{jitterterm}^2)}}$$

$$\mathcal{L} = \exp(-\chi^2 /2)$$

Correlated noise (non-random, "red" noise):

$$\mathcal{L} \neq \exp(-\chi^2 /2)$$

Refer back to Tom Loredo's talk yesterday

Uncorrelated noise ("white" noise, "jitter")

All data points are completely independent of each other

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Uncorrelated noise (“white” noise, “jitter”)

All data points are completely independent of each other

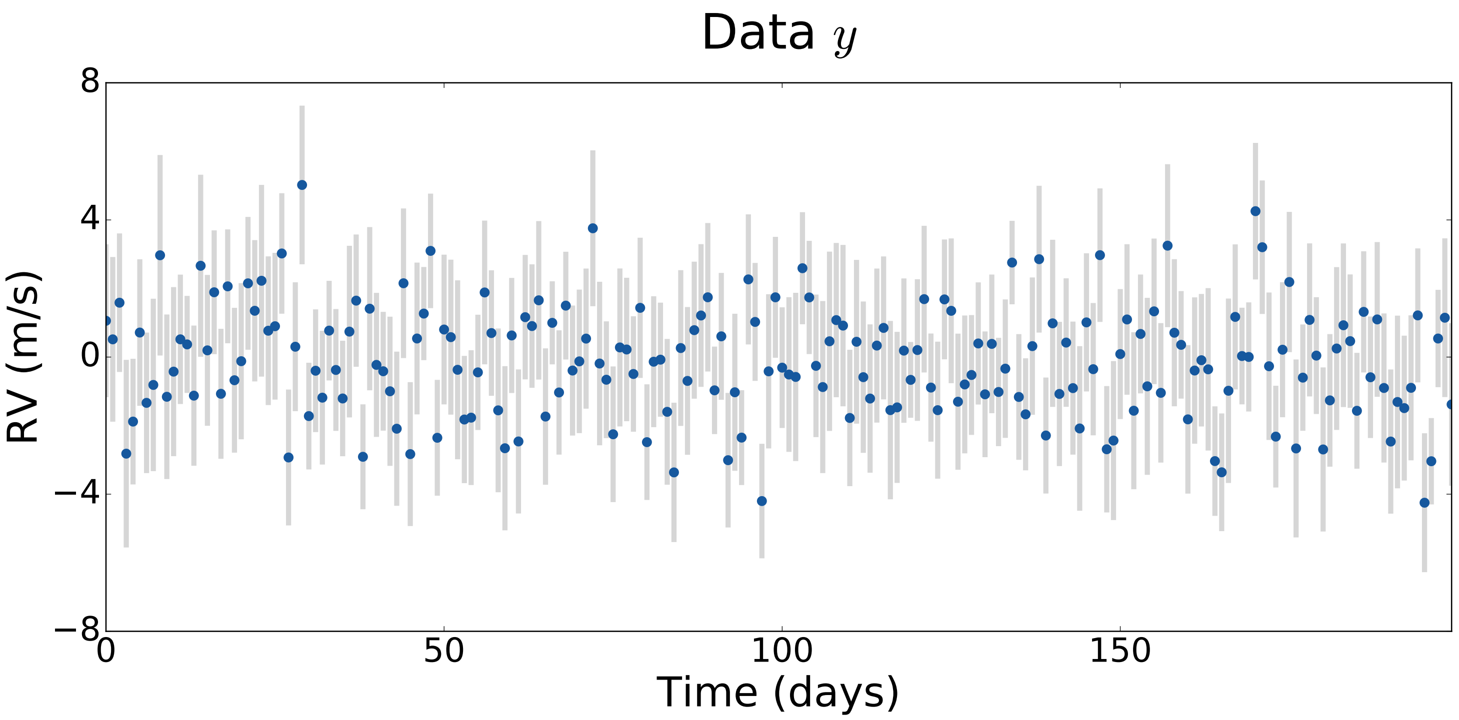

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Uncorrelated noise ("white" noise, "jitter")

All data points are completely independent of each other

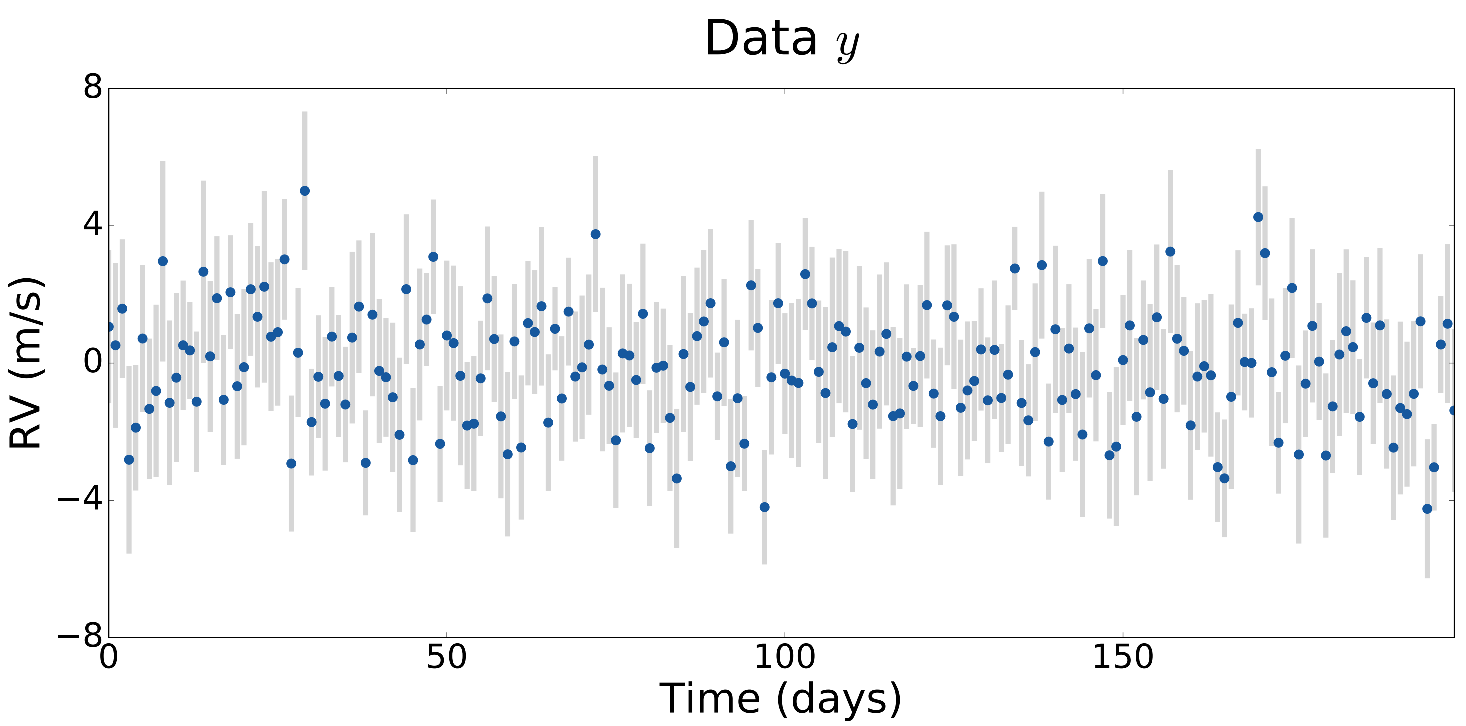

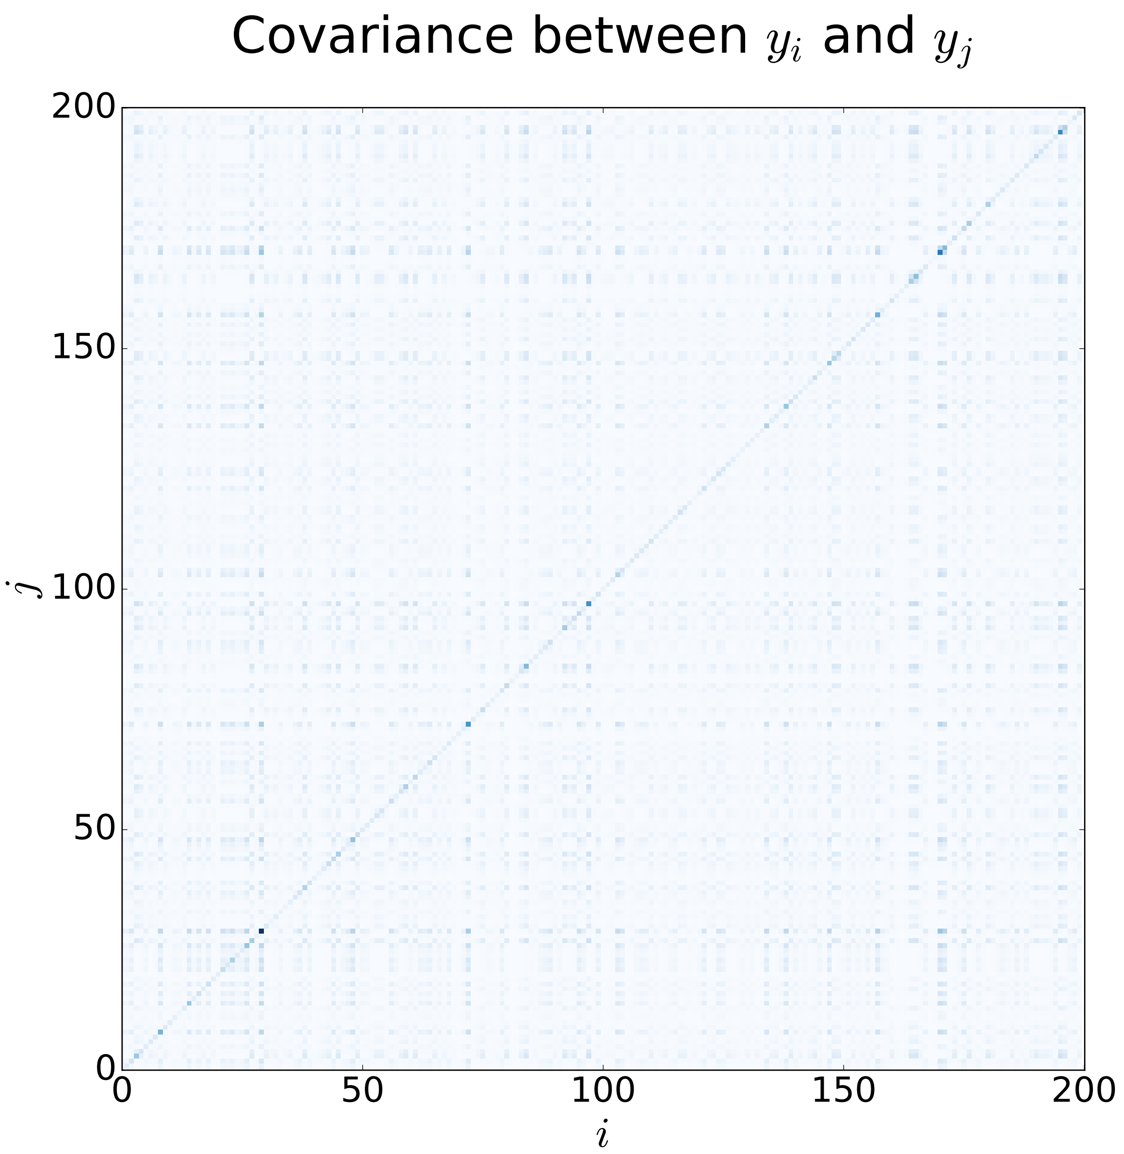

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Uncorrelated noise ("white" noise, "jitter")

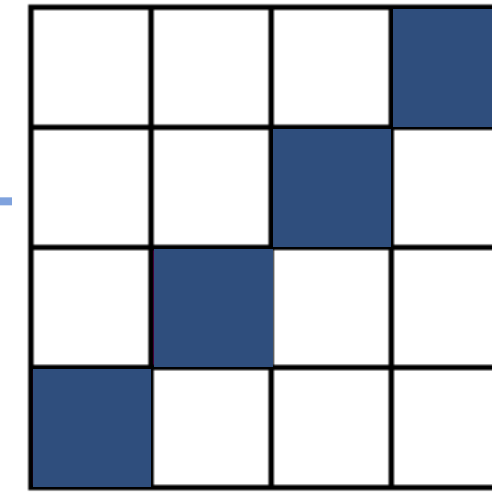

All data points are completely independent of each other

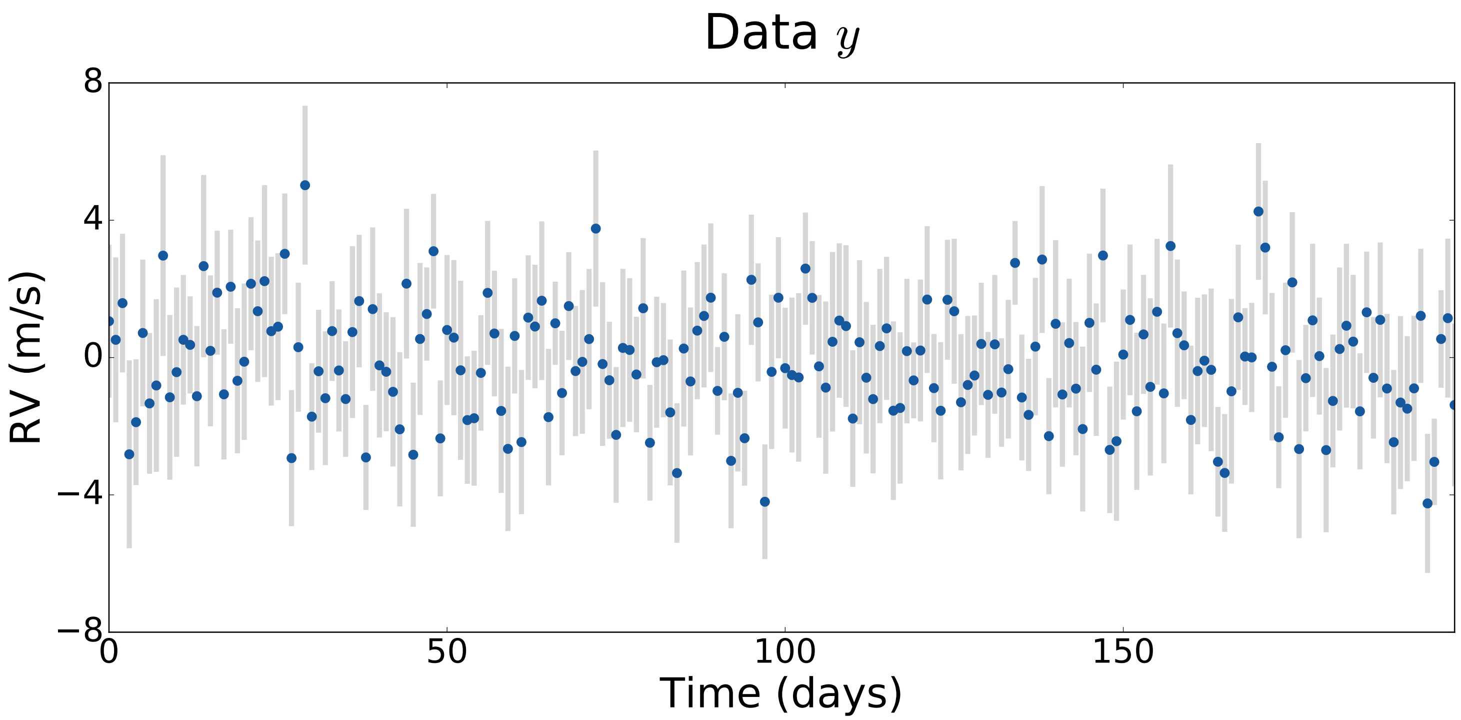

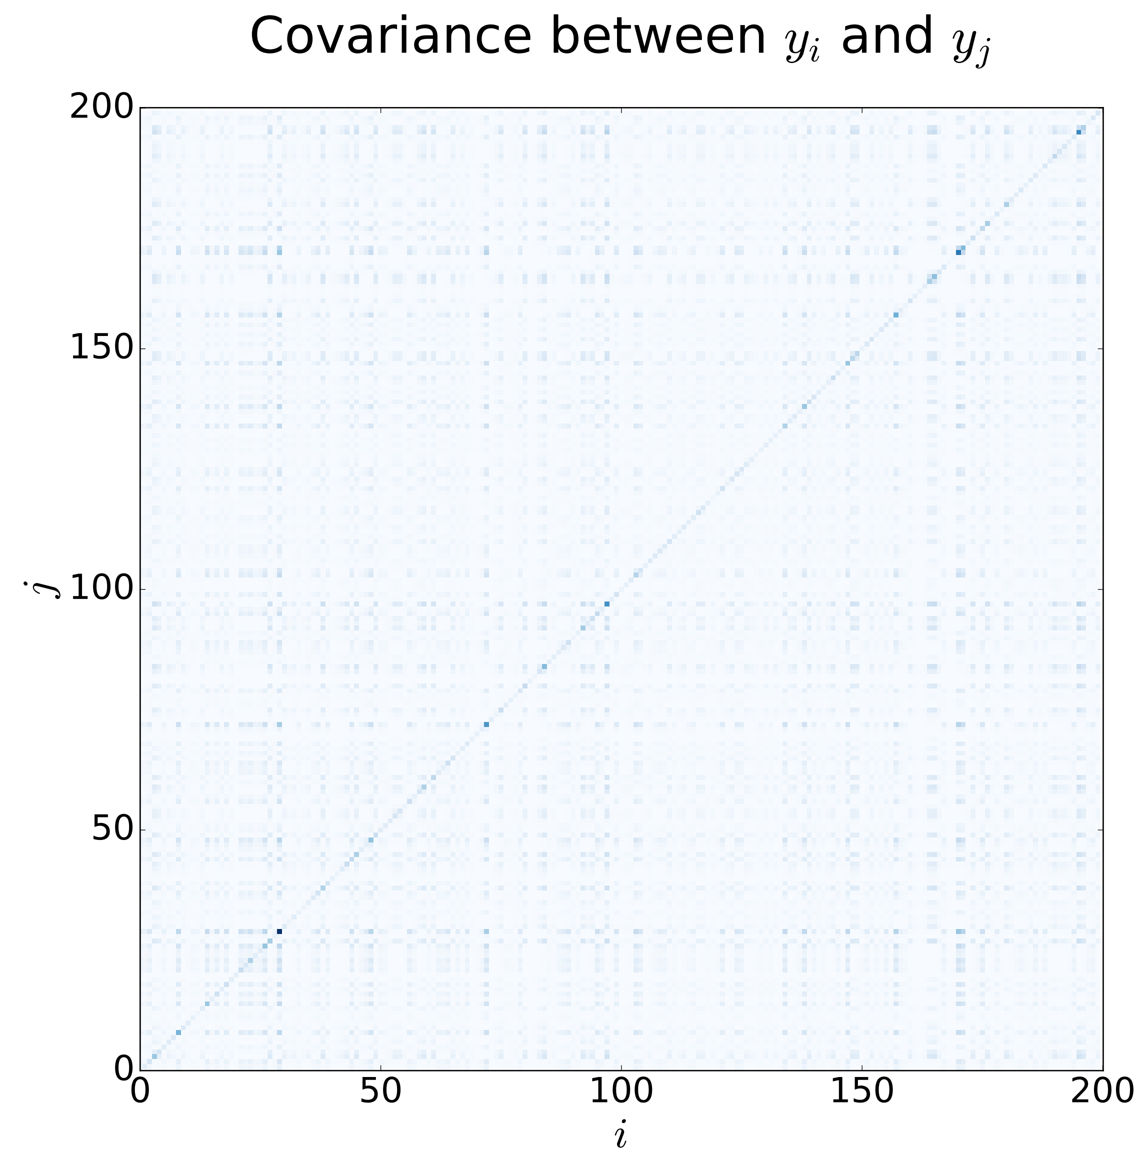

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Uncorrelated noise ("white" noise, "jitter")

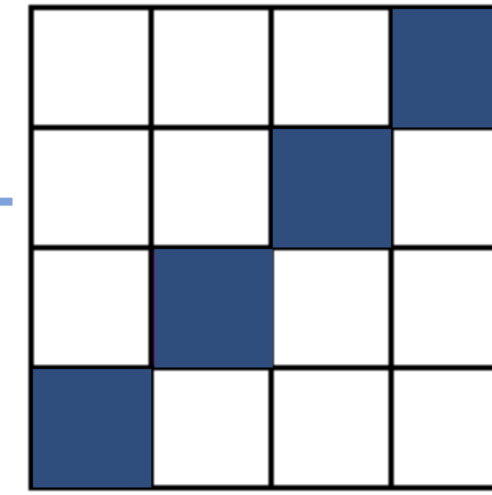

All data points are completely independent of each other

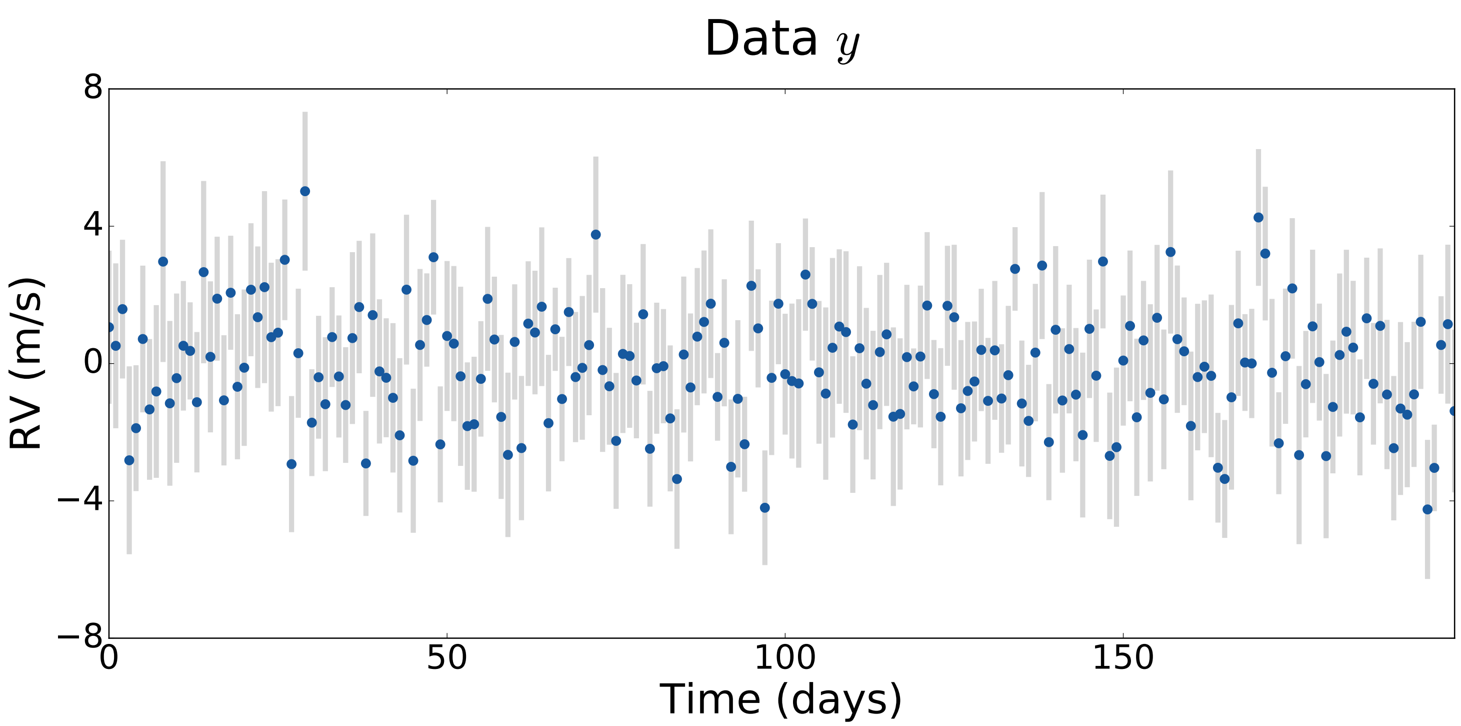

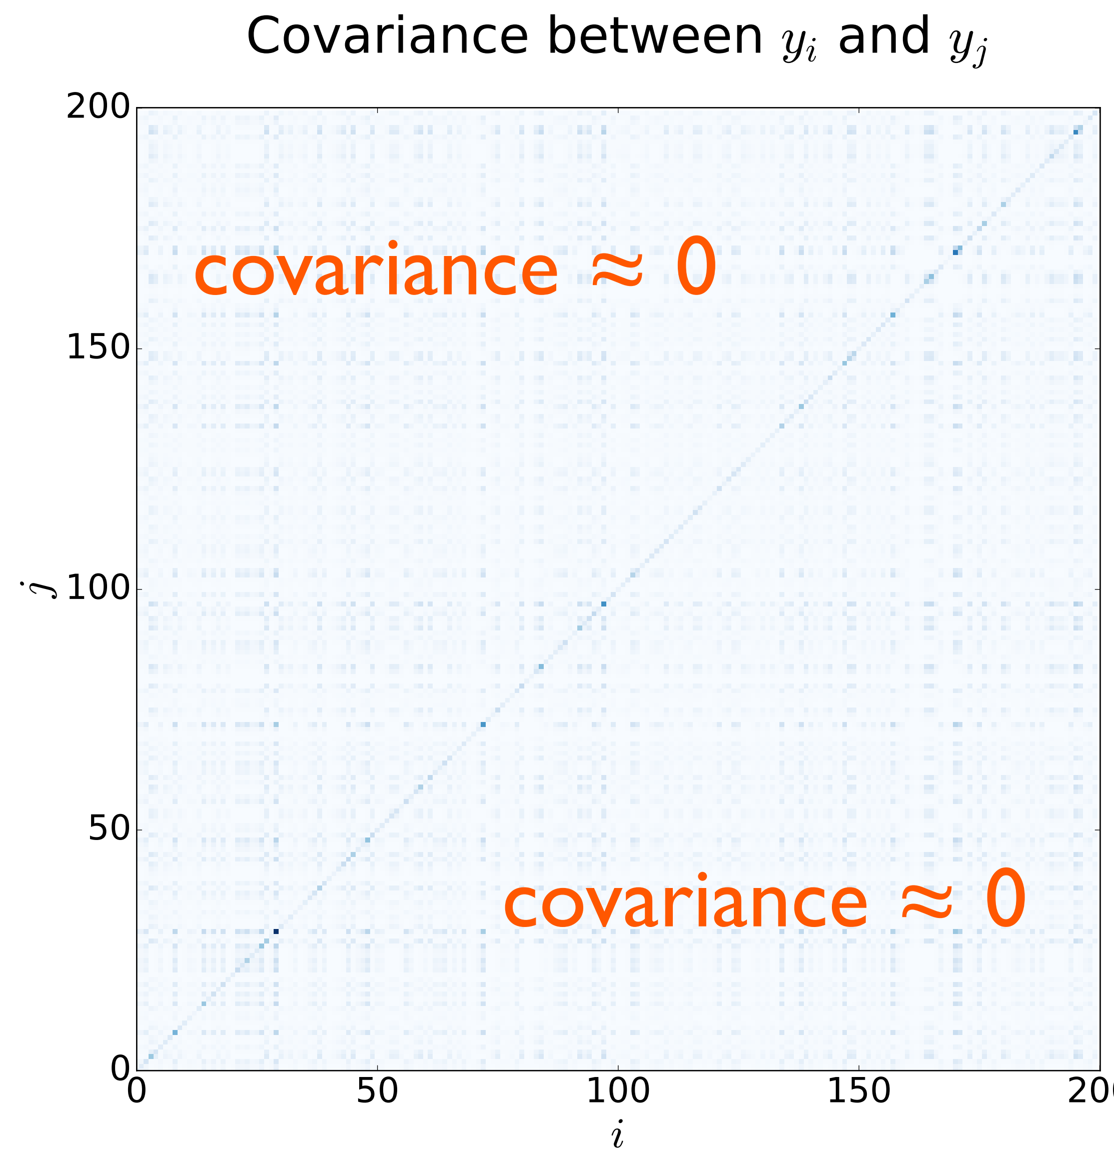

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Uncorrelated noise (“white” noise, “jitter”)

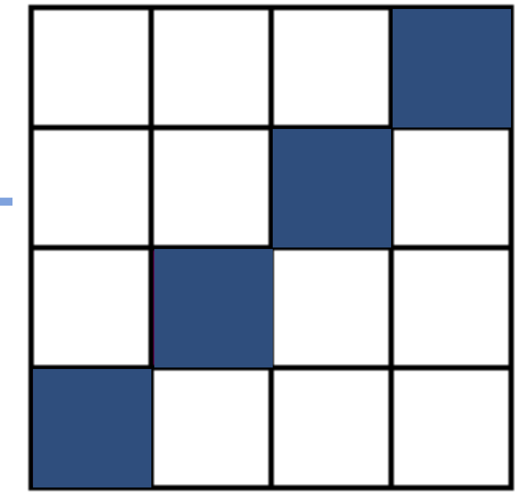

All data points are completely independent of each other

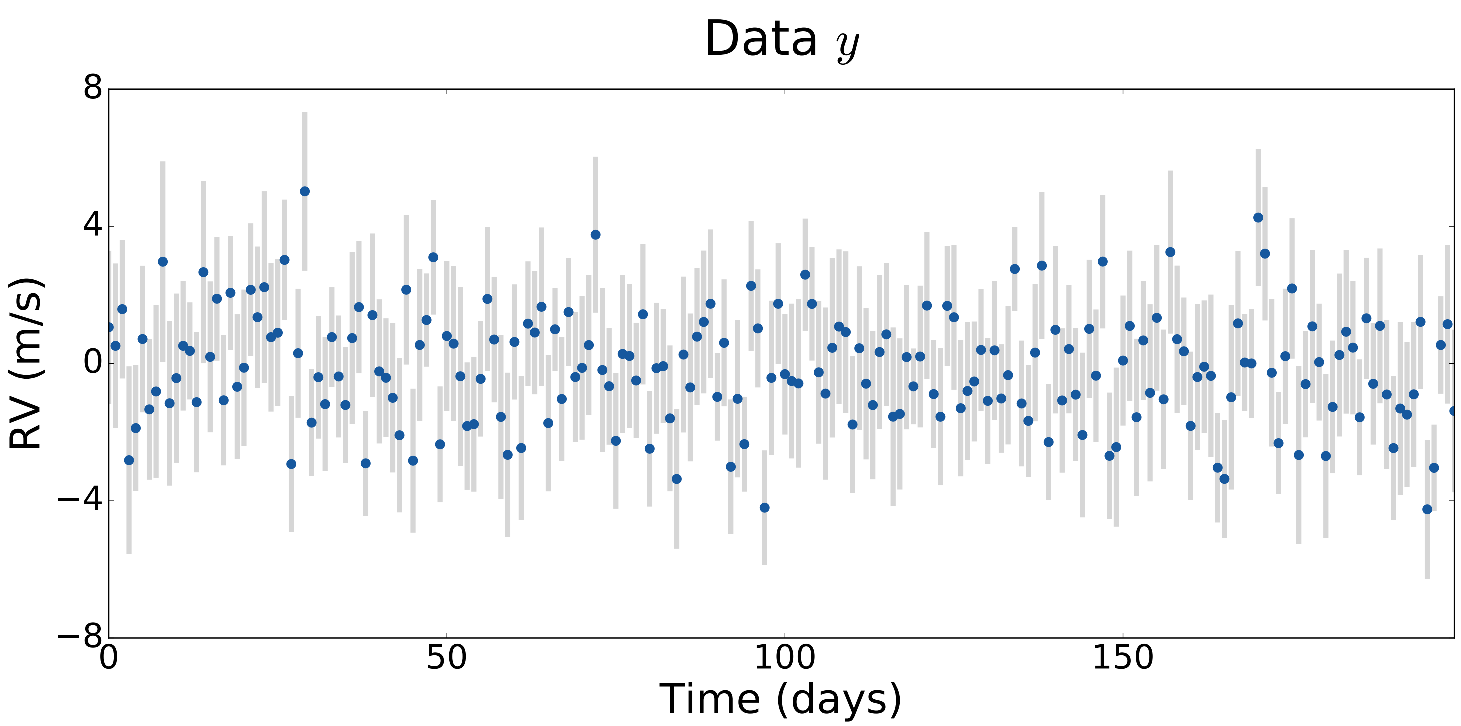

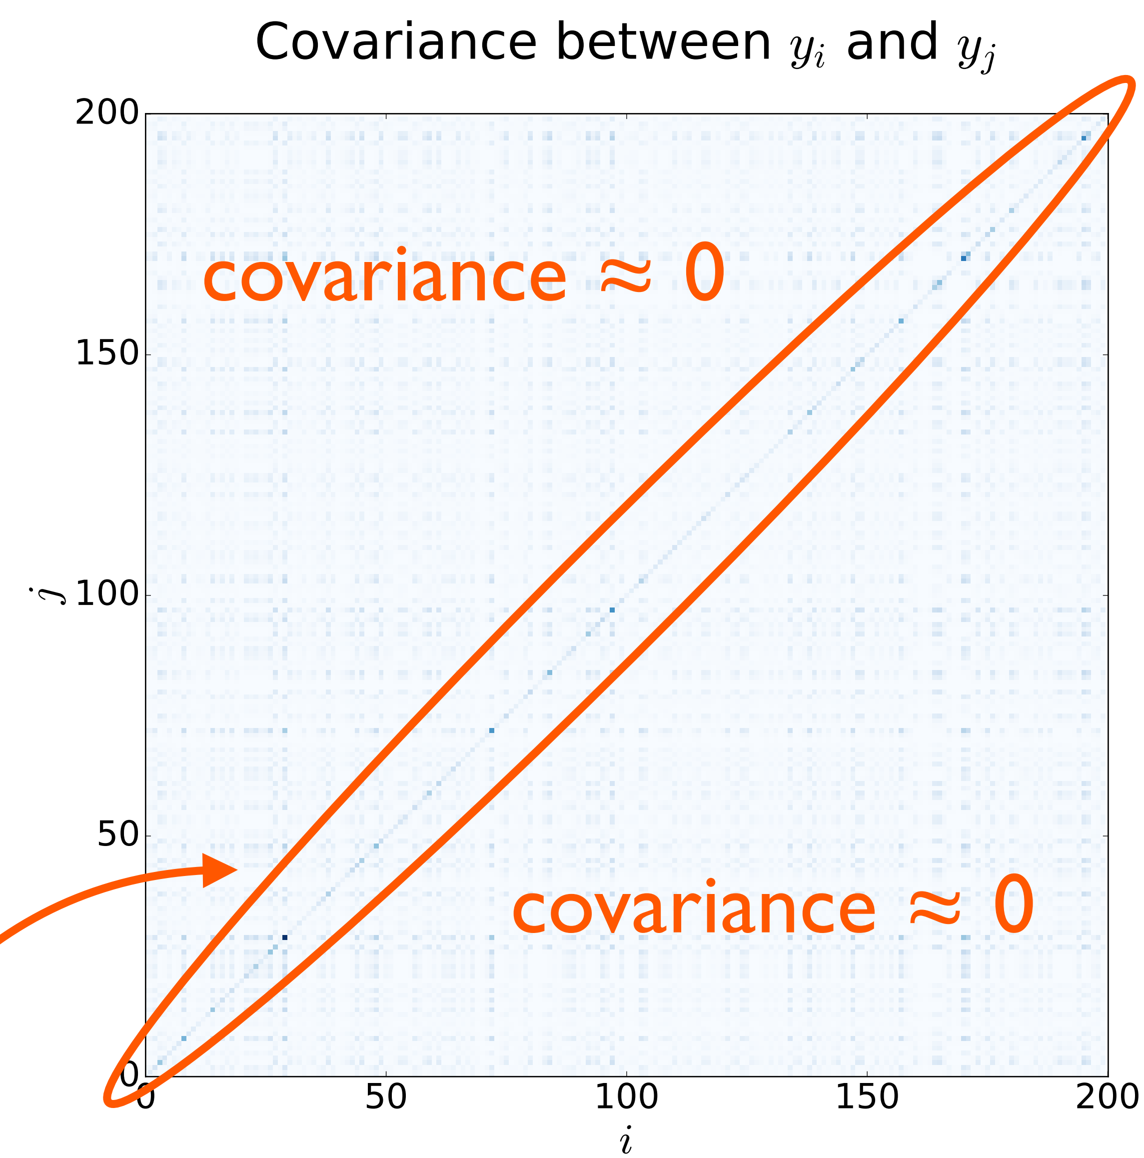

There is significant covariance only in the diagonal elements: they are the variance of each observation, *i.e.* their error bar!

