Anomalies in Cosmology

Marco Raveri

University of Genova

The plan

Is this the end of cosmic concordance?

* The standard cosmological model
* The Hubble constant tension
* Tension detection

  …when things get complicated…
* Growth tensions

The Standard Cosmological Model

The universe as a physics laboratory

The Planck scale 10^{28} eV — Energy — The Horizon 10^{-33} eV

$\sim 10^{-32}$ seconds | 1 second | 380 000 years | 400 million years | 13.8 billion years

Inflation | **Hydrogen Recombination** | **Dark Matter** | **Dark Energy**

(Based on ESA cosmic history)

High resolution pictures of our Universe

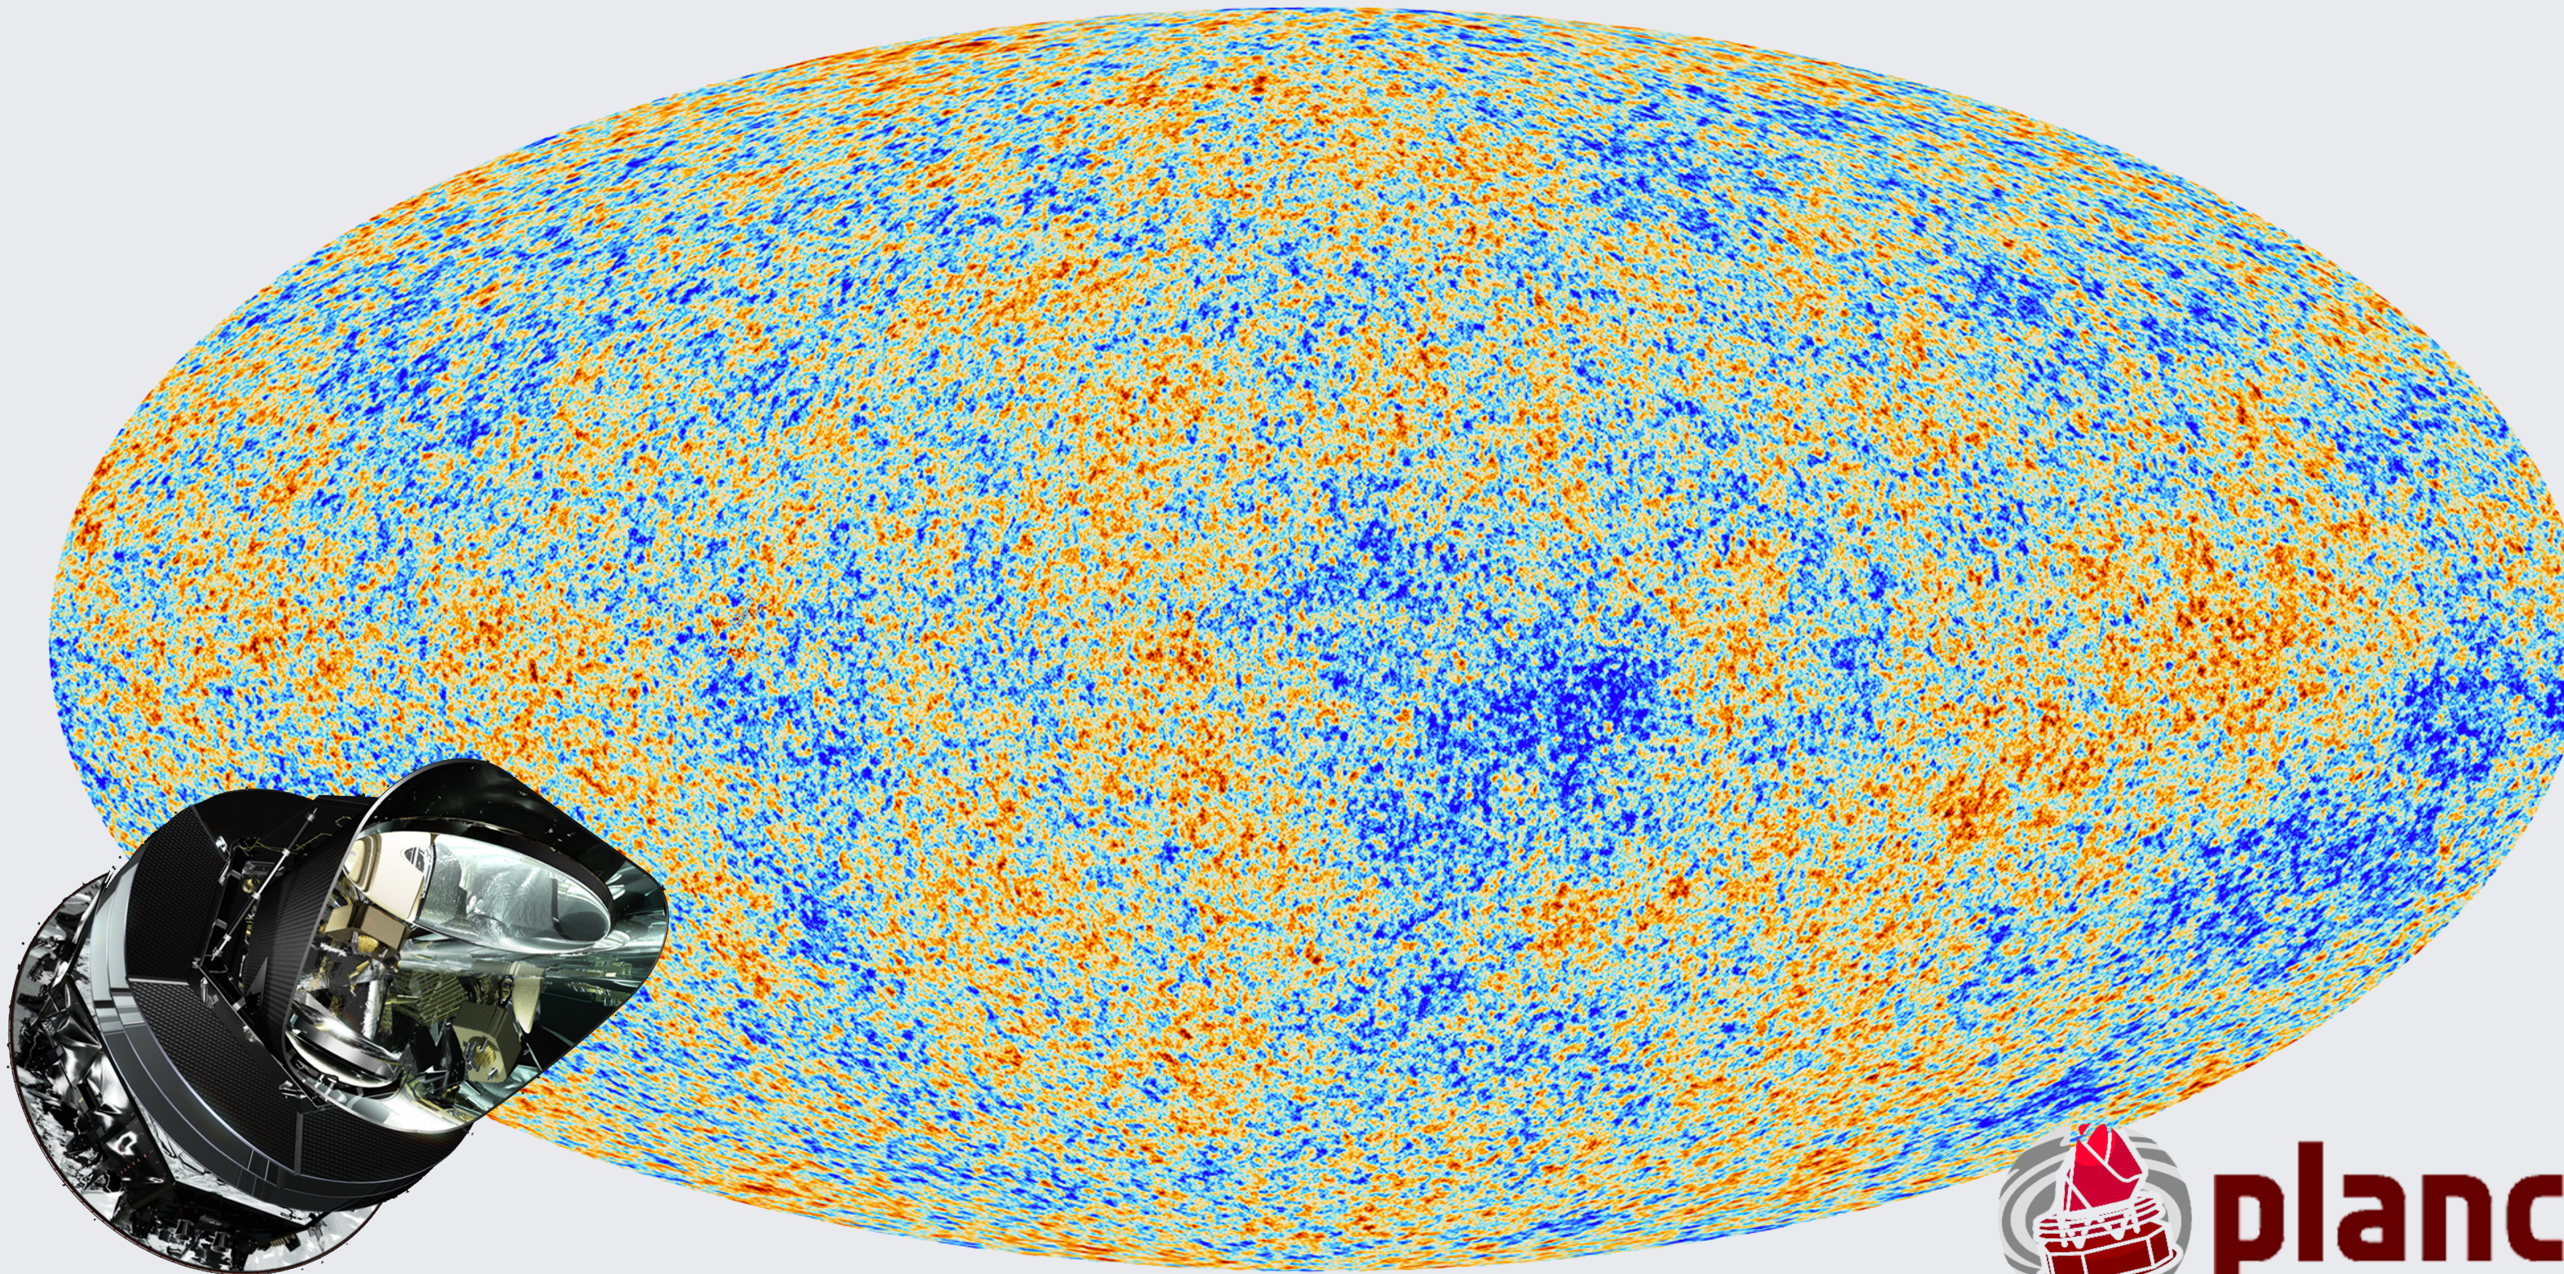

planck

Snapshot of the burning universe 300k years after the big bang

(from the Planck satellite)

High resolution pictures of our Universe

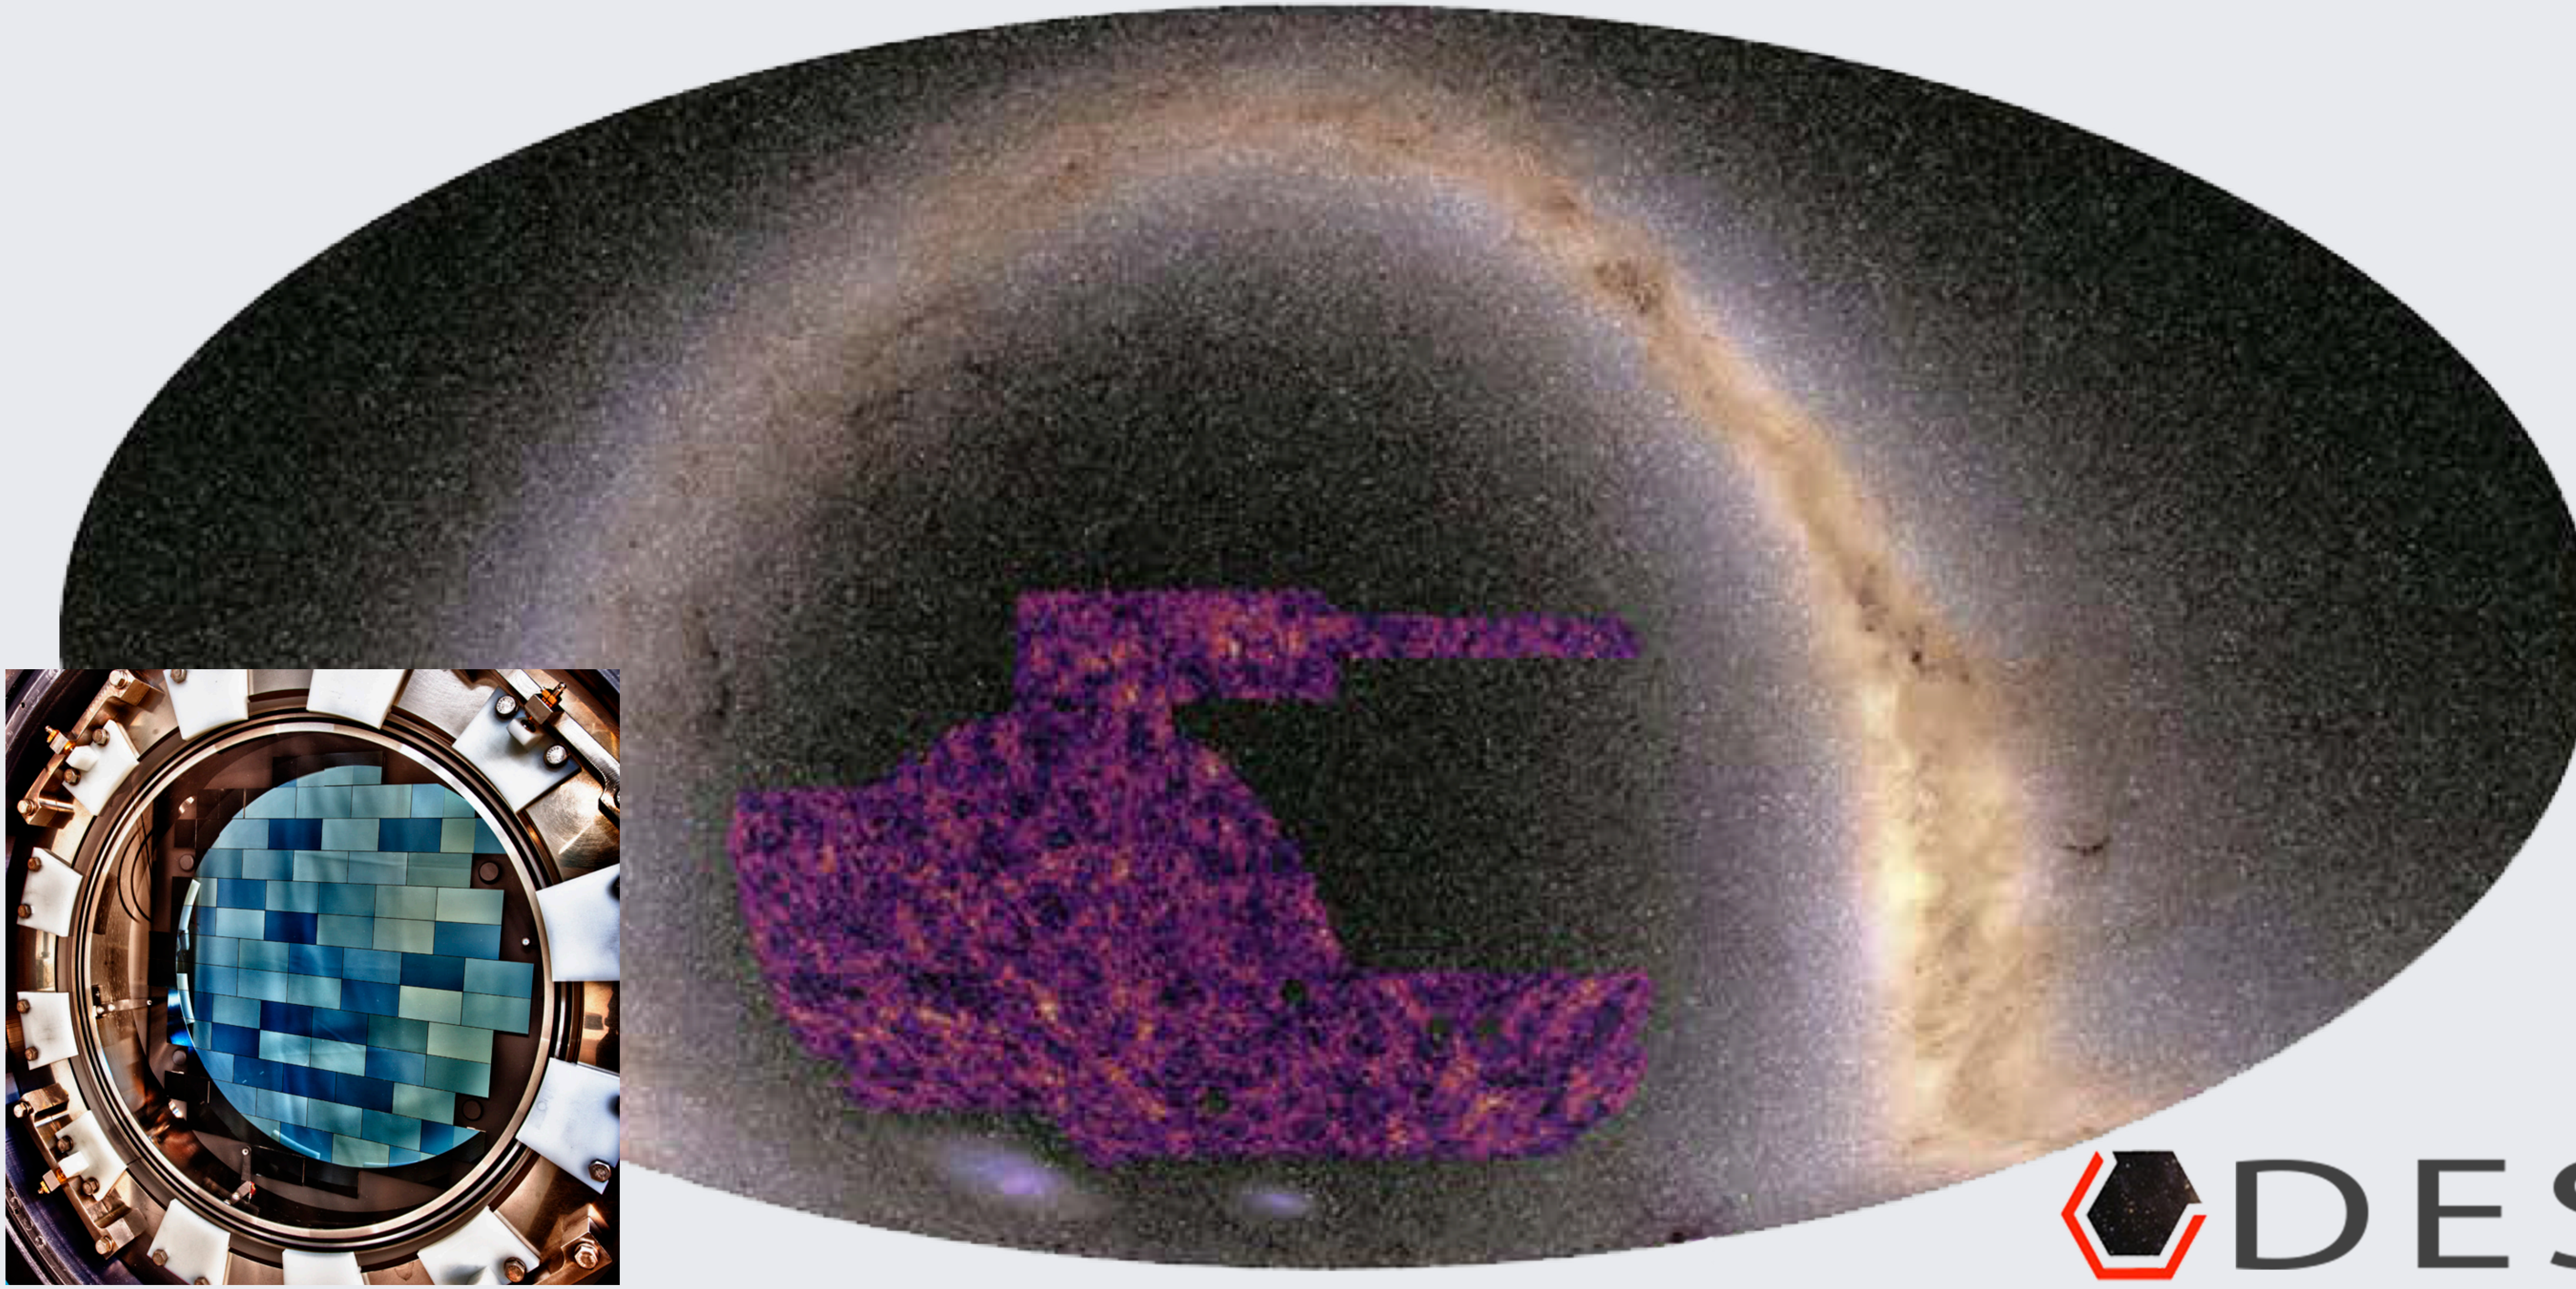

DES

Dark Energy Survey

Snapshot of cosmic structures ~14 billion years later…

(from the Dark Energy Survey)

The modern picture of our Universe

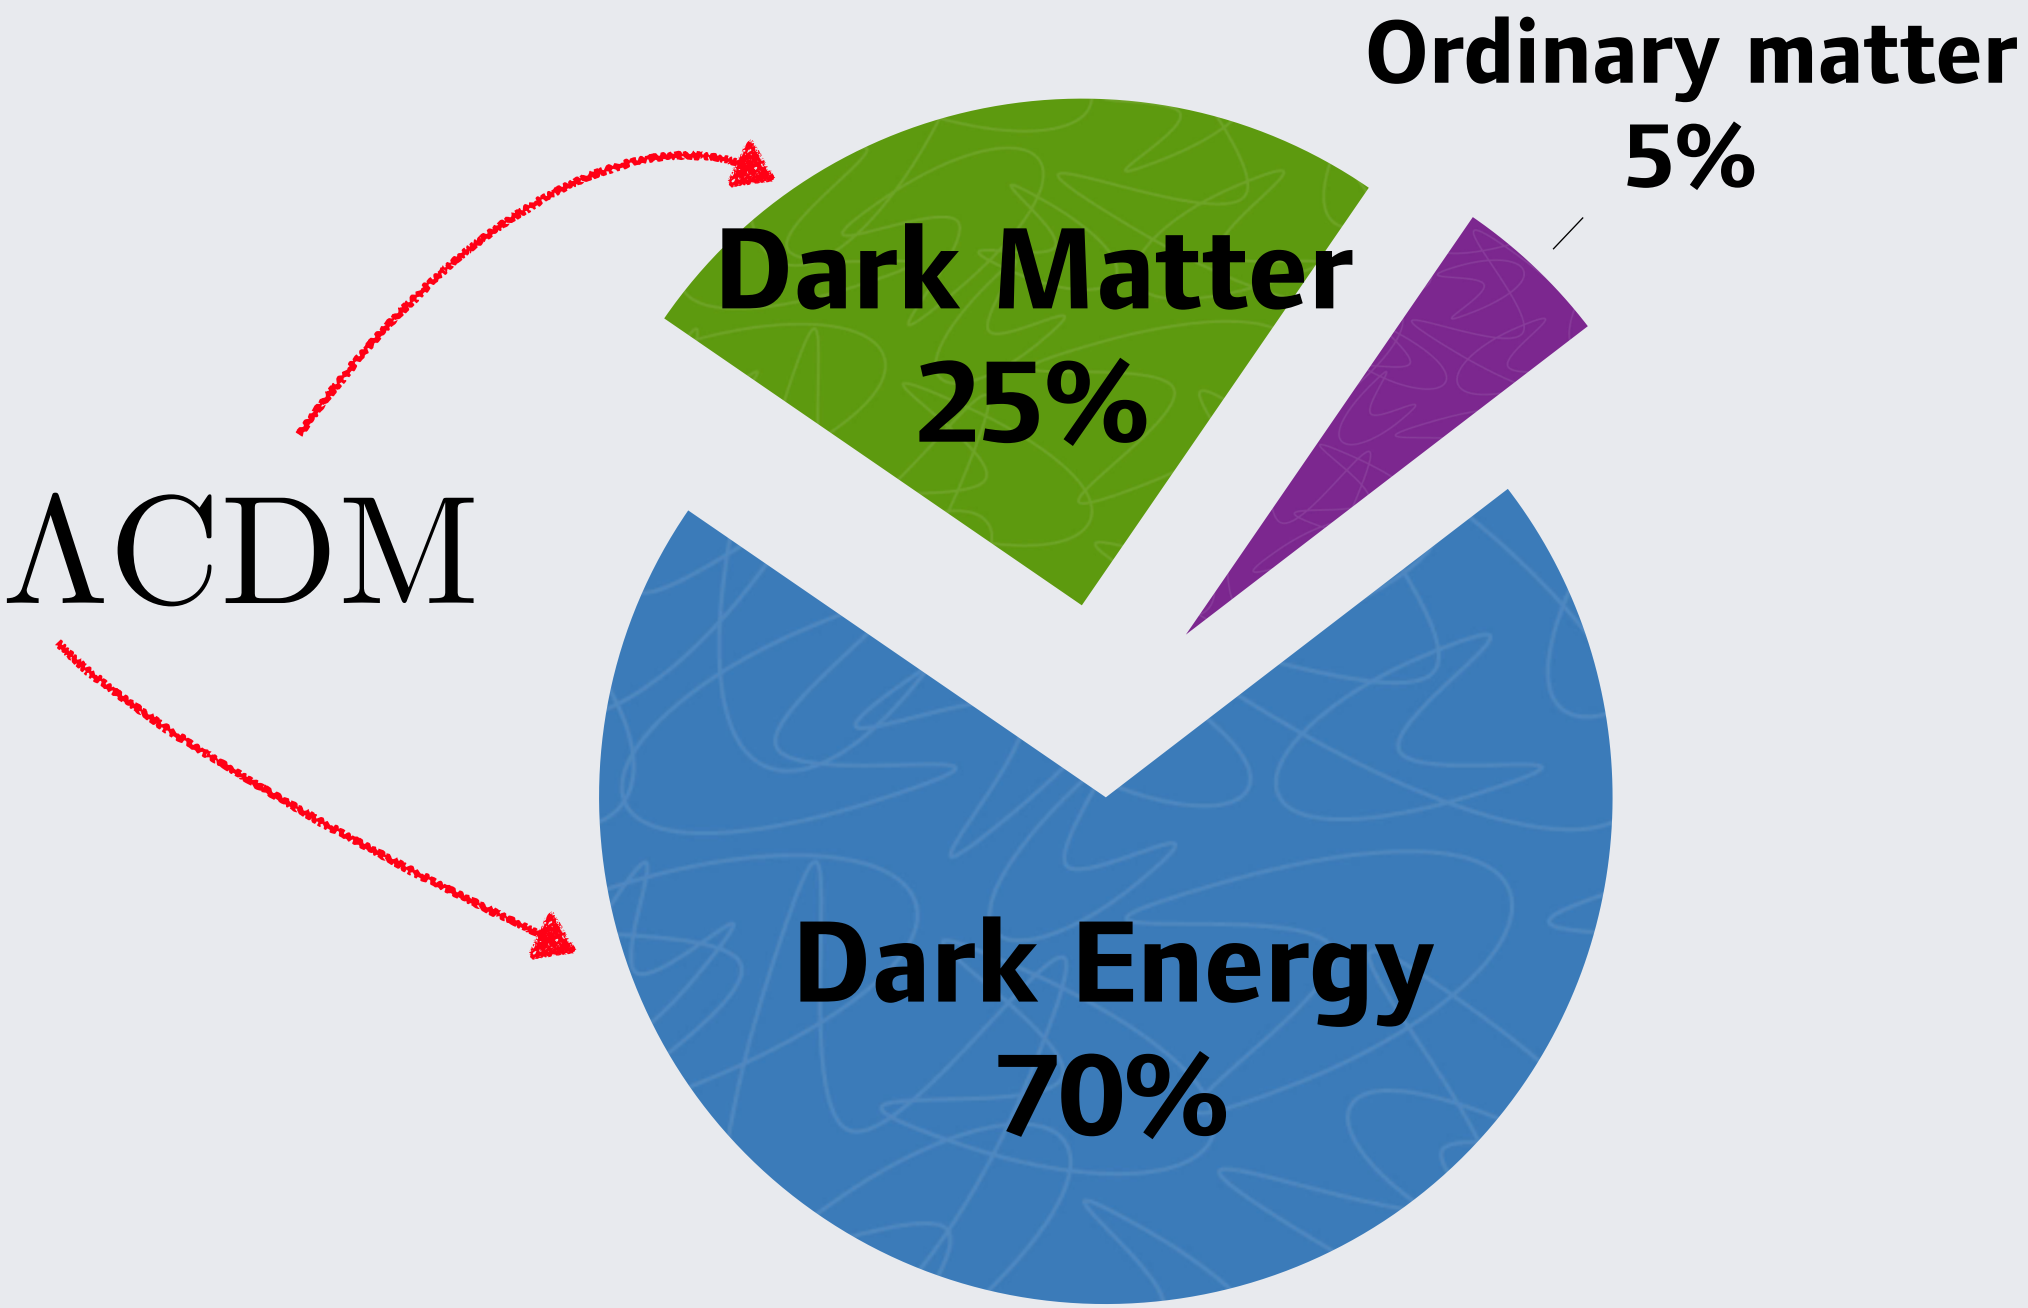

Inferred at about 1% precision

(according to all cosmological data sets we have now…)

The modern picture of our Universe

planck

Dark Energy Survey

DES

euclid

SIMONS OBSERVATORY

Dark Energy Spectroscopic Instrument

CMB-S4

WMAP

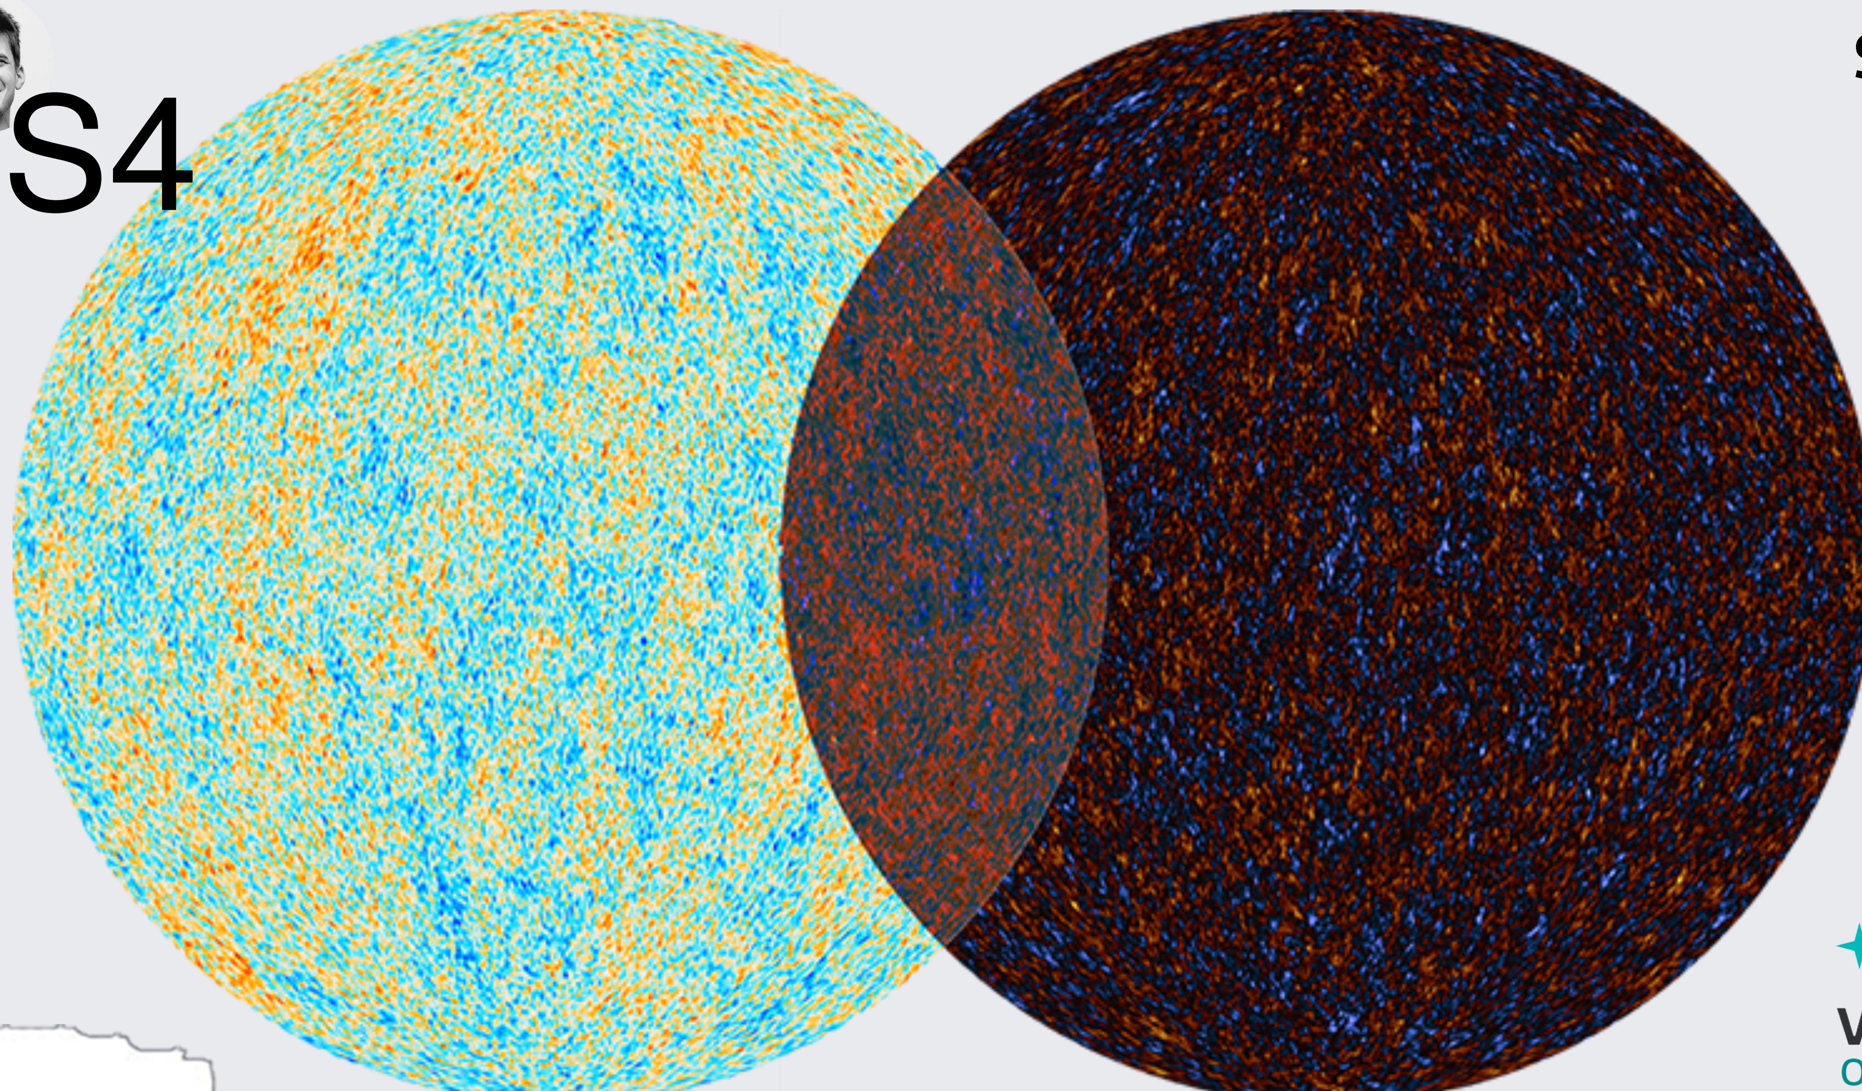

LiteBIRD

JWST

ACT

VERA C. RUBIN OBSERVATORY

SPT

KiDS

HSC

SDSS

(a small Marco close to the surveys I am involved in…)

