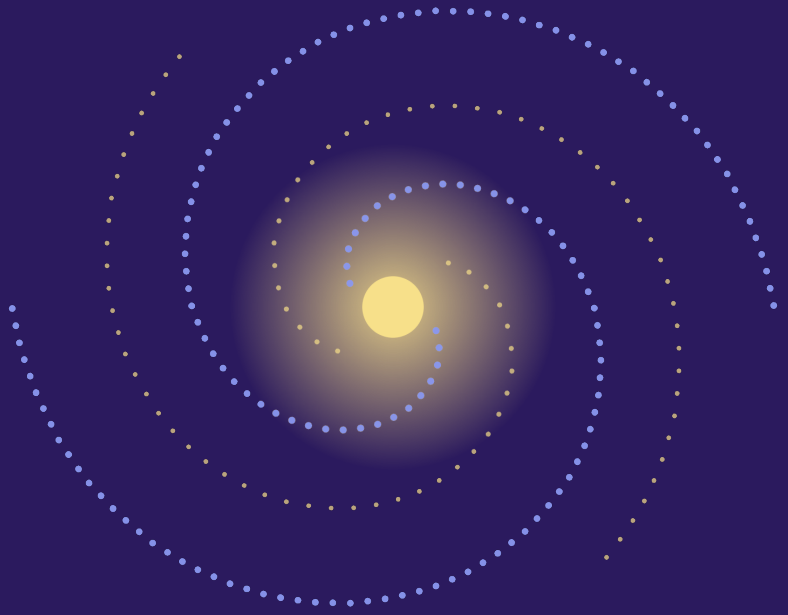
LBNL

No Mode Left Behind: combining experiments to measure inflation



Sofia Chiarenza · October 13 2026

UNIVERSITY OF WATERLOO

PERIMETER INSTITUTE

OVERVIEW

Outline

01 Motivation

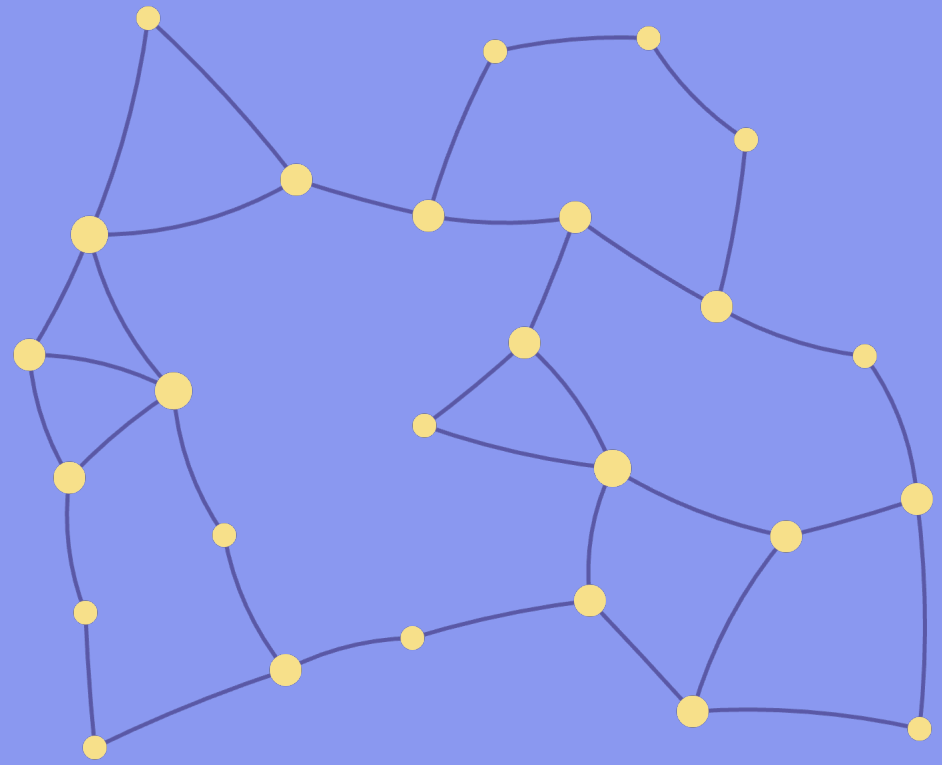
02 Blast

03 PNG constraints from DESIxCMB lensing

04 Future Directions

05 Conclusions

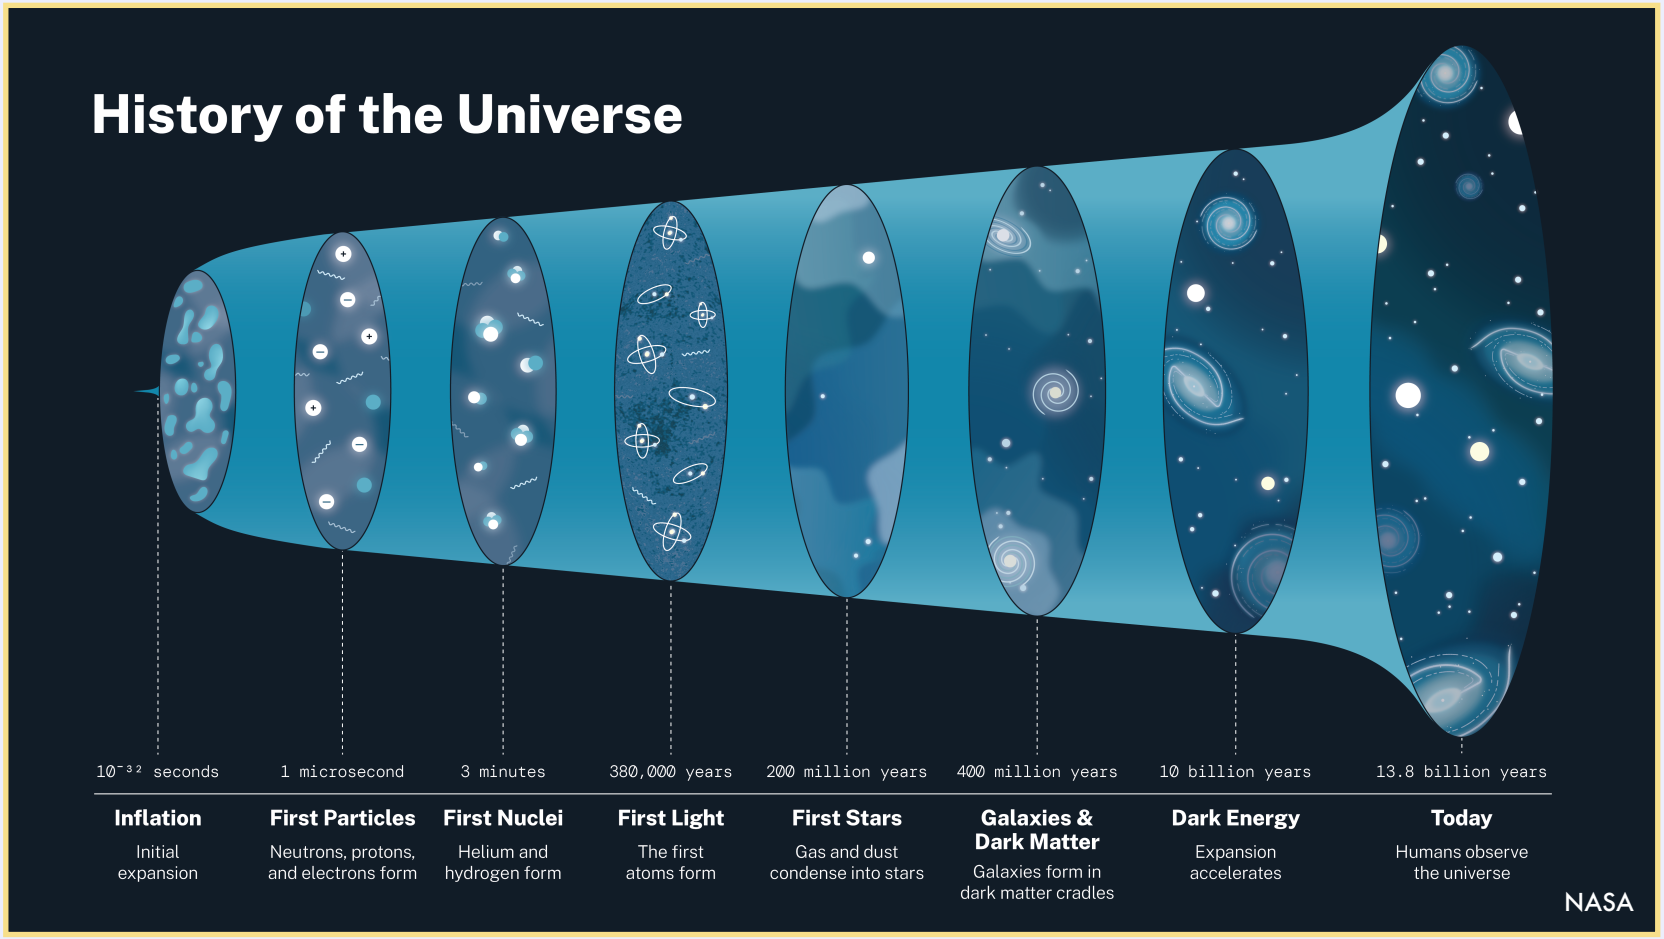
SECTION



01 Motivation

MOTIVATION

Inflation



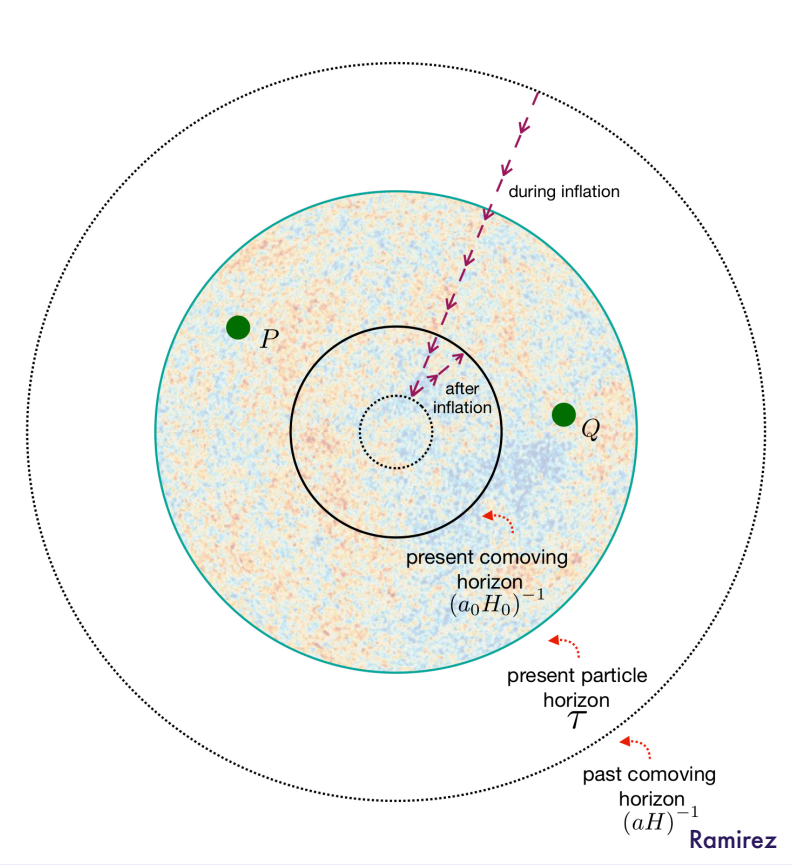
Why inflation?

- Horizon problem: causally disconnected regions with the same CMB temperature
- Flatness problem: why $\Omega_k \approx 0$?
- Initial conditions: what seeded structure formation?

MOTIVATION

Inflation

Quantum fluctuations generated during inflation are stretched to cosmological scales



Simplest model: an almost scale-invariant spectrum

$$\Delta^2\zeta(k) = k^3 P_\zeta(k) \approx A_s \left(\frac{k}{k_\star}\right)^{n_s - 1}$$

Planck 2018

$$A_s = (2.101 \pm 0.034) \times 10^{-9}$$

$$n_s = 0.9649 \pm 0.0042$$

MOTIVATION

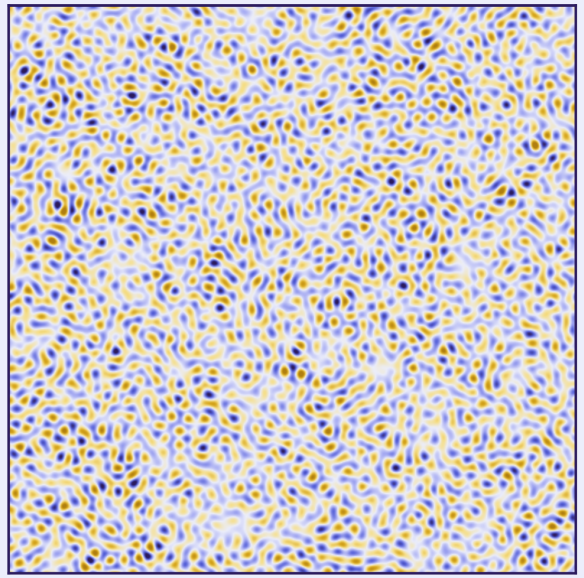
Inflation

The power spectrum completely describe a Gaussian field

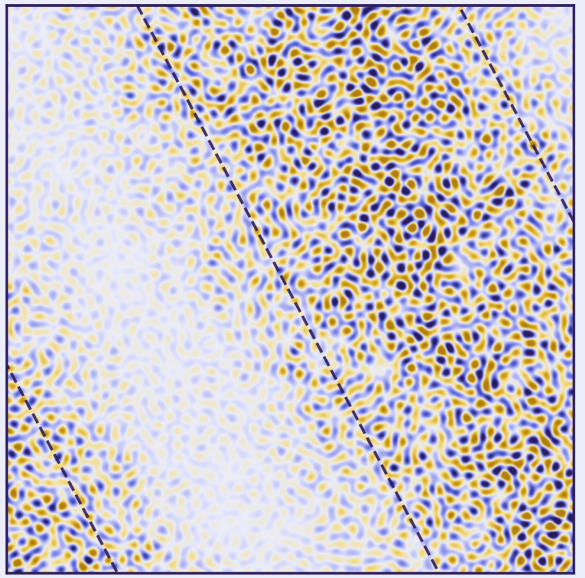
But is there more?

$$\Phi_{NG} := \Phi_G + f_{NL}(\Phi_G^2 - \langle \Phi_G \rangle^2)$$

Different inflation models leave different non-Gaussian signatures



Gaussian Field: $f_{NL} = 0$

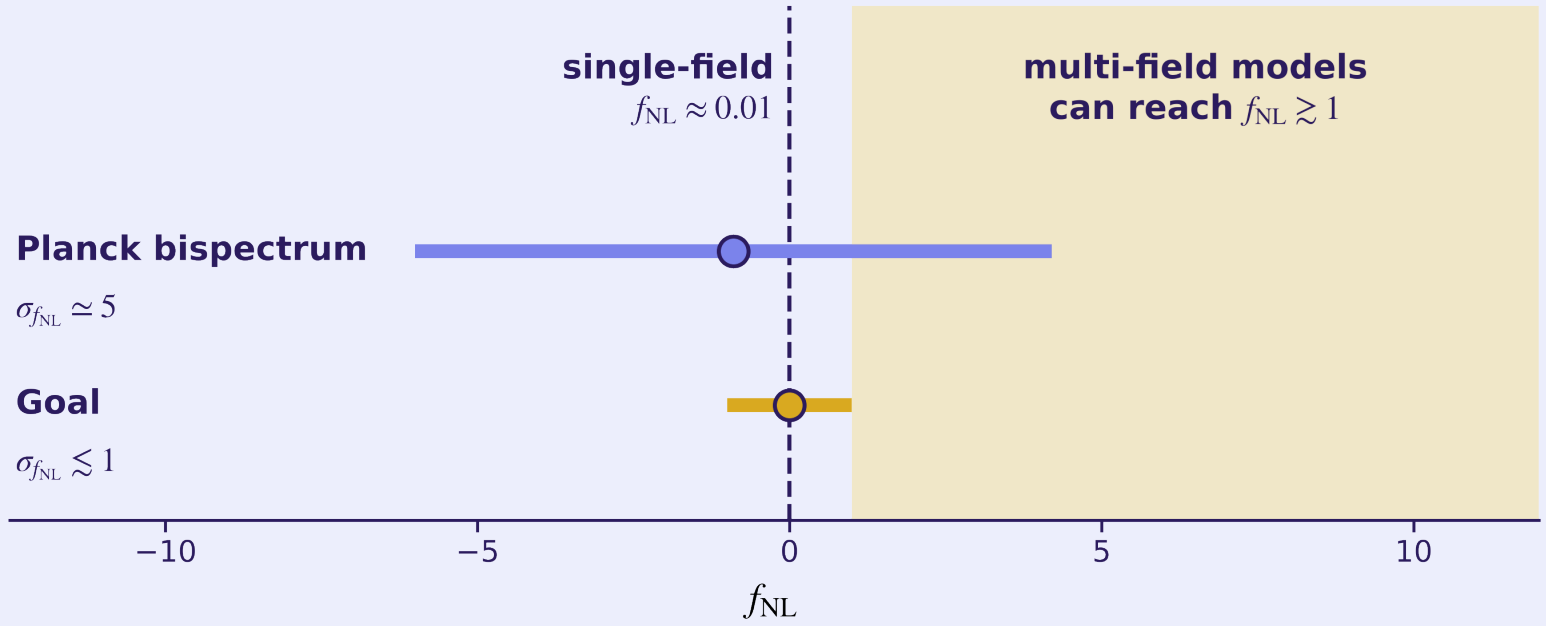


Non Gaussianity: $f_{NL} > 0$

MOTIVATION

How to measure inflation?

The current best measurement is $f_{\mathrm{NL}} = -0.9 \pm 5.1$ from **Planck bispectrum**



The goal is $\sigma_{f_{\mathrm{NL}}} \leq 1$

<u>But</u> CMB measurements are <u>cosmic variance limited</u>

MOTIVATION

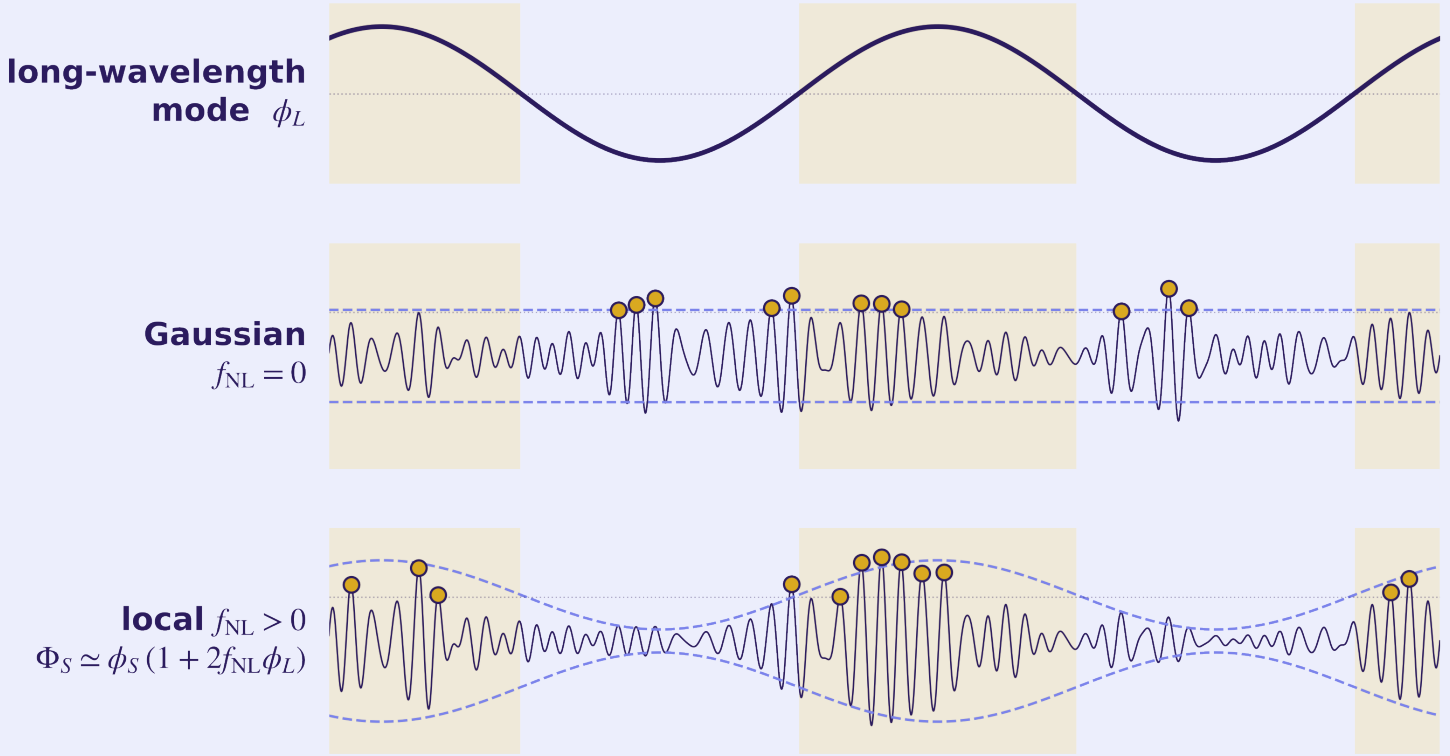
Measure PNG from galaxies

Dalal et al. (2008): Primordial non-Gaussianities, parametrized by $f_{\rm NL}$, leave a signature on the clustering of galaxies

$$\Phi_{\rm NG} := \Phi_{\rm G} + f_{\rm NL}(\Phi_{\rm G}^2 - \langle \Phi_{\rm G} \rangle^2)$$

$$\delta_{\rm NG} = \delta_{\rm G}(1 + 2 f_{\rm NL} \phi_G)$$

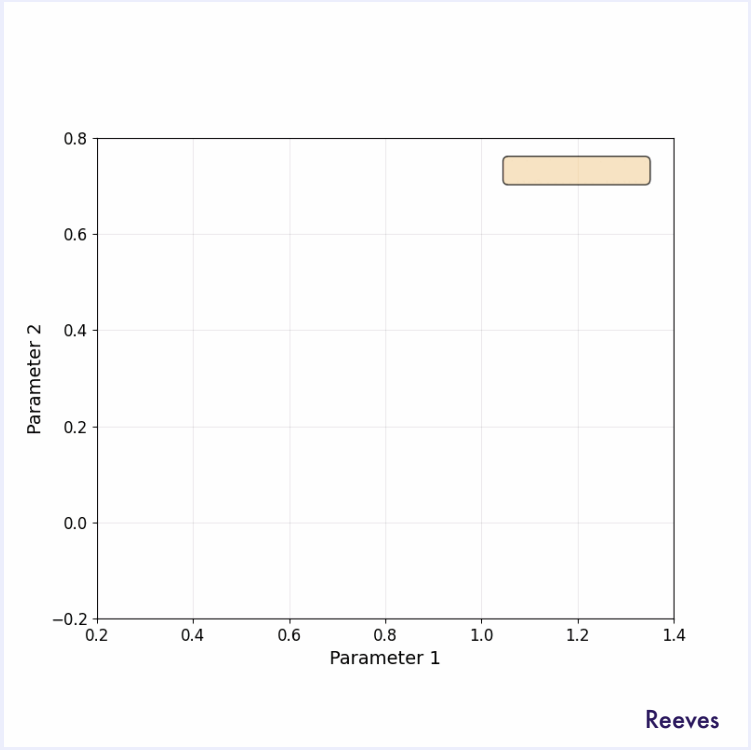
$$\Delta b(k,z) = \frac{b_\Phi(z) f_{\rm NL}}{T_m(k,z)} \propto \boxed{\frac{1}{k^2}}$$



