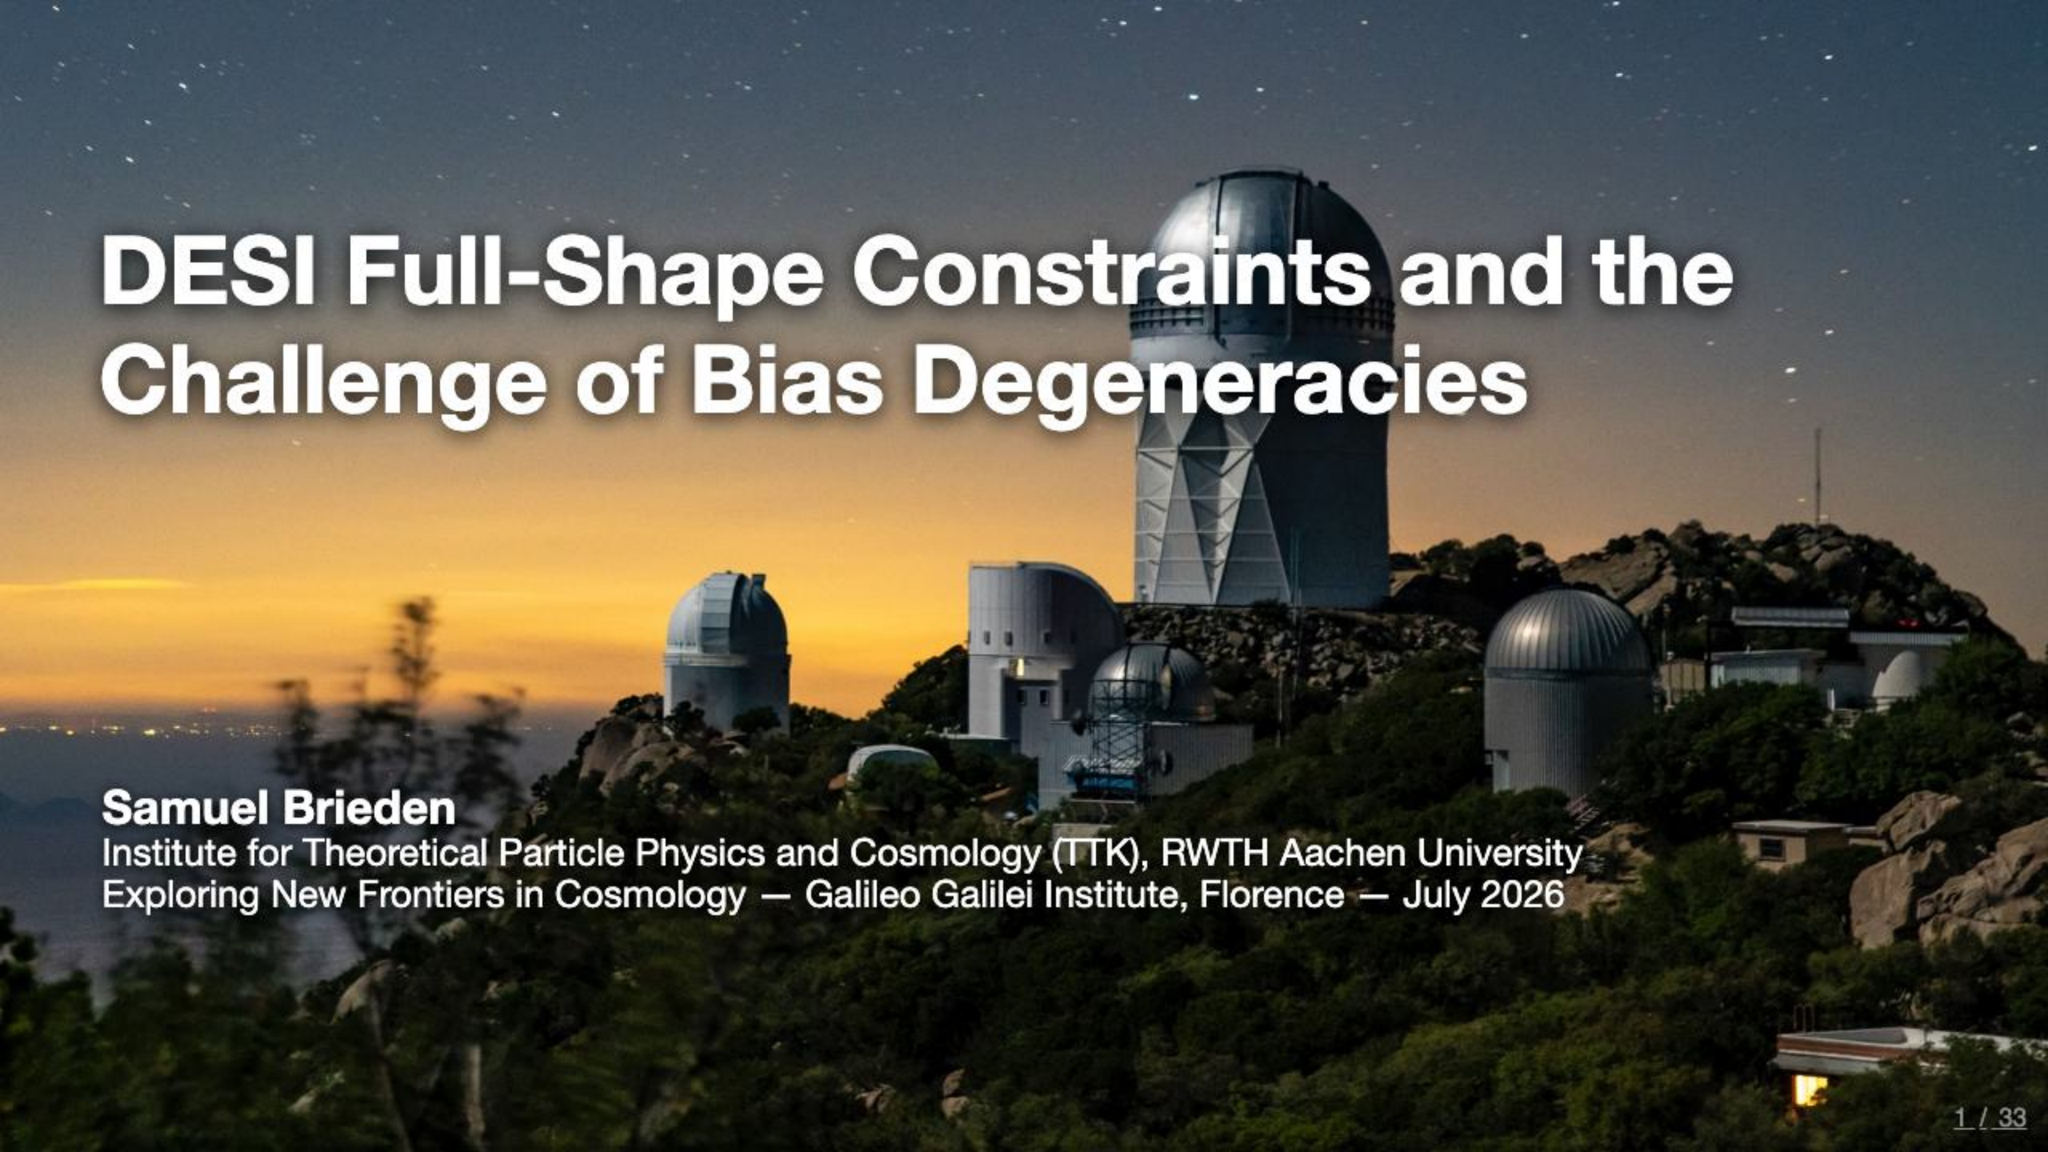

DESI Full-Shape Constraints and the Challenge of Bias Degeneracies

Samuel Brieden
Institute for Theoretical Particle Physics and Cosmology (TTK), RWTH Aachen University
Exploring New Frontiers in Cosmology — Galileo Galilei Institute, Florence — July 2026

Outline

- I. Where we stand — full shape & ShapeFit; DR1×DR2 = "DESI1.5": dark energy, neutrinos, curvature — and growth at $z < 0.1$ from peculiar velocities
- II. The challenge — ~8 nuisance parameters per tracer: degeneracies, prior volume, projection effects
- III. A physical way out — teaser: bias from the cosmic web (→ full story in my week-4 talk)

Dark Energy Spectroscopic Instrument (DESI)

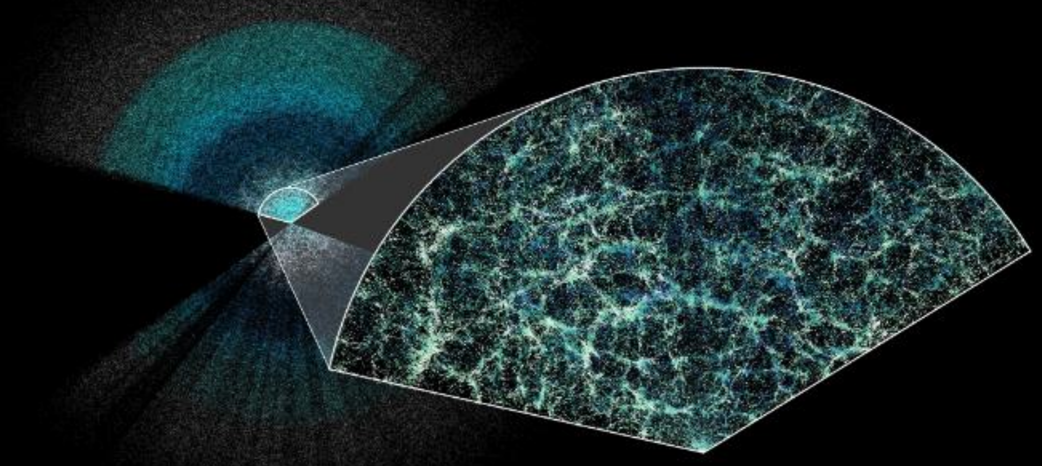

- Mayall 4-m telescope, Kitt Peak — **5000 robotic fibres**
- DR1 (2024): >6M redshifts for cosmology → BAO + first full-shape analysis
- DR2 (2025): **>14M** galaxy & quasar redshifts + Lyα forest of 0.8M QSOs (3 yr) → most precise BAO to date
- 6 (7) tracer samples: BGS, LRG1–3, ELG, QSO (+ Lyα)

slice: C. Lamman / DESI Collaboration

What we measure

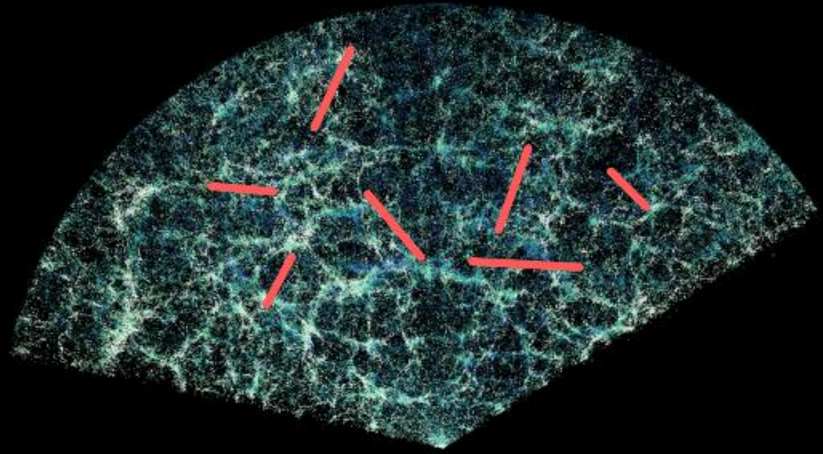

3D map → **pair counts** → two-point statistics:

$$\xi(r,\cos\theta) \xrightarrow{\mathcal{F}} P(k,\cos\theta) \xrightarrow{\text{Legendre}} \xi_\ell(s),\ P_\ell(k)$$