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Correlated noise ("red" noise)

Data points are correlated with each other

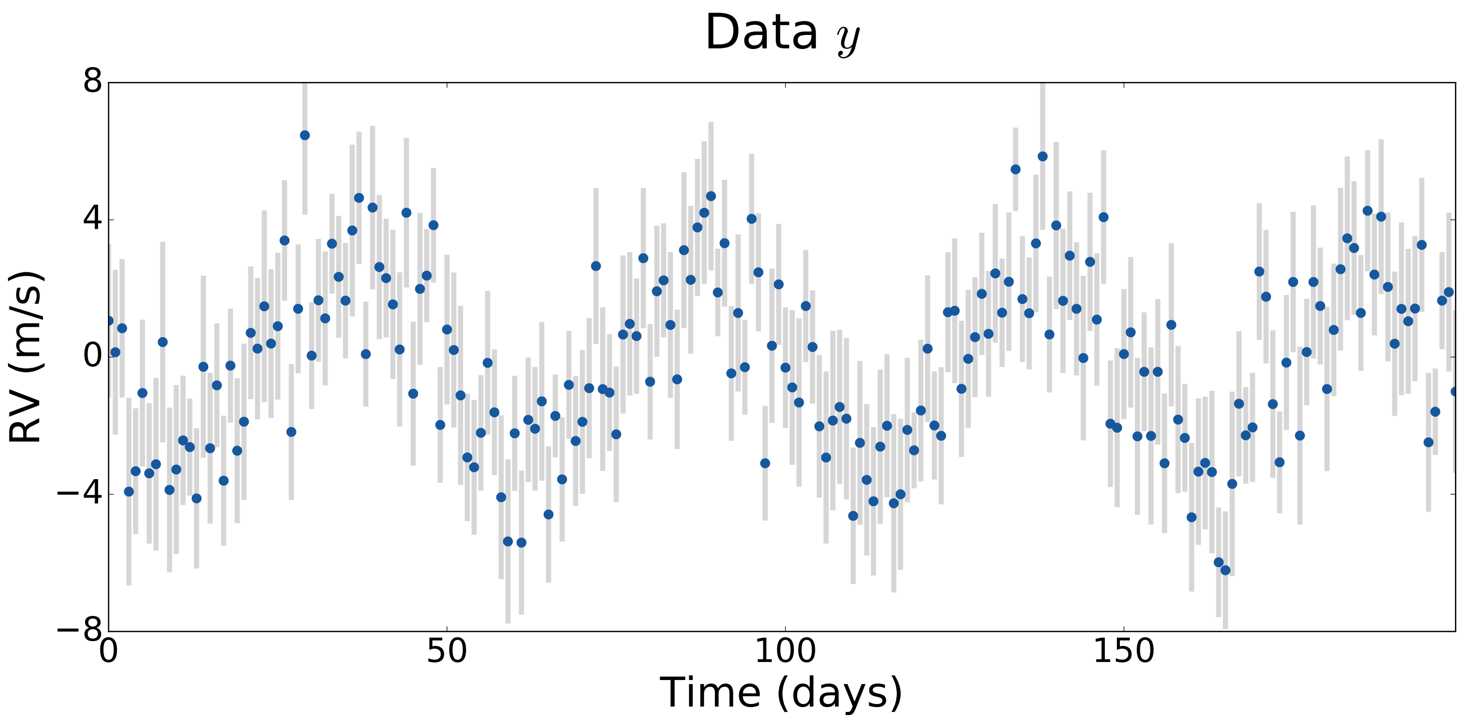

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Correlated noise ("red" noise)

Data points are correlated with each other

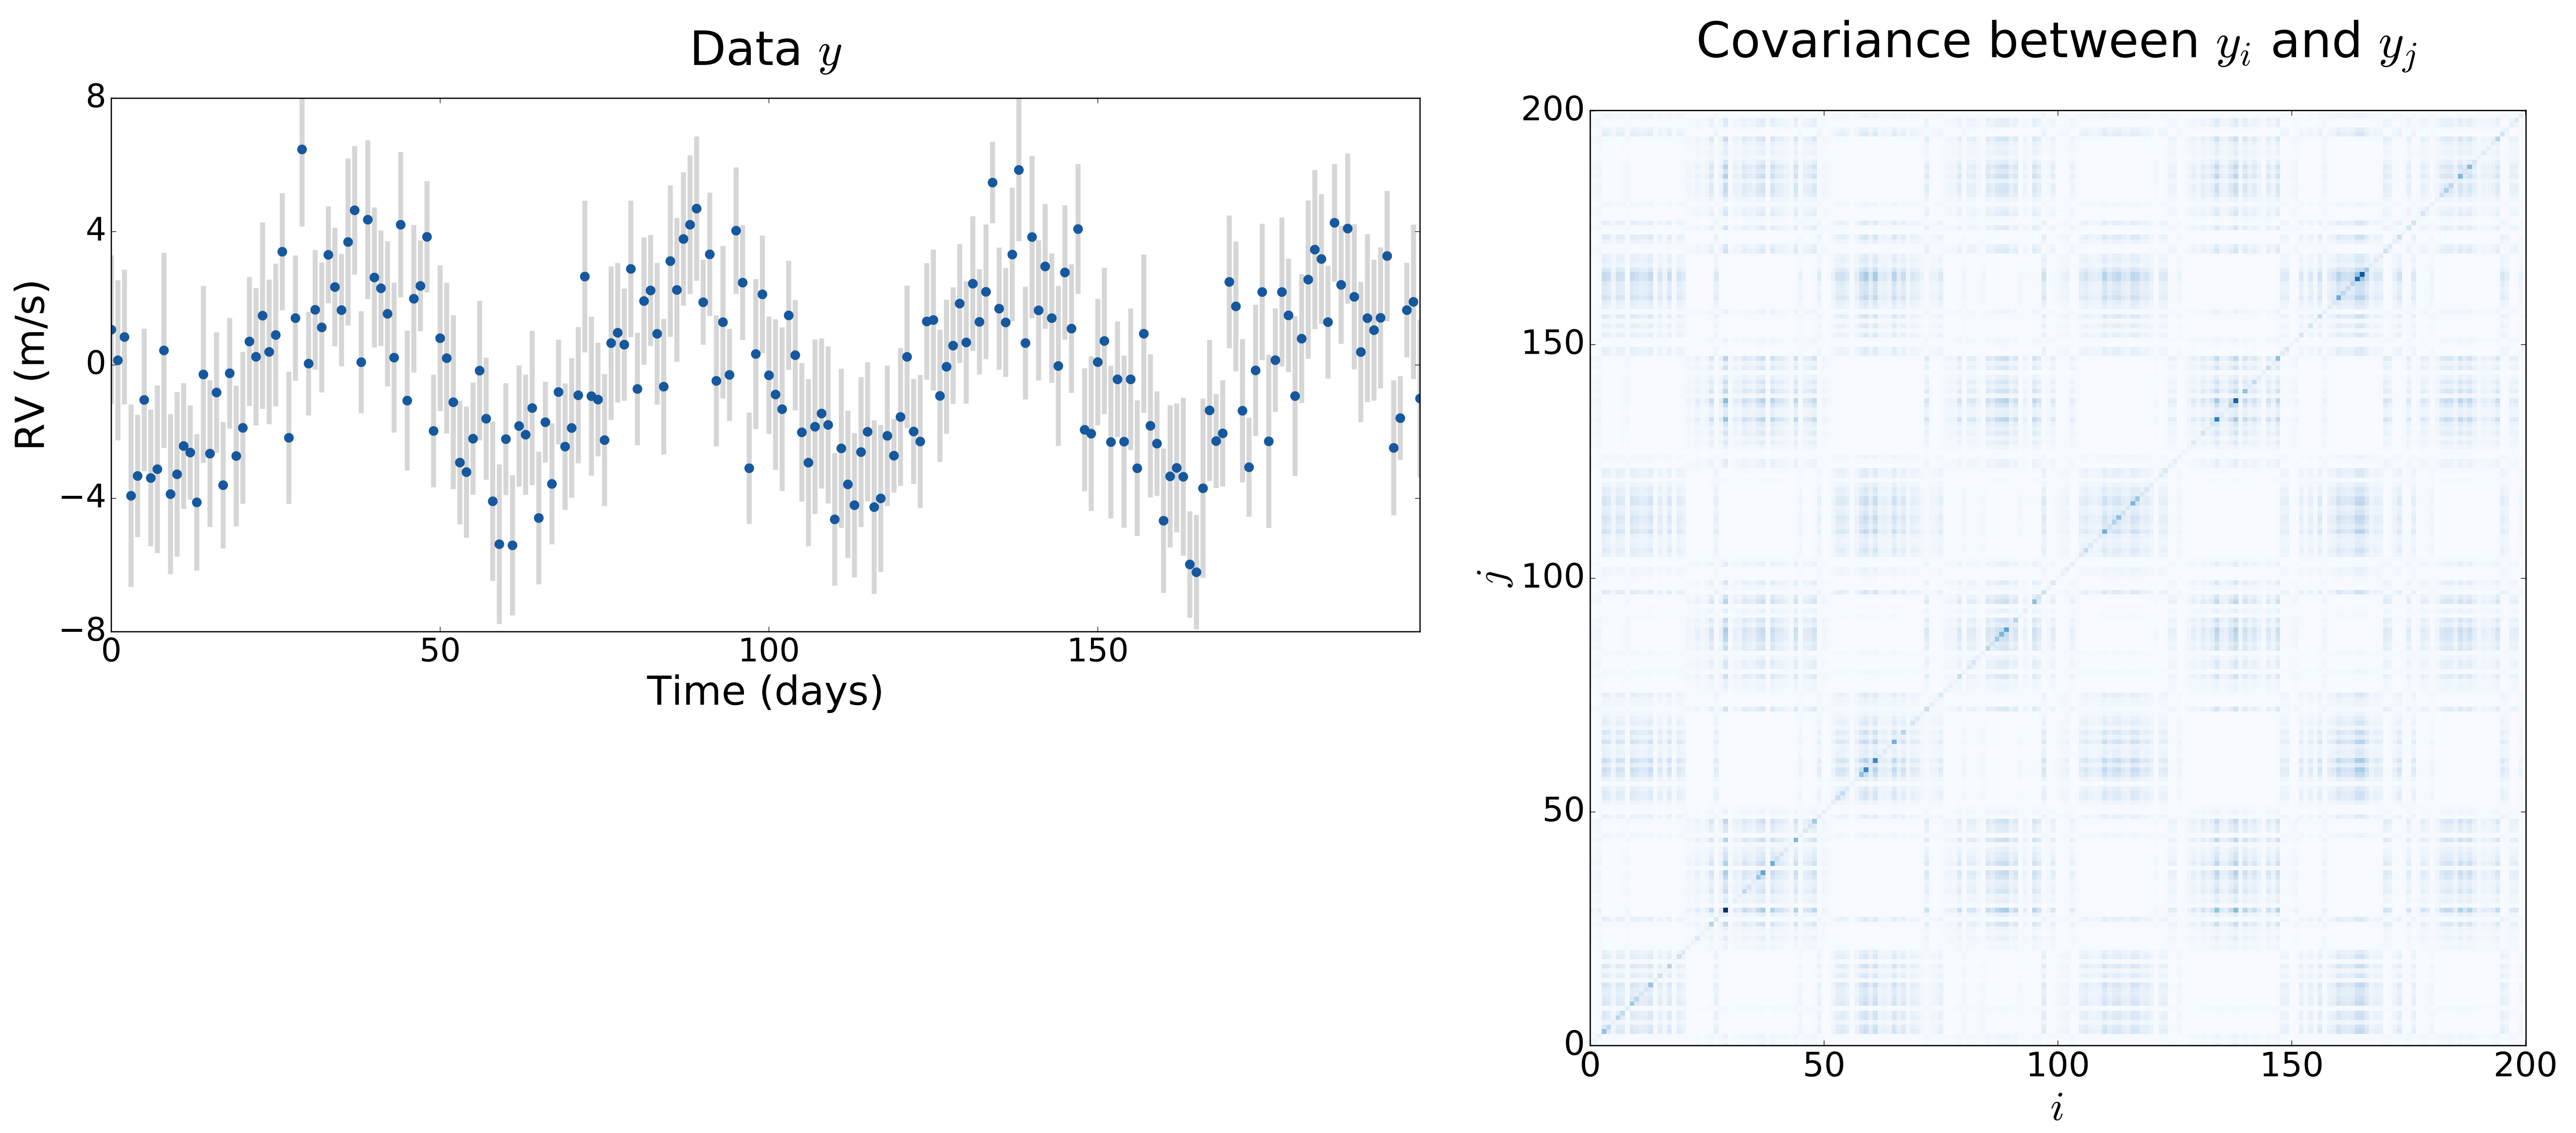

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Correlated noise ("red" noise)

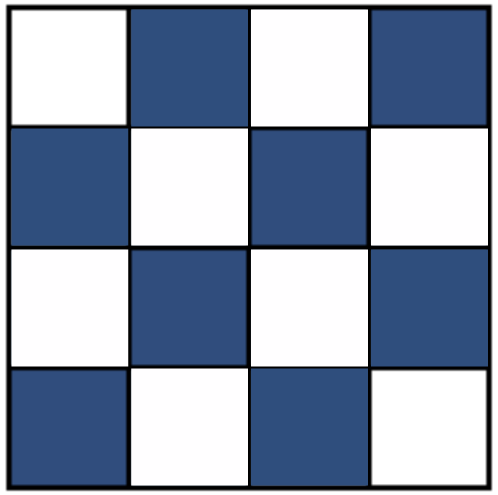

Data points are correlated with each other

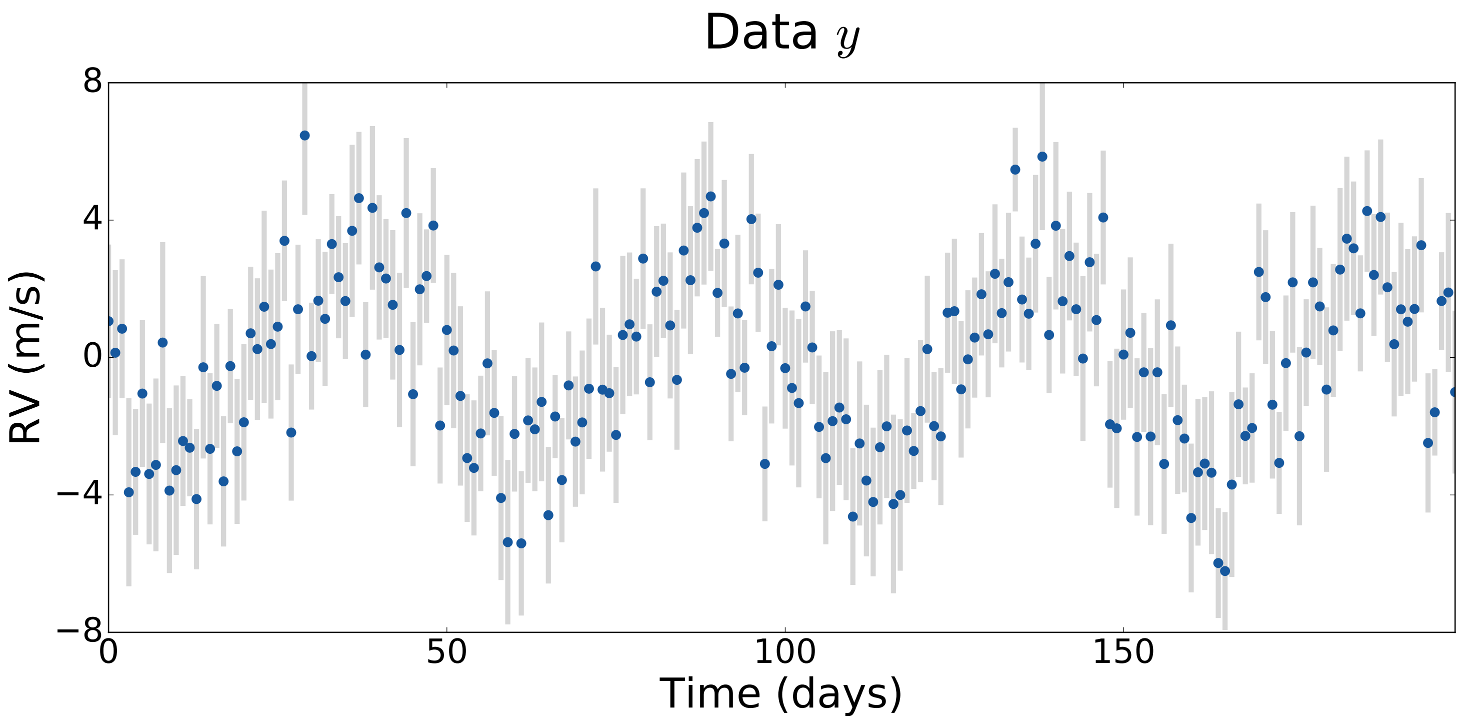

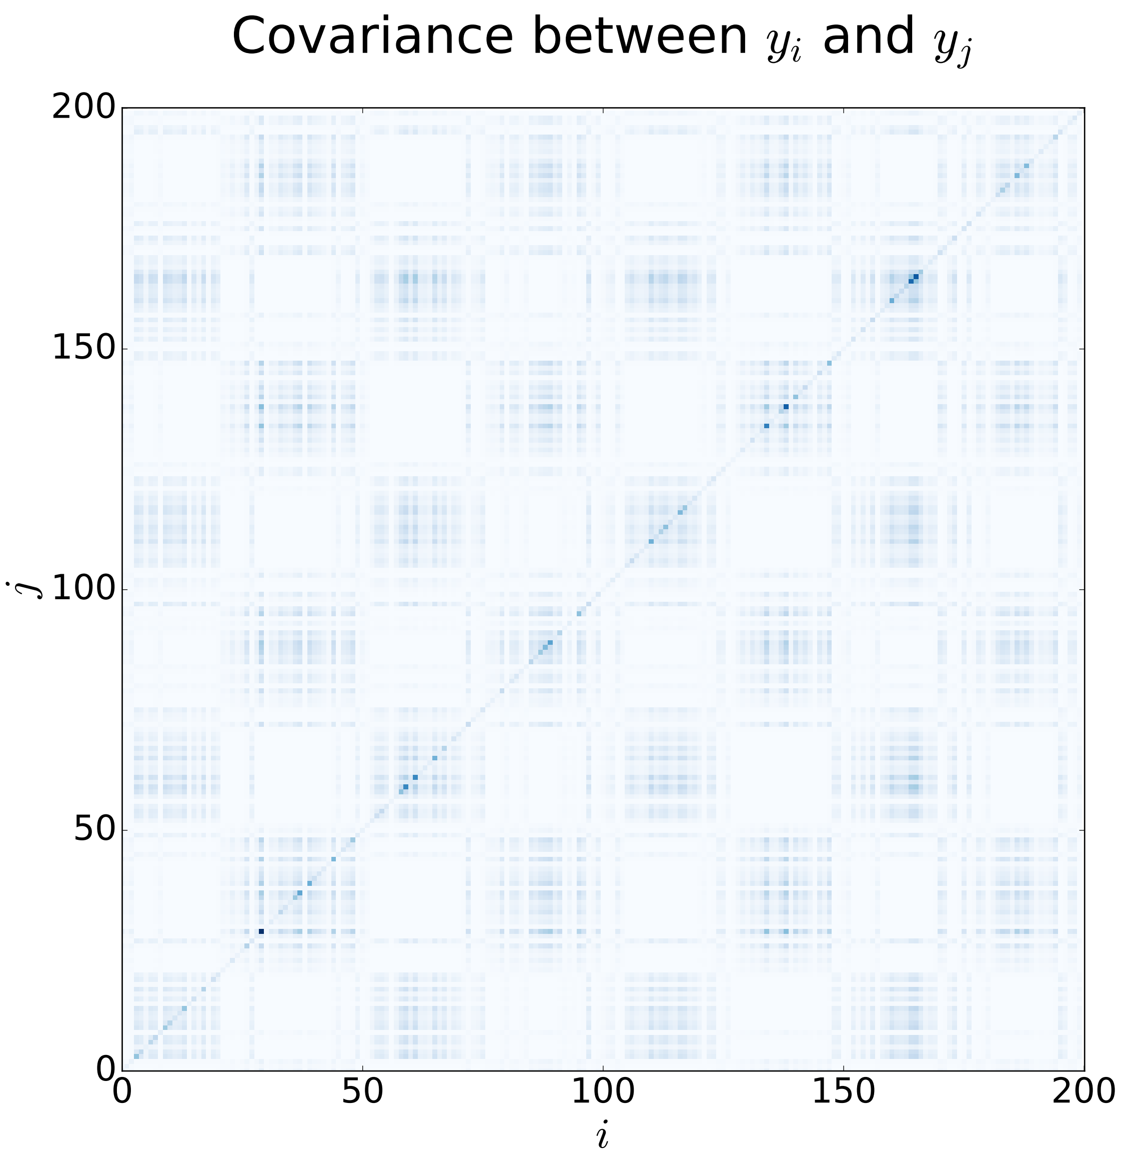

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Correlated noise ("red" noise)

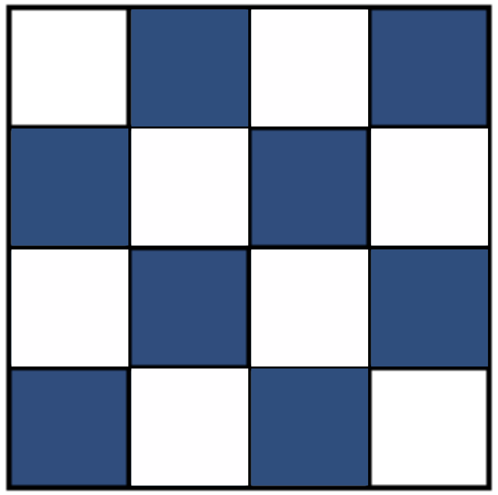

Data points are correlated with each other

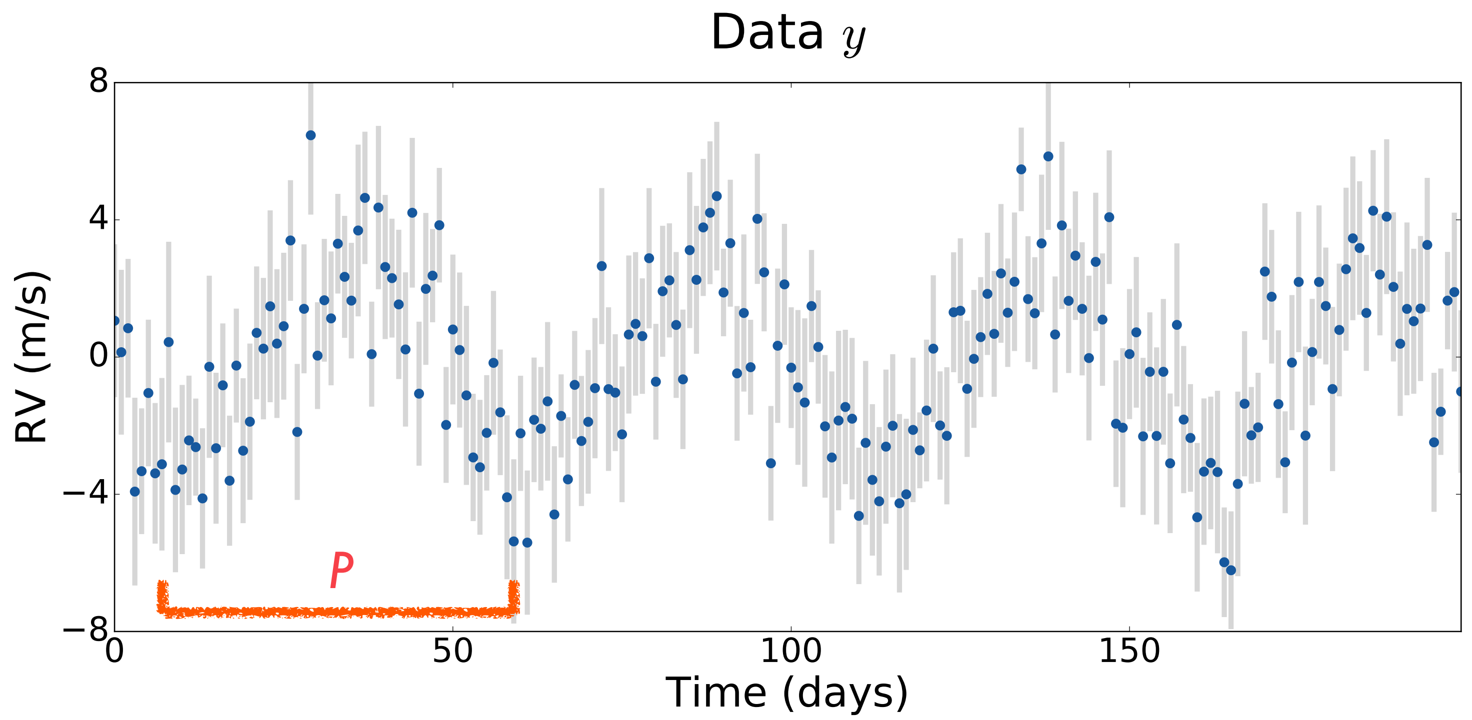

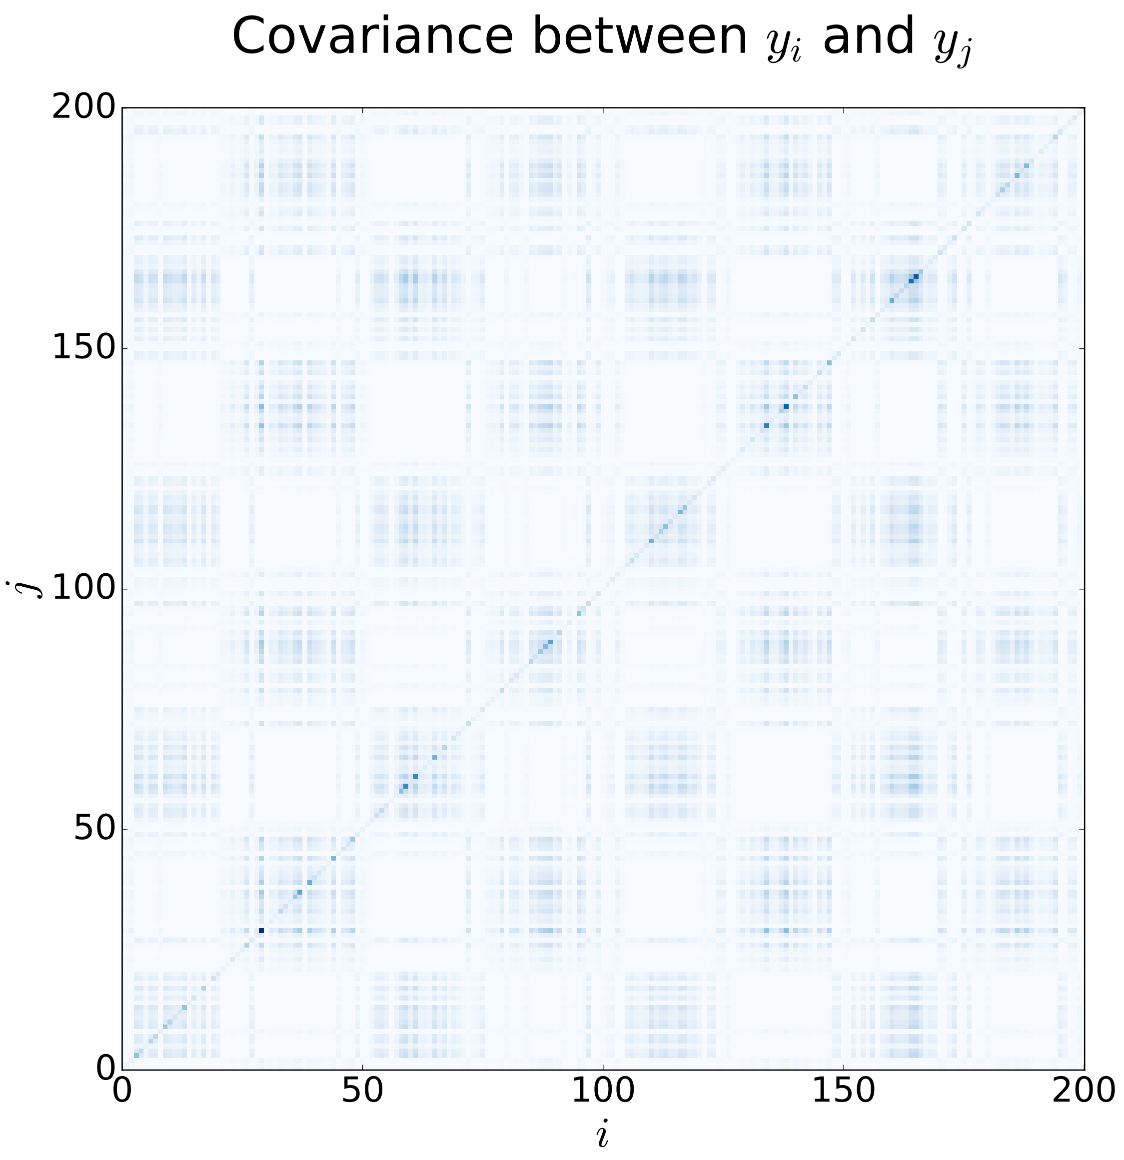

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Correlated noise ("red" noise)