MOTIVATION

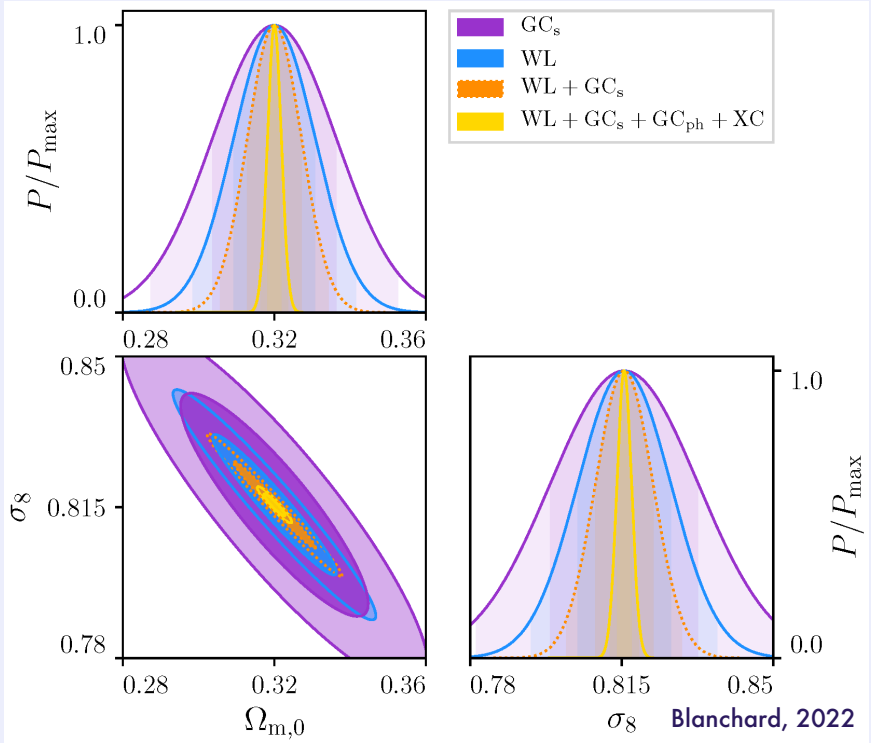
The strategy: multi-probe cosmology

- Identify and mitigate systematics



Reeves

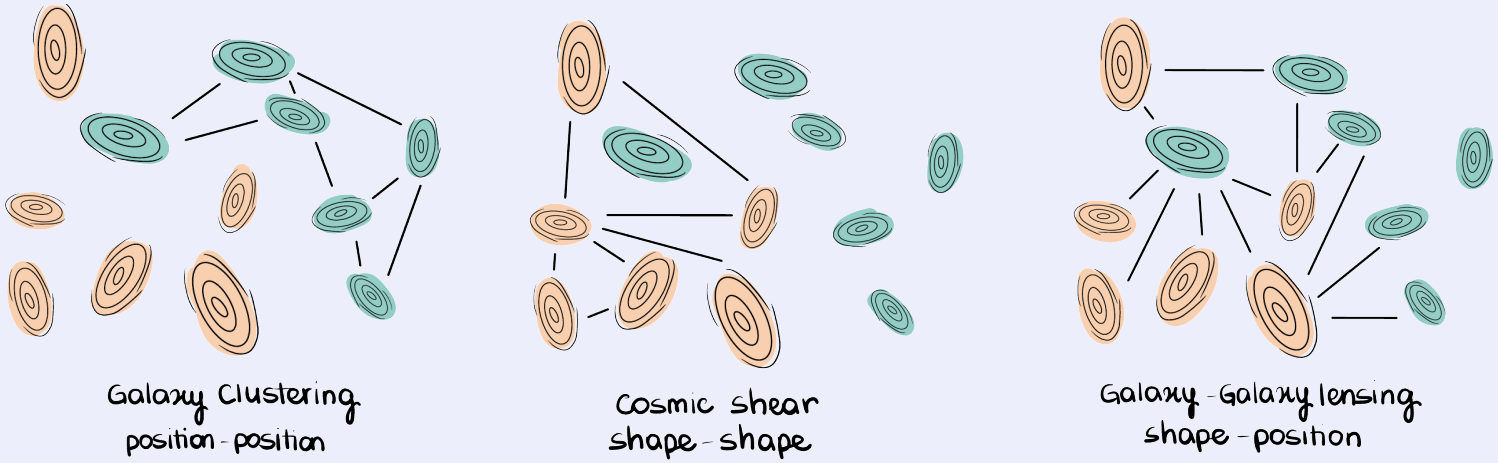
- Access genuinely new information

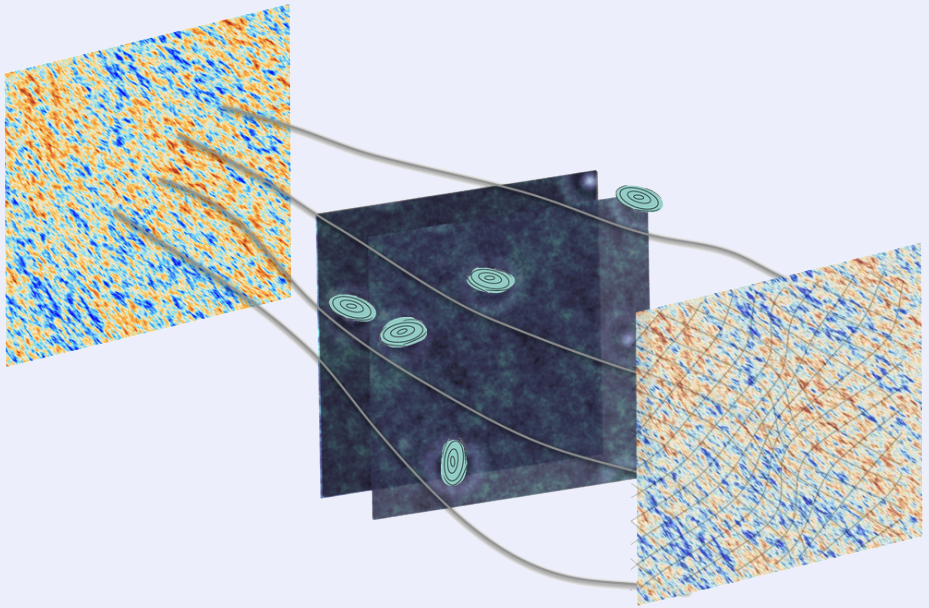


Blanchard, 2022

MOTIVATION

Why angular clustering?



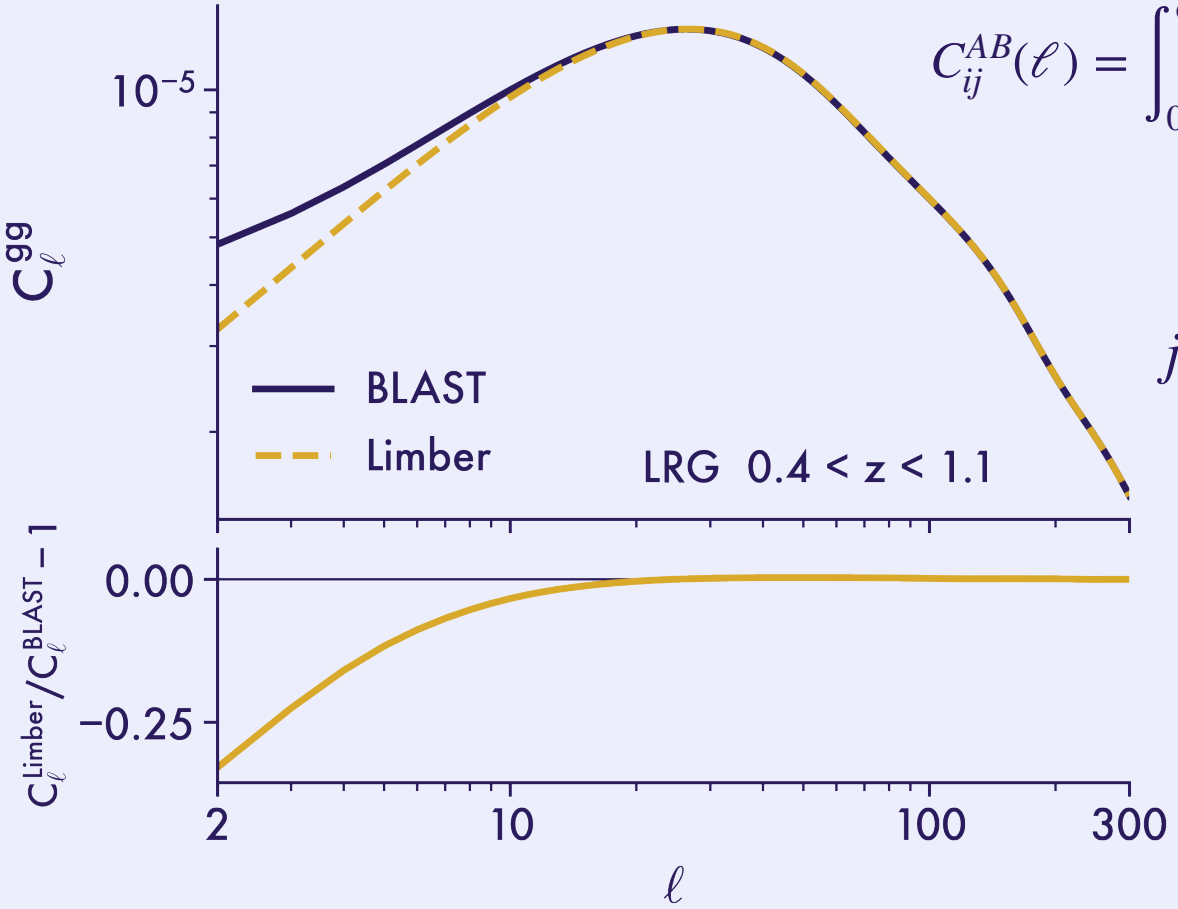


All these fields naturally live on the sphere!

$$C_{ij}^{AB}(\ell) = \int_0^\infty \mathrm{d}\chi_1 W_i^A(\chi_1) \int_0^\infty \mathrm{d}\chi_2 W_j^B(\chi_2) \int_0^\infty \mathrm{d}k\, k^2 P(k, \chi_1, \chi_2)\, j_\ell(k\chi_1)\, j_\ell(k\chi_2)$$

MOTIVATION

Large scales needs the full calculation!



$$C_{ij}^{AB}(\ell) = \int_0^\infty \mathrm{d}\chi_1 W_i^A(\chi_1) \int_0^\infty \mathrm{d}\chi_2 W_j^B(\chi_2) \int_0^\infty \mathrm{d}k\, k^2 P(k, \chi_1, \chi_2)\, j_\ell(k\chi_1)\, j_\ell(k\chi_2)$$

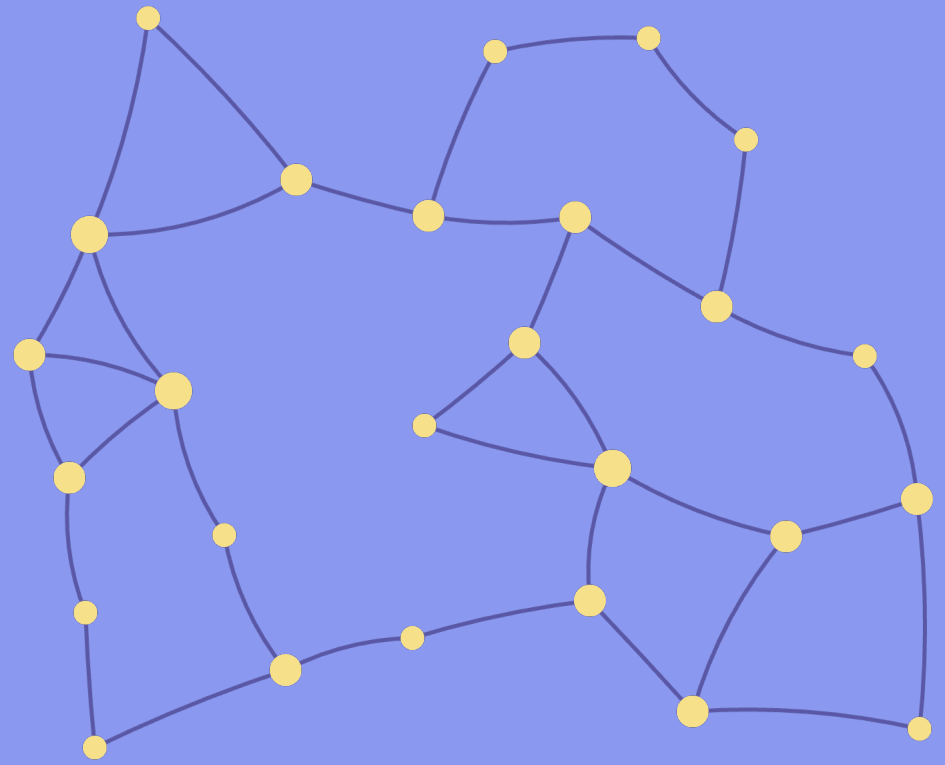
$$j_\ell(x) \sim \delta_D\left(\ell + \frac{1}{2} - x\right) \Longrightarrow \text{Inner } k \text{ integral} \propto \delta_D(\chi_1 - \chi_2)$$

$$C_{AB}(\ell) = F_\ell^{AB} \int \frac{d\chi}{\chi^2} W_A(\chi) W_B(\chi) P_{AB}\left(k = \frac{\ell + 0.5}{\chi}, z\right)$$

SECTION

02

Blast



BLAST

The algorithm



Decomposition of the power spectrum on the basis of the Chebyshev polynomials to solve the inner k integral

- Close to optimal polynomial approximant

$$P\left(k,\chi_1,\chi_2;\theta\right) \approx \sum_{n=0}^{N_{\max}} c_n\left(\chi_1,\chi_2;\theta\right) T_n(k)$$

- The hardest part of the integral can be precomputed

$$w_\ell(\chi_1,\chi_2) = \sum_{n=0}^{N_{\max}} c_n\left(\chi_1,\chi_2;\theta\right) \int_{k_{min}}^{k_{max}} \mathrm{d}k\, k^2 T_n(k) j_\ell\left(k\chi_1\right) j_\ell\left(k\chi_2\right) \equiv \sum_{n=0}^{N_{\max}} c_n\left(\chi_1,\chi_2;\theta\right) \hat{T}_n(\chi_1,\chi_2)$$

Chiarenza 2024, Chiarenza 2026a

BLAST

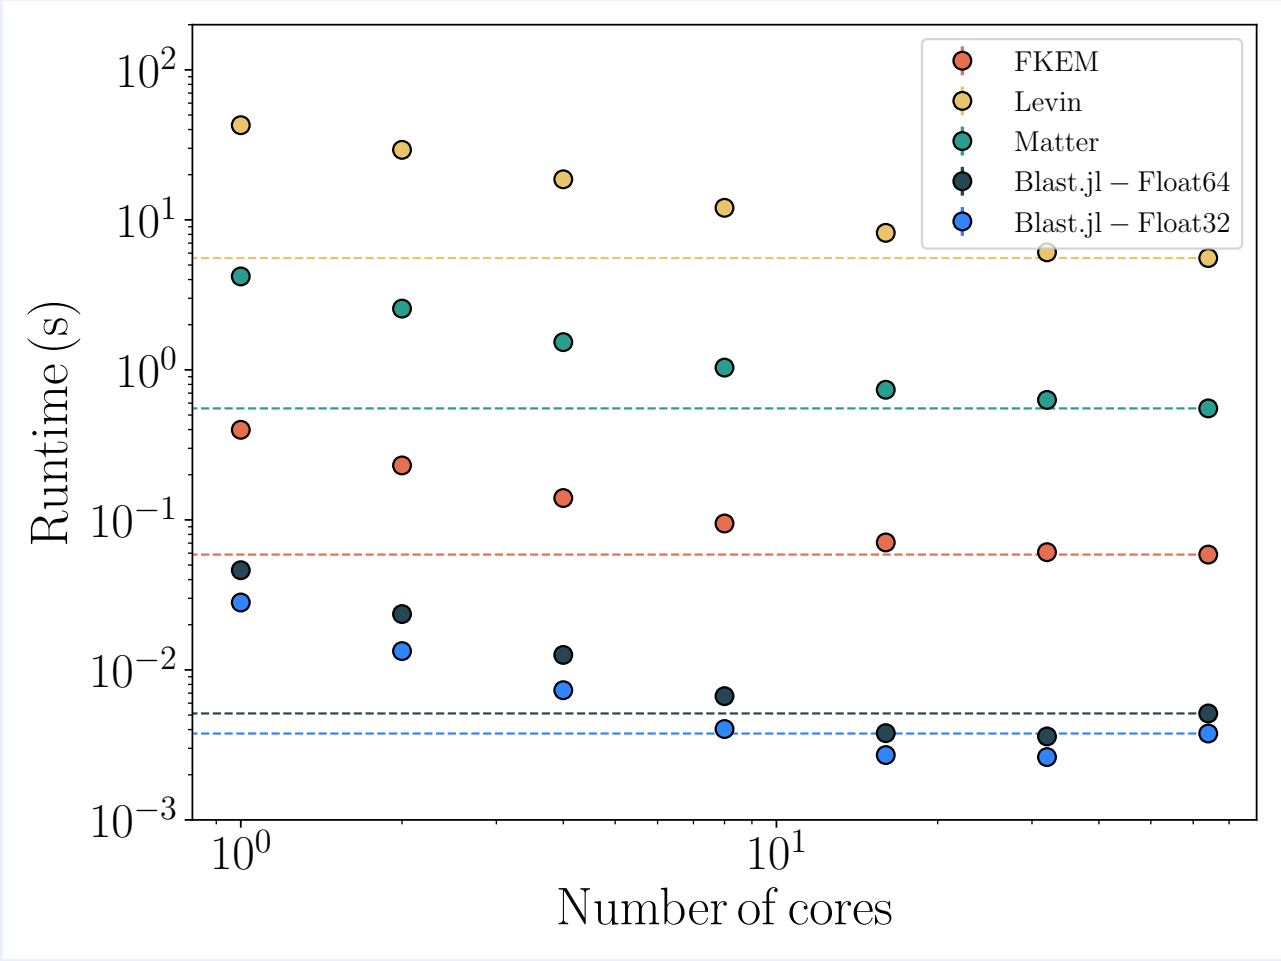
Blast 3×2pt

Computes 3×2pt angular power spectra: clustering, galaxy-galaxy lensing, shear

10x faster than the existing non-Limber algorithms

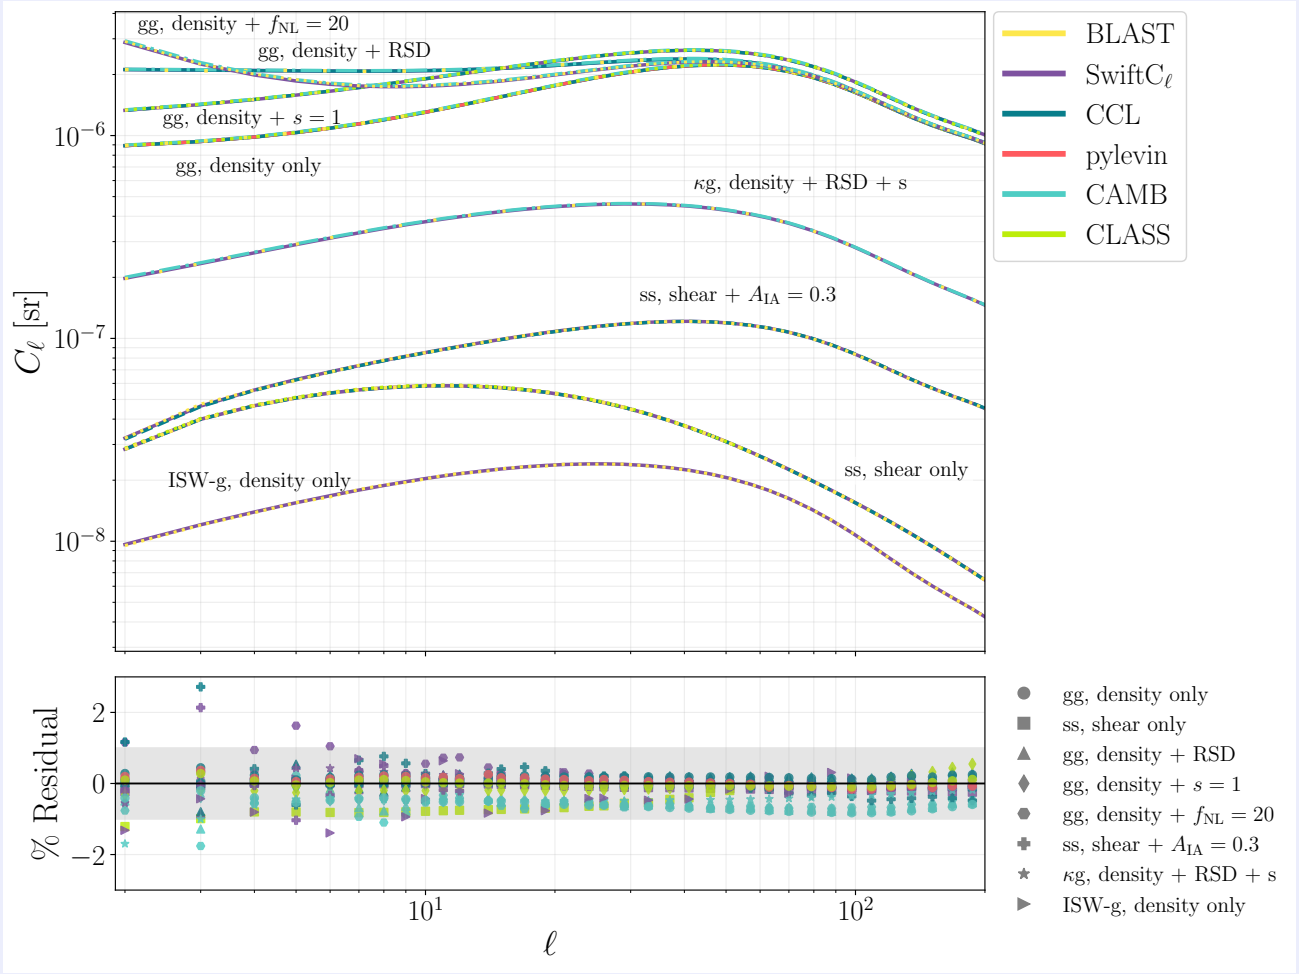
Runtime independent on the number of tomographic bins

Accuracy within LSST Y10 requirements



BLAST

Blast 6×2pt: all the physics, with gradients



Full physics: add RSD, mag bias, PNG, IA, CMB Lensing and ISW

Same trick, but extended: there are just more **precomputed** cosmology independent integrals

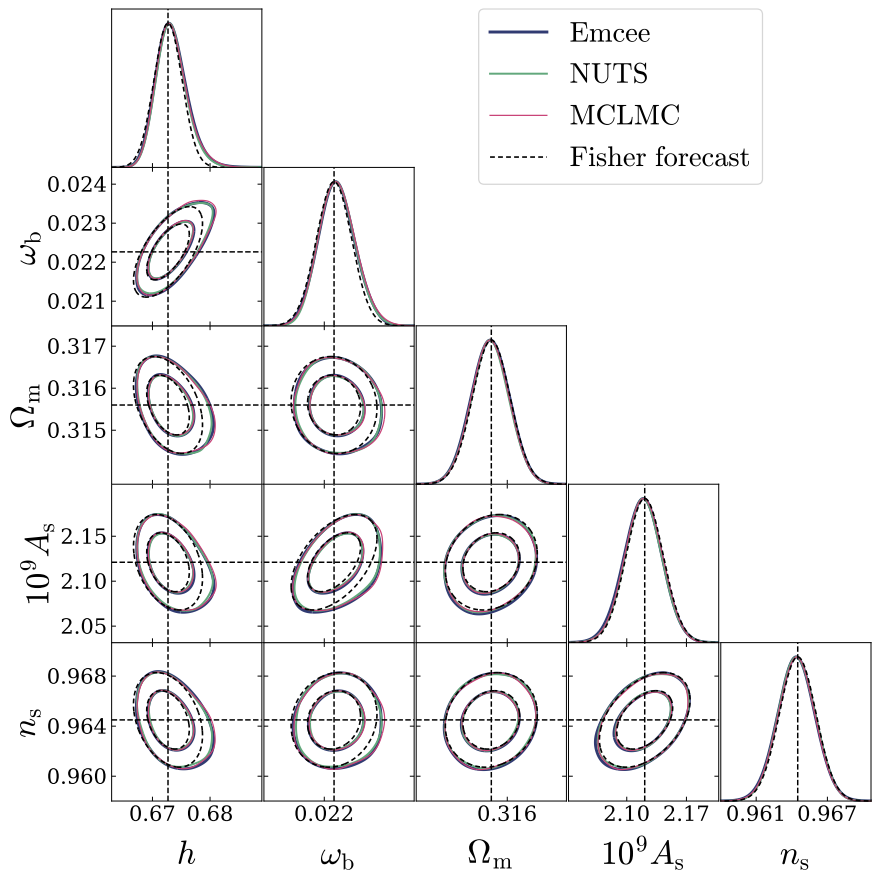
Fully differentiable with custom reverse-mode AD: the gradient cost independent on the number of parameters