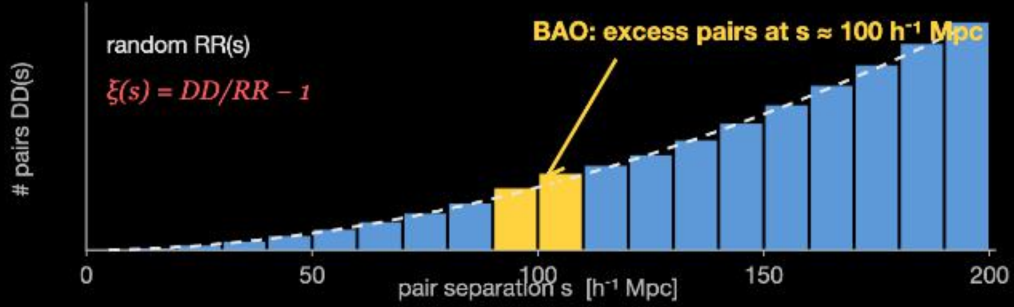

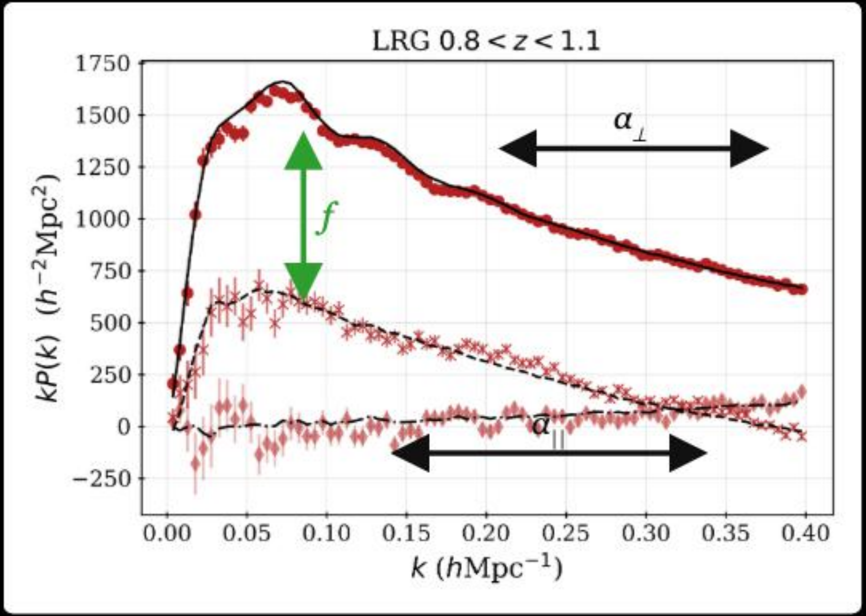

$$\mathbf{A} = \{\alpha_\perp,\ \alpha_\parallel,\ f\}$$

dilation ⊥ / ∥ to the line of sight + growth

map: C. Lamman / DESI Collaboration · multipoles: DESI DR1, LRG 0.8<z<1.1

From map to cosmology: three routes

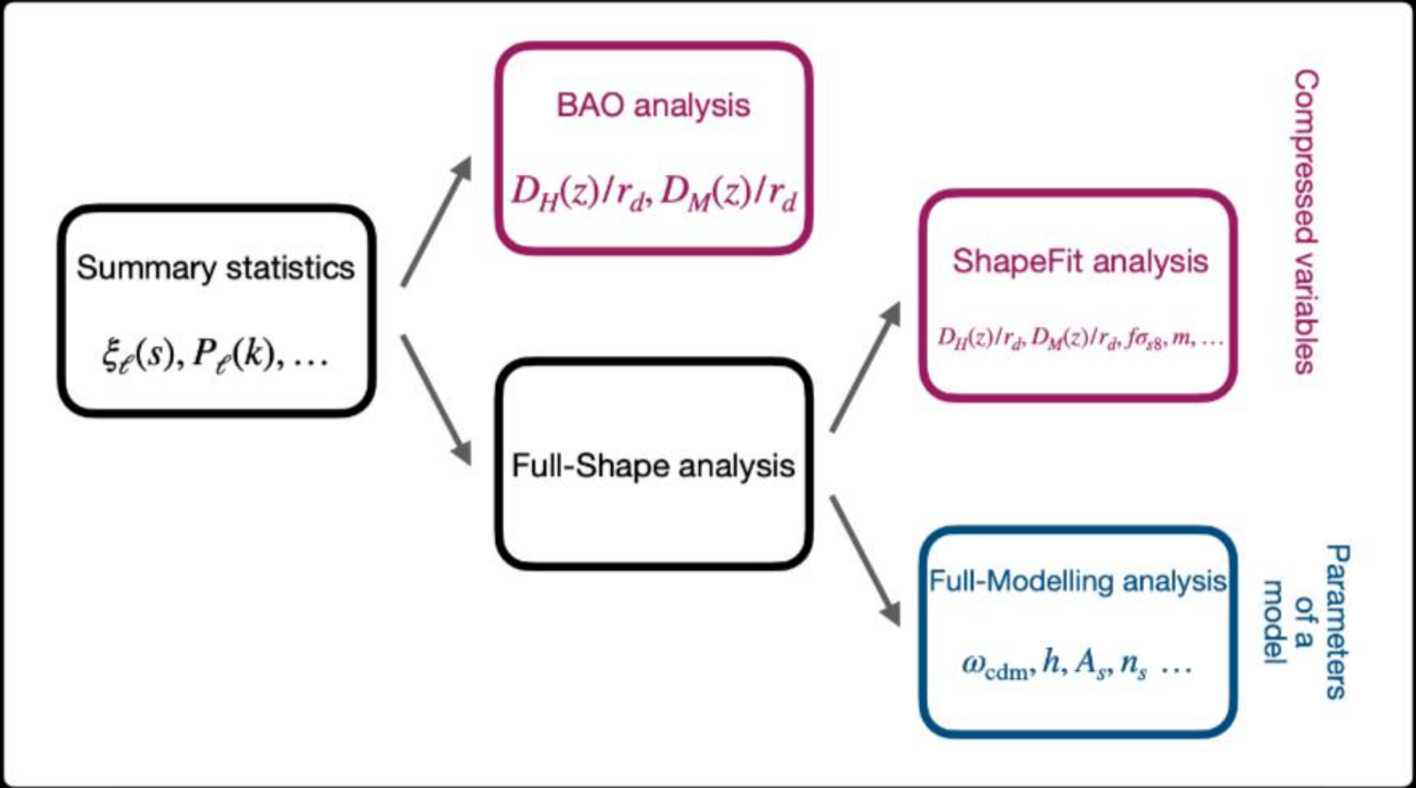

compressed variables — cosmology-agnostic, interpret later

parameters of a model — direct fit ("full modelling")

What is ShapeFit?

Residual parameter dependence of $P(k)$ after r_{d}-rescaling — the derivative response:

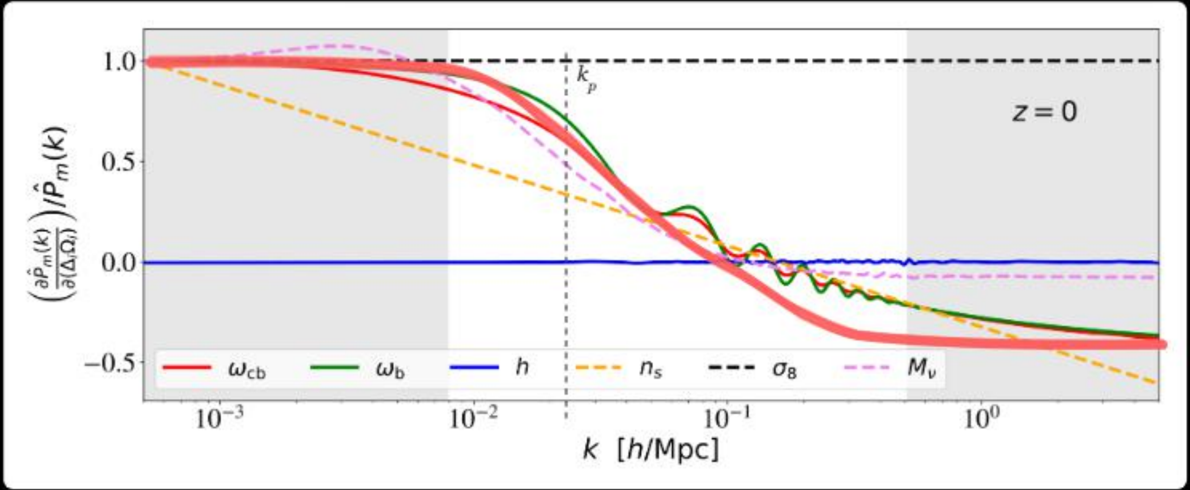

One dominant effect at DESI scales: a **scale-dependent slope** m at pivot k_p

$$\mathbf{A} = \{\alpha_\perp,\ \alpha_\parallel,\ f,\ m\}$$

BAO + RSD + Shape: captures full-modelling information, **model-independently** — and the compressed posteriors are nearly Gaussian.

SB, Gil-Marín & Verde 2021 (arXiv:2106.07641)

Full-shape modelling: the EFTofLSS

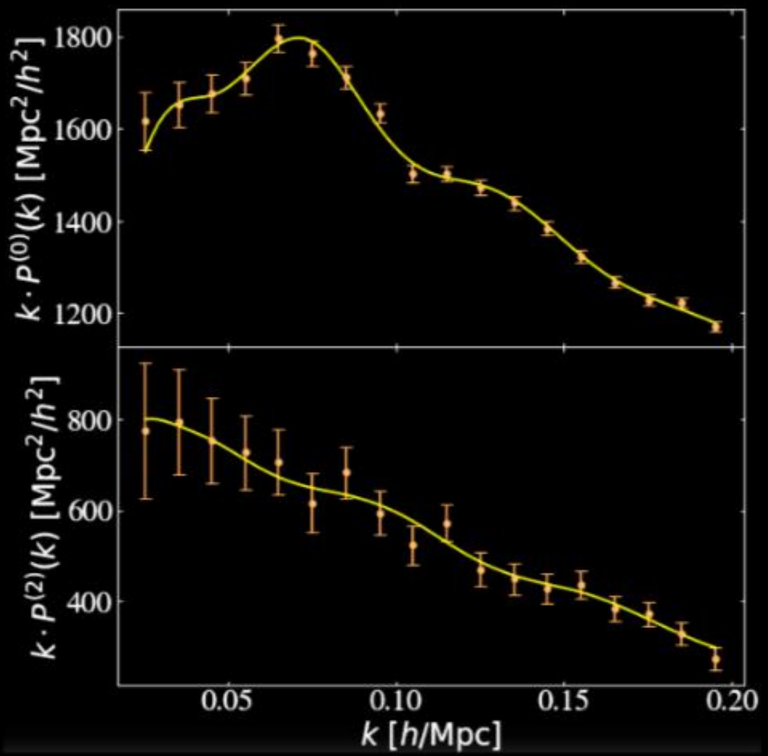

- $P_m(k)$ template
- galaxy bias: $b_1,\ b_2,\ b_{s^2},\ b_{\Gamma_3}$
- one-loop PT + counterterms $\alpha_0,\ \alpha_2$ + stochastic $\mathrm{SN}_0,\ \mathrm{SN}_2$
- redshift-space distortions (velocileptors)

≈ 8 nuisance parameters × 6 tracers ~ **48 nuisance** vs. **5–8 cosmological (4 ShapeFit)**