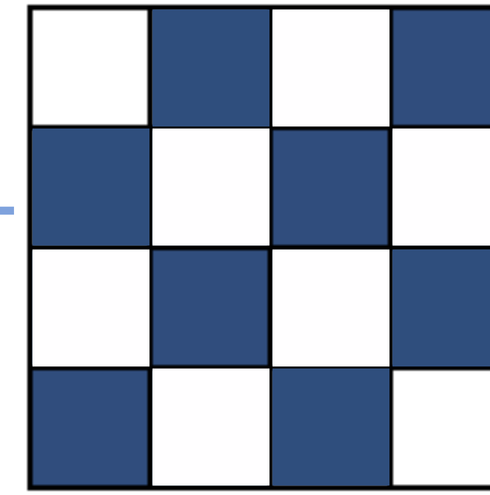

Data points are correlated with each other

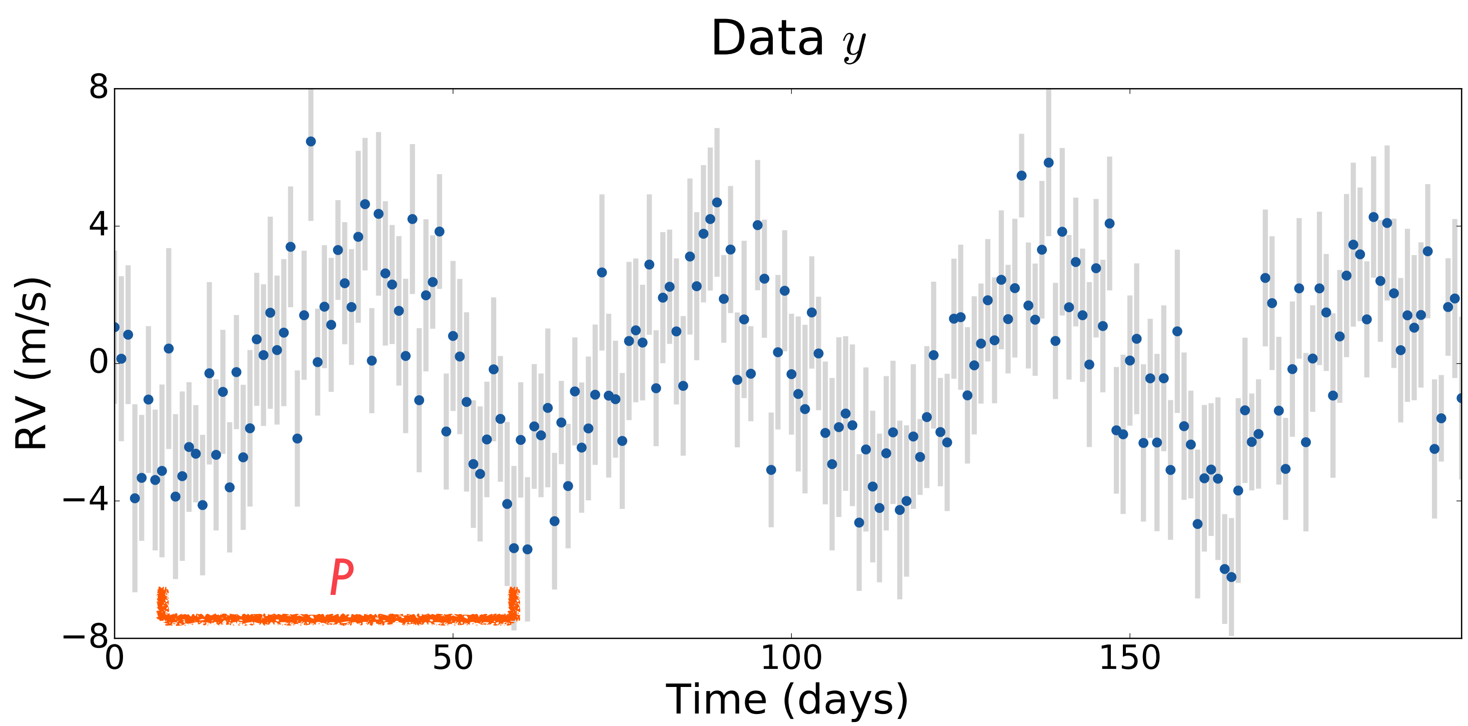

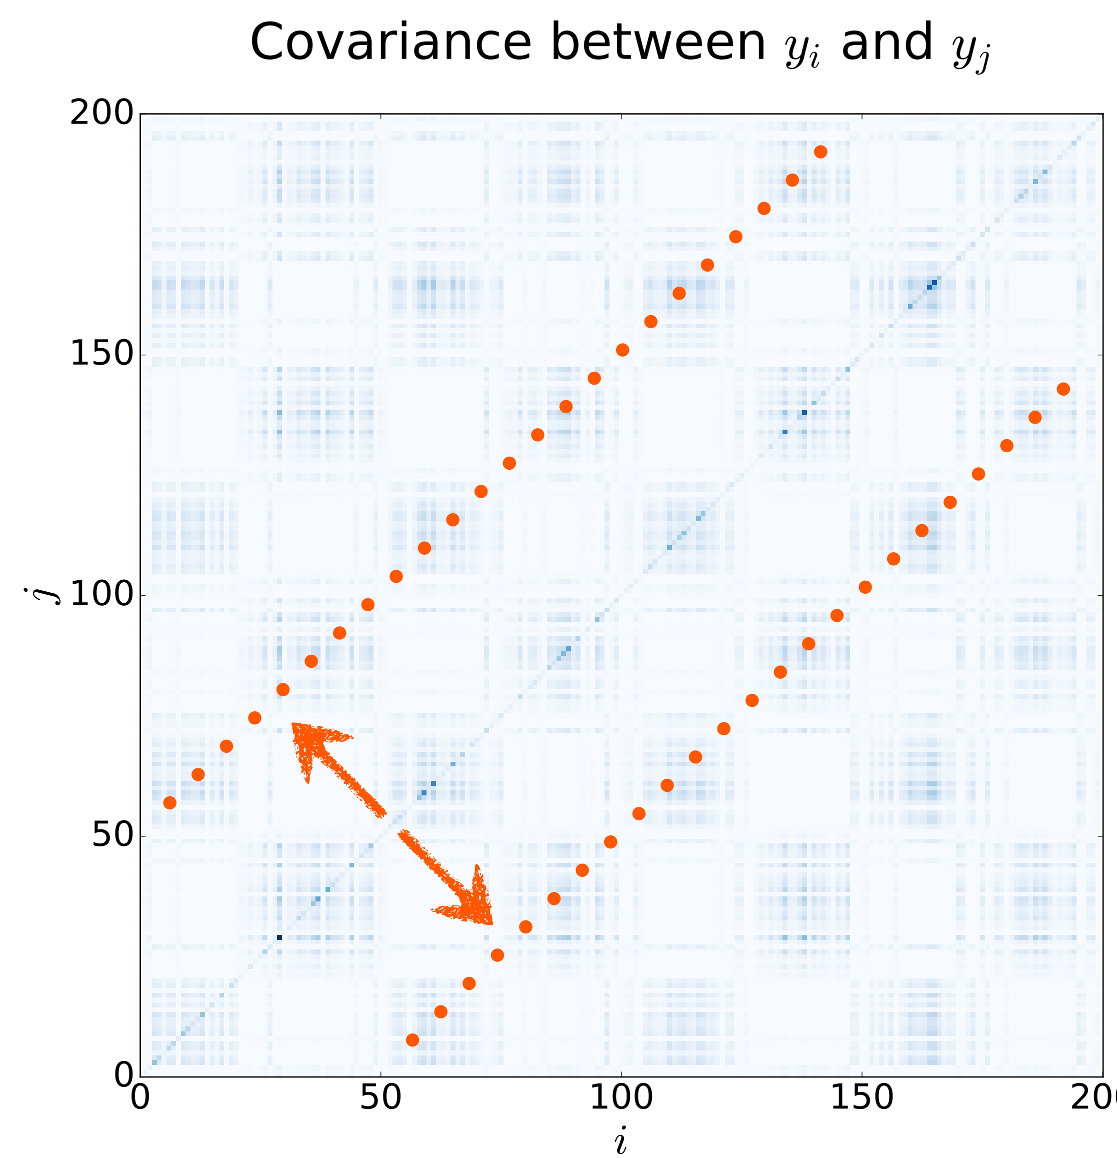

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Correlated noise (“red” noise)

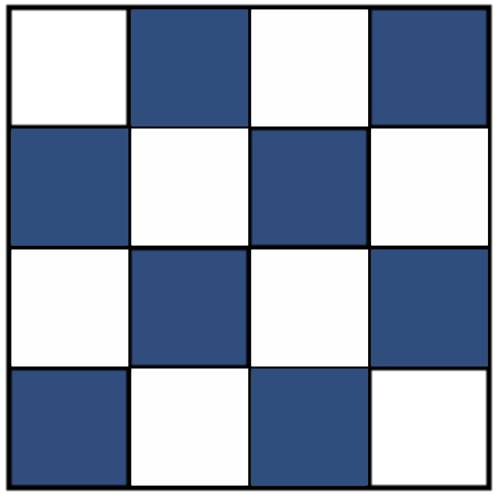

Data points are correlated with each other

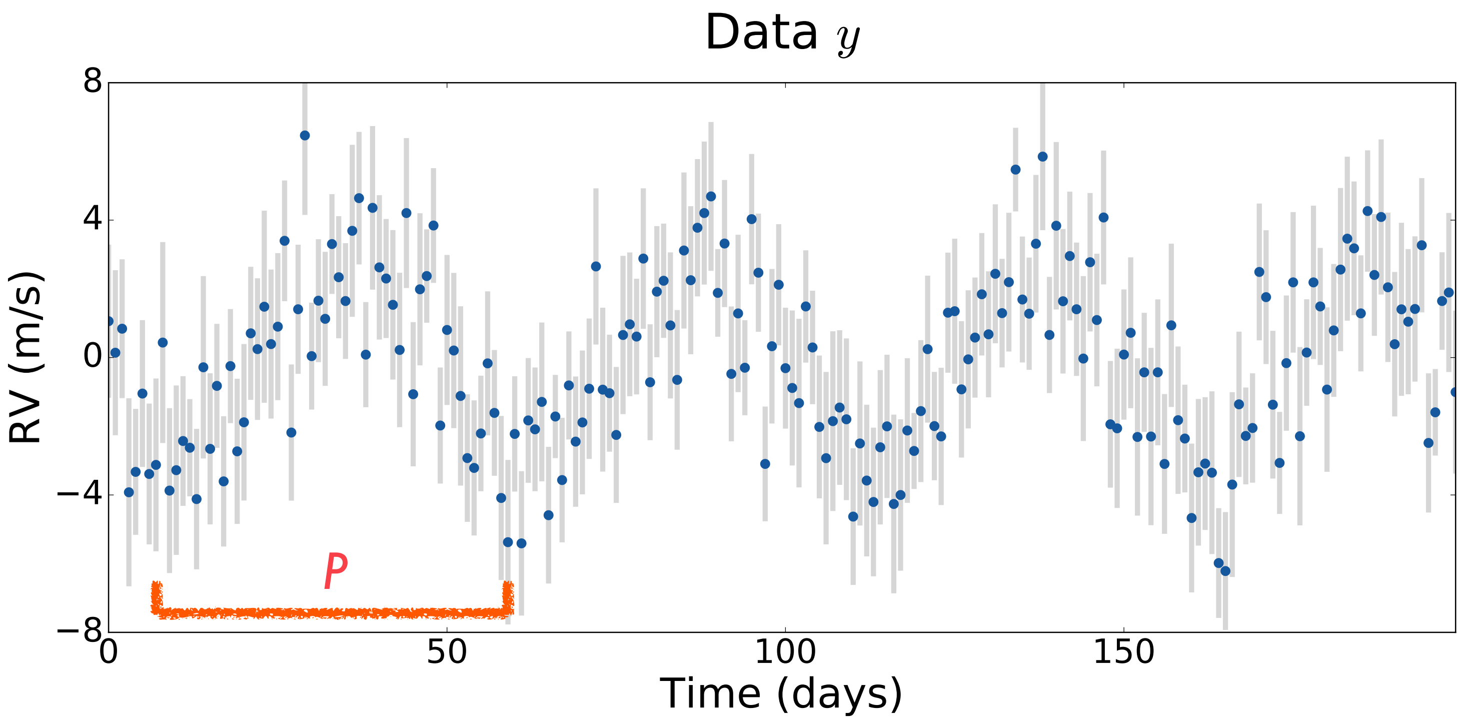

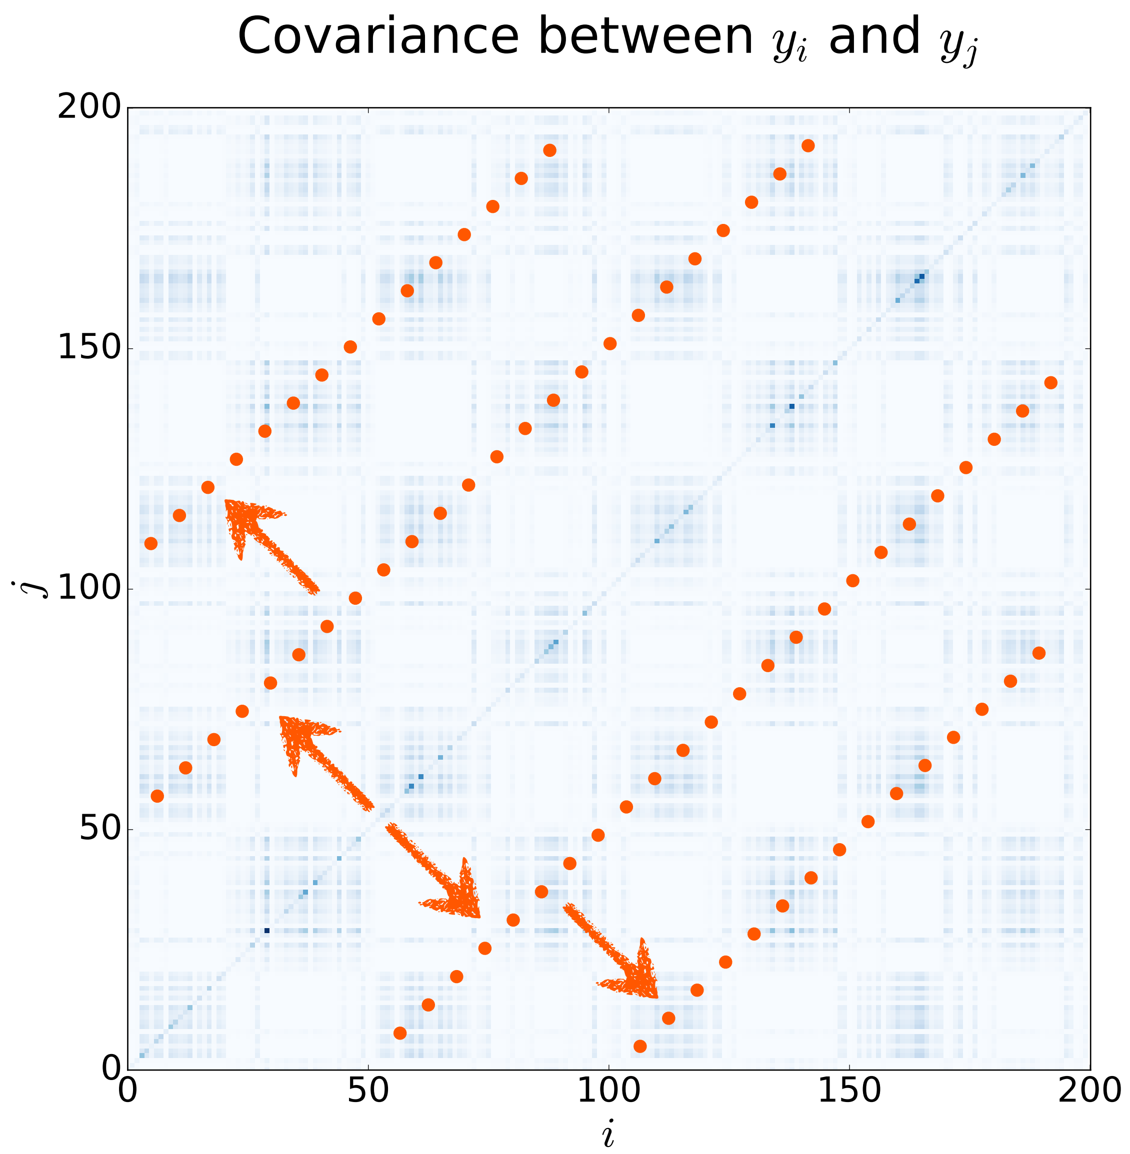

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Real-life example: CoRoT-7 lightcurve

Quasi-periodic variations

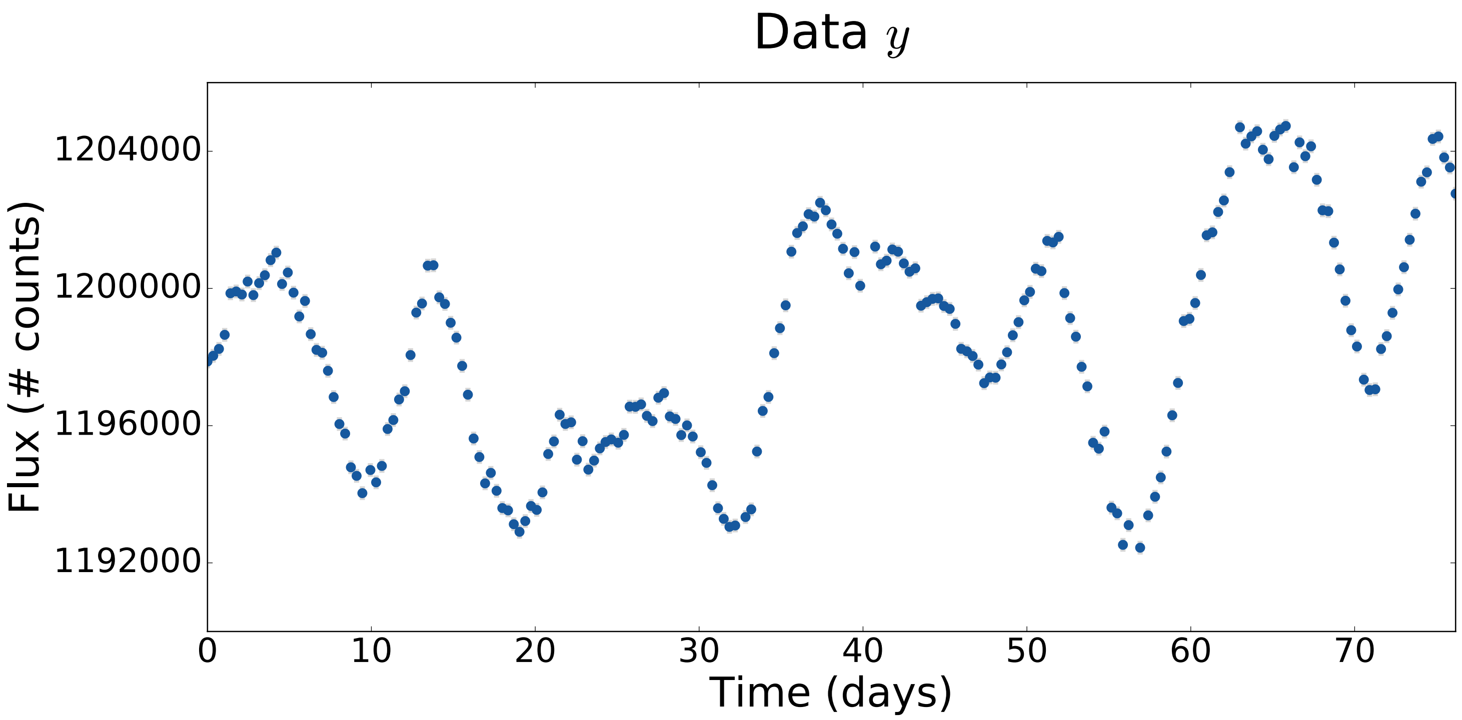

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Real-life example: CoRoT-7 lightcurve

Quasi-periodic variations

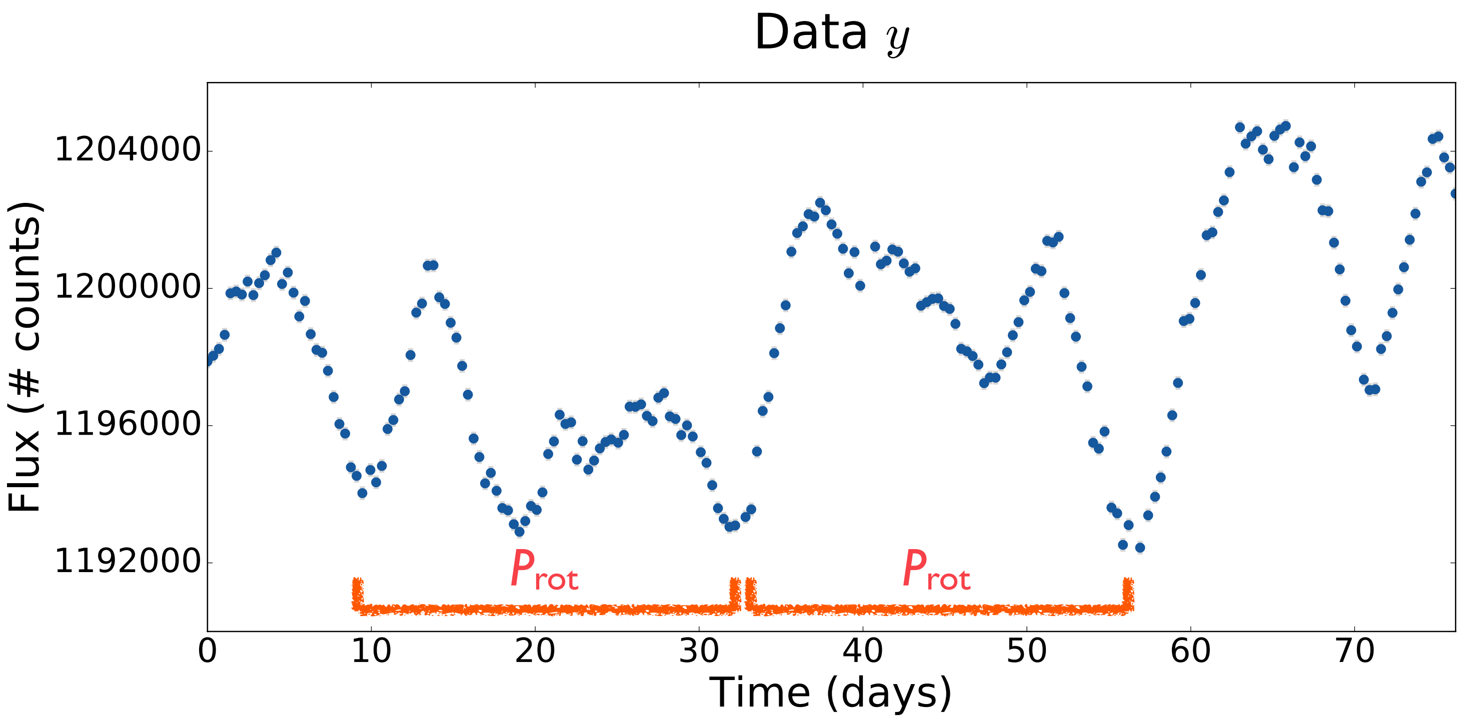

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Real-life example: CoRoT-7 lightcurve

Quasi-periodic variations

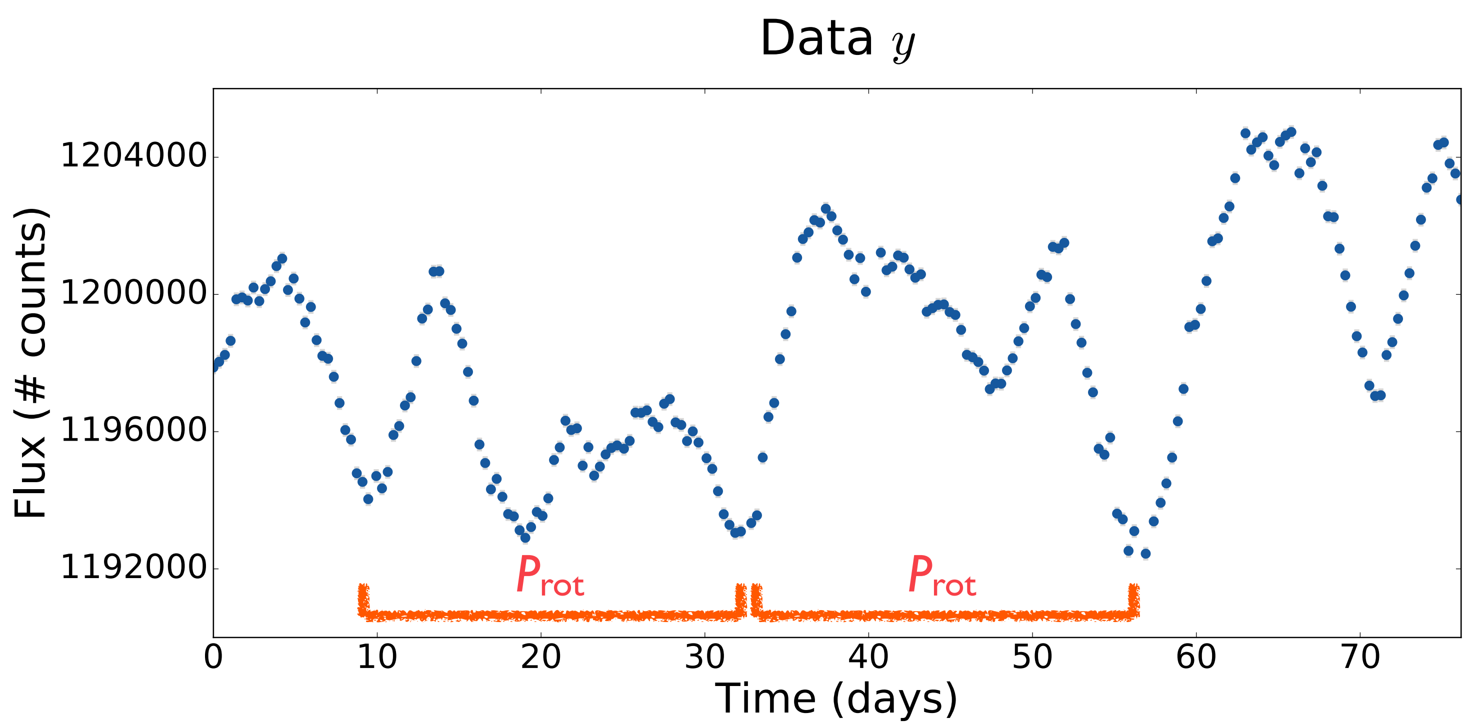

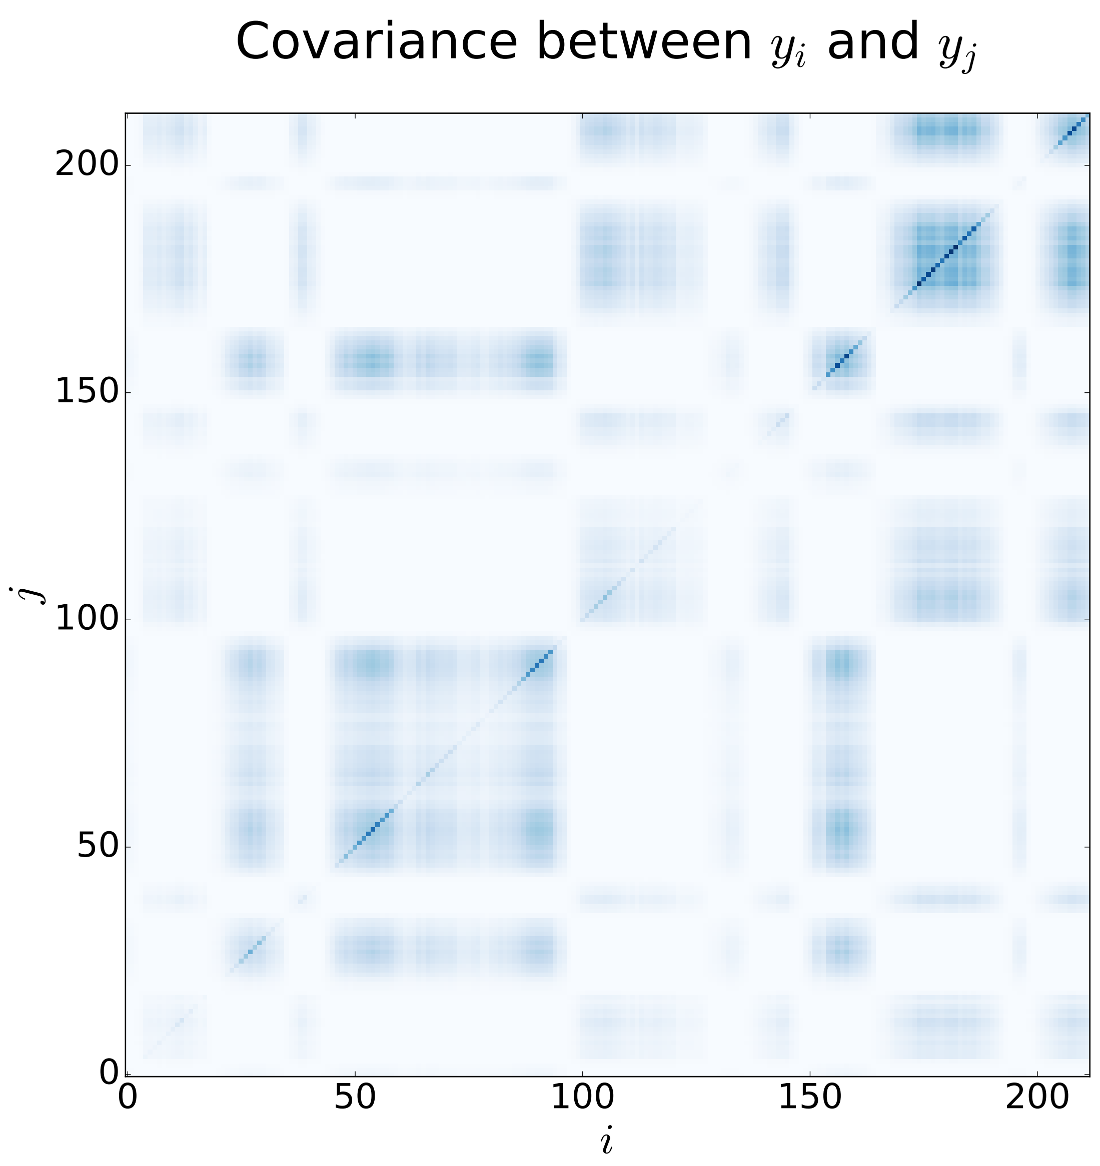

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Real-life example: CoRoT-7 lightcurve

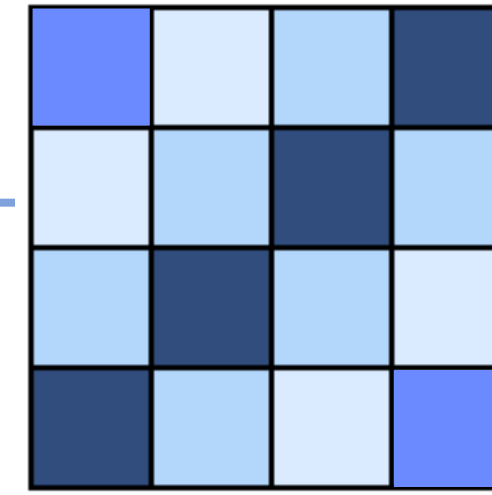

Quasi-periodic variations

Data y

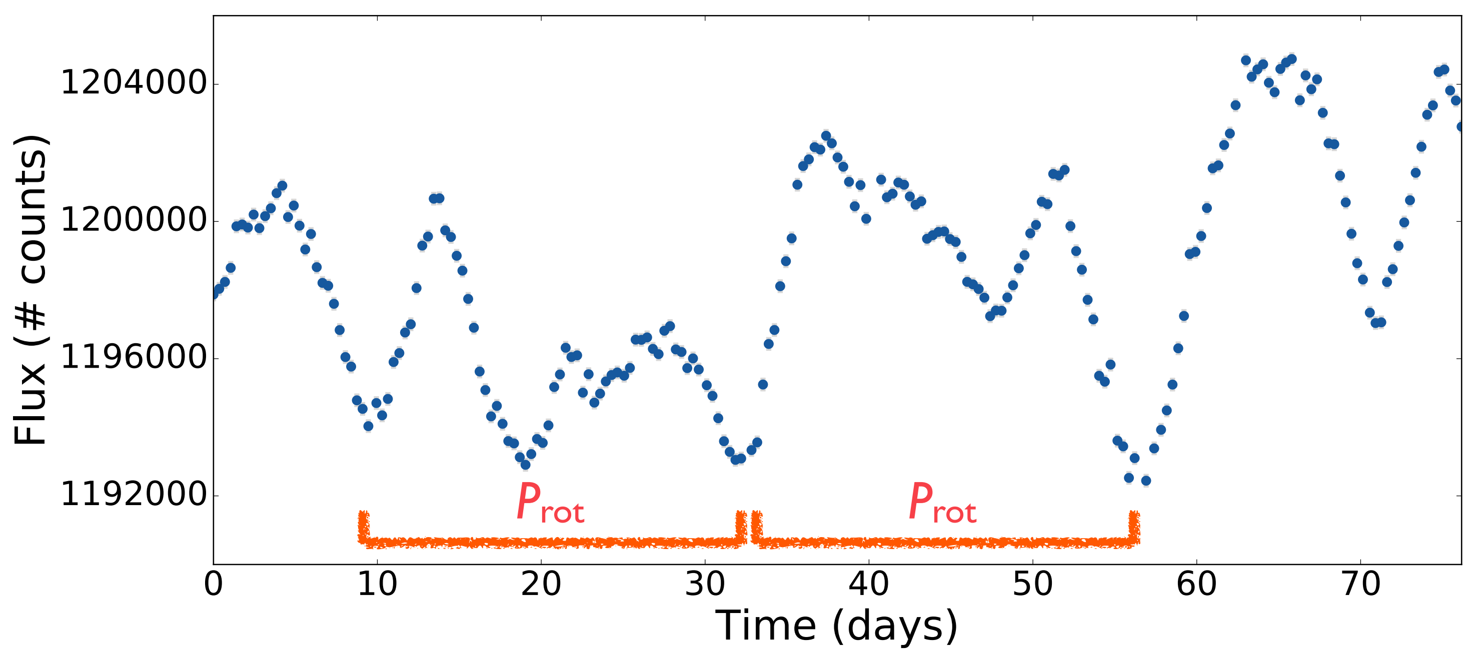

Covariance between y_i and y_j

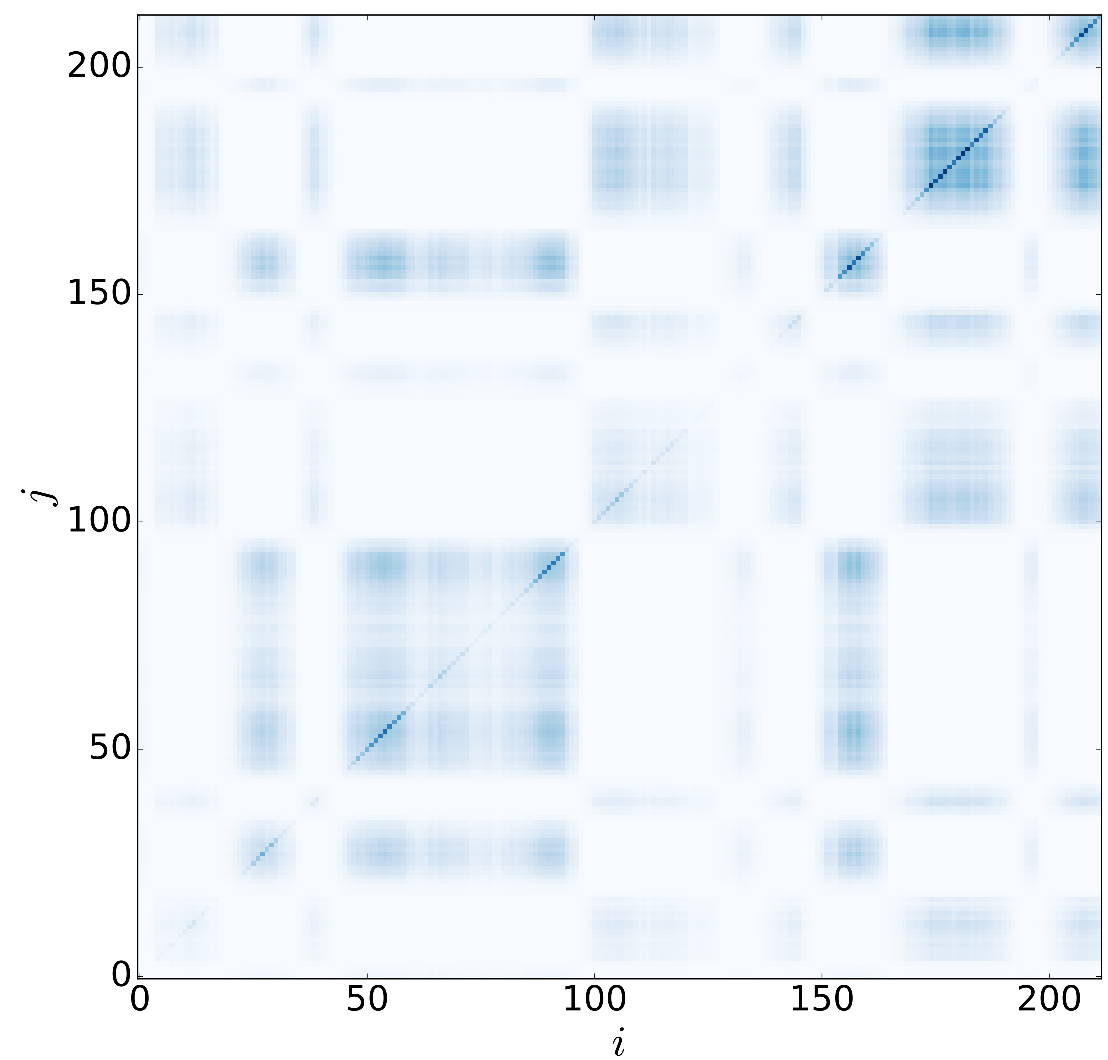

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Real-life example: CoRoT-7 lightcurve