Validated: sub-percent agreement with brute force integration and other codes

BLAST

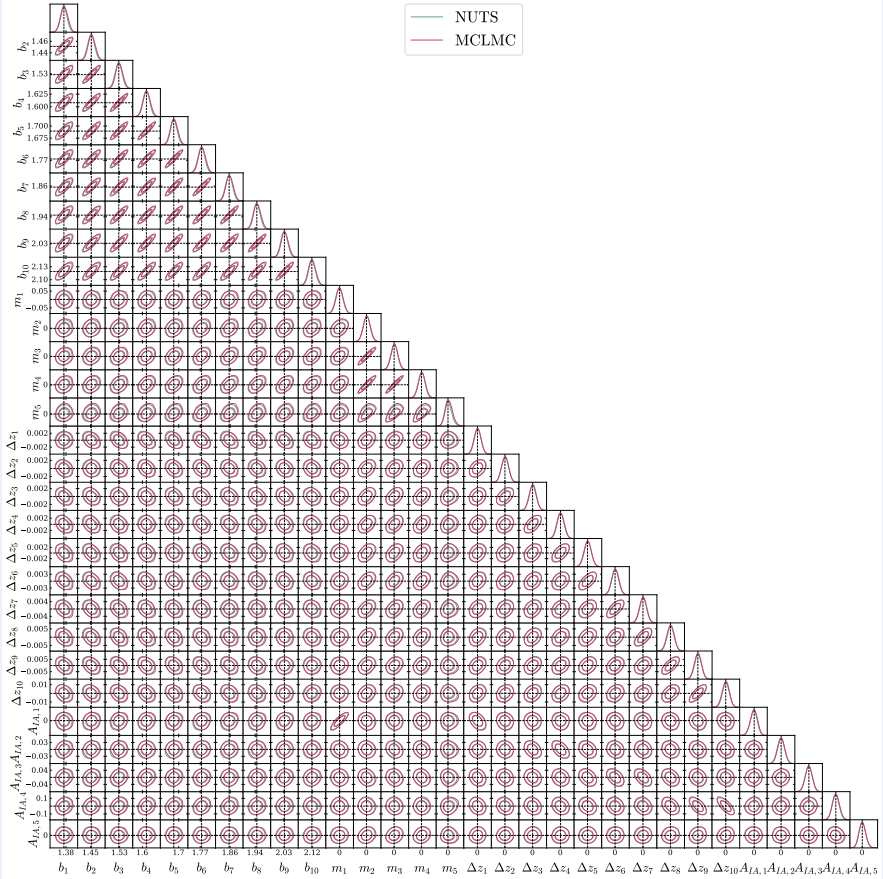
Sampling performance

LSST Y10: 10 clustering + 5 shear bins



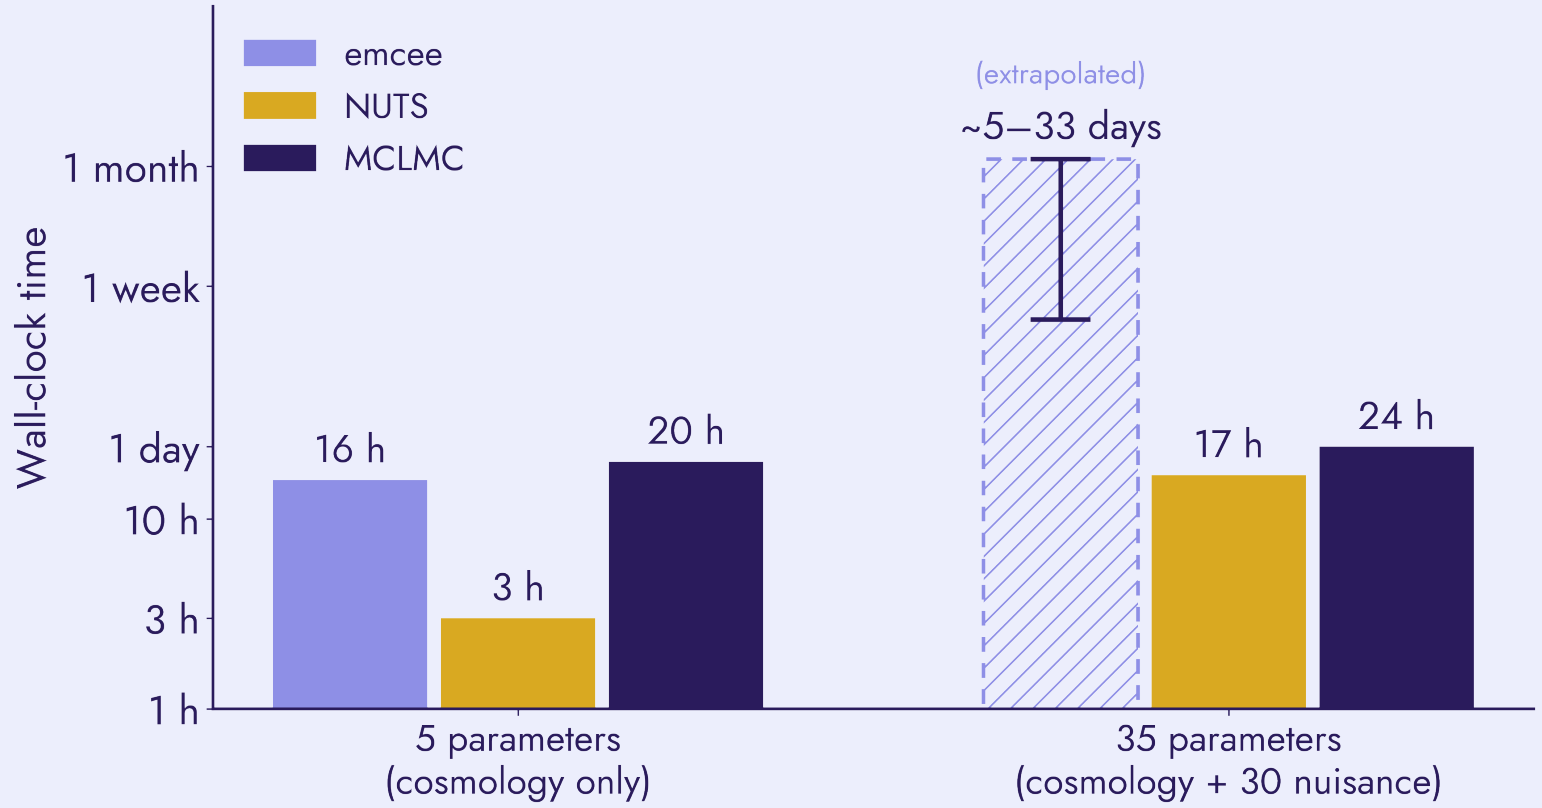
Gradient-based samplers recover the same posteriors as traditional samplers

But the sampling is much more efficient, enabling exploration of large parameter spaces



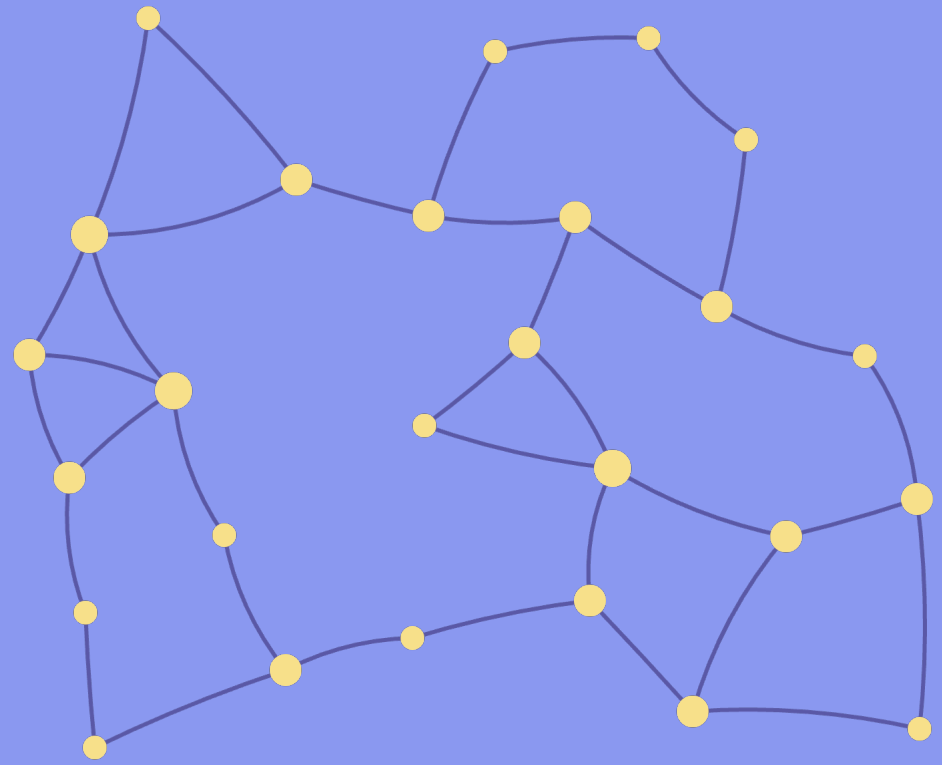
BLAST

Sampling performance in numbers



Inference takes **hours** instead of **weeks**

SECTION



03

PNG constraints from DESIxCMB lensing

The Dark Energy Spectroscopic Instrument





ACKNOWLEDGMENTS

DESI is an international science collaboration managed by the Department of Energy's (DOE's) Laurence Berkeley National Laboratory (Berkeley Lab) with primary funding for construction and operations from DOE's Office of Science. Kitt Peak National Observatory is a program of the National Science Foundation's (NSF's) NOIRLab, which the DOE contracts with to operate the Nicholas U. Mayall 4-meter Telescope for the DESI survey.

U.S. DEPARTMENT OF ENERGY | Office of Science

This exhibit was made possible thanks to the generous support of the following partners and the whole DESI collaboration:

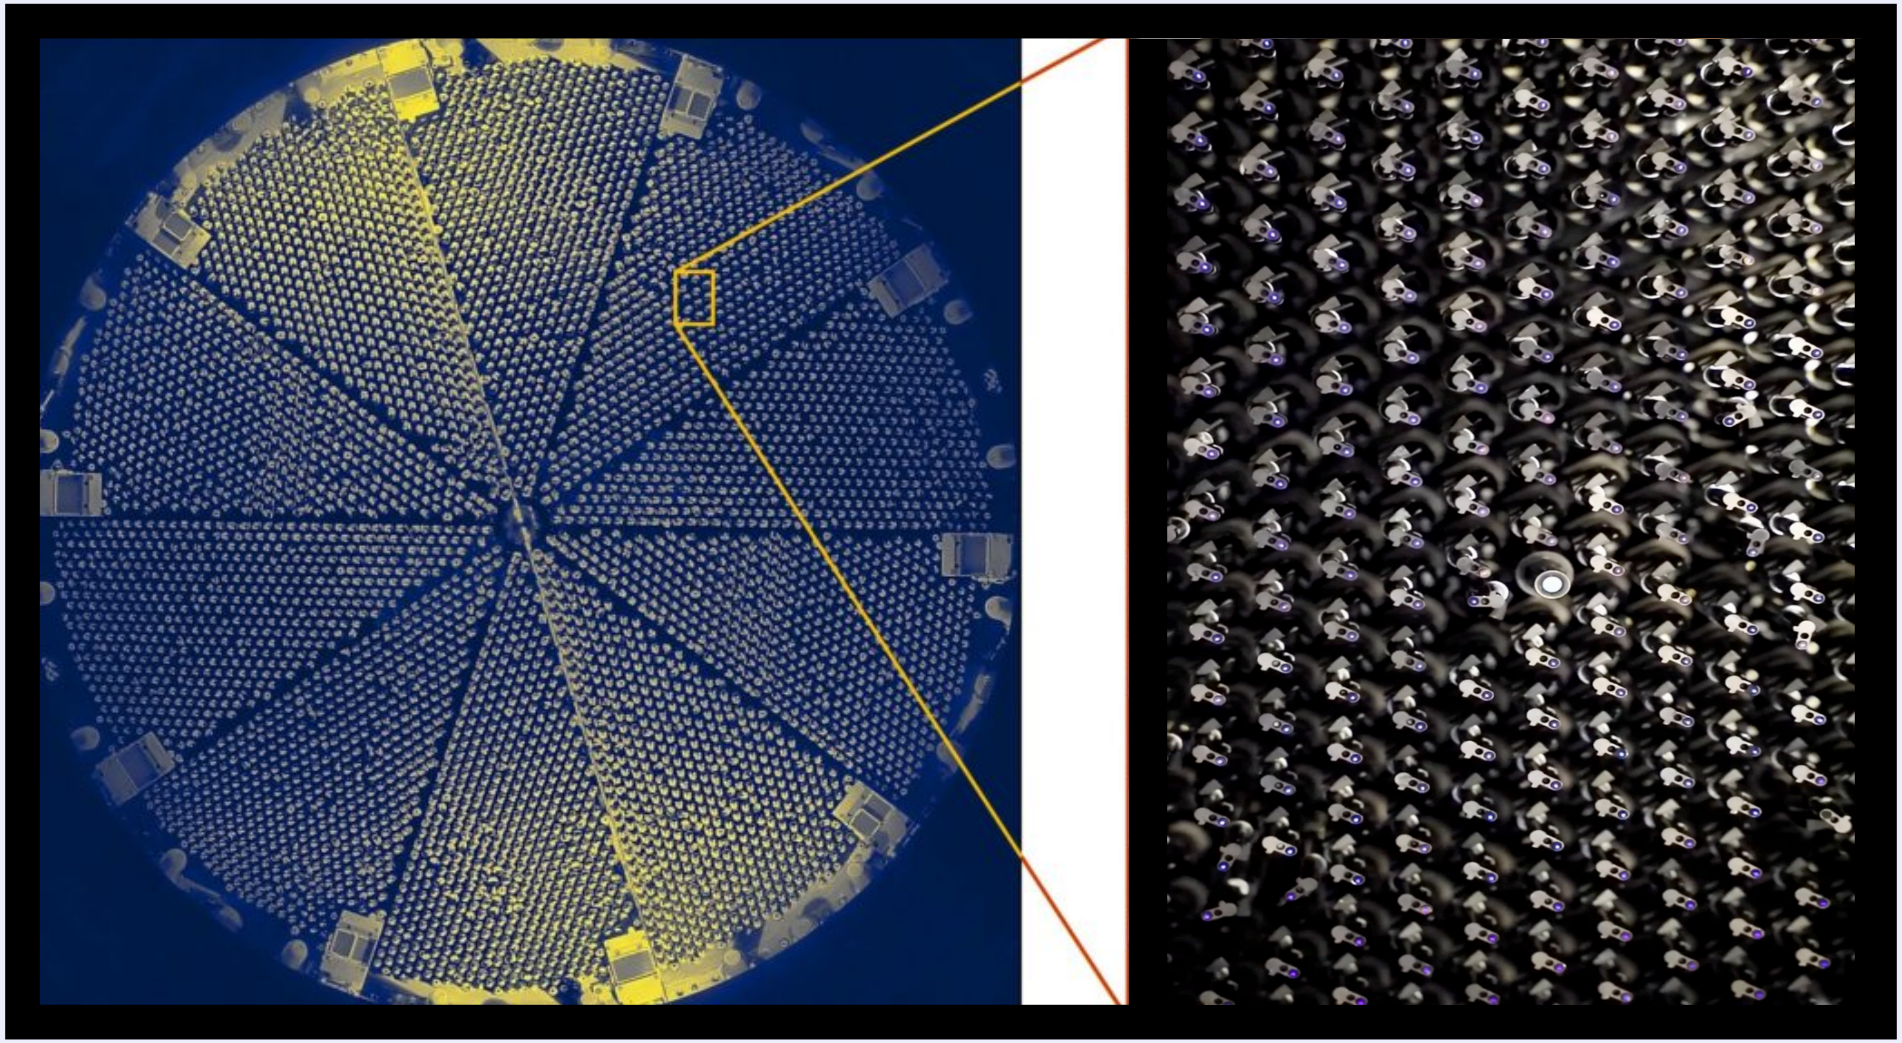
BERKELEY LAB · UNIVERSITY OF MICHIGAN · BOSTON UNIVERSITY · Berkeley UNIVERSITY OF CALIFORNIA · University of Pittsburgh