DESI DR1 full shape (2024)

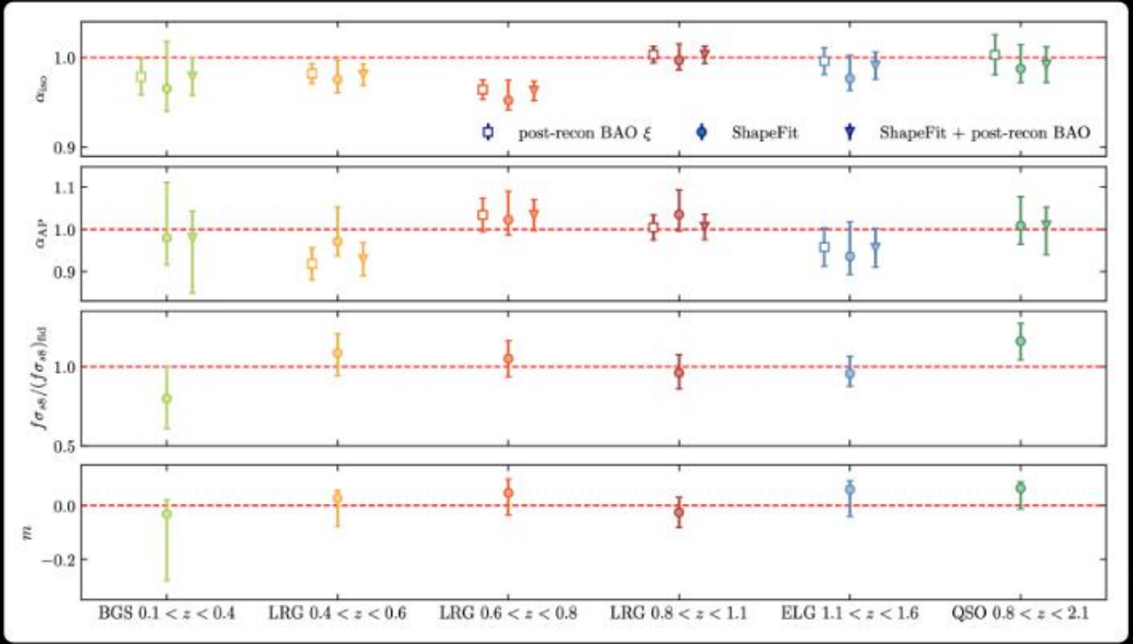

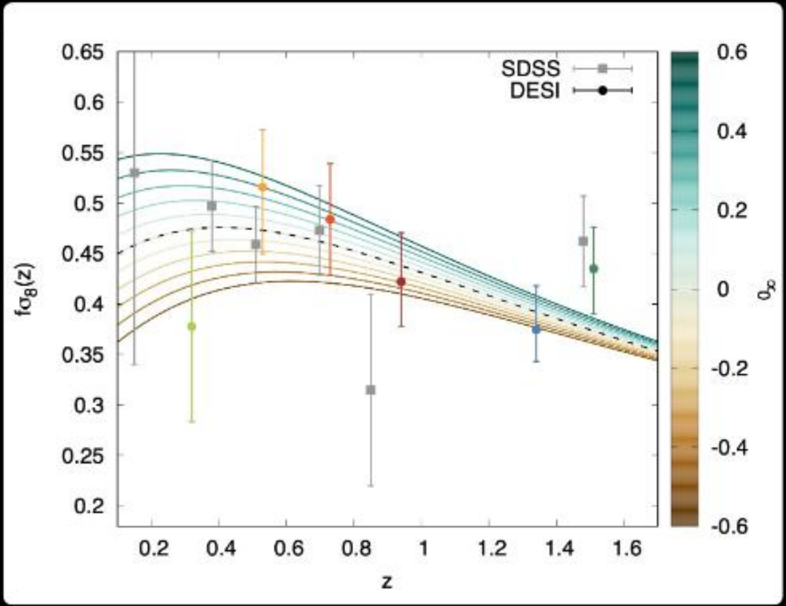

Growth of structure consistent with **GR** (μ0 ≈ 0) and with ΛCDM.

DESI 2024 V & VII; ShapeFit: SB+ 21 (arXiv:2106.07641)

DESI DR2 BAO (2025): the dark-energy hint

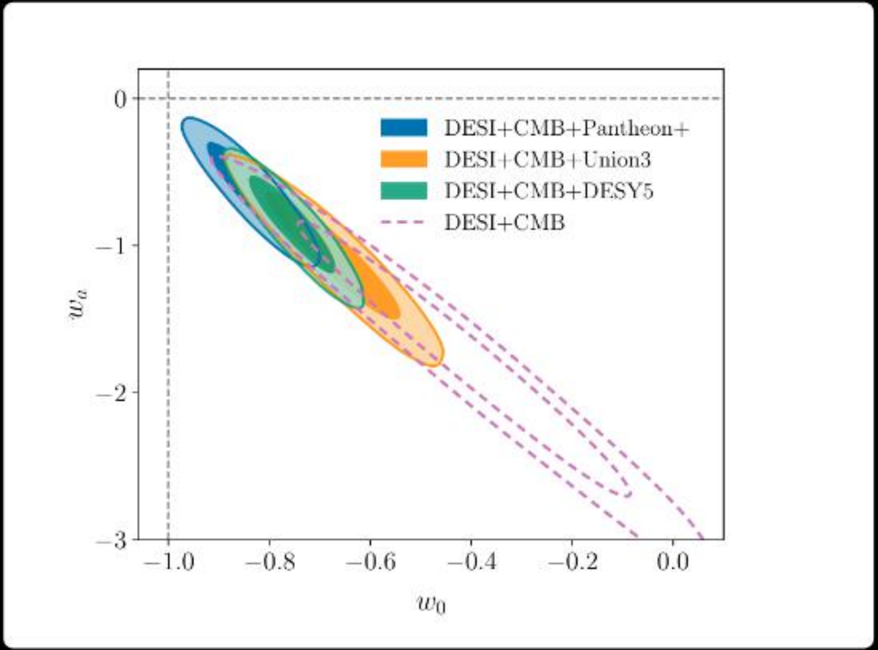

- Most precise BAO to date: >14M galaxies & quasars, sub-percent in every tracer bin (down to 0.47%)
- $$w_0 = -0.42 \pm 0.21, \quad w_a = -1.75 \pm 0.58$$

  (DESI + CMB, $w_0 w_a$CDM)
- preference over ΛCDM: **2.8–4.2σ** depending on SNe sample
- this is expansion history only — growth & broadband shape not yet in the game → **full shape**

DESI Collaboration 2025 (arXiv:2503.14738)

“DESI1.5”: DR1 full shape × DR2 BAO

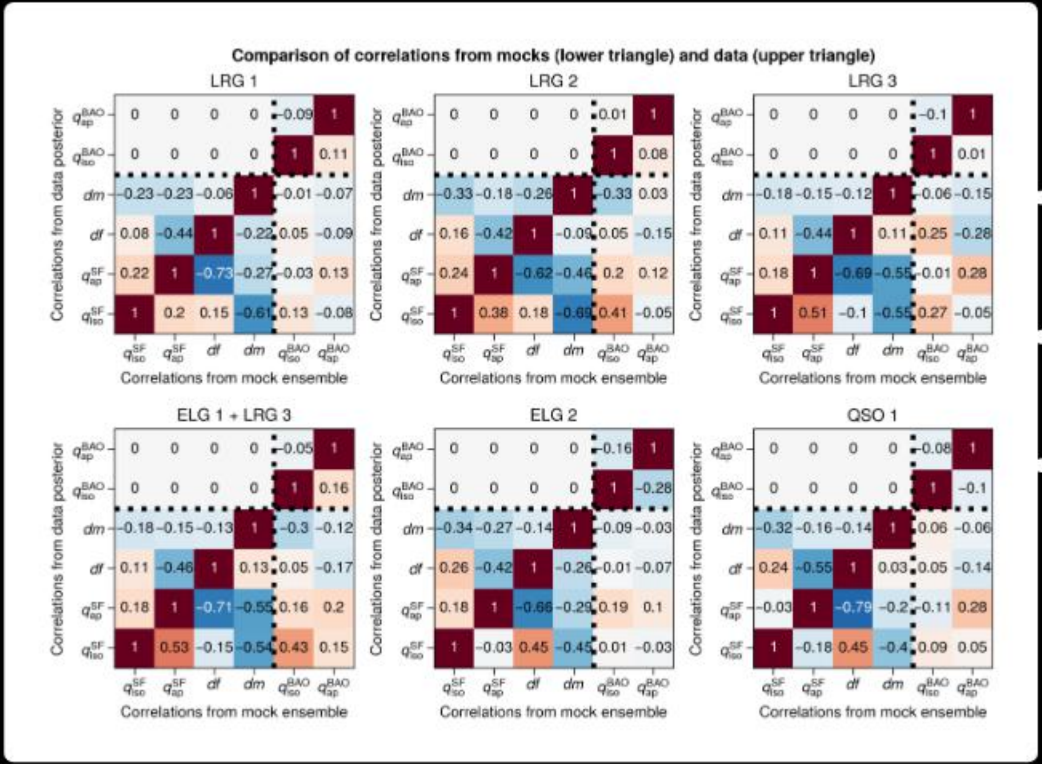

- DR1 full shape via **ShapeFit**: $(q_{\rm iso}, q_{\rm AP}, dm, df)$ per tracer + DR2 BAO
- cross-covariance from 100 EZmocks re-projected onto the DR2 footprint
- compression ⇒ only **8 cross-covariance numbers** per tracer needed — the full-modelling equivalent is untractable

Forero-Sánchez, Gil-Marín, Verde et al. 2026 (arXiv:2602.18761)

DESI1.5: ΛCDM

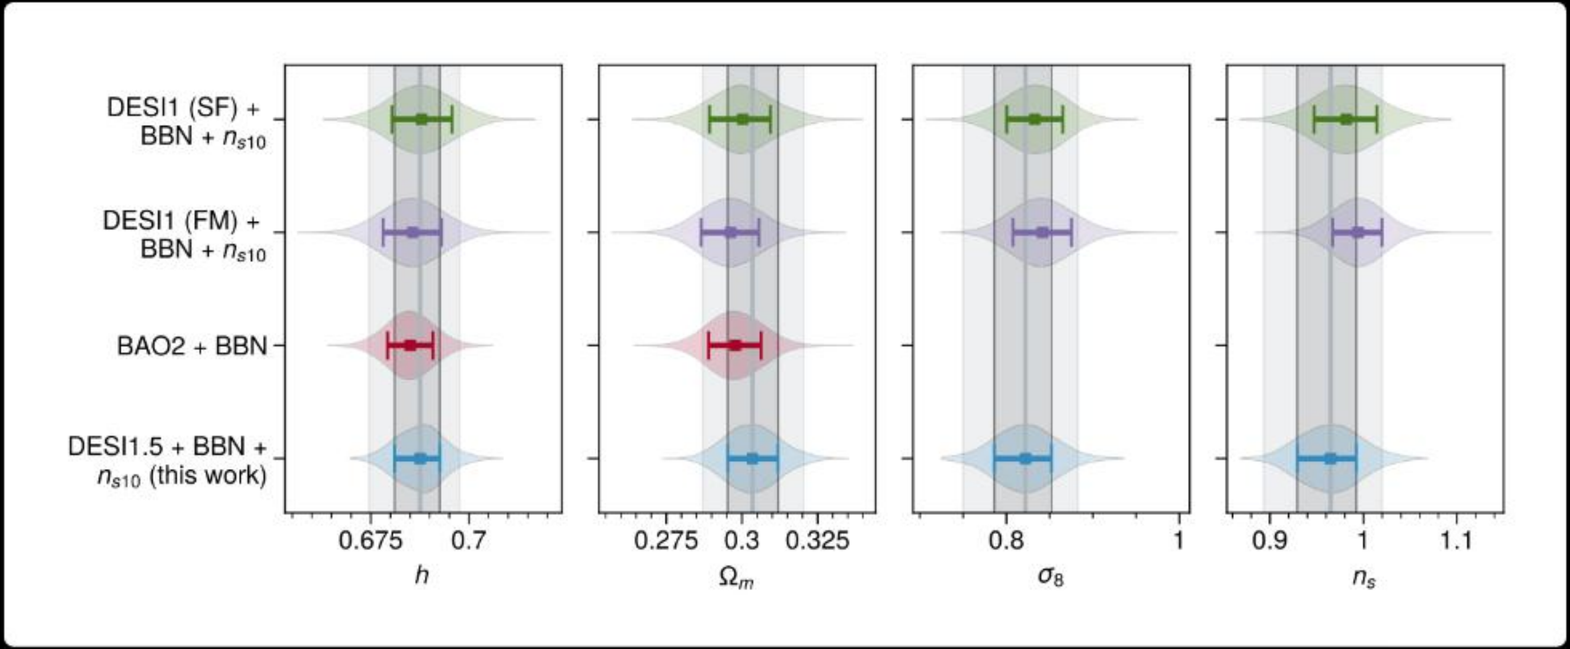

$$\Omega_m = 0.3035 \pm 0.0085, \quad h = 0.6876 \pm 0.0059, \quad \sigma_8 = 0.822 \pm 0.034$$

with a BBN prior + a 10×-Planck-width n_s prior — **no CMB**.