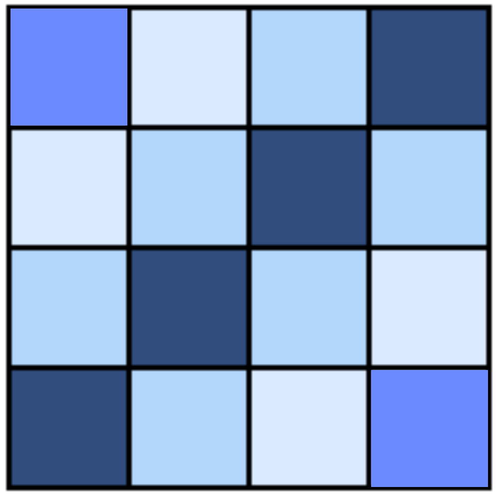

Quasi-periodic variations

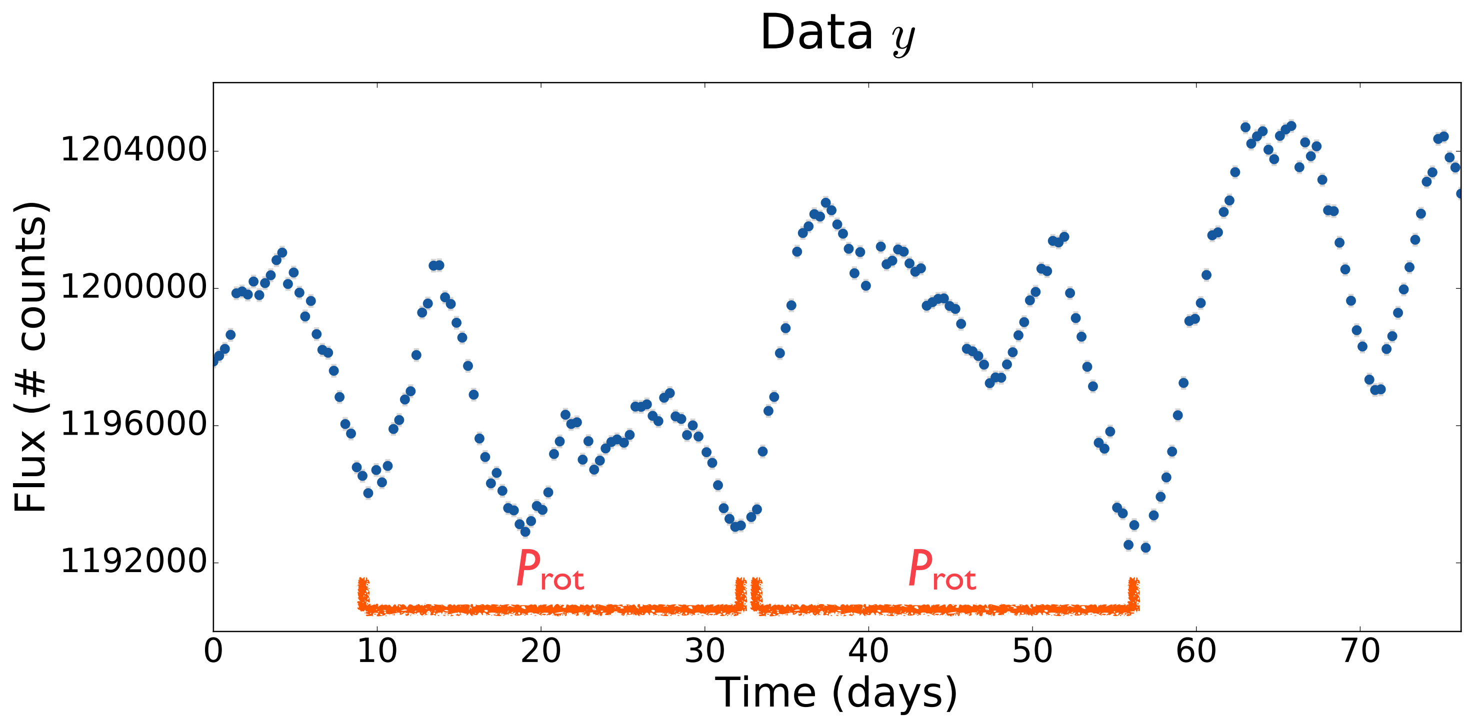

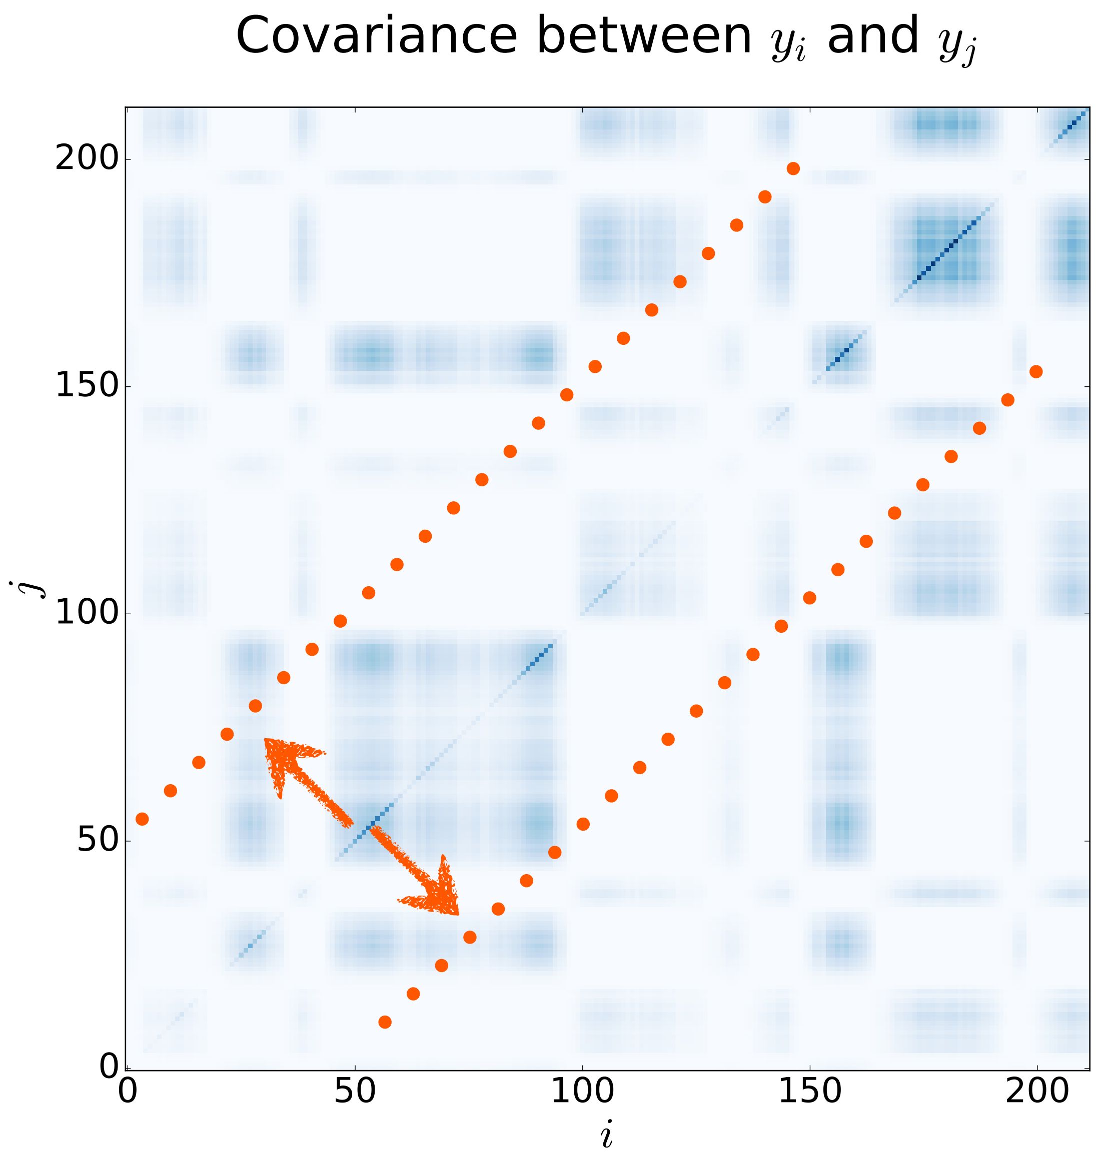

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Real-life example: CoRoT-7 lightcurve

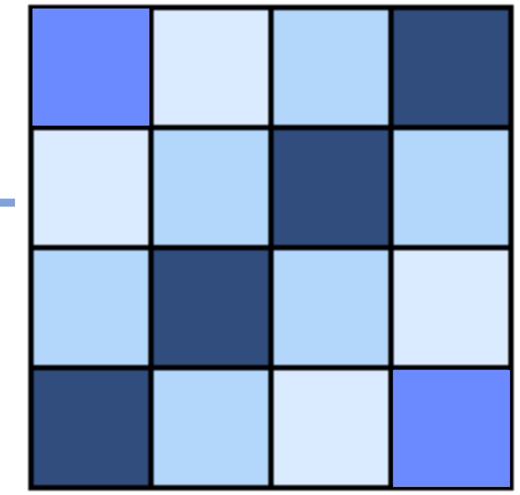

Quasi-periodic variations

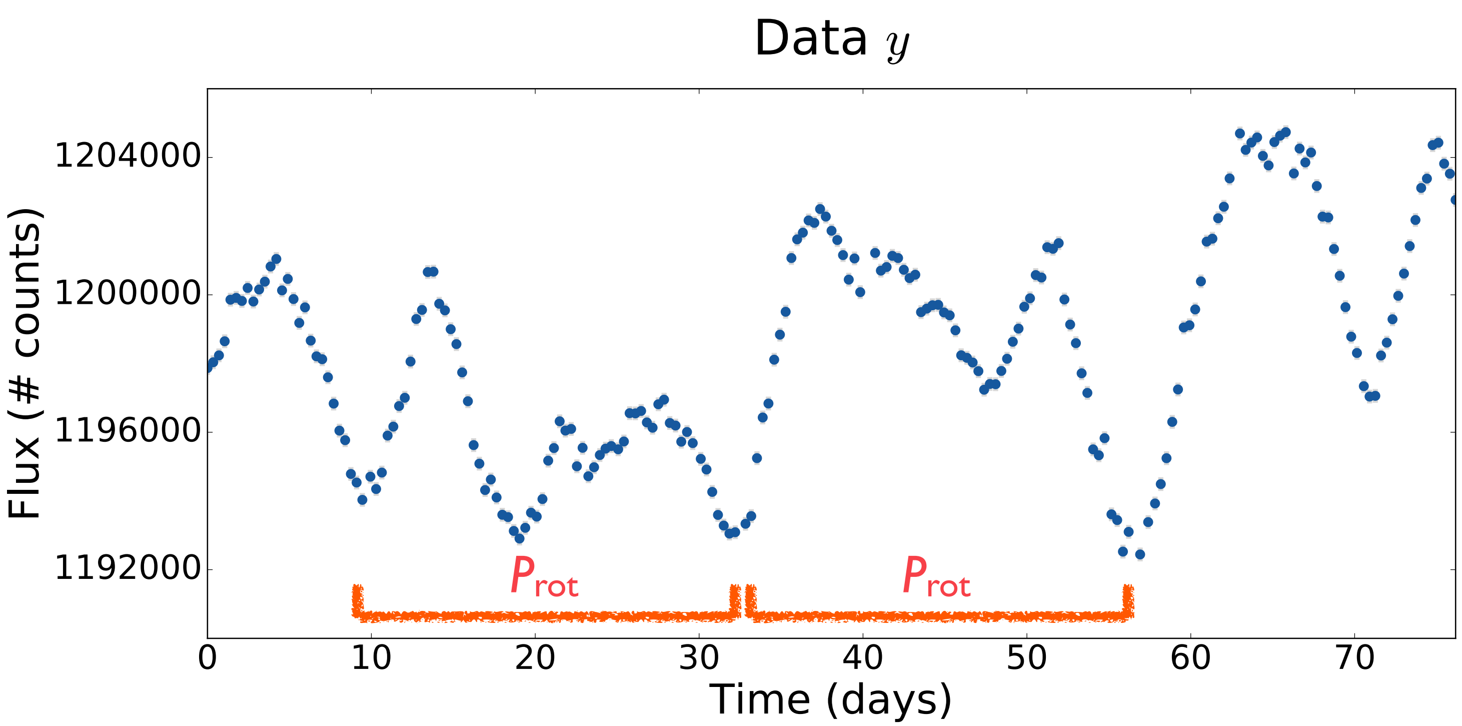

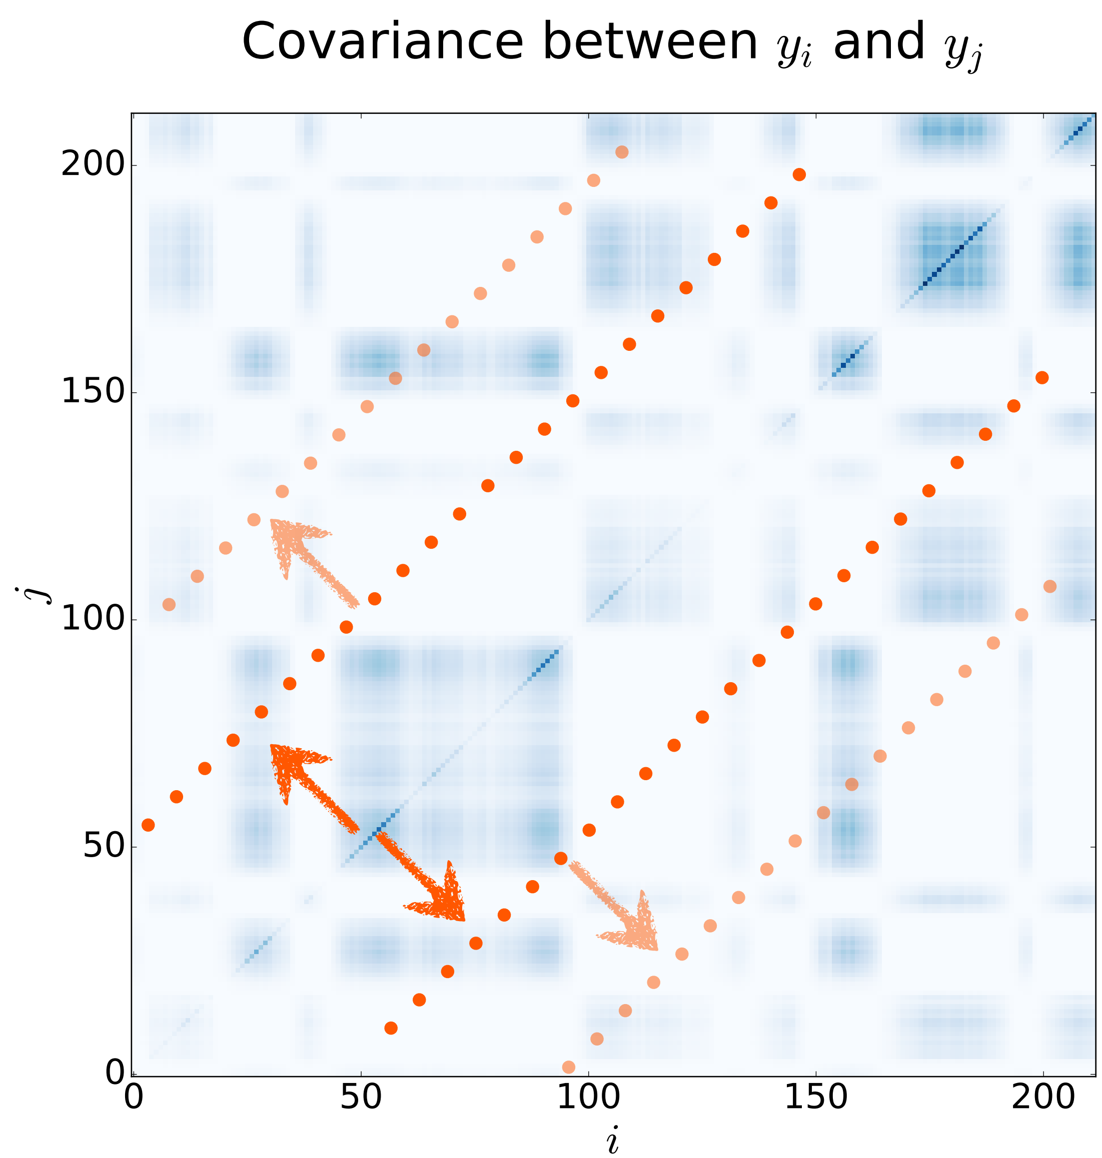

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Non-parametric frameworks: we fit in "correlation (covariance) space"

Covariance matrix

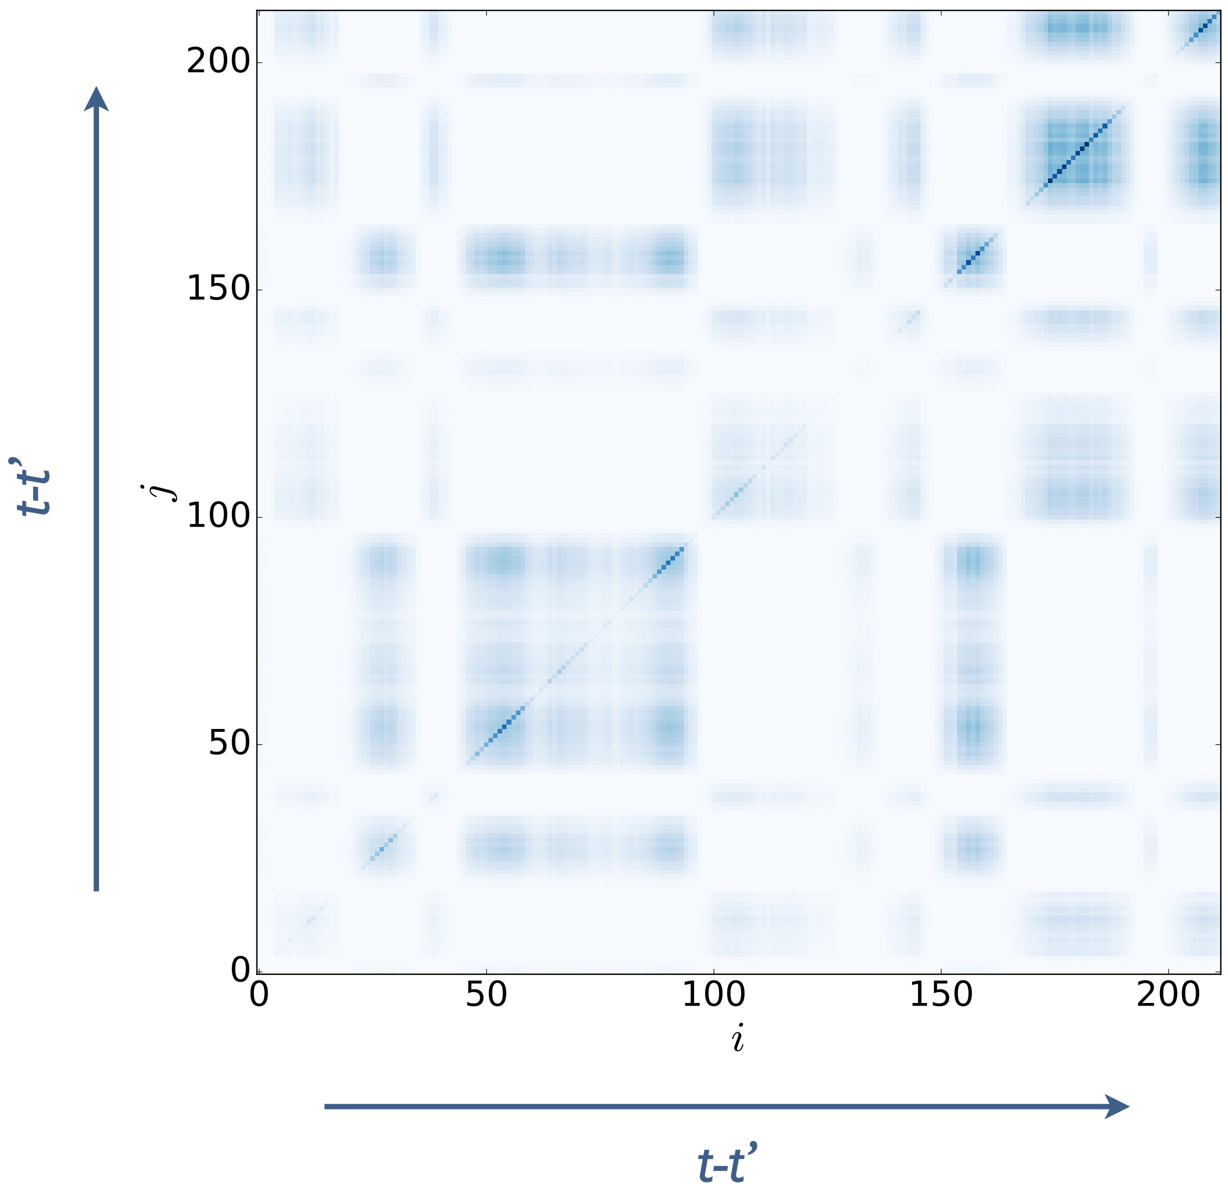

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Non-parametric frameworks: we fit in "correlation (covariance) space"

Covariance matrix $\mathbf{K}_{i,j} = k(t_i, t_j) = k(t,t')$

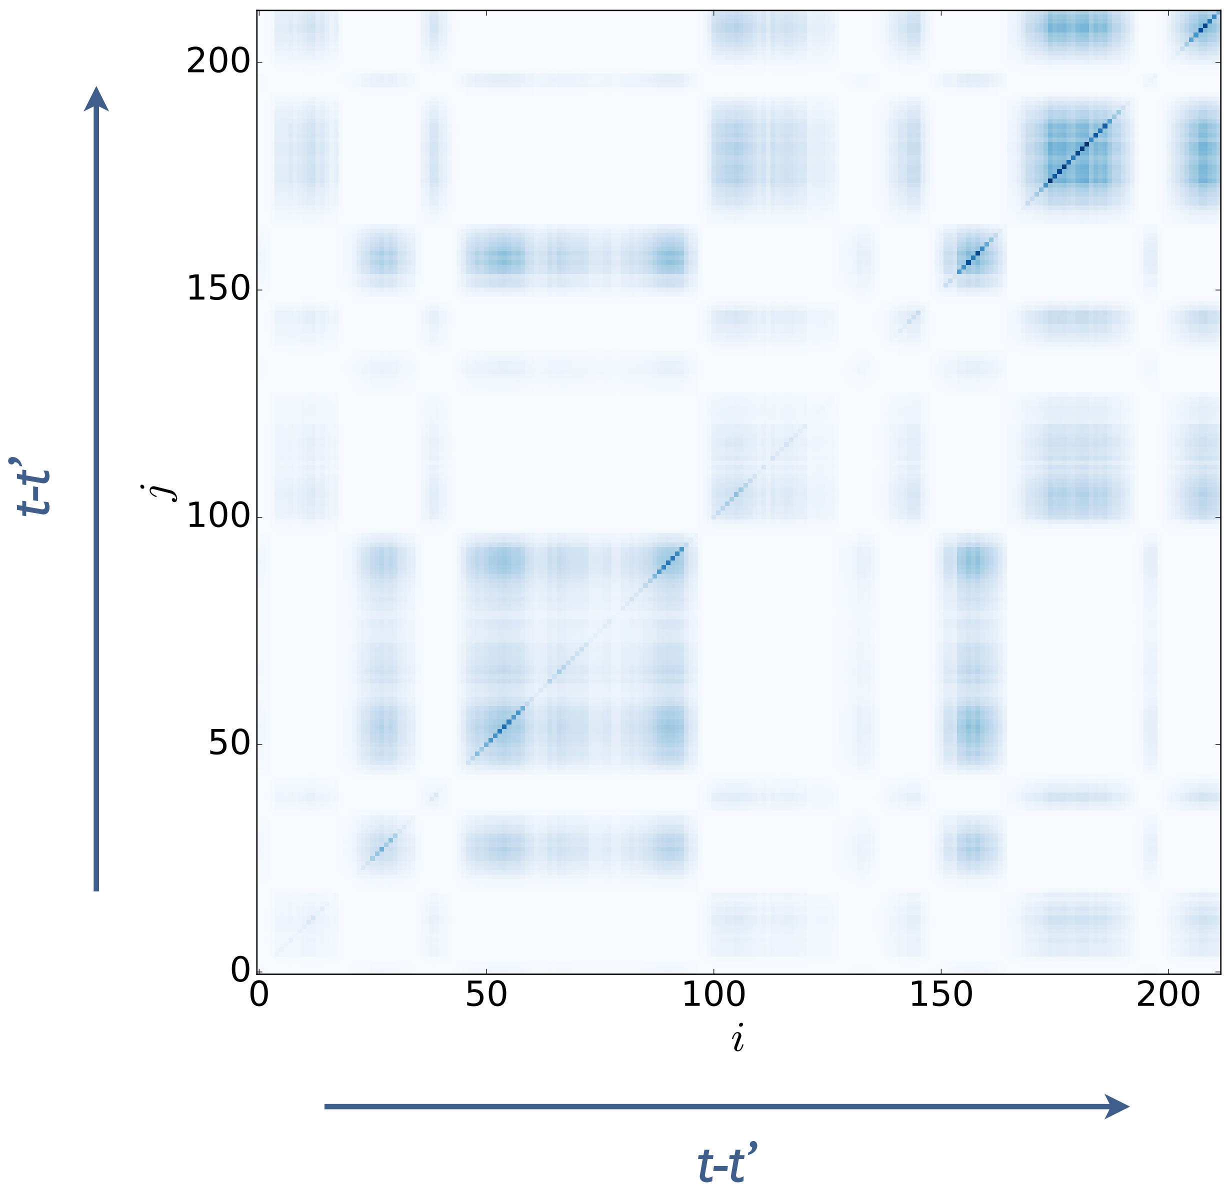

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Non-parametric frameworks: we fit in “correlation (covariance) space”

Covariance matrix $\mathbf{K}_{i,j} = k(t_i, t_j) = k(t,t')$

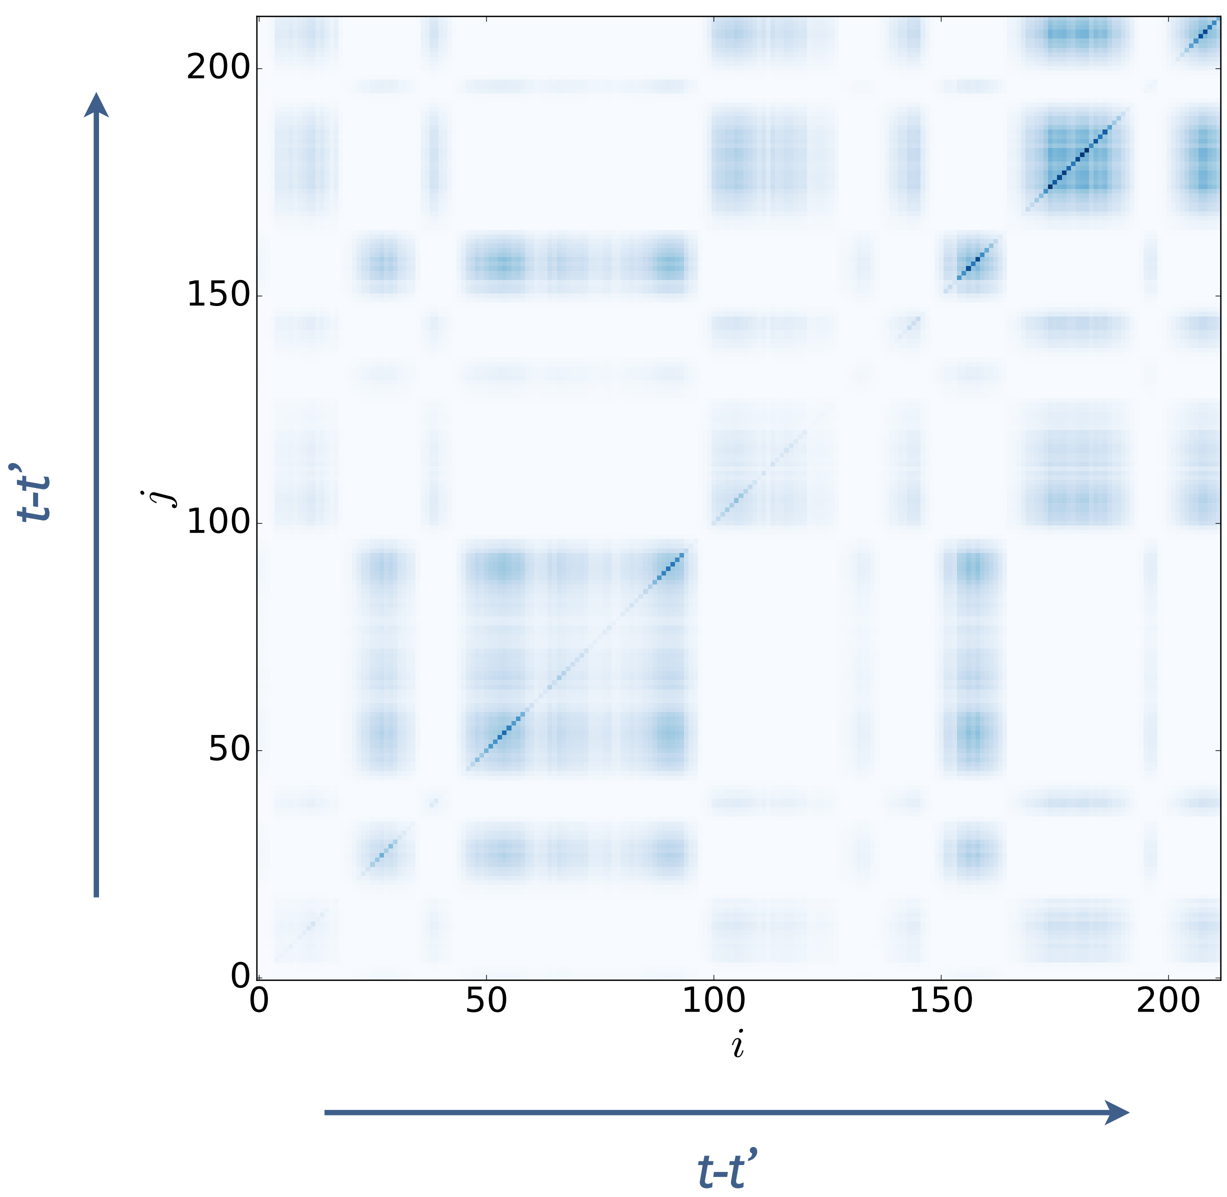

Covariance function $k(t, t')$

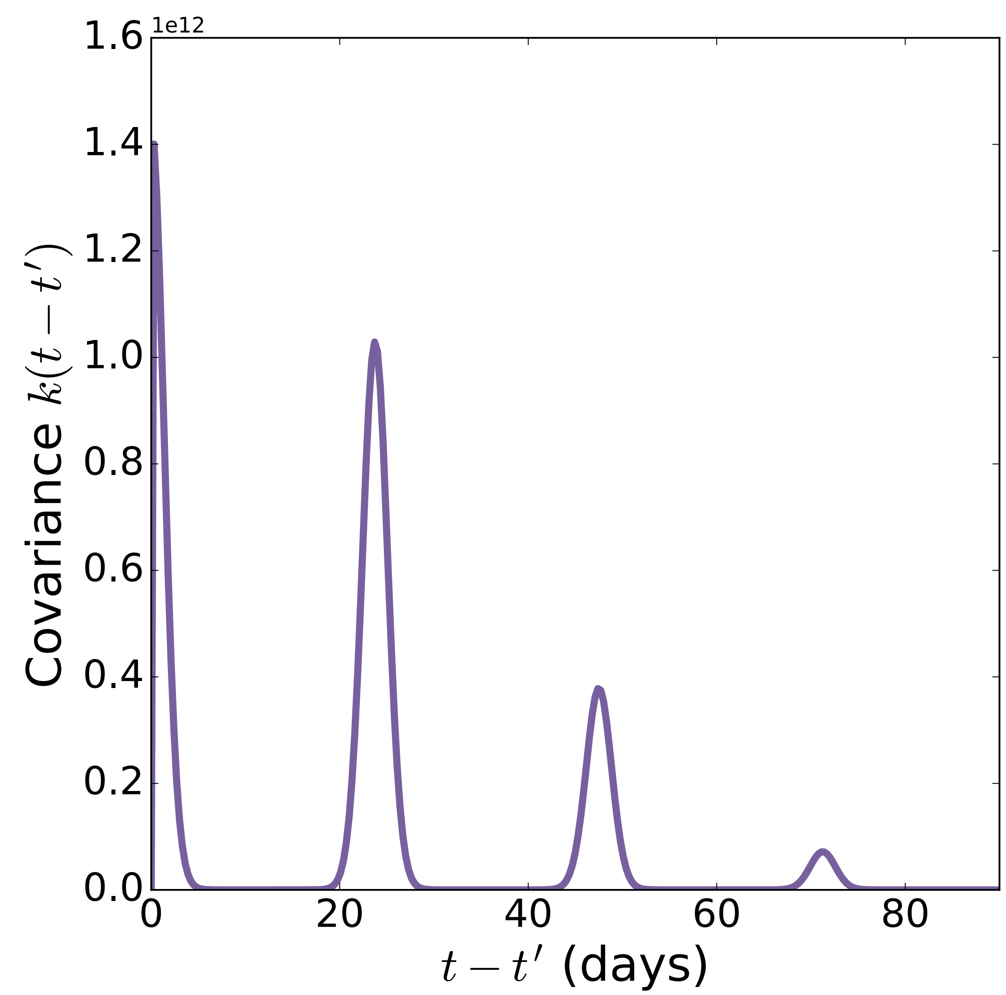

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Non-parametric frameworks: we fit in “correlation (covariance) space”