The Hubble tension

The Hubble constant tension:

H_0 [km/s/Mpc]

The Hubble constant has physical dimensions $[T^{-1}]$

Absolute (not relative) calibration of the expansion rate

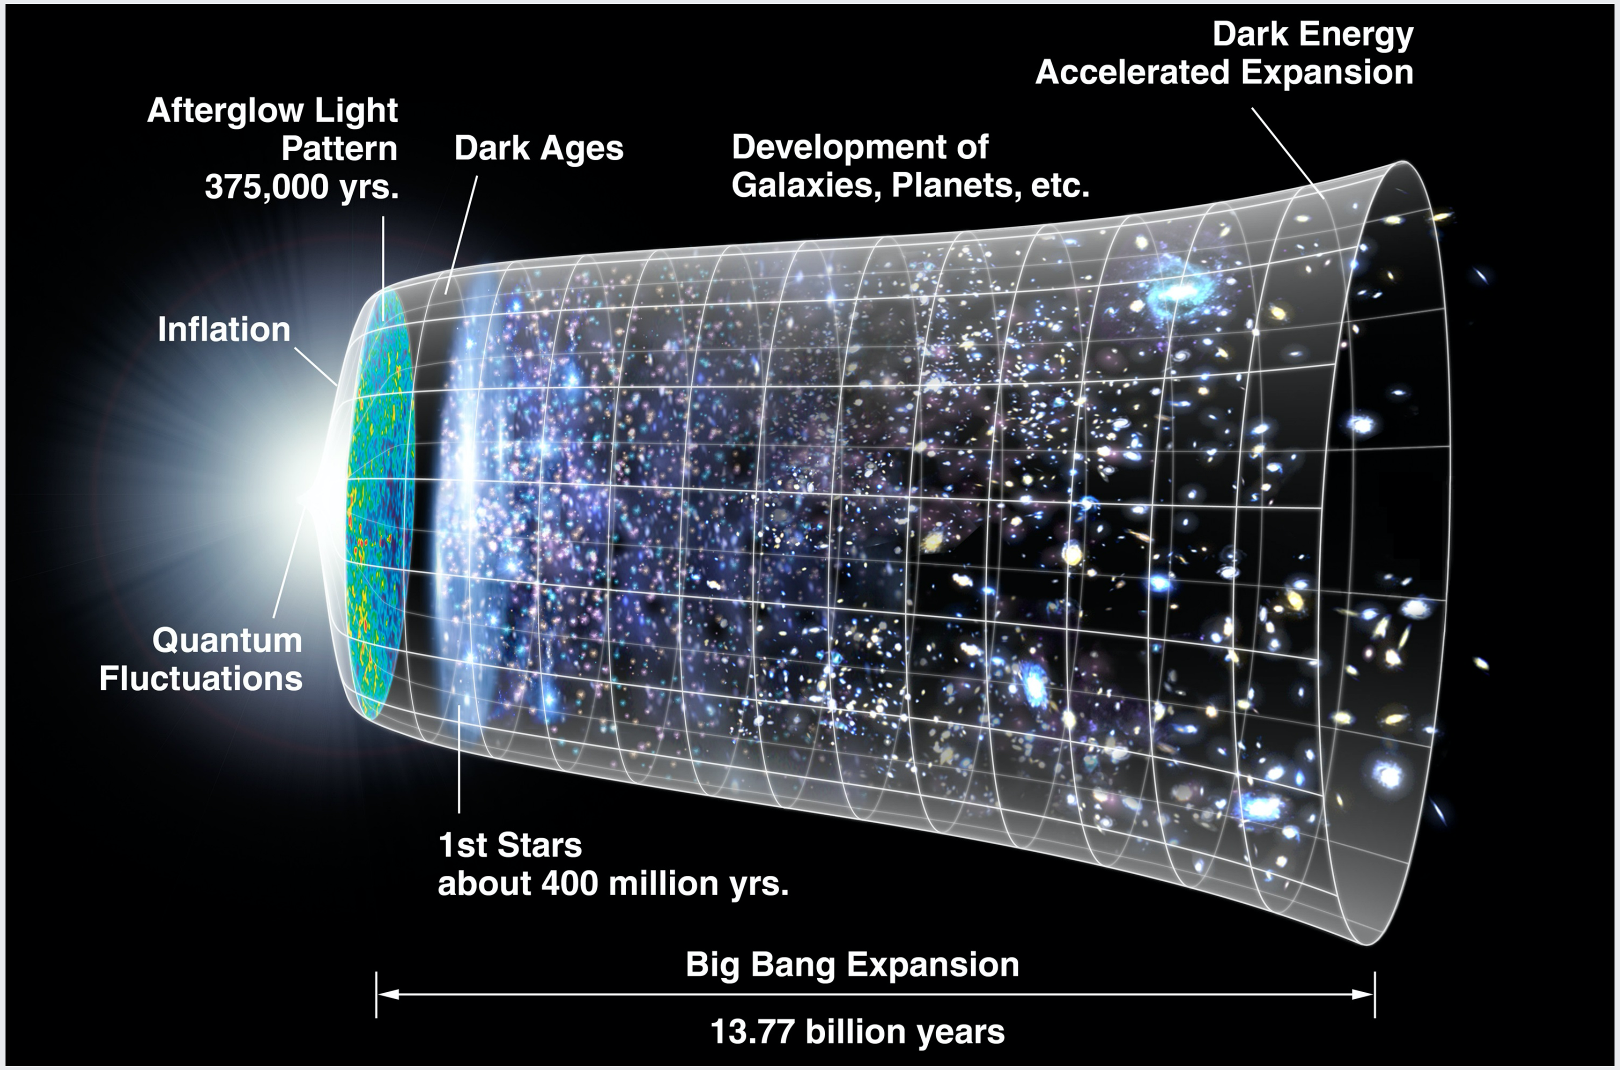

(image credit NASA/WMAP Science Team)

The Hubble constant tension: direct measurement

Measure distance of ~ nearby objects, measure redshifts
$\rightarrow H_0 = 73.15 \pm 0.97$ km/s/Mpc

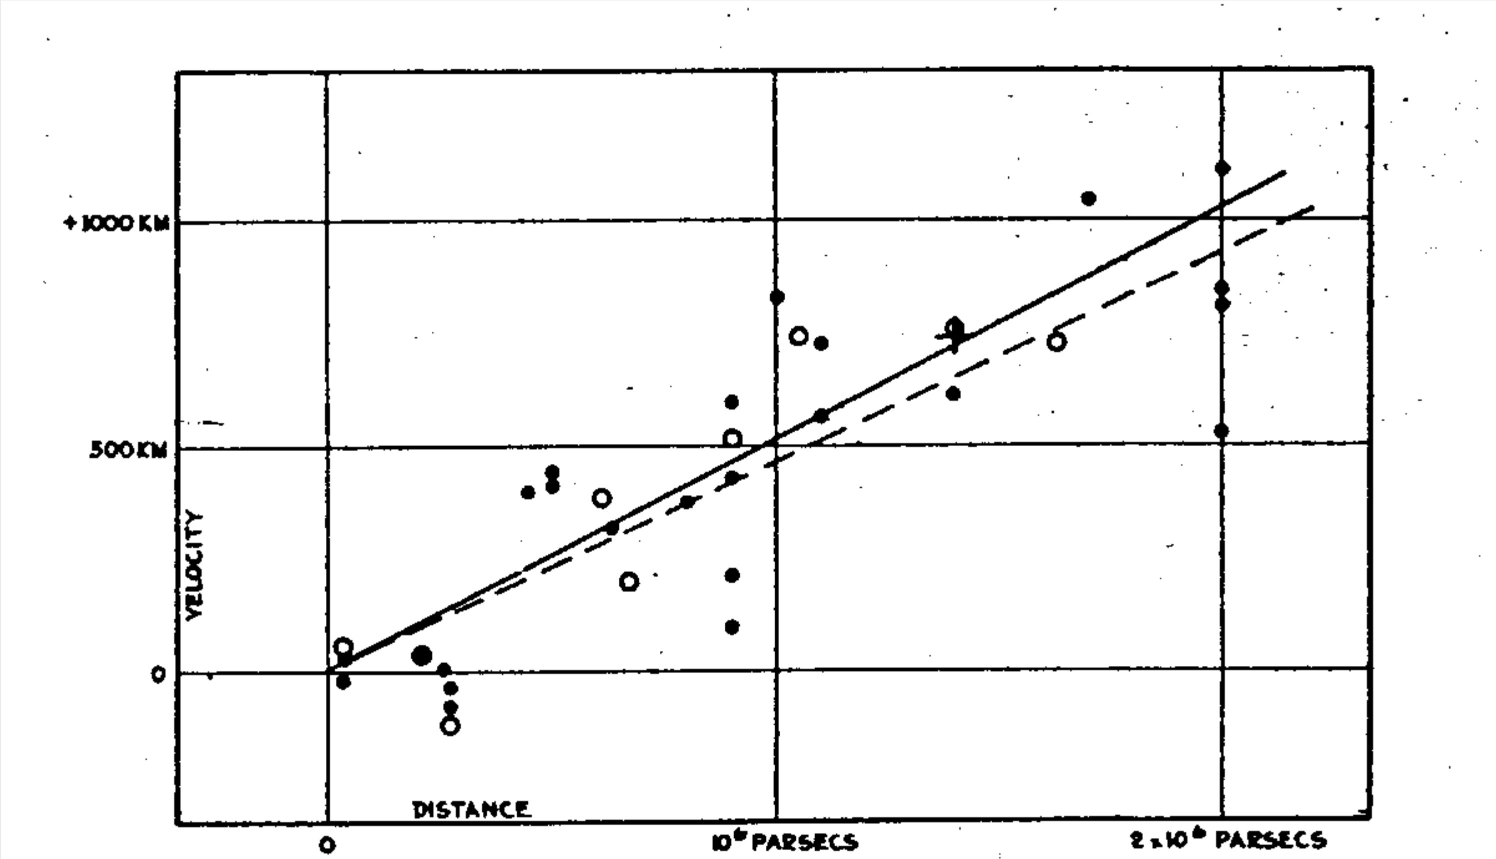

FIGURE 1

Conceptually not (too) different from original method

(Hubble's 1929 paper and arXiv:2208.01045)

The Hubble constant tension: direct measurement

Today's more complicated distance ladder to mitigate systematics: tradeoff between going far and going faint

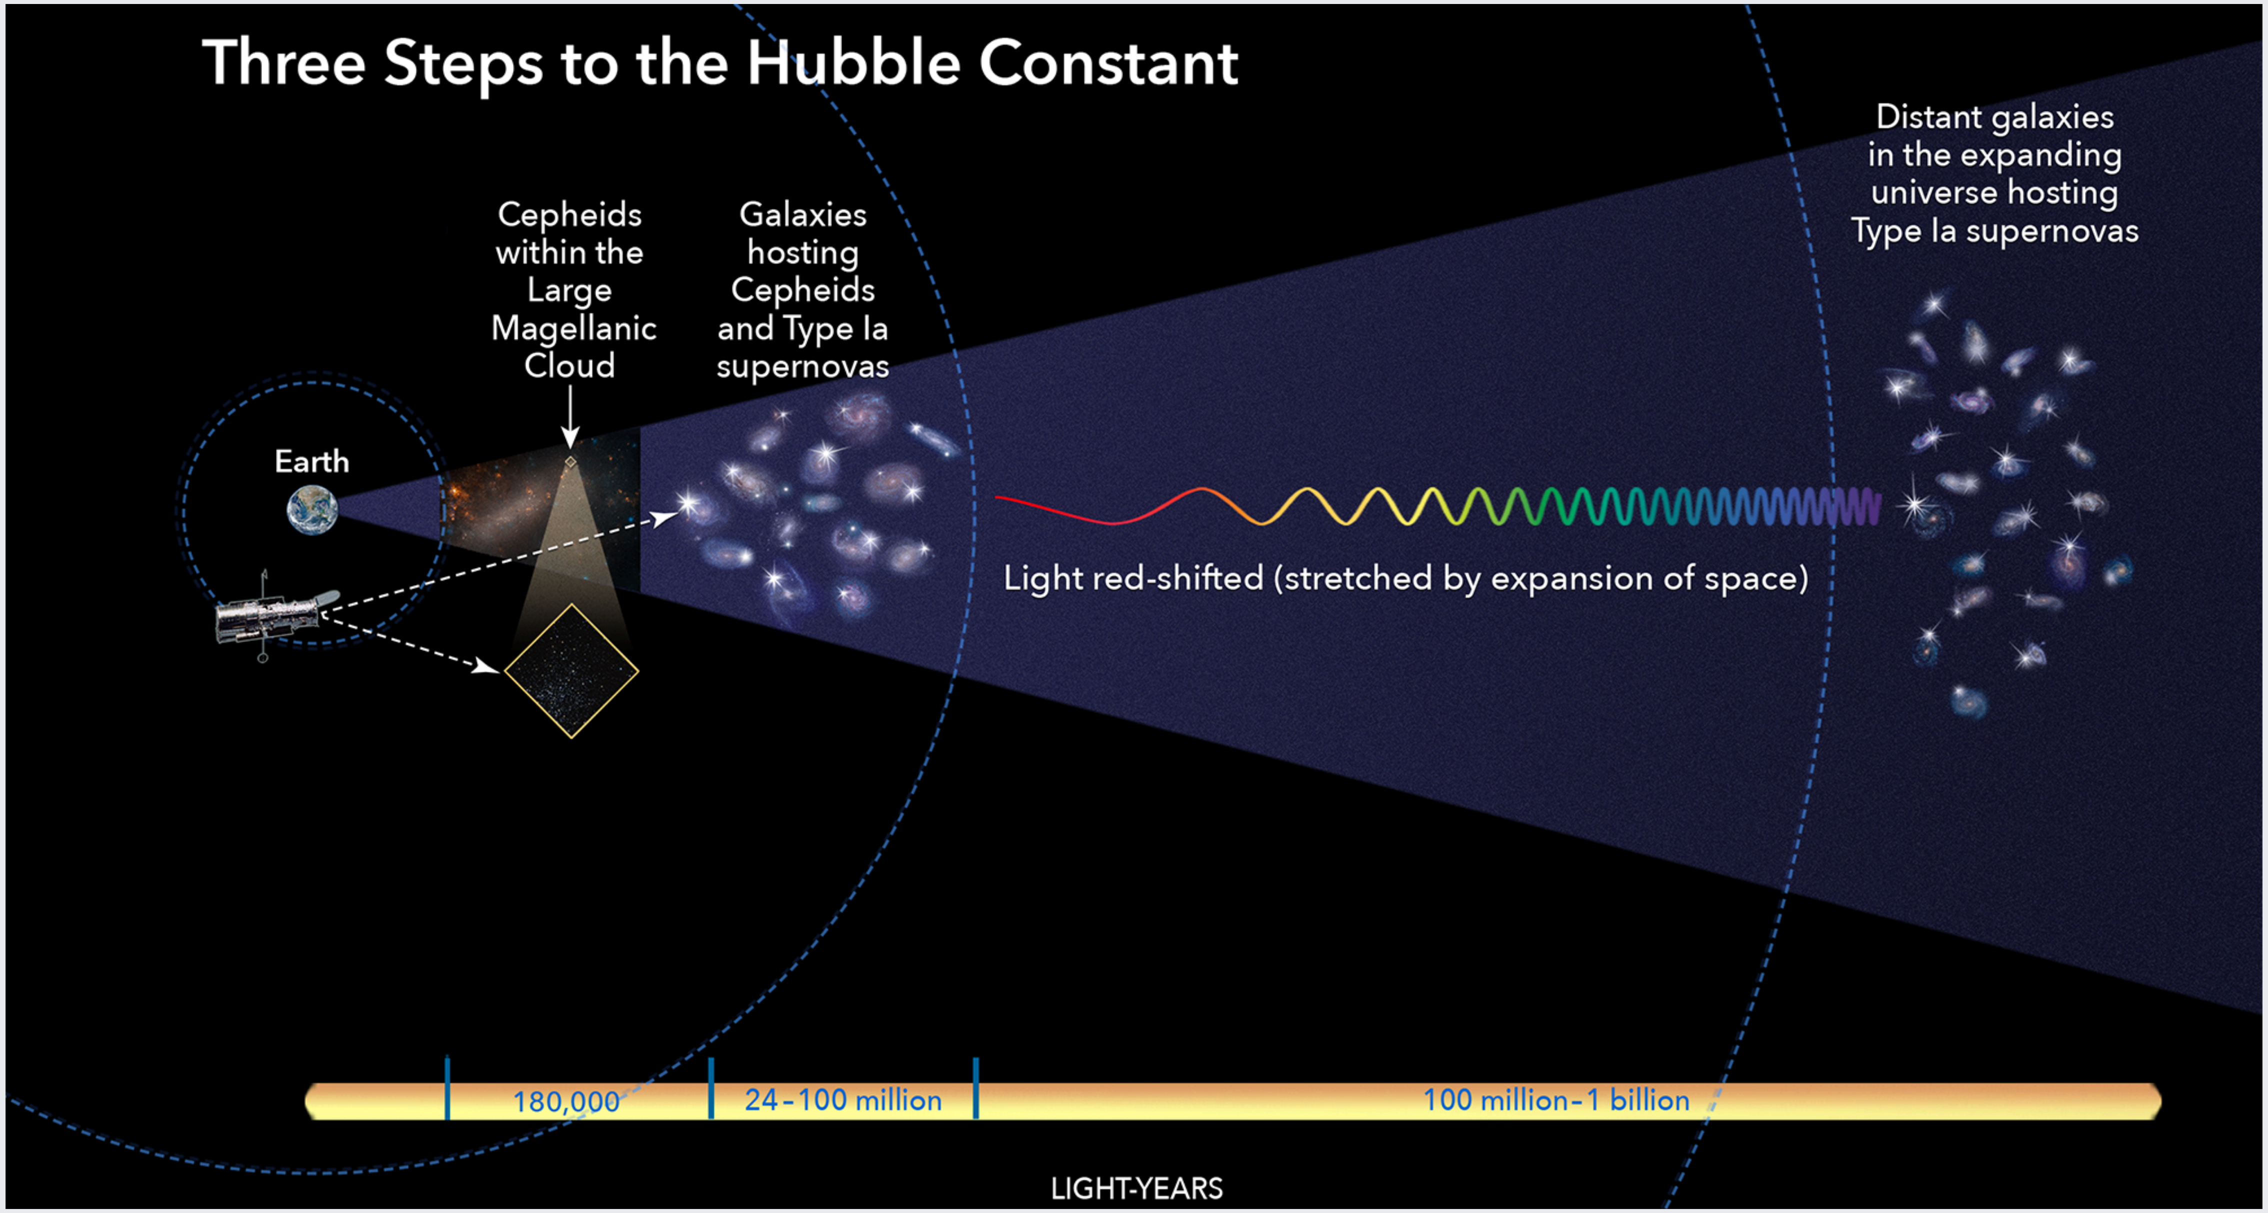

(image credit NASA)

The Hubble constant tension: CMB inference

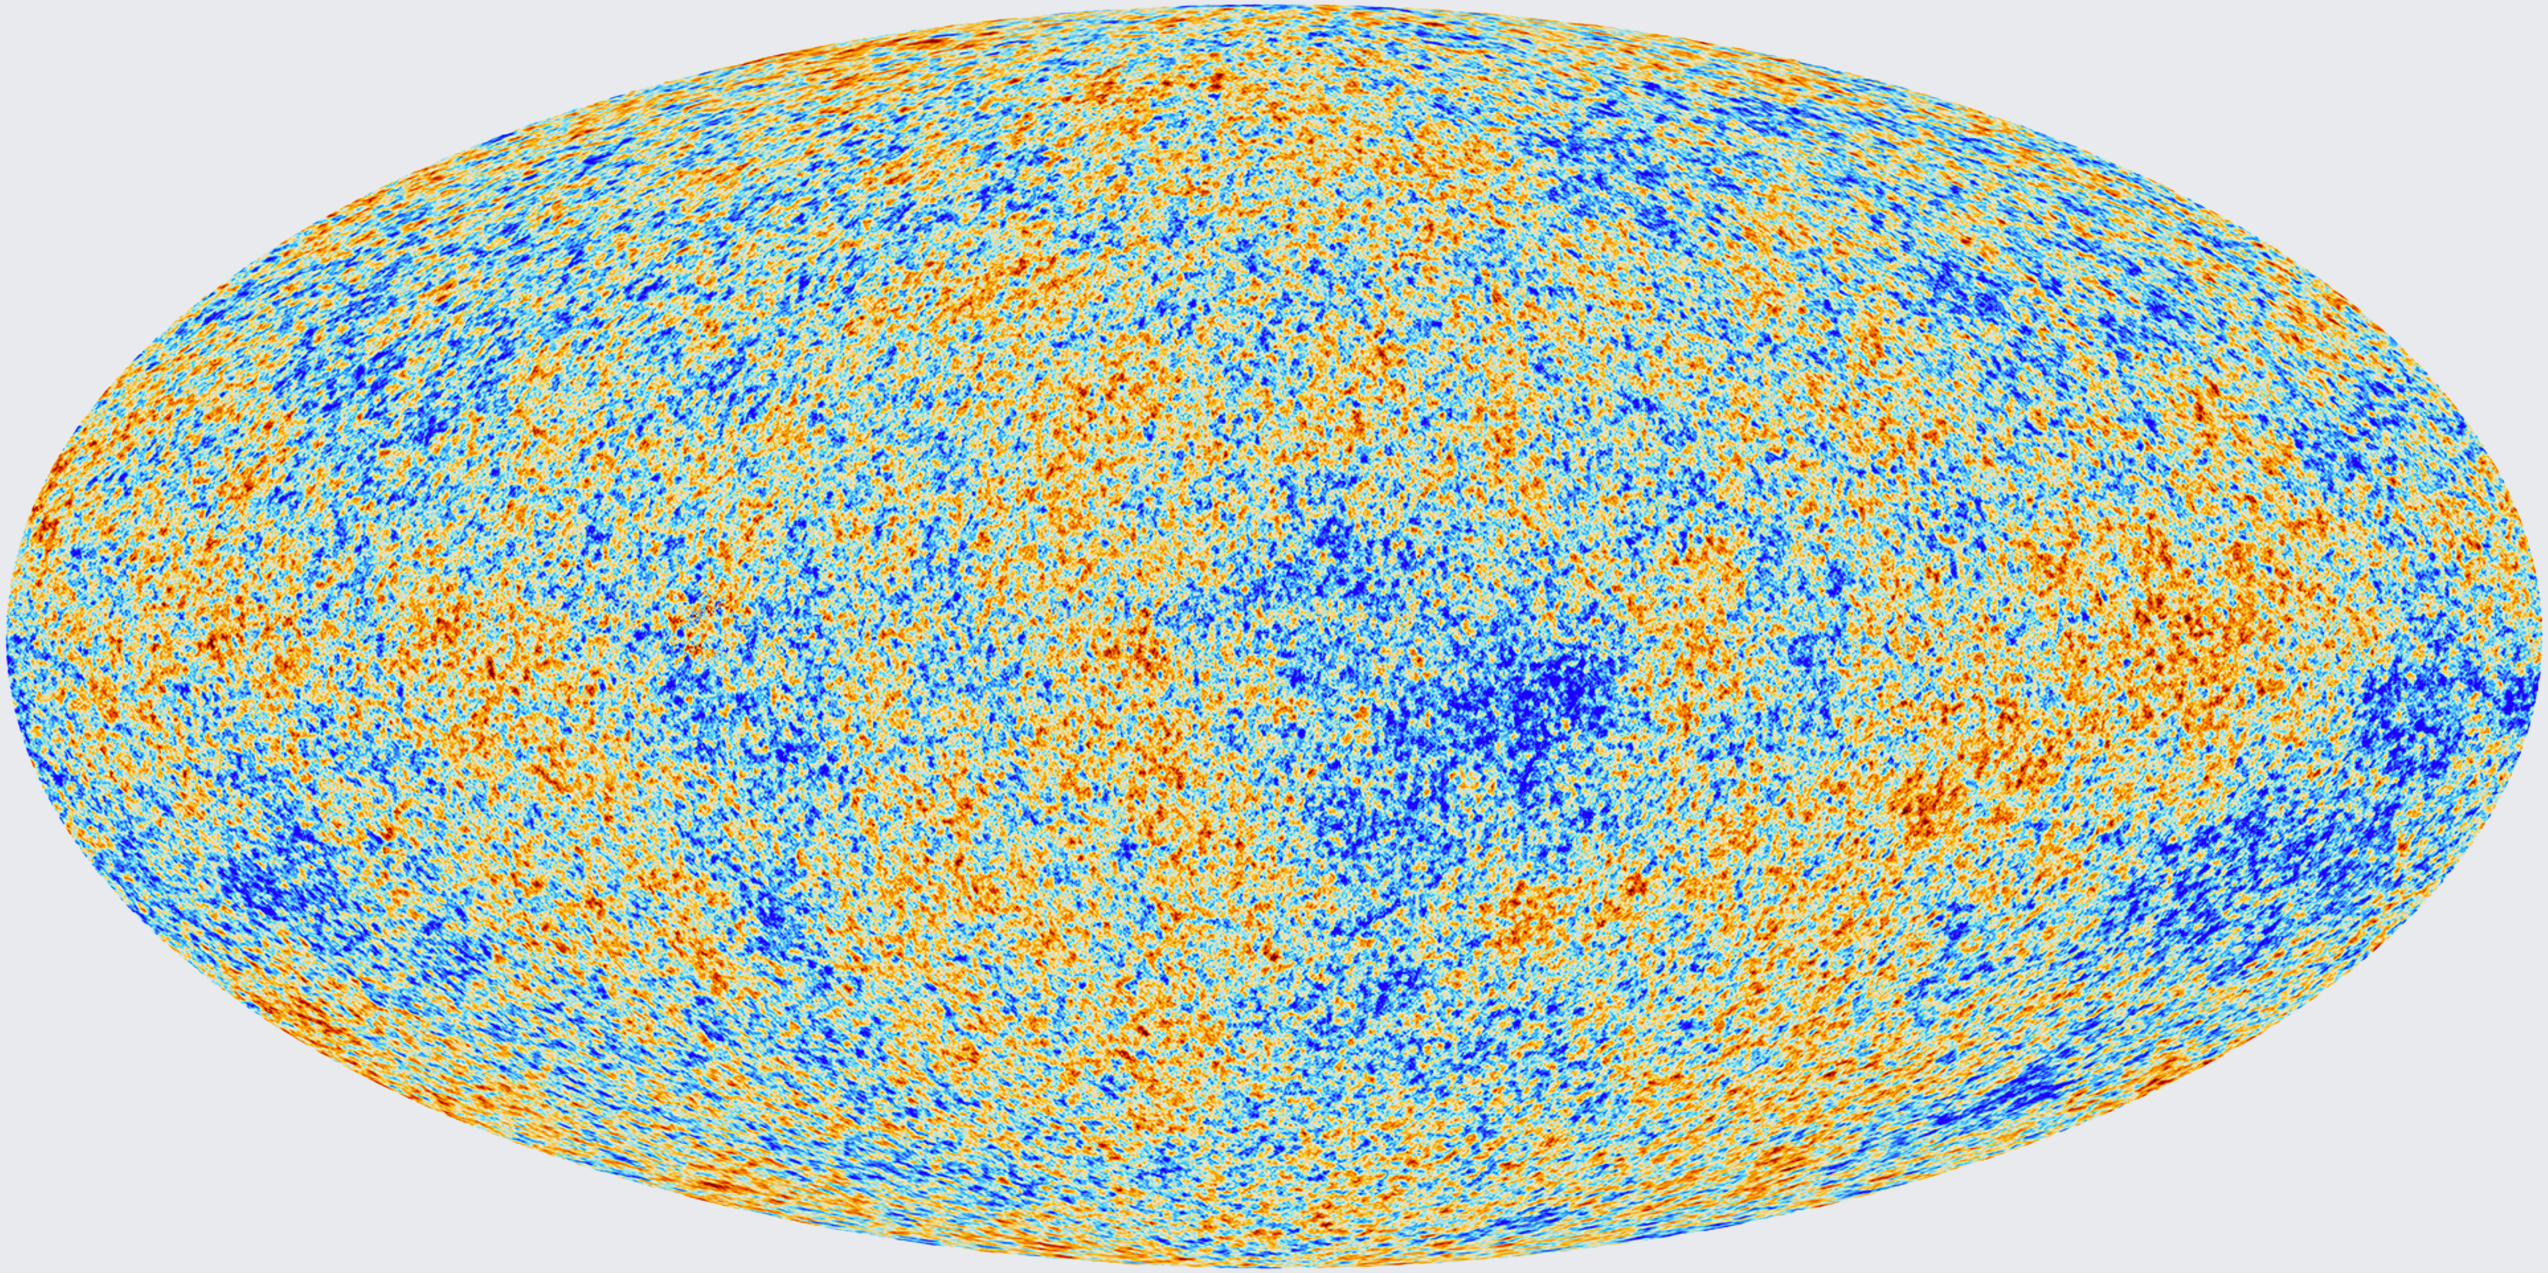

Take a Fourier transform (on the sphere..)

(image credit NASA)

Consonance and dissonance

Has physical dimensions

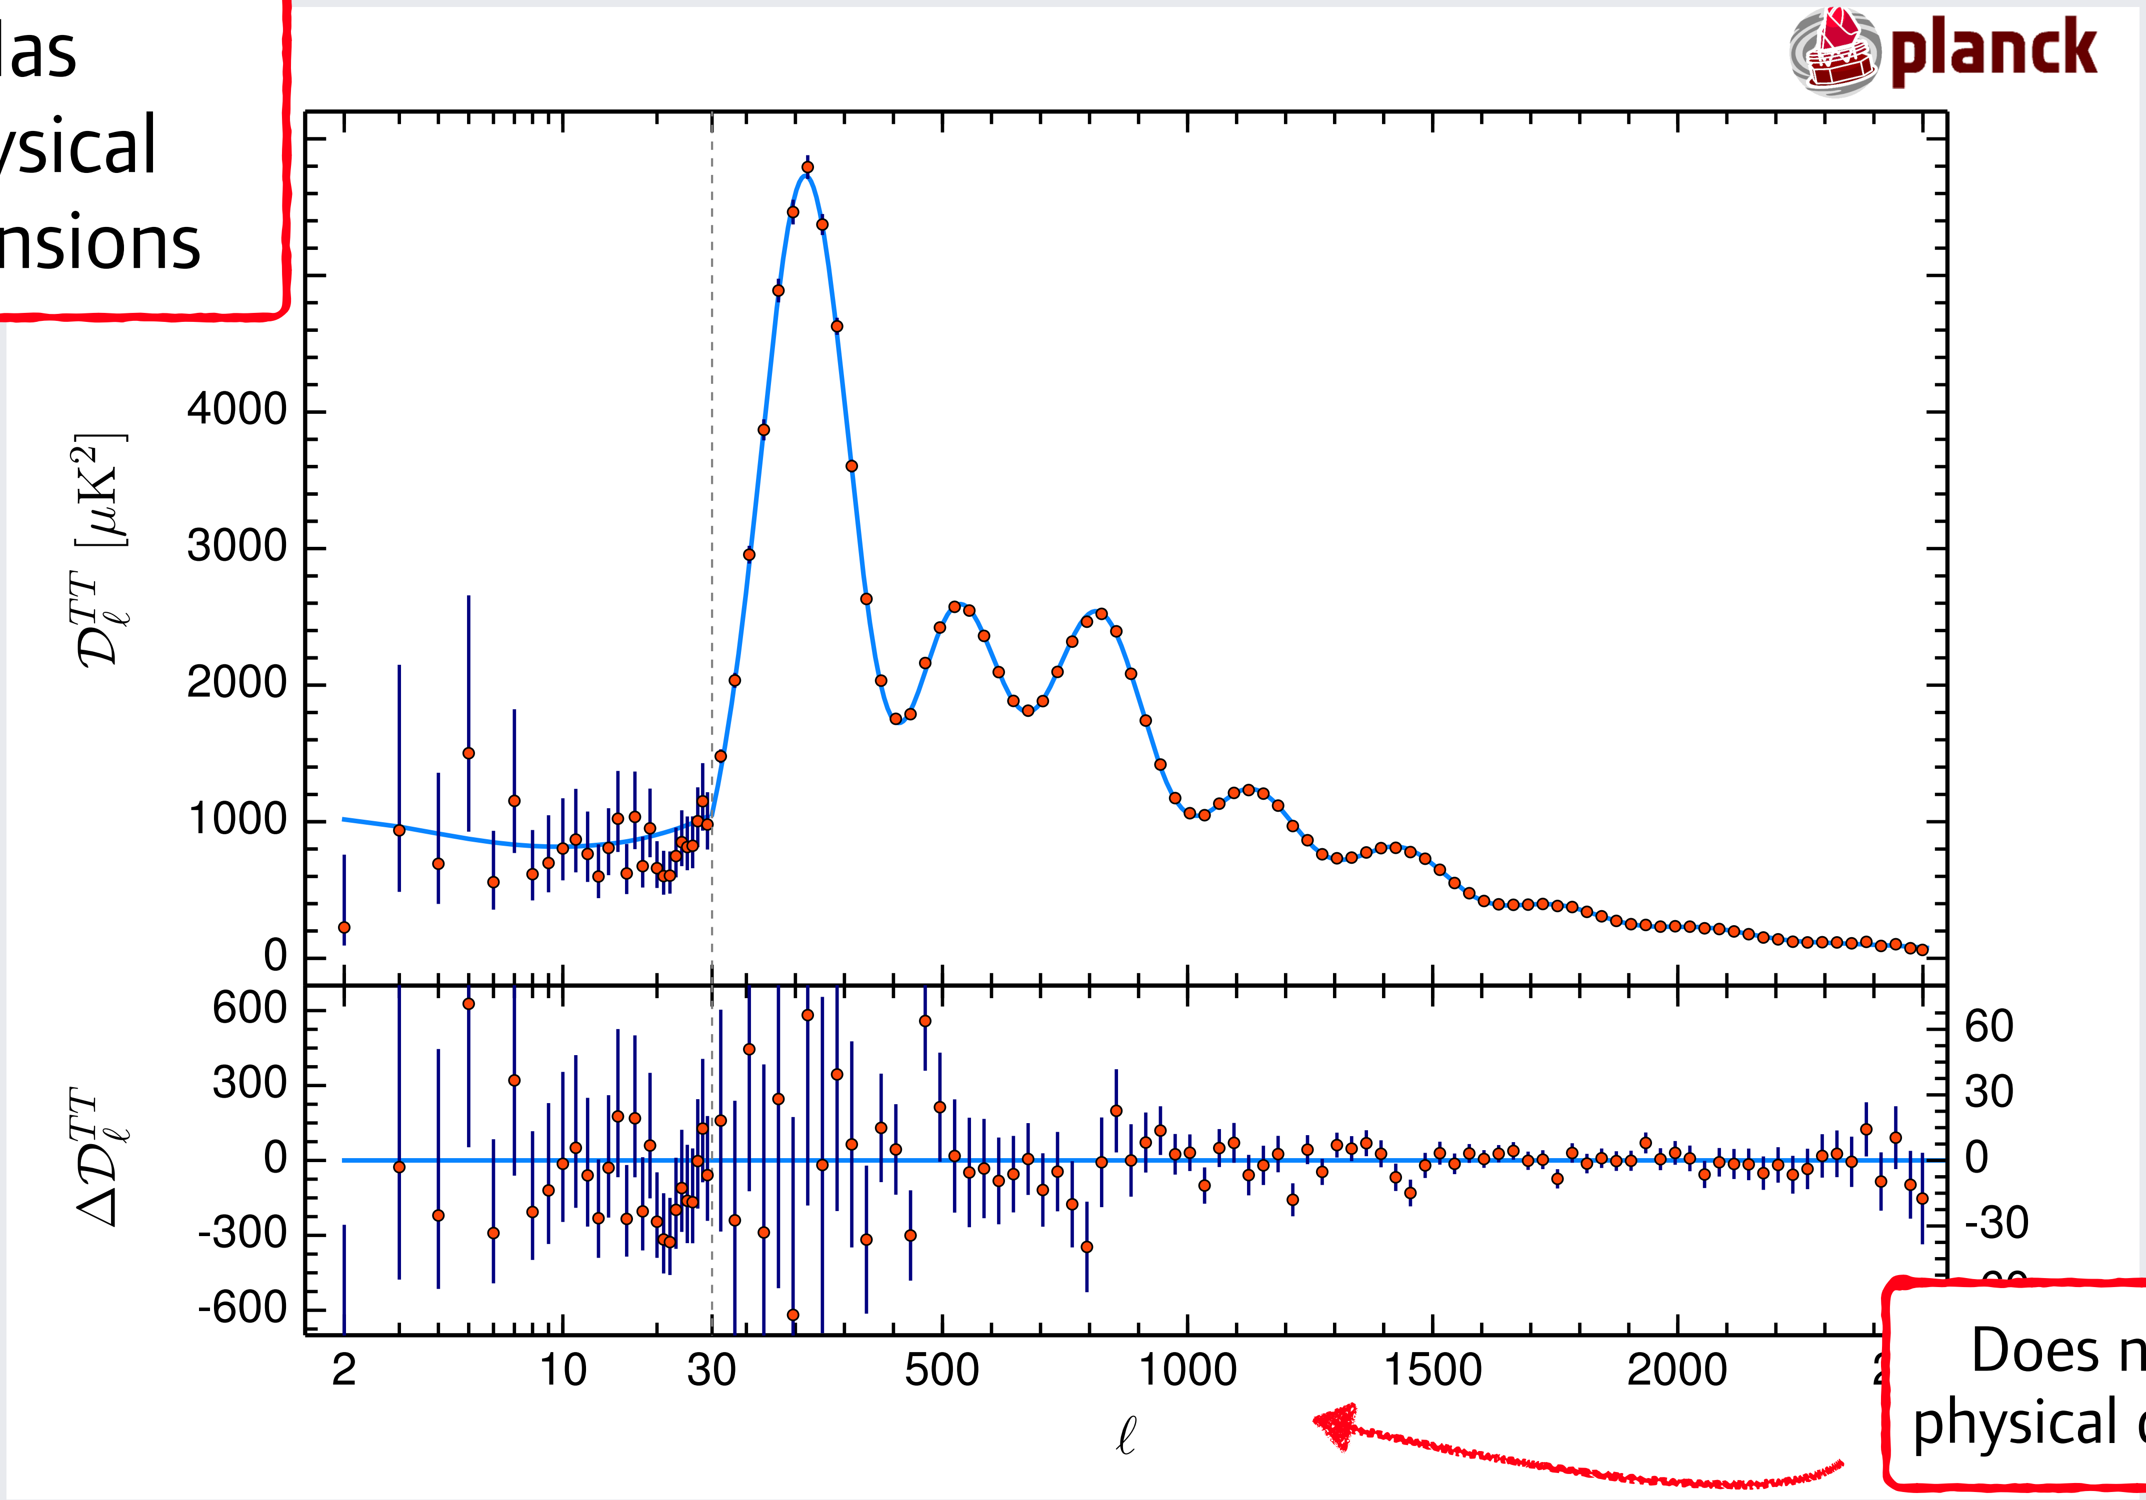

Does not have physical dimension

Second calibrator: power spectrum of CMB temperature fluctuations

(Planck 2018 cosmological parameters, arXiv:1807.06209)