Forero-Sánchez et al. (arXiv:2602.18761), Fig. 4

DESI1.5: evolving dark energy

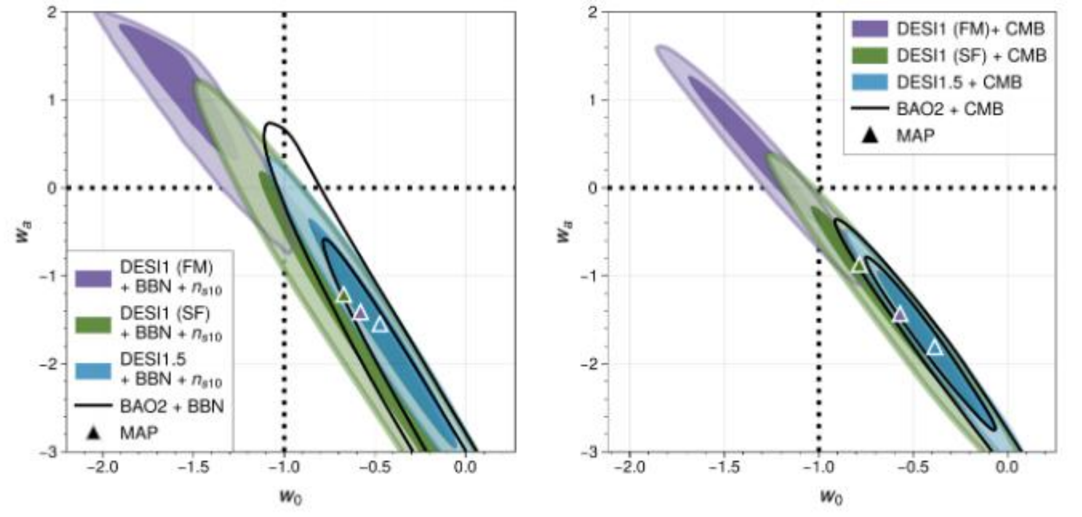

$$w_0 = -0.49 \pm 0.25$$
$$w_a = -1.52 \pm 0.77$$

DESI-only (+BBN+n_{s10})

- ~30% tighter than DR1 full shape + DR1 BAO
- ΛCDM at **1.4σ** DESI-only (1.7σ BAO alone); → **~3σ** with CMB
- (Note the MAP triangles on this plot)

Forero-Sánchez et al. (arXiv:2602.18761), Fig. 6

DESI1.5: neutrino mass

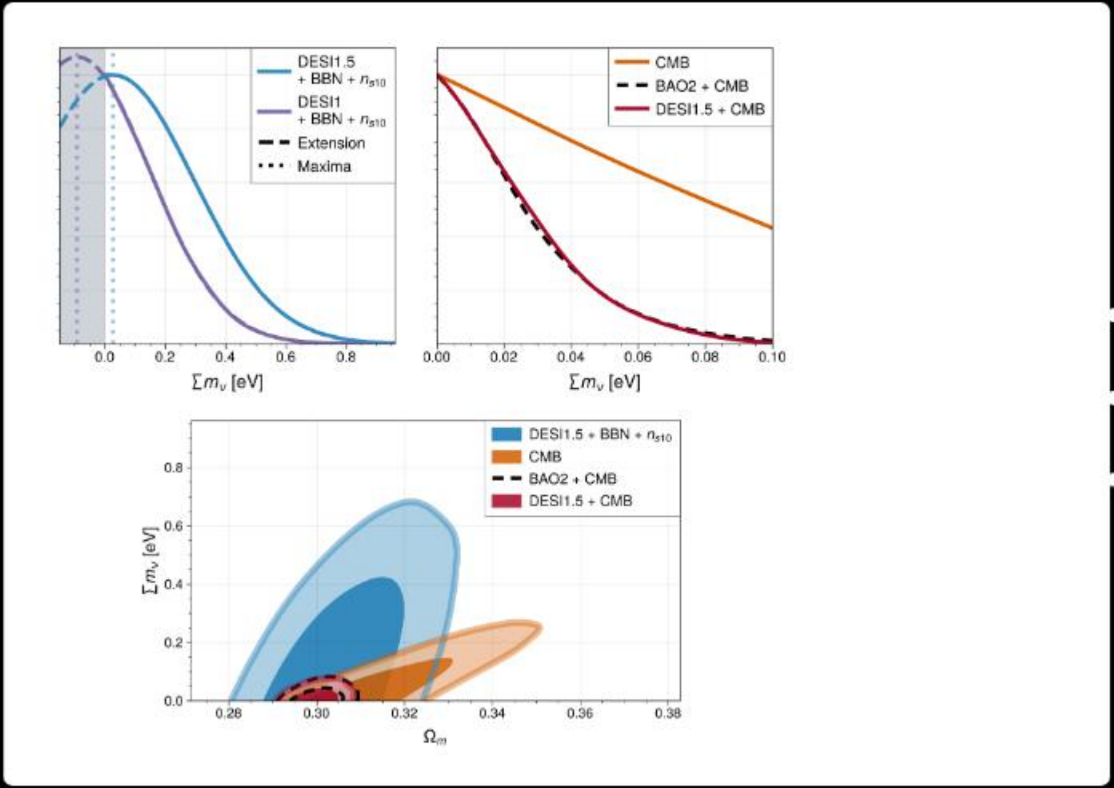

- $\sum m_\nu < 0.542\,\mathrm{eV}$ (95%, DESI-only)
- posterior peaks at **positive** mass
- + CMB: $\sum m_\nu < 0.0624\,\mathrm{eV}$

Forero-Sánchez et al. (arXiv:2602.18761), Fig. 7

DESI1.5: curvature

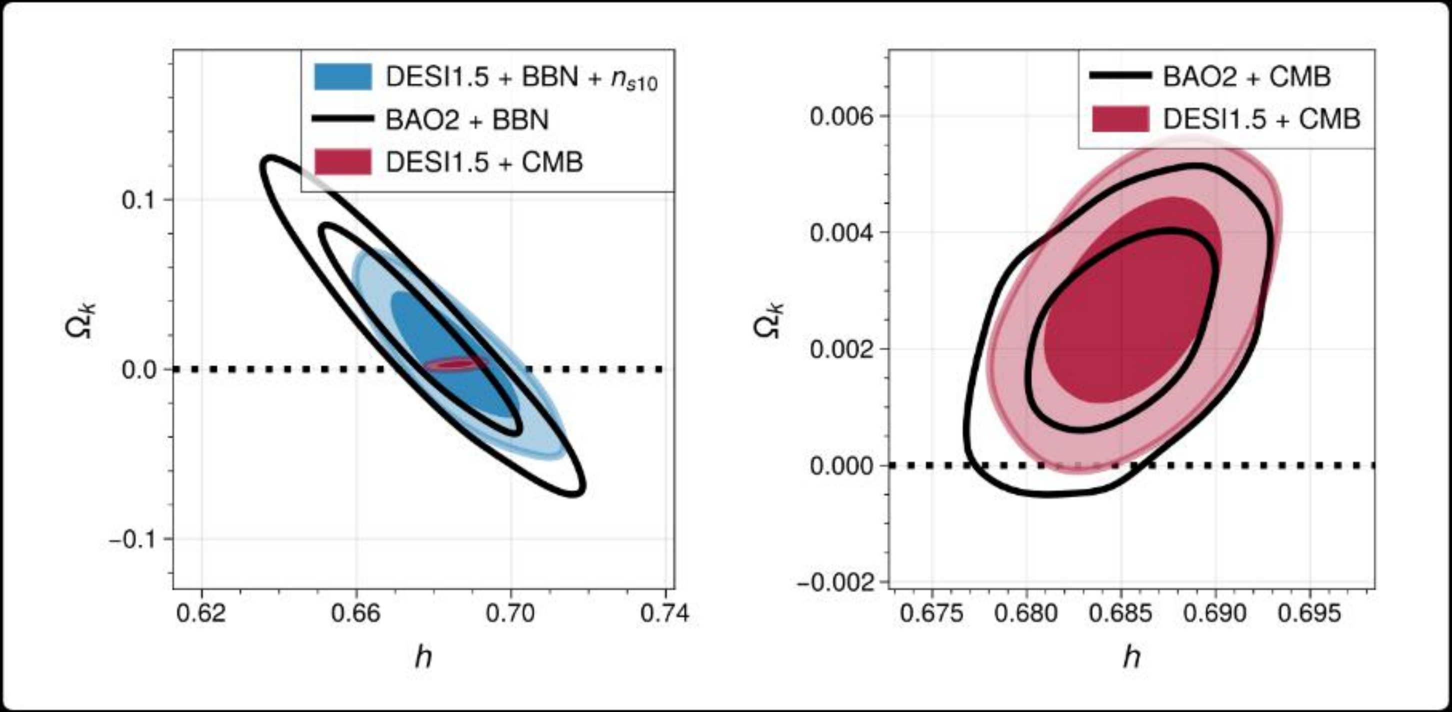

DESI-only: $\Omega_k = 0.007 \pm 0.024$ – 2× tighter than BAO alone

+ CMB: $\Omega_k = 0.0028 \pm 0.0011$

Forero-Sánchez et al. (arXiv:2602.18761), Fig. 10

Update: DESI1.5 + the LRG bispectrum

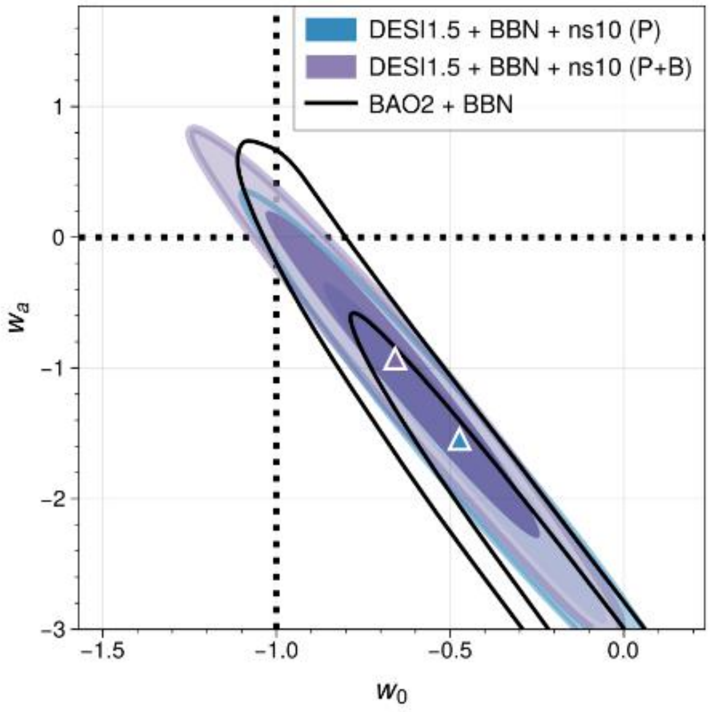

- DR1 LRG bispectrum monopole joins the compression ($k \leq 0.12$): f and σ_{s8} now measured separately
- amplitude rises & tightens: $\sigma_8^{P+B} = 0.868 \pm 0.025$ (from $\sigma_8^{P} = 0.822 \pm 0.034$)
- dark energy relaxes toward ΛCDM: it sits inside the 1σ contour
- preference for $w_0 w_a$: ~2.8σ w/ CMB; 3.05σ w/ CMB + **DES-Dovekie** SNe (SNe alone: 1.6σ)
- $\sum m_\nu = 0.26 \pm 0.17$ eV (DESI-only) — a positive peak