Covariance matrix $\mathbf{K}_{i,j} = k(t_i, t_j) = k(t,t')$

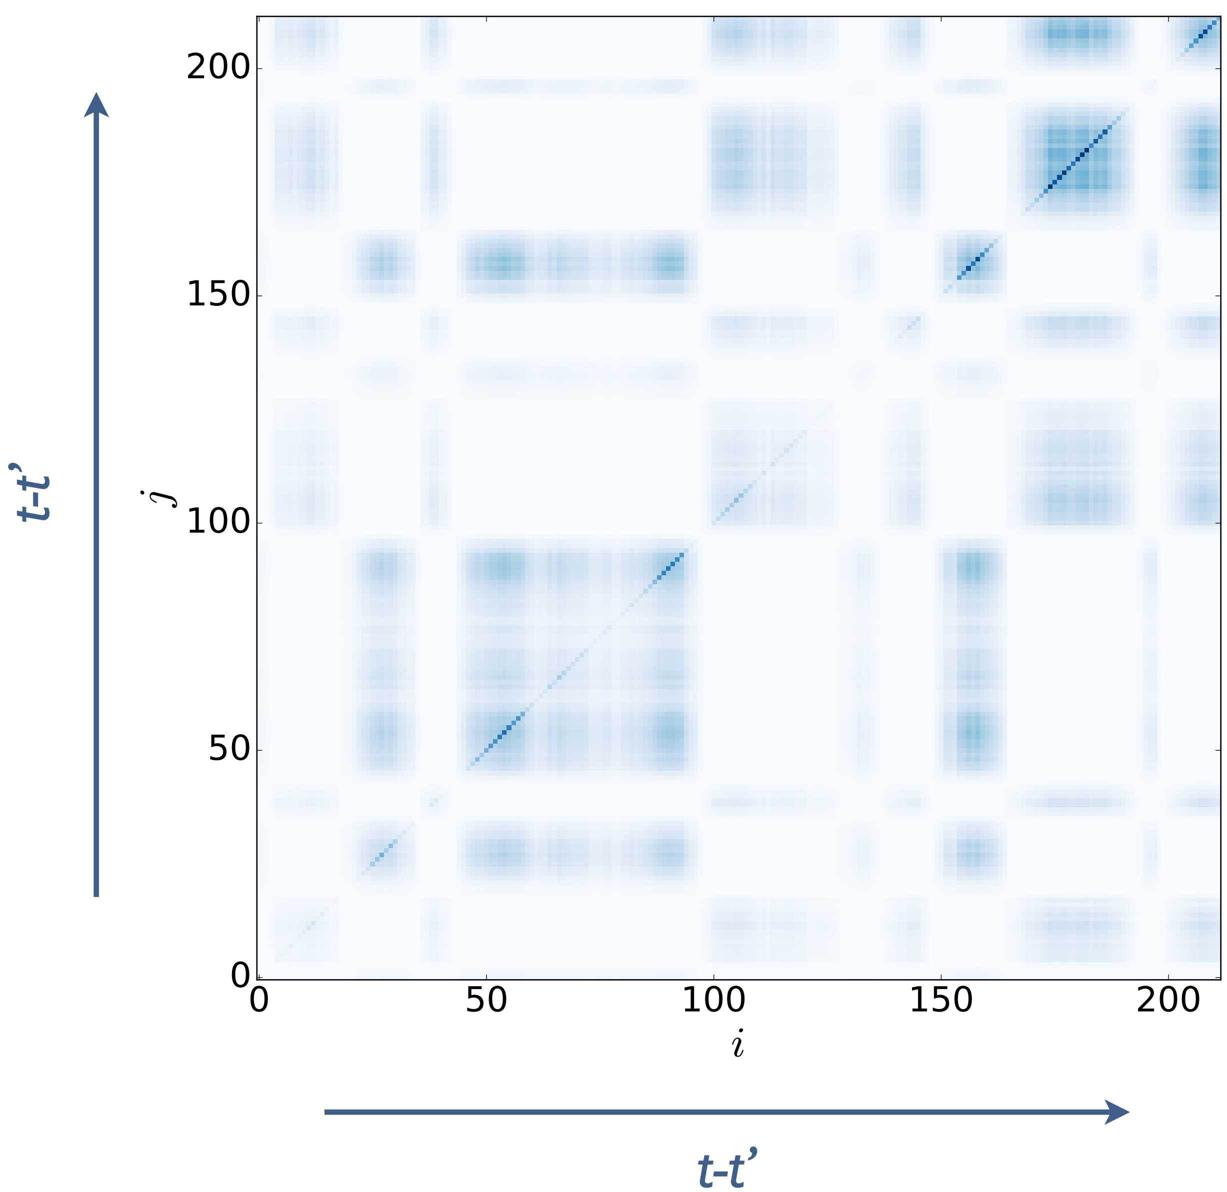

Covariance function $k(t, t')$

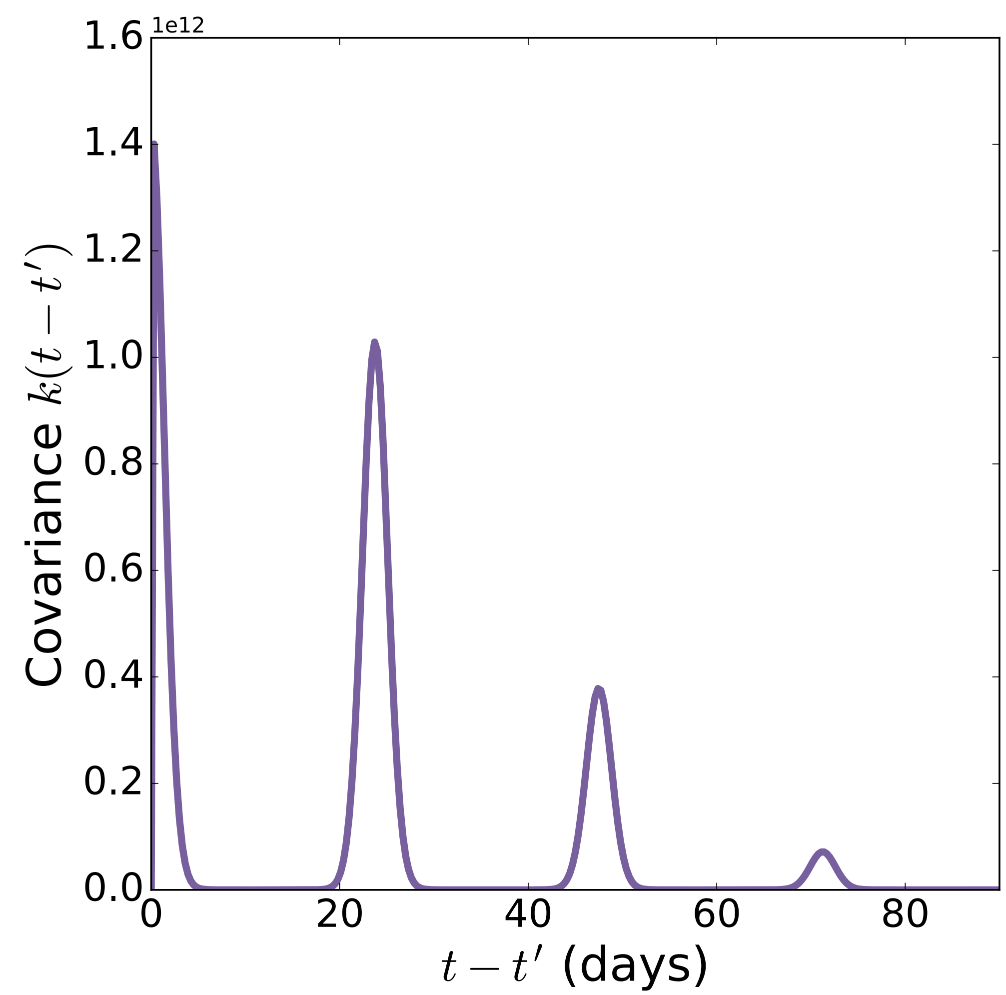

Quasi-periodic form:

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

Choose a covariance function to fill your covariance matrix

Covariance matrix $\mathbf{K}_{i,j} = k(t_i, t_j) = k(t,t')$

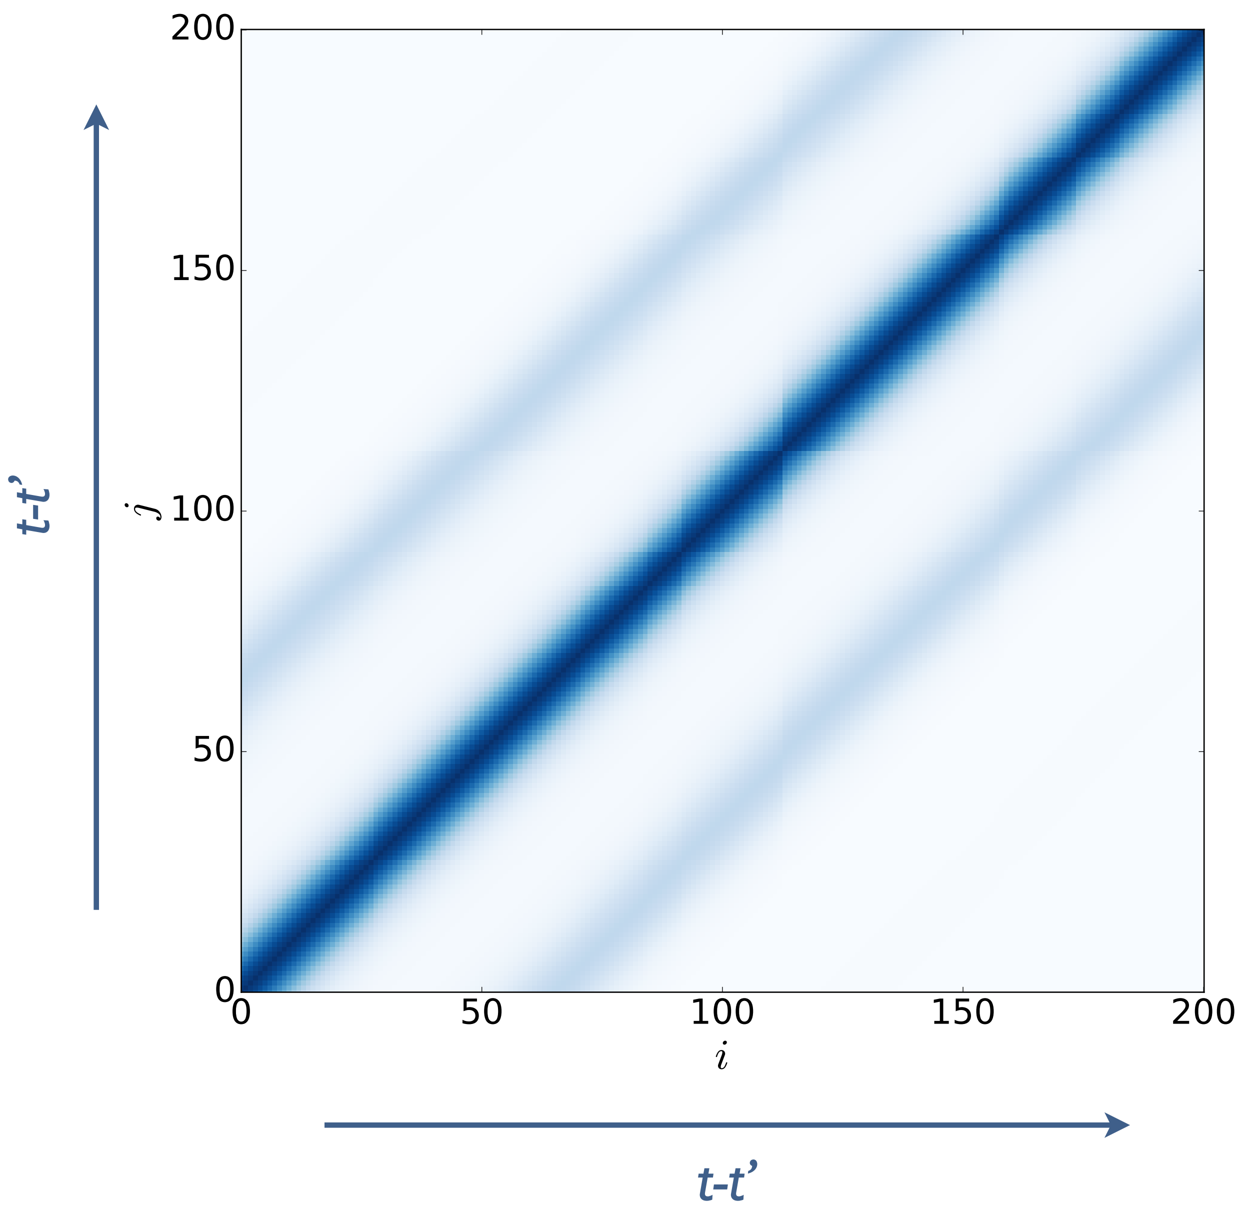

Covariance function $k(t, t')$

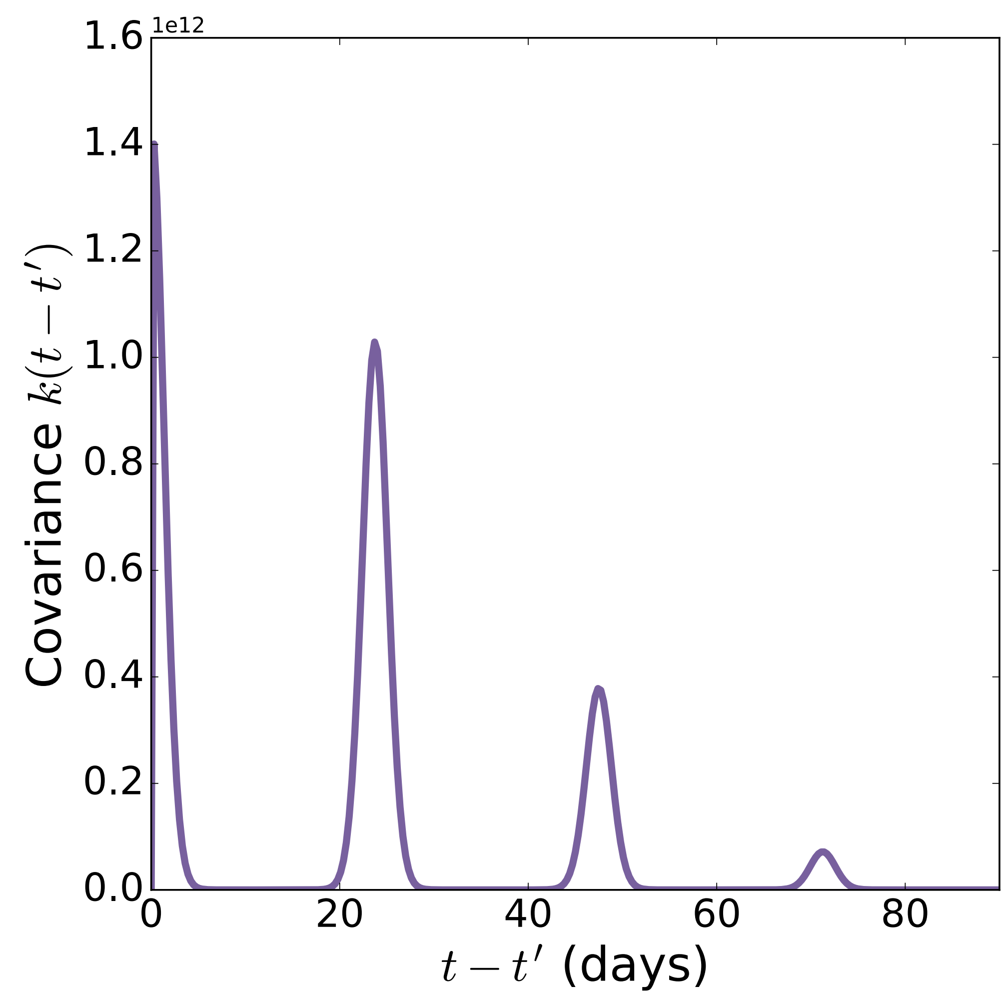

Quasi-periodic form: $k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2} \right]$

See Rassmussen & Williams (2006), Haywood (2015, Chap. 2) and others

The covariance function encodes the correlation/covariance/frequency structure of the data series

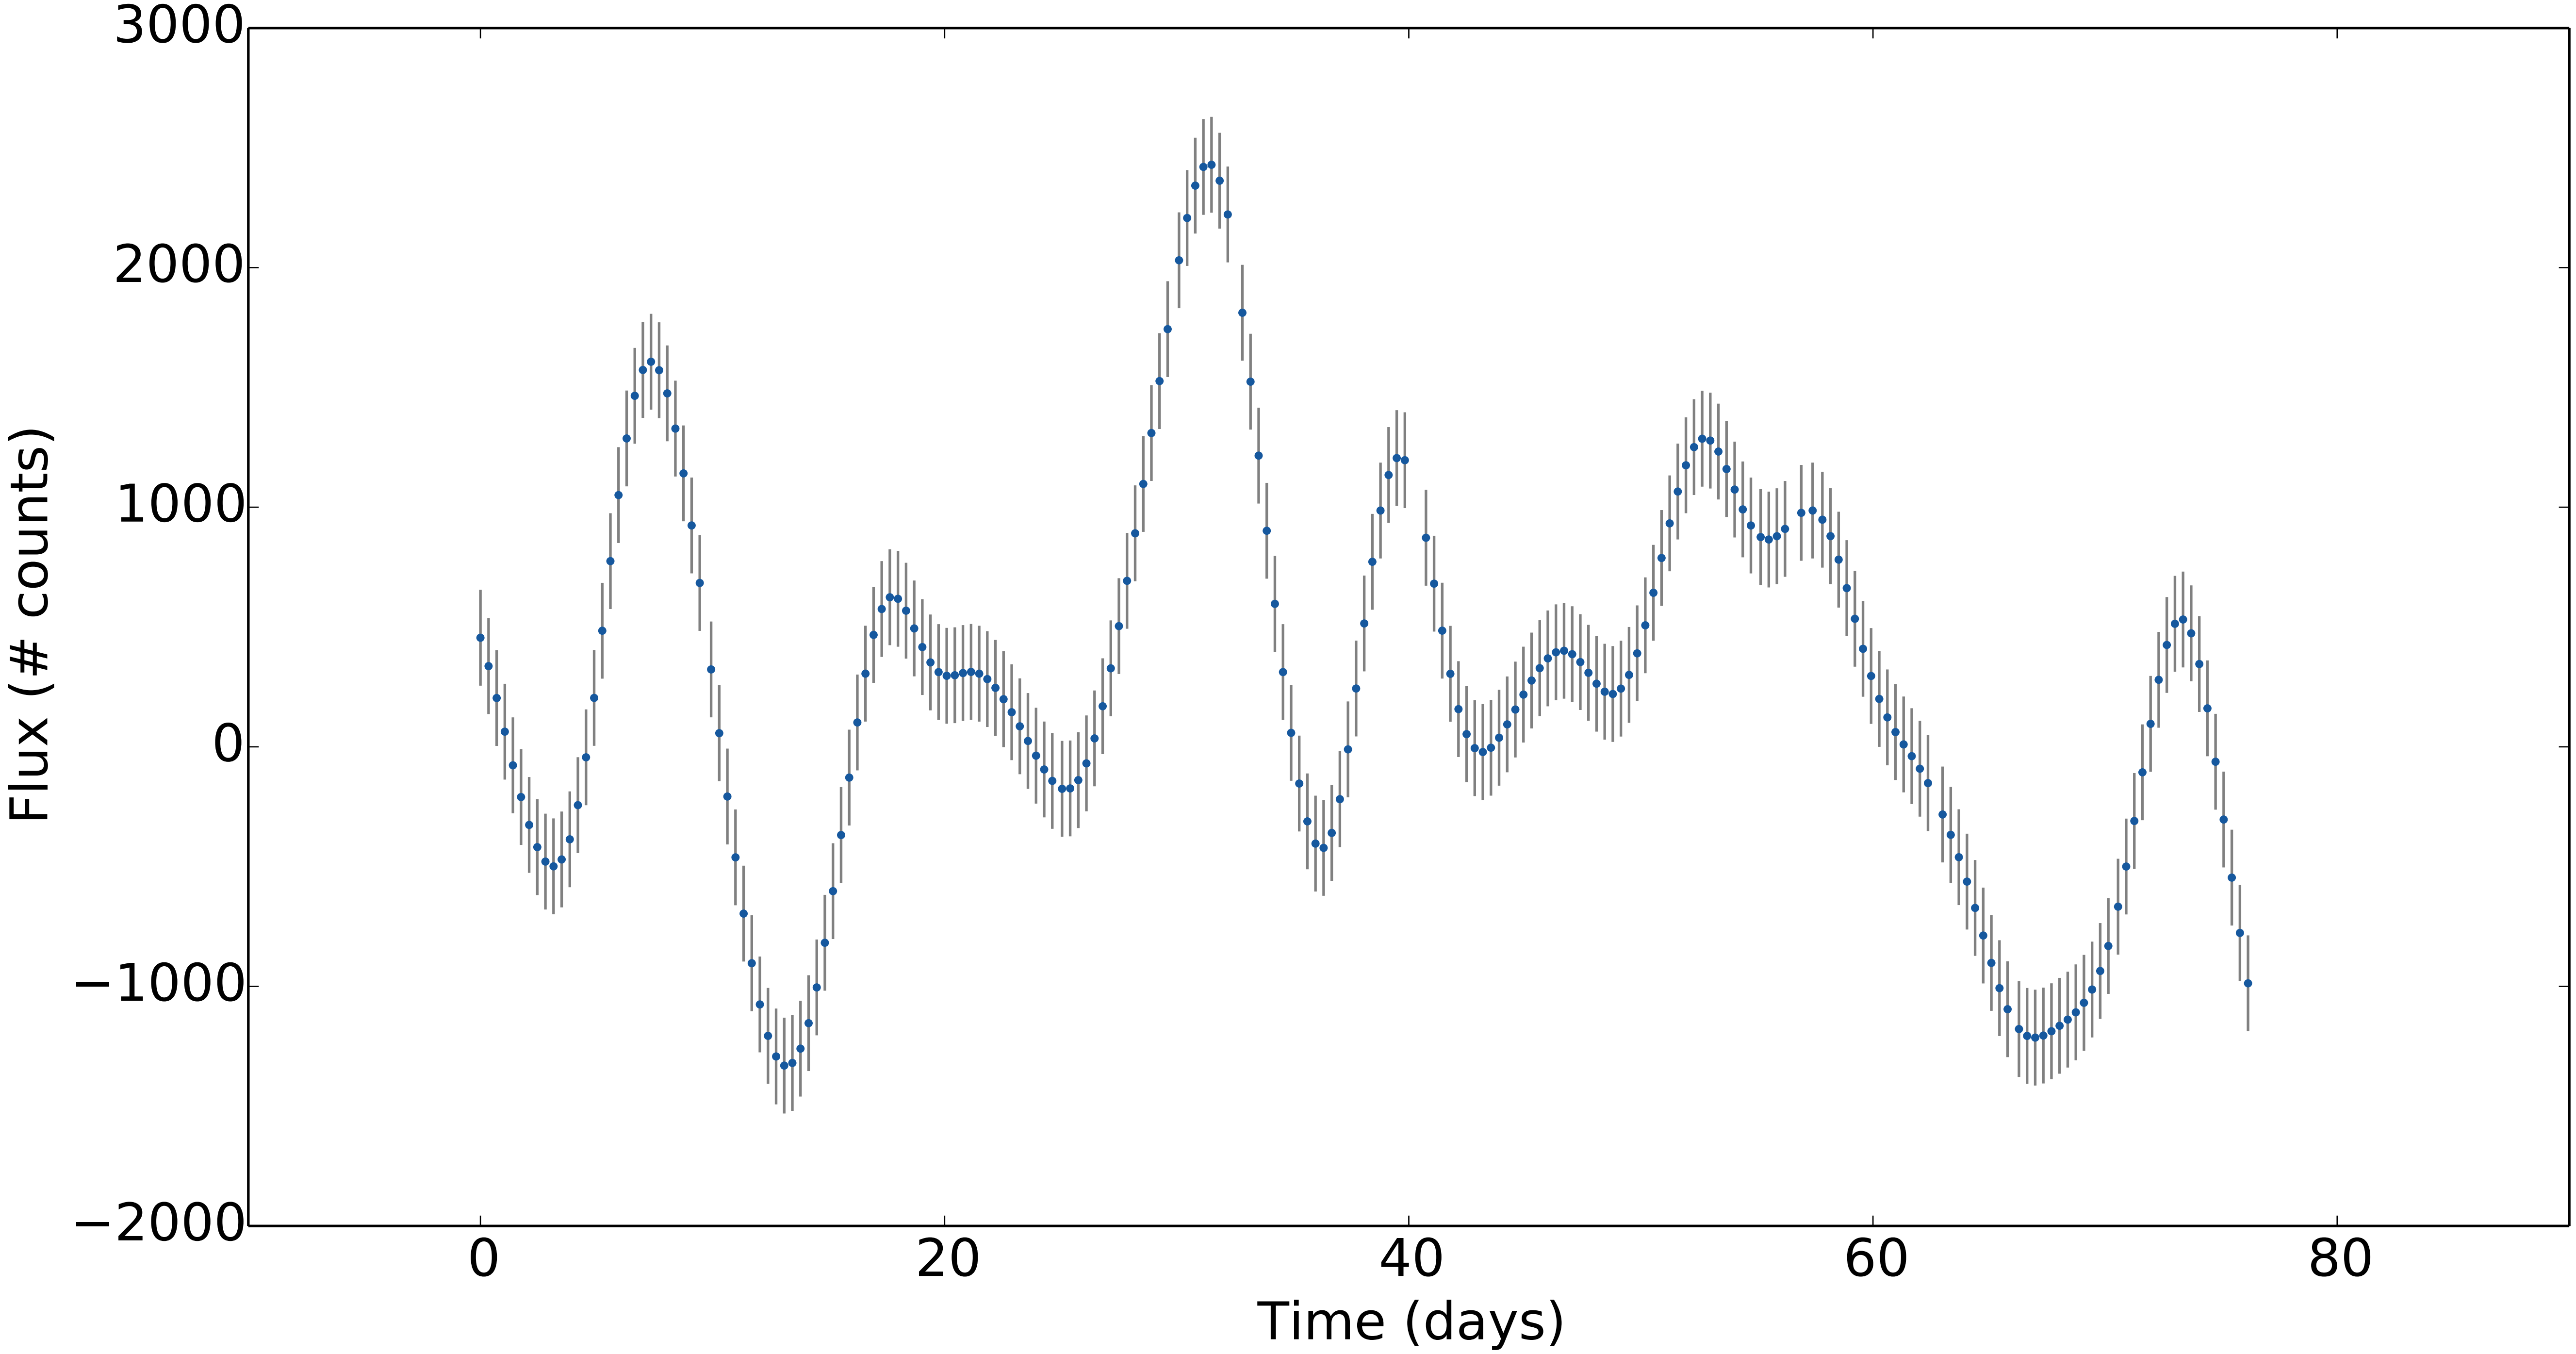

Image credit: C. Nava

The covariance function encodes the correlation/covariance/frequency structure of the data series