NOIR Lab · DARK ENERGY SPECTROSCOPIC INSTRUMENT

U.S. Department of Energy Office of Science

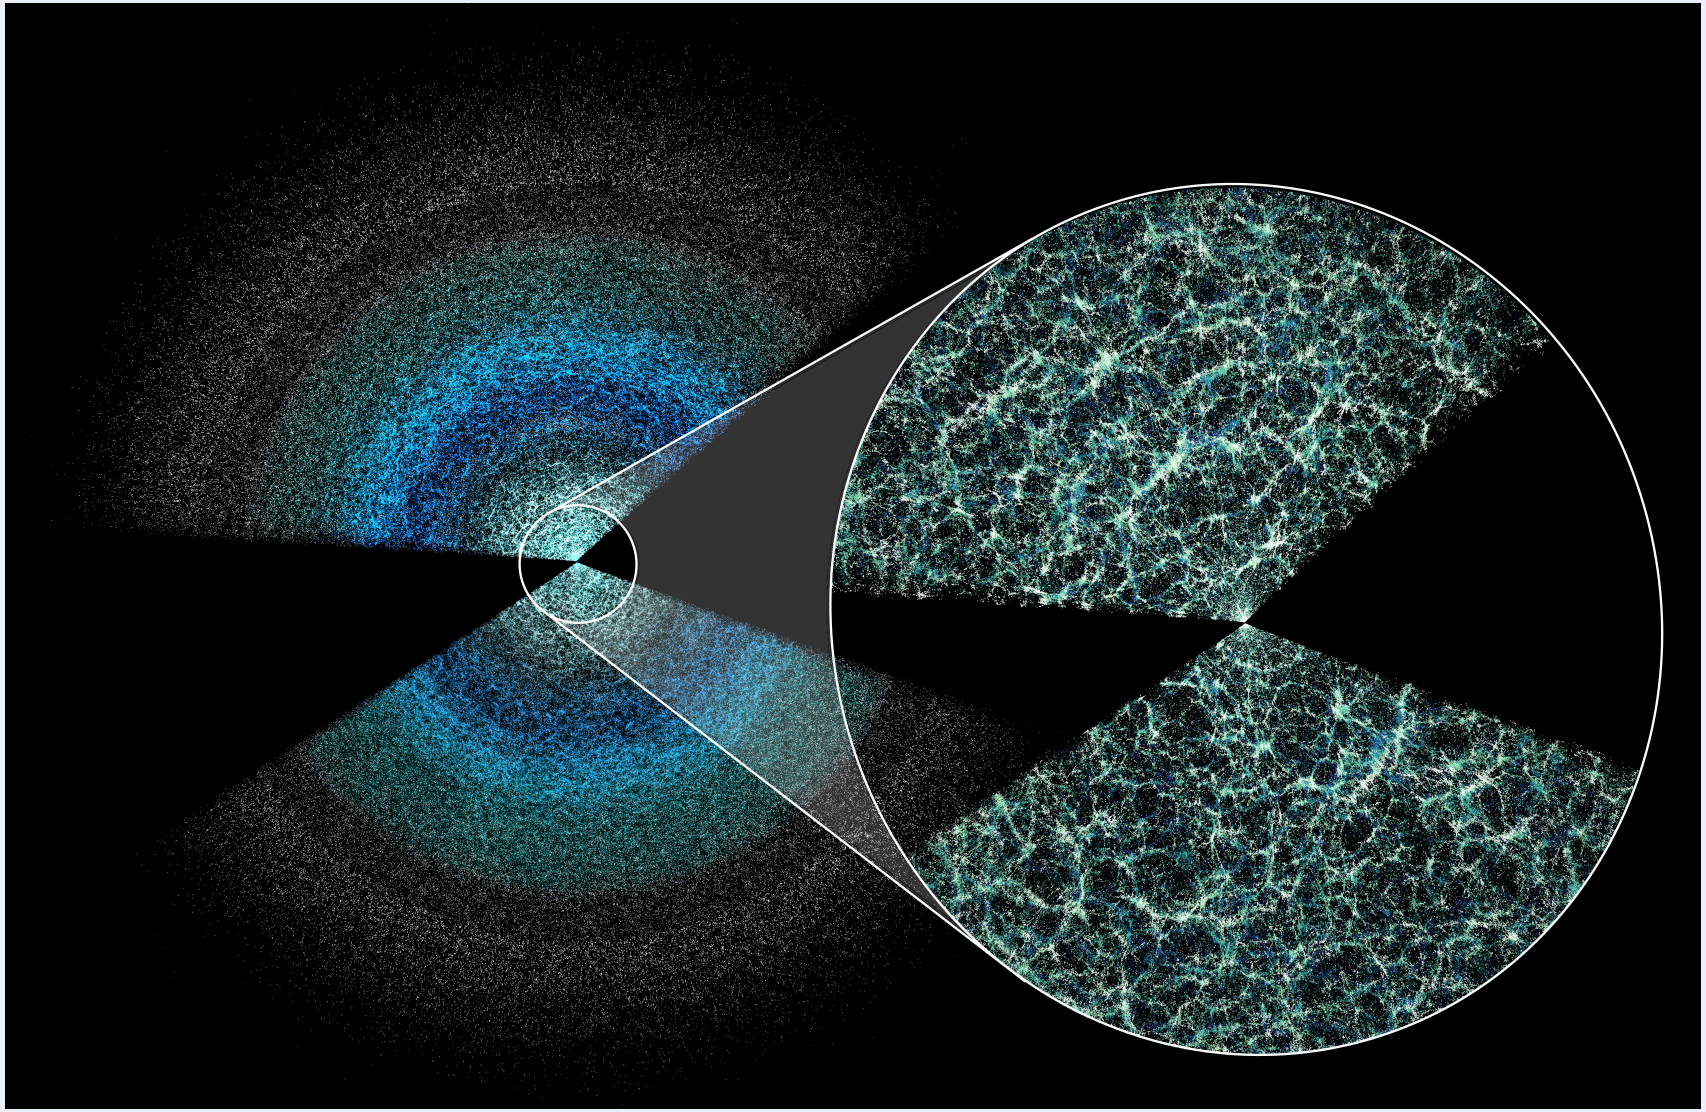
PNG & DESI

The Dark Energy Spectroscopic Instrument



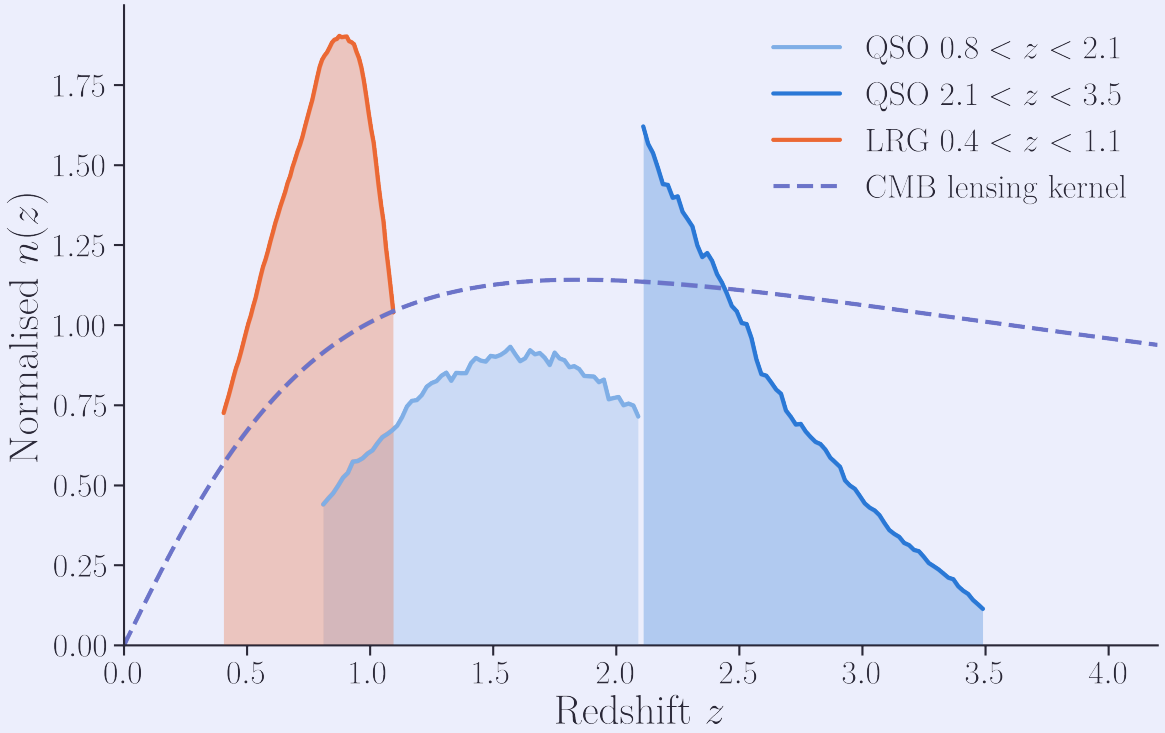
PNG & DESI

The Dark Energy Spectroscopic Instrument

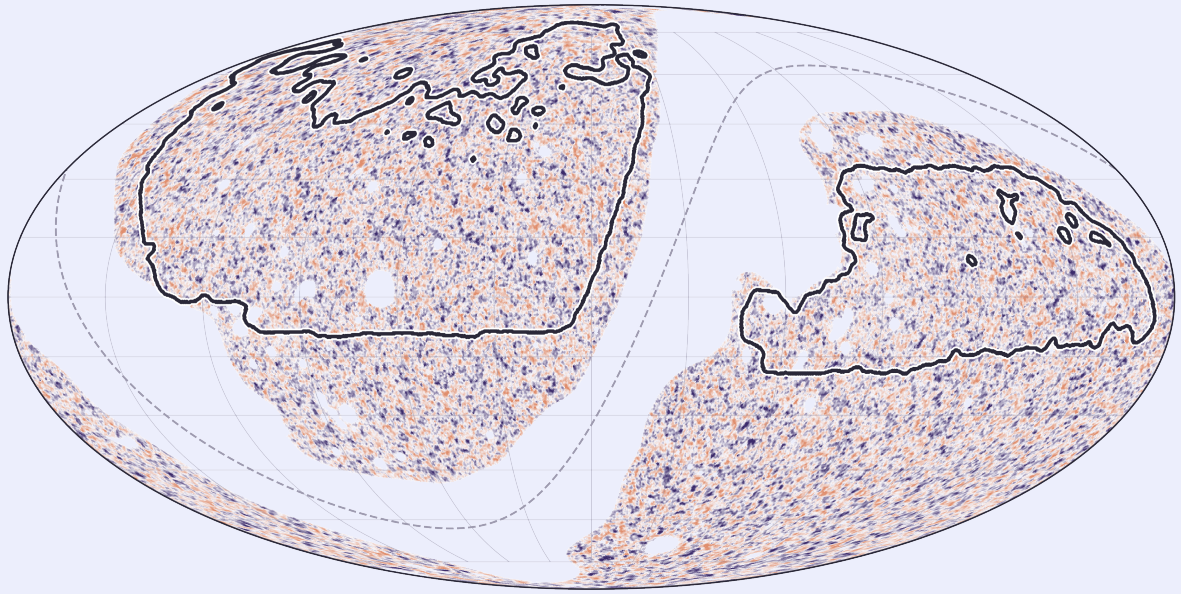


PNG & DESI

The datasets



×

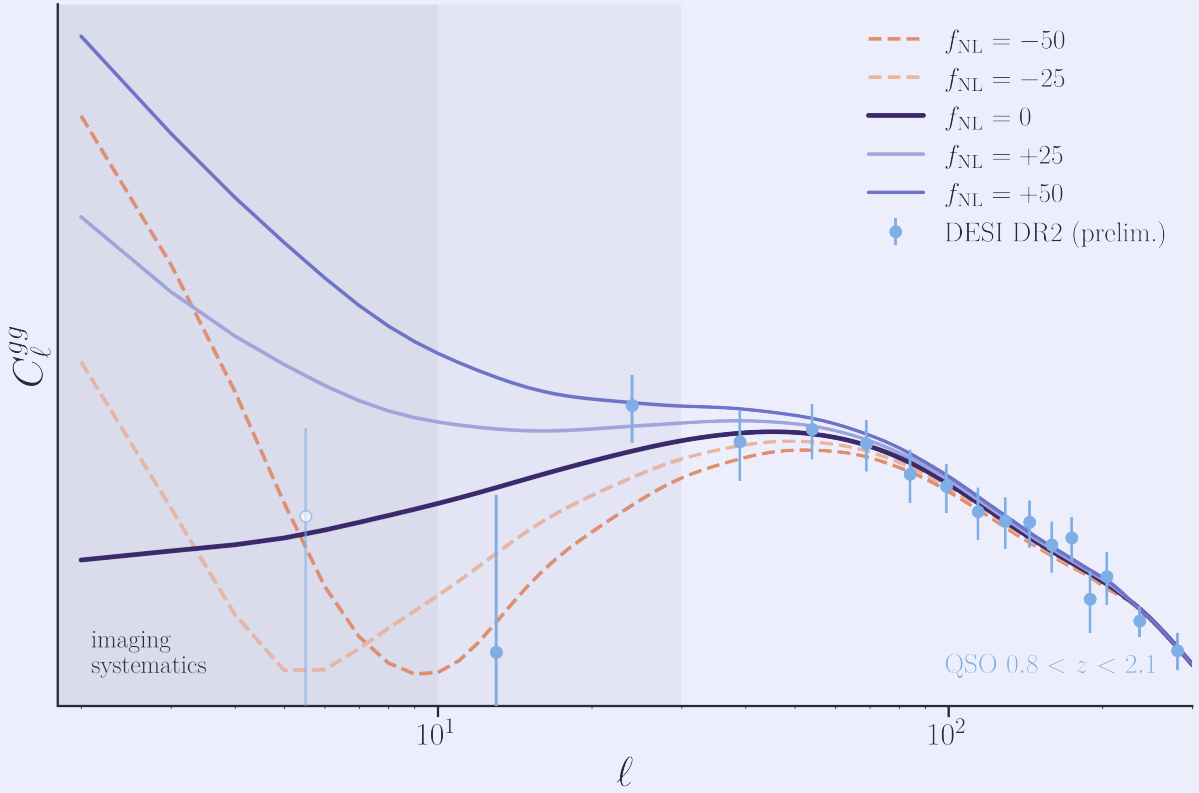


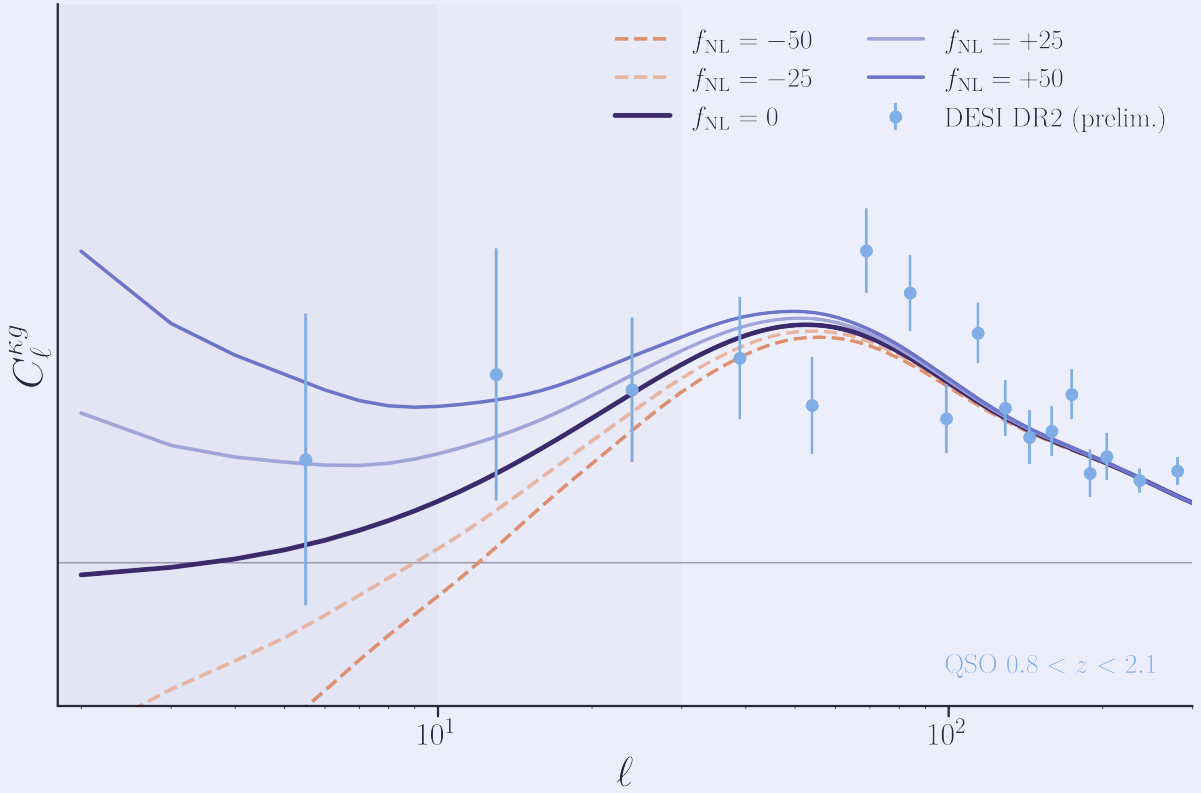
DESI DR2 LRG and QSO

Planck PR4 CMB Lensing

PNG & DESI

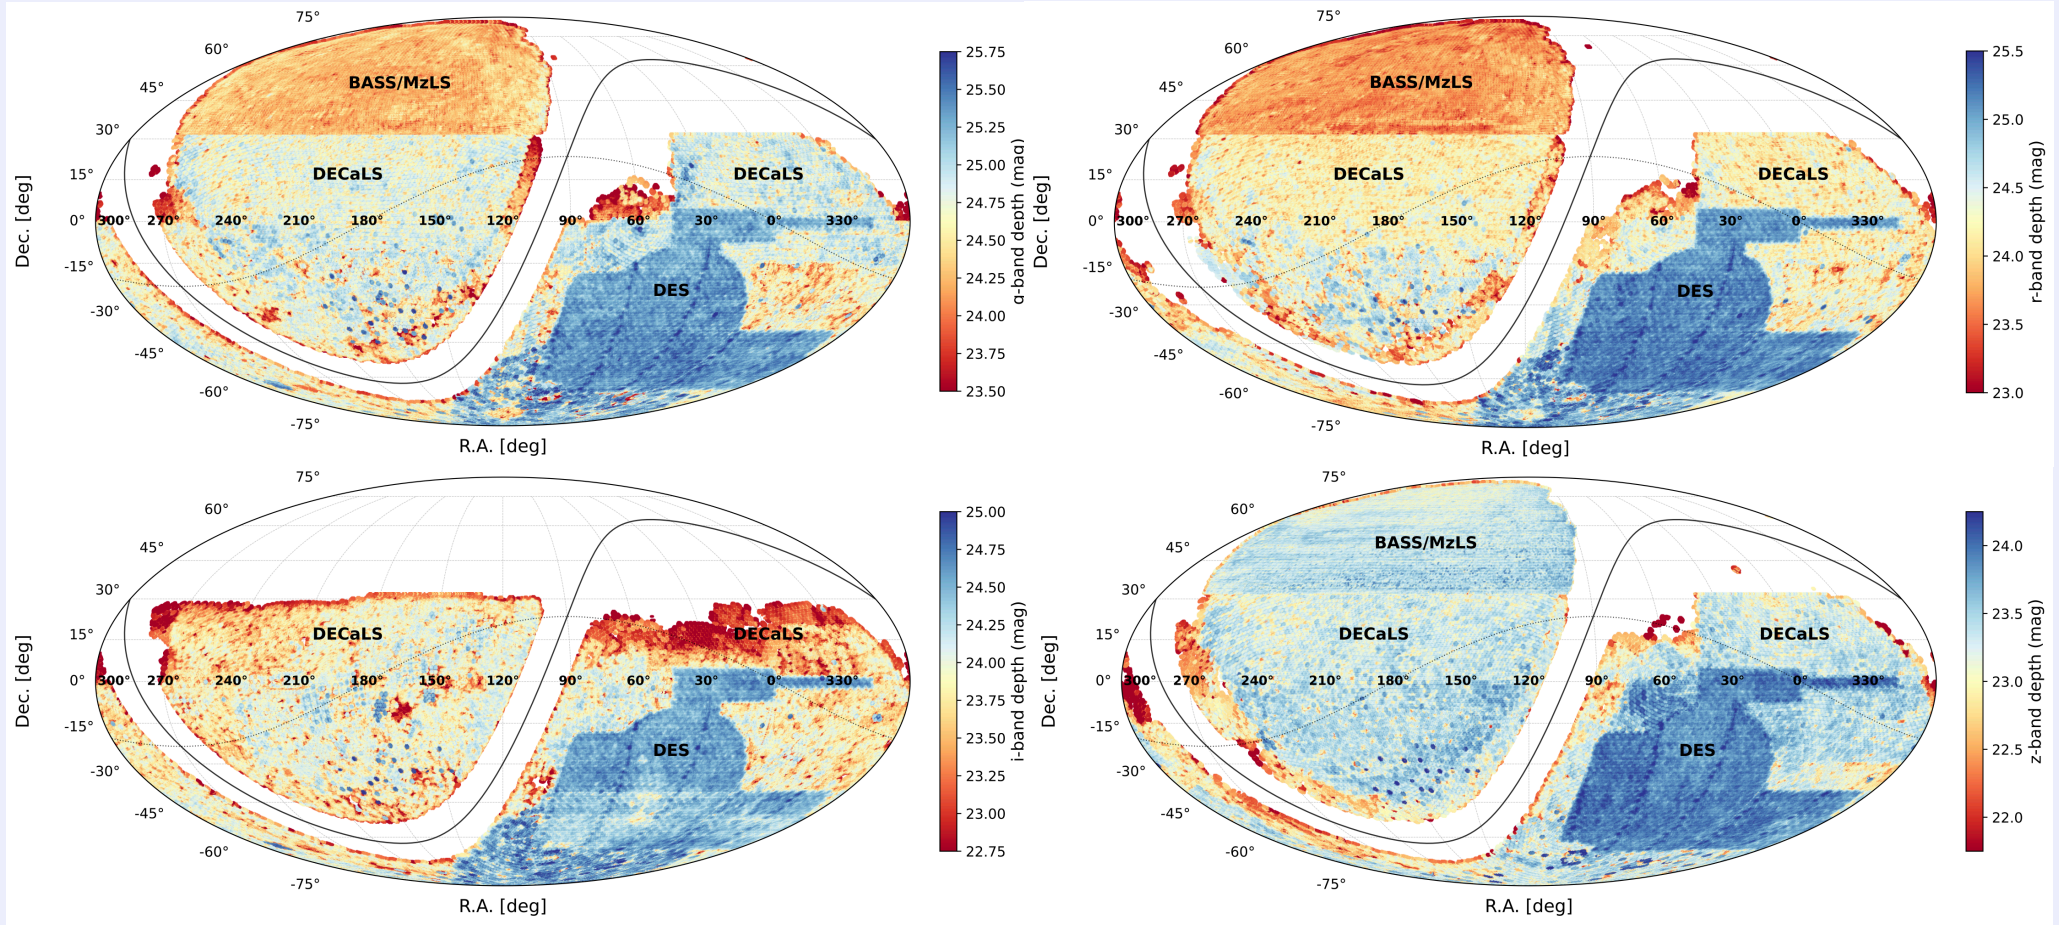
The signal





PNG and imaging systematics are very degenerate!

Dealing with imaging systematics

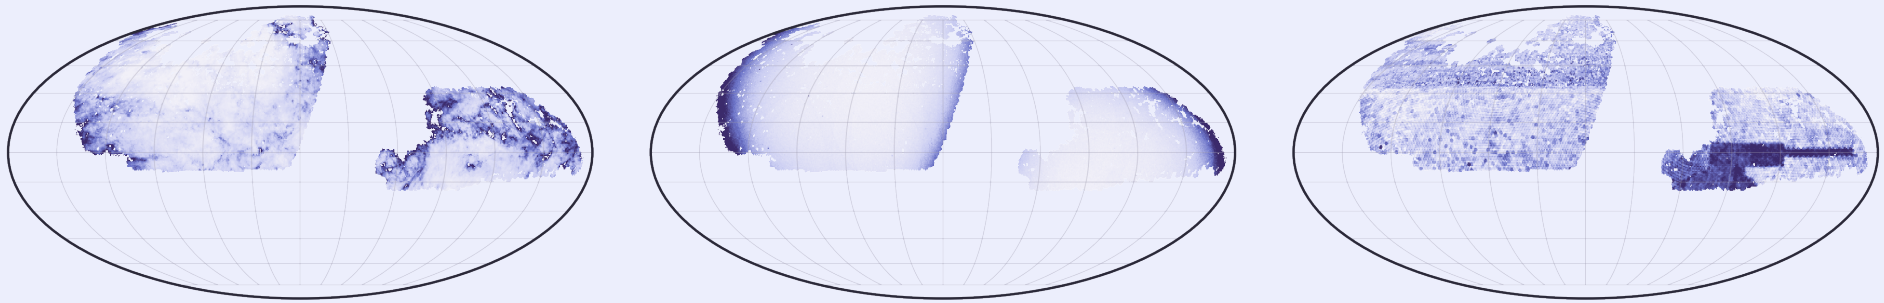


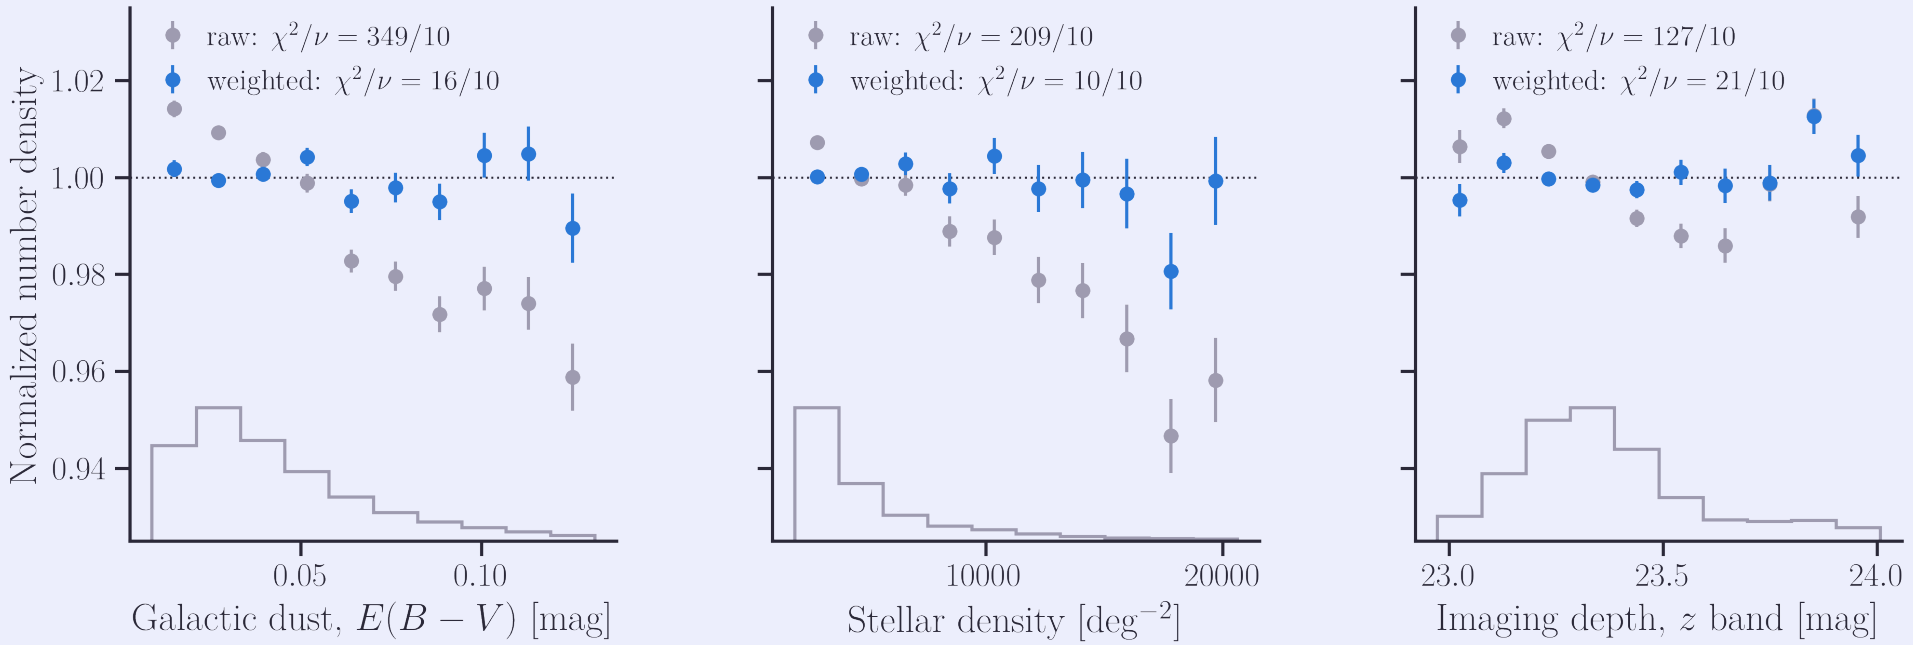
Legacy survey used for target selection

https://www.legacysurvey.org/status/

PNG & DESI

Defining systematic weights





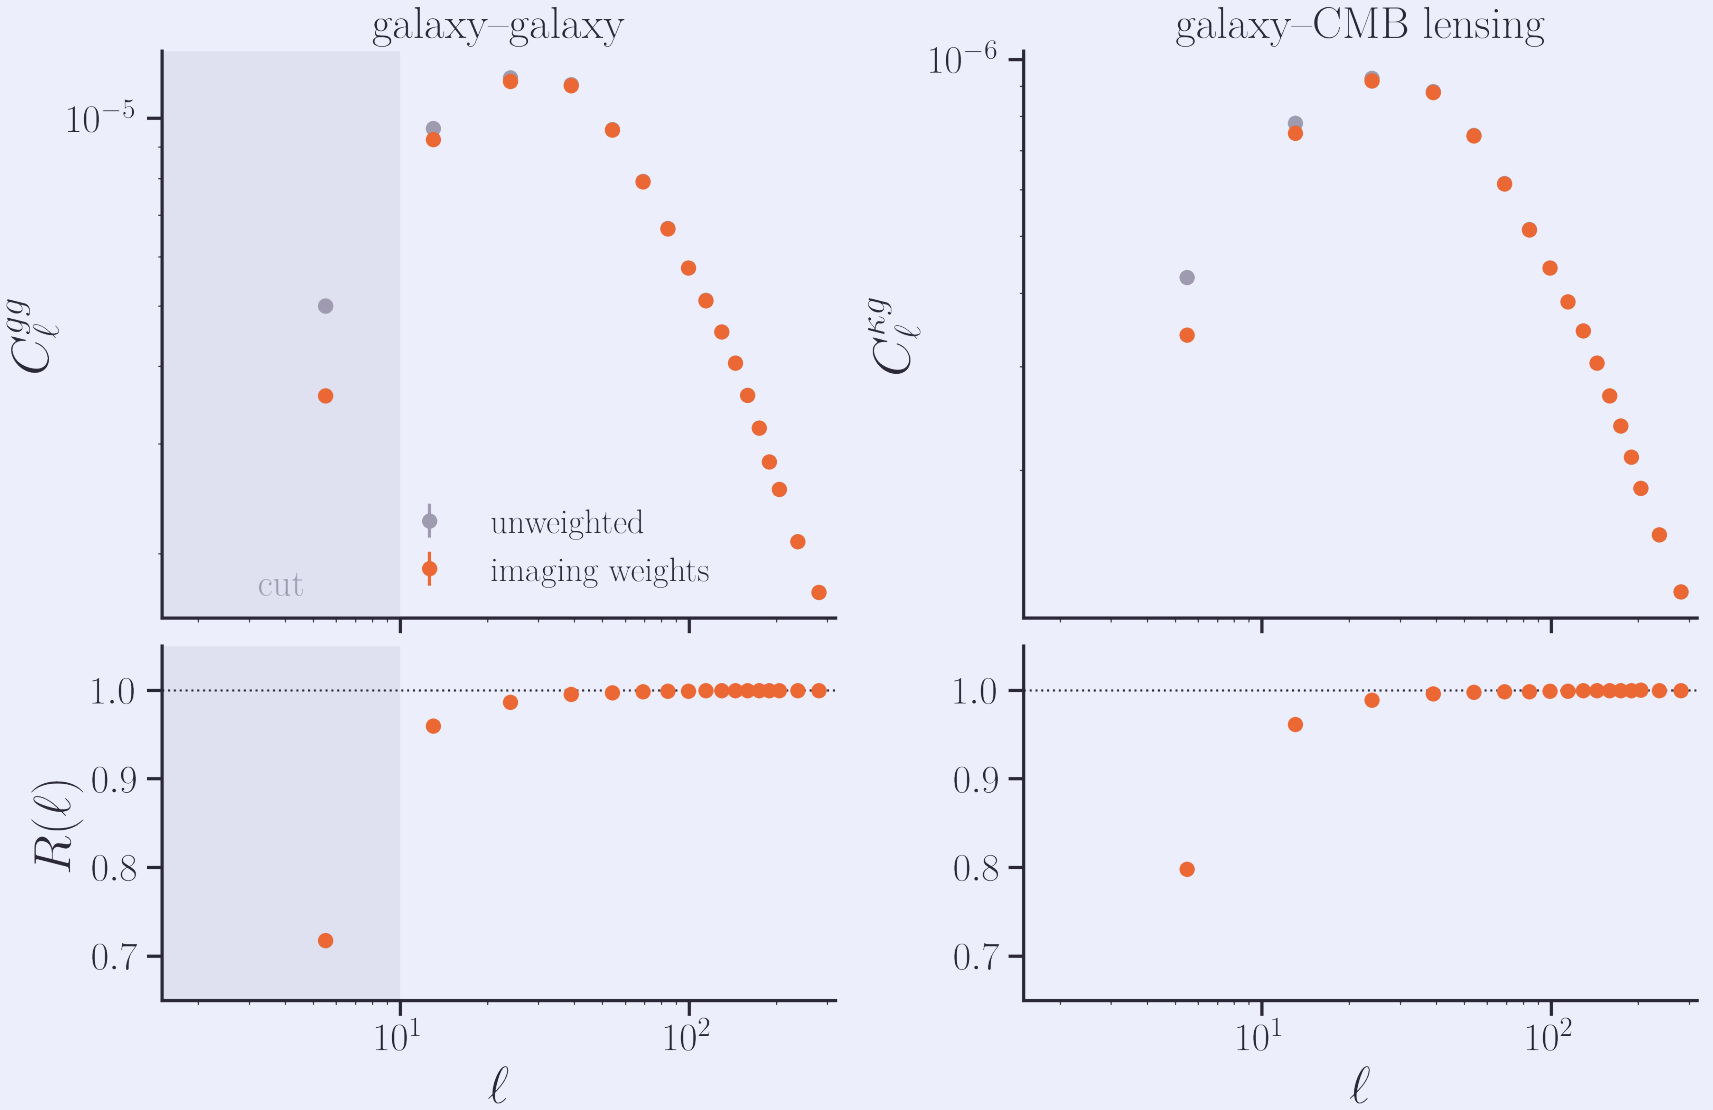
Without correction, the quasar density drops by up to 5% where there is more dust or more stars; the imaging weights remove this unwanted trend.