Forero-Sánchez et al. 2026 (arXiv:2606.23936), Fig. 5 (DESI-only panel)

PART II

The challenge:
bias degeneracies & projection effects

The nuisance sector

$$P_g(k,\mu) \;=\; \left(b_1 + f\,\mu^2\right)^2 P_m(k) \;+\; P_{\text{1-loop}}\left[b_2,\, b_{s^2},\, b_{\Gamma_3}\right] \\ +\; \left(\alpha_0 + \alpha_2\mu^2\right) k^2 P_m(k) \;+\; \mathrm{SN}_0 + \mathrm{SN}_2\,\mu^2 k^2$$

- amplitude degeneracy: the data constrain only the products $b_1\sigma_8$, $f\sigma_8$ (+ shot-noise mixtures) — A_s on its own is not measured
- loop terms: $P(k)$ alone barely constrains $b_2, b_{s^2}, b_{\Gamma_3}$ (prior-dominated) — their loop contributions are smooth, mutually near-degenerate broadband shapes
- counterterms mimic shape changes (m, n_s, ω_m); stochastic terms mimic small-scale power
- ⇒ prior volume effects …

Projection effects: toy example

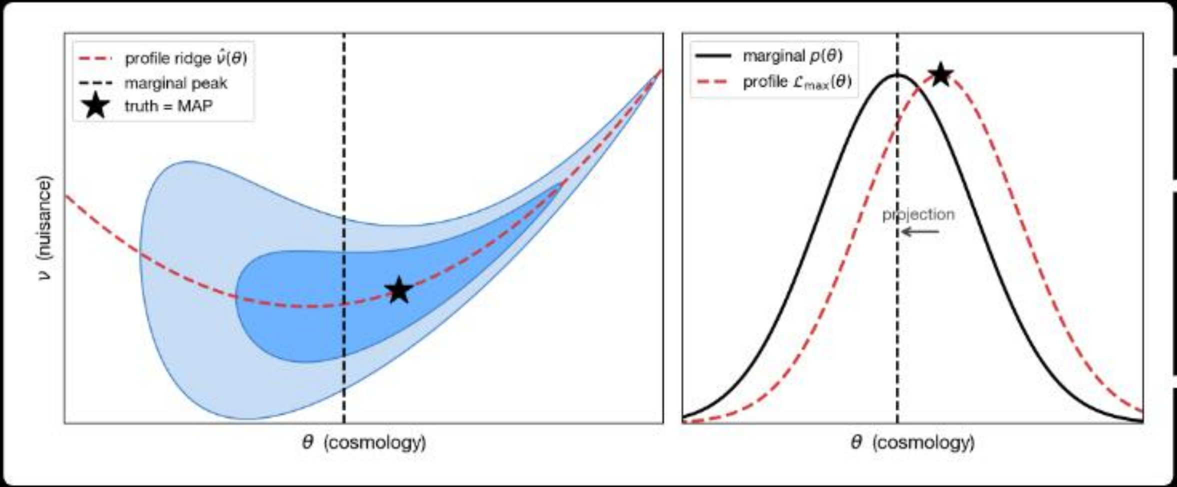

- wide flat priors × curved degeneracies ⇒ posterior mass ≠ best fit
- even uniform priors are informative — the parametrization matters (b_1 vs $b_1\sigma_8$)
- symptom: MAP can be outside its own 68% contour

toy illustration made for this talk · mechanism: Forero-Sánchez et al. (arXiv:2602.18761), §4.5 & App. A

Prior sensitivity in practice (DR1)

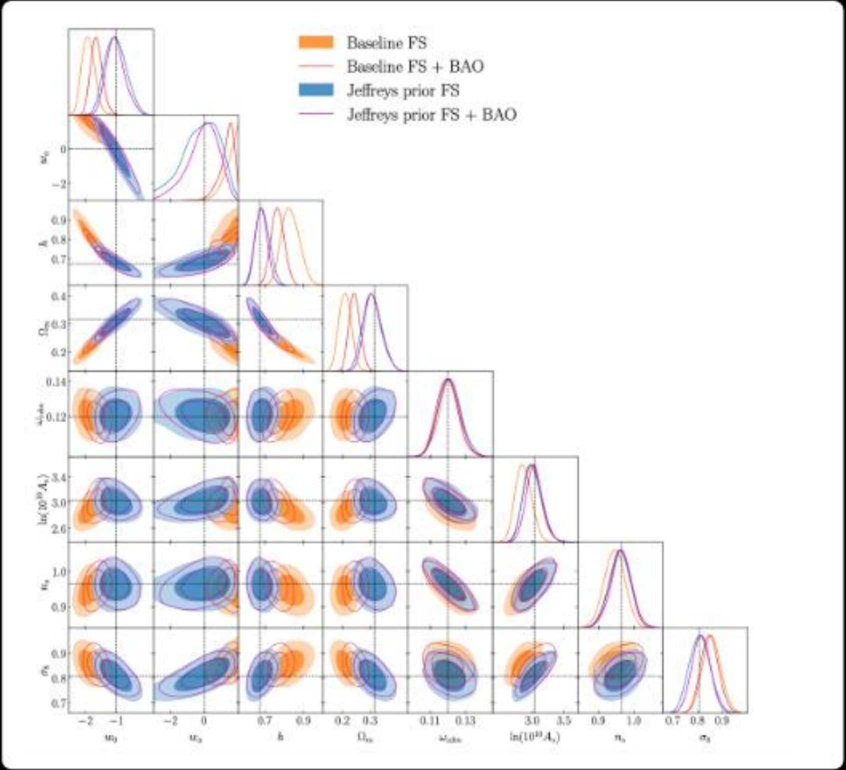

flat → Jeffreys priors on counterterms + stochastic: the cosmology contours move (noiseless synthetic data)

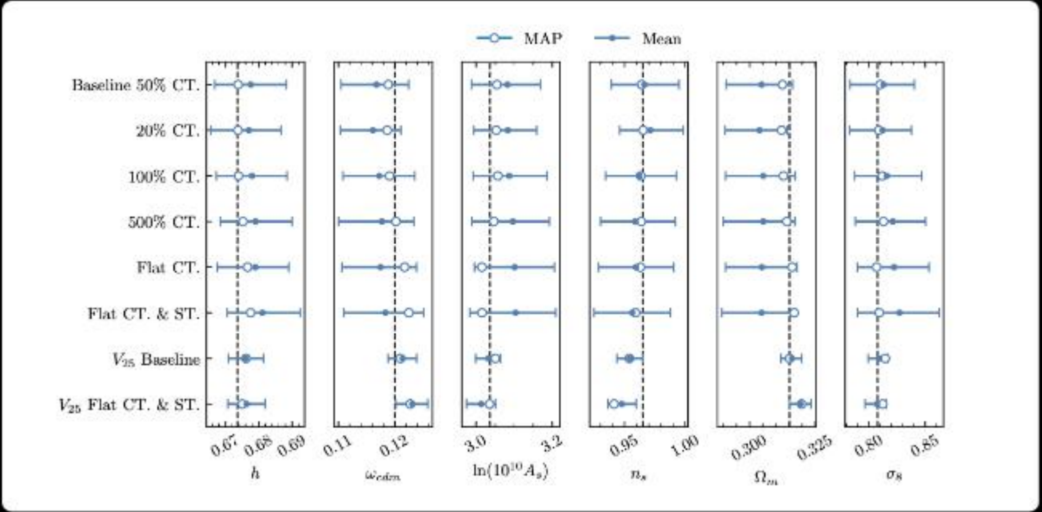

counterterm/stochastic prior widths (20%–500%–flat): MAP–mean drift — which fades at 25× the DR1 volume (V25 rows)

DESI 2024 V (arXiv:2411.12021), Figs. 3 & 2

Why compression helps — pros and cons

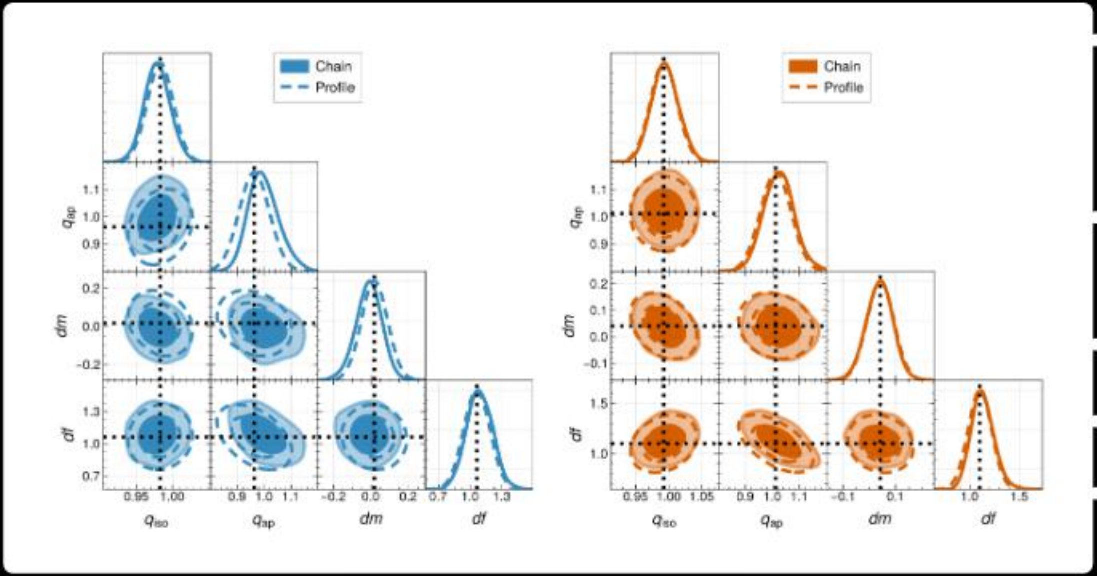

- ShapeFit compressed posteriors: chain ≈ profile ≈ **Gaussian** → MAPs inside, DESI-only extensions become reliable
- 4 physical numbers per tracer absorb the broadband — the compression regularizes
- but: needs an external n_s prior;
- still discards information
- need to think about systematics (Ghaemi et al. 2025, Gsponer et al. 2026)

Forero-Sánchez et al. (arXiv:2602.18761), Fig. 13

Full Modelling: P(k)+B(k)

Two options: shrink the prior (physics) — or **add data that measures the nuisances**.