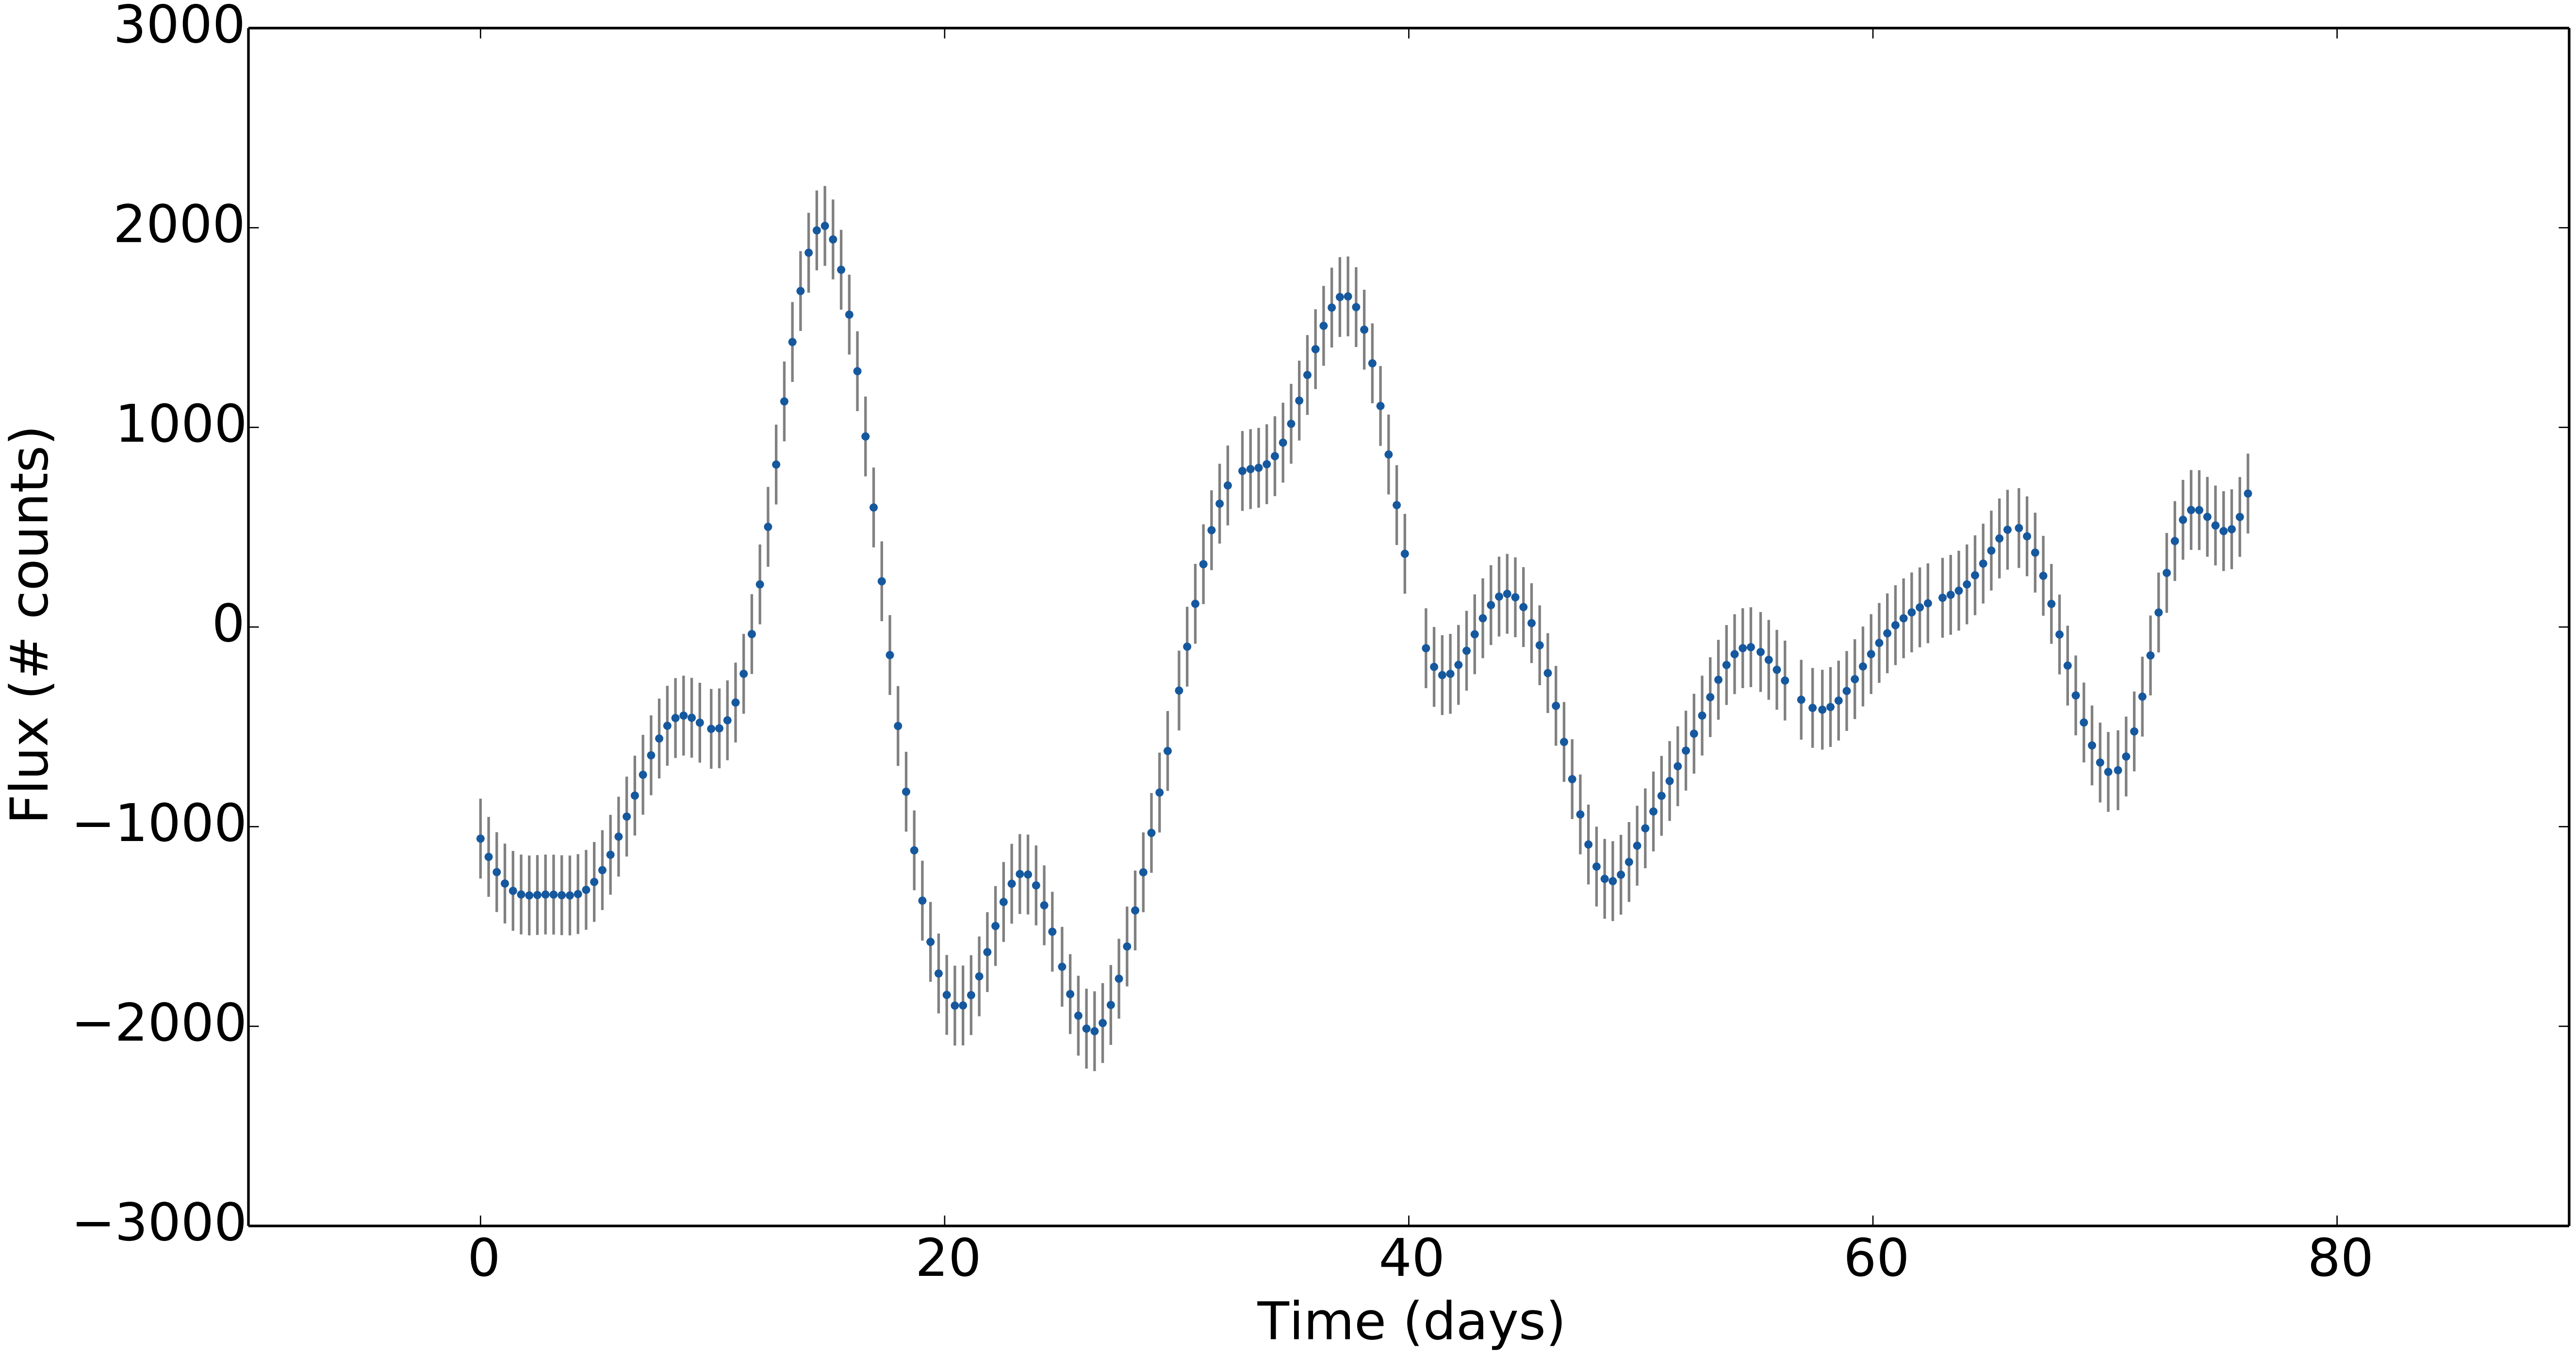

Image credit: C. Nava

The covariance function encodes the correlation/covariance/frequency structure of the data series

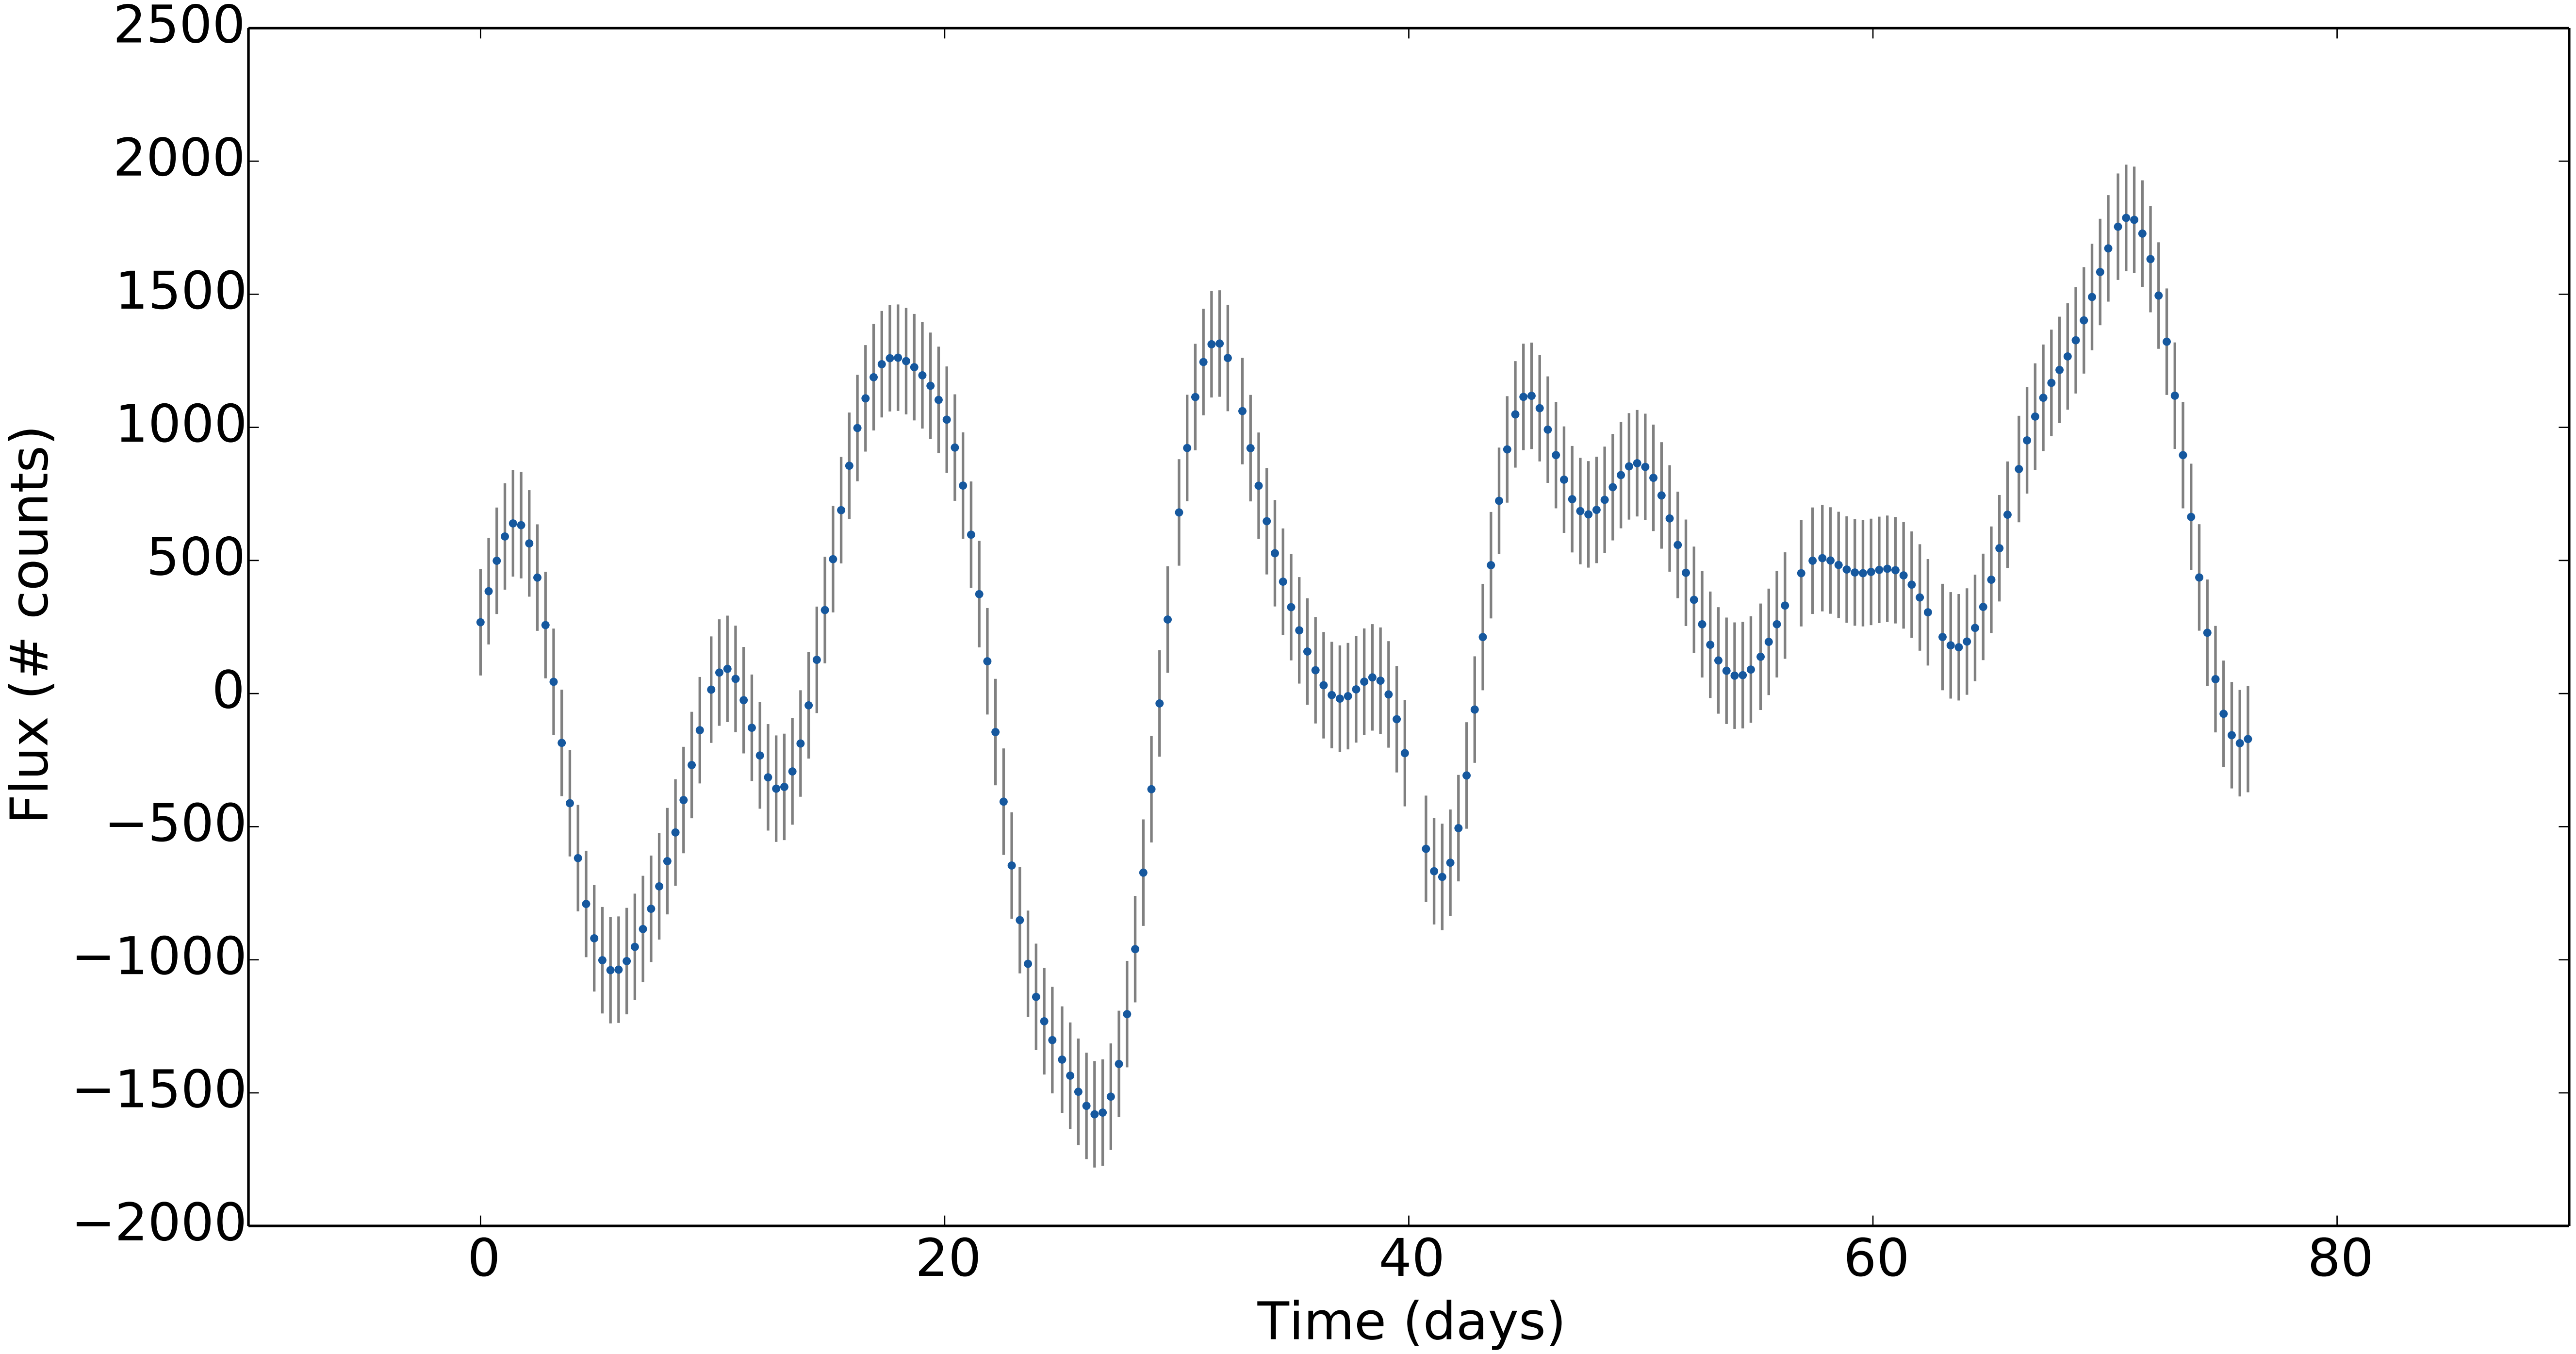

Image credit: C. Nava

Why is this form of covariance function adequate to represent stellar activity?

Why is this form of covariance function adequate to represent stellar activity?

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

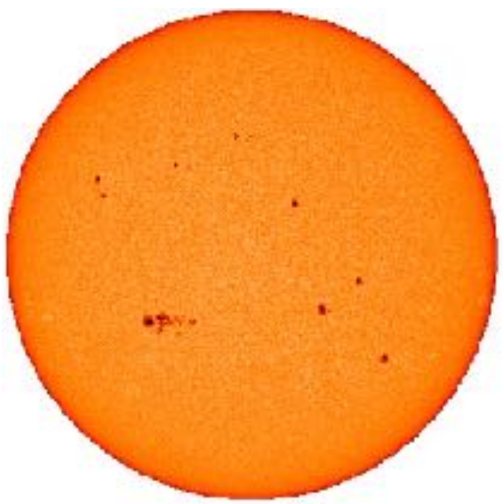

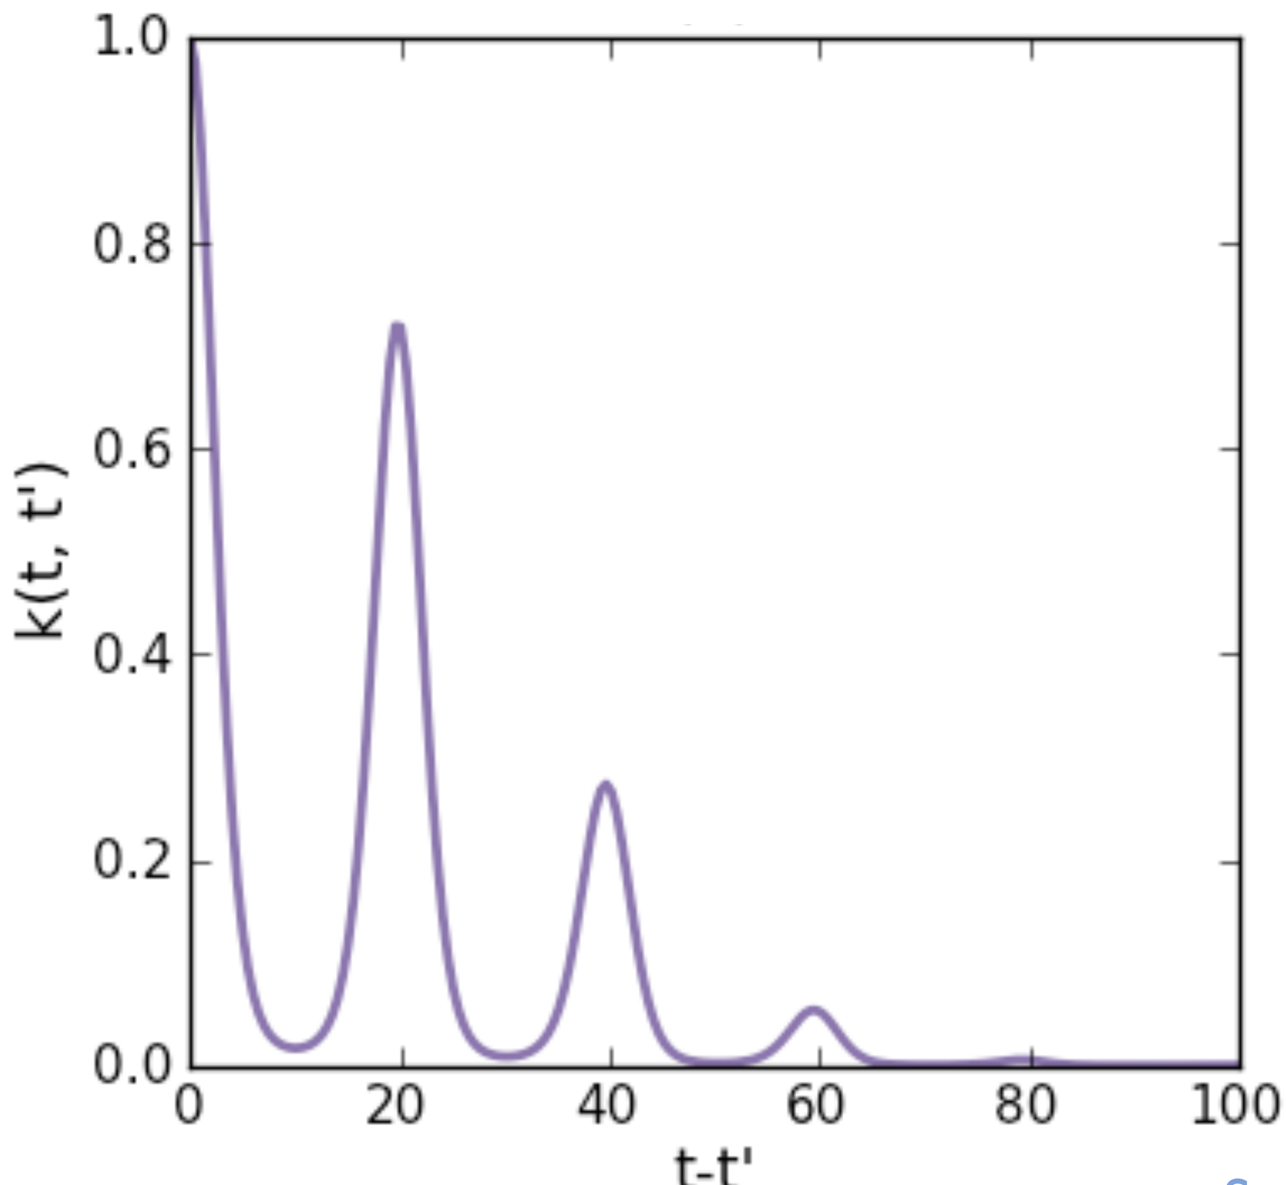

See Haywood et al. (2014, 2018),
López-Morales et al. (2016) and others

Why is this form of covariance function adequate to represent stellar activity?

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2} \right]$$

recurrence timescale P_{rot}

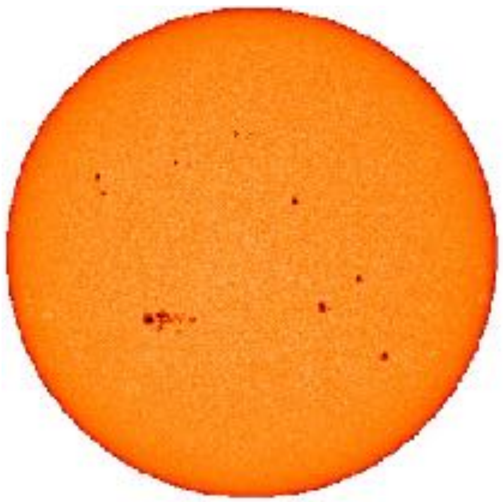

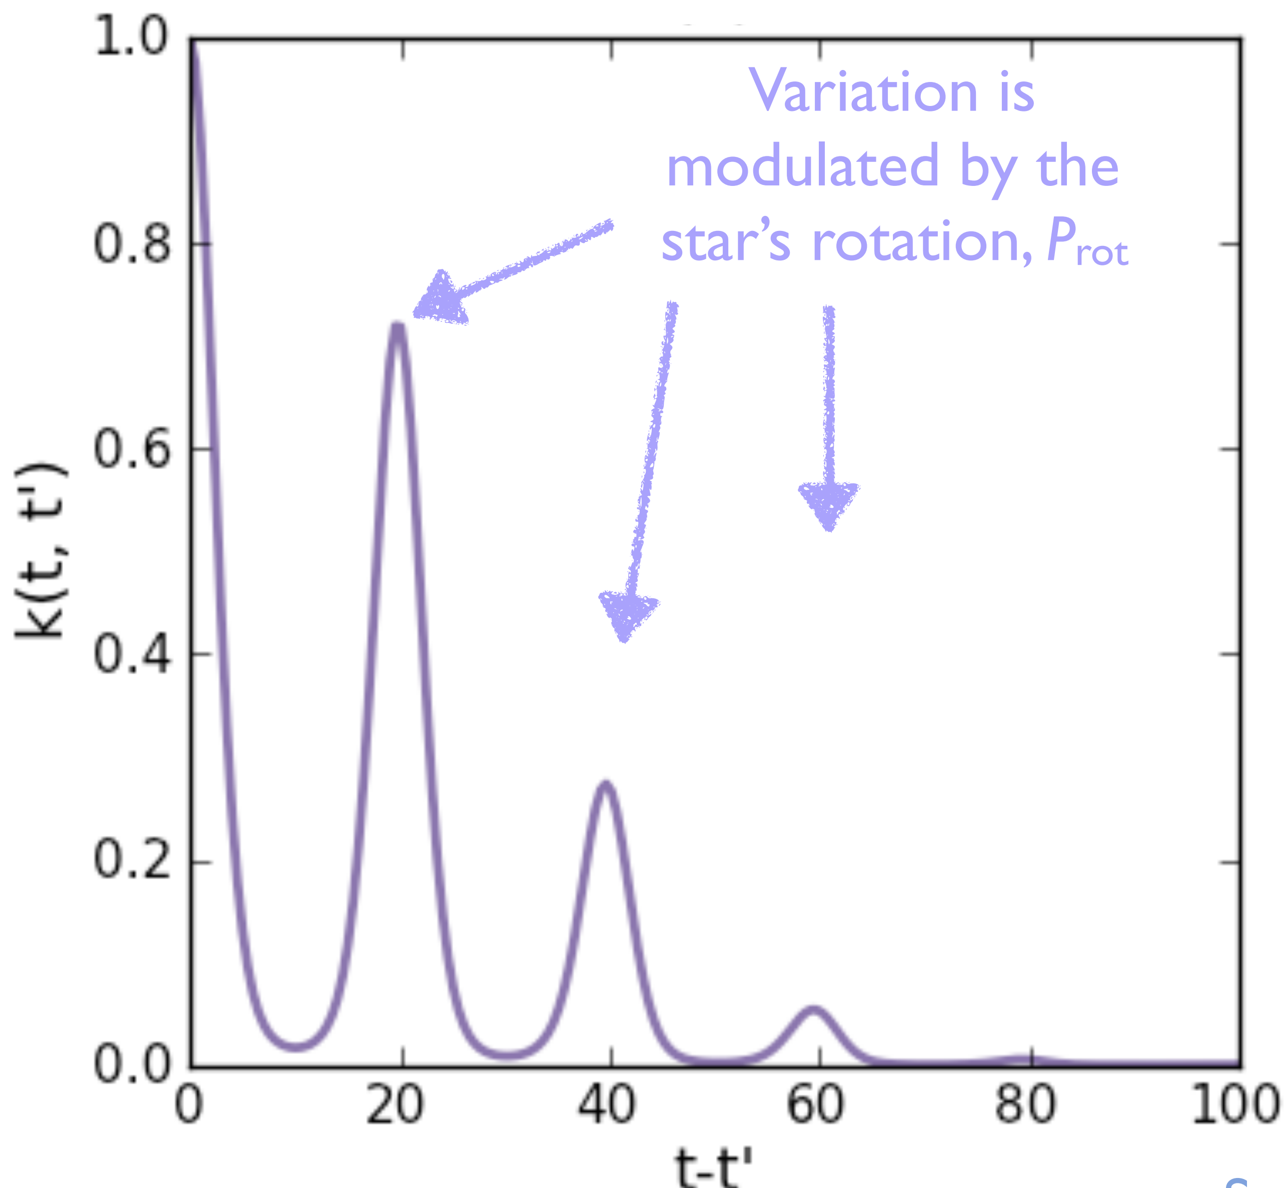

See Haywood et al. (2014, 2018), López-Morales et al. (2016) and others

Why is this form of covariance function adequate to represent stellar activity?

$$k(t, t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2} \right]$$

decay timescale

recurrence timescale P_{rot}

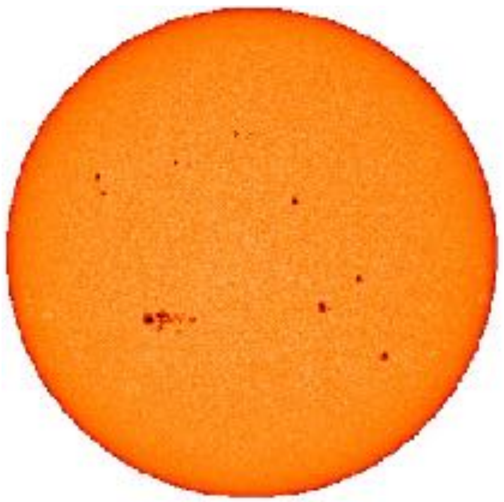

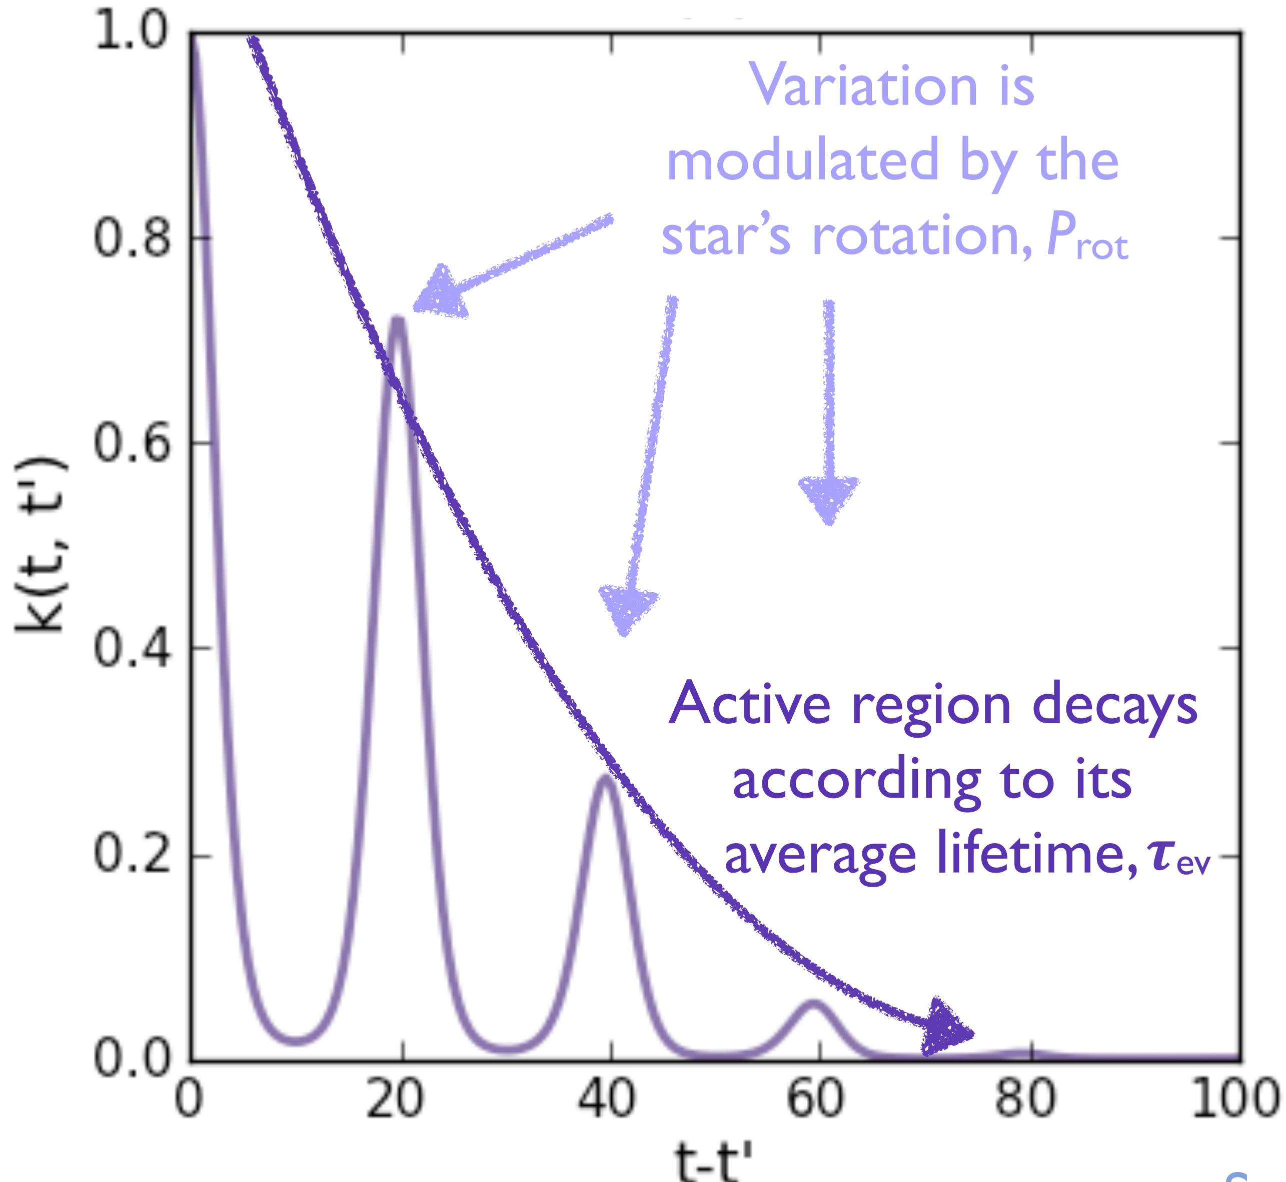

See Haywood et al. (2014, 2018), López-Morales et al. (2016) and others

Why is this form of covariance function adequate to represent stellar activity?