PNG & DESI

Angular Mode Removal

By correcting for the weights, some real **signal** is **lost**

We quantify this using **uncontaminated** mocks



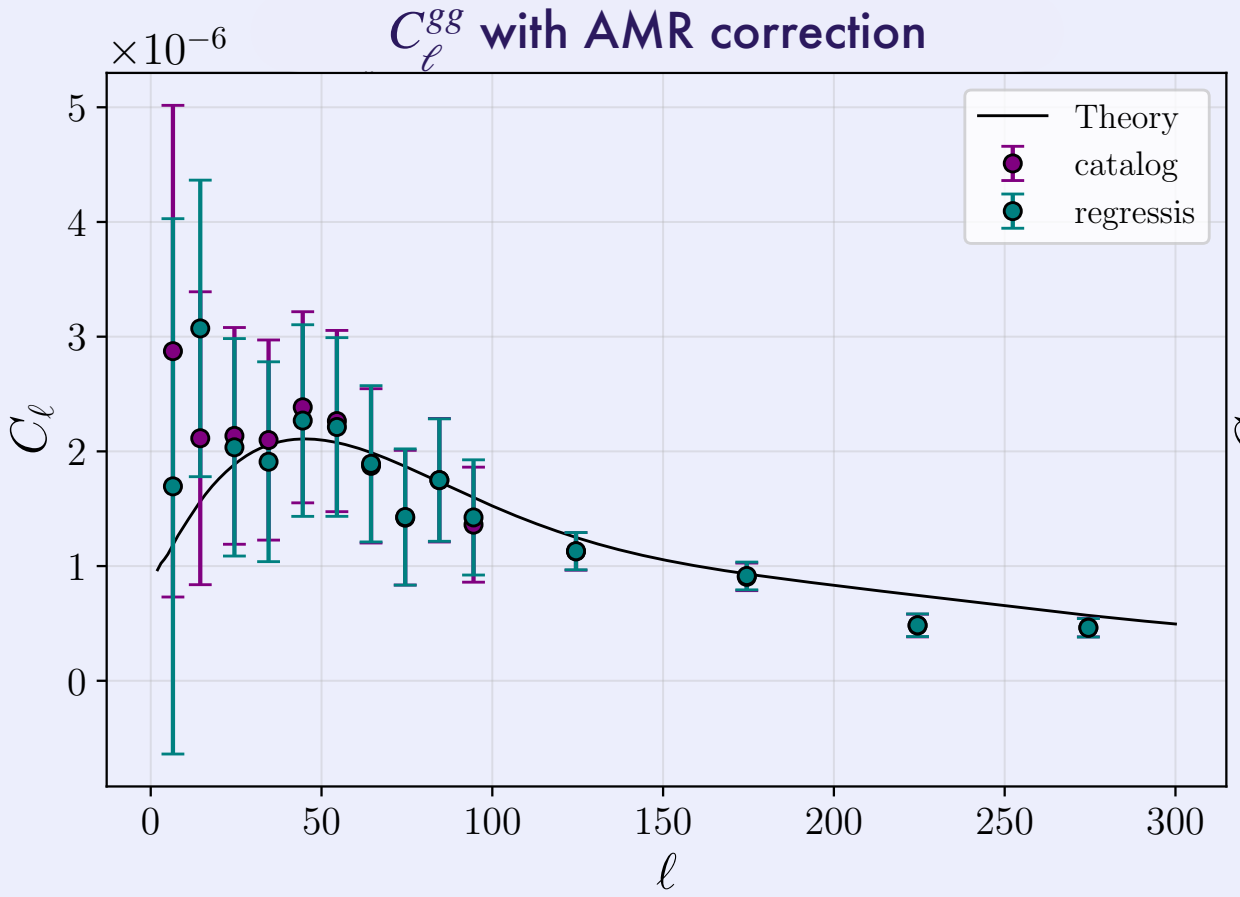
DR1 QSO x Planck PR4

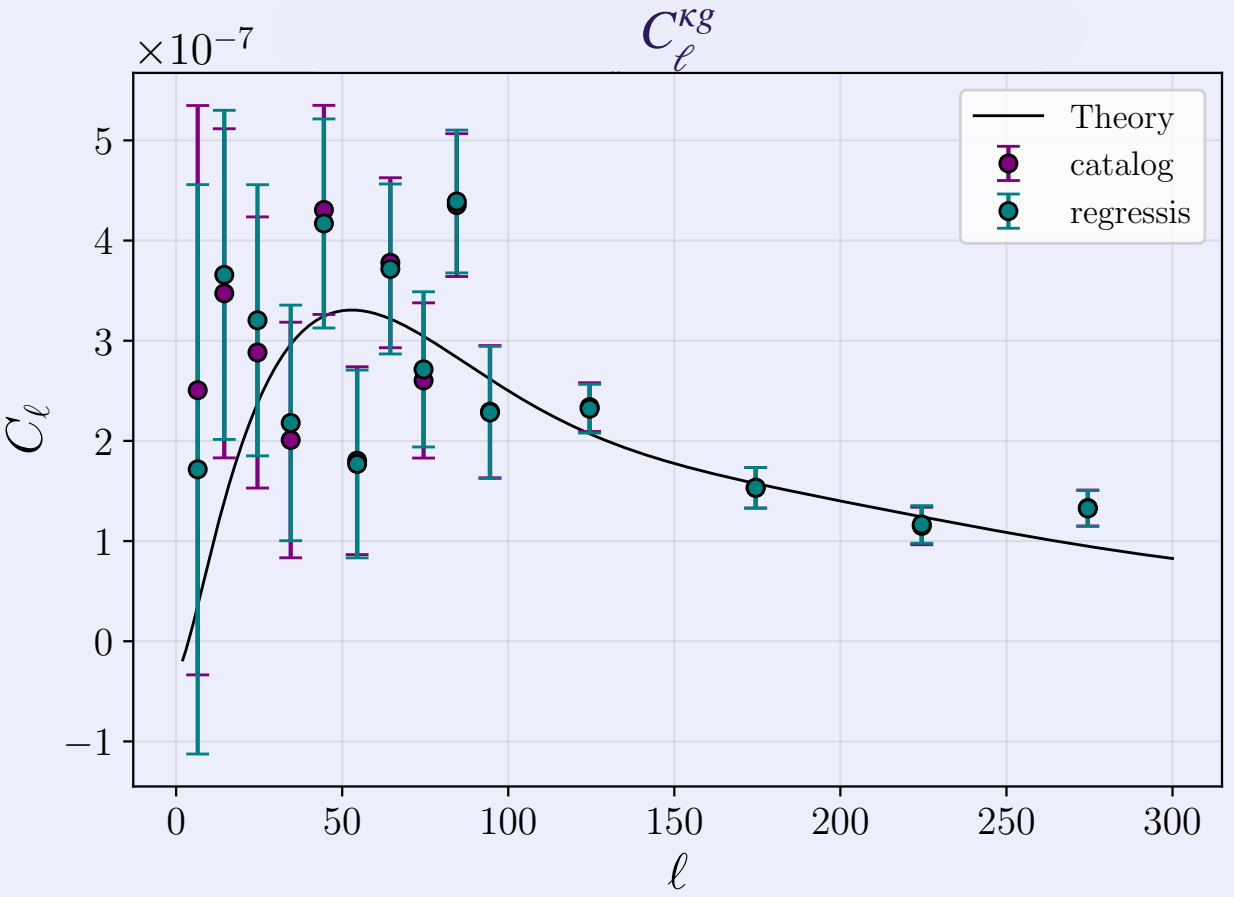
Constraining $f_{\rm NL}$ via the cross-correlation of DESI quasar targets and Plank 2018 CMB lensing (Krolewski 2024)

- DESI Legacy Survey Imaging → DESI Y1 spectroscopic quasars sample
  - → Smaller area, but higher density
  - → Cleaner sample with fewer imaging systematics
  - → Apply optimal $f_{\rm NL}$ weights
- Constraints for $f_{\rm NL}$, A_s and n_s instead of Planck 2018 Flat ΛCDM values
  - → Better error propagation

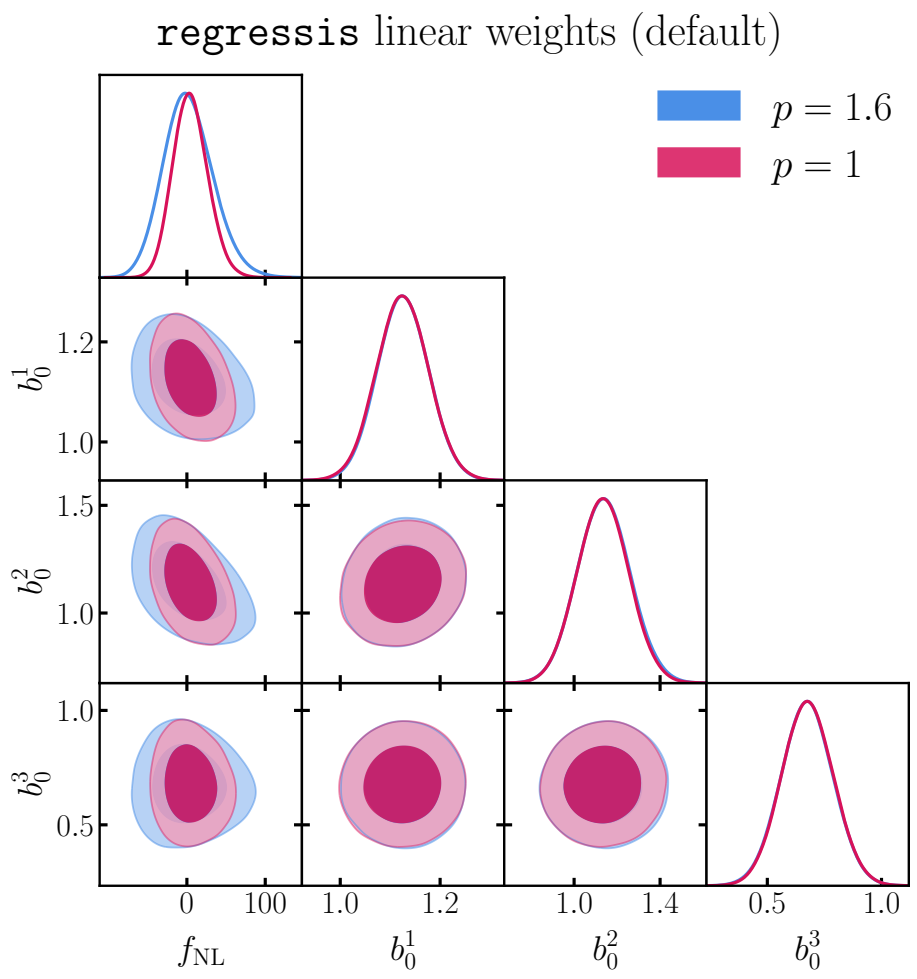
DR1 QSO x Planck PR4

The **auto power spectrum**, where systematic power usually shows up, **looks clean**!





DR1 QSO x Planck PR4



→ New catalog-based pipeline, now default in DESI DR2

→ $\approx 35\,\%$ tighter than Legacy Survey QSOs (Krolewski+23), comparable to Quaia (Fabbian+25)