Consonance and dissonance: gravitational ringing

Potential wells = inflationary seeds of structure

Photons and baryons fall into wells
pressure due to Thompson scattering resists:
acoustic oscillations

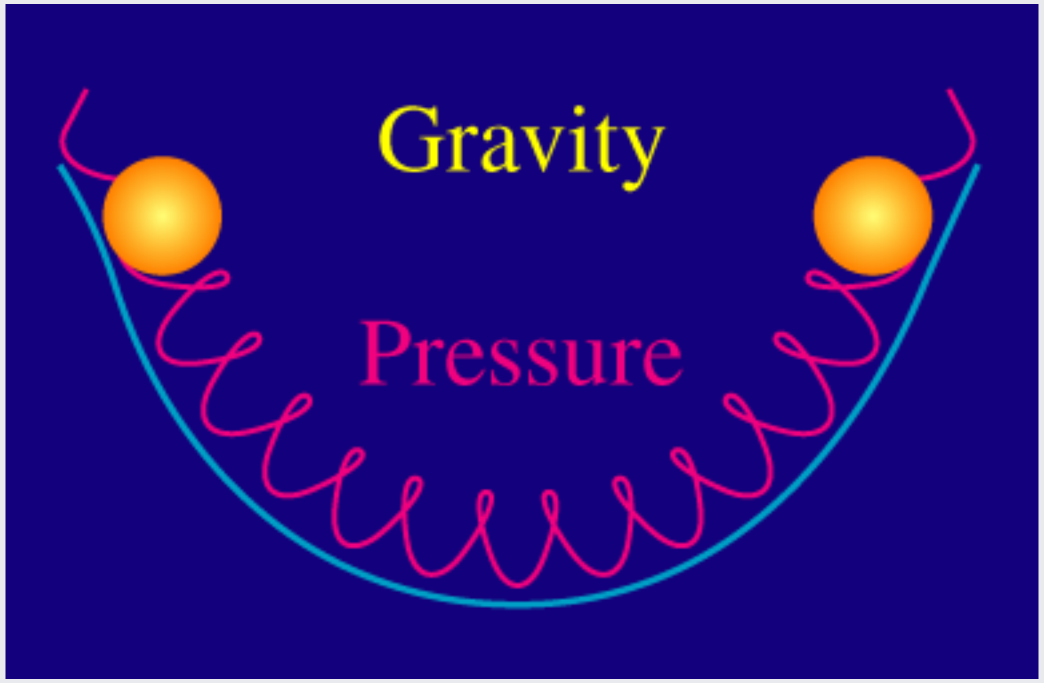

(credit Wayne Hu)

Consonance and dissonance: seeing sound

Has physical dimensions

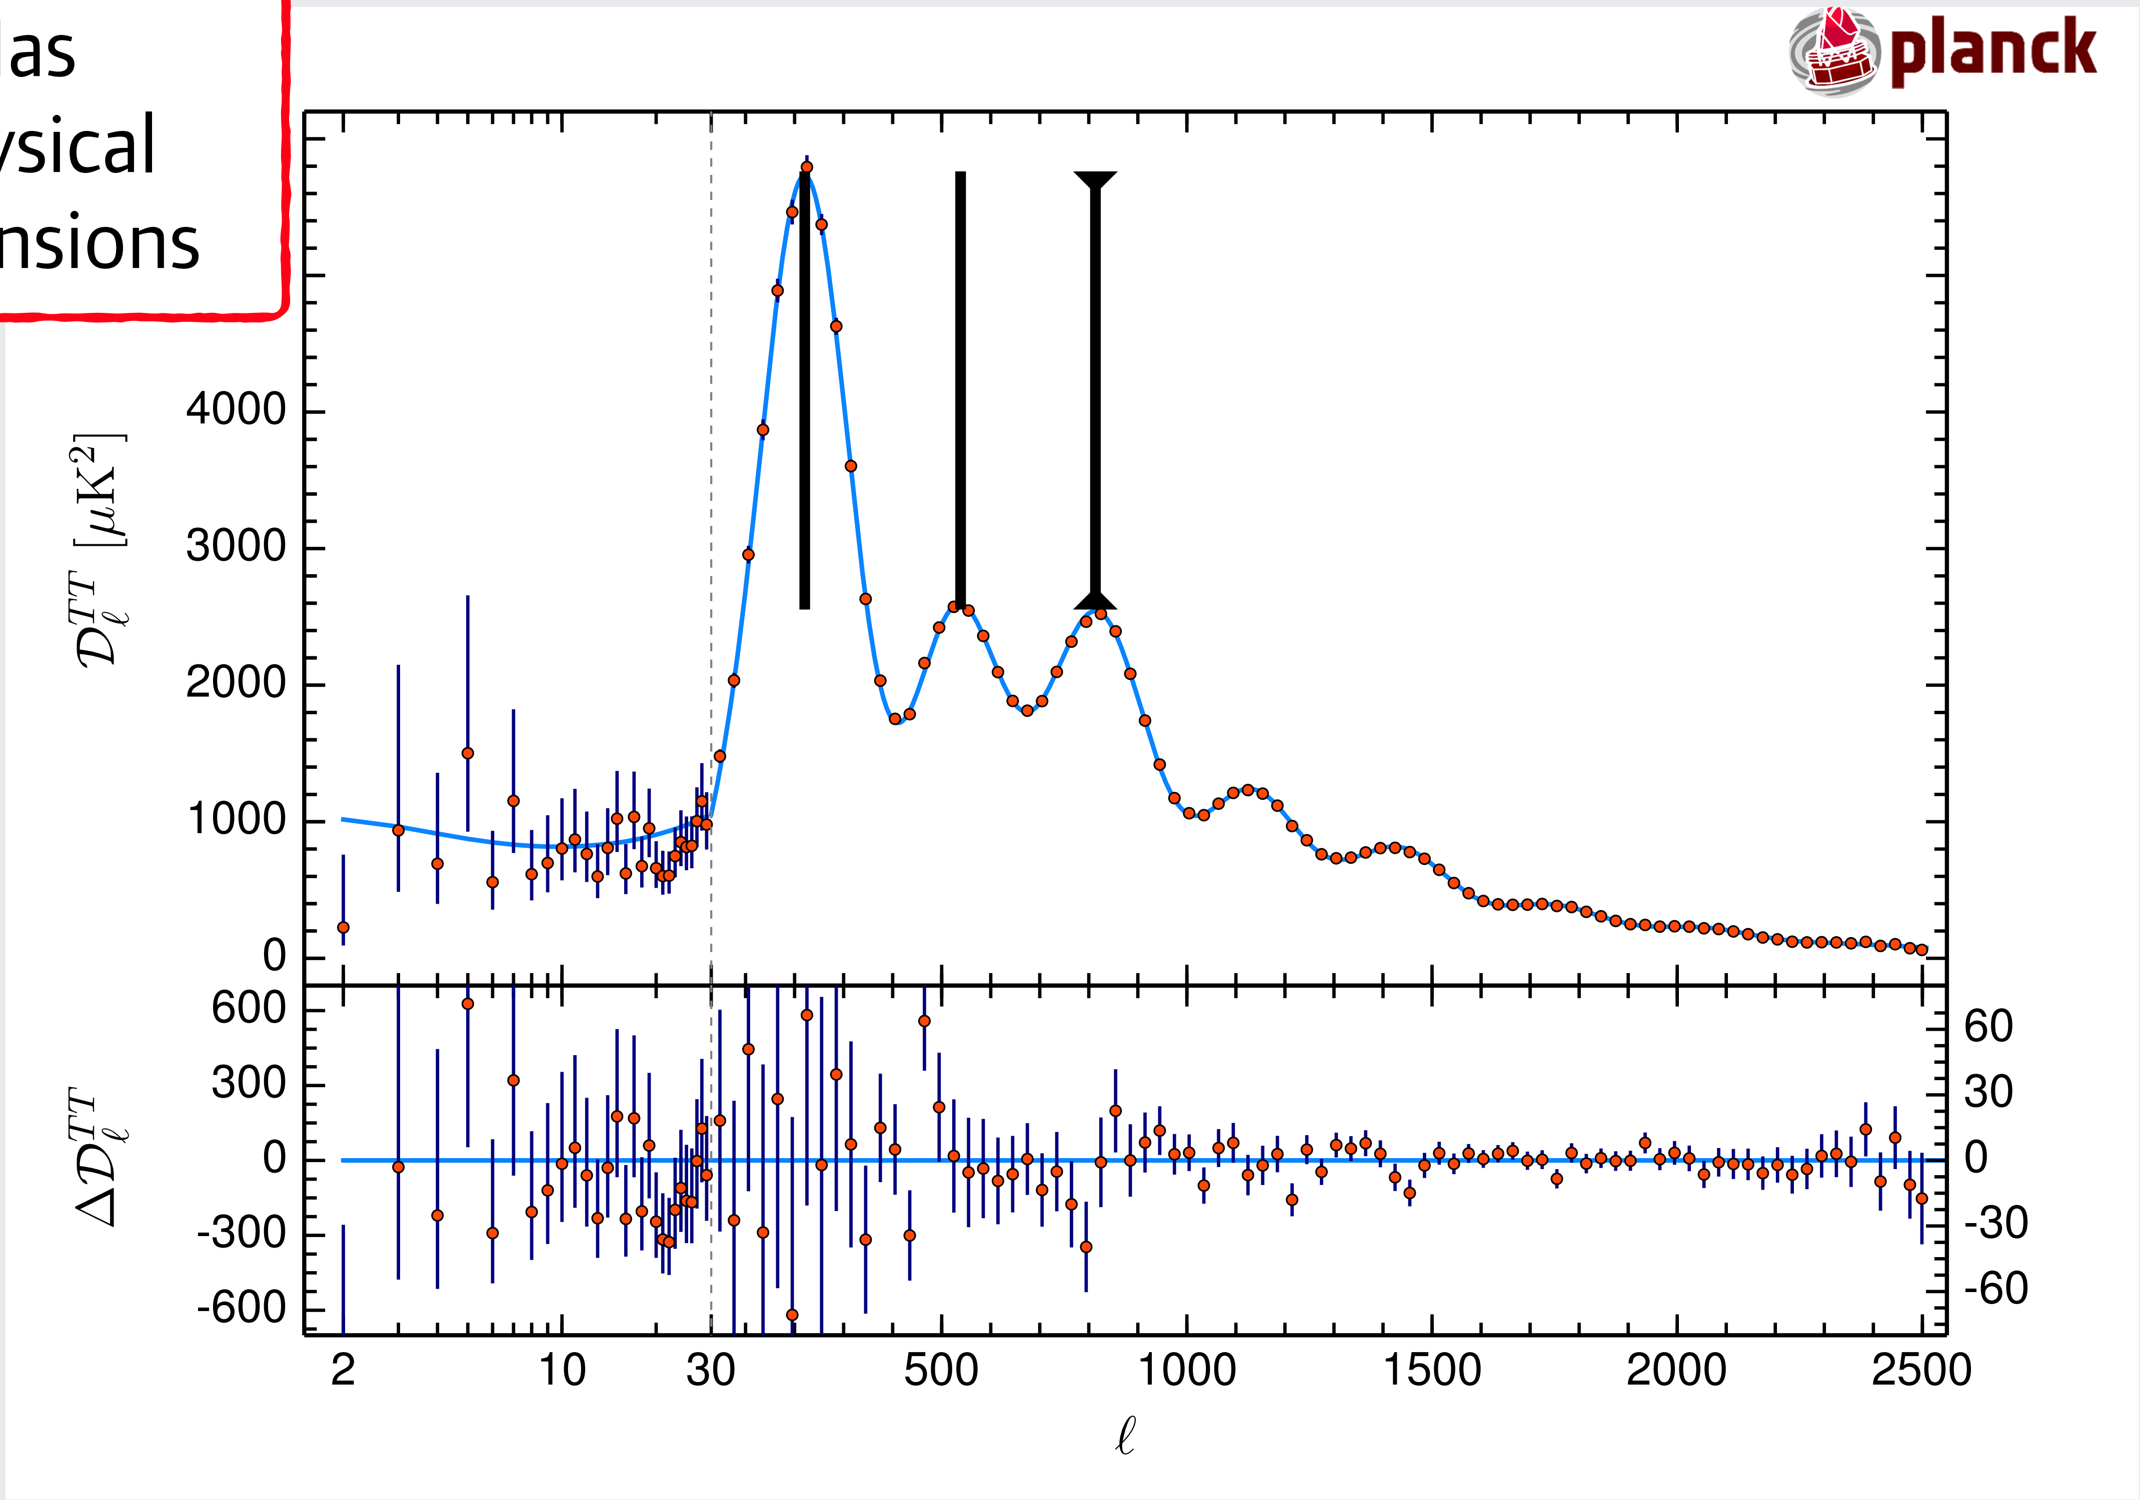

(Planck 2018 cosmological parameters, arXiv:1807.06209)

Consonance and dissonance: timbre of the Universe

$\rho_X \sim \Omega_X h^2$ (physical and relative densities)

Peak heights and pre-recombination physics:
$\Omega_r h^2, \Omega_b h^2$ and $\Omega_{dm} h^2 \;\rightarrow r_s$ sound horizon

Conceptually like a musical instrument:
different instruments (models) playing the same pitch are distinguished by the structure of the spectrum

Consonance and dissonance: timbre of the Universe

Angular scale and distance to recombination: $\Omega_{dm} + \Omega_b$

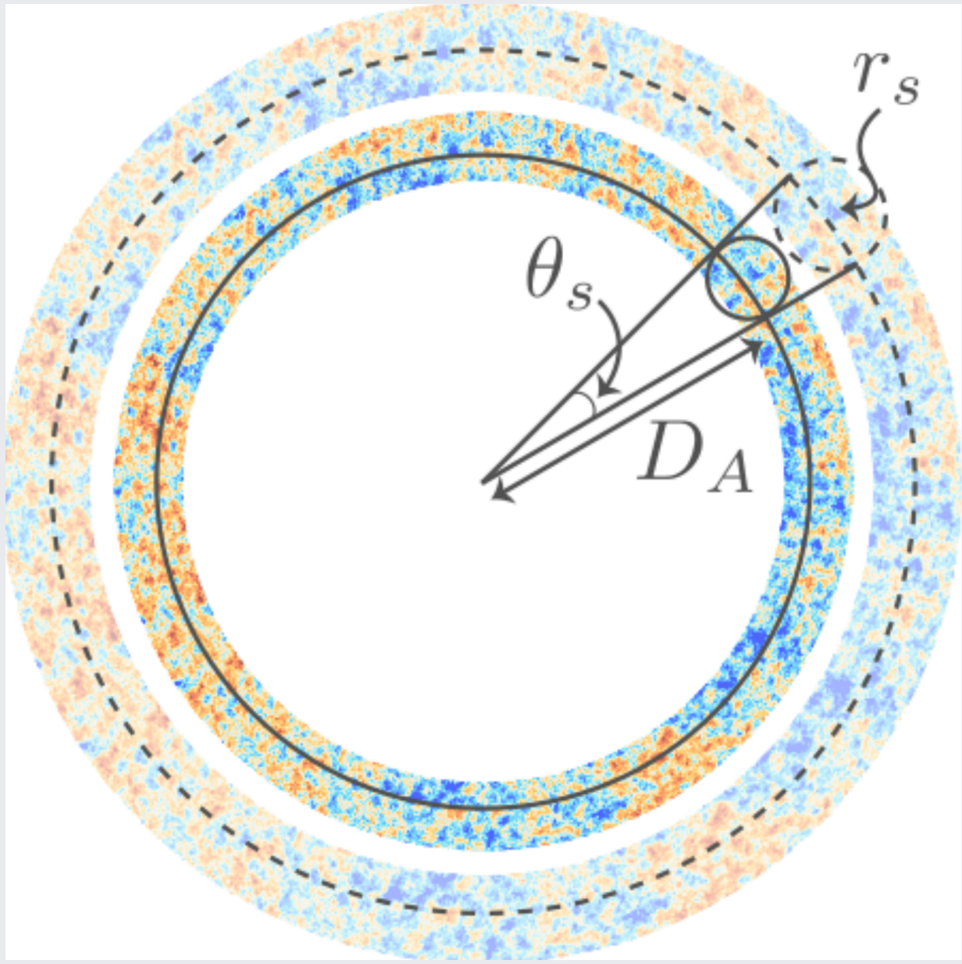

$$\theta_s = r_s / D \simeq r_s H_0$$

$$H_0 = 67.4 \pm 0.5 \text{ km/s/Mpc}$$

(image credit Tristan Smith)

Gaussian estimators of 1D tensions

The usual recipe

(exact for Gaussian models with only 1 parameter)

1D parameter shifts

$$T_1 \equiv \frac{|\theta_1 - \theta_2|}{\sqrt{\sigma_1^2 + \sigma_2^2}}$$

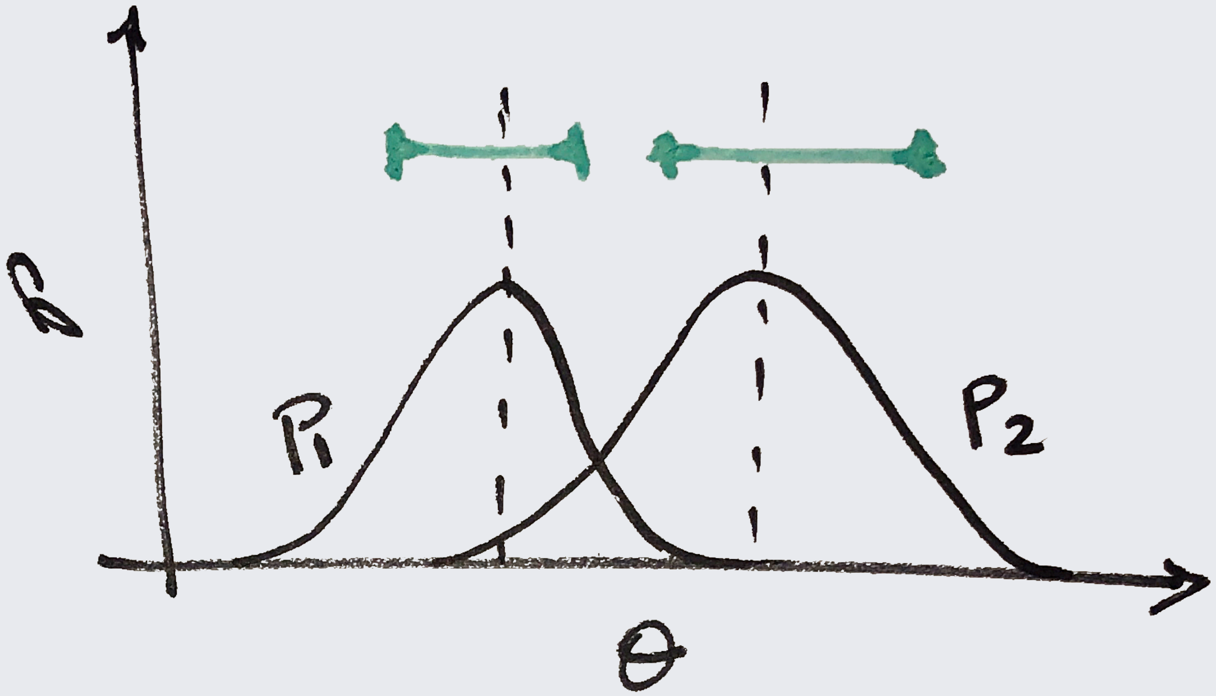

(can be shown to be unique in 1D)

The Hubble constant tension:

From CMB inference $H_0 = 67.4 \pm 0.5$ km/s/Mpc

From direct measurement $H_0 = 73.15 \pm 0.97$ km/s/Mpc

$\sim 8\,\%$ discrepancy between early and late time calibration of distances detected at $\sim 5\sigma$

Measuring tensions

Agreement or disagreement?

Seemingly easy question: do experiments agree?

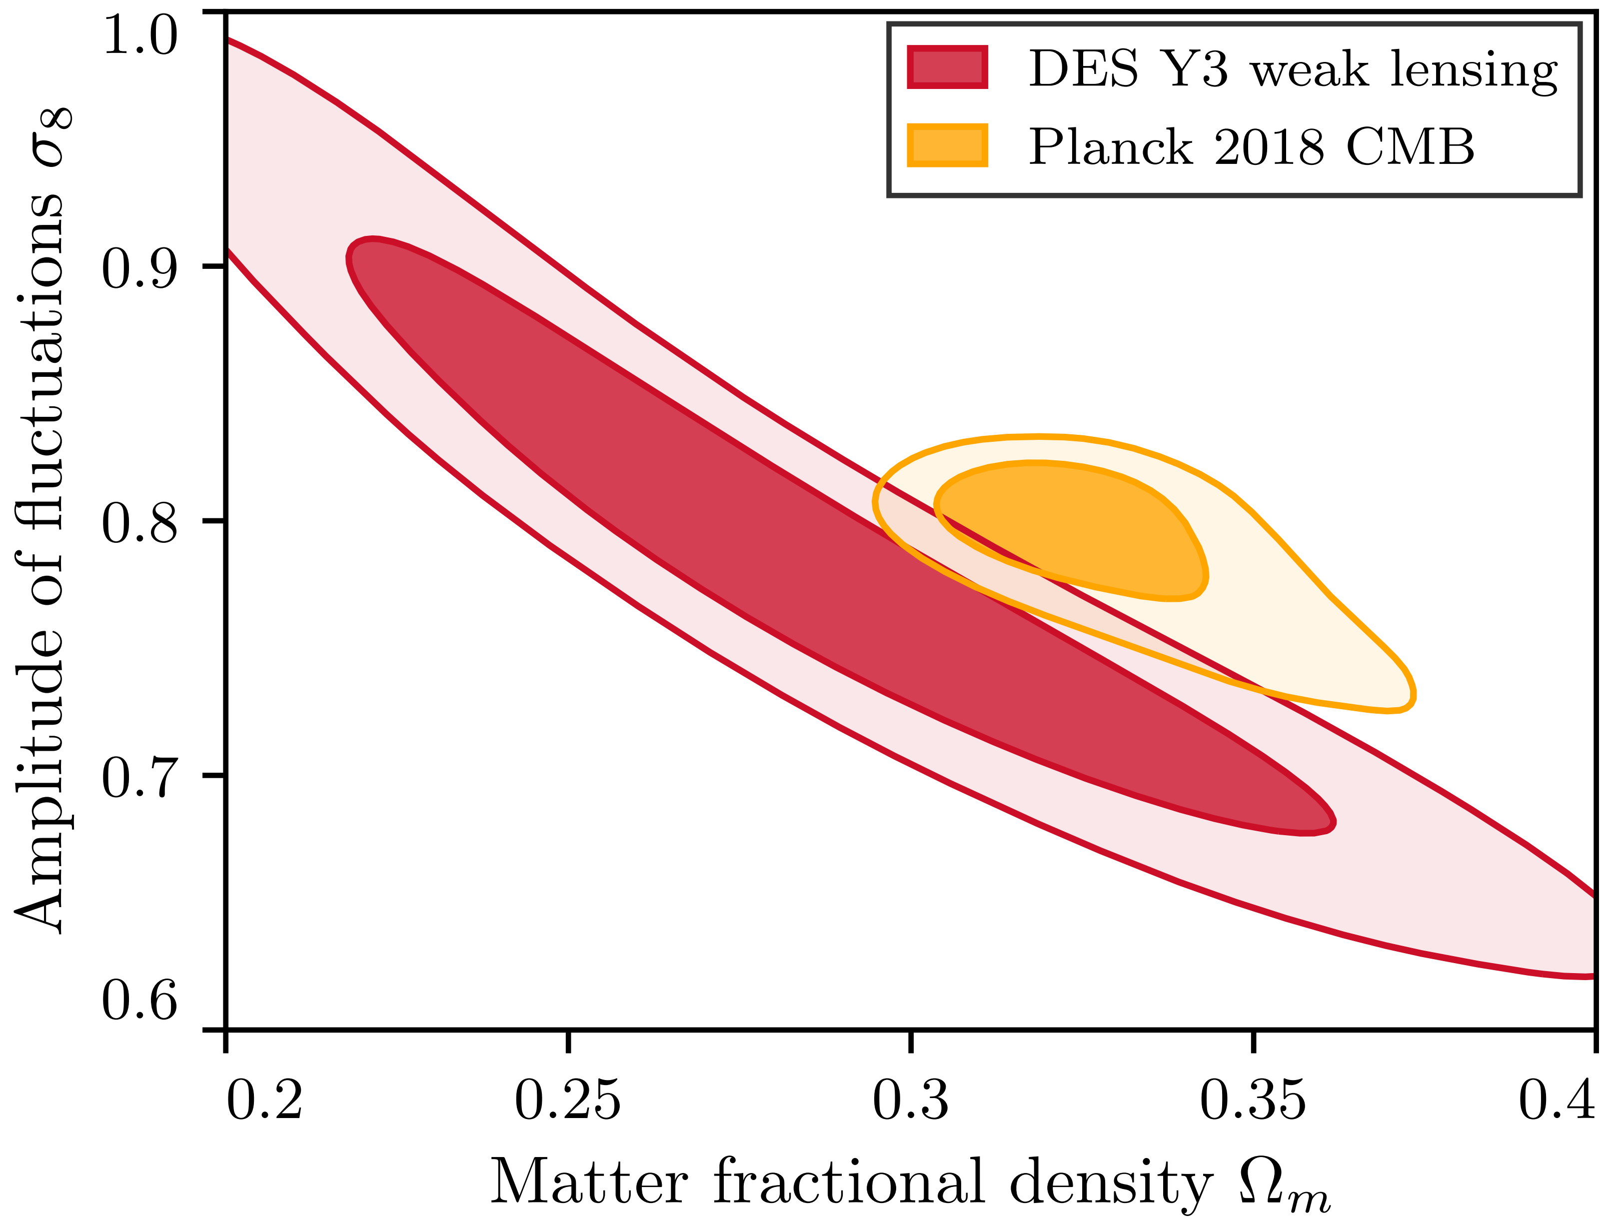

(Based on some work in DES)

Testing concordance: the challenges

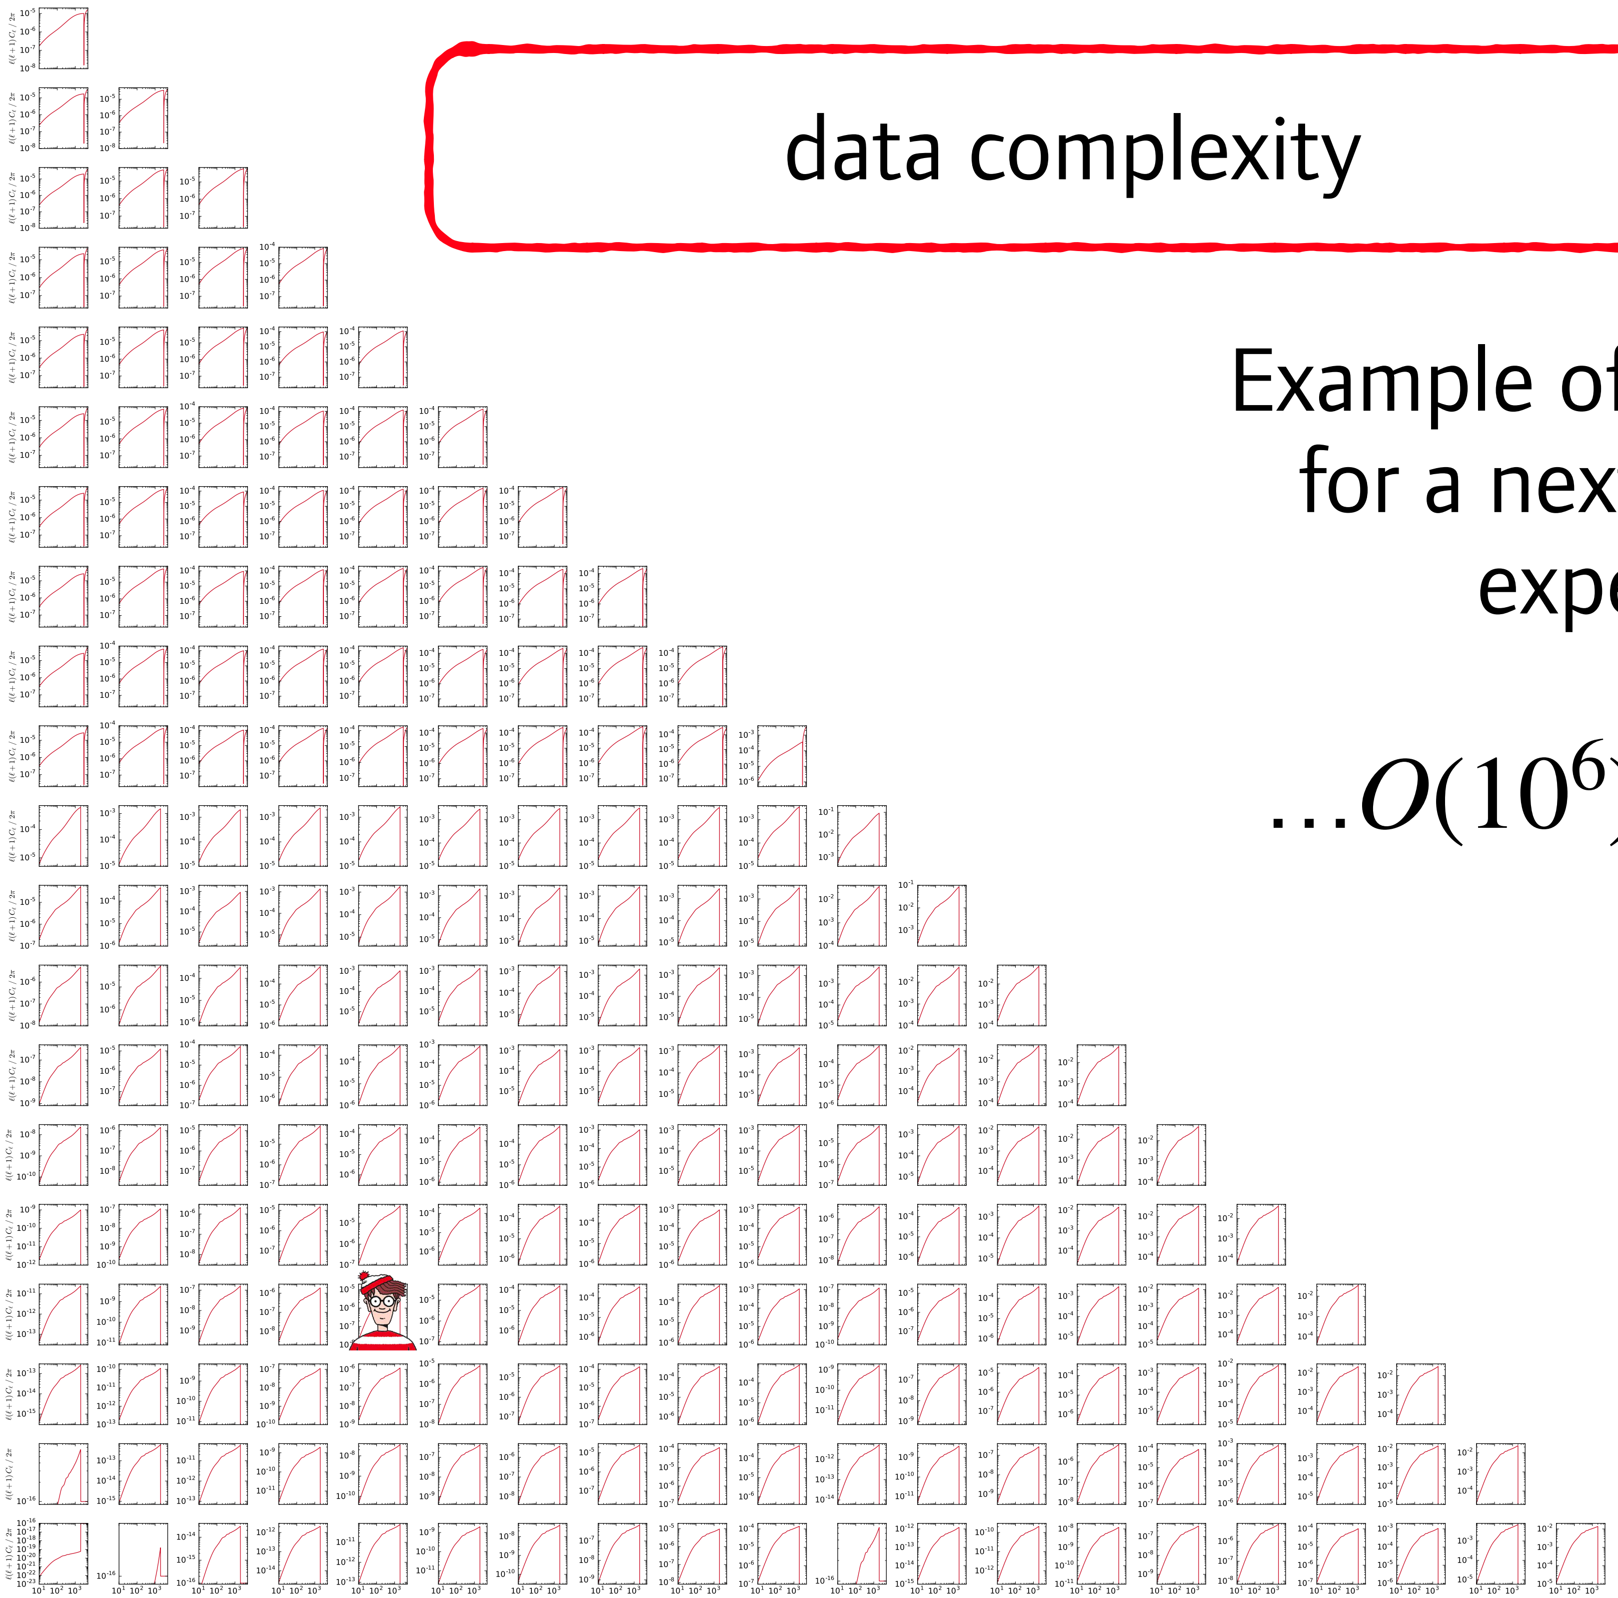

data complexity

Example of a data vector for a next generation experiment

$\ldots O(10^6)$ data points!