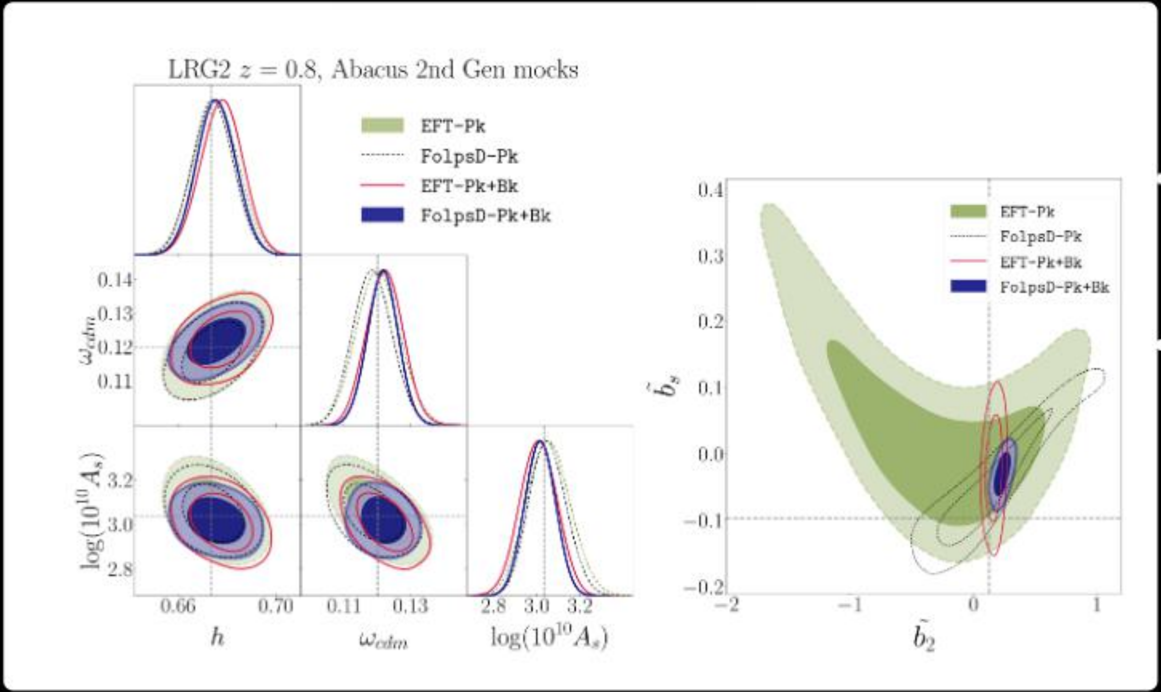

- b_2, b_{s^2} enter the **tree-level bispectrum** directly — in $P(k)$ only at one loop: the open $\tilde{b}_2$–$\tilde{b}_s$ degeneracy collapses
- FolpsD: its IR+FoG damping replaces counterterm/stochastic freedom → fewer flat directions, keeps small-scale info — P-only already ~10% tighter than EFT-P; damping+B up to ~30%

Bansal and DESI 2026 (arXiv:2604.08895; Fig. 4)

Fixes on the market

HOD-informed priors
Ivanov+24, Zhang+24
inherits galaxy-formation assumptions; simulation-tuned

Simulation-calibrated bias relations
Eggemeier+21, Zennaro+22, Akitsu+24
accurate where calibrated; emulator-bound

Frequentist profile likelihood
Morawetz+25
prior-free diagnosis; costly, and no posterior to propagate

Reparametrization / Jeffreys priors
Hadzhiyska+23, Paradiso+24, Tsedrik+25 (2509.09562), Bonici+26 (2607.02498)
reshapes the prior volume — specific choices can be physically motivated, but the volume problem is relocated rather than removed

Wanted: **analytic, physical, parameter-free** priors.
Idea: use the cosmic web hierarchy to construct such priors

PART III

A physical way out: bias from the cosmic web

Web Halo Model Peak Background Split (WHM-PBS)

Ellipsoidal collapse Excursion Set

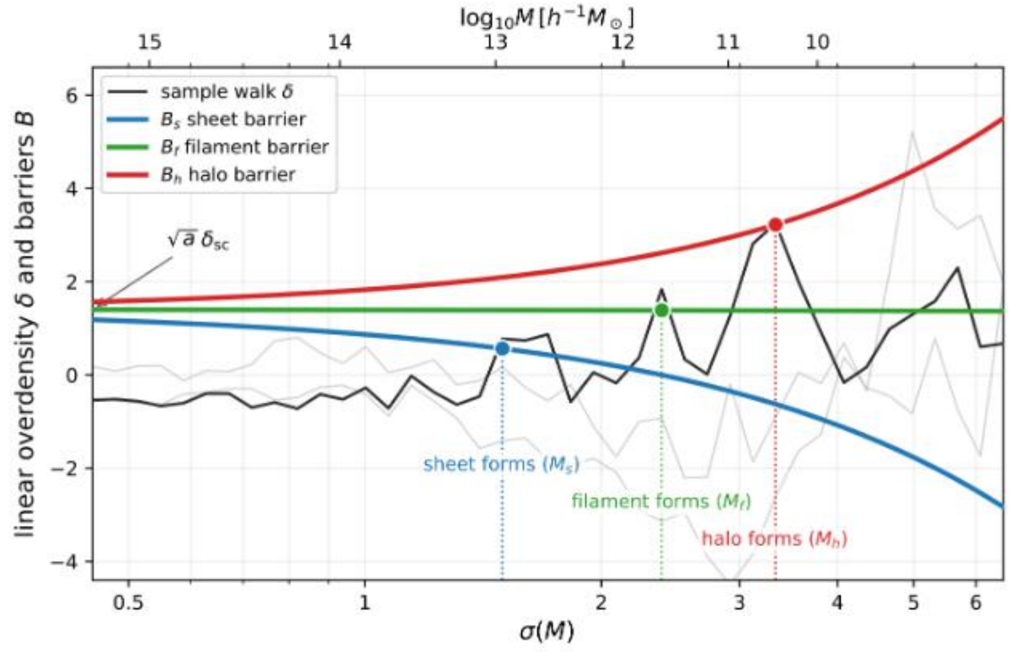

- haloes form inside filaments inside sheets: one excursion-set walk, three nested moving barriers (Shen+06)
- the walk fixes the **host distribution** $P(M_f \mid M_h)$ — broad and skewed; a single calibrated constant ($a = 0.707$)
- peak–background split on the environment ⇒ halo bias = host-averaged bias: **a distribution, not a number**

> EFT nuisance parameters = **moments of one host distribution** — computable, not free.

SB & Tipp, in prep.; Web–Halo Model: SB, Beutler & Pellejero-Ibañez 2025 (arXiv:2508.10902); Barriers from Shen et al. 2006

The peak–background split in one picture

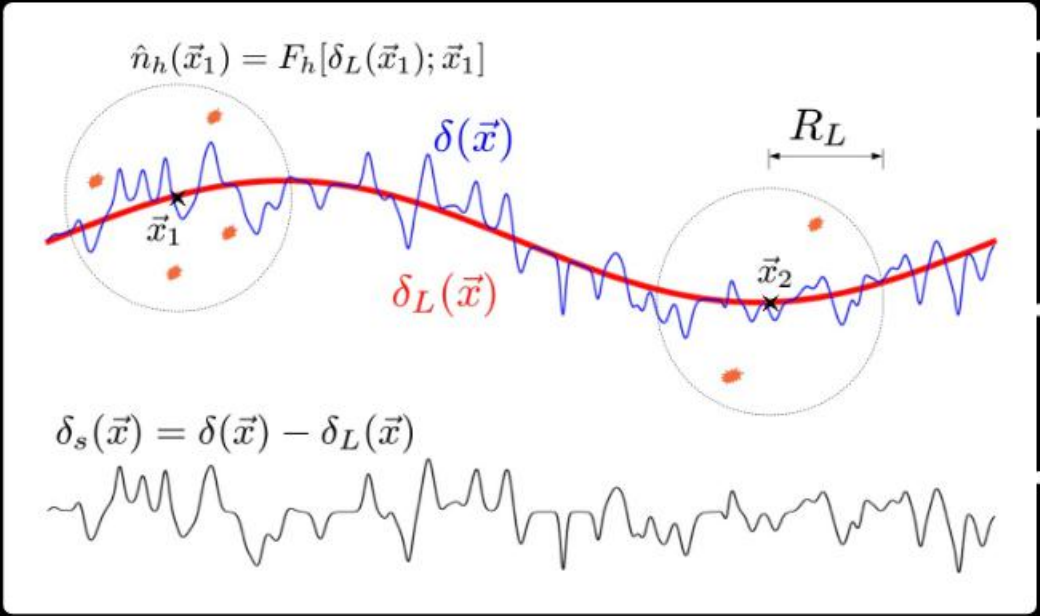

- split the density field: $\delta = \delta_L + \delta_s$
- a long mode raises the local mean ⇒ small-scale peaks reach the collapse threshold earlier: $B \to B - \delta_b$
- the abundance responds ⇒ that is bias: $b_N(M) = \frac{1}{n(M)} \frac{\partial^N n(\,M \mid \delta_b)}{\partial \delta_b^N}\Big|_{\delta_b=0}$
- classic PBS evaluates this on the halo mass function — WHM-PBS evaluates it on the halo's environment → next slide

sketch: Schmidt, Jeong & Desjacques 2012 (arXiv:1212.0868), Fig. 1 · PBS: Kaiser 84; Bardeen+86; Cole & Kaiser 89; Mo & White 96; Sheth & Tormen 99

WHM-PBS: the environment-averaged bias

1. PBS as a barrier shift — bias = first-crossing response:

$$b_N(M) = \frac{1}{n(M)} \frac{\partial^N n(\,M \mid \delta_b)}{\partial \delta_b^N}\bigg|_{\delta_b=0}$$

2. Chapman-Kolmogorov decomposition

$n_h(M_h) = \int N(M_h|M_f)\; n_f(M_f) dM_f$ — closure ⇒ normalised

3. δ_b lowers halo and host barriers equally ⇒ $N(M_h|M_f)$ passes through $\partial/\partial\delta_b$:

$$\begin{aligned} b^L_{h,N}(M_h) &= \frac{1}{n_h(M_h)} \int N(M_h|M_f)\, \frac{\partial^N n_f(M_f|\delta_b)}{\partial \delta_b^N}\bigg|_0 dM_f \\ &= \int \underbrace{\frac{N(M_h|M_f)\, n_f(M_f)}{n_h(M_h)}}_{P(M_f|M_h)}\, b^L_{f,N}(M_f)\; dM_f \;\equiv\; \langle b^L_{f,N} \rangle_{M_f} \end{aligned}$$

the mean reproduces standard PBS bias **exactly** (closure) — but the width of $P(M_f|M_h)$ makes bias **a distribution, not a number**

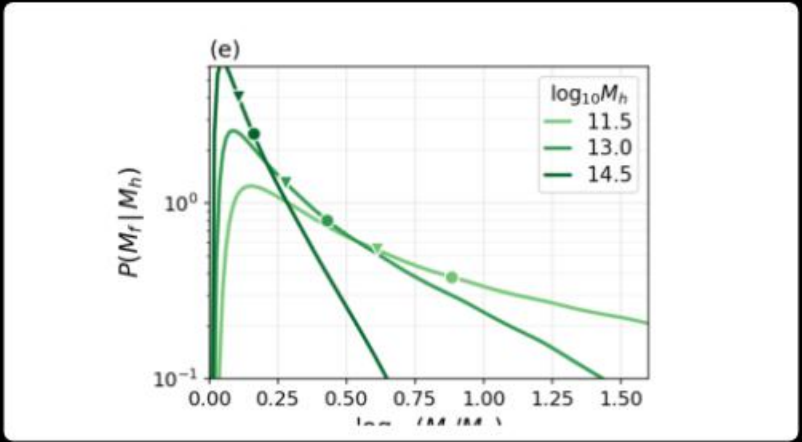

the host distribution at fixed halo mass: broad, right-skewed — and every moment of it is a prediction: widths → the prior bands (next two slides); scatter → stochasticity & assembly bias (week 4)