→ Optimal $f_{\rm NL}$ weights, chains with free $A_{\rm s}$, $n_{\rm s}$

$$f_{\rm NL} = 2^{+28}_{-32}$$

DR1 QSO x Planck PR4

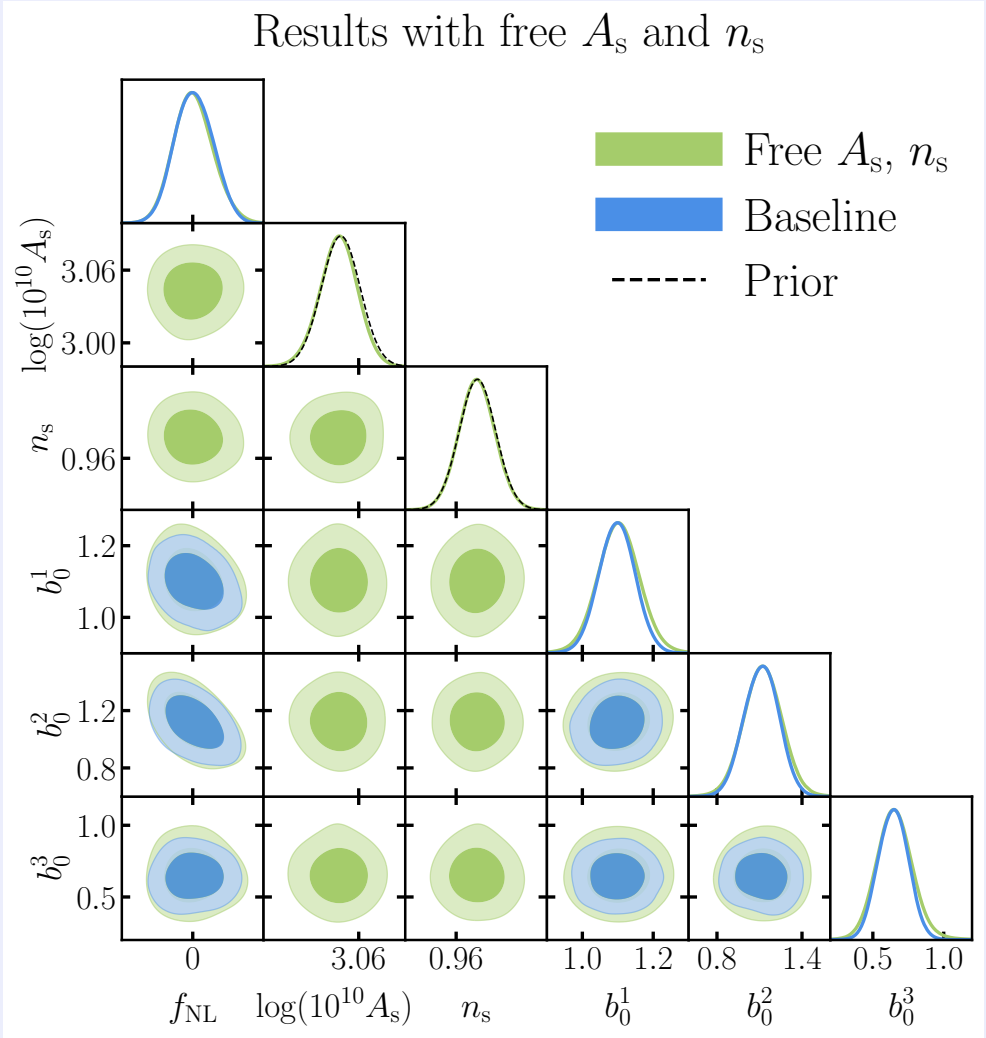
Impact of the primordial power spectrum

Planck 2018 priors on A_s an and n_s

$$A_s = (2.101 \pm 0.034) \times 10^{-9}$$

$$n_s = 0.9649 \pm 0.0042$$

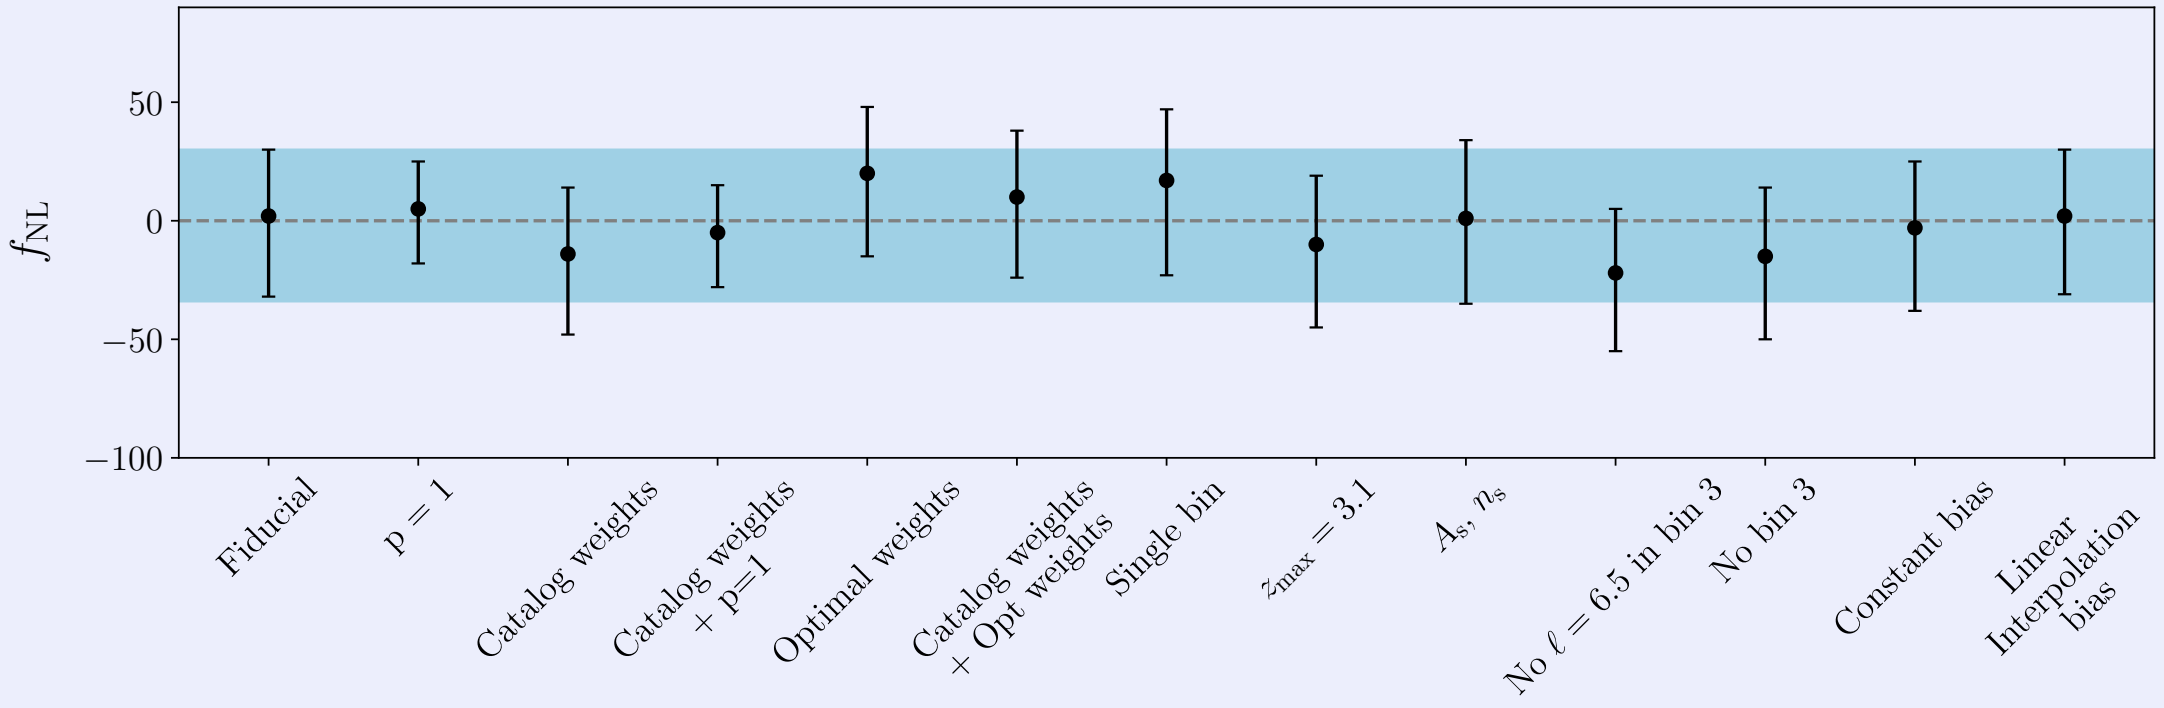
$$\boxed{f_{NL} = 1^{+33}_{-36}}$$



PNG & DESI

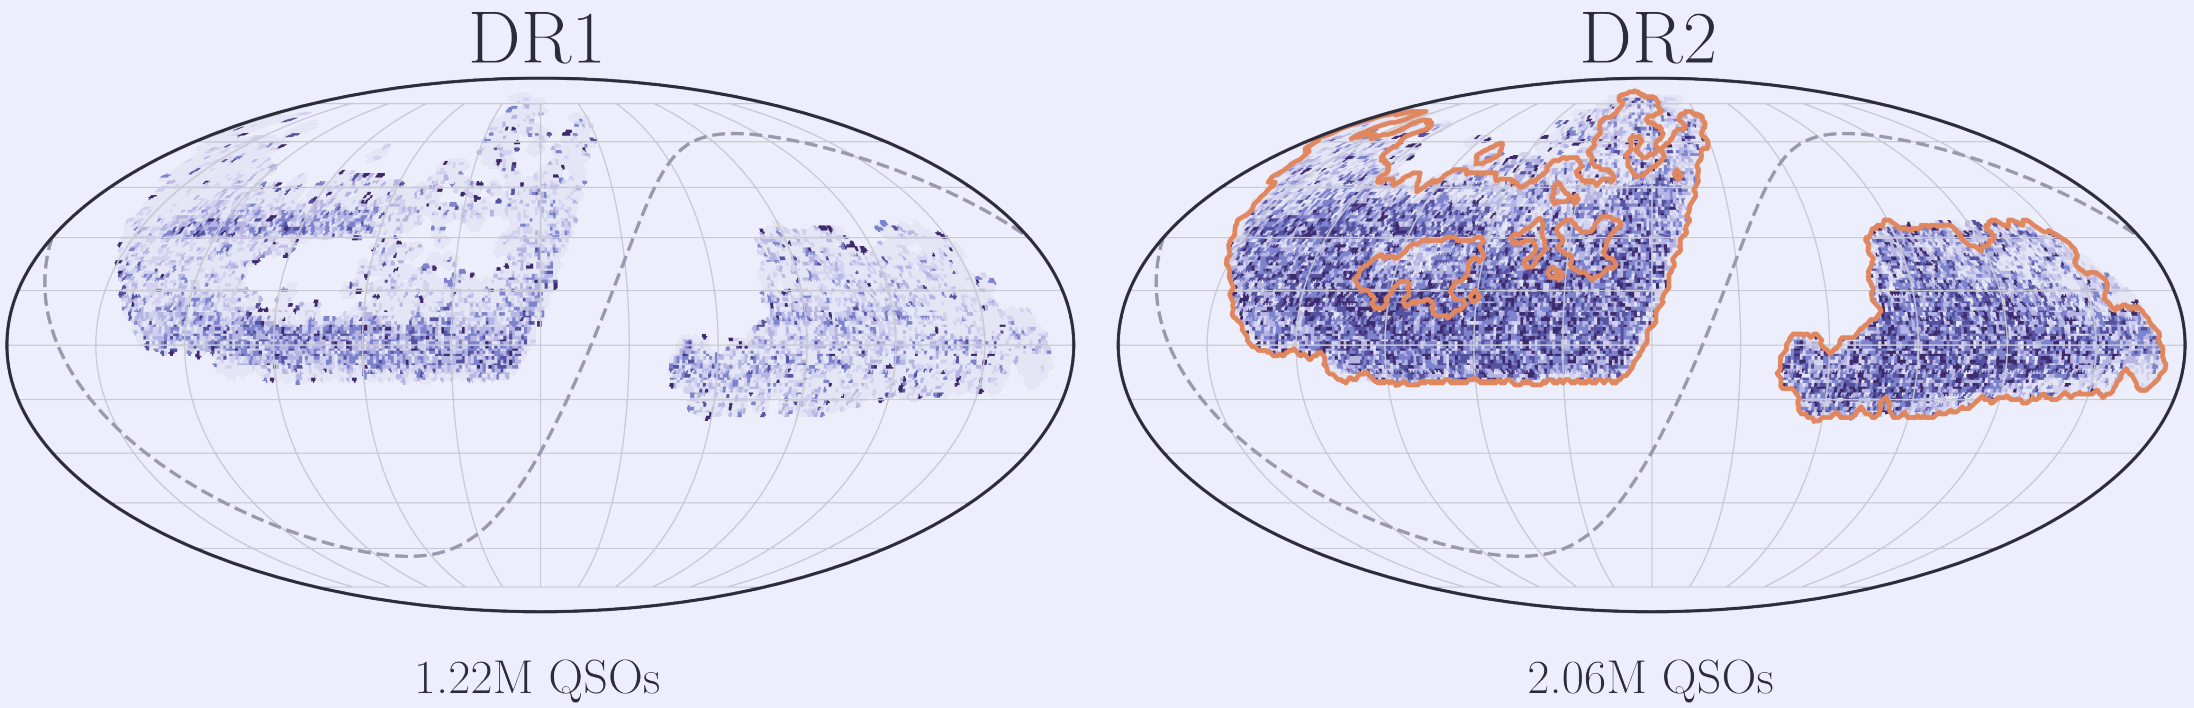
DR1 QSO x Planck PR4

Constraint is **robust** to analysis variations



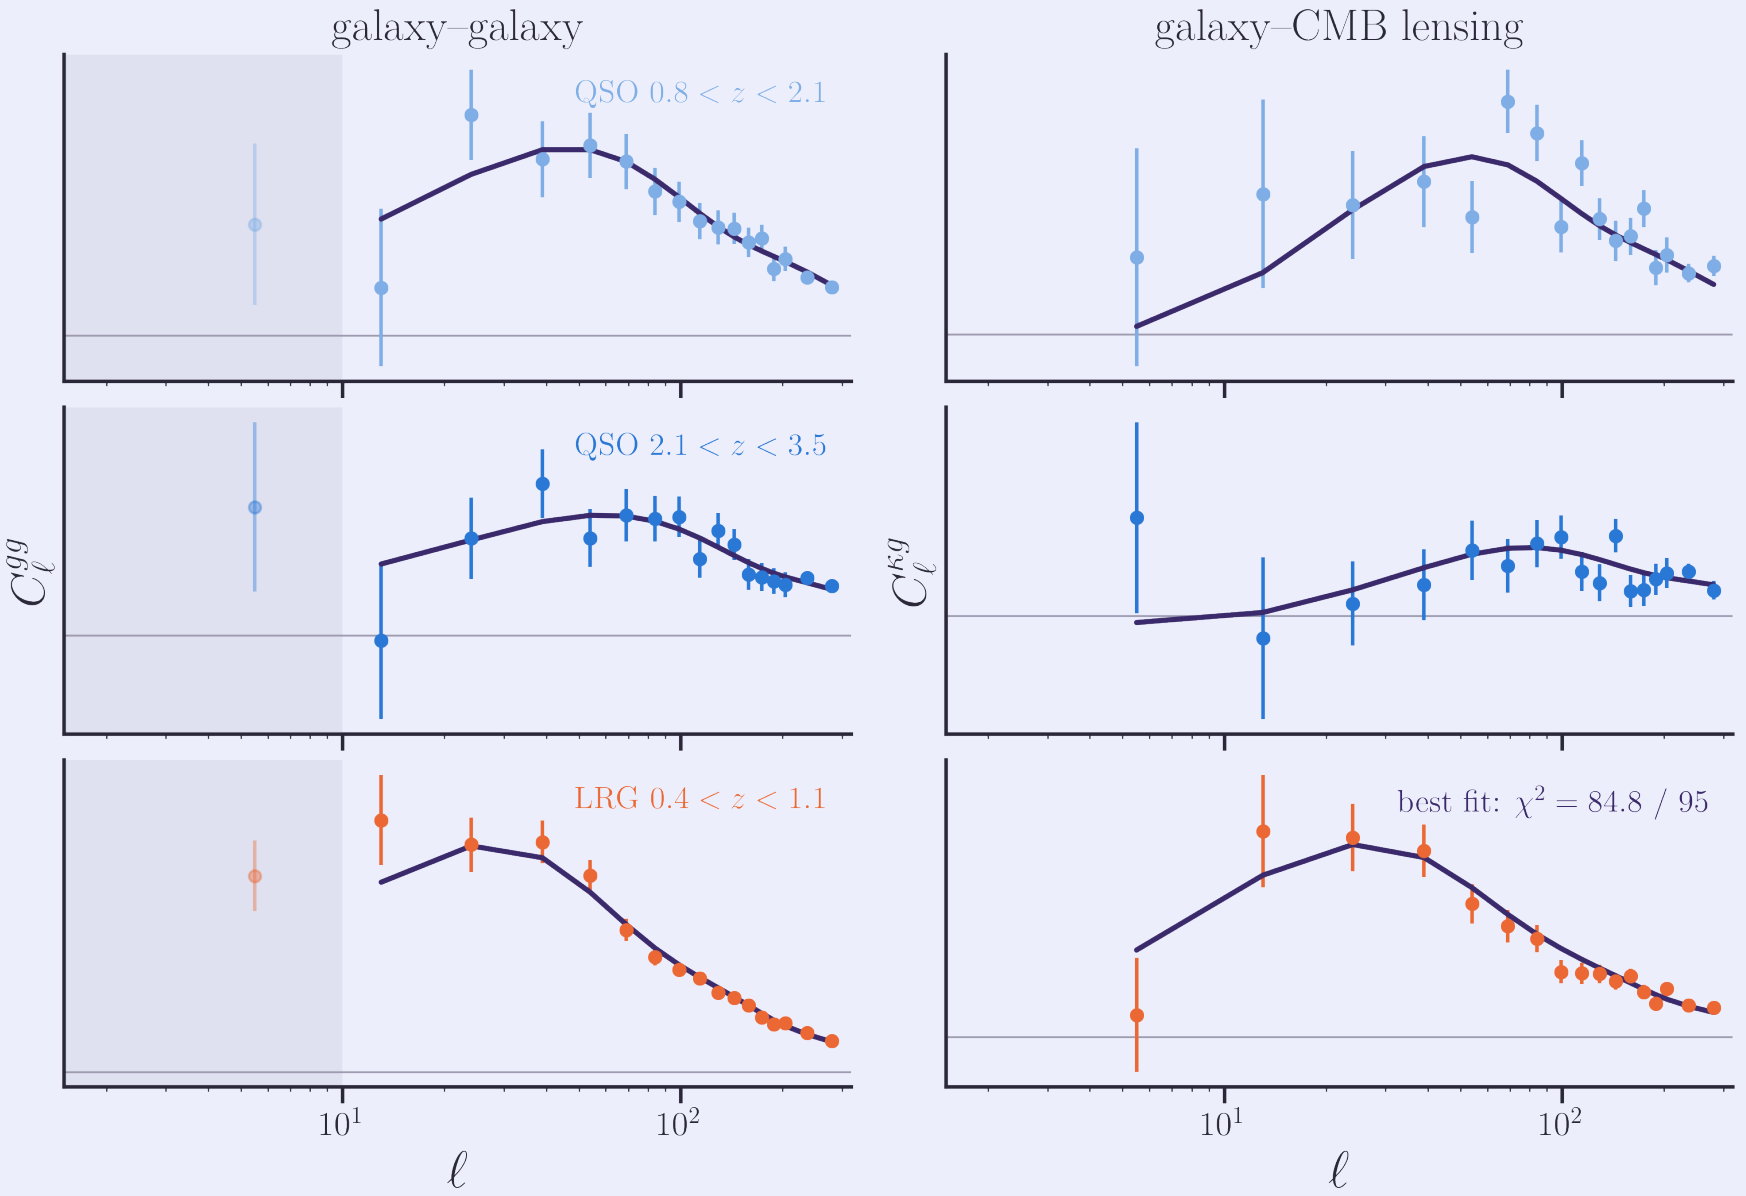
PNG & DESI

From DR1 to DR2



PNG & DESI

Measurements



DR2 setup

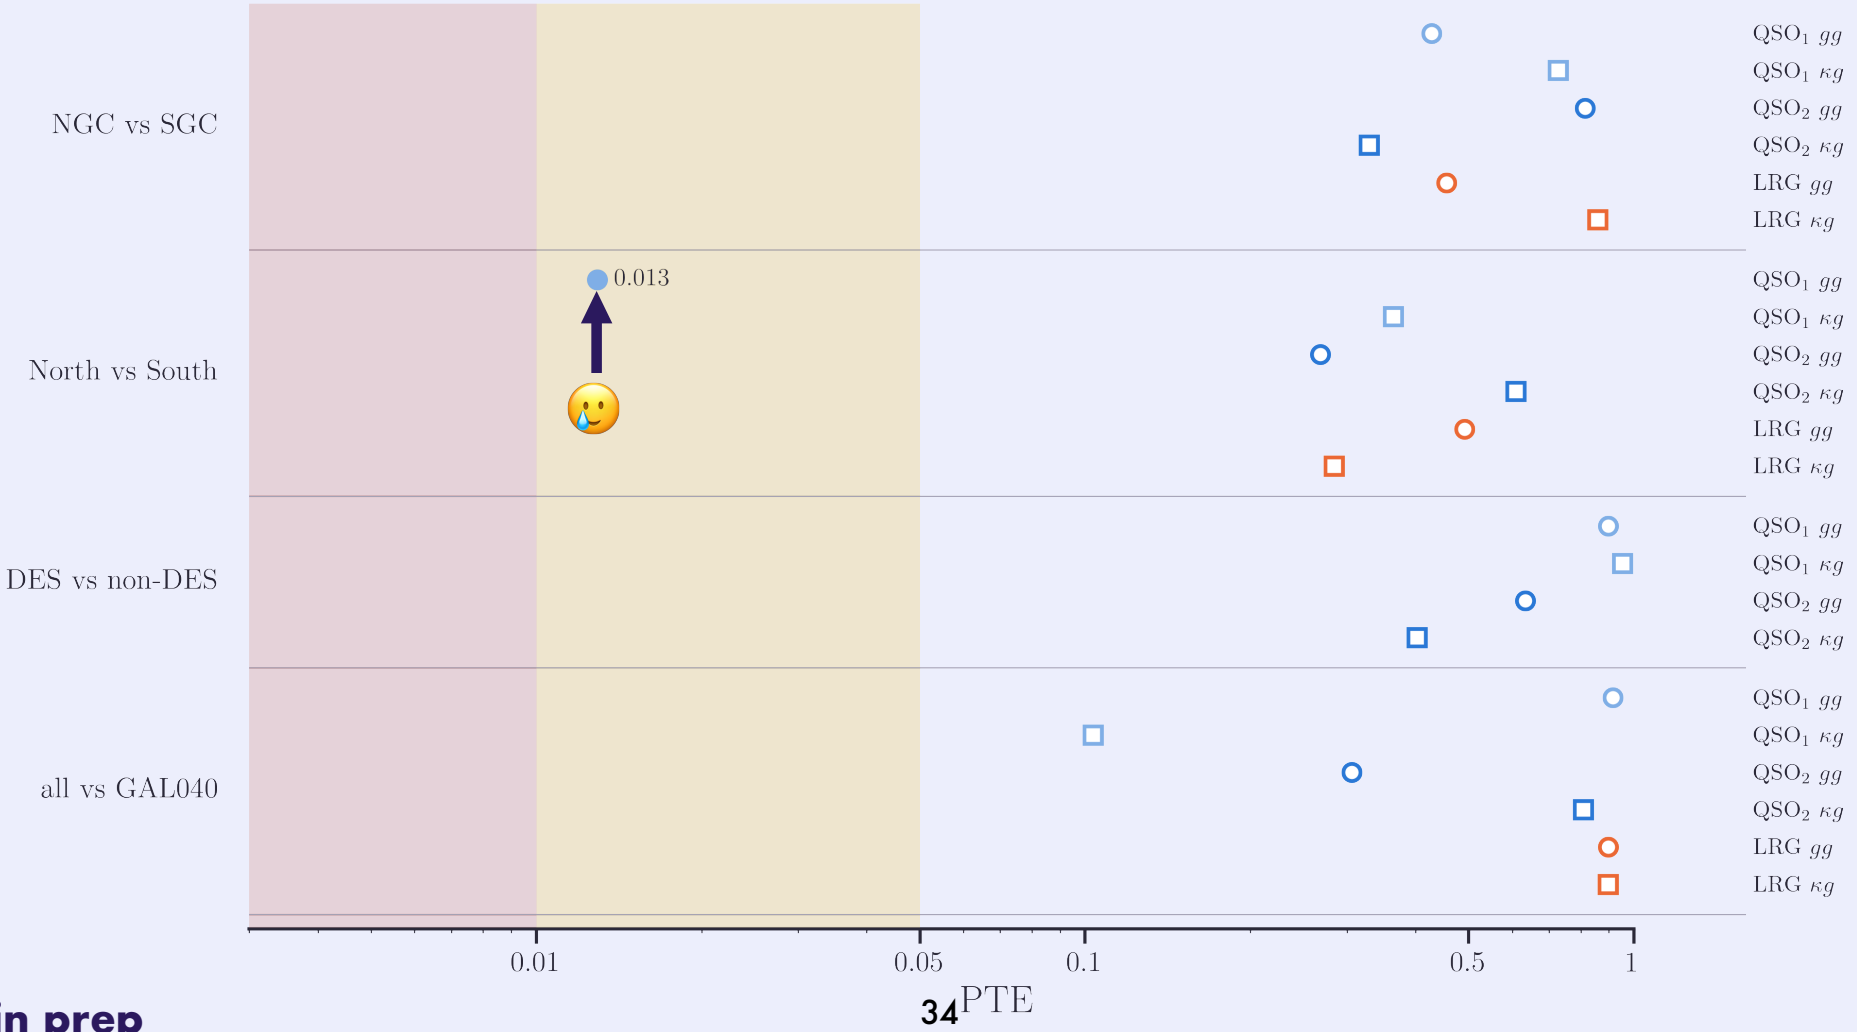
- 2 tracers, 3 bins
- Both C_ℓ^{gg} and $C_\ell^{\kappa g}$ are fitted
- Covariance from 1440 LogNormal Mocks
- Scale cut: $C_\ell^{gg}(\ell > 10)$

Chiarenza 2026b, in prep

PNG & DESI

Null tests

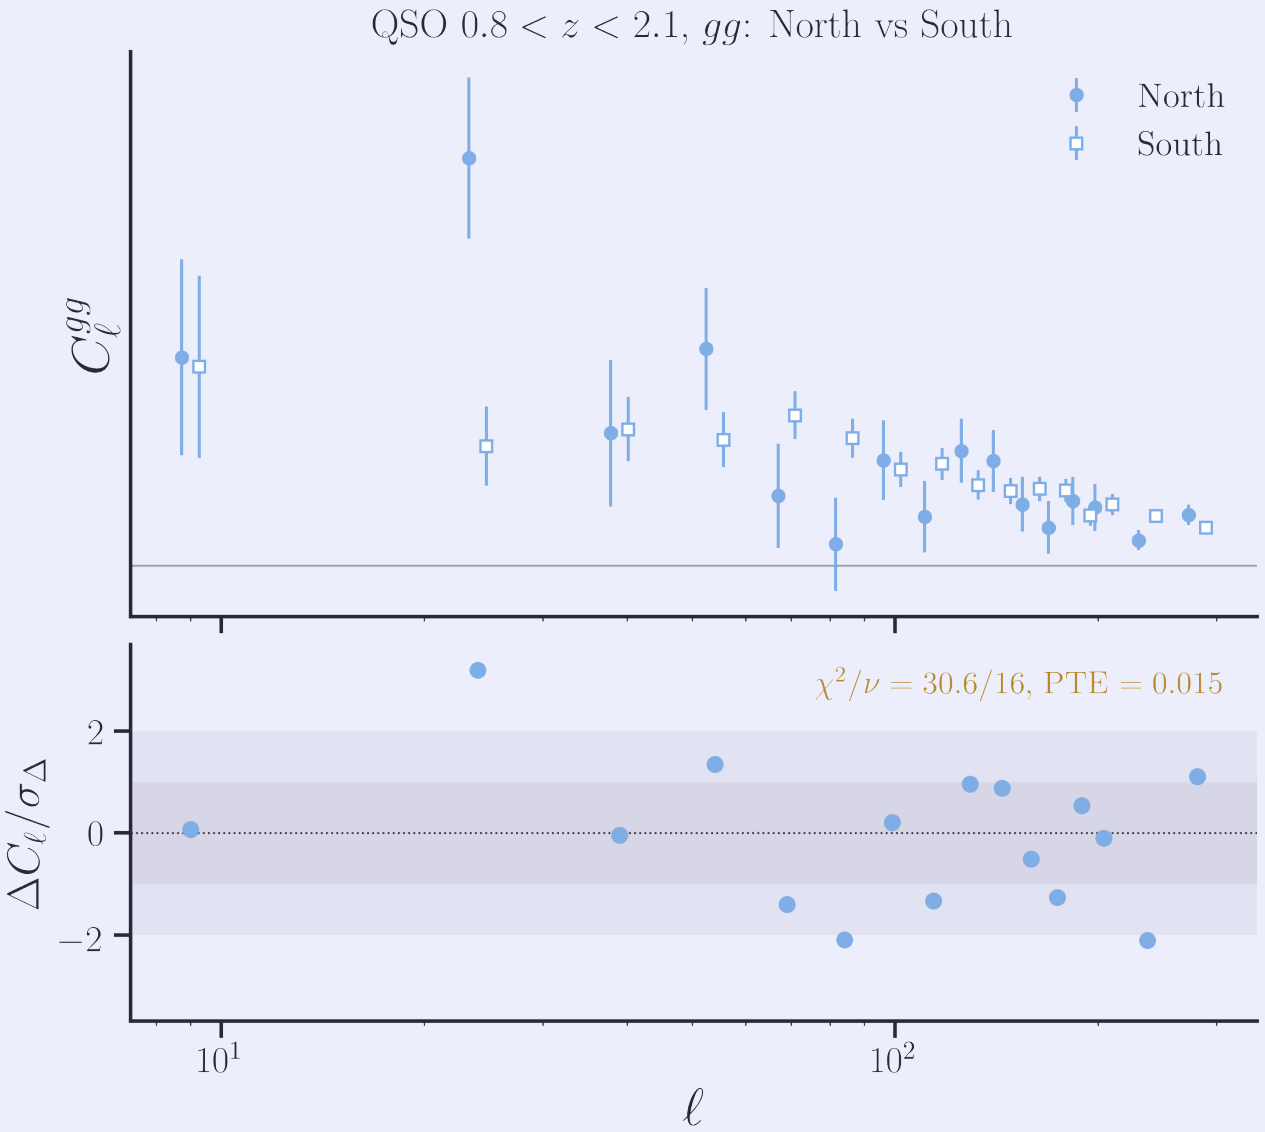
Checking for extra residuals systematics: **split the data** so the signal is identical but the potential systematics differ



Chiarenza 2026b, in prep

PNG & DESI

Null tests



Is the difference between the two regions consistent with what we expect?

$$\langle \Delta C_\ell \rangle = (R^a - R^b) \cdot C_\ell^{model}$$

$$Cov[\hat{C}^a - \hat{C}^b] \rightarrow \text{from mocks}$$

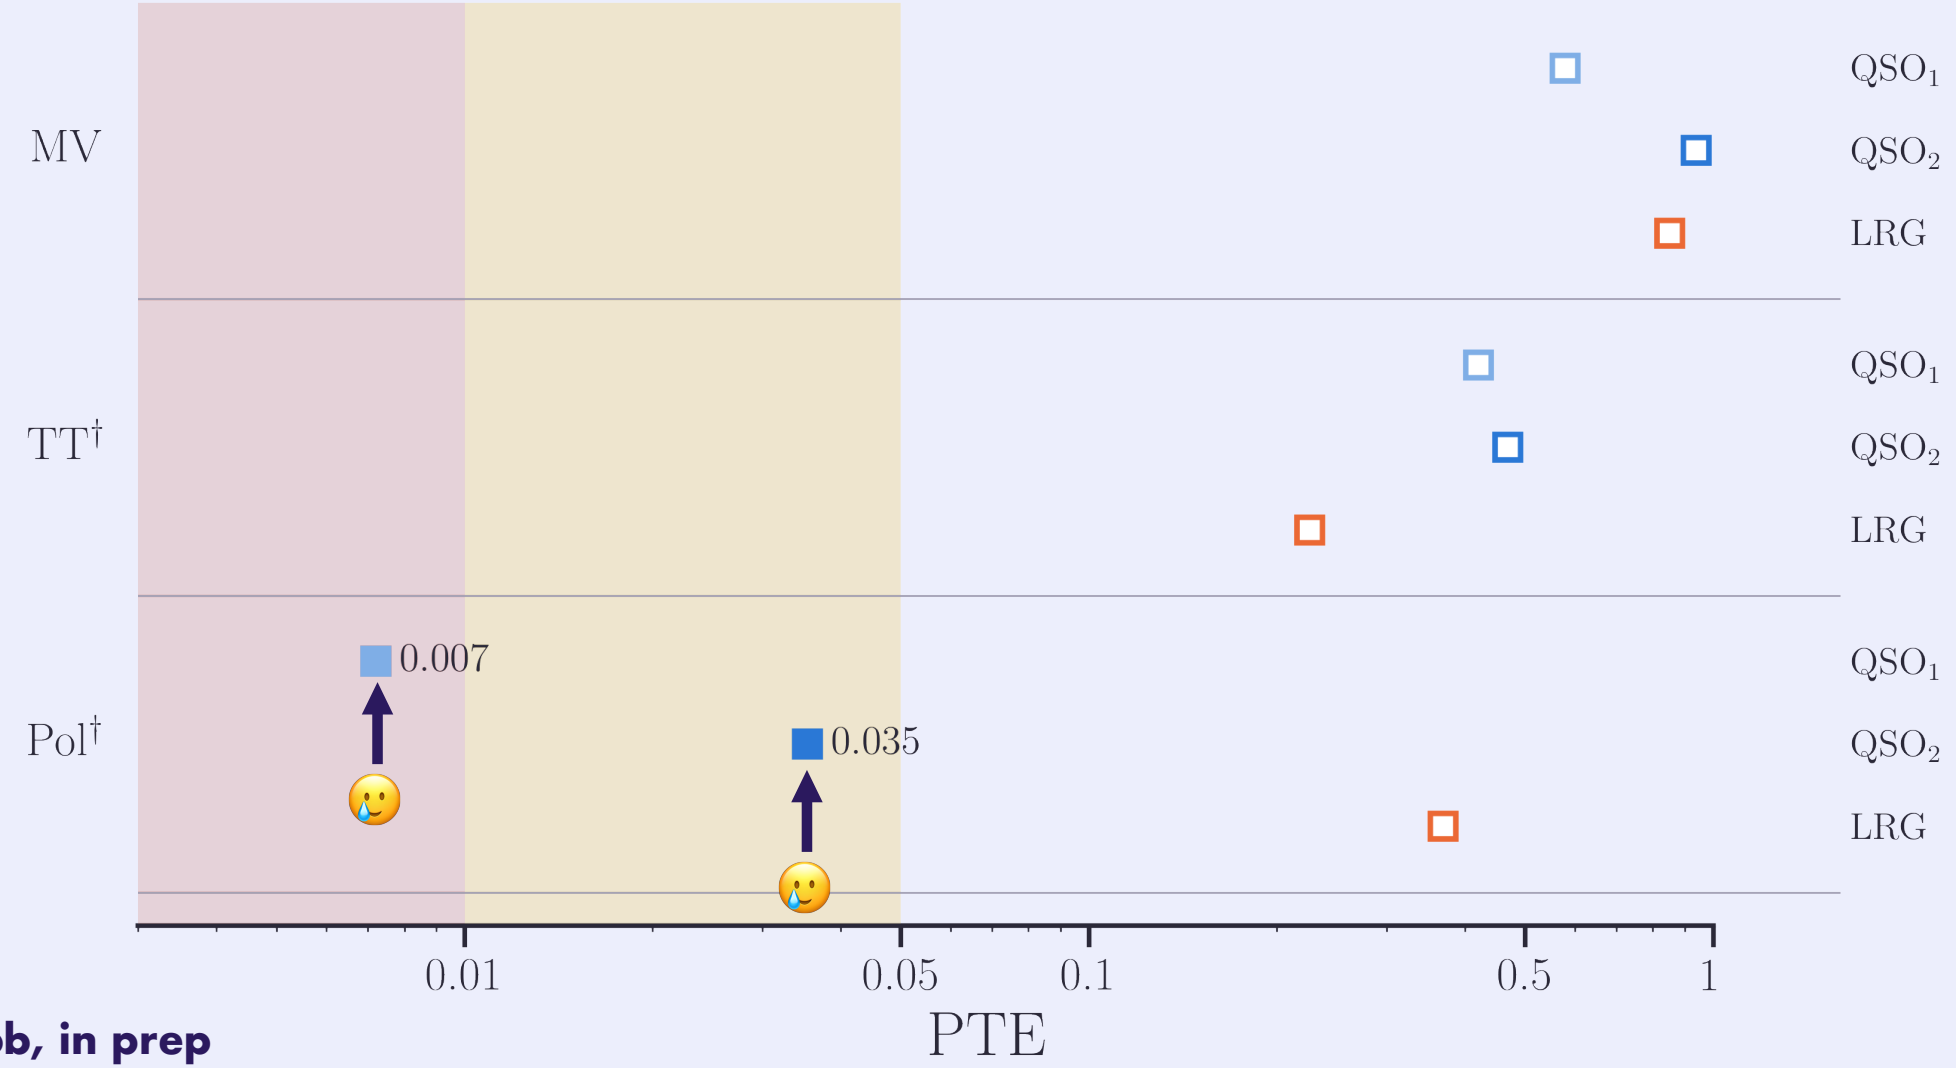
$$\chi^2 = (\Delta - \langle \Delta \rangle)^T C_\Delta^{-1} (\Delta - \langle \Delta \rangle)$$

Chiarenza 2026b, in prep

PNG & DESI

Null tests

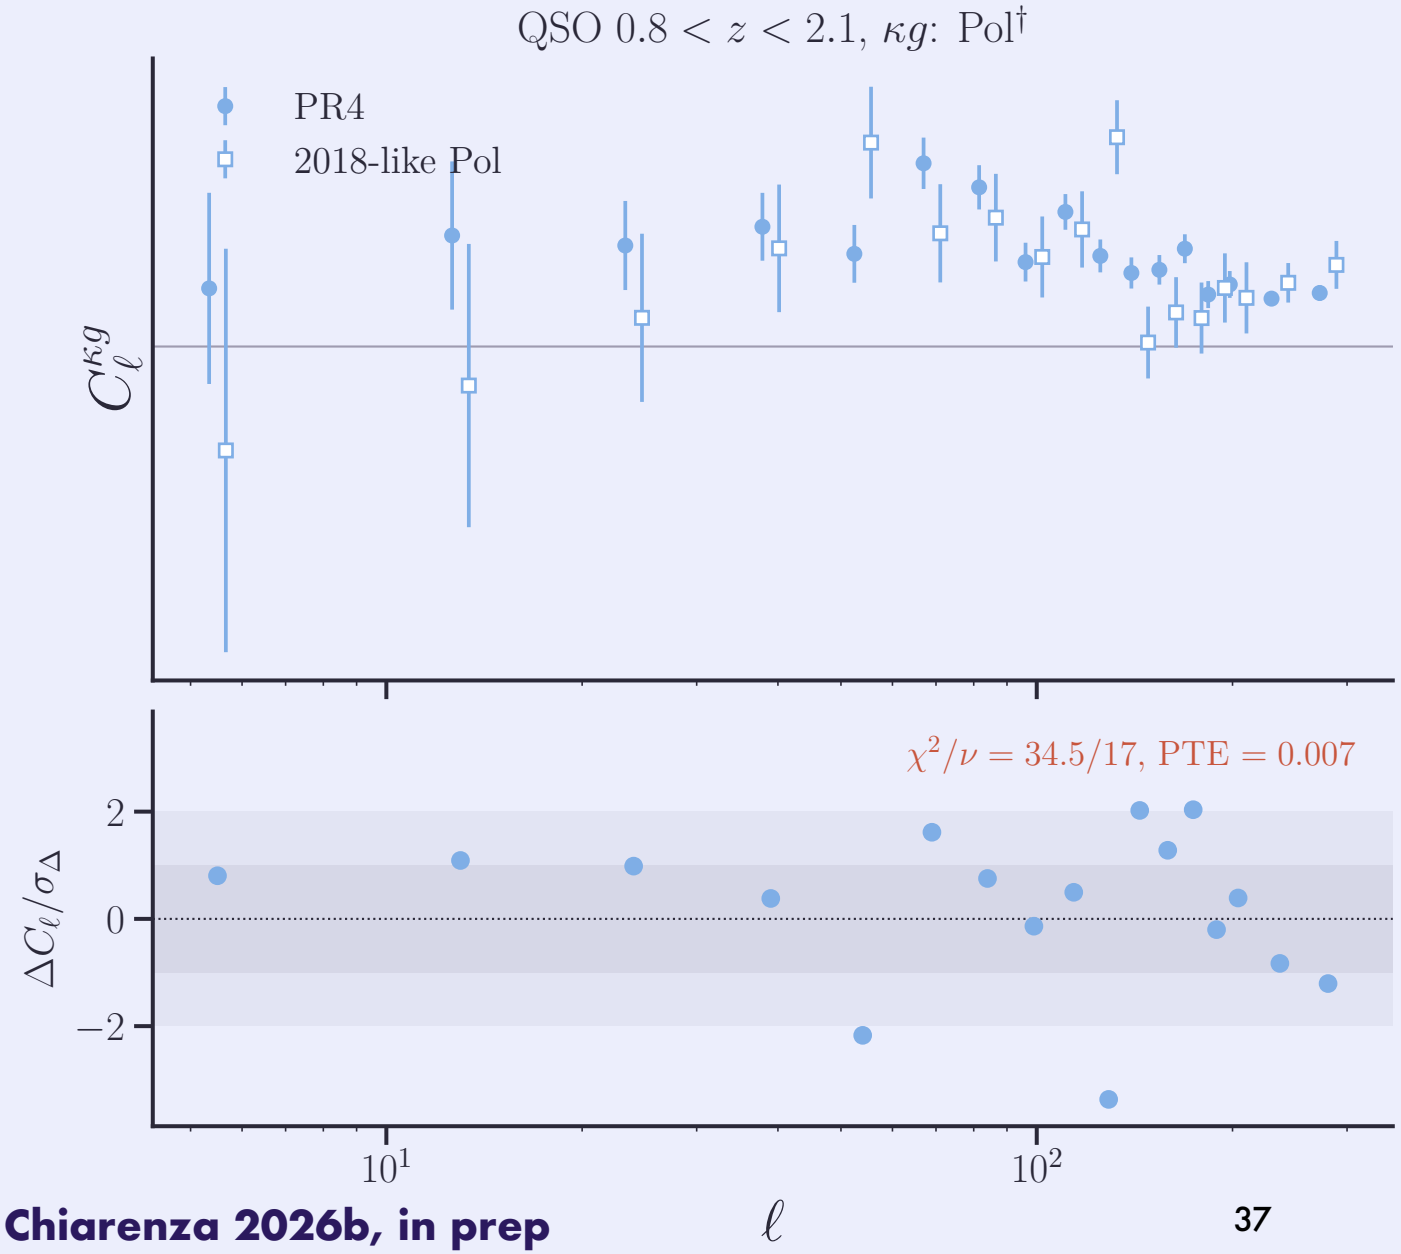
We also test different **lensing reconstructions**: this tests possible **foregrounds** in the lensing map (CIB, tSZ)



Chiarenza 2026b, in prep

PNG & DESI

Null tests



The different lensing reconstructions are strongly correlated

$$\mathbf{C}_\Delta = \mathbf{C}^{\mathrm{PR4,PR4}} + \mathbf{C}^{X,X} - 2 \cdot \mathbf{C}^{\mathrm{PR4},X}$$

$$\chi^2 = (\Delta - \langle\Delta\rangle)^T C_\Delta^{-1} (\Delta - \langle\Delta\rangle)$$