(for the cosmologists, it is Euclid)

Testing concordance: the challenges

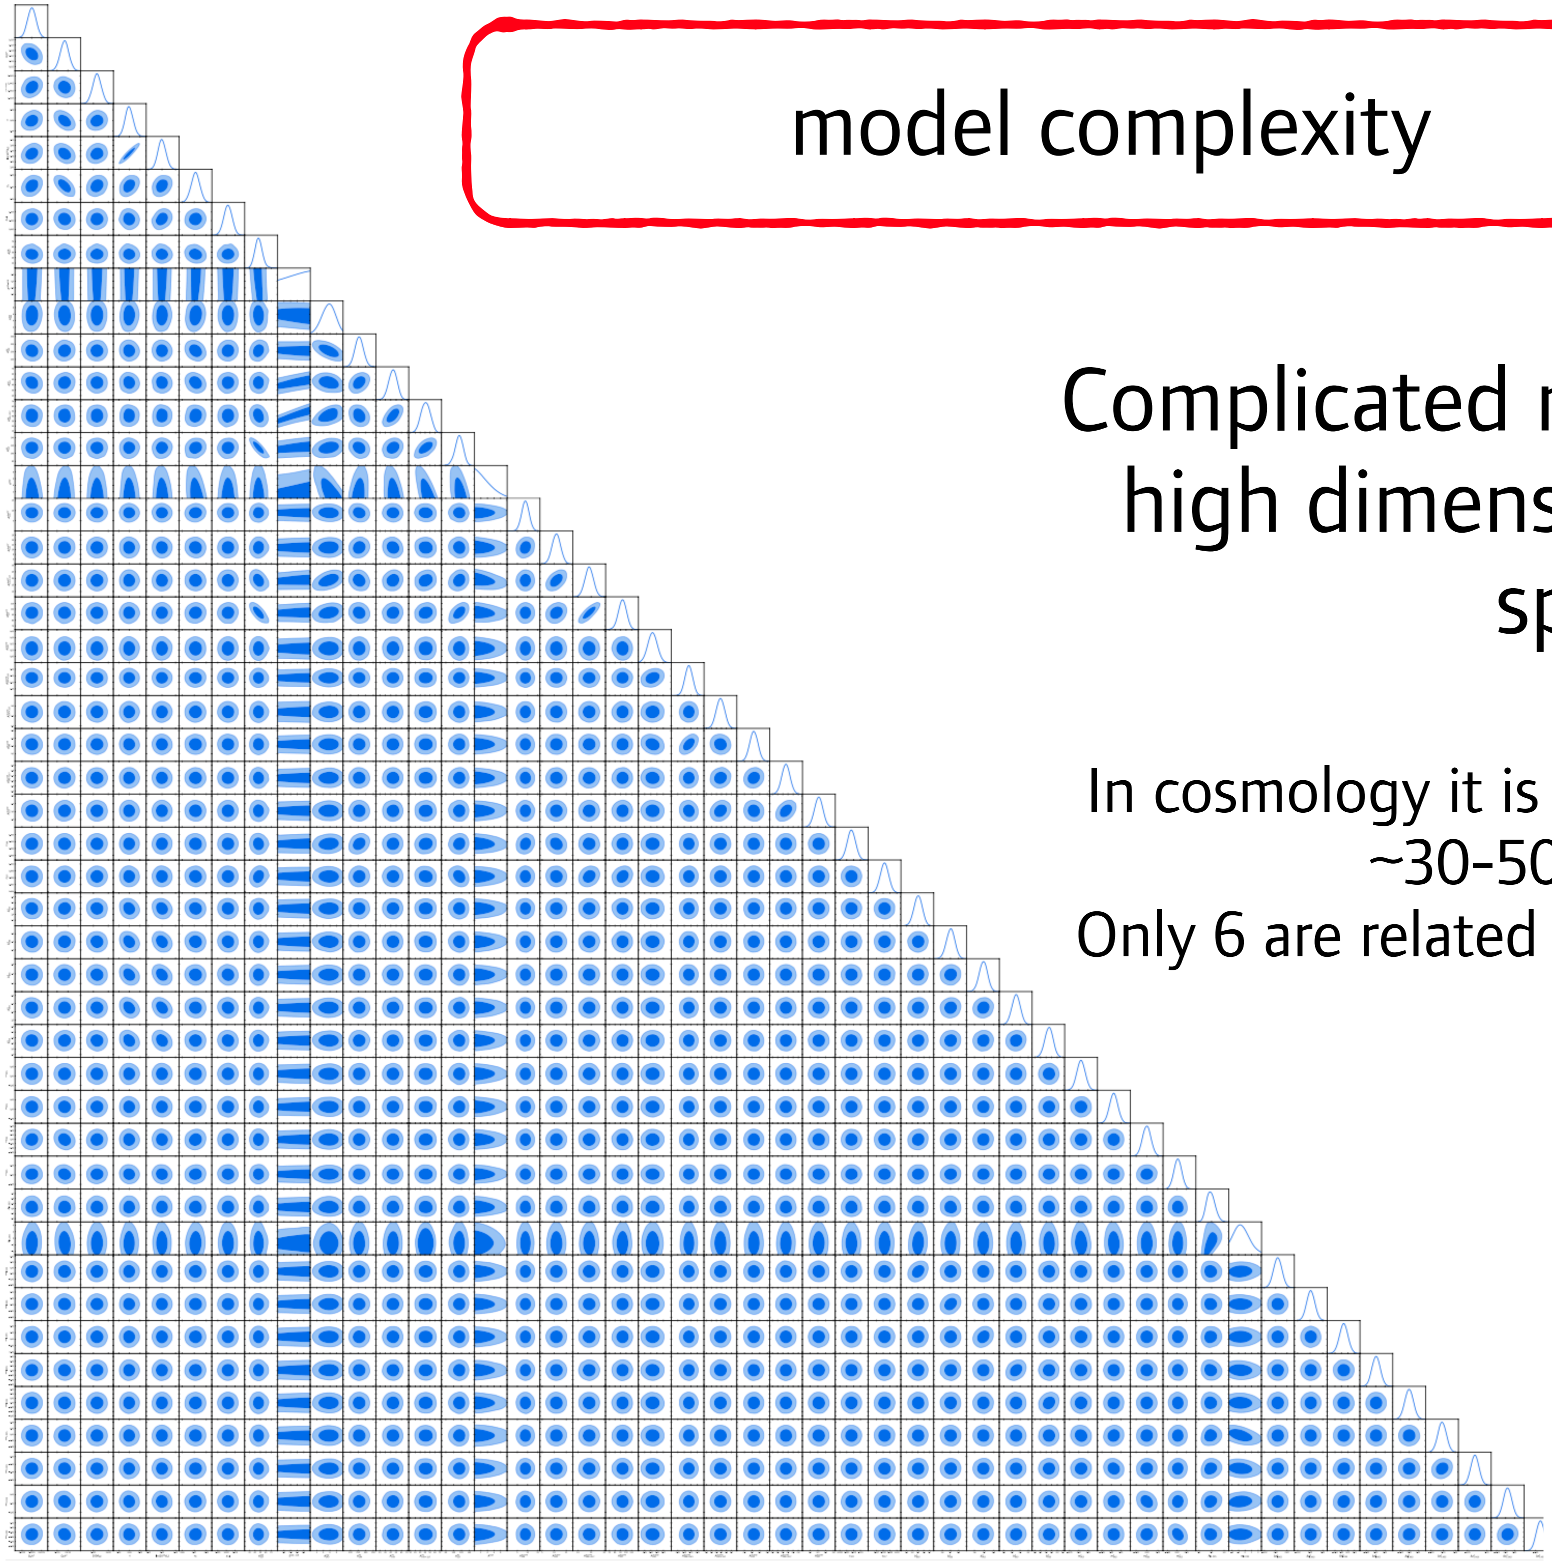

model complexity

Complicated models come with high dimensional parameter spaces.

In cosmology it is customary to work with ~30-50 parameters.
Only 6 are related to fundamental physics.

Testing concordance: the challenges

Projections of the parameter space might hide discrepancies

Do they agree?

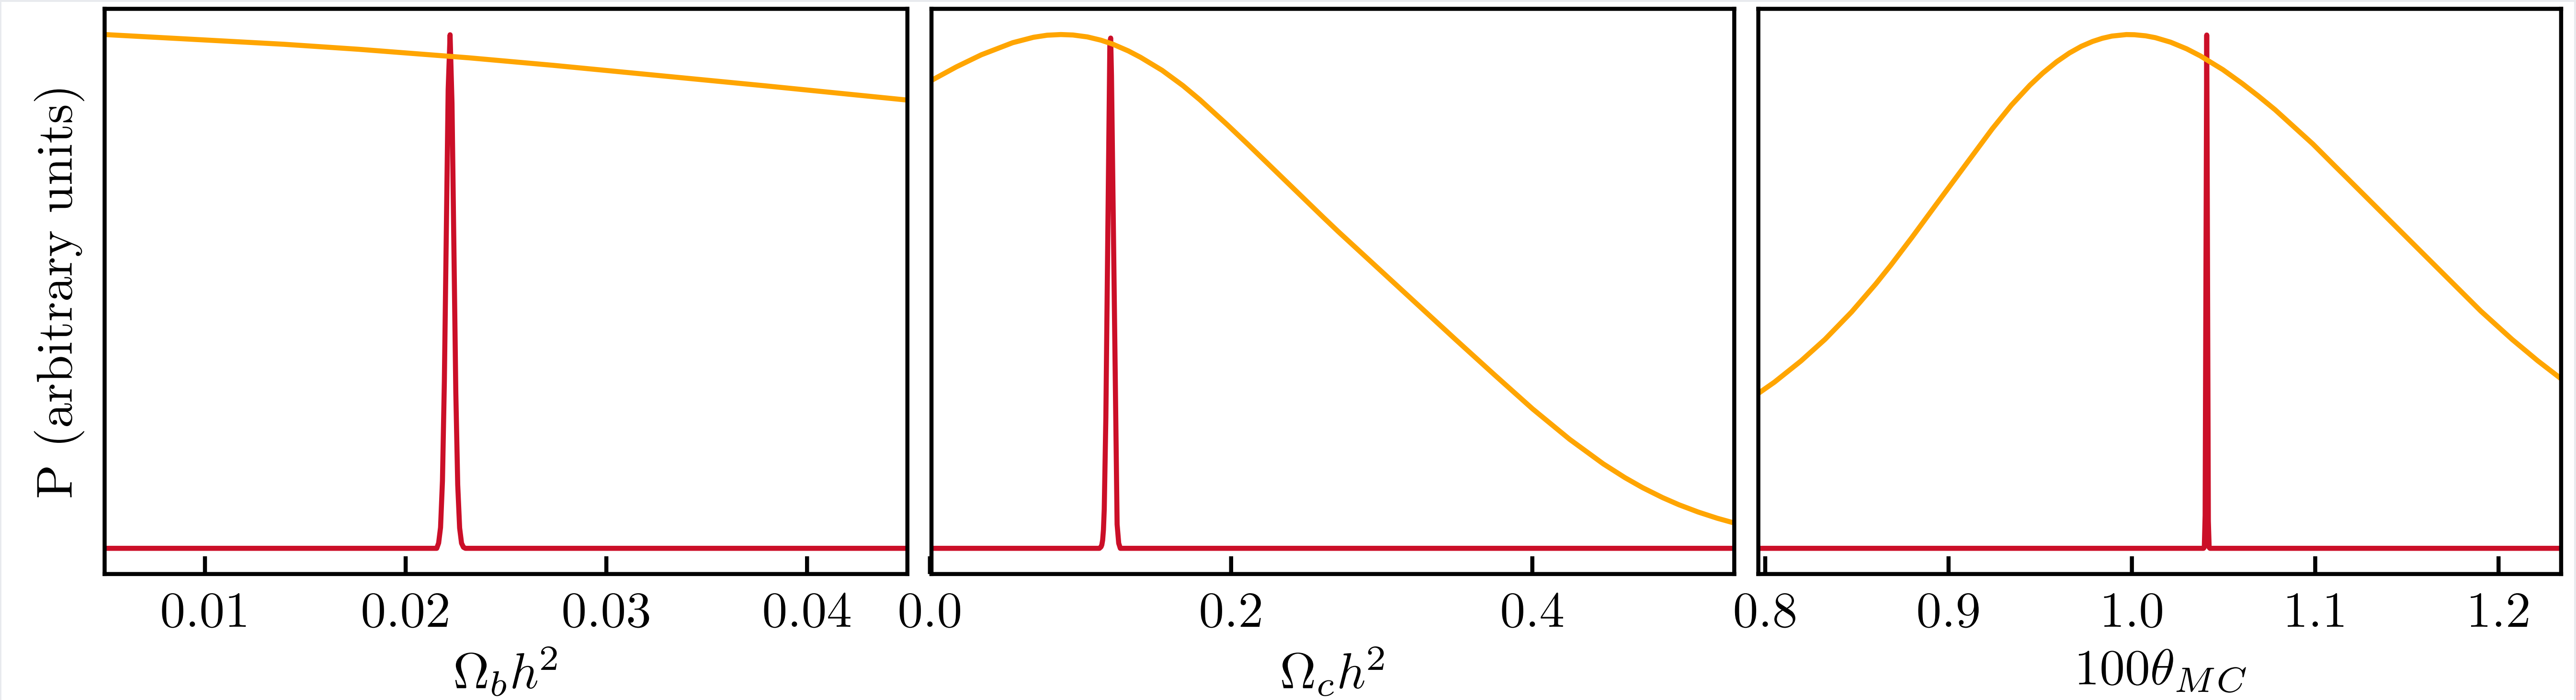

What are these data sets?

Testing concordance: the challenges

Projections of the parameter space might hide discrepancies

Do they agree? No, to 5 sigma

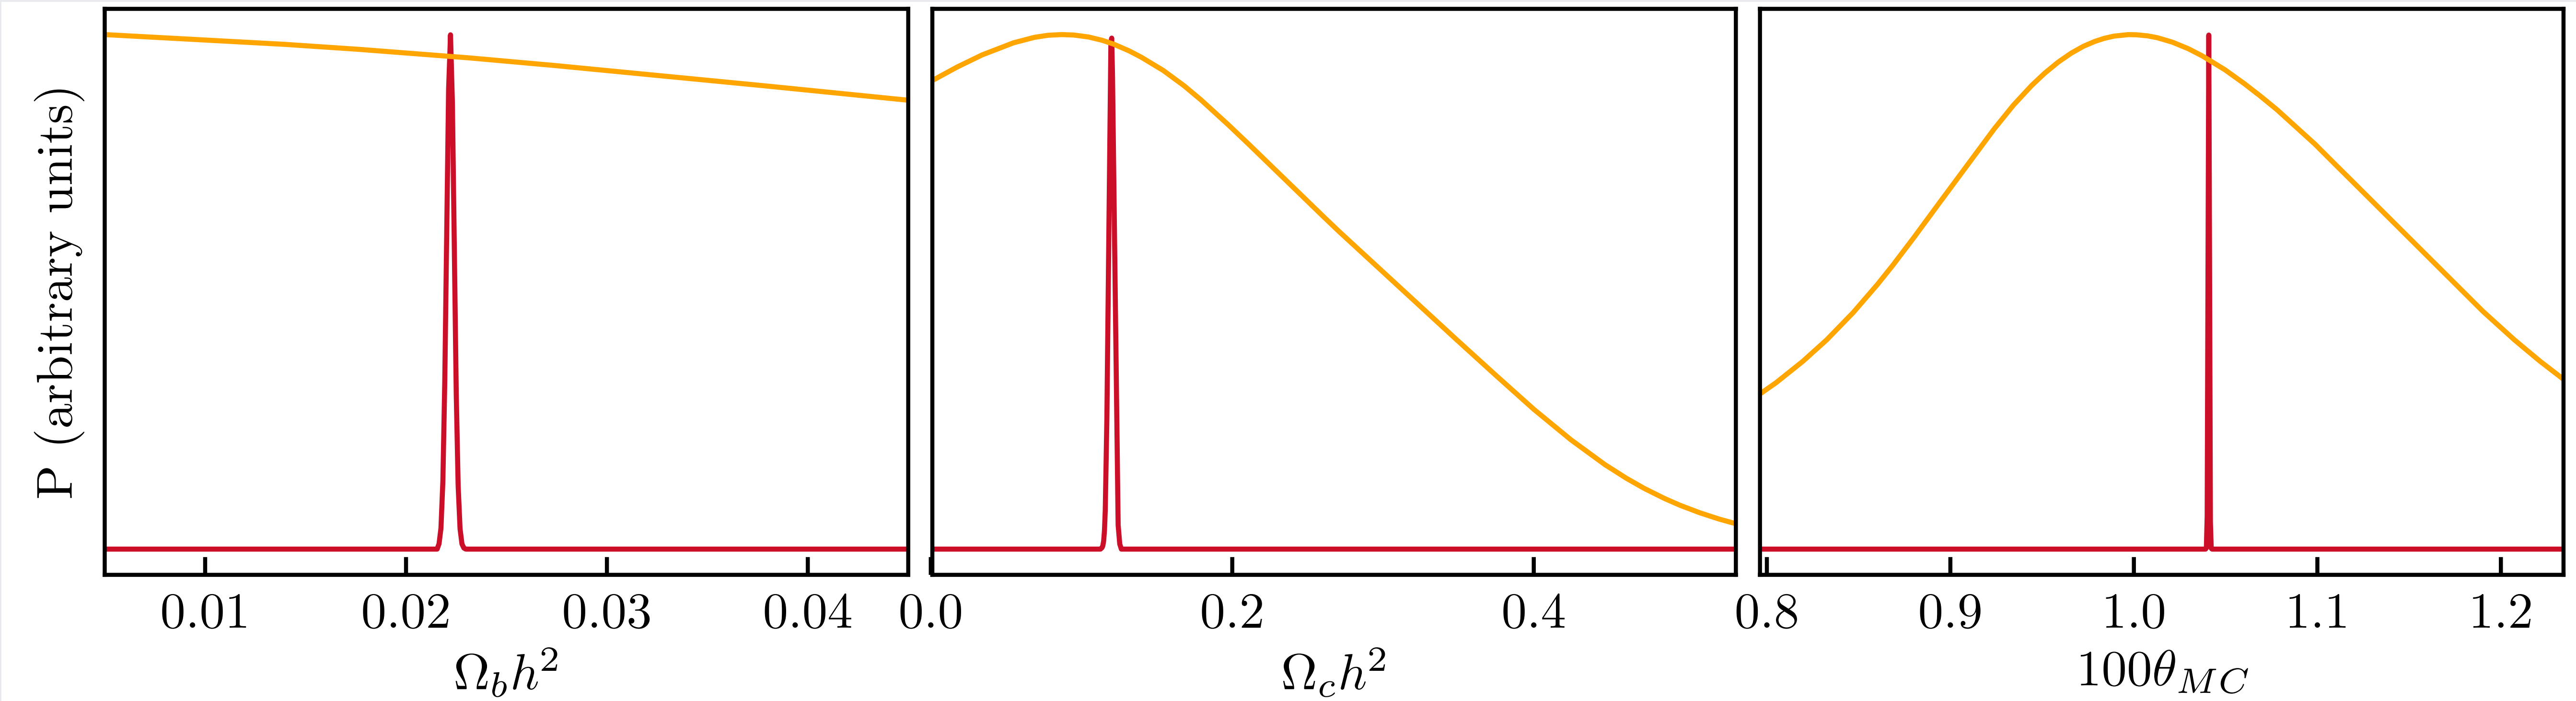

What are these data sets? Planck CMB and local Hubble constant

Testing concordance: the challenges

What is going on?

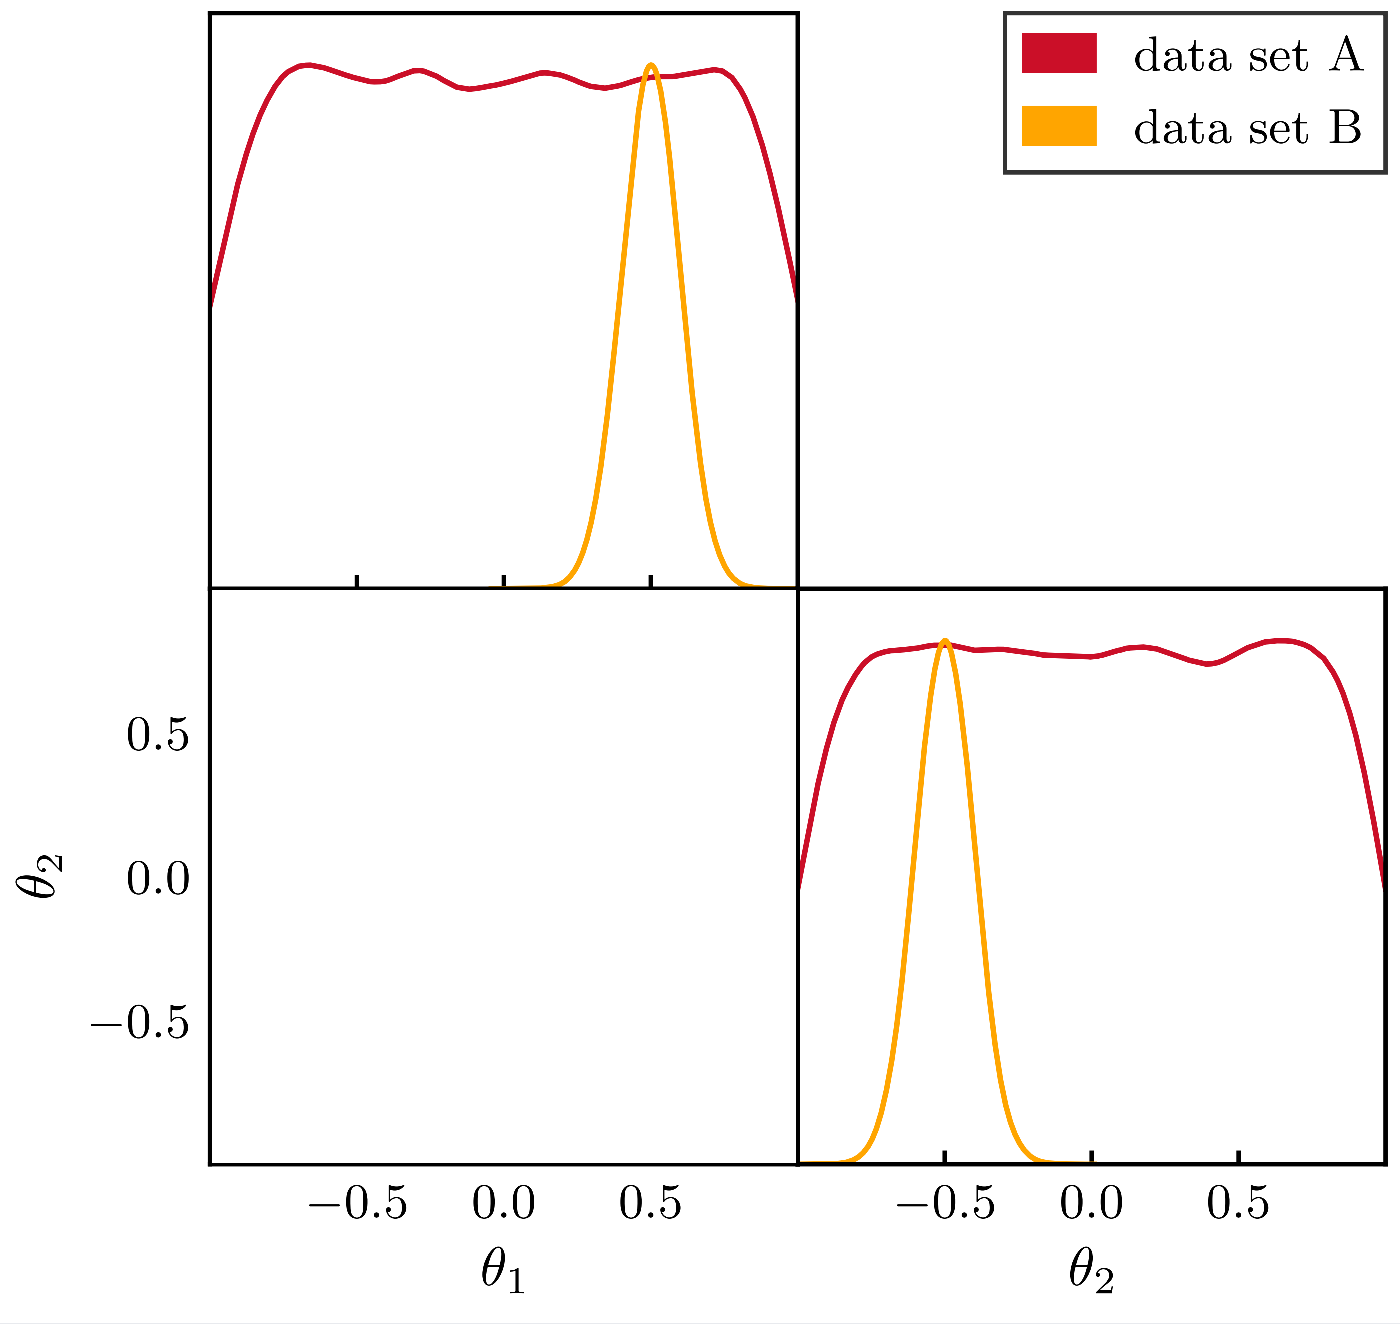

(made up posteriors from Lemos, MR et al arXiv:2012.09554)

Testing concordance: the challenges

We are being tricked by low dimensional projections

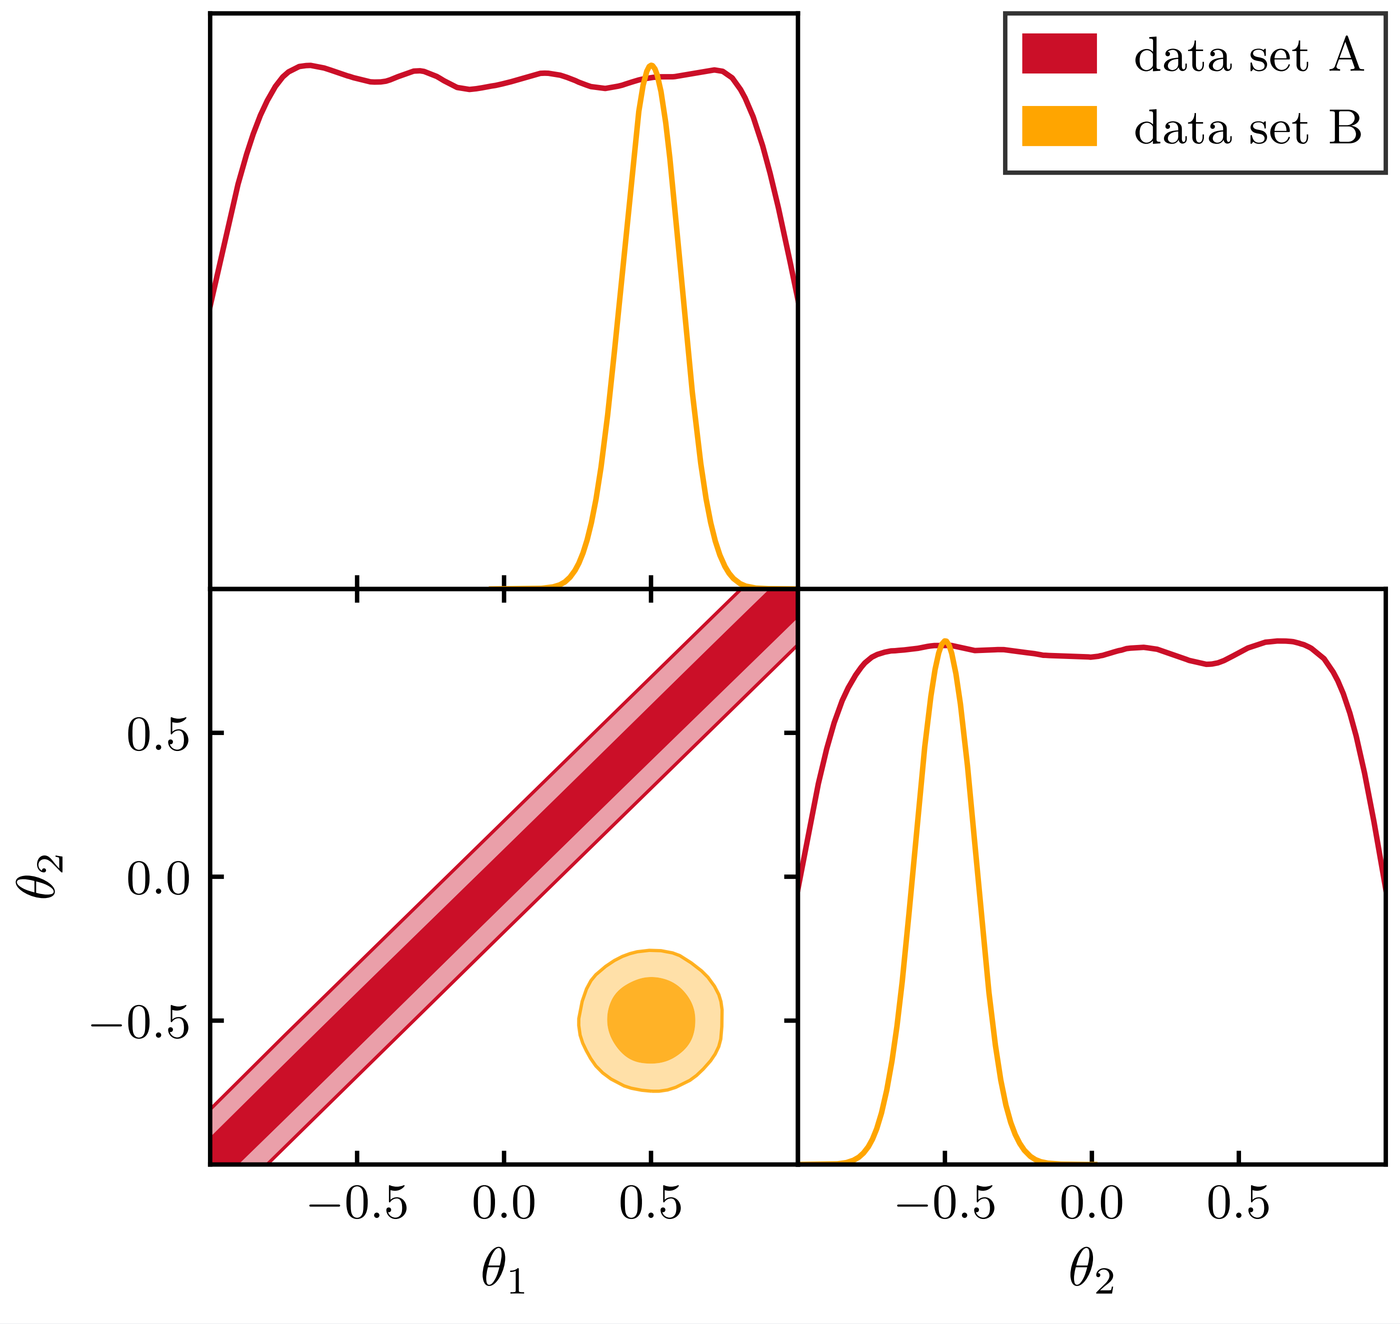

(made up posteriors from Lemos, MR et al arXiv:2012.09554)

Testing concordance: the tools of the trade

Theoretical papers that I will talk about:

- Raveri and Hu, "Concordance and Discordance in Cosmology" **arXiv 1806.04649**
- Raveri, Zacharegkas, Hu, "Quantifying concordance of correlated cosmological data sets" **arXiv 1912.04880**
- Raveri and Doux, "Non-Gaussian estimates of tensions in cosmological parameters" **arXiv 2105.03324**

Code implementation (in Python, with several example notebooks)

```
~ pip install tensiometer
```

Relies on GetDist for handling parameter distributions

```
~ pip install getdist
```

…and tensorflow for ML methods … which complicates installation…

Gaussian estimators: parameter differences 1D

The usual recipe

(exact for Gaussian models with only 1 parameter)