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

decay timescale

recurrence timescale P_{rot}

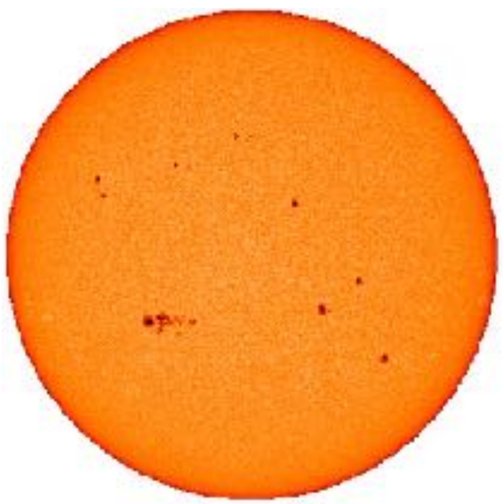

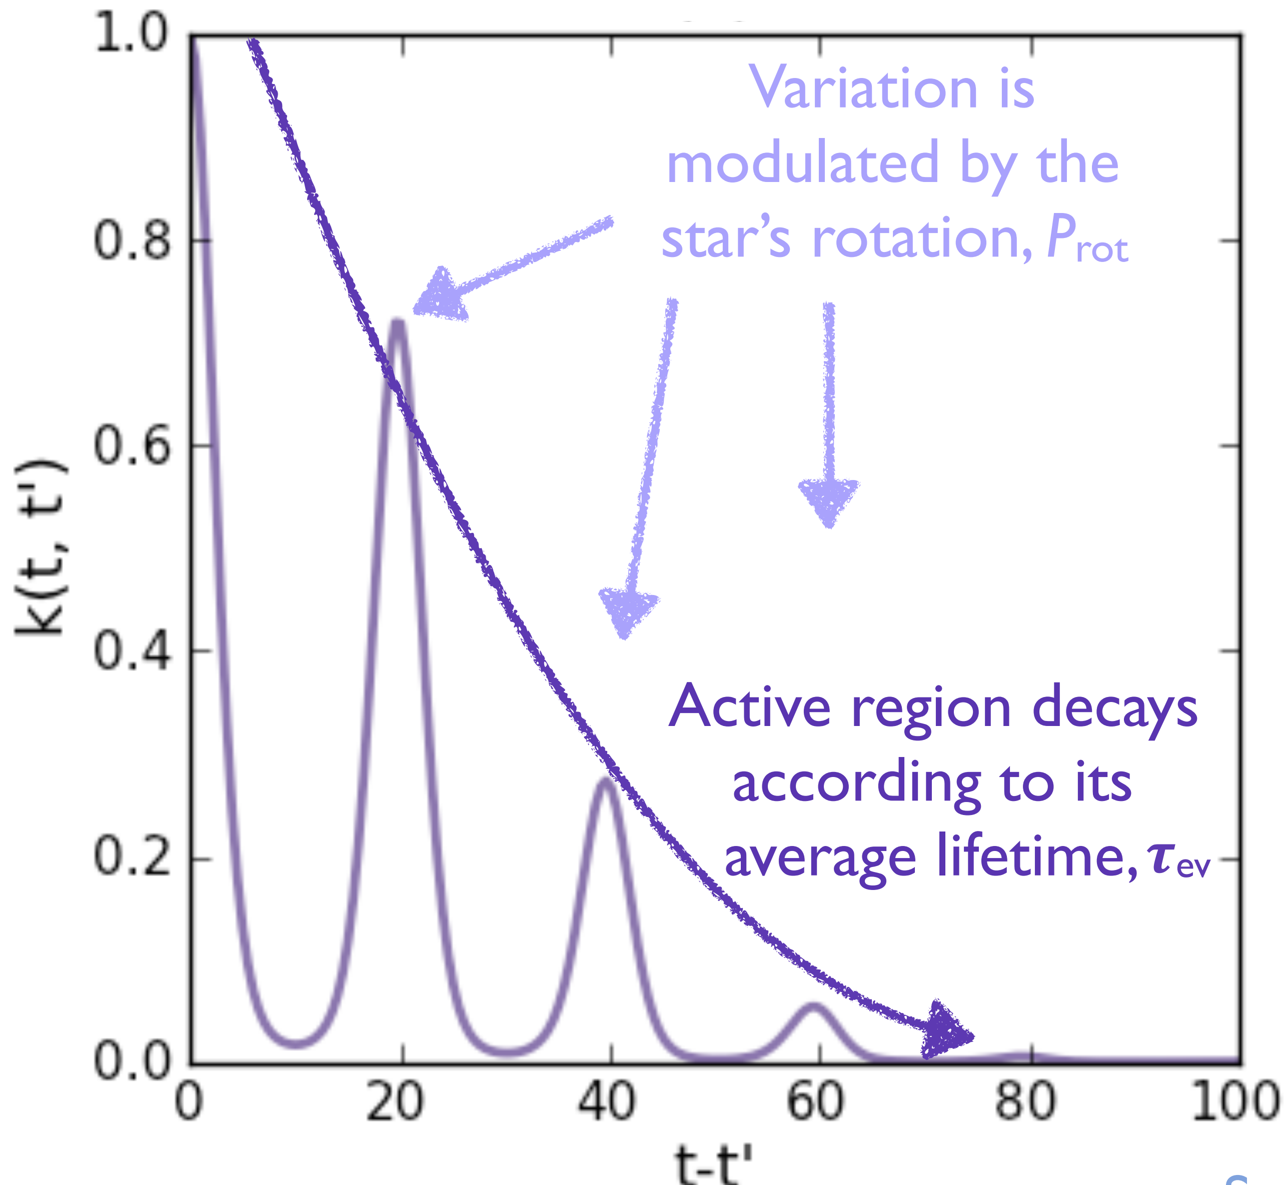

Active-region lifetimes of FGKM stars: see Giles, Collier Cameron & Haywood (2017)

See Haywood et al. (2014, 2018), López-Morales et al. (2016) and others

Why is this form of covariance function adequate to represent stellar activity?

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

decay timescale

recurrence timescale P_{rot}

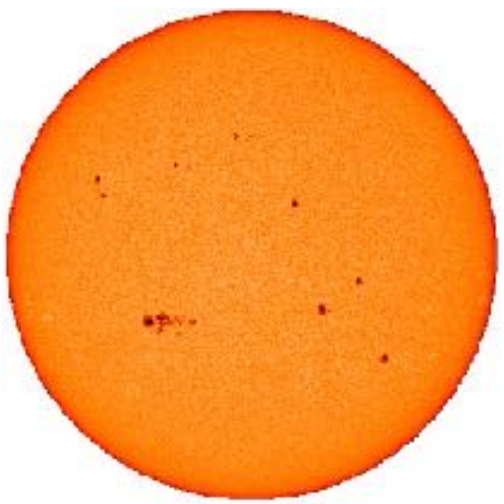

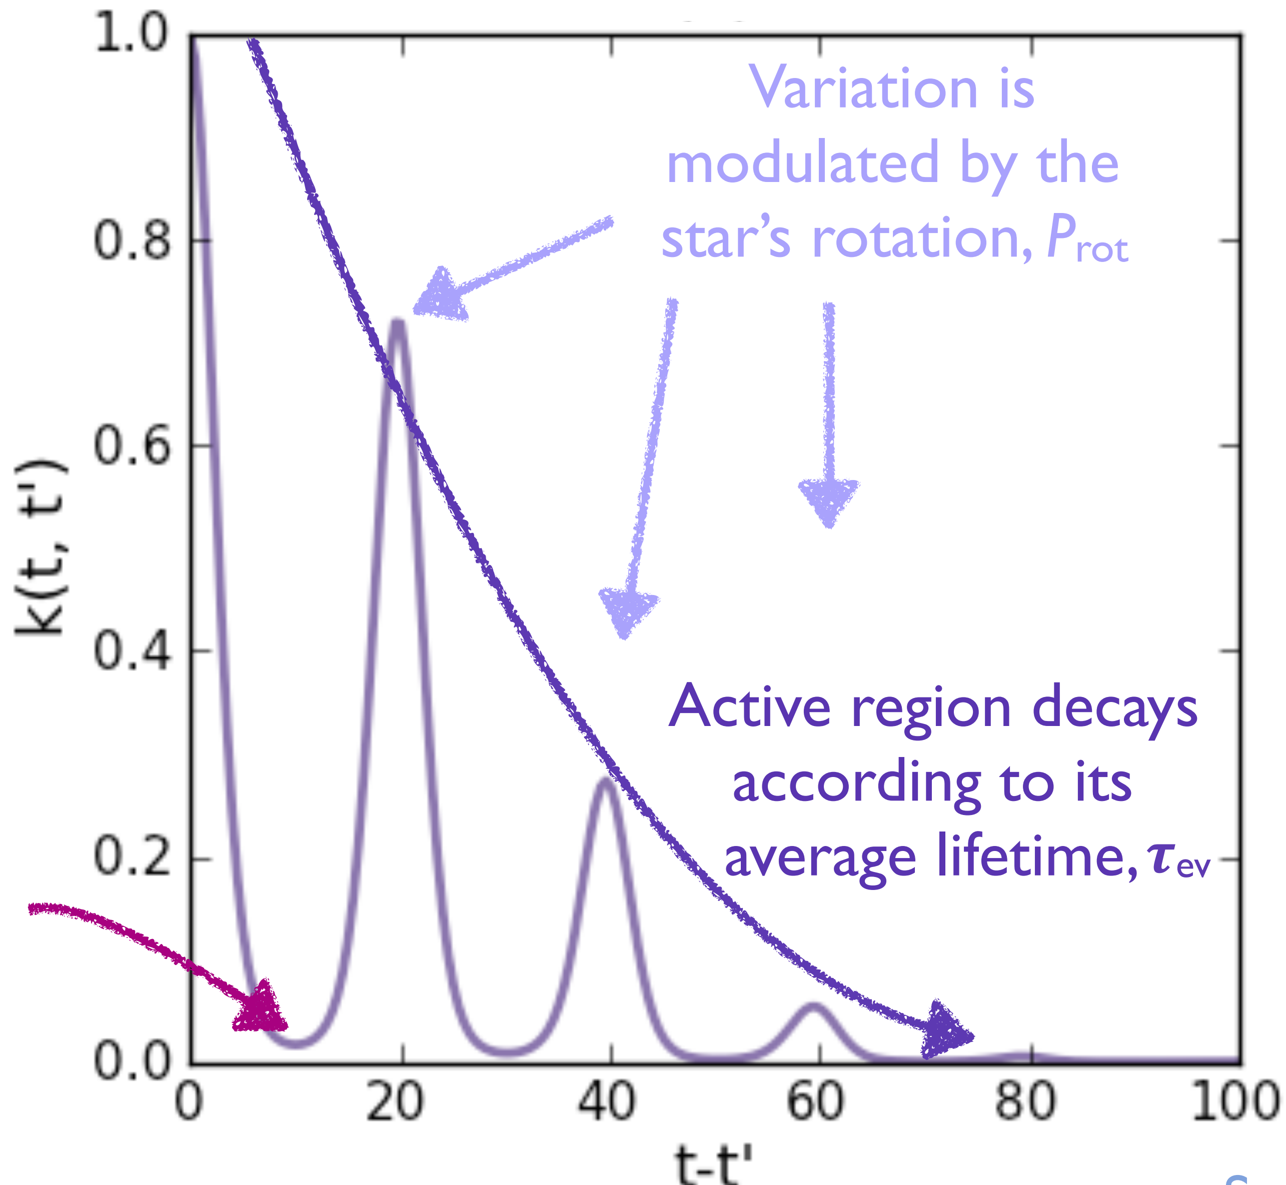

Active region is on the invisible side of the star

See Haywood et al. (2014, 2018), López-Morales et al. (2016) and others

Why is this form of covariance function adequate to represent stellar activity?

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

decay timescale

"smoothing" factor

recurrence timescale P_{rot}

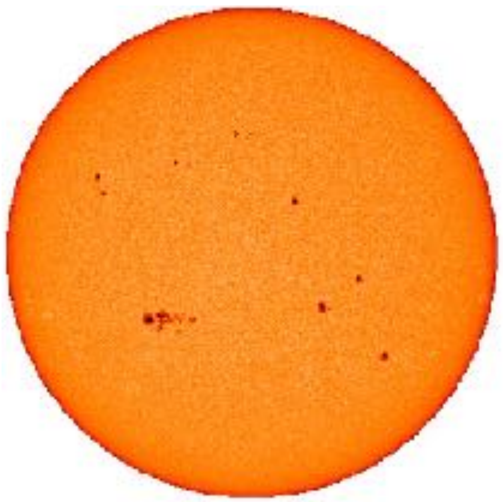

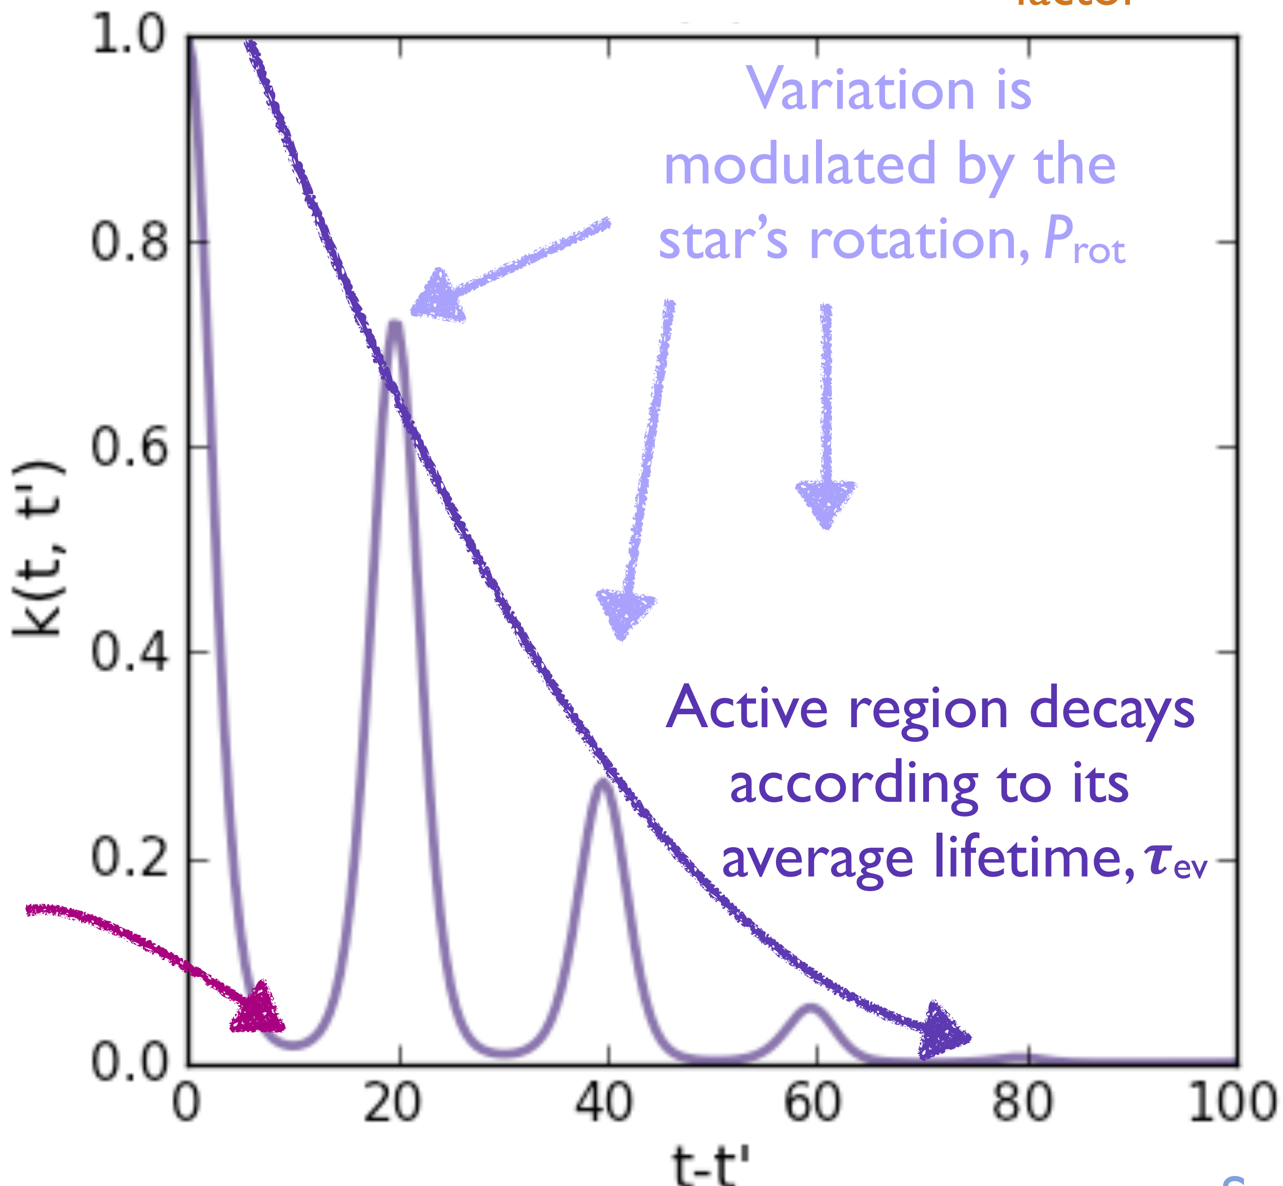

Active region is on the invisible side of the star

See Haywood et al. (2014, 2018), López-Morales et al. (2016) and others

How can we constrain η_4 ?

Jeffers & Keller (2009)

Synthetic stellar surface

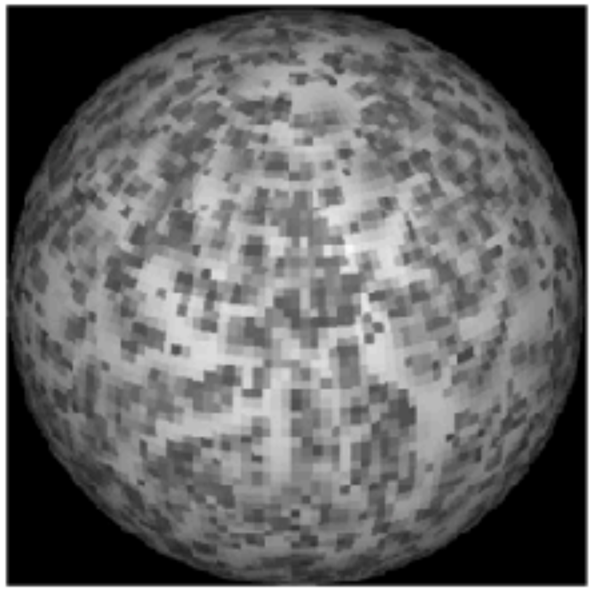

How can we constrain η_4 ?

Jeffers & Keller (2009)

Synthetic stellar surface

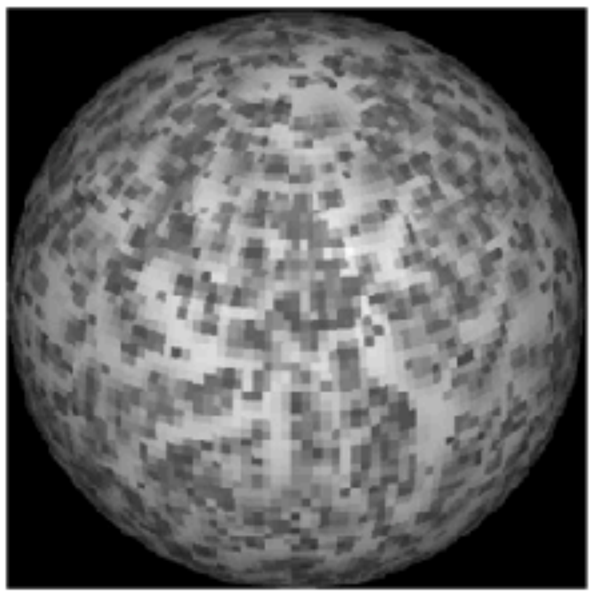

Reconstructed lightcurve

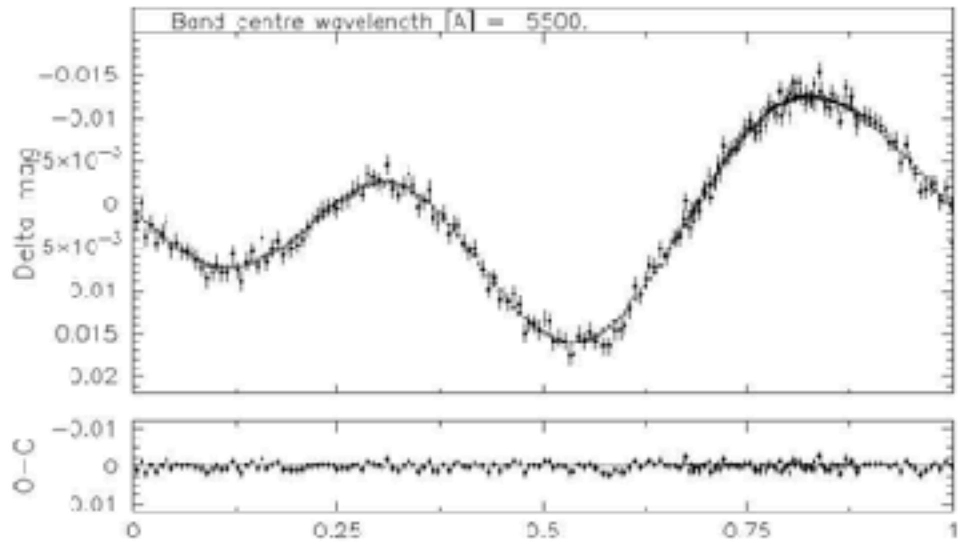

One rotation period

How can we constrain η_4 ?

Jeffers & Keller (2009)

Synthetic stellar surface

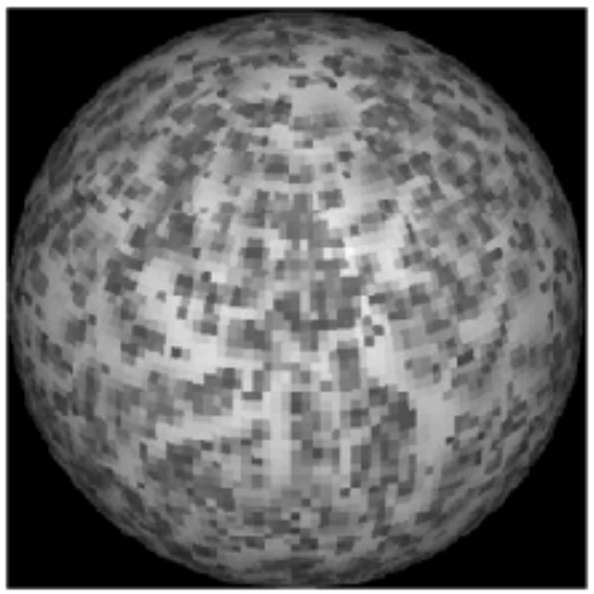

Reconstructed lightcurve

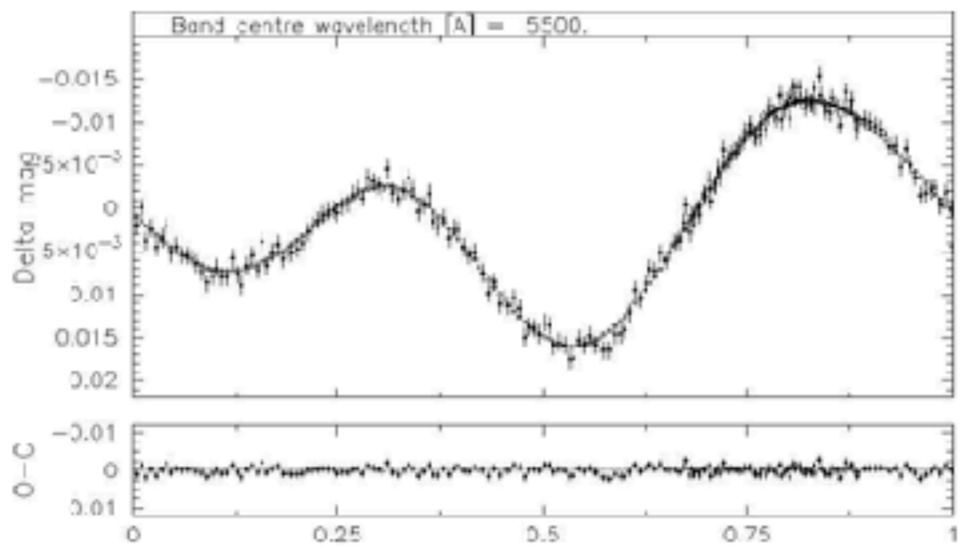

← One rotation period →

Reconstructed surface map

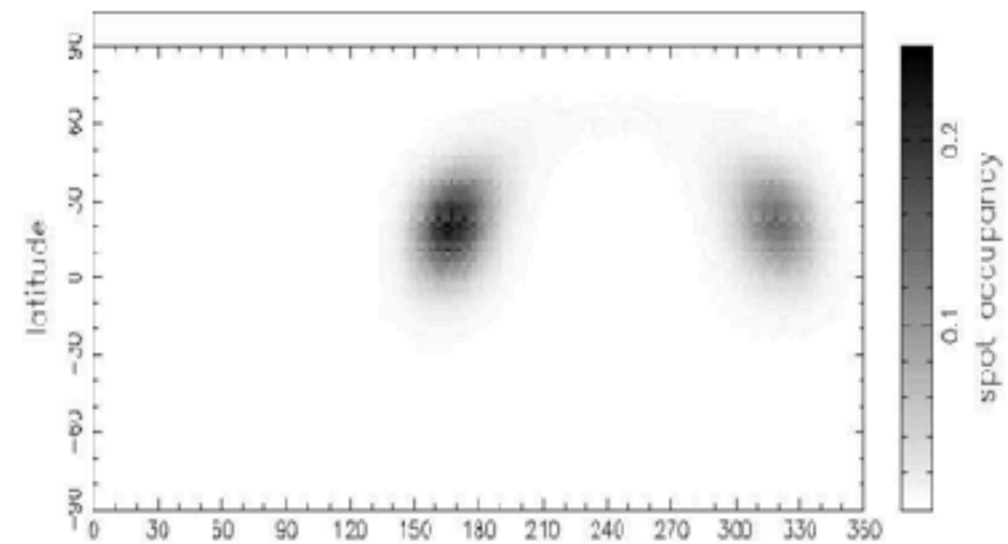

← Longitude →

How can we constrain η_4 ?

Jeffers & Keller (2009)

Synthetic stellar surface

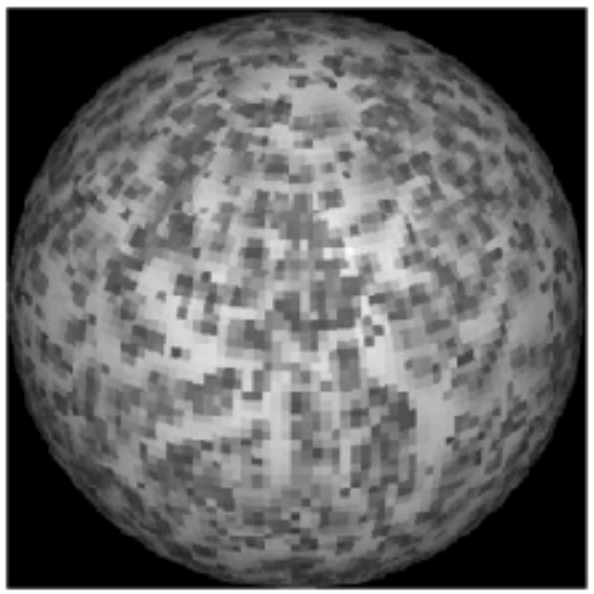

Reconstructed lightcurve

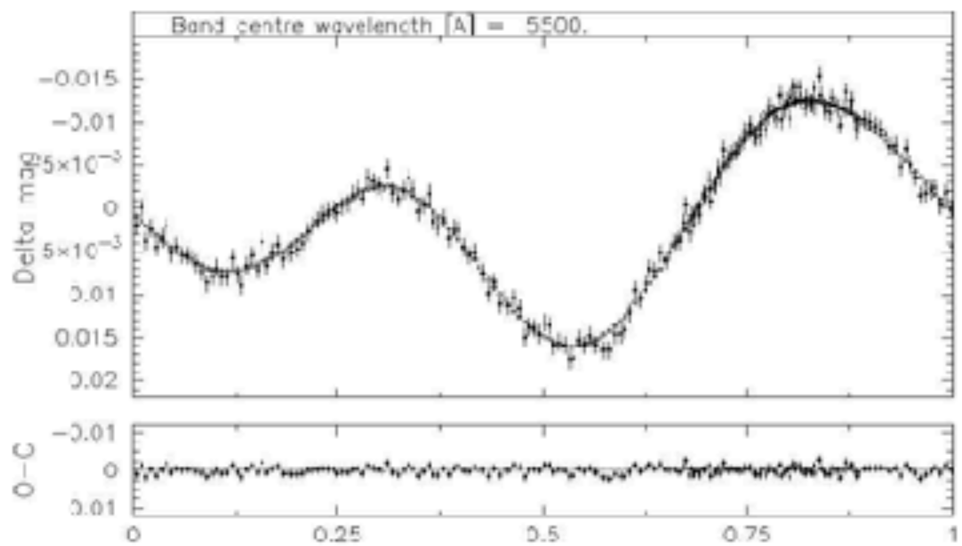

One rotation period

Reconstructed surface map

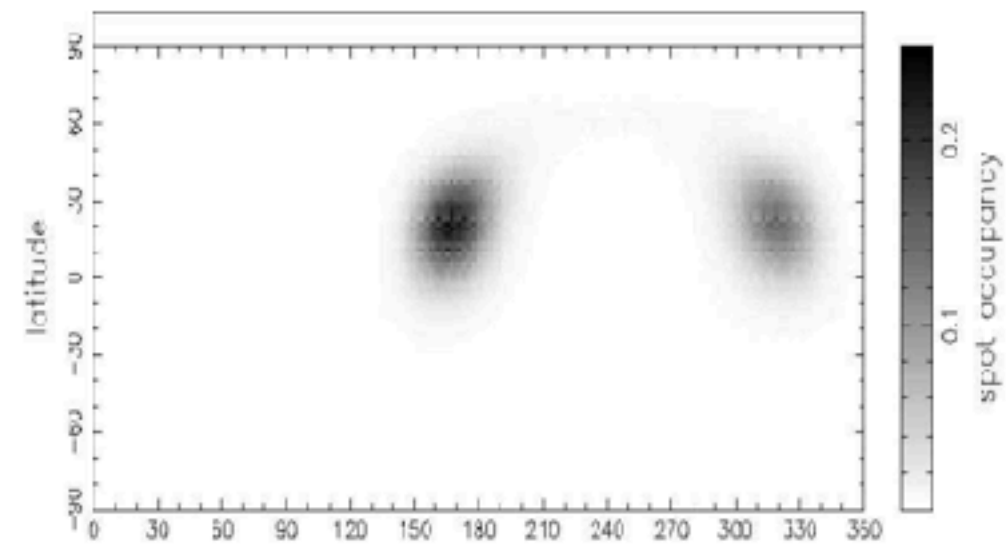

Longitude

A lightcurve, or an RV curve, will only ever show 2-3 peaks per stellar rotation.

This is equivalent to $\eta_4 \approx 0.5$.

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

Why is this form of covariance function adequate to represent stellar activity?