Chiarenza 2026b, in prep

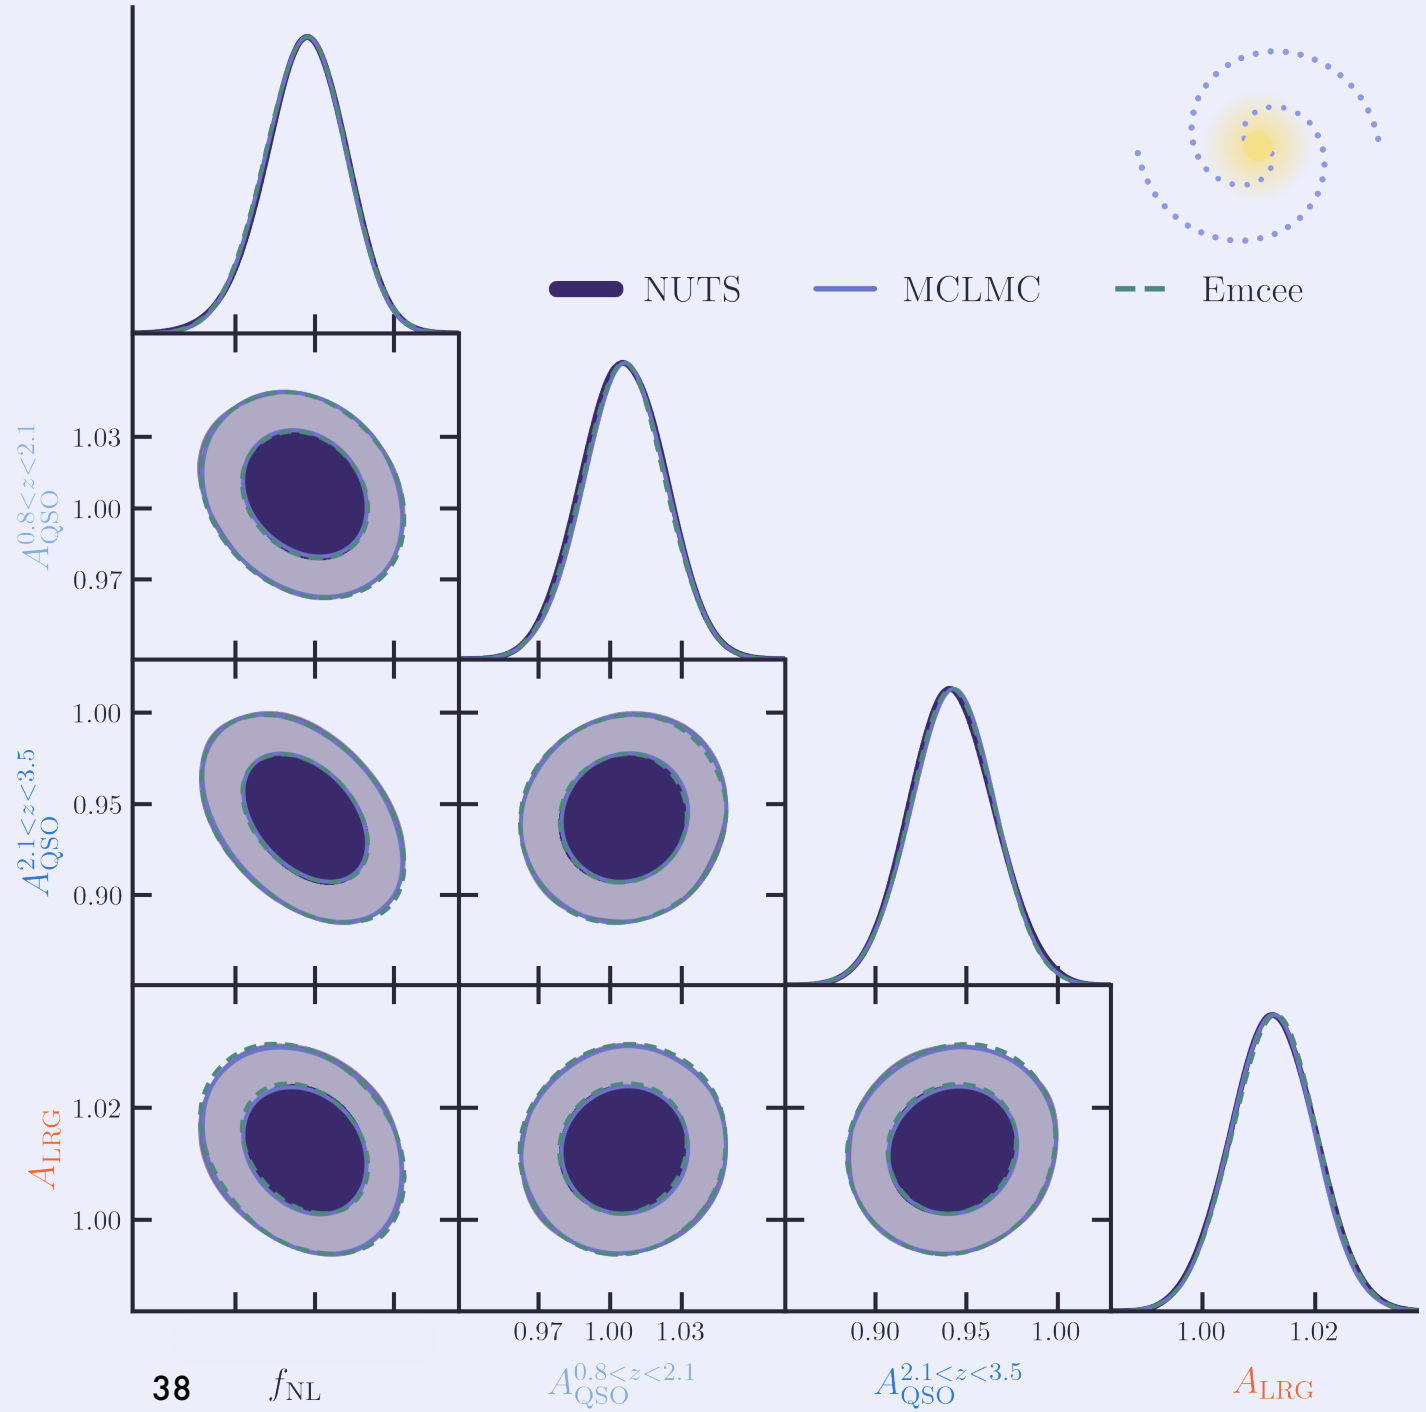
PNG & DESI

Results!!!

DR2 QSO+LRG x Planck PR4

$$C_\ell^{gg}(\ell > 10) + C_\ell^{\kappa g}(\ell > 2)$$

$$\sigma_{f_{\rm NL}} = 10$$



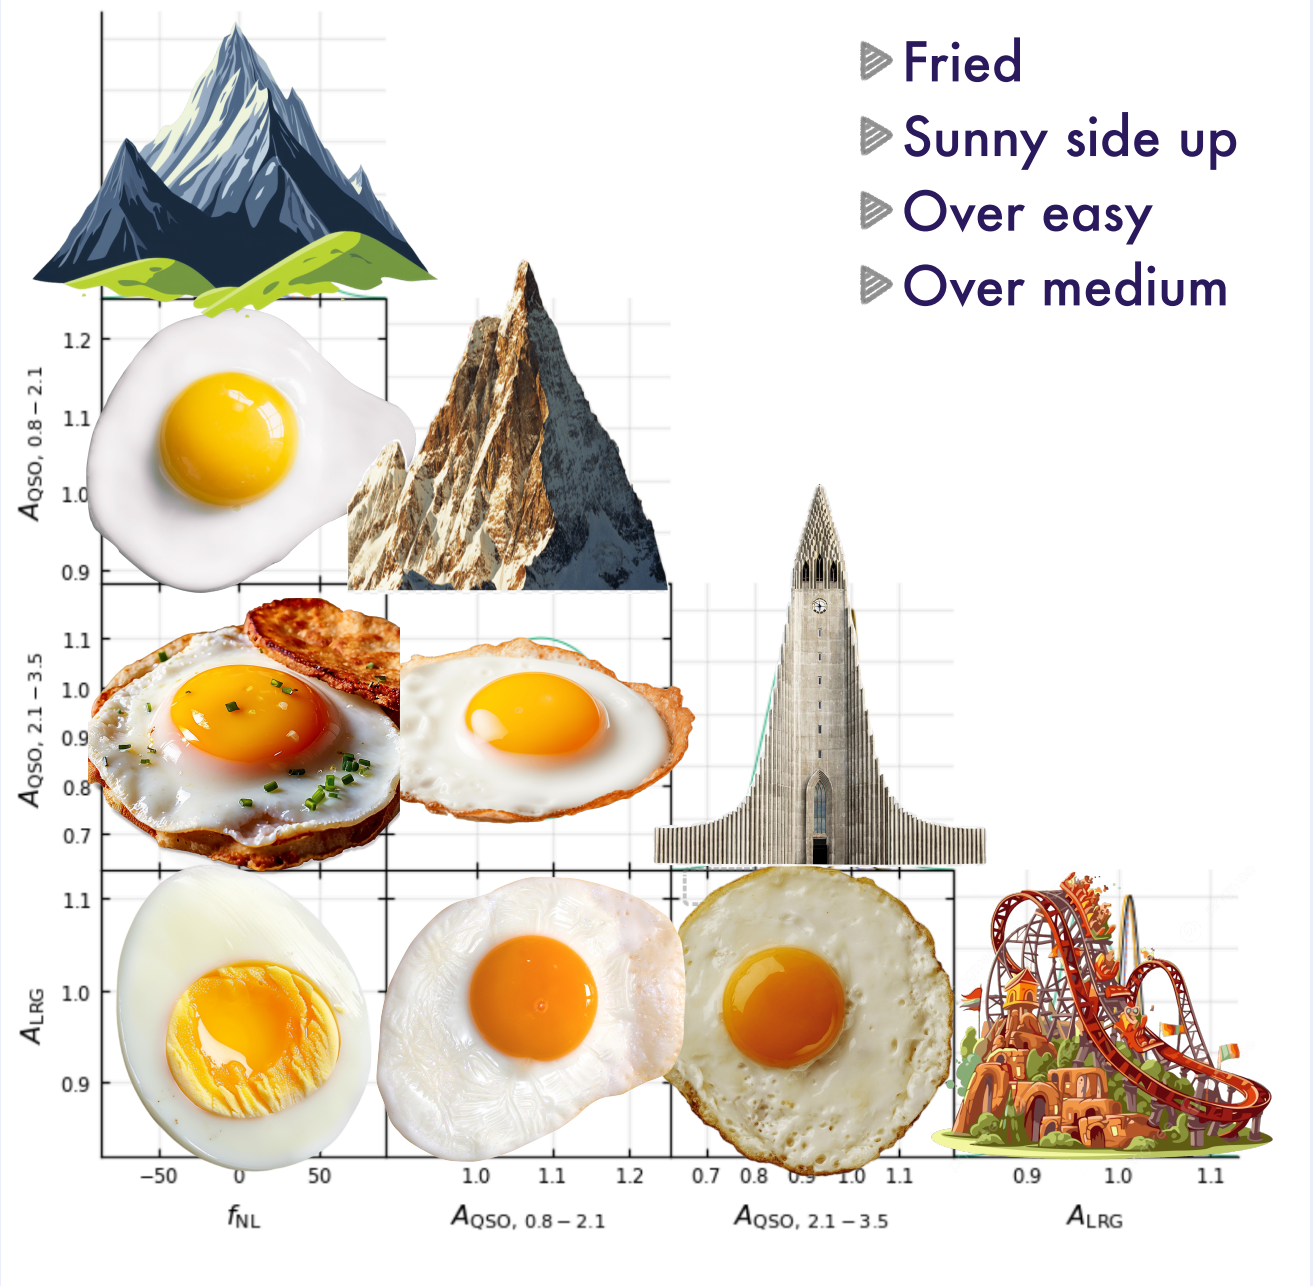
Chiarenza 2026b, in prep

PNG & DESI

Results!!!

$$P(k) + C_\ell^{\kappa g}(\ell > 2)$$

$$\sigma_{f_{\rm NL}} = 5$$



Chiarenza 2026b, in prep

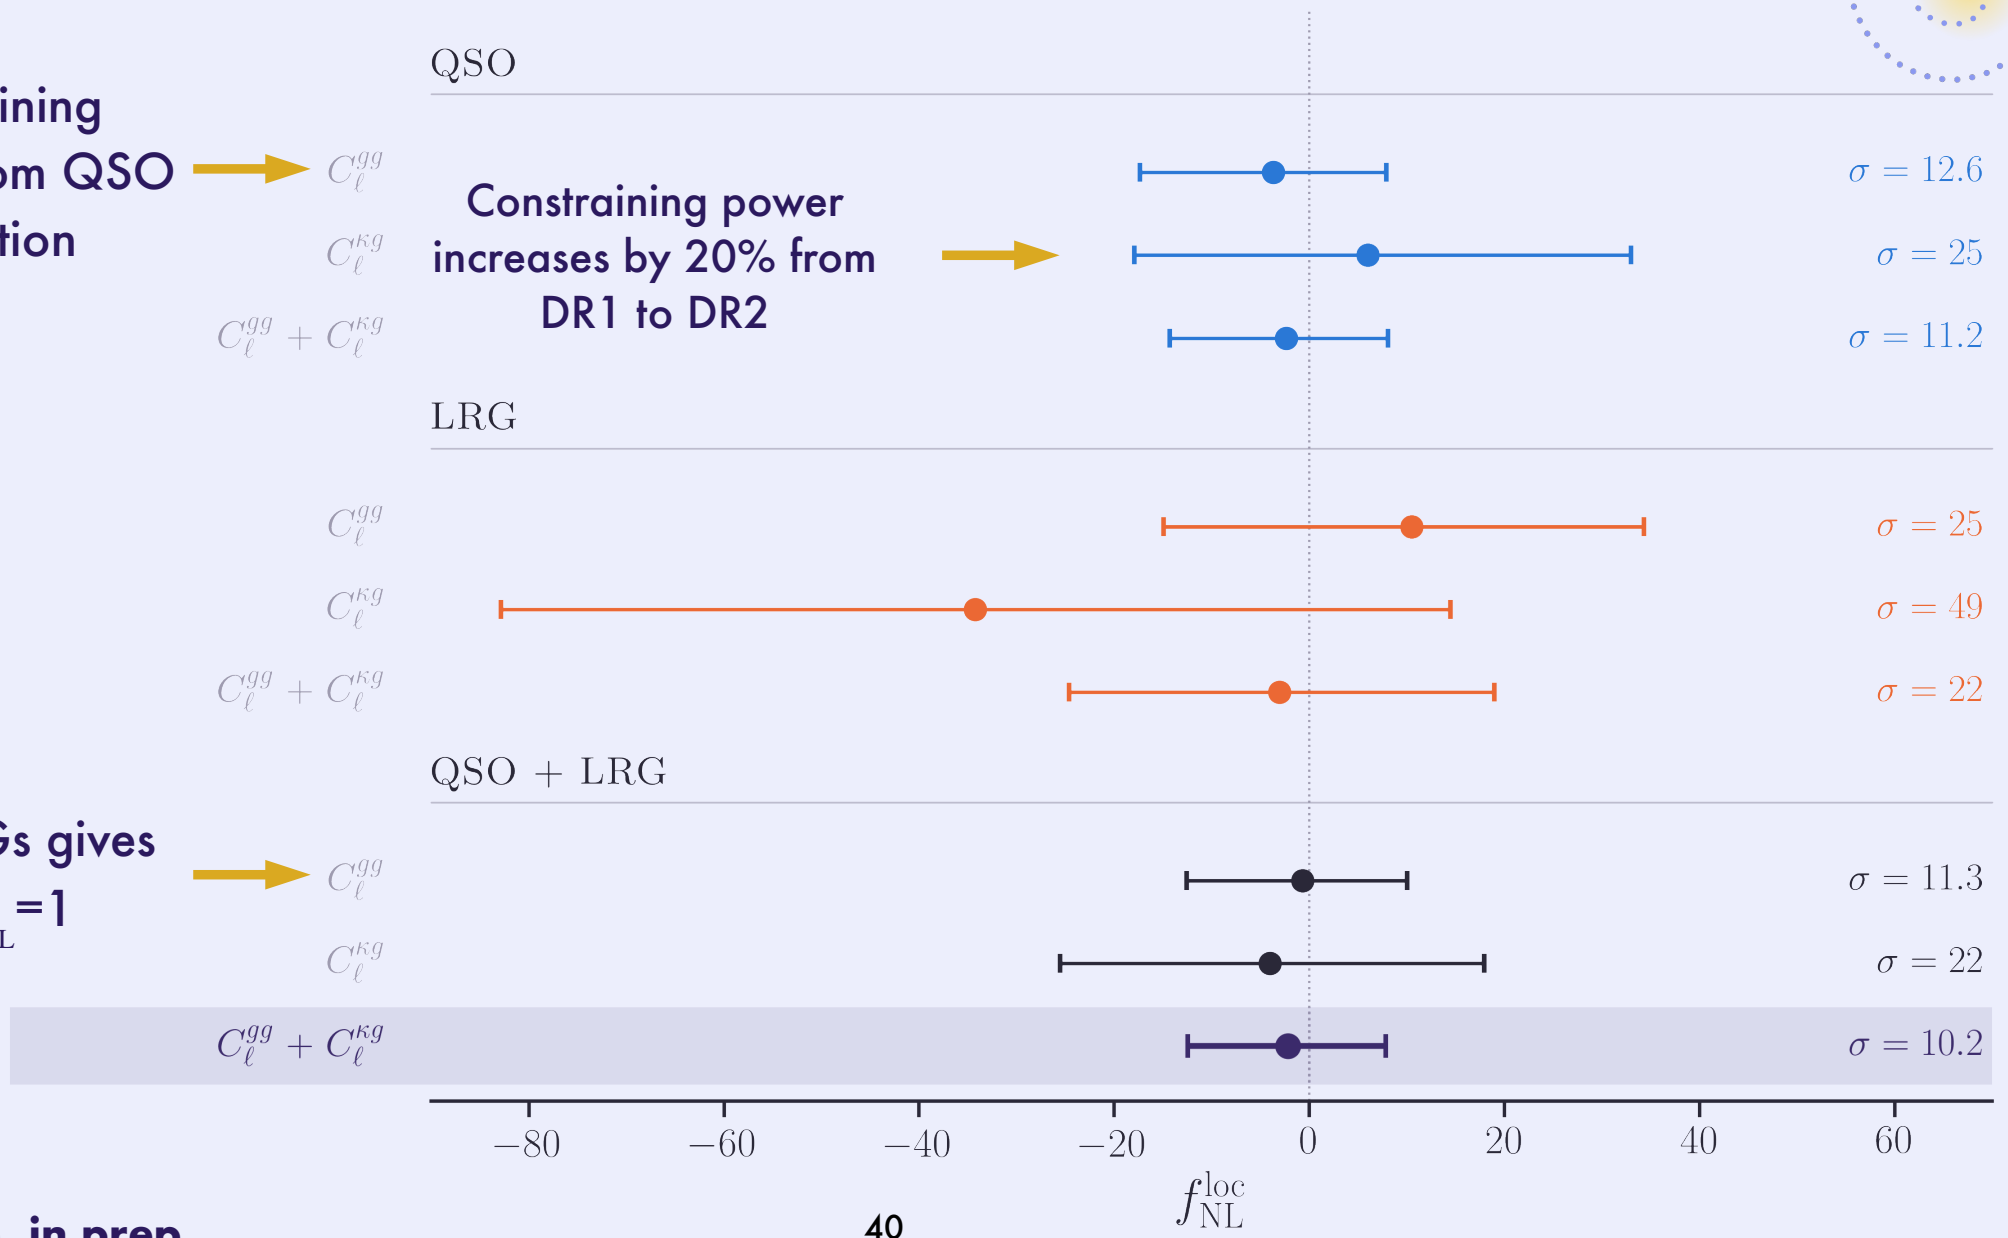
PNG & DESI

Analysis variations

Most constraining power comes from QSO autocorrelation

Constraining power increases by 20% from DR1 to DR2

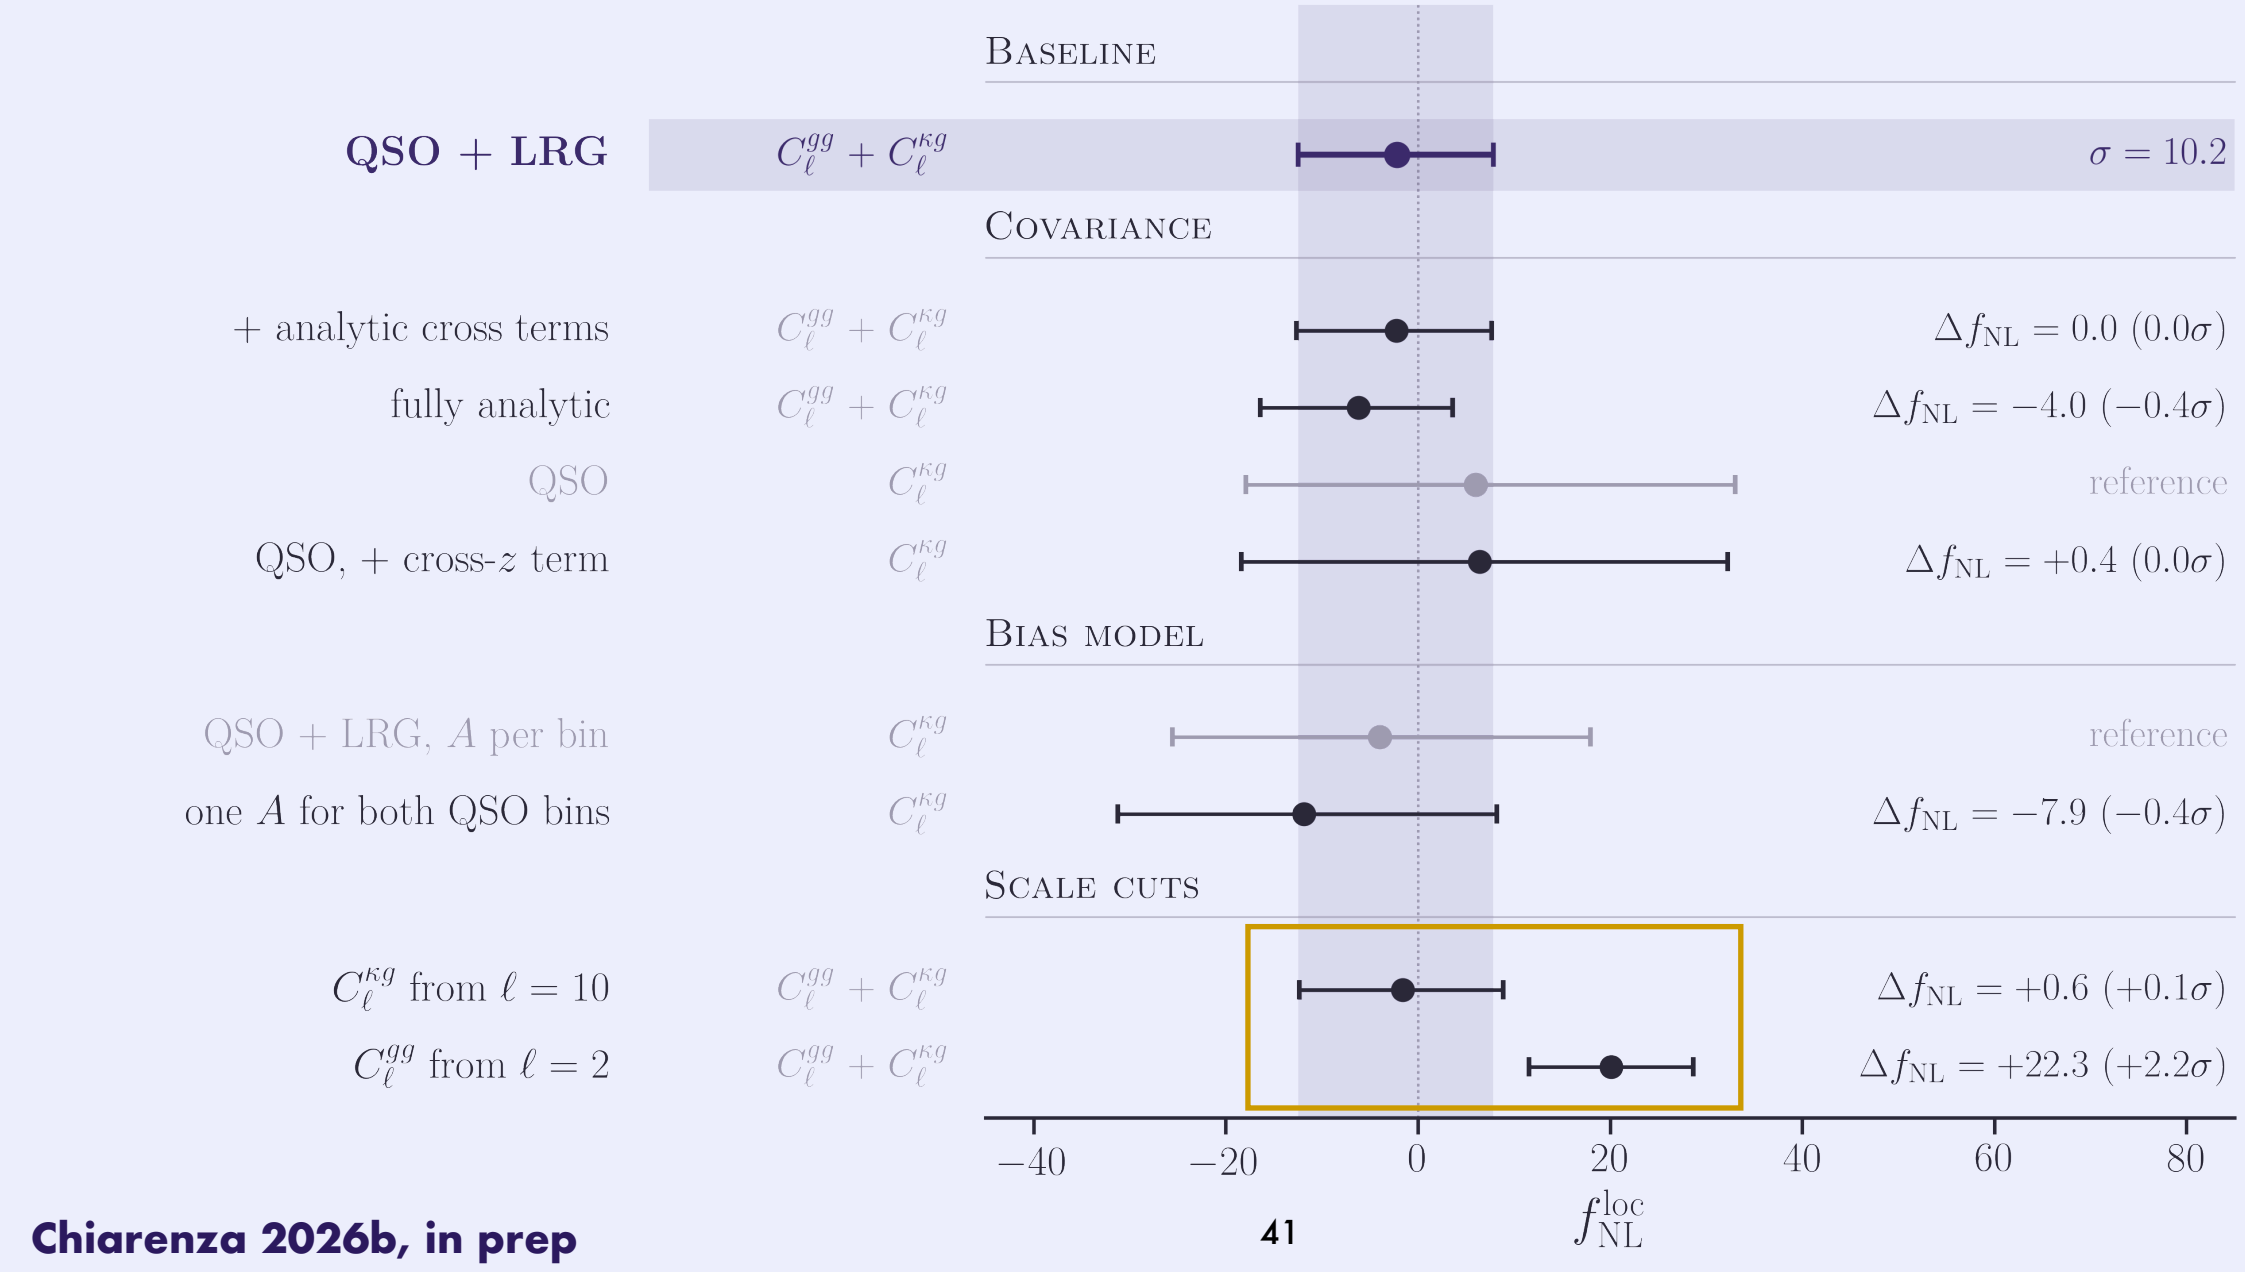
Adding the LRGs gives an extra $\sigma_{f_{\rm NL}}$=1