1D parameter shifts

$$T_1 \equiv \frac{|\theta_1 - \theta_2|}{\sqrt{\sigma_1^2 + \sigma_2^2}}$$

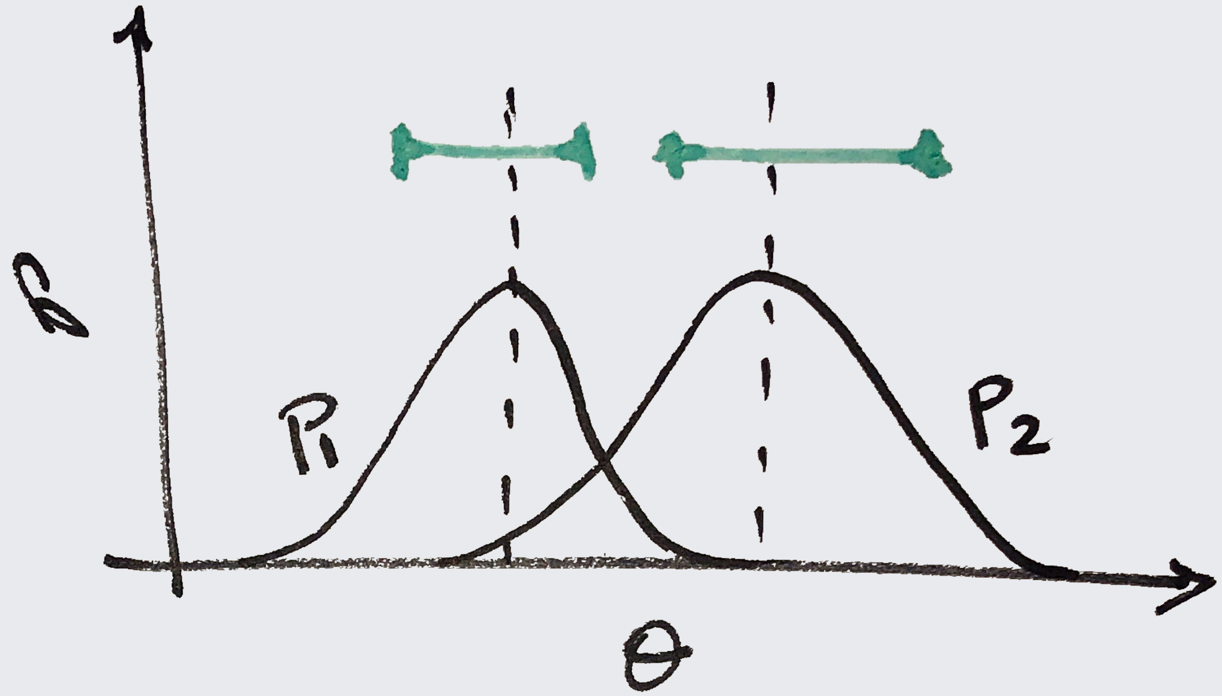

Gaussian estimators: parameter differences ND

ND parameter shifts

$$Q_{\rm DM} \equiv (\theta_p^1 - \theta_p^2)^T (\mathcal{C}_{p1} + \mathcal{C}_{p2} - \mathcal{C}_{p1}\mathcal{C}_{\Pi}^{-1}\mathcal{C}_{p2} - \mathcal{C}_{p2}\mathcal{C}_{\Pi}^{-1}\mathcal{C}_{p1})^{-1} (\theta_p^1 - \theta_p^2)$$
$$\sim \chi^2(\mathrm{rank}[\mathcal{C}_{p1} + \mathcal{C}_{p2} - \mathcal{C}_{p1}\mathcal{C}_{\Pi}^{-1}\mathcal{C}_{p2} - \mathcal{C}_{p2}\mathcal{C}_{\Pi}^{-1}\mathcal{C}_{p1}])\,.$$

Significantly more complicated…

The complication arises from the removal of the prior.
A prior constrained parameter combination cannot contribute to tensions…

I am still searching for a numerically reliable way of computing this…

(see R&H 1806.04649 for the proof)

Gaussian estimators: parameter differences ND

ND parameter shifts (in update form)

$$Q_{\rm UDM} \equiv (\theta_p^1 - \theta_p^{12})^T (\mathcal{C}_{p1} - \mathcal{C}_{p12})^{-1} (\theta_p^1 - \theta_p^{12}) \sim \chi^2(\text{rank}[\mathcal{C}_{p1} - \mathcal{C}_{p12}]) \,.$$

Same as standard parameter shifts, kind of magic…

Weights directions that dataset 2 improves over 1, in this way can be computed reliably

Offers mitigation of non-Gaussianities
(1 can be the most Gaussian of the two datasets)

(see R&H 1806.04649 for the proof)

Gaussian estimators: compatibility with the prior

The update parameter shift estimator offers a nice additional internal consistency check

Sometimes we have Gaussian priors and the posterior is an update of the prior

ND prior compatibility

$$Q_{\rm UDM} \equiv (\theta_\Pi - \theta_p)^T (\mathcal{C}_\Pi - \mathcal{C}_p)^{-1} (\theta_\Pi - \theta_p)$$

Non-gaussian estimators: parameter differences

Full dimensional distribution of differences in parameters

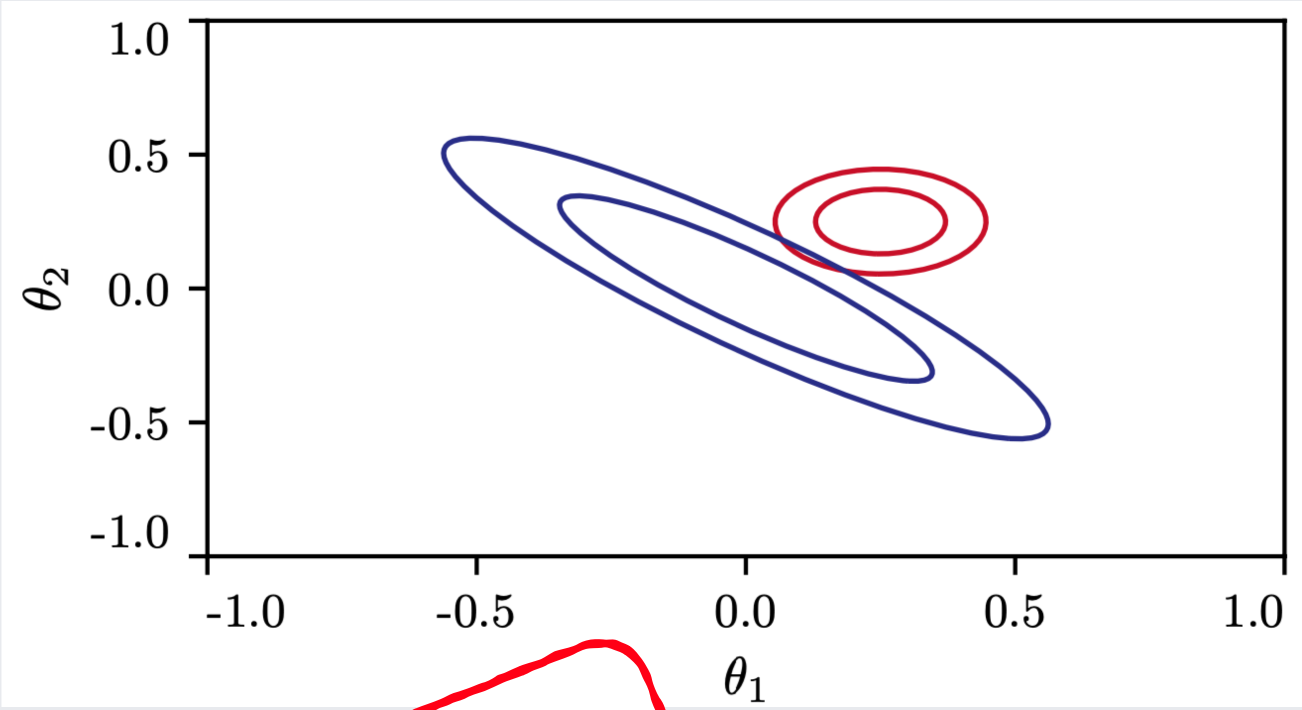

Start with:

$$P_1(\theta_1)P_2(\theta_2)$$

change variables:

$$P_1(\theta_1)P_2(\theta_1 - \Delta\theta)$$

integrate out

Distribution of parameter differences

$$P(\Delta\theta) = \int P_1(\theta_1)P_2(\theta_1 - \Delta\theta)d\theta_1$$

(Raveri & Hu arXiv:1806.04649, Raveri & Doux arXiv:2105.03324)

Non-gaussian estimators: parameter differences

$$P(\Delta\theta) = \int P_1(\theta_1) P_2(\theta_1 - \Delta\theta) d\theta_1$$

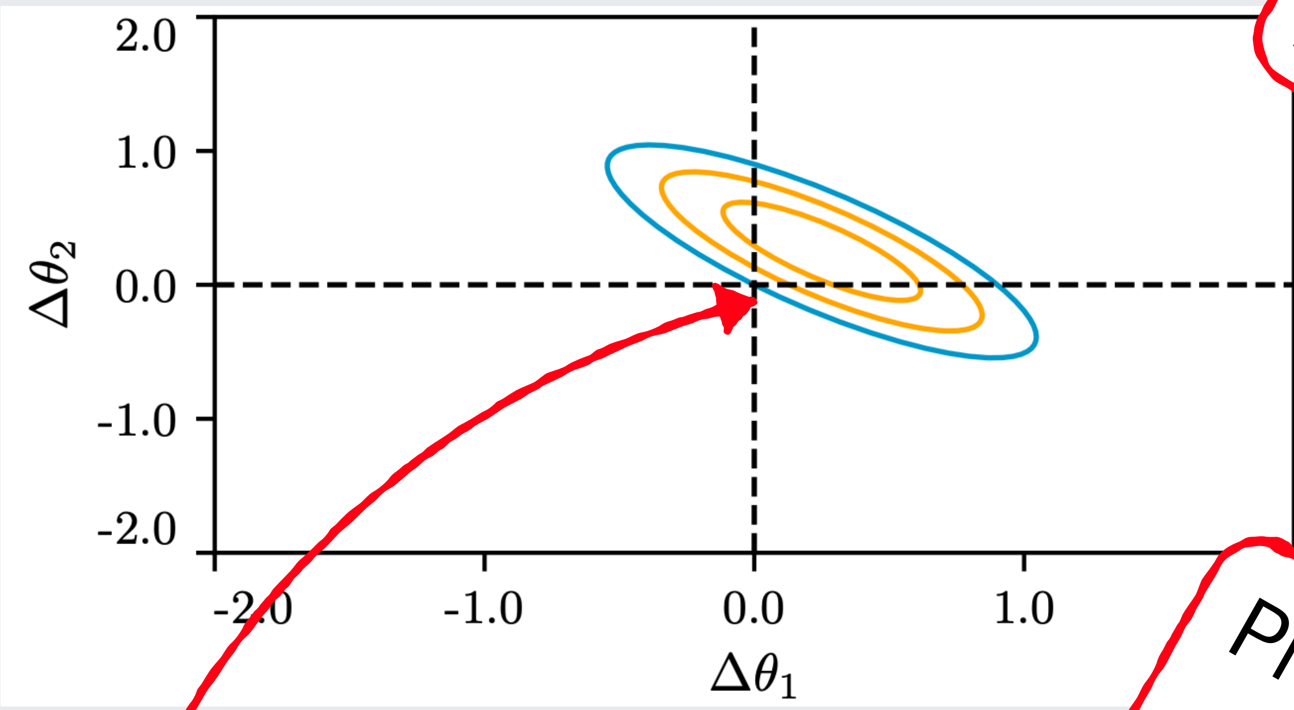

Distribution of parameter differences

Our goal:

$$\Delta = \int_{P(\Delta\theta)>P(0)} P(\Delta\theta) d\Delta\theta$$

Probability of non-zero difference

(Raveri & Hu arXiv:1806.04649, Raveri & Doux arXiv:2105.03324)

Parameter differences: the bad news...

Distribution of parameter differences

$$P(\Delta\theta) = \int P_1(\theta_1)P_2(\theta_1 - \Delta\theta)d\theta_1$$

Usually **very high dimensional** integral

Probability of non-zero difference

$$\Delta = \int_{P(\Delta\theta)>P(0)} P(\Delta\theta)d\Delta\theta$$

Usually **very high dimensional** integral

Parameter differences in practice

Distribution of parameter differences

$$P(\Delta\theta) = \int P_1(\theta_1) P_2(\theta_1 - \Delta\theta) d\theta_1$$

Usually **very high dimensional** integral

First integral can be done from MCMC chains:
difference of samples is a sample of parameter differences

(Raveri and Doux arXiv:2105.03324)

Parameter differences in practice

Probability of non-zero difference

$$\Delta = \int_{P(\Delta\theta)>P(0)} P(\Delta\theta) d\Delta\theta$$

Second integral can be done with KDE but is very expensive

Naive algorithm is N^2 (not doable)
R&D arXiv:2105.03324 has the NlogN algorithm which is still very expensive (curse of dimensionality of KDEs)

(Raveri and Doux arXiv:2105.03324)

Parameter differences and machine learning

$$\Delta = \int_{P(\Delta\theta)>P(0)} P(\Delta\theta) d\Delta\theta$$

Probability of non-zero difference

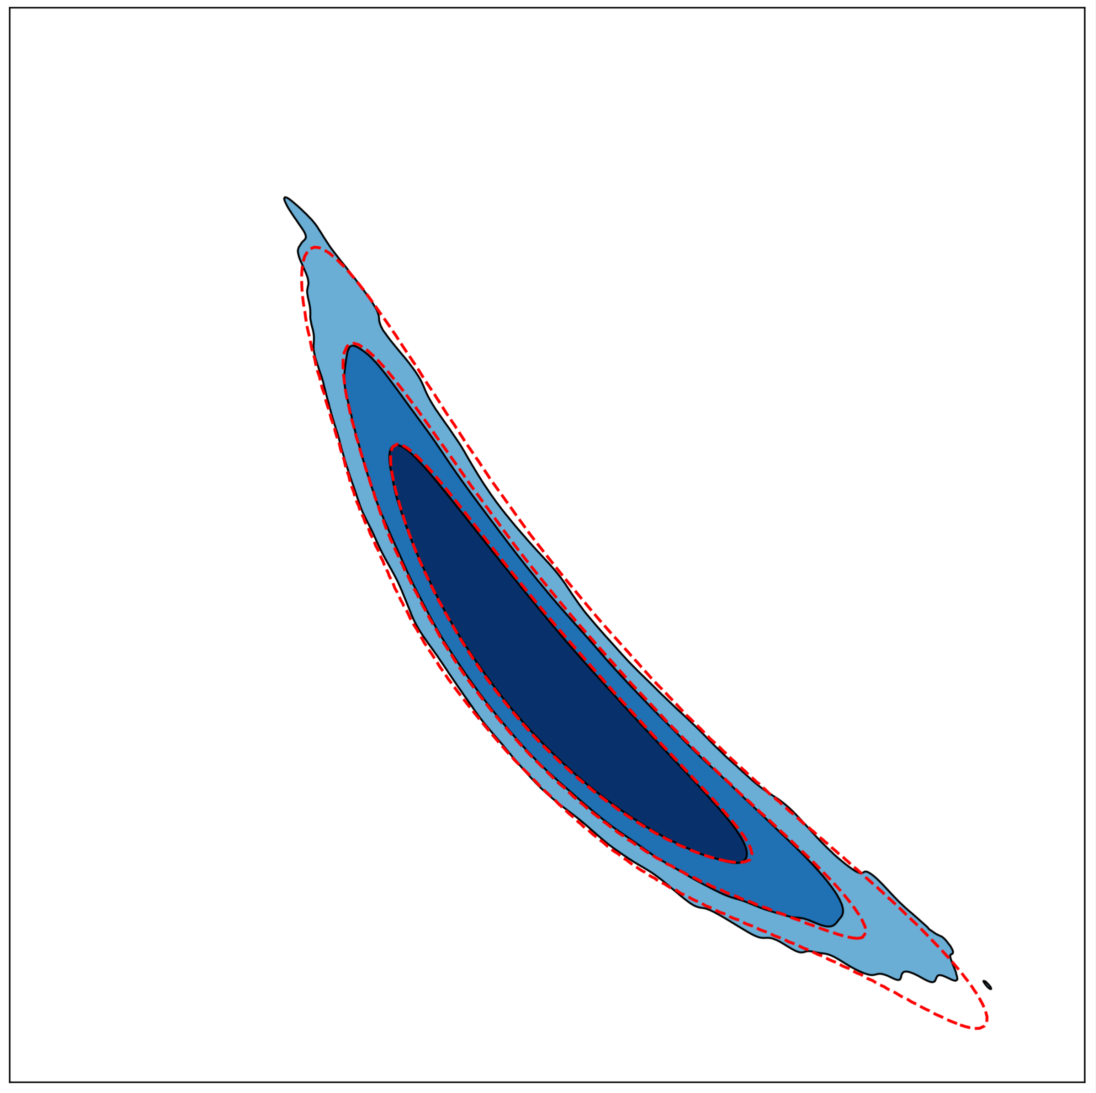

Solution:
learn the parameter difference posterior

$$P(\Delta\theta)$$

…then MC integrate

(Raveri and Doux arXiv:2105.03324)

Normalizing flows

MAF normalizing flows

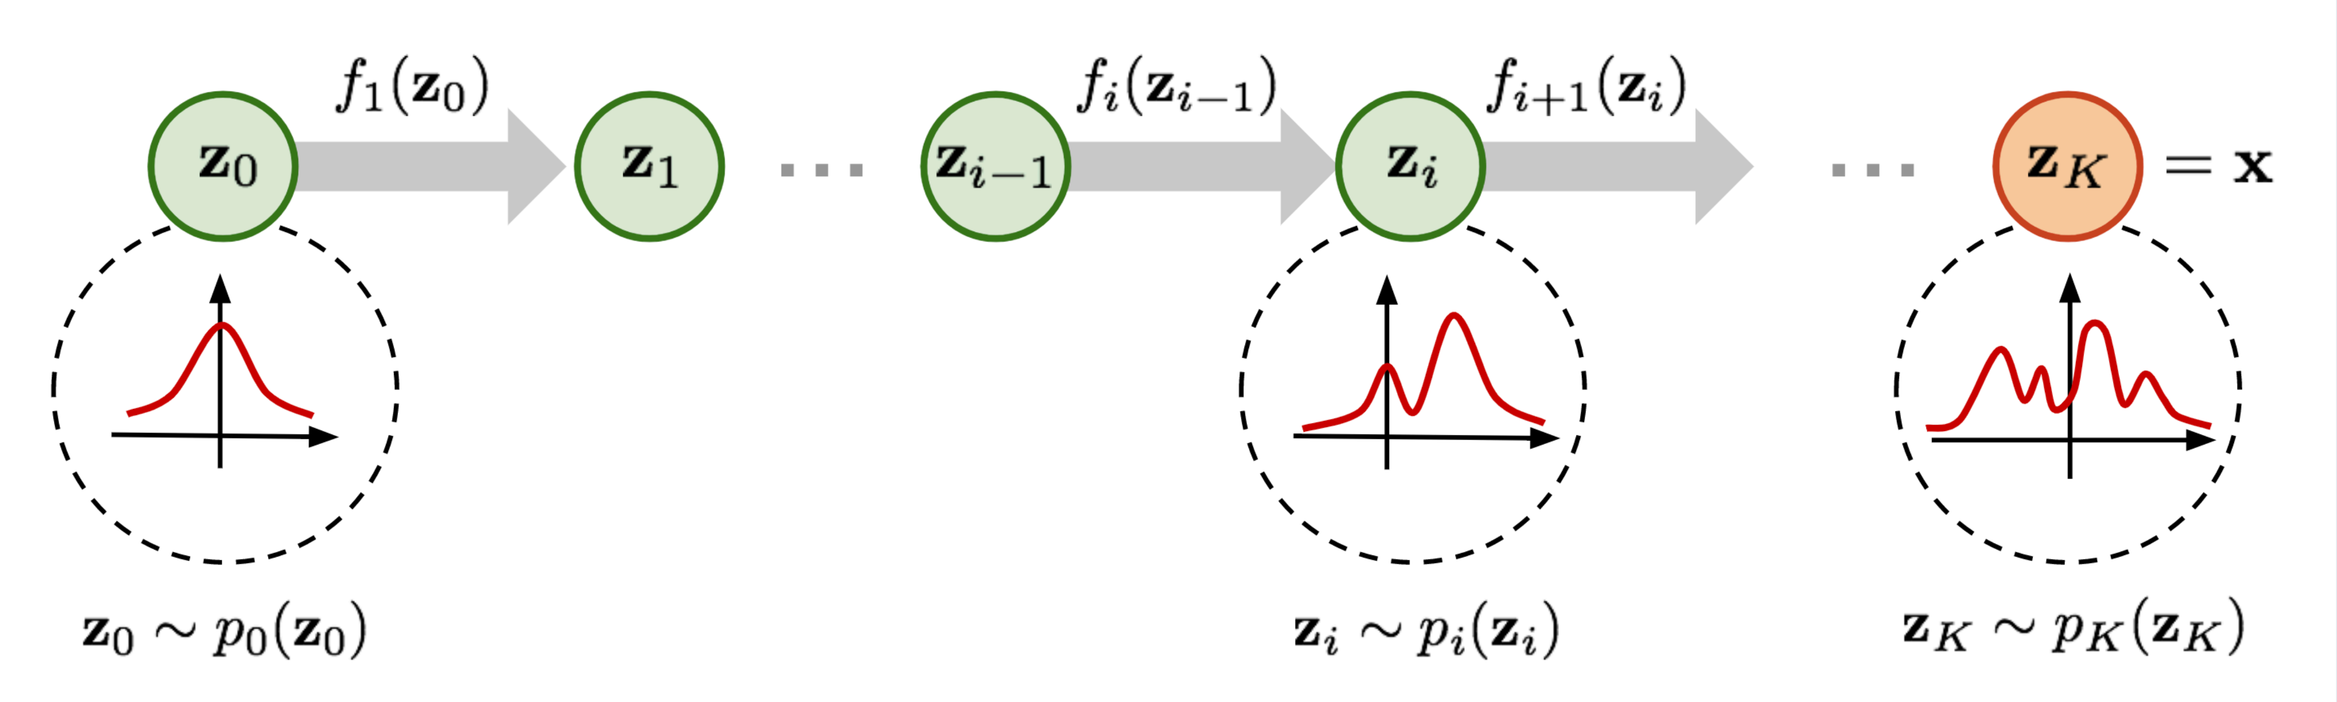

$$
\begin{aligned}
y_1 &= \mu_1 + \sigma_1 z_1 \\
y_i &= \mu(y_{1\ldots i-1}) + \sigma(y_{1\ldots i-1}) z_i
\end{aligned}
$$

Stacked with permutations to ensure no coordinate is unlucky

(Papamakarios, Pavlakou, Murray arXiv:1705.07057)

Normalizing flows performances

…trained PDFs are indistinguishable from real (KDE) ones…

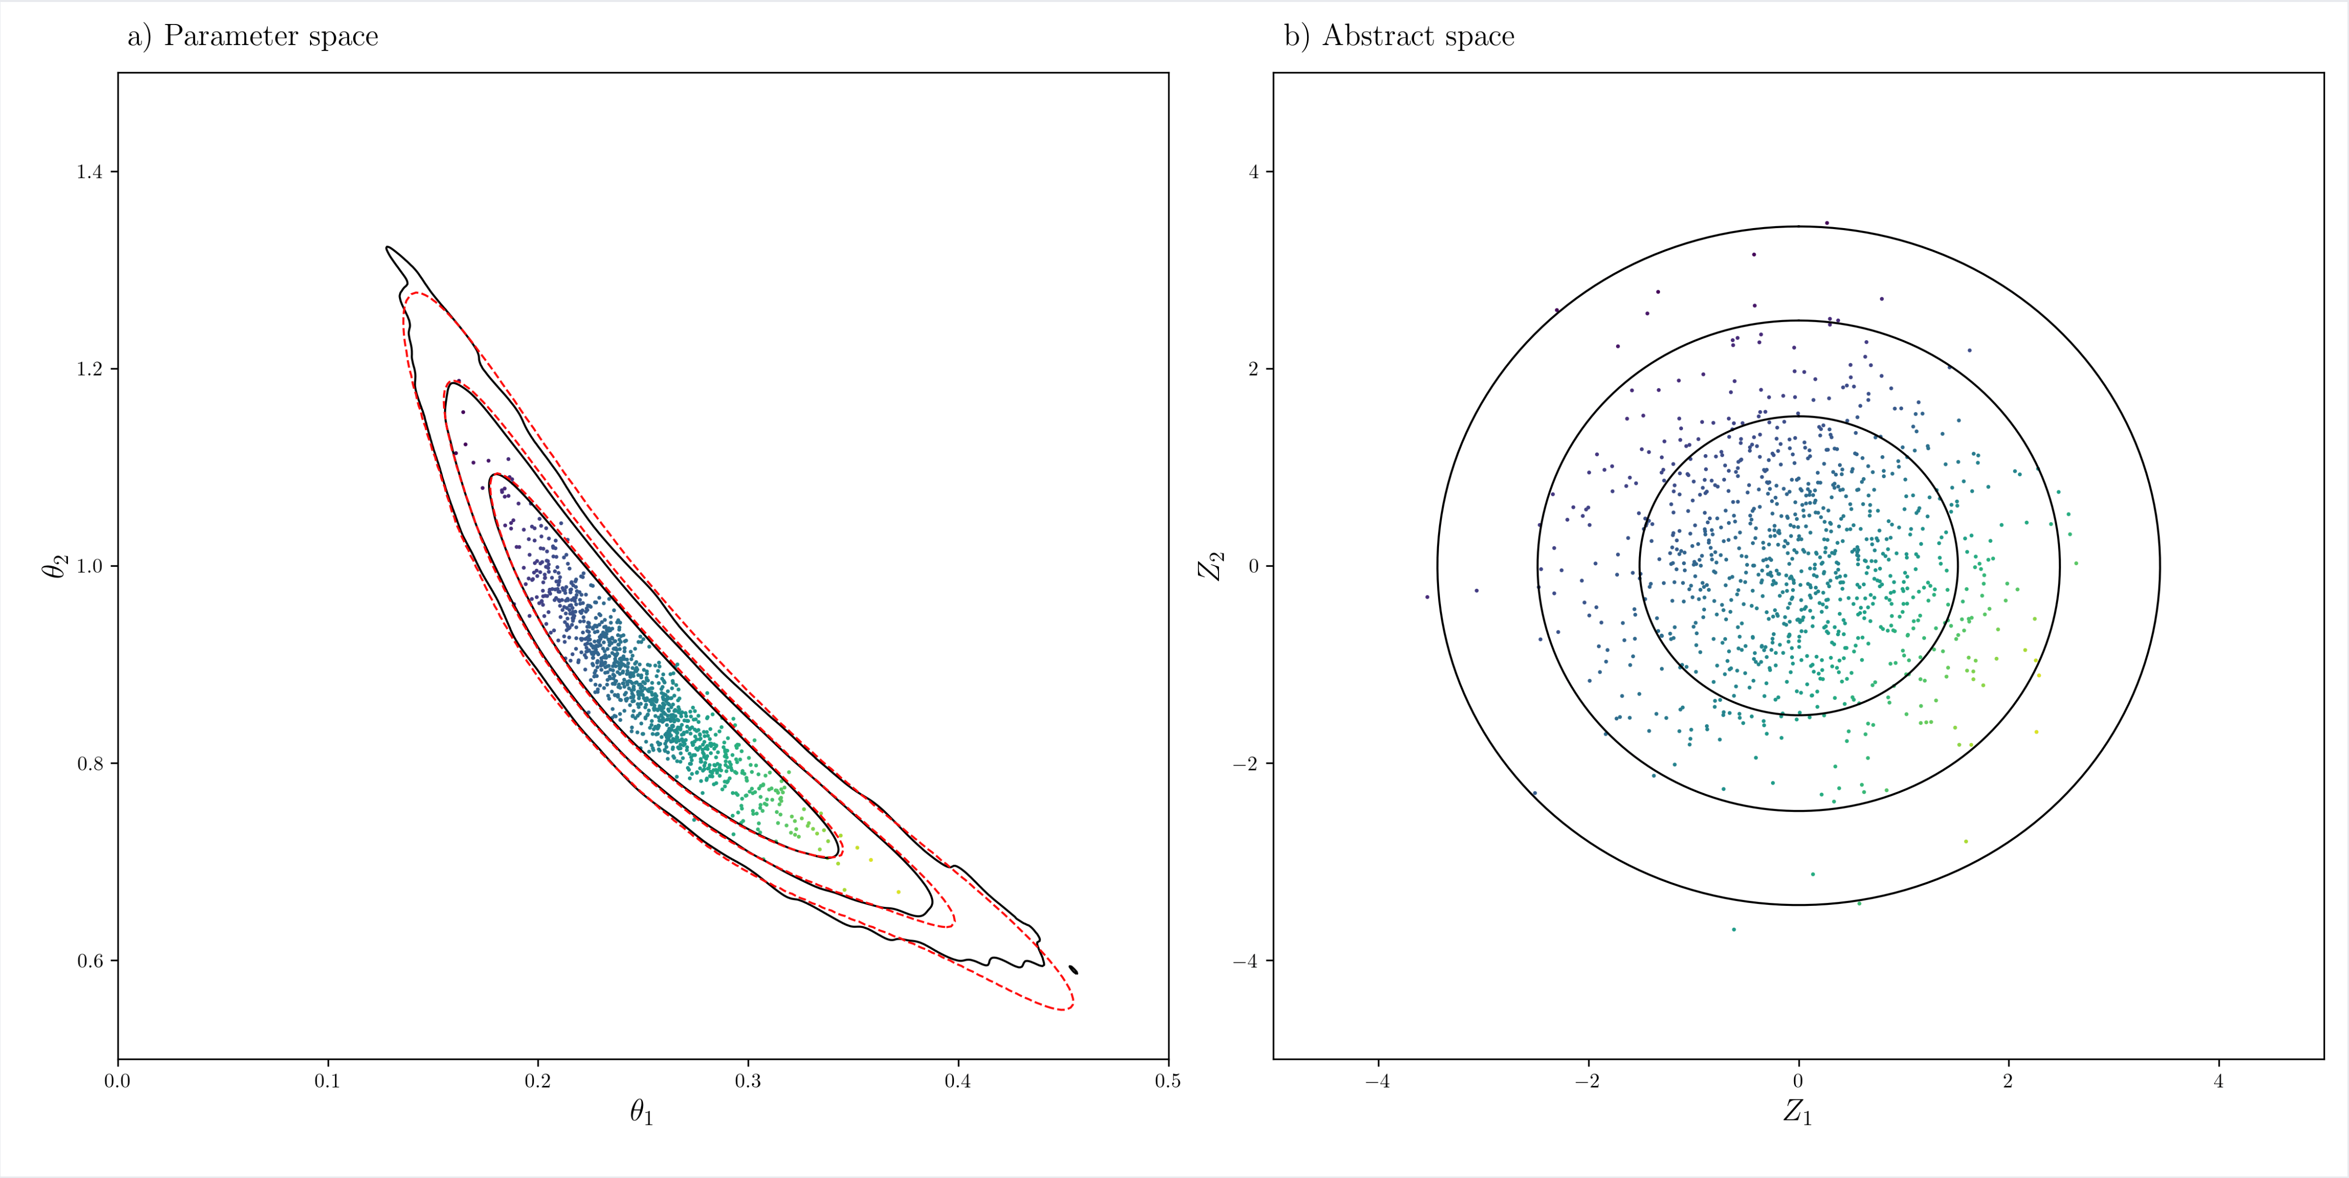

(Raveri and Doux arXiv:2105.03324)

Normalizing flows performances

Matches or outperforms KDE based methods in all our tests

KDE

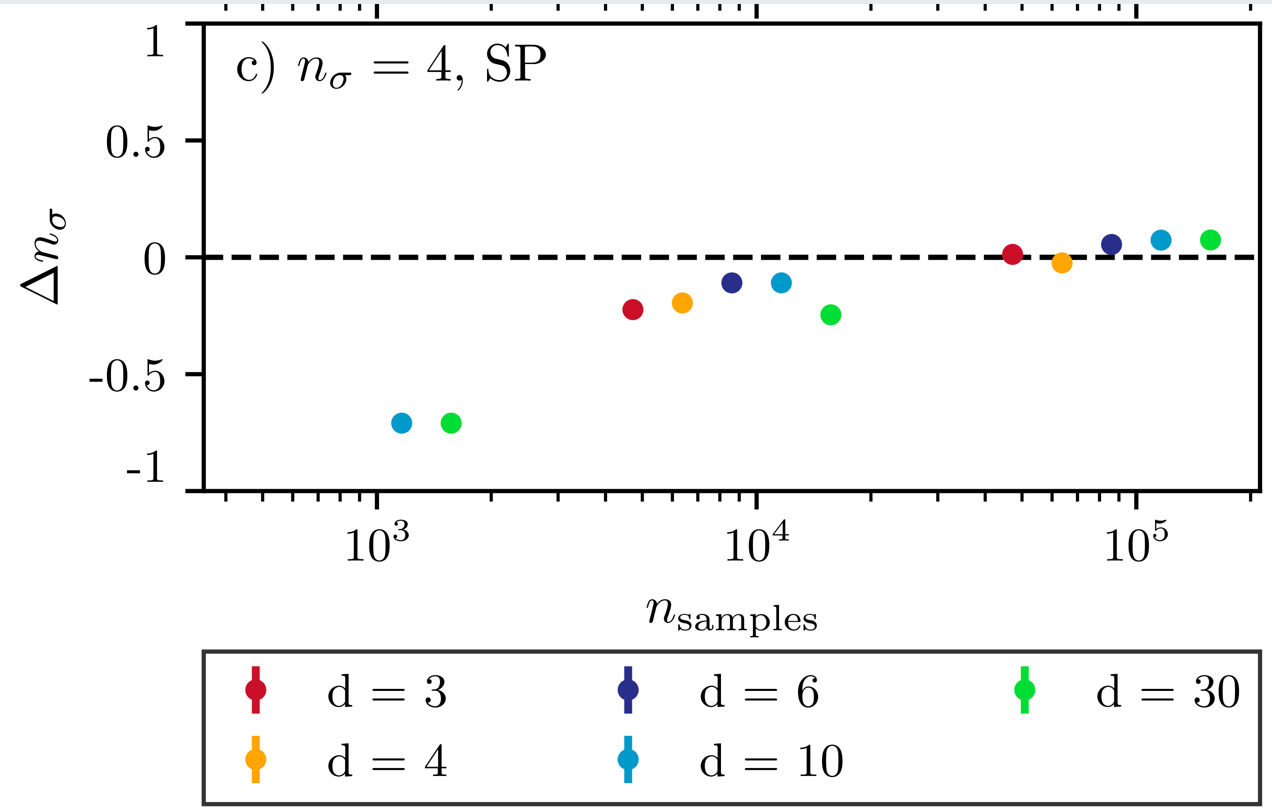

Flow

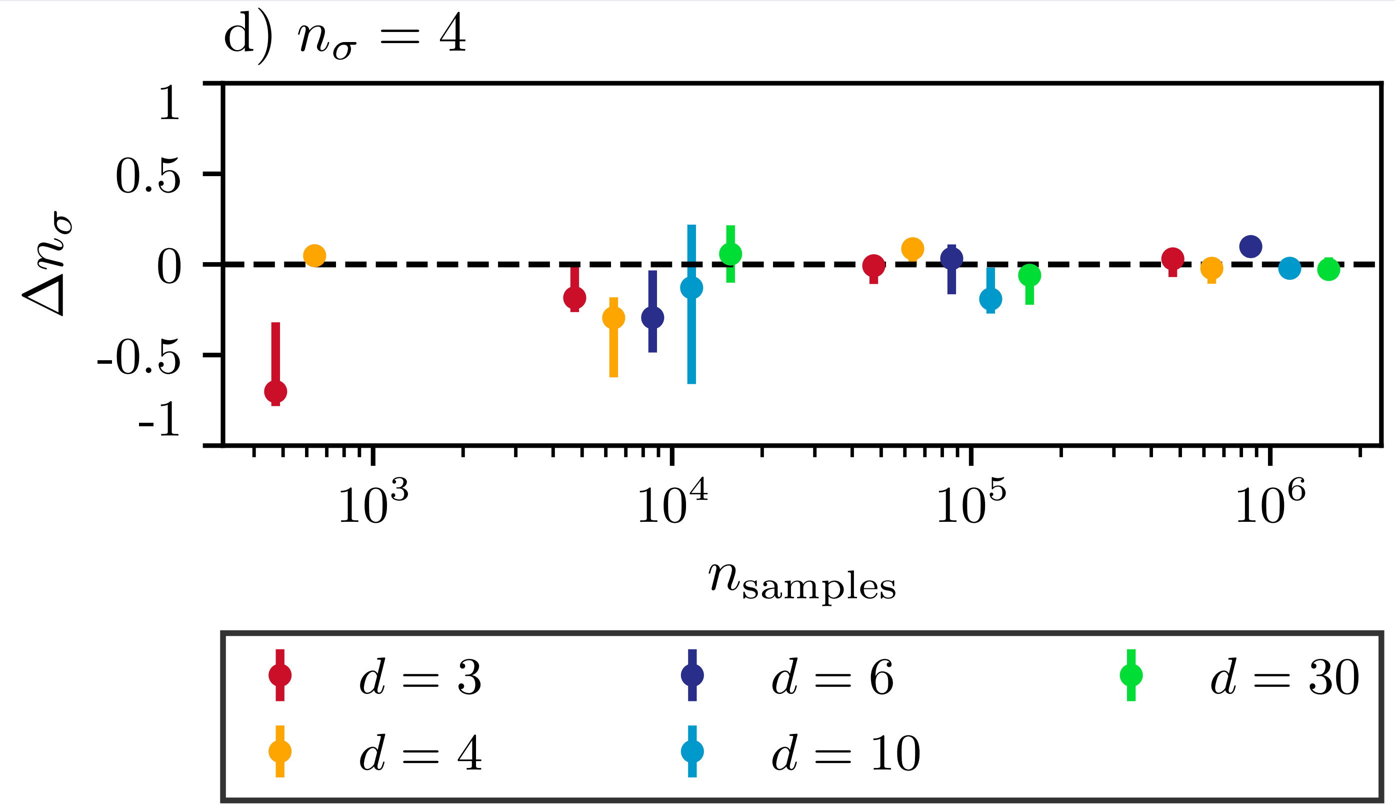

(Raveri and Doux arXiv:2105.03324)

End to end consistency check pipeline

Computational tools

```
~ pip install tensiometer
```

With several examples

Example calculation of agreement and disagreement between posterior distributions

Marco Raveri (mraveri@sas.upenn.edu)

This notebook shows an end to end calculation of the agreement and disagreement between two experiments with different methods, in an idealized case and in a realistic case involving two cosmological experiments.