SB & Tipp, in prep. (Eqs. 3.3, 2.6–2.7, 3.4; Fig. 2) · Eulerian frame via co-evolution

Bias-relation priors: $b_2(b_1)$, $b_3(b_1)$

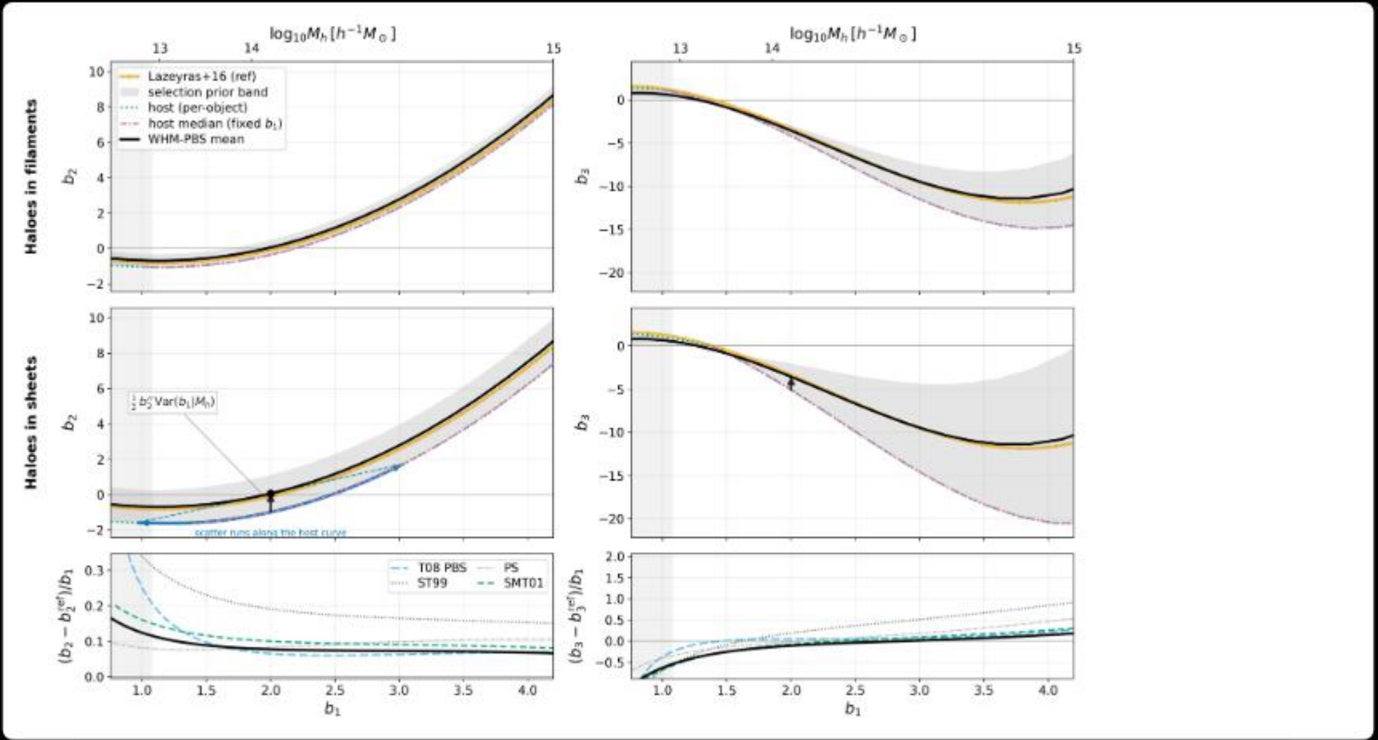

tight bands (σ ≈ 0.2) with a hard lower envelope — tracking the Lazeyras+16 N-body fits with no fitted parameter; the width is inherited from the host distribution

SB & Tipp, in prep. (Fig. 5)

The tidal sector: $b_{s^2}(b_1)$ and $b_{\Gamma_3}(b_1)$

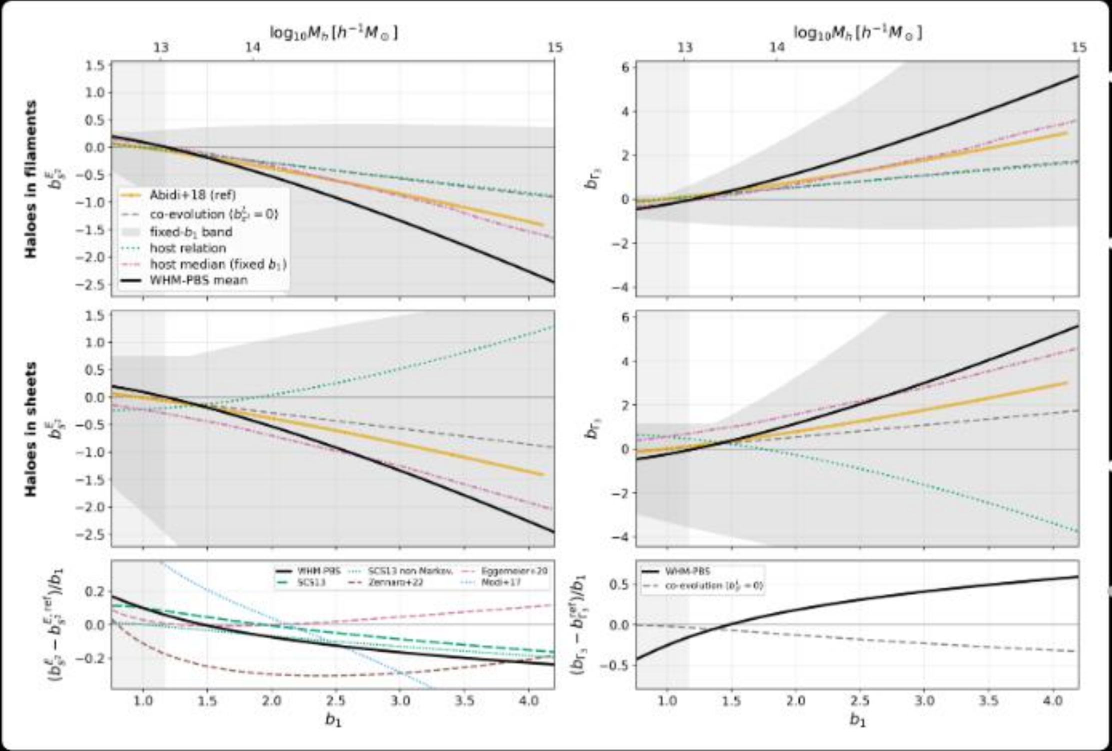

- a long-wavelength shear raises the moving part of the barrier: $\delta B = \frac{3}{2}\, s_L^2\, \mathrm{d}B/\mathrm{d}S \Rightarrow b_{s^2}$ with **no new parameter**
- sits below pure co-evolution (grey dashed, $b^L_{s^2} = 0$) — and carries **genuine scatter** at fixed b_1: the assembly-bias asymmetry of Lazeyras+21
- b_{Γ_3} follows by co-evolution-with-tides: lifted above the standard EFT prior $\frac{23}{42}(b_1 - 1)$
- the tidal imprint cancels in one-loop $P(k)$ — but survives in the **bispectrum**: testable with DESI P+B

SB & Tipp, in prep. (Fig. 6)

Proof of concept: a DESI-like demo

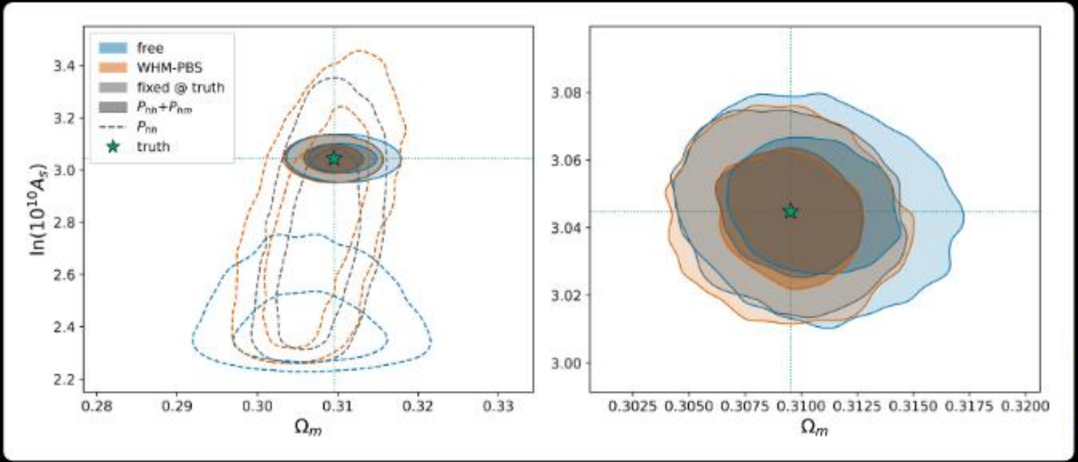

one-loop EFT fit (velocileptors), 3 DESI-like tracer bins, mock = model at truth, $\chi^2(\mathrm{truth}) = 0$ ⇒ any drift is pure inference geometry. Dashed = P_{hh}; filled = $P_{hh} + P_{hm}$; right: zoom onto the recovered truth.

figure of bias FoB (σ, combined over $\Omega_m, h, \ln 10^{10} A_s$) − P_{hh} only:

free (wide flat)	3.15
fixed @ WHM-PBS +1σ	2.18
fixed @ truth (oracle)	0.71
WHM-PBS priors	0.47

widths, not centres: the band beats even the oracle — while a pin displaced by ±1σWHM blows up (2.18 / 1.85): the band absorbs the very error that breaks a pin

SB & Tipp, in prep. (Table 2, Fig. 11)

The demo: full posteriors

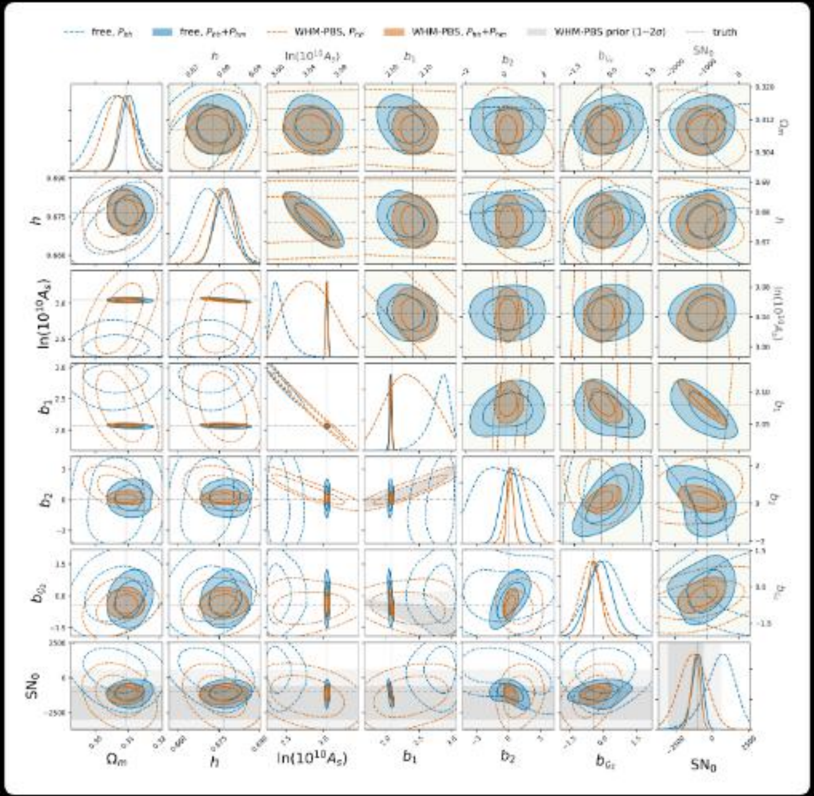

LRG2 bin: free (blue) vs WHM-PBS (orange); dashed P_{hh}, filled $P_{hh} + P_{hm}$; grey = WHM-PBS prior bands (1–2σ) on $b_2, b_{\mathcal{G}_2}, \mathrm{SN}_0$; dotted = truth

SB & Tipp, in prep.