Chiarenza 2026b, in prep

PNG & DESI

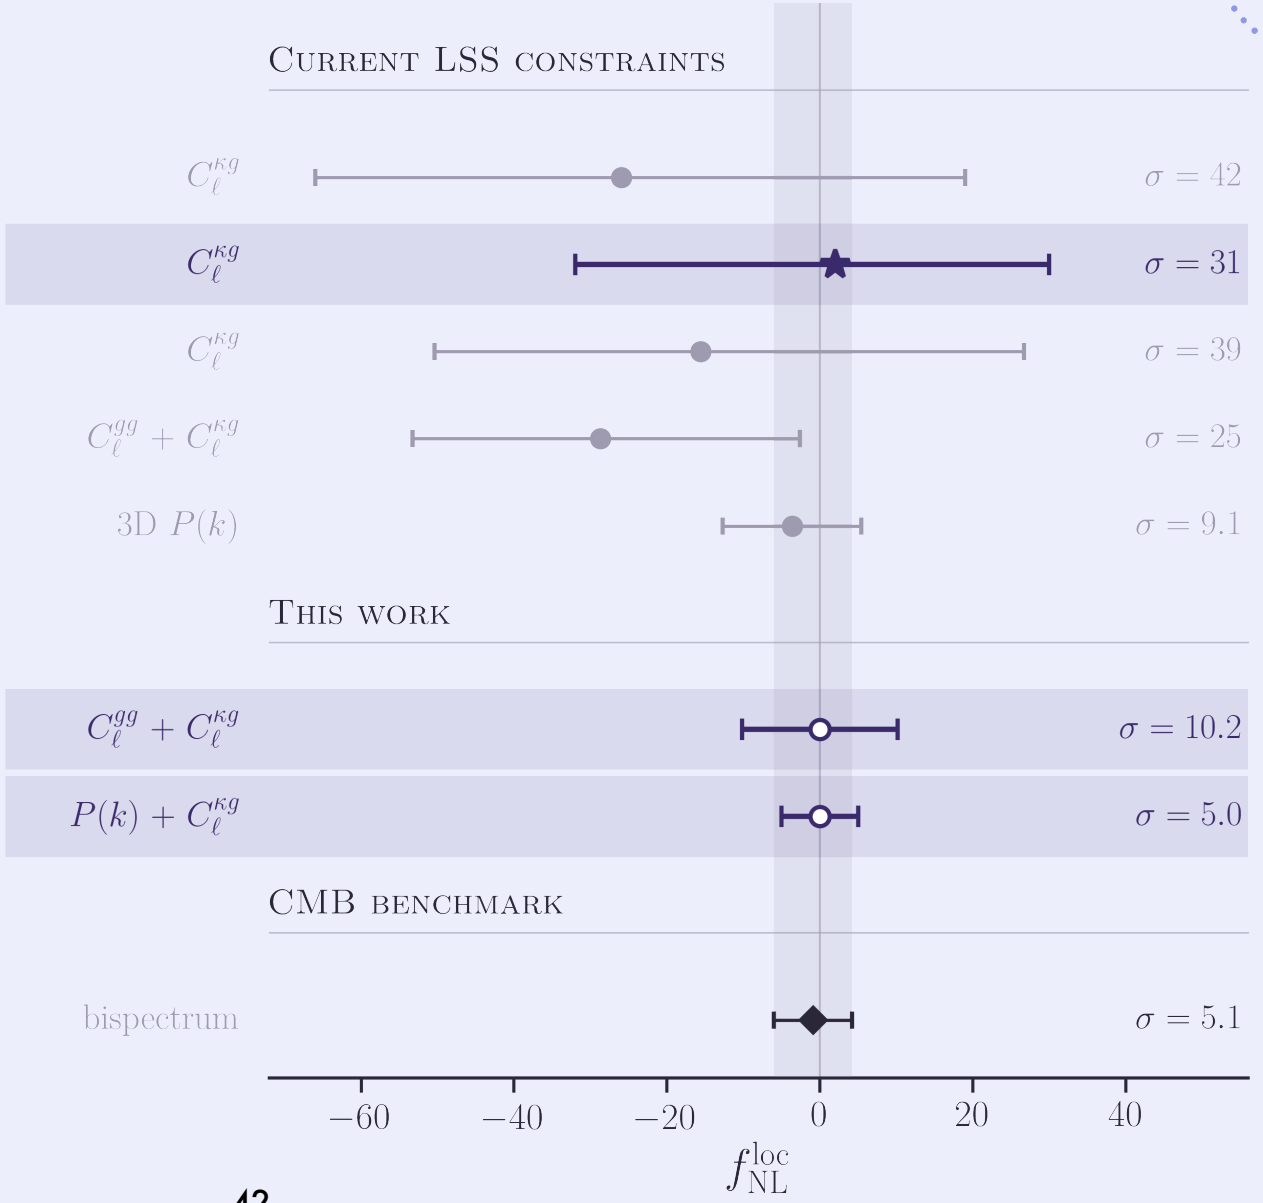
Robustness



Chiarenza 2026b, in prep

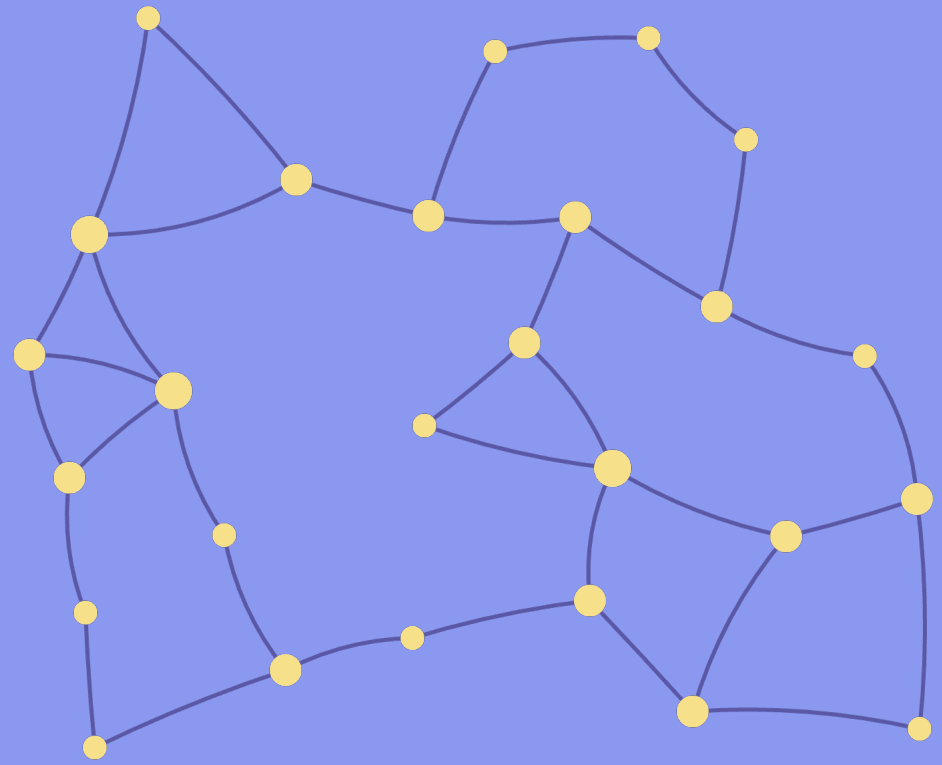
PNG & DESI

PNG constraint overview



Chiarenza 2026b, in prep

SECTION



04

Future Directions

QML estimator

Problem: NaMaster is not optimal, especially on large scales.

What QML does: the QML estimator (Tegmark, 1997) is an inverse-covariance-weighted quadratic estimator: unbiased and minimum-variance, it reaches the Fisher limit.

$$y_\ell = \mathbf{x}^T \mathbf{E}_\ell \mathbf{x} - b_\ell, \qquad \mathbf{E}_\ell = \frac{1}{2} \mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial C_\ell} \mathbf{C}^{-1}$$

$$\hat{C}_\ell = \sum_{\ell'} \left(F^{-1}\right)_{\ell\ell'} y_{\ell'}, \qquad F_{\ell\ell'} = \frac{1}{2} \mathrm{Tr}\left[\mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial C_\ell} \mathbf{C}^{-1} \frac{\partial \mathbf{C}}{\partial C_{\ell'}}\right]$$

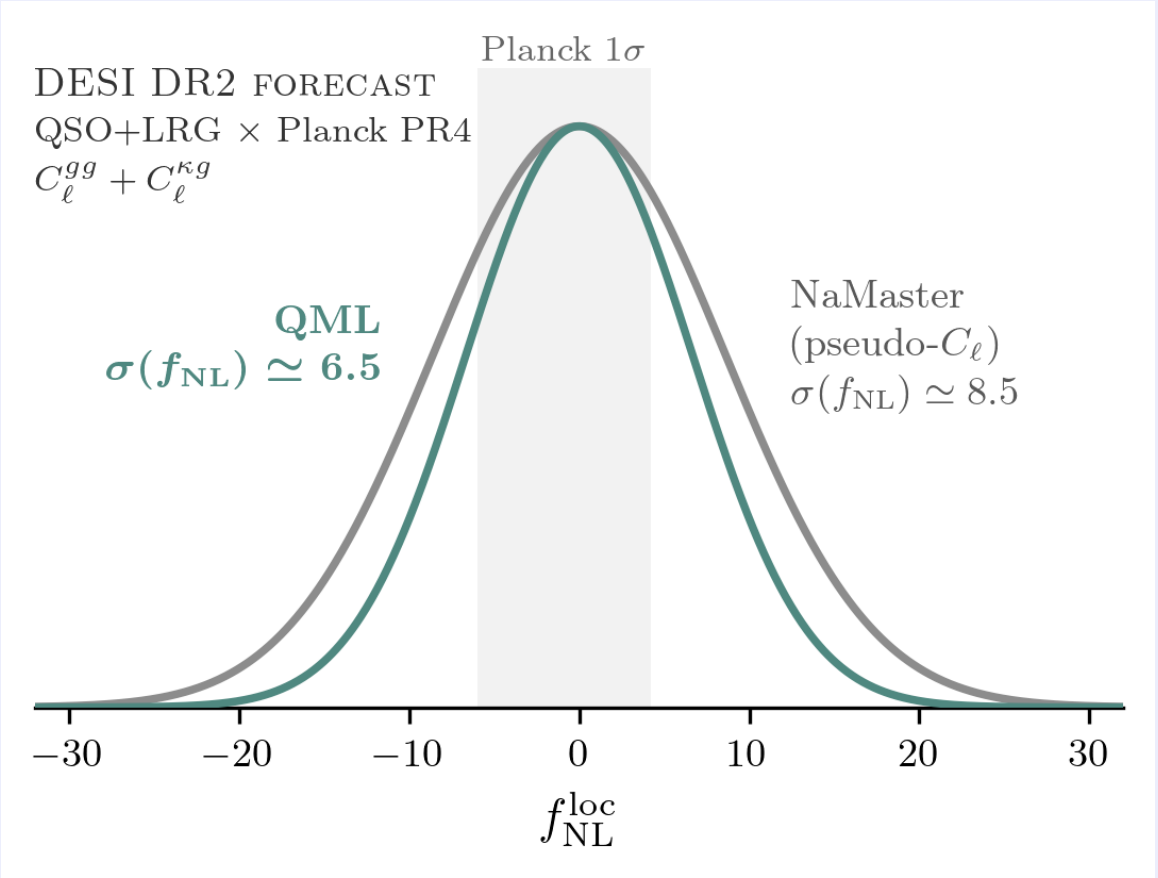
FUTURE DIRECTIONS

QML estimator

The catch: it work in pixel space, scales as O(N_pix³). That's only feasible at low resolution.

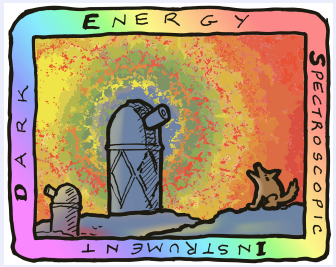
The fix: many optimized codes are being developed (Kvasiuk 2026, Galloni 2026)

The prize: Same data, tighter $f_{\rm NL}$: the gain comes for free from the largest scales.



FUTURE DIRECTIONS

New data, analyzed together



DESI DR3

- $14{,}000\ \mathrm{deg}^2$
- DR3 x CMB lensing, final year PNG analysis

Euclid

- Leading CMBX paper
- Euclid spec sample x (ACT , Planck)
- Measure $w_0 - w_a$, $\sigma_8(z)$

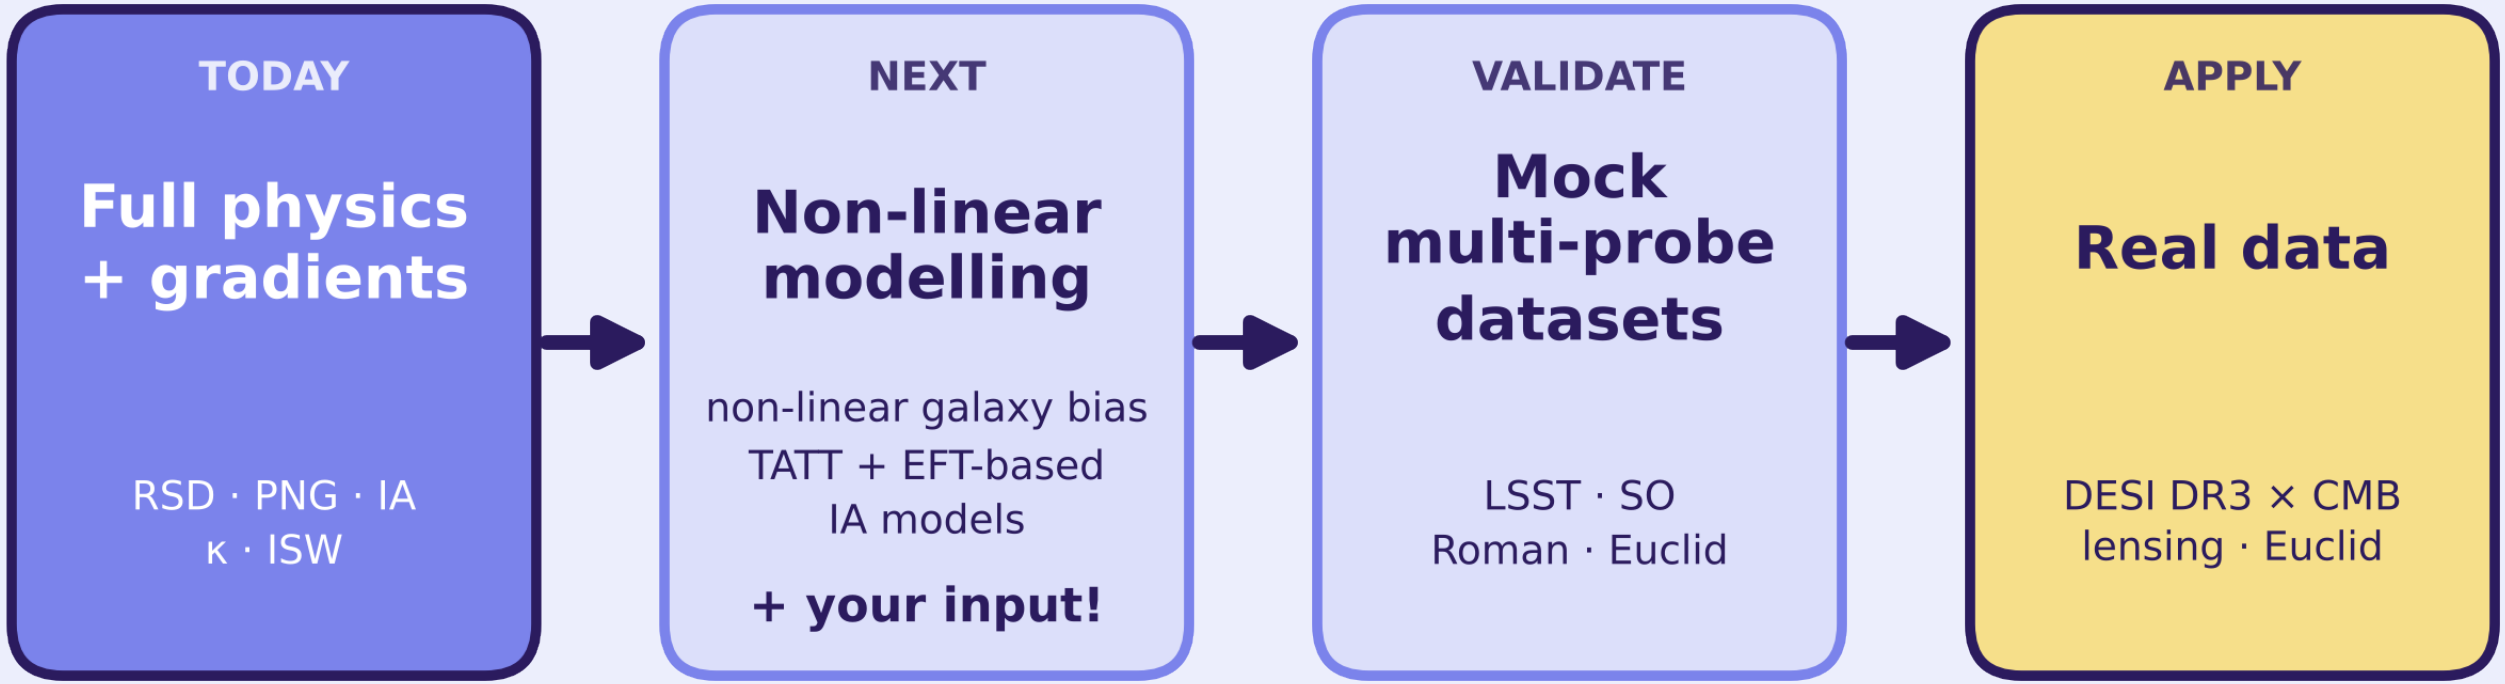
Simons Observatory

- Lower-noise and cleaner CMB lensing maps will sharpen every CMBX measurement

A lot of joint analysis to be done!

FUTURE DIRECTIONS

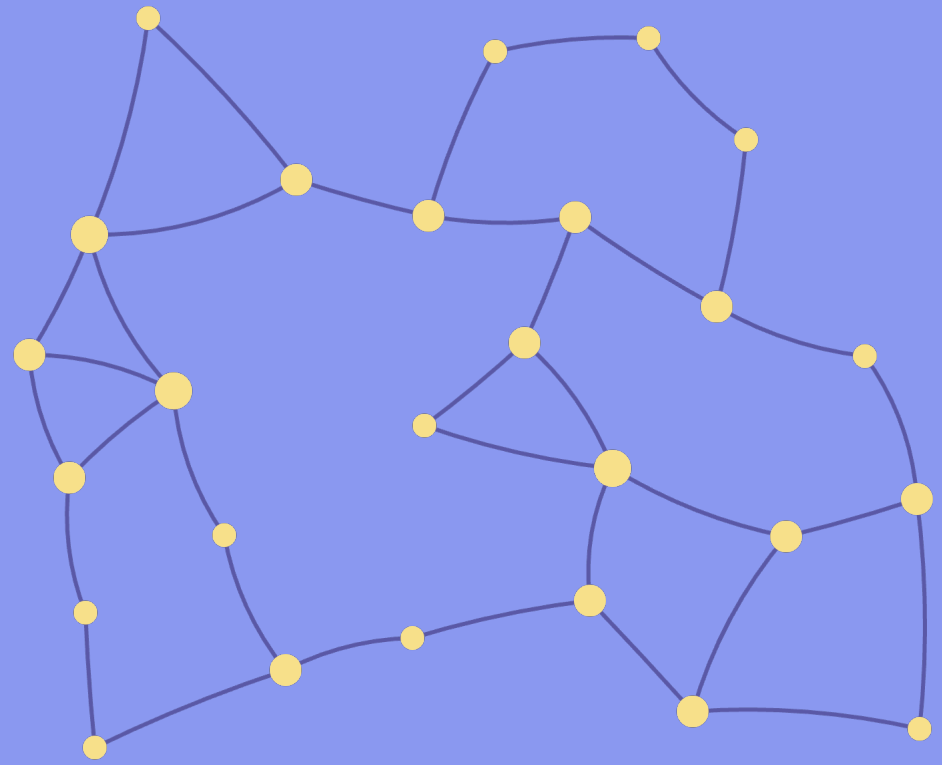
Further Blast Expansions



https://github.com/sofiachiarenza/Blast.jl

SECTION

05 Conclusions



CONCLUSIONS

In summary...

BUILD: Blast.jl

- Exact and full effect non-Limber C_ℓ, 10× faster than existing algorithms and fully differentiable
- Gradient-based samplers bring a 35-parameter inference down to hours instead of weeks

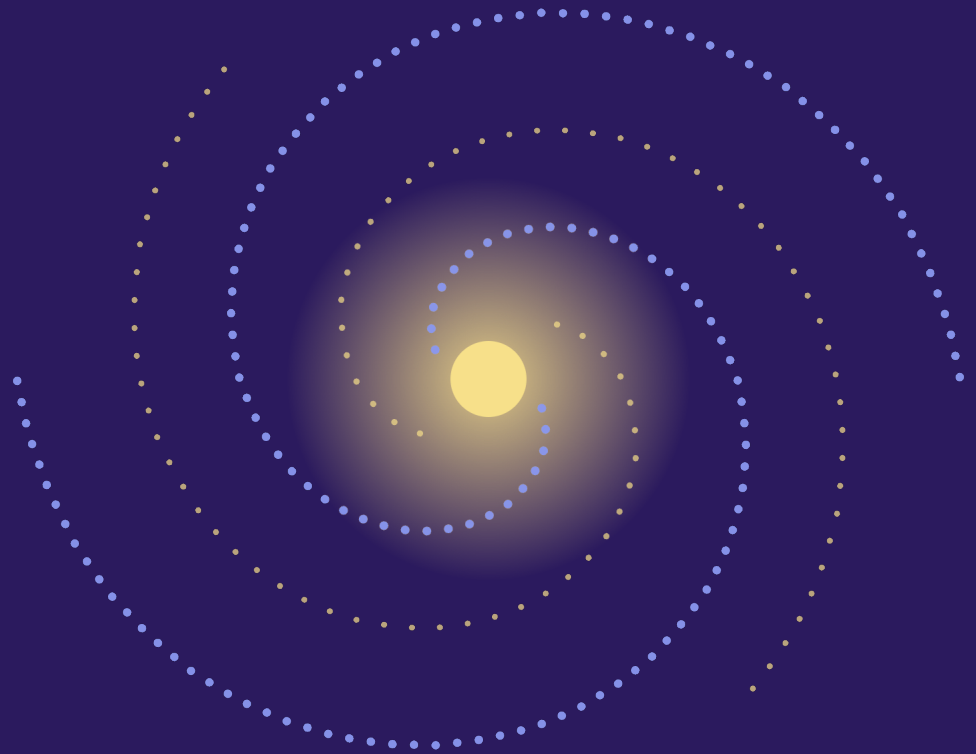
APPLY: PNG from DESI DR2 × Planck PR4

- $f_{\rm NL} = -2.2 \pm 10$ from 2D QSO + LRG
- Multi-tracer, multi-probe analysis, validated with null tests and 10+ robustness variants

COMING SOON

- Optimal large scales: NaMaster vs QML vs field level
- DESI DR3 × CMB lensing
- Euclid spec-z × CMB lensing for $\sigma_8(z)$ and $w_0 - w_a$

QUESTIONS



Thank you!