In [1]:
```
# Initial setup:
%matplotlib inline
%config InlineBackend.figure_format = 'retina'
%load_ext autoreload
%autoreload 1
# import libraries:
import sys, os
```

Growth tension(s)

Gravitational lensing

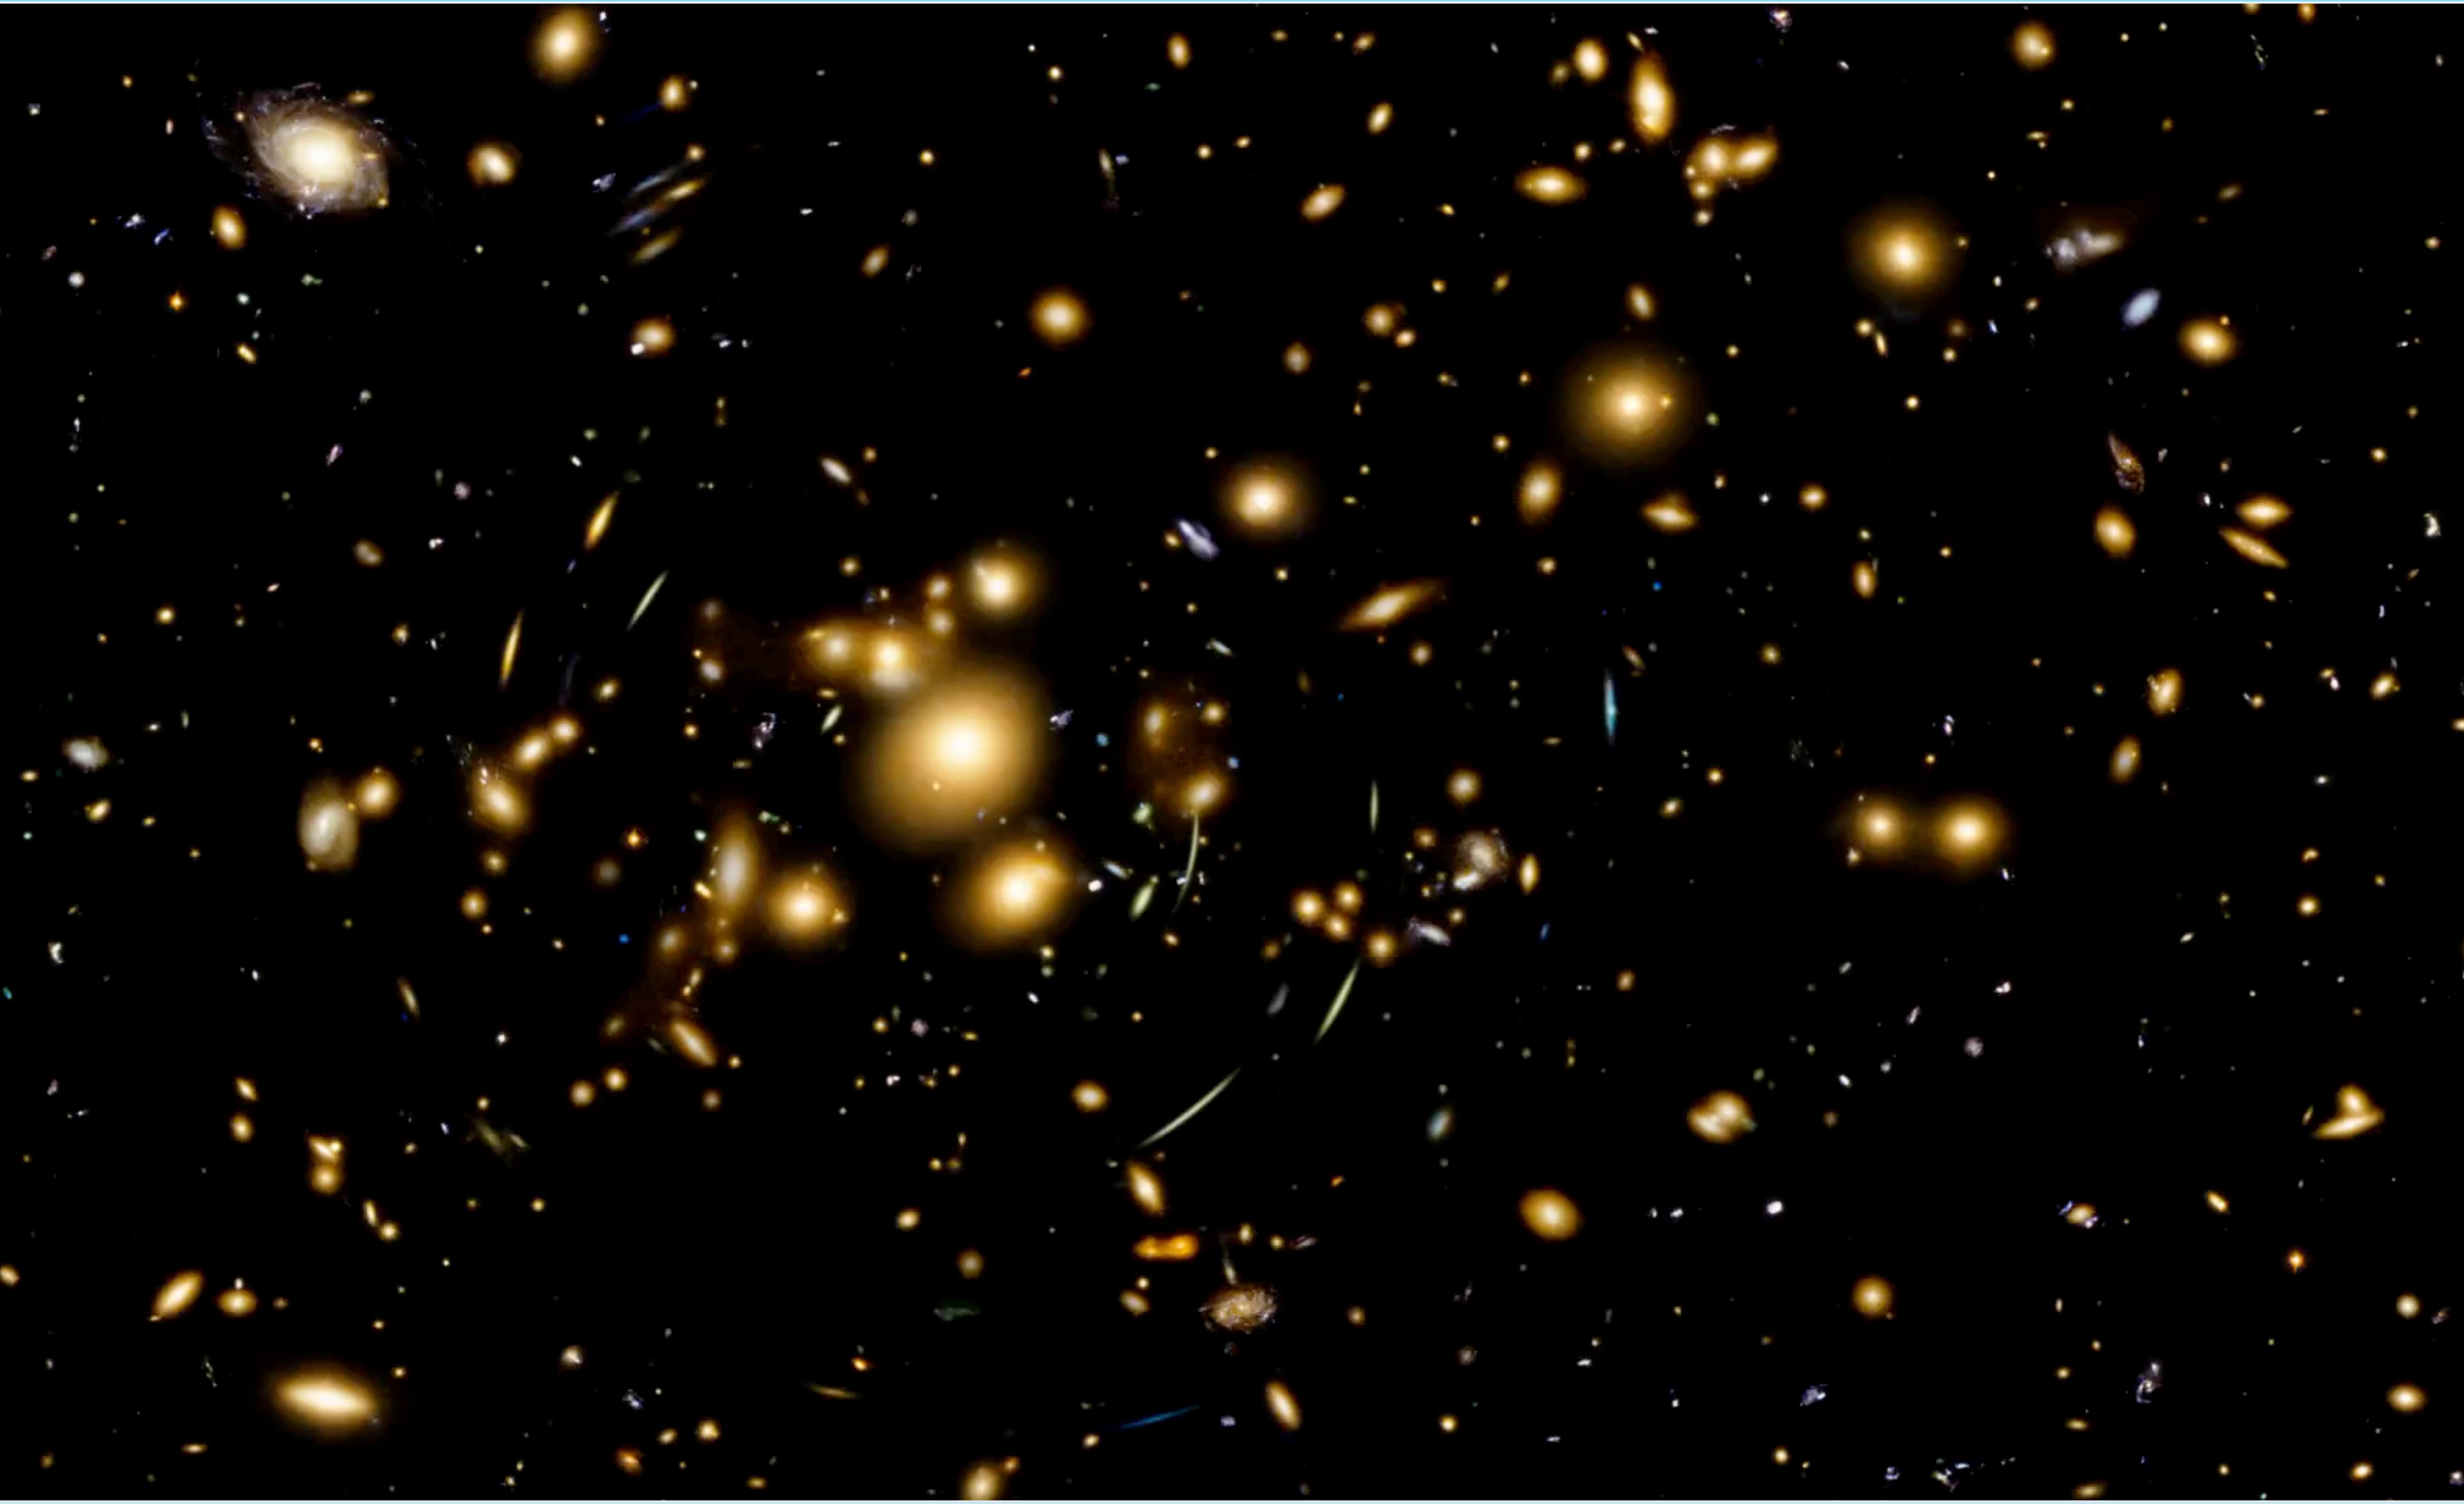

(video credit ESA)

Gravitational lensing

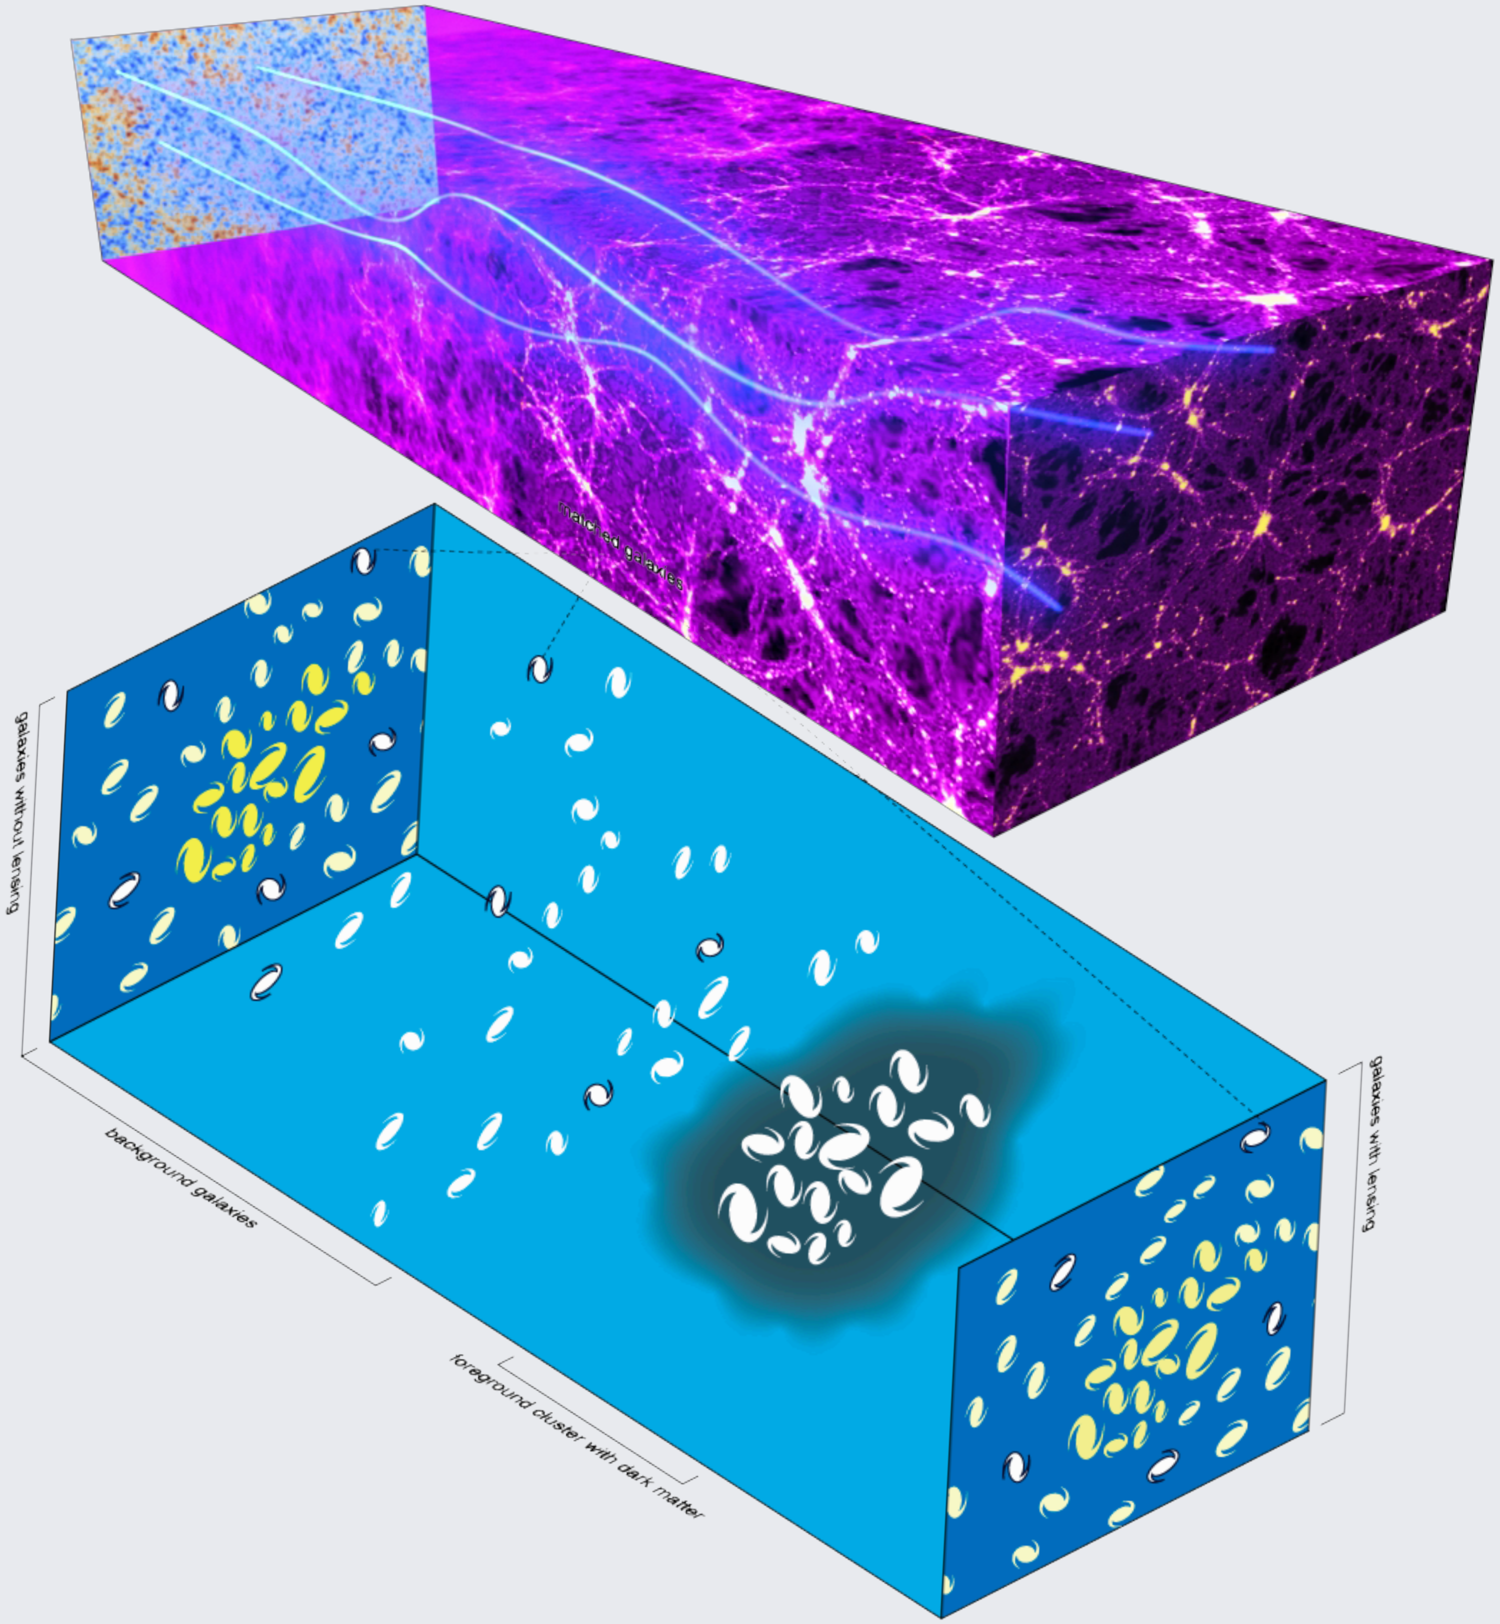

Both the CMB and galaxy shapes are distorted

Weak lensing surveys try to detect the lensing effect from large scale structures in front of other galaxies

Structure's probe 1: galaxies

Dark Energy Survey

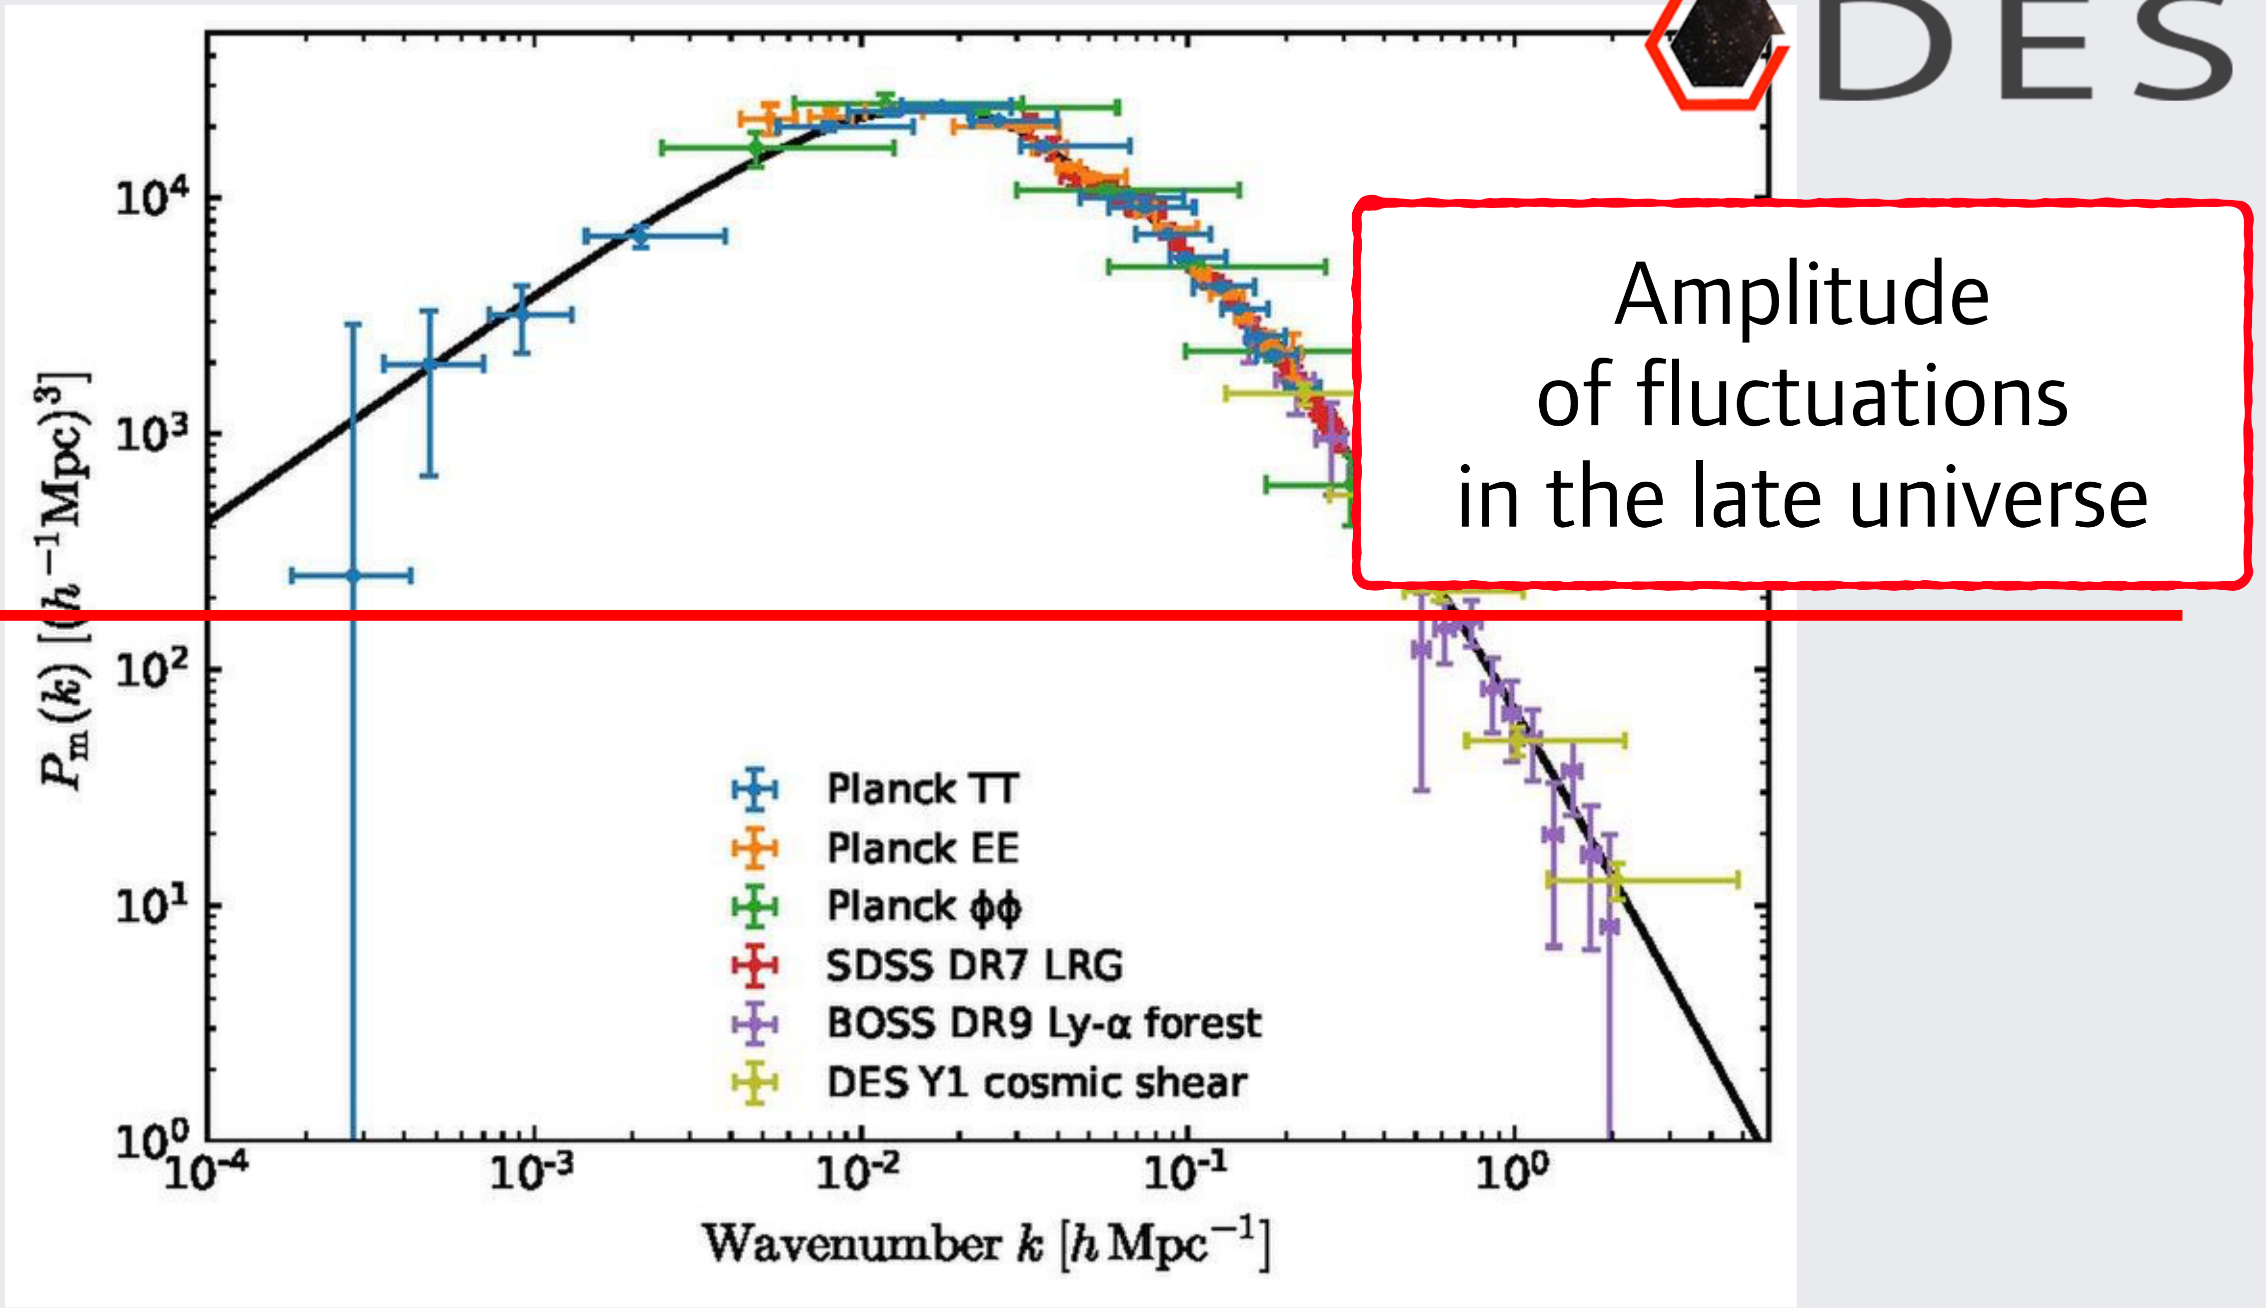

First calibrator: power spectrum of galaxy clustering

Structure's probe 1: galaxies

Physical intuition is easy → amplitude of the signal
Model parameter dependency is complicated

clumpiness of the Universe

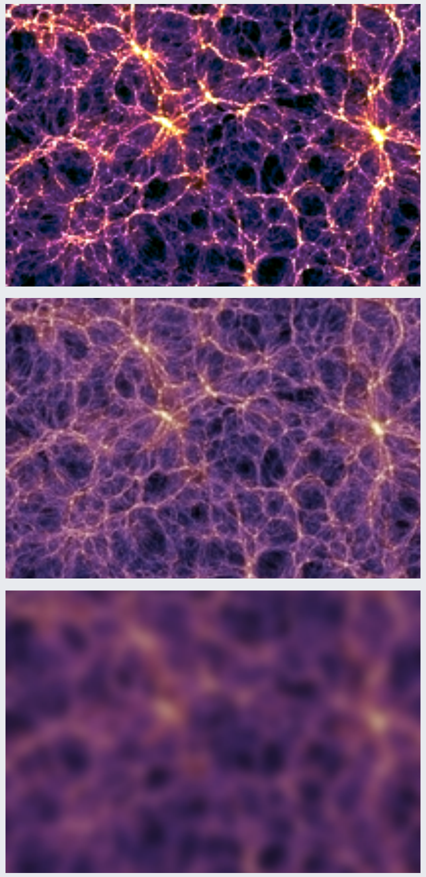

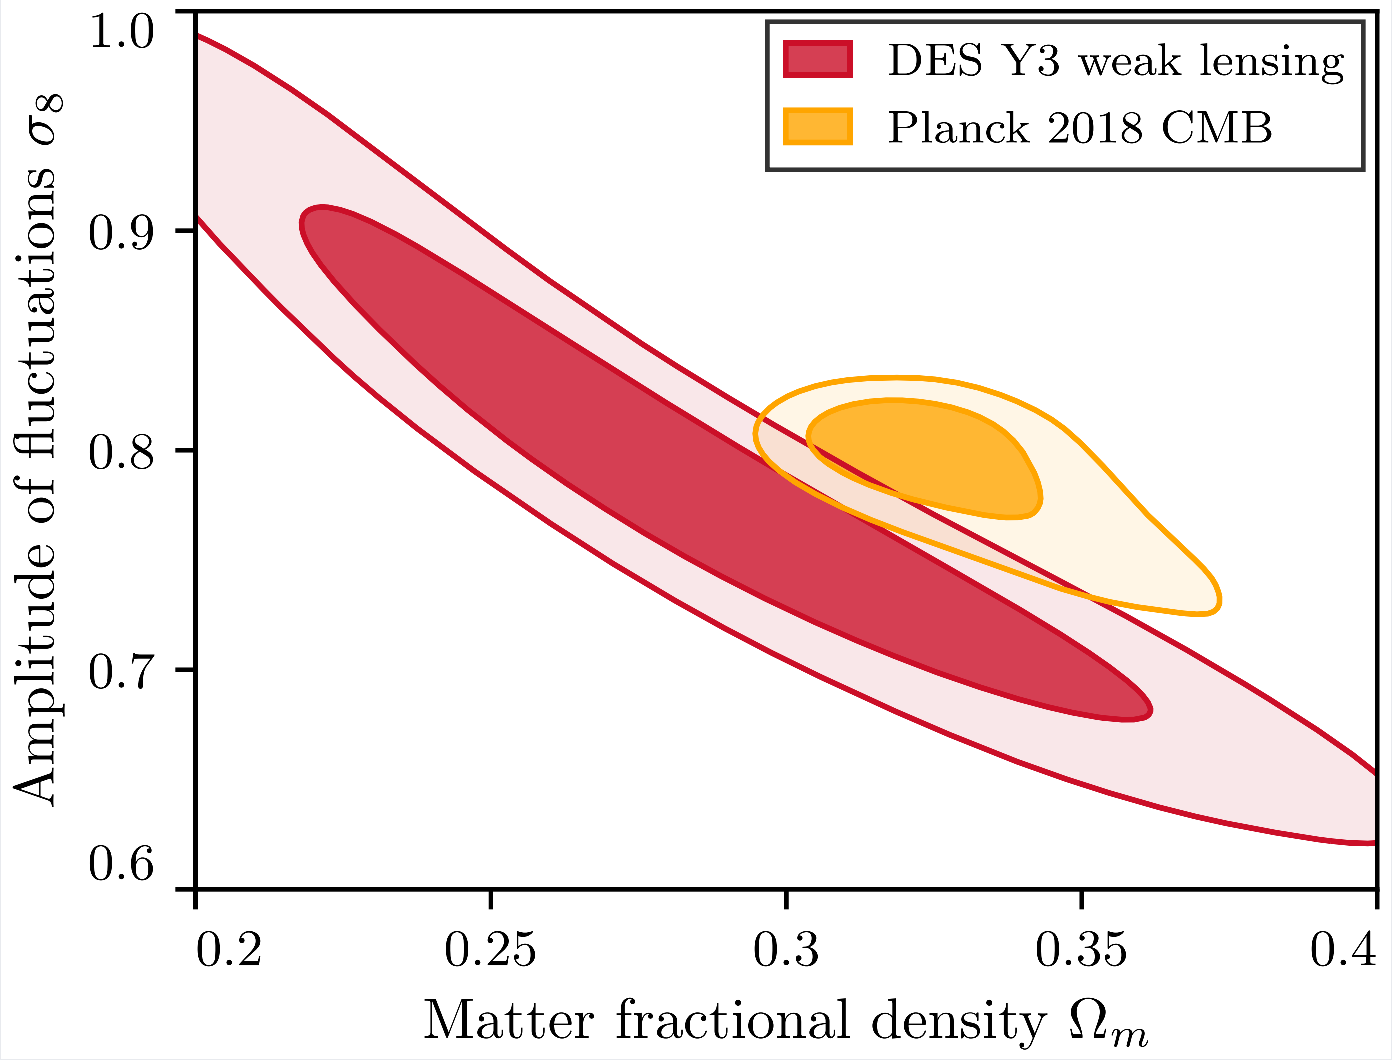

fractional matter abundance
(everything else is Dark Energy)