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2} \right]$$

amplitude (η_1)

decay timescale (η_2)

"smoothing" factor (η_4)

recurrence timescale P_{rot} (η_3)

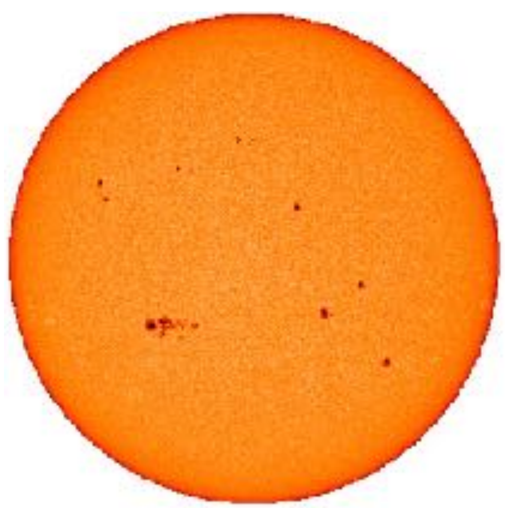

Active region is on the invisible side of the star

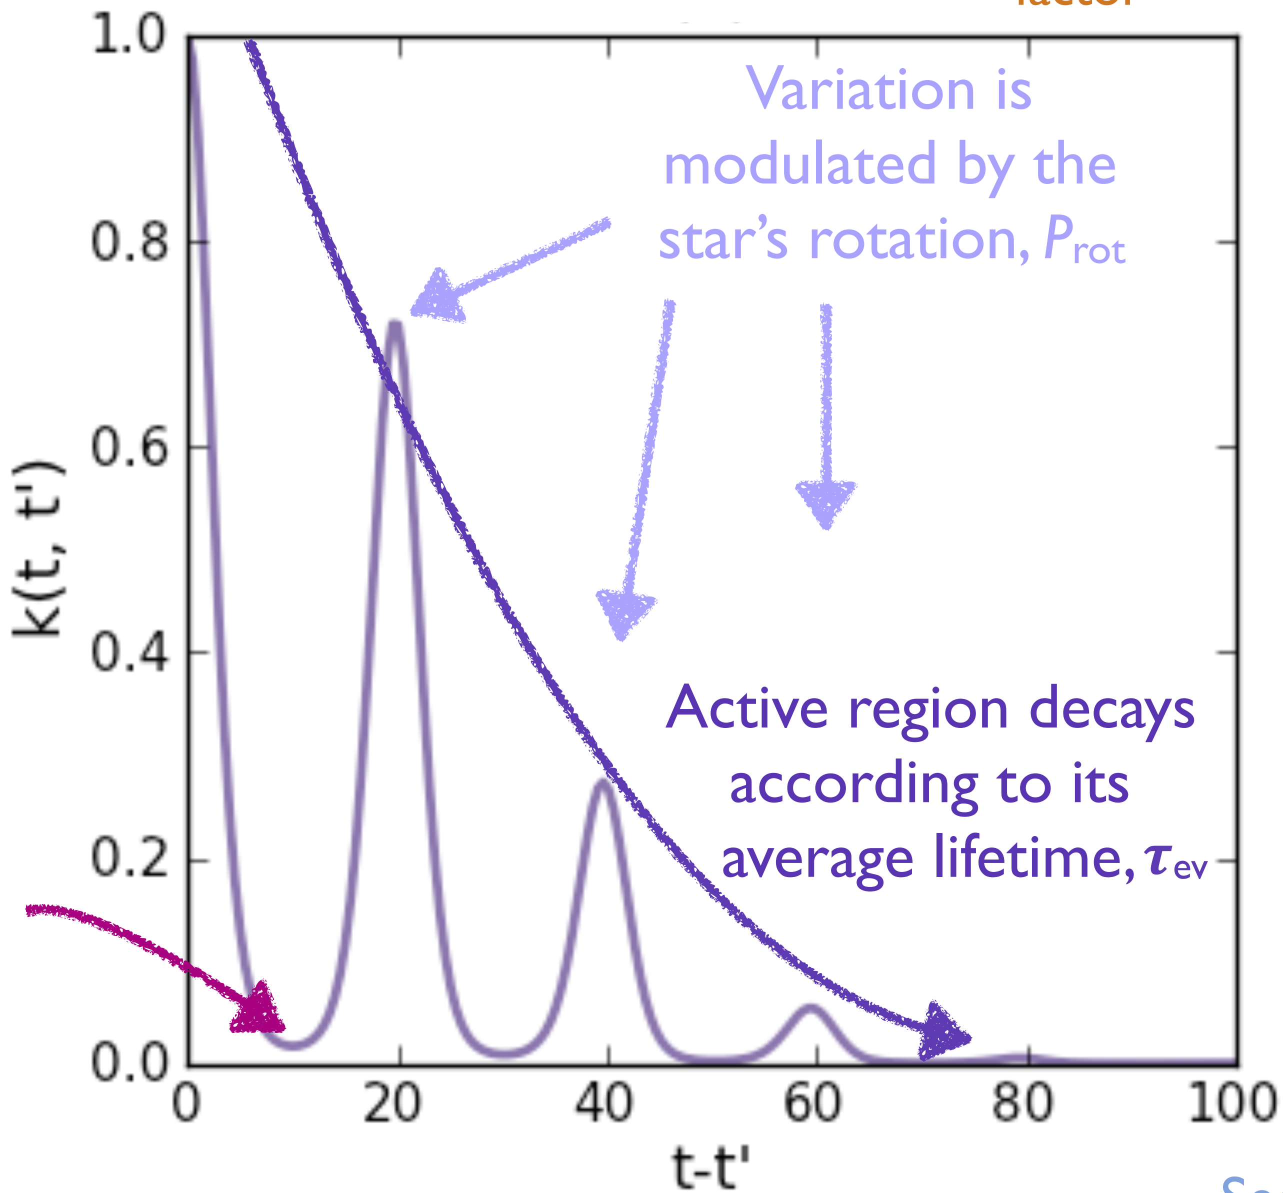

See Haywood et al. (2014, 2018), López-Morales et al. (2016) and others

Determining the hyperparameters

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

amplitude (η_1)

decay timescale (η_2)

"smoothing" factor (η_4)

recurrence timescale P_{rot} (η_3)

Determining the hyperparameters

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\,\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

amplitude

decay timescale

"smoothing" factor

recurrence timescale $P_{\rm rot}$

- Can "train" the GP on an auxiliary dataset, eg. the lightcurve (Haywood et al. 2014, see also Grunblatt et al. 2015, Cloutier et al. 2017 and others), the spectroscopic indicators like FWHM, BIS (Rajpaul et al. 2015, Jones et al. 2017 and others), in some cases even the RVs themselves (Faria et al. 2016, Barros et al. 2017)

Determining the hyperparameters

$$k(t,t') = \eta_1^2 \cdot \exp\left[-\frac{(t-t')^2}{2\,\eta_2^2} - \frac{2\sin^2\left(\frac{\pi(t-t')}{\eta_3}\right)}{\eta_4^2}\right]$$

amplitude

decay timescale

“smoothing” factor

recurrence timescale P_{rot}

- Can “train” the GP on an auxiliary dataset, eg. the lightcurve (Haywood et al. 2014, see also Grunblatt et al. 2015, Cloutier et al. 2017 and others), the spectroscopic indicators like FWHM, BIS (Rajpaul et al. 2015, Jones et al. 2017 and others), in some cases even the RVs themselves (Faria et al. 2016, Barros et al. 2017)
- Can “fix” the hyperparameter values using Gaussian priors, based on prior knowledge/analysis (López-Morales et al. 2016, Dittmann et al. 2017, Haywood et al. 2018 and others)

Will a Gaussian process eat my planet?

Haywood et al. (2014), see also Faria et al. (2016)

Will a Gaussian process eat my planet?

It shouldn't.

Haywood et al. (2014), see also Faria et al. (2016)

Will a Gaussian process eat my planet?

It shouldn't.

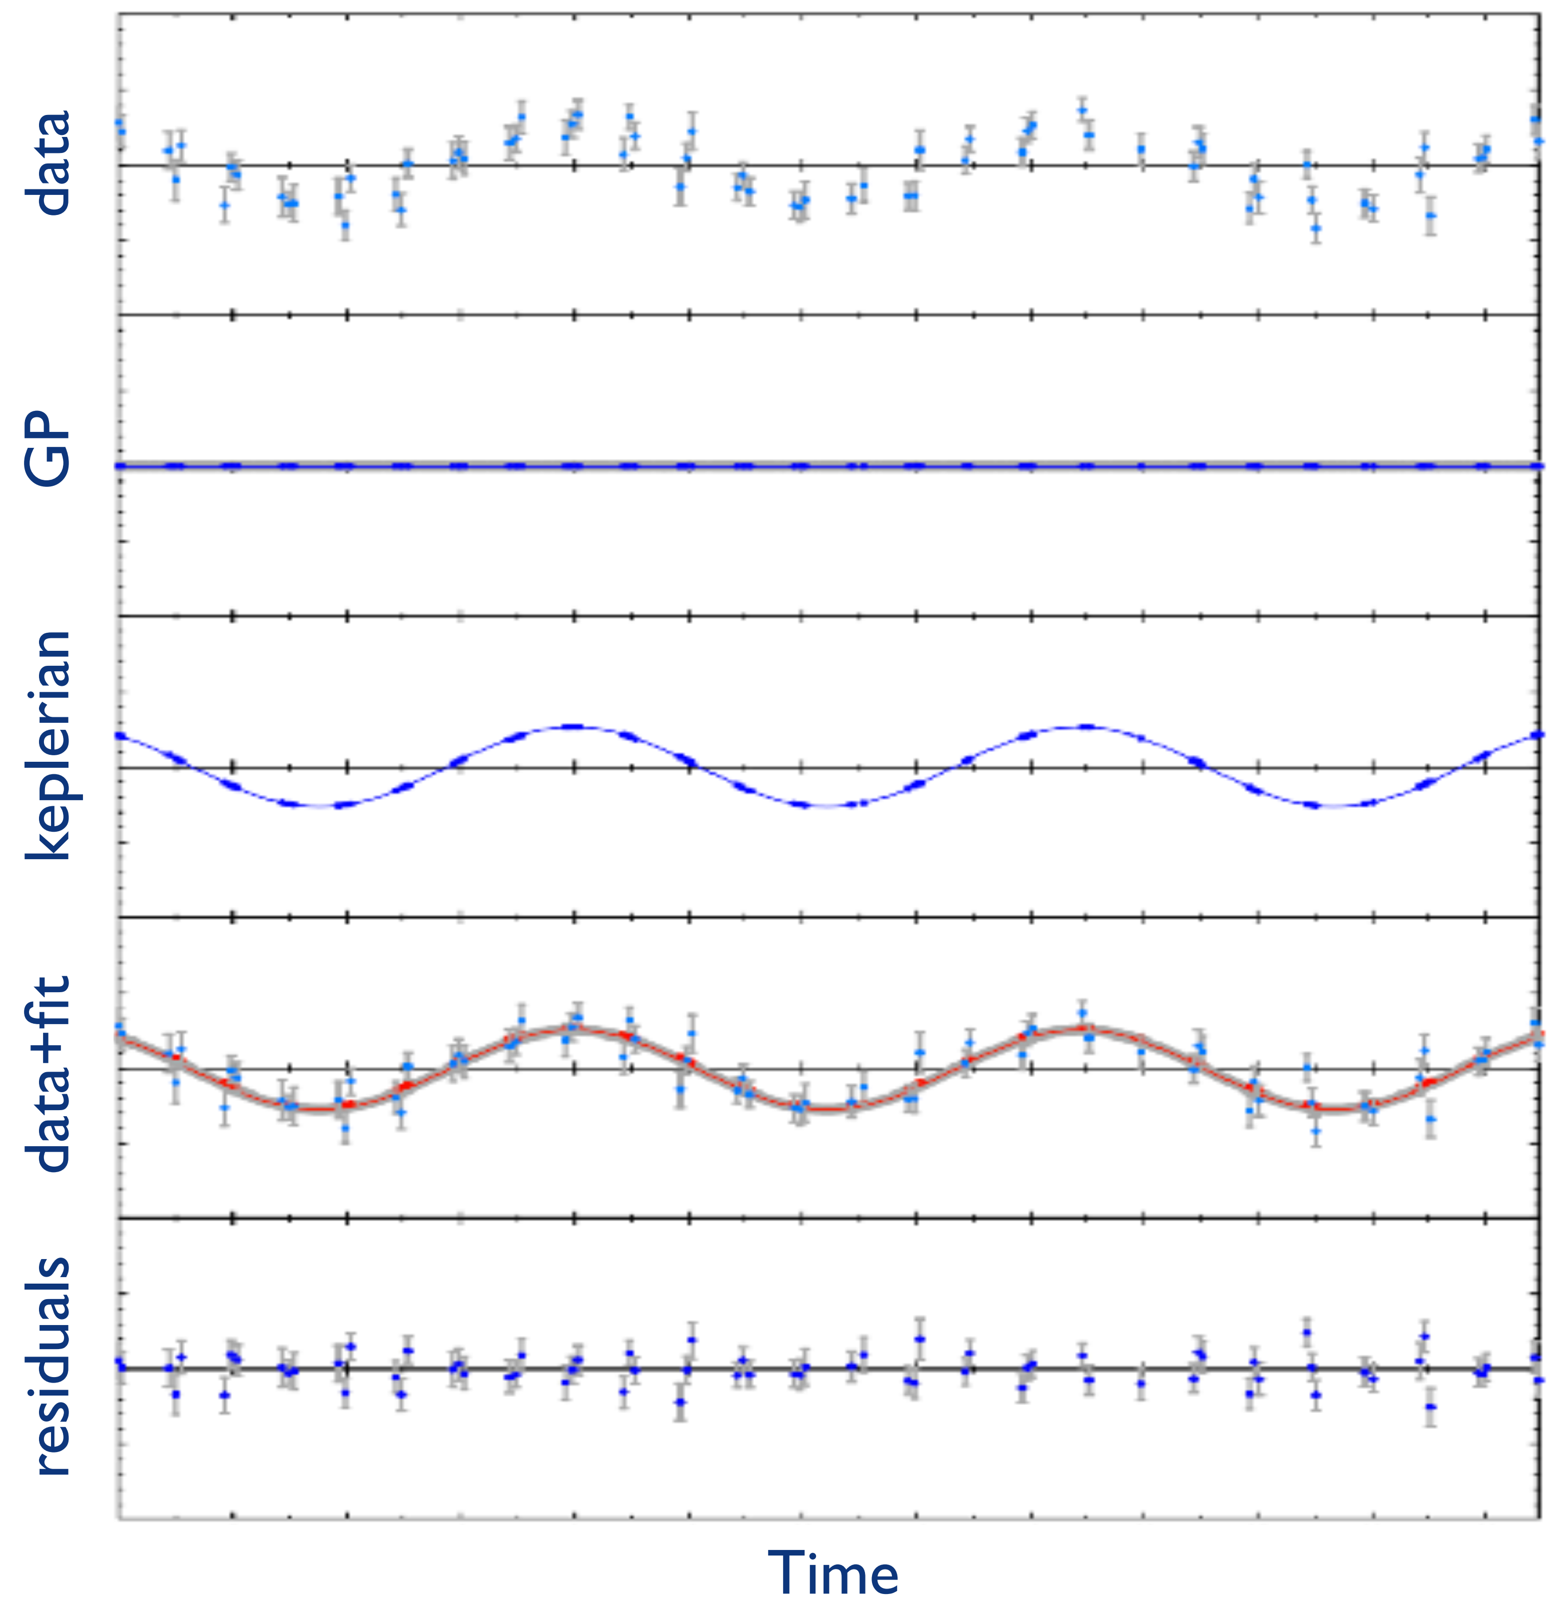

Haywood et al. (2014), see also Faria et al. (2016)

Will a Gaussian process eat my planet?

It shouldn't.

Uncorrelated noise
+ injected Keplerian

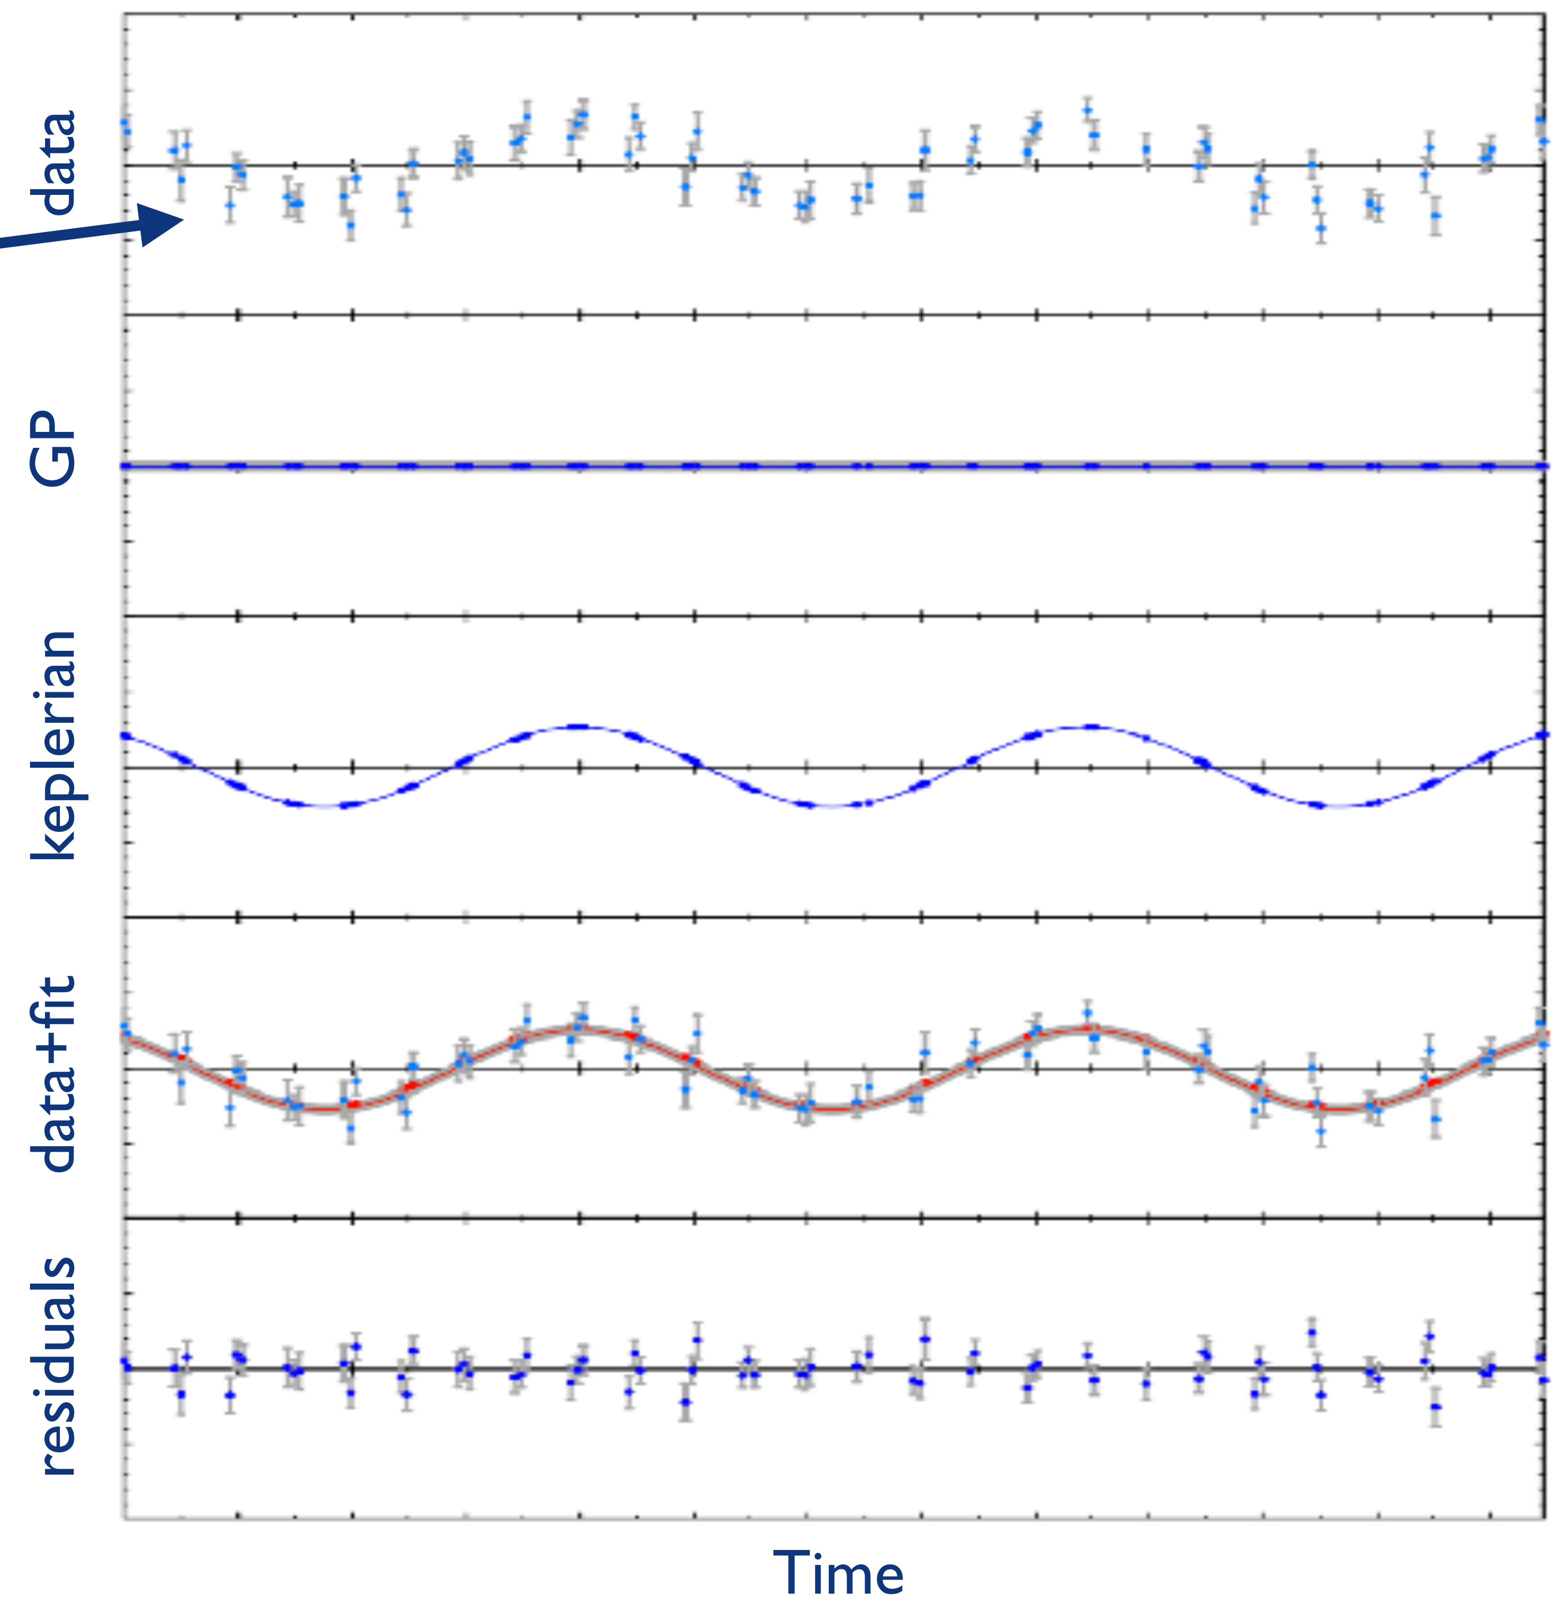

Haywood et al. (2014), see also Faria et al. (2016)

Will a Gaussian process eat my planet?

It shouldn't.

Uncorrelated noise
+ injected Keplerian

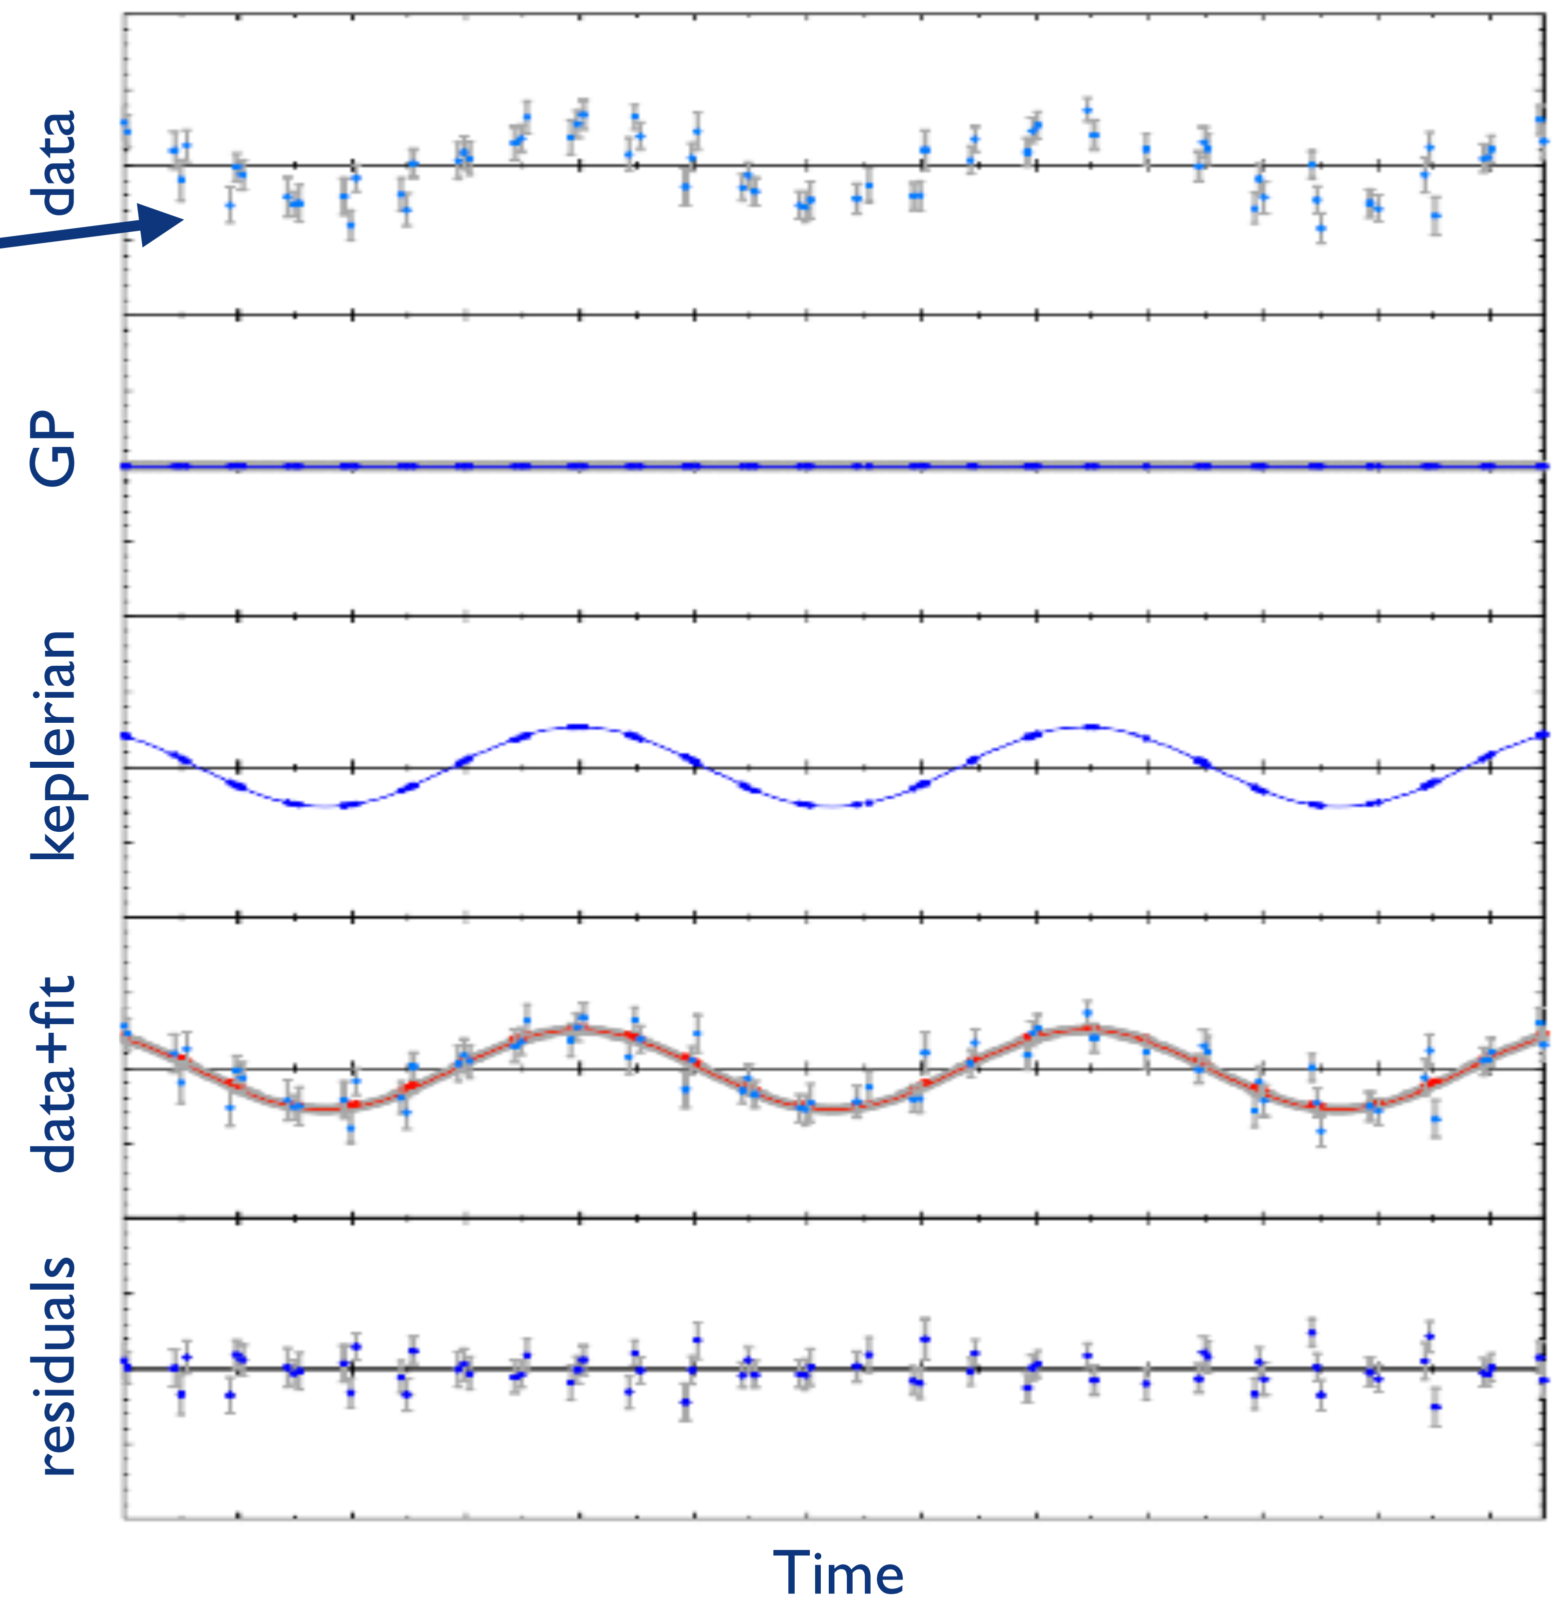

Occam's razor:

The flexibility of the GP is balanced by a penalty term in the likelihood function. The Keplerian is a simple model (less flexible but perfectly adequate for a keplerian signal) so it is favoured by the likelihood function.

Haywood et al. (2014), see also Faria et al. (2016)

Conclusions

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.
- The error bars "live" with the model (not the data).

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.

- The error bars "live" with the model (not the data).

- In the absence of correlated noise, the likelihood "simplifies down" to the chi square.

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.

- The error bars "live" with the model (not the data).

- In the absence of correlated noise, the likelihood "simplifies down" to the chi square.

- Accounting for stellar activity in RV observations using a Gaussian process (or any other correlated noise framework) can make your results more accurate — but not necessarily more precise.

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.
- The error bars "live" with the model (not the data).
- In the absence of correlated noise, the likelihood "simplifies down" to the chi square.
- Accounting for stellar activity in RV observations using a Gaussian process (or any other correlated noise framework) can make your results more accurate — but not necessarily more precise.
- Our choice of covariance function/hyperparameters is crucial. We must think carefully of what physical phenomena/instrumental sources we are trying to account for.

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.

- The error bars "live" with the model (not the data).

- In the absence of correlated noise, the likelihood "simplifies down" to the chi square.

- Accounting for stellar activity in RV observations using a Gaussian process (or any other correlated noise framework) can make your results more accurate — but not necessarily more precise.

- Our choice of covariance function/hyperparameters is crucial. We must think carefully of what physical phenomena/instrumental sources we are trying to account for.

- Statistics tell you about the significance of a signal. Satistics (no matter how advanced or robust) does NOT tell you about the physical origin or nature of a signal.

Conclusions

- Aim for an accurate (*i.e.* unbiased) measurement. Precision is a great advantage, *only if* the measurement is also accurate.

- The error bars "live" with the model (not the data).

- In the absence of correlated noise, the likelihood "simplifies down" to the chi square.

- Accounting for stellar activity in RV observations using a Gaussian process (or any other correlated noise framework) can make your results more accurate — but not necessarily more precise.

- Our choice of covariance function/hyperparameters is crucial. We must think carefully of what physical phenomena/instrumental sources we are trying to account for.

- Statistics tell you about the significance of a signal. Satistics (no matter how advanced or robust) does NOT tell you about the physical origin or nature of a signal.

Papers & textbooks:

- The classic textbook reference (in which you will find all the equations and statistical jargon): C. E. Rasmussen & C. K. I. Williams, Gaussian Processes for Machine Learning, the MIT Press, 2006, (online: www.GaussianProcess.org/gpml).
- A very clear introduction to GPs (not specific to exoplanets): Roberts et al. (2012)
- GPs to account for stellar activity in RV analyses: Baluev (2013), Haywood et al. (2014), Grunblatt et al. (2015), Rajpaul et al. (2015), Faria et al. (2016), Anglada-Escudé et al. (2016), López-Morales, Haywood et al. (2016), Barros et al. (2017), Jones et al. (2017), Cloutier et al. (2017), Haywood et al. (2018) and others.
- GPs applied to transmission spectroscopy for the study of planetary atmospheres: Gibson et al. (2011) and Czekala et al. (2014) and others.
- Machine learning to detrend Kepler and K2 lightcurves: Foreman-Mackey et al. (2015), Crossfield et al. (2015), Foreman-Mackey et al. (2014), Ambikasaran et al. (2014), Aigrain et al. (2015) and Barclay et al. (2015), Armstrong et al. (2016) and others.

Develop your GP intuition:

- This is a fantastic lecture on the nature of Gaussian processes by David MacKay. I thoroughly recommend watching it! http://videolectures.net/gpip06_mackay_gpb/
- Suzanne Aigrain and her group have given many talks and tutorials, all available here: http://splox.net/tag/gps/
- Read Chapter 2 of my PhD thesis (Haywood 2015): https://research-repository.st-andrews.ac.uk/handle/10023/7798?mode=full&submit_simple=Show+full+item+record
- Discussion on "astrophysically-motivated" GPs: Haywood et al. (2018), López-Morales et al. (2016)

Useful codes:

- Dan Foreman-Mackey's *George*, *celerite* (and emcee) codes are publicly available at: http://dan.iel.fm/research/.
- João Faria's *kima*, for exoplanet detection in RVs with DNest4 and GPs: https://github.com/j-faria/kima
- Radvel, a radial velocity modelling toolkit co-written by BJ Fulton, Erik Petigura, Sarah Blunt and Evan Sinukoff: https://radvel.readthedocs.io/en/latest/

The use of GPs in exoplanet science is growing fast; this is only a small selection of papers/codes/etc. and is in no way exhaustive.