The full story: week 4

GGI · WEEK 4

Modeling Higher-Order Bias via the Cosmic Web: Implications for DESI Full Shape Analyses

- the Web–Halo Model & the moving-barrier hierarchy
- bias relations with widths: $b_2(b_1),\ b_3(b_1),\ b_{s^2}(b_1)$
- stochasticity & assembly bias — two faces of one predicted variance
- the DESI priors recipe, and the demo in depth

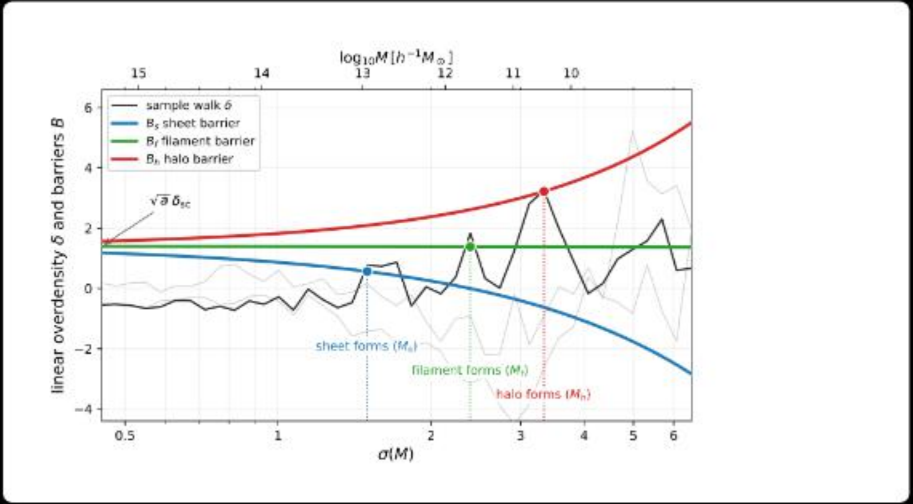

Conclusions

- **DESI full shape delivers.** DESI1.5 = DR1 full shape × DR2 BAO: reliable DESI-only constraints at last — $w_0 w_a$CDM ~30% tighter, first credible DESI-only $\sum m_\nu$. With the LRG **bispectrum**: $\sigma_8 = 0.868 \pm 0.025$ and the DESI-only DE hint relaxes toward ΛCDM. And from **peculiar velocities**, growth at $z = 0.07$ to 12% — $\gamma = 0.58 \pm 0.11$, no bias model required.

- **The bottleneck is the EFT nuisance sector.** Wide flat priors + degeneracies ⇒ projection effects: stranded MAPs, artificial neutrino bounds, prior-dependent contours. Compression regularizes the statistics — it cannot restore the marginalized-away information.

- **Two levers, one goal.** Deepen the likelihood — the bispectrum pins b_2, b_{s^2} (DESI's P+B pipeline is coming). And make the priors physical — the cosmic web makes the rest computable: WHM-PBS priors reach oracle-level accuracy in a DESI-like demo. Theory & recipe: week 4.

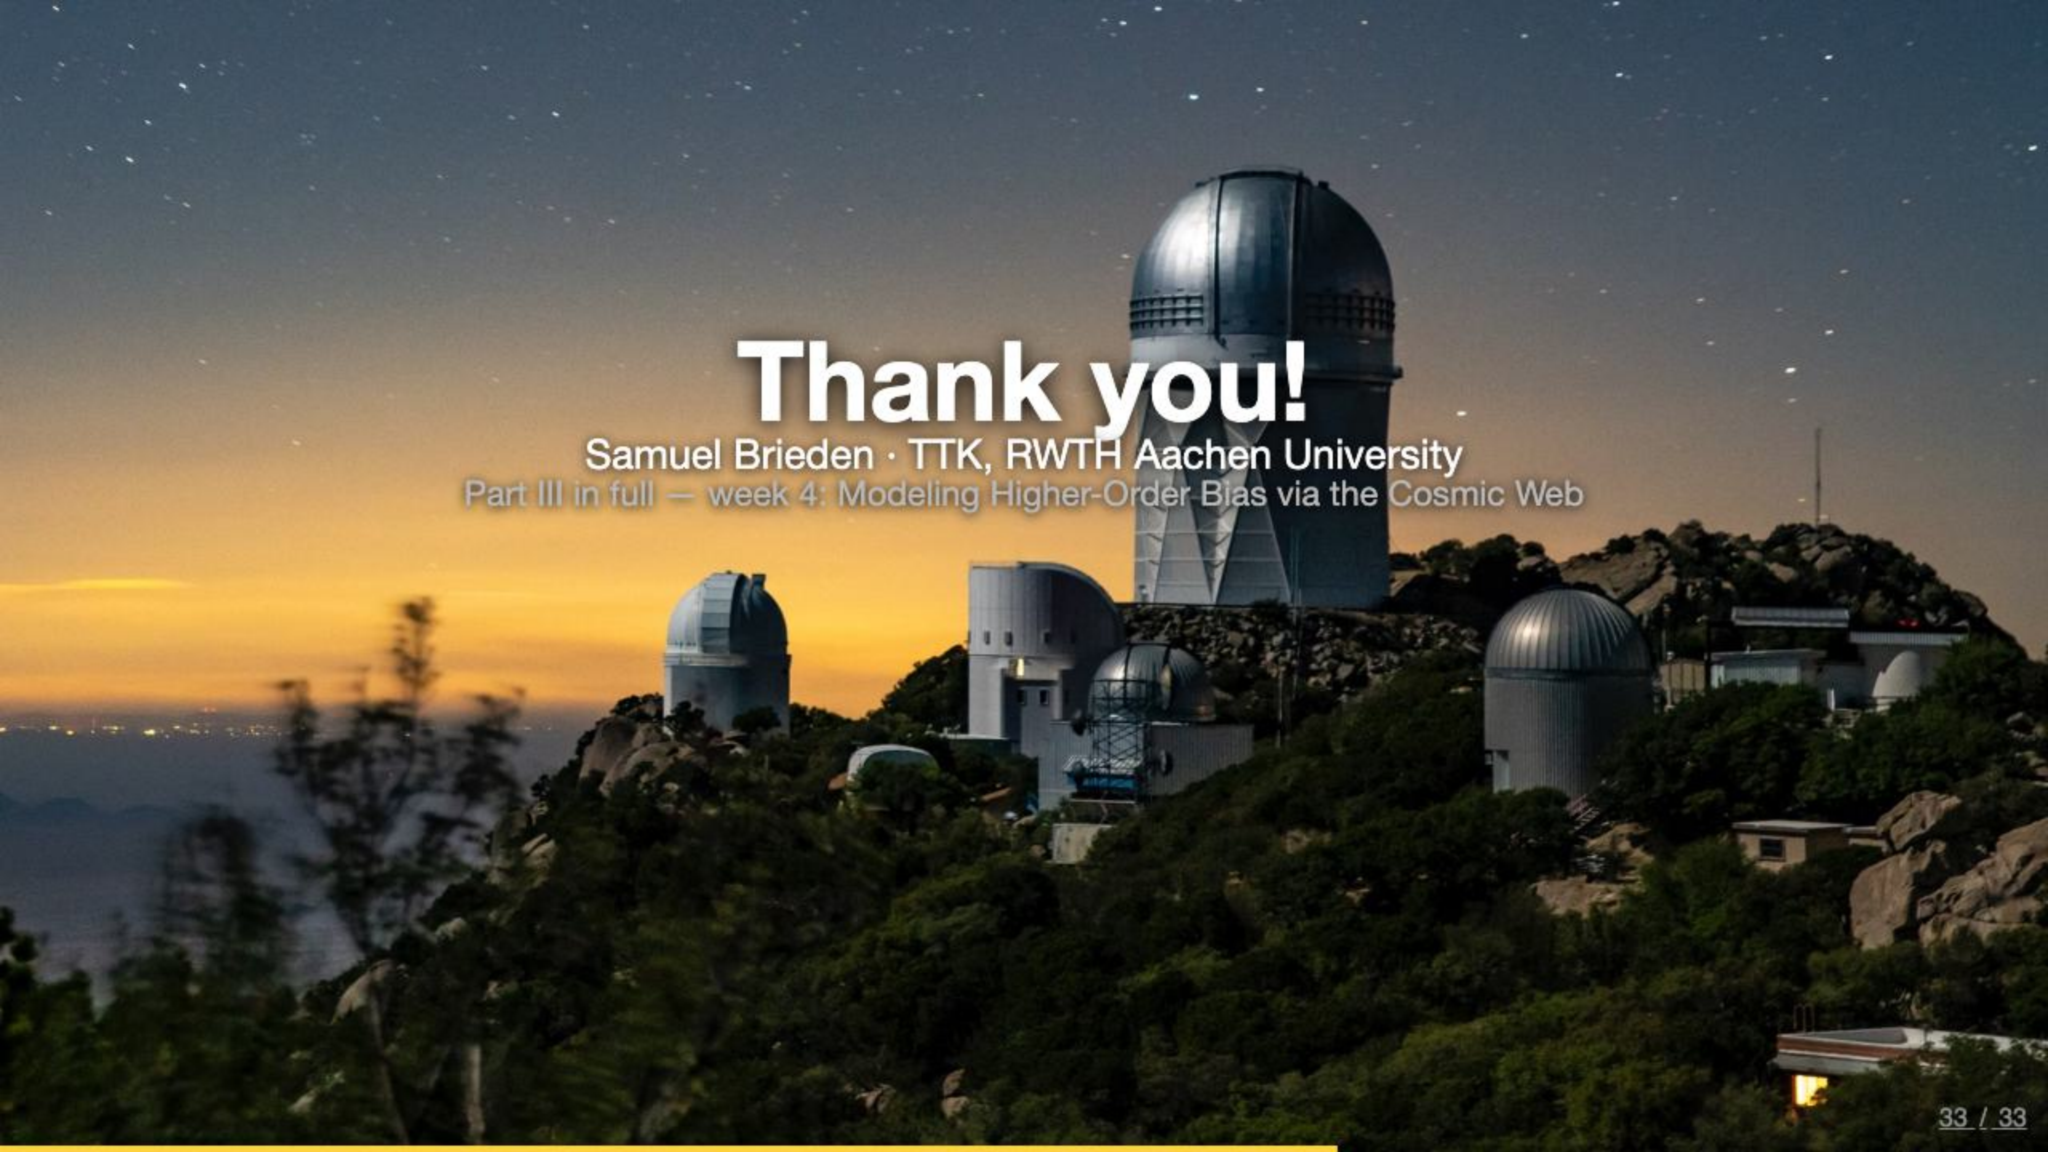

Thank you!

Samuel Brieden · TTK, RWTH Aachen University

Part III in full — week 4: Modeling Higher-Order Bias via the Cosmic Web