(see arXiv:2112.05737 for best constrained parameters)

Structure's probe 2: CMB

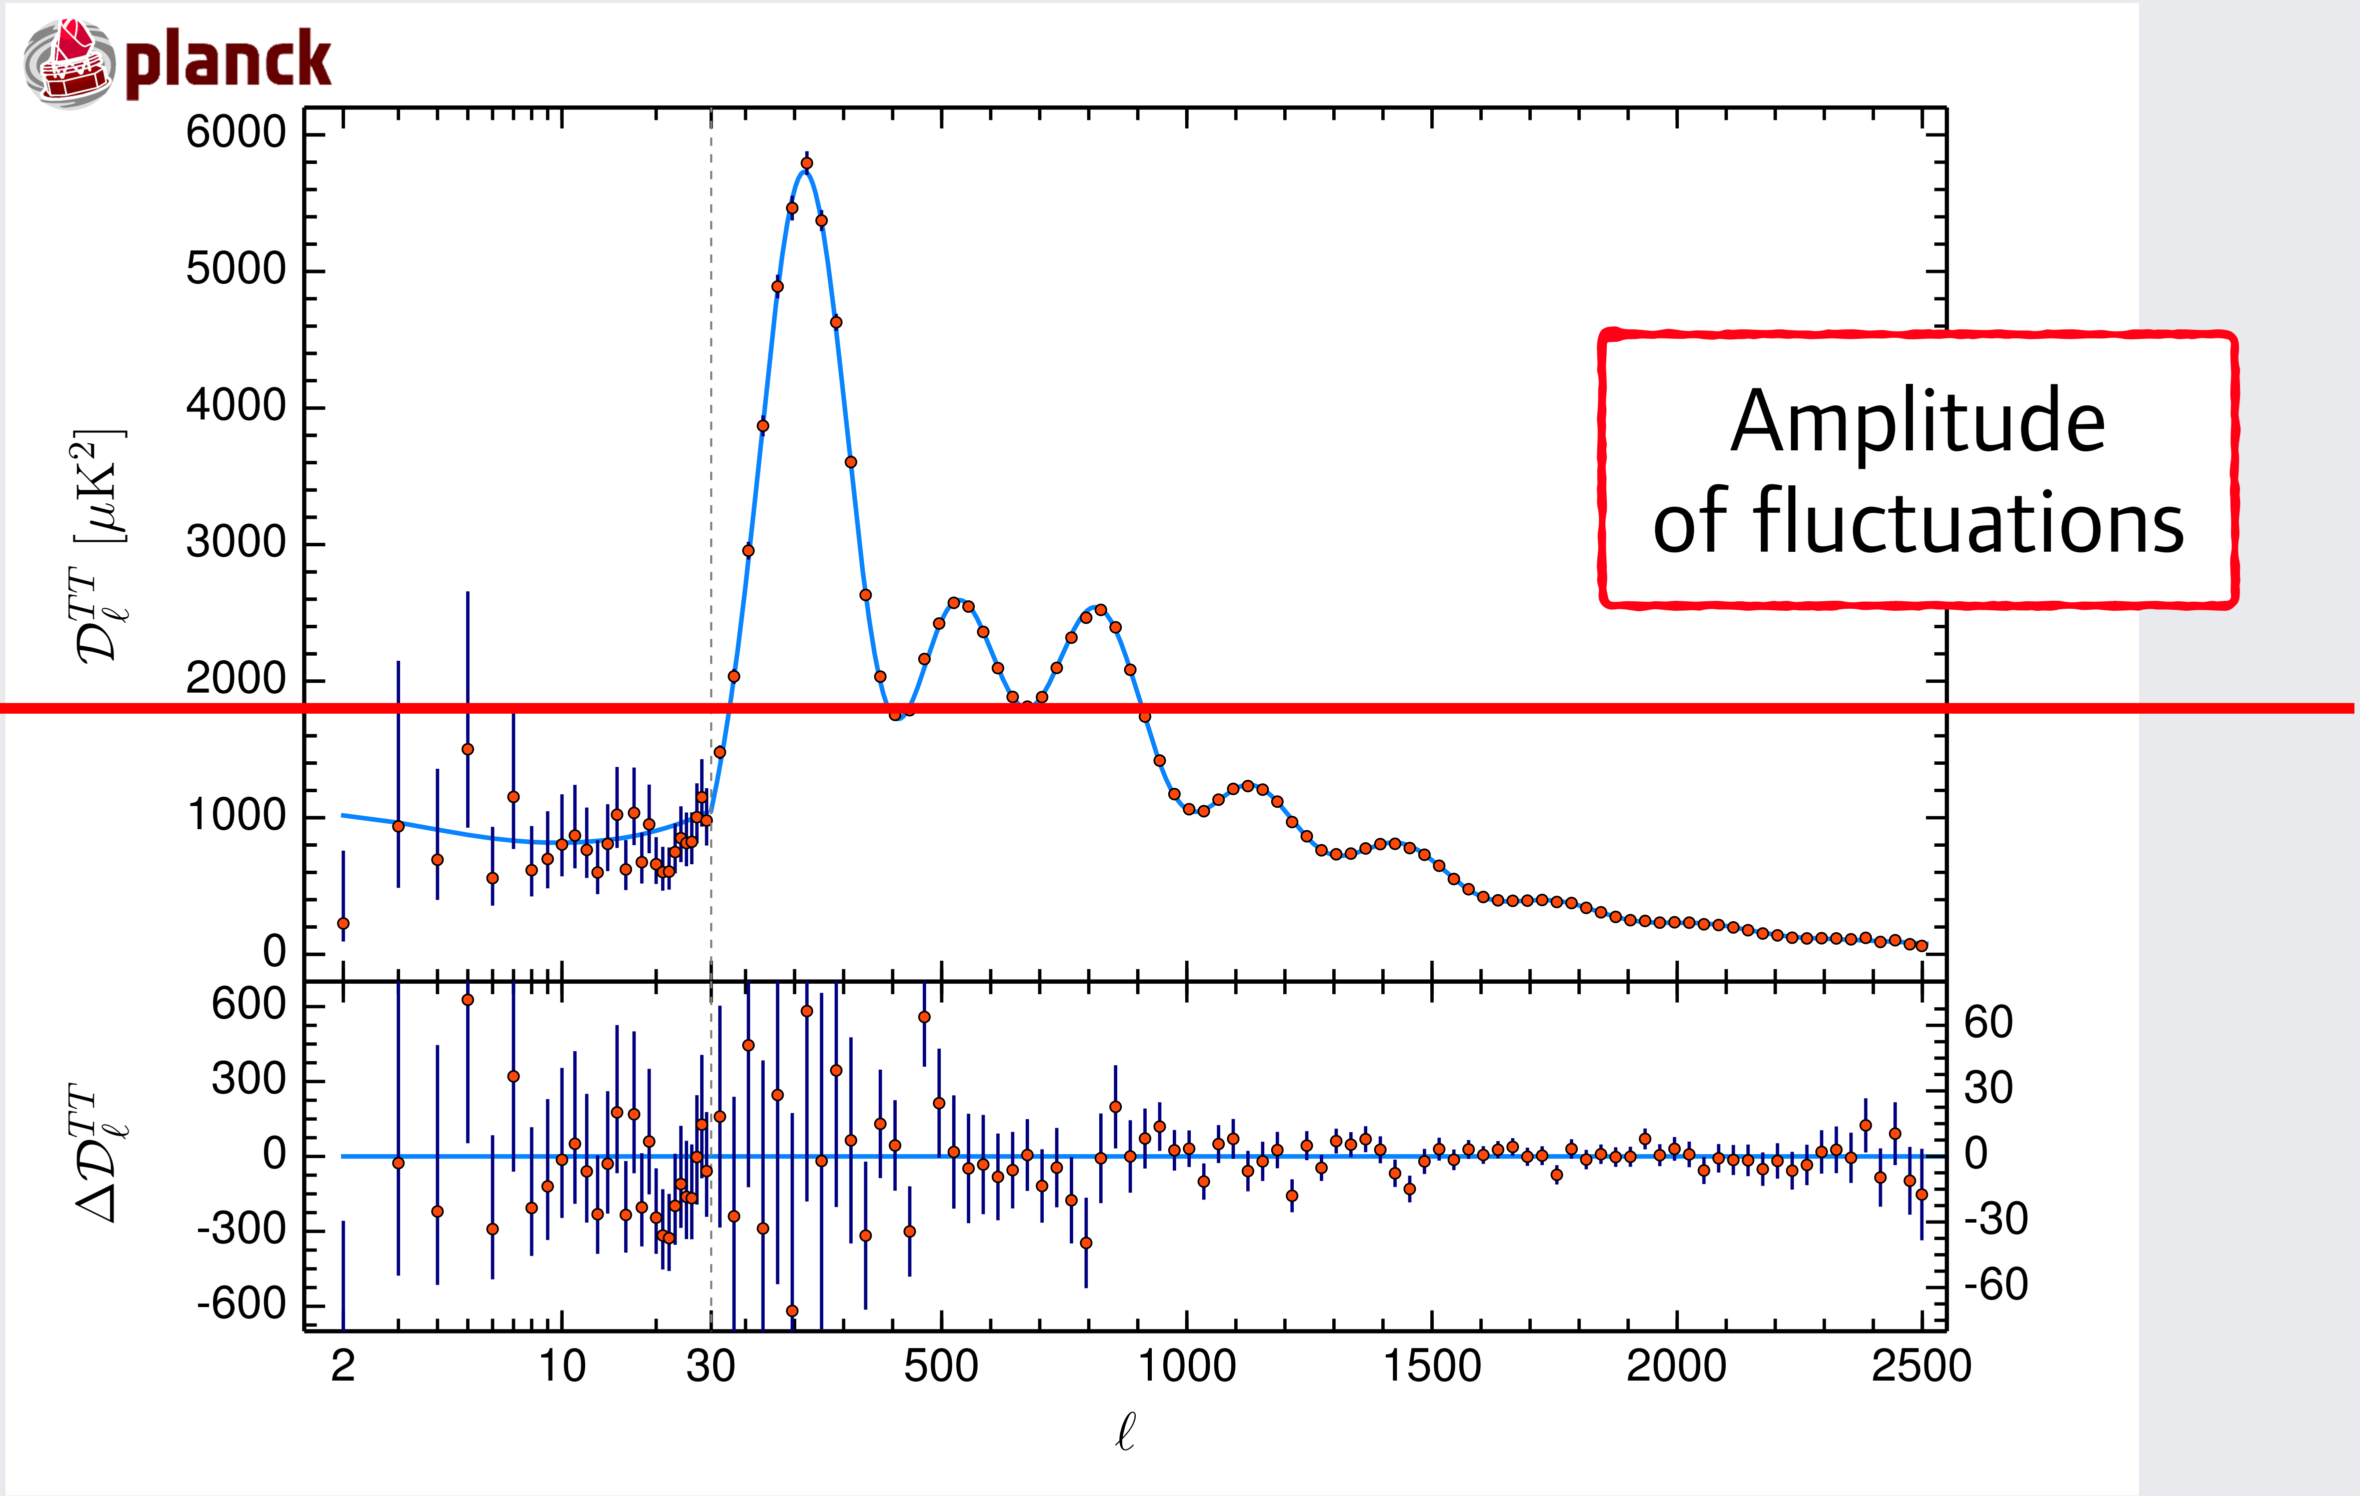

Second calibrator: power spectrum of CMB temperature fluctuations

(Planck 2018 cosmological parameters, arXiv:1807.06209)

Structure's probe 2: CMB

The amplitude of the CMB signal is a composite of the primordial amplitude of fluctuations (A_s) and optical depth from us to recombination (τ) and is $\propto A_s \cdot e^{-2\tau}$

The CMB can self-calibrate in two ways:
1- CMB lensing
2- large scale polarization

Structure's probe 2: CMB

1- lensing: smears the CMB acoustic peaks

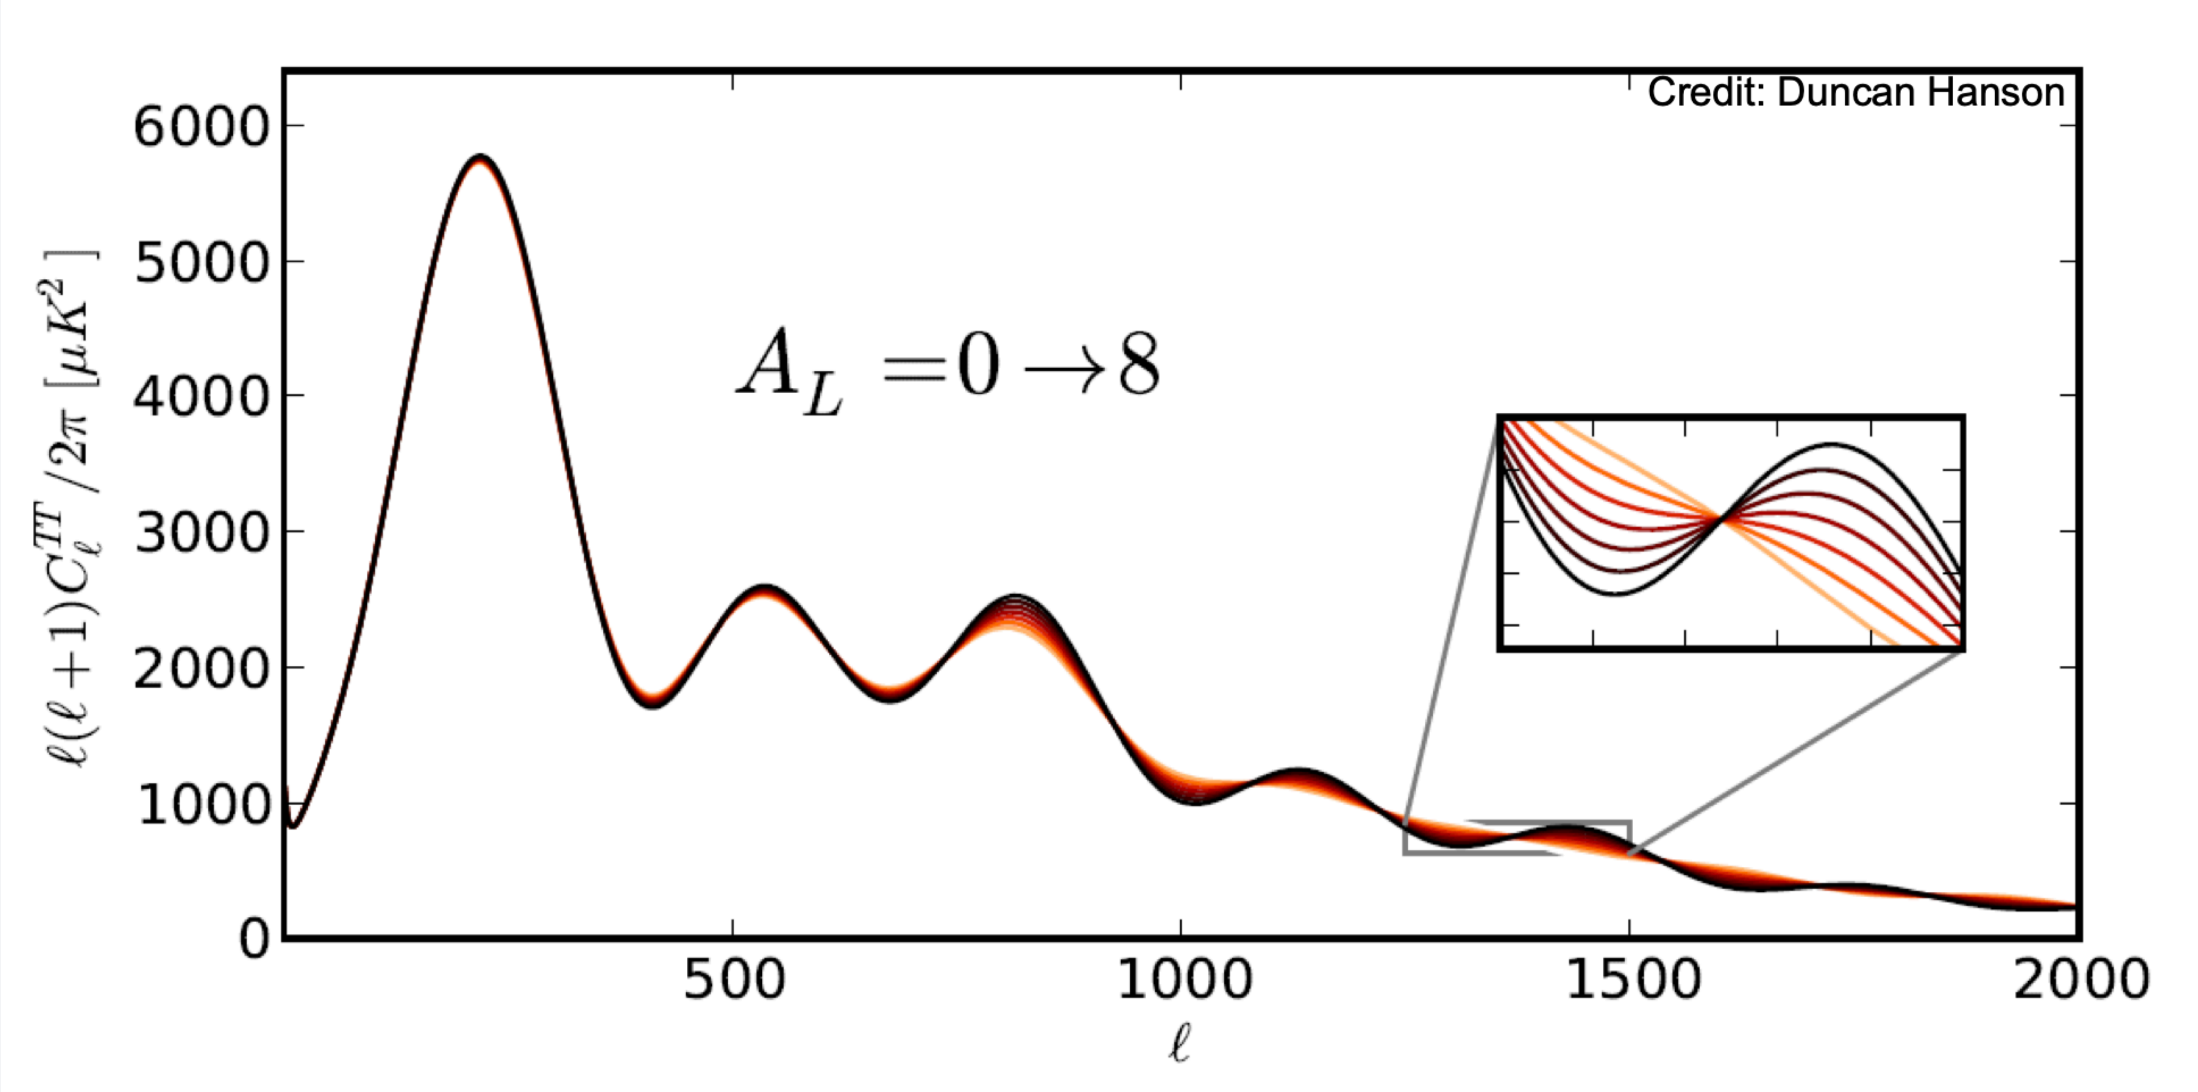

Structure's probe 2: CMB

2- large scale polarization: re-scattering of photons generates polarization anisotropies

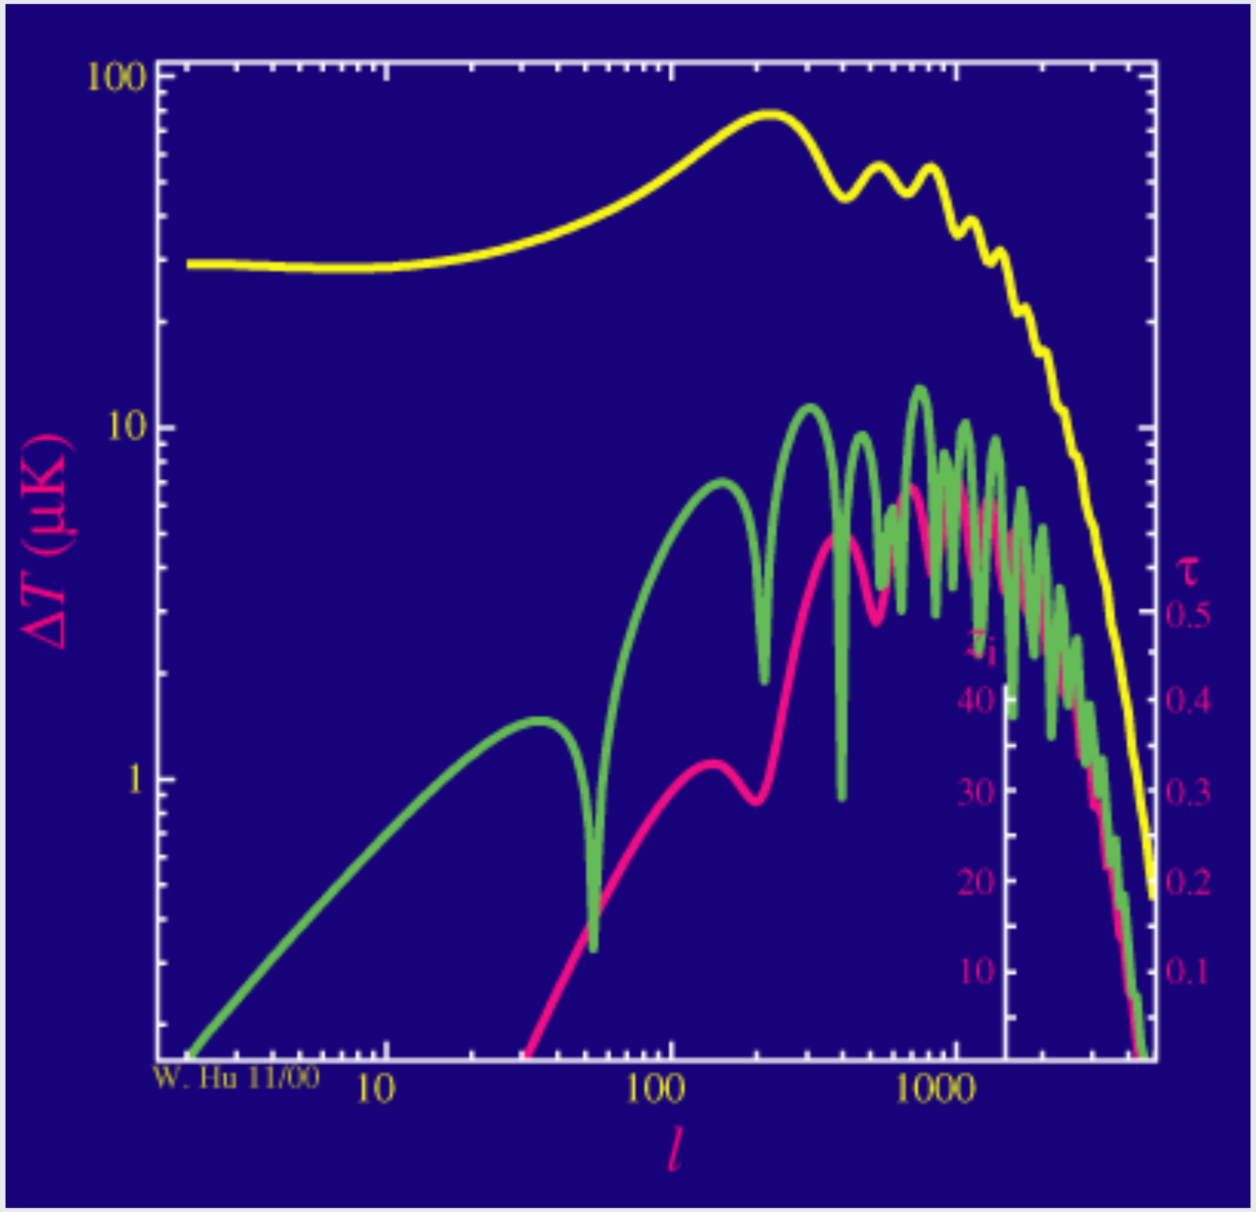

(credit Wayne Hu tutorials)

Exercise with current data sets

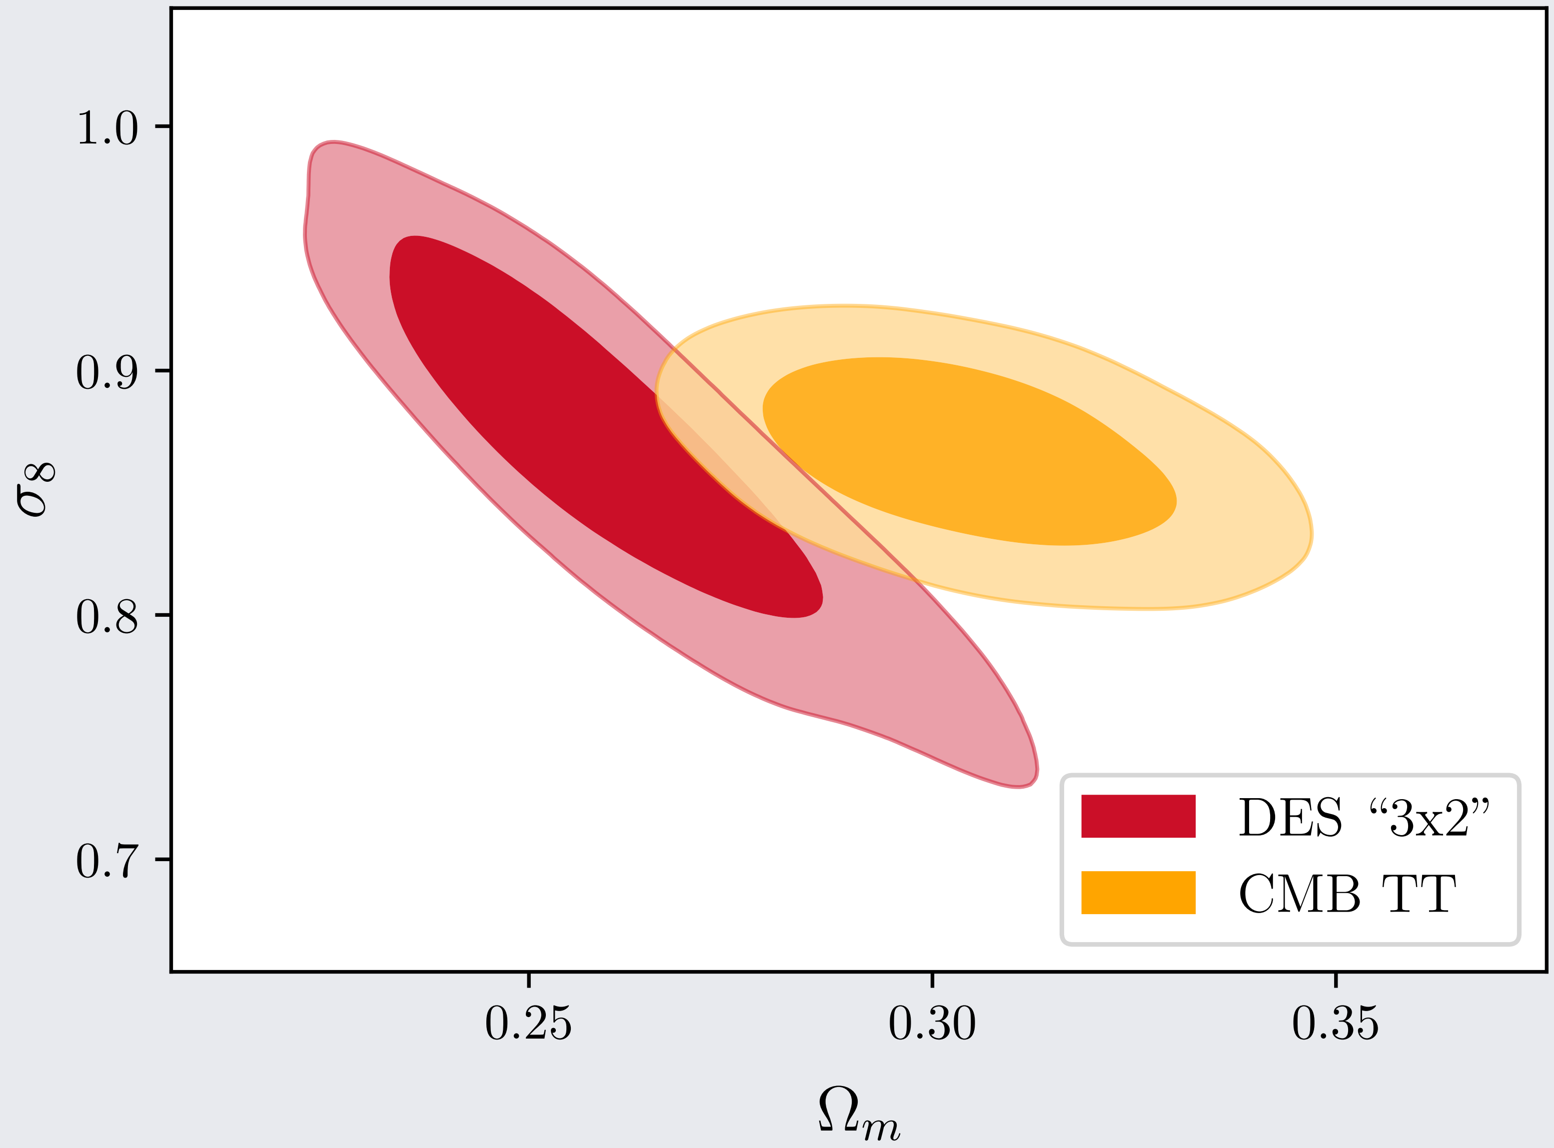

DES linear vs CMB temperature $\Rightarrow P = 0.03\ (2.1\sigma)$
if considering 1d then $\Rightarrow 2.6\sigma \Rightarrow$ **look elsewhere effect**

(based on work in DES)

Growth tension 1: CMB-galaxies

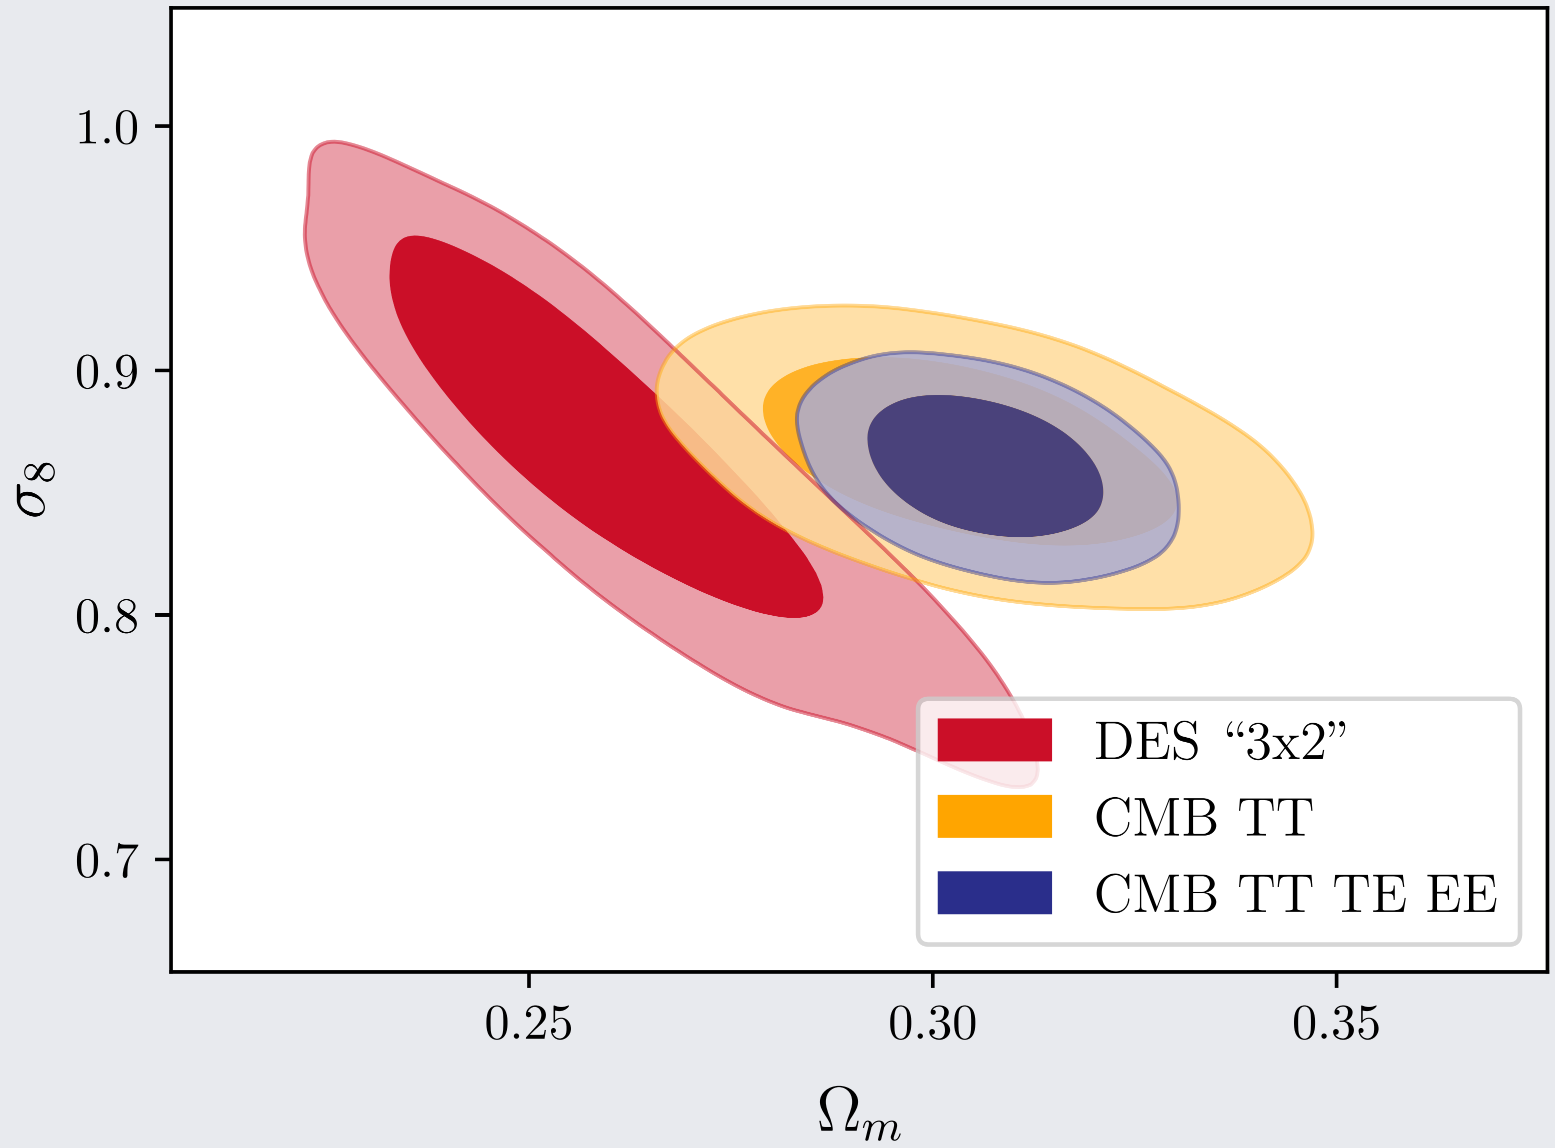

DES linear vs CMB temp + pol $\Rightarrow P = 0.001\ (3.2\sigma)$
Update (temp vs temp + pol) in full agreement

(based on work in DES)

Exercise with current data sets

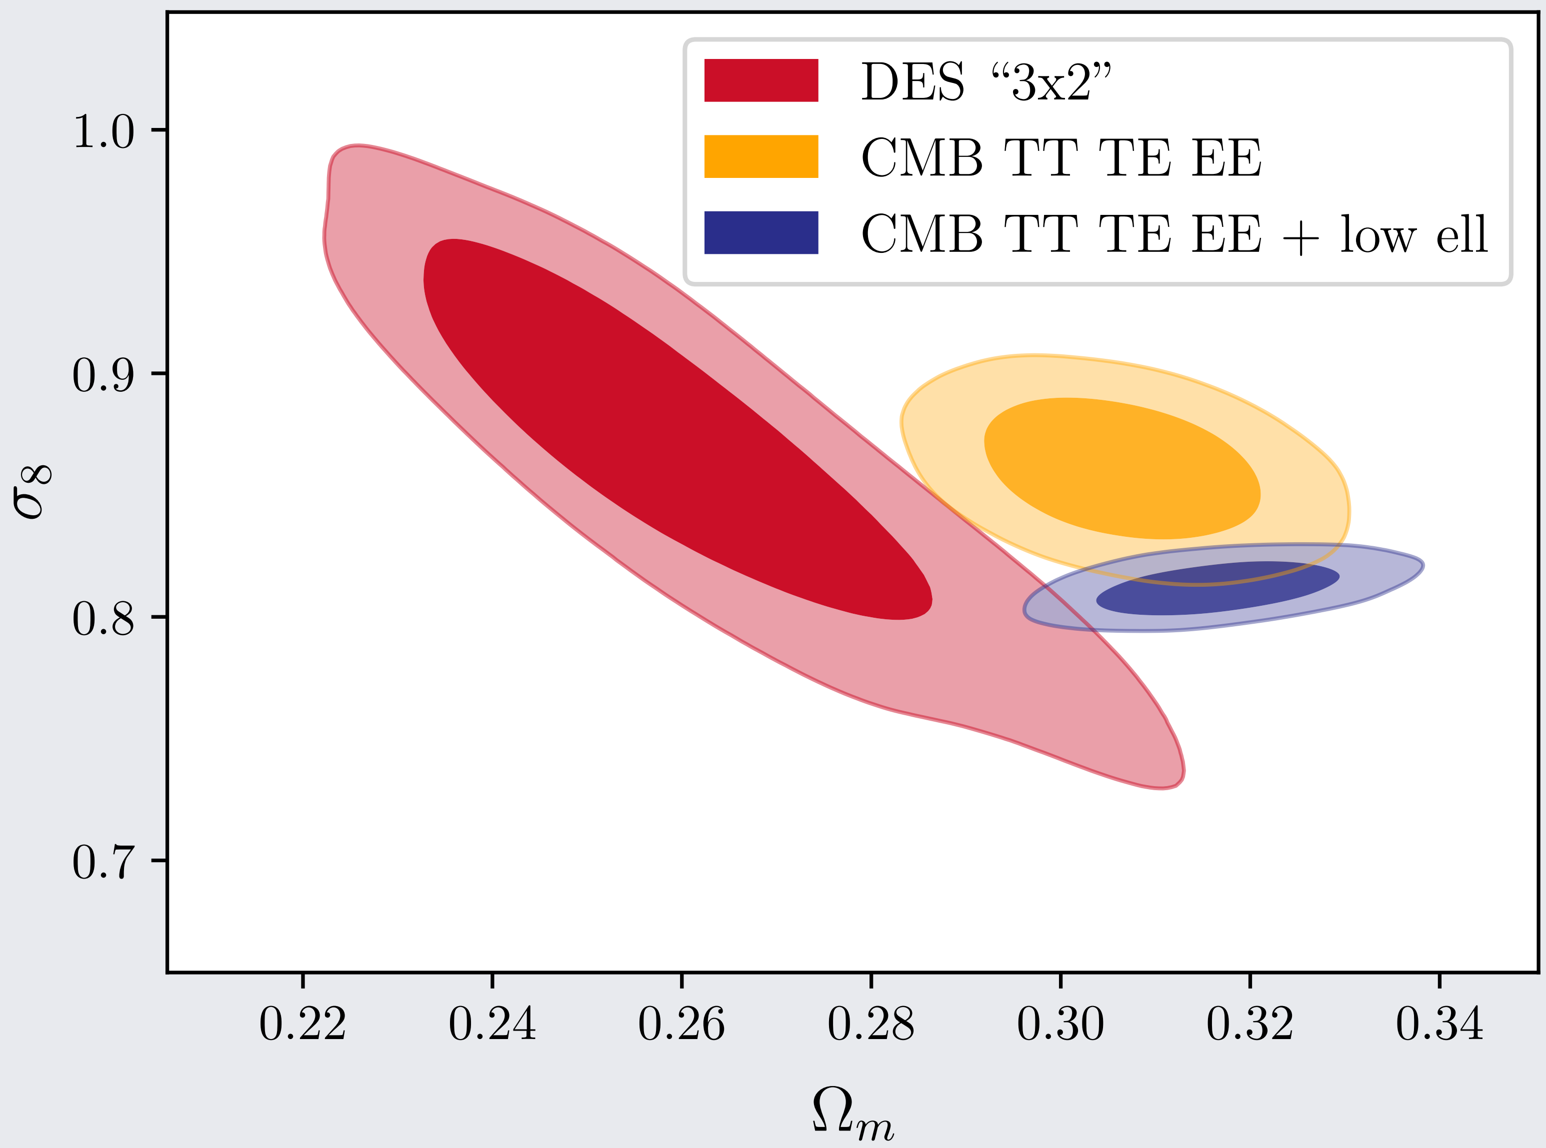

DES vs CMB temp+pol+large scale pol $\Rightarrow P = 0.006\ (2.7\sigma)$

(based on work in DES)

Growth tension 2: CMB lensing

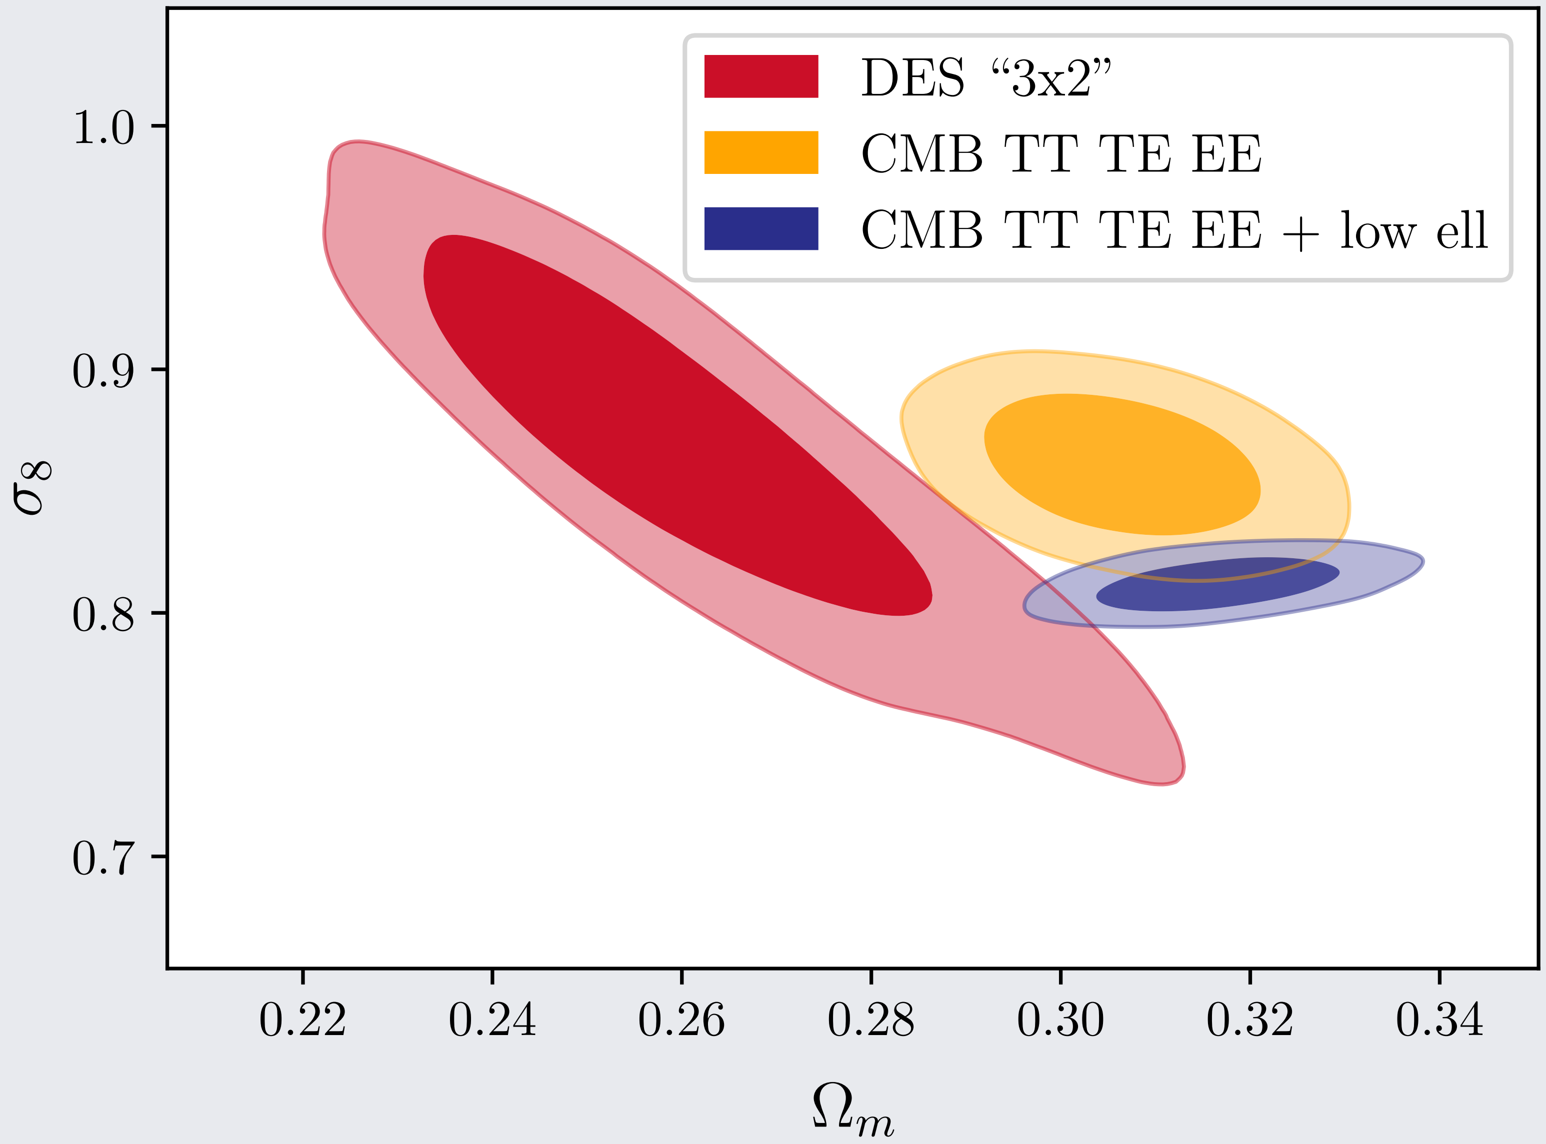

Update adding large scale polarization is in some tension
$\Rightarrow P = 0.003\ (3\sigma)$

(based on work in DES)

Growth tension 2: CMB lensing

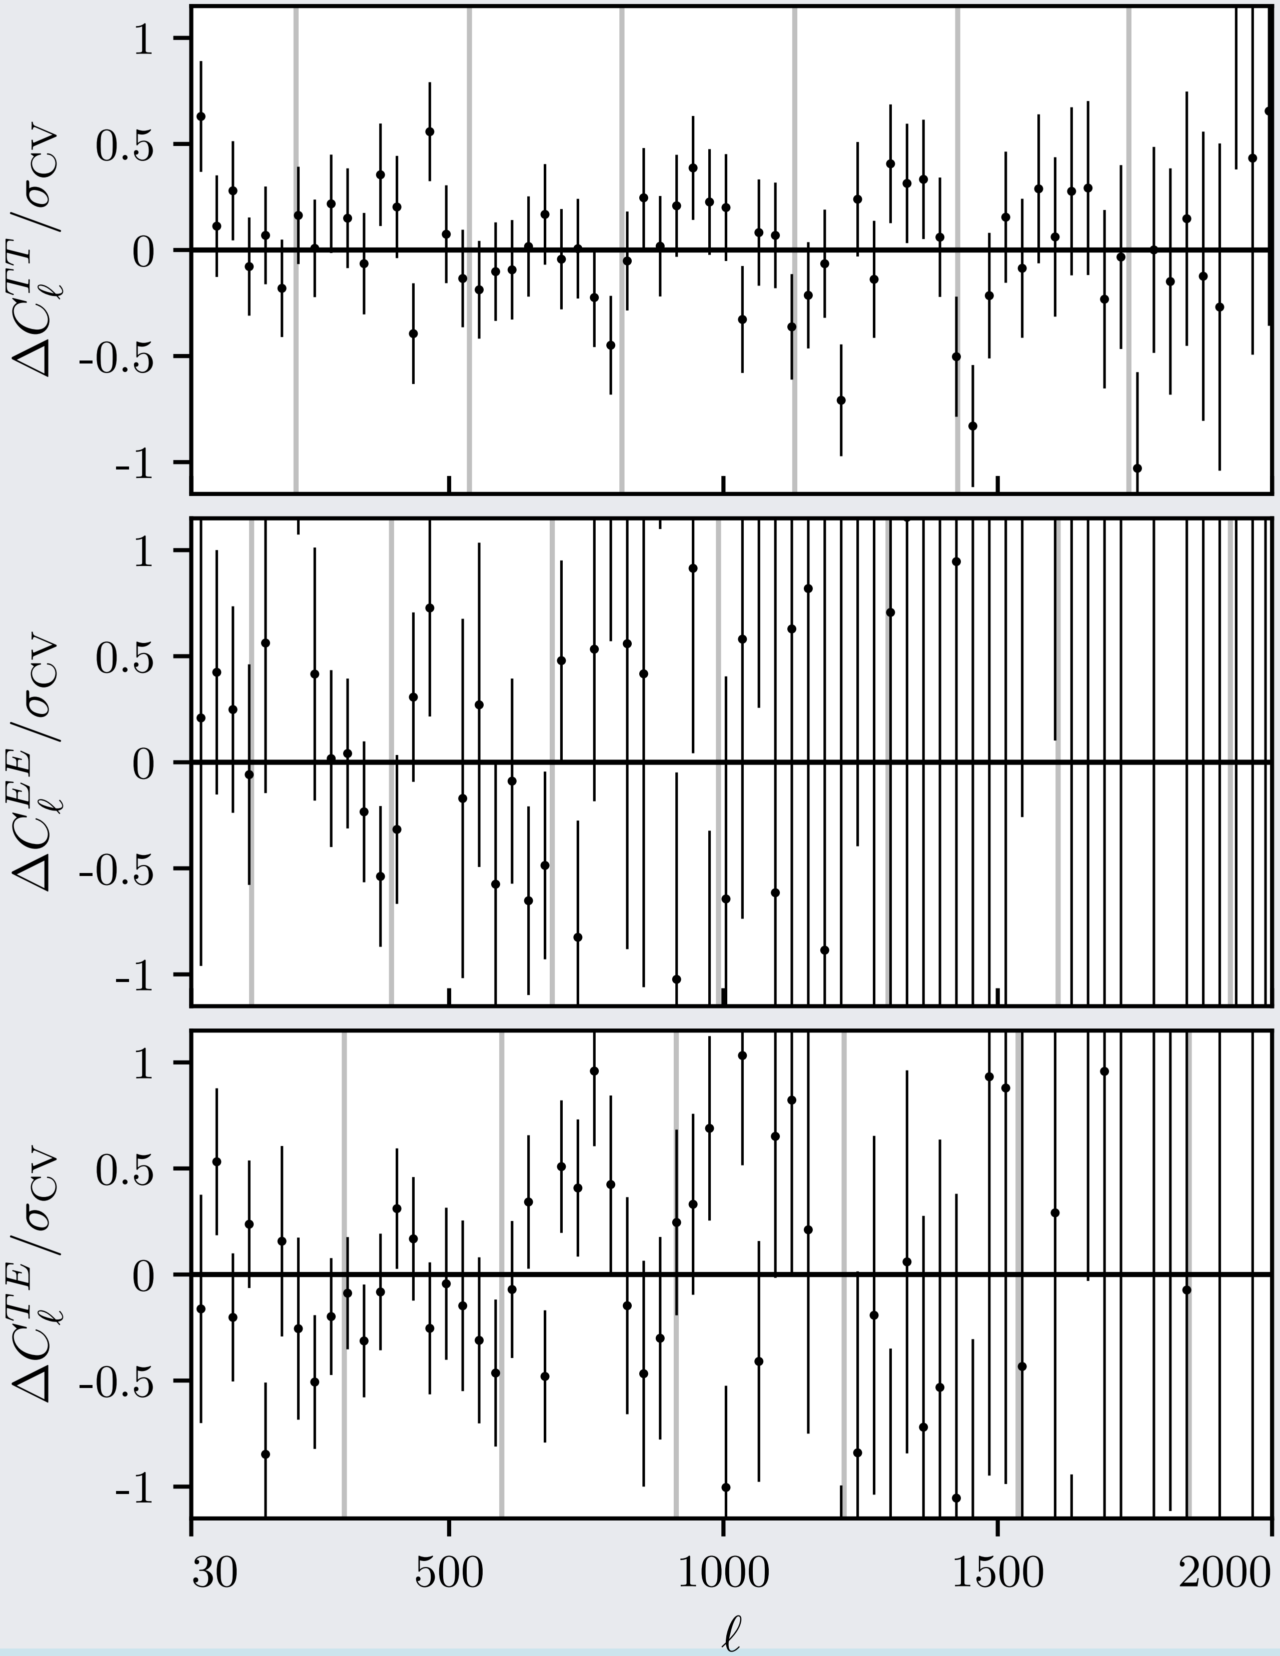

Oscillatory residuals in CMB best fit

Application to growth probes

Best constrained parameters see arXiv:2112.05737

DES 3x2 linear
$$S_8 \equiv \sigma_8(\Omega_m/0.3)^{0.9} = 0.762 \pm 0.023$$

CMB all
$$S_8 = 0.852 \pm 0.025$$

Both 3% measurements

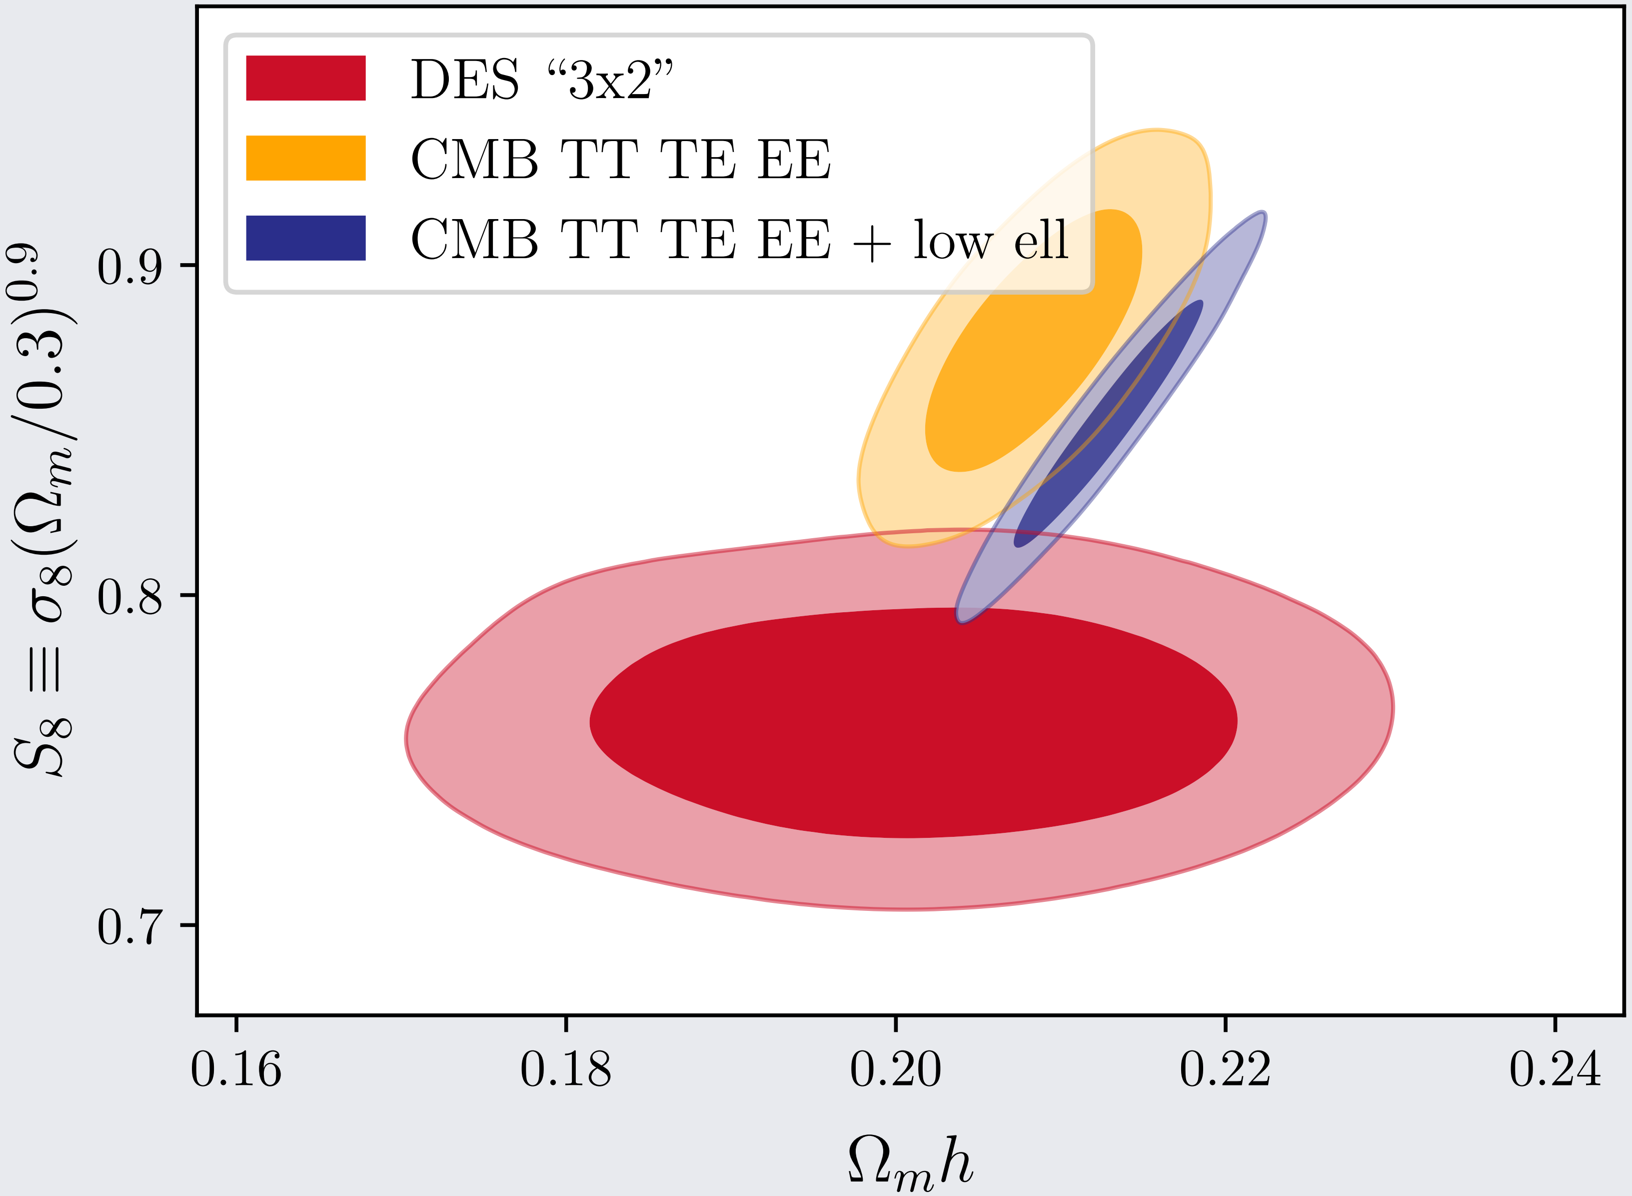

(based on work in DES)

Conclusions

Outlook

* **Concordance in cosmology**: a puzzle with multiple pieces

* **Three main (calibration) tensions on the way to 1%**:

    * Hubble constant

    * Growth of structures

    * Lensing of the CMB

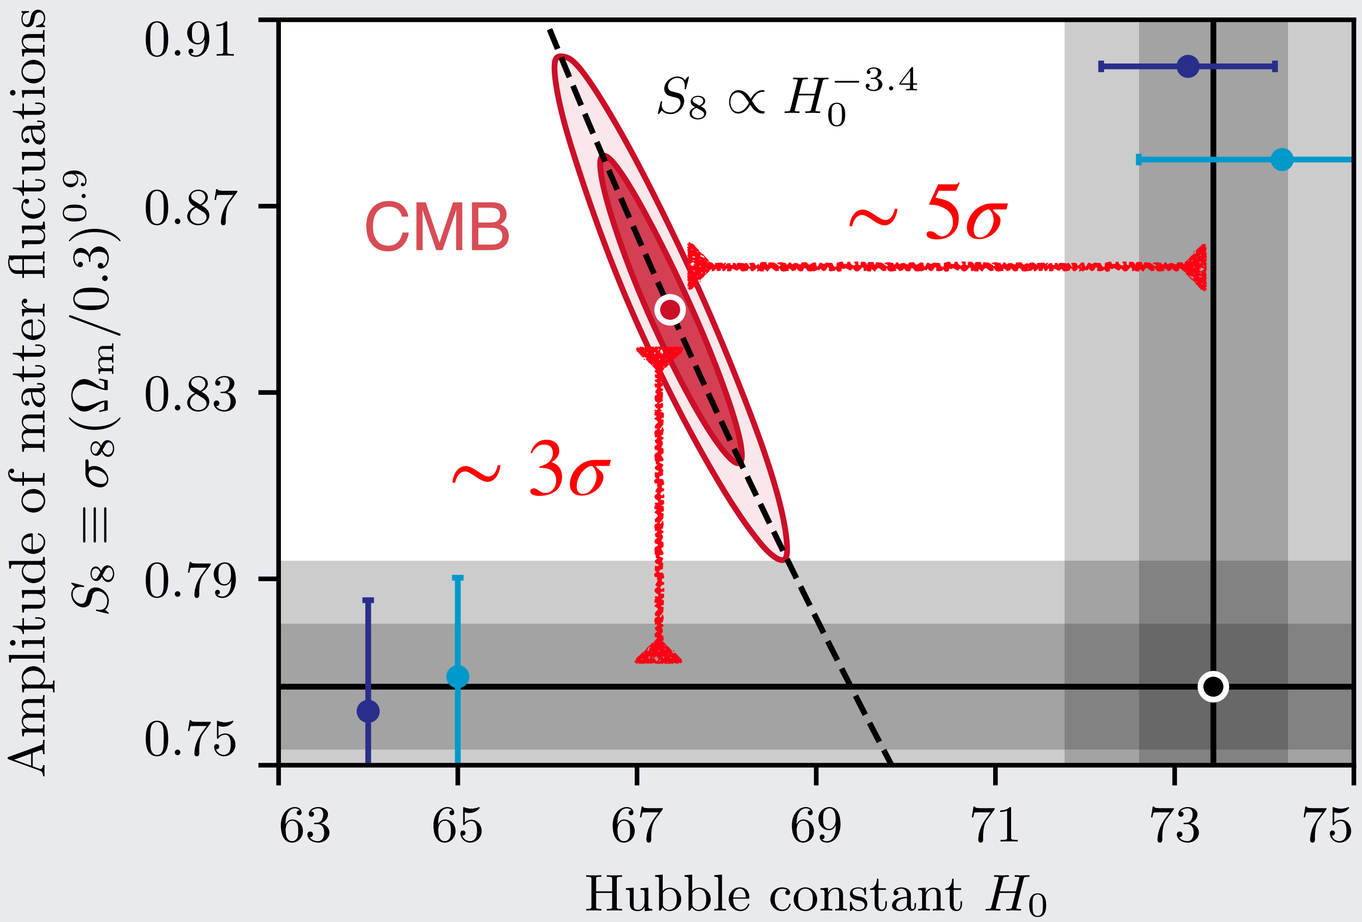

* They all seem to involve hi vs low redshift…

Join the search for tensions

(whatever field you work in)

```
~ pip install tensiometer
```

Contributions are very welcome!

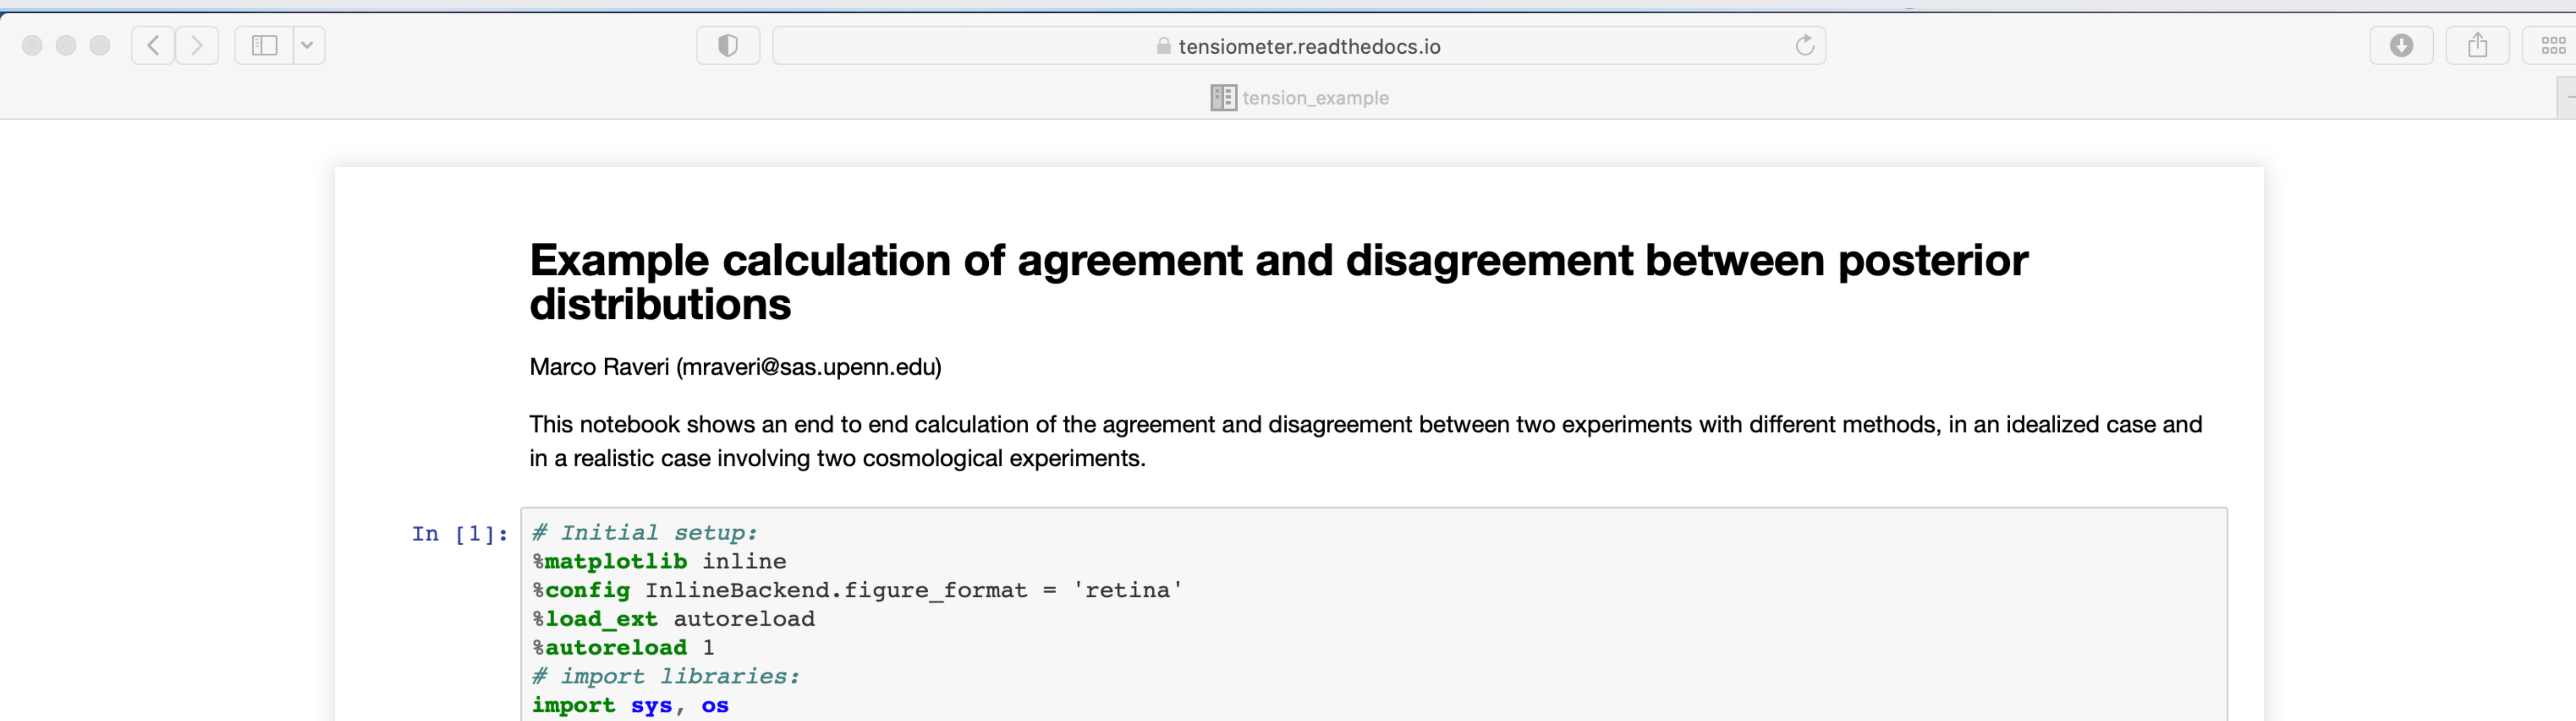

Example calculation of agreement and disagreement between posterior distributions

Marco Raveri (mraveri@sas.upenn.edu)

This notebook shows an end to end calculation of the agreement and disagreement between two experiments with different methods, in an idealized case and in a realistic case involving two cosmological experiments.

In [1]:
```
# Initial setup:
%matplotlib inline
%config InlineBackend.figure_format = 'retina'
%load_ext autoreload
%autoreload 1
# import libraries:
import sys, os
```