Coherence–Filtered BECO Seeding at the Drag Epoch:
Drag–latched memory snapshot, BAO inheritance, and ν–parametrized rare peaks in ΦBSU

Foorumipläjäys Eusalta
(internal draft; method+simulation insert for the ΦBSU Part II/III track)

March 3, 2026

Abstract

We update a simulation–ready closure for *primordial* BECO seeding (Buoyant Exotic Compact Objects) within the ΦBSU programme. The modelling stance is "process, not identity": in the tightly coupled photon–baryon plasma, proto–BECO structures are not persistent identified masses but *coherently fluctuating interference foci* in the null–threading network. Only when comparator coherence opens around the baryon drag epoch can these foci write a long–lived "vacuum memory" which later condenses into compact centres.

This revision inserts a "production" pipeline result based on CLASS outputs and implements two key technical upgrades: (i) the focusing proxy is gauge safe in linear theory (built from the Bardeen/Weyl potential and made dimensionless by $\mathcal{H}^{-2}$); (ii) seeding is *drag–latched* by a narrow impulse kernel $\hat{g}(\eta) \propto \partial_\eta G(\eta)$, rather than by indefinite post–drag accumulation.

Most importantly, the tracer abundance and bias are no longer hand–set: we provide a publishable map from the drag–latched transfer $S_d(k)$ to $\{\bar{n}(\nu), b(z;\nu)\}$ using rare–peak statistics. Once a single rarity parameter ν is chosen, the spectral moments of the implied stochastic memory field determine $\bar{n}$ and b without ad hoc normalizations. We summarize BAO/CMB hygiene conditions and report the acceptance checks for the current reference run (Planck–2018 baseline): $r_d = 147.114\,\mathrm{Mpc}$ is unchanged at CLASS tolerance; the seeding window peaks near $k_{\mathrm{peak}} \simeq 0.147\,\mathrm{Mpc}^{-1}$, i.e. $k_{\mathrm{peak}}/k_{\mathrm{BAO}} \simeq 3.45$; the drag impulse kernel is localized (peak $z \simeq 943$, FWHM $\simeq 51\,\mathrm{Mpc}$ in conformal distance); BBKS spectral moments of the implied memory field give $R_* \simeq 10.49\,\mathrm{Mpc}$, $\gamma \simeq 0.855$; and the configuration–space BAO bump matches matter to $\Delta r \simeq 0.18\,\mathrm{Mpc}$ once shot noise is removed.

Contents

1 Programme context and target

This note is a method+simulation insert for the ΦBSU research track. We assume the following programme components are documented elsewhere and do not repeat them in detail: (i) the ΦBSU Part I field dictionary (global phase, vacuum 4–density, identity/curvature split), (ii) the 4–dimensional cosmological principle and "buoyant spacetime" picture, and (iii) the hypersymmetry/projection stance (operational charts vs. primary null–threading base).

The narrow objective here is quantitative and falsifiable:

> *Provide a gauge–safe, BAO/CMB–hygienic seeding map at the drag epoch, and derive the seed tracer abundance $\bar{n}$ and bias $b(z)$ directly from that map, so the model can be confronted with BAO/CMB/RSD constraints without hand tuning.*

2 Minimal ΦBSU dictionary and operational proxy

We adopt the ΦBSU dictionary: a global phase $\Phi = e^{i\alpha}$ on the Spin$(3,1)$ bundle defines an invariant vacuum 4–density

$$\rho(x) \equiv \|\nabla\alpha\|, \tag{1}$$

and the buoyancy (four–acceleration) field

$$a_\mu(x) \equiv -\partial_\mu \ln \rho. \tag{2}$$

Local radiative/force content is carried by a curvature sector separated from identity/holonomy by

$$A = A_{\rm geom} + A_{\rm id}, \qquad A_{\rm id} := d\alpha, \qquad F_{\rm geom} := dA_{\rm geom}. \tag{3}$$

In this note the metric is used only as an operational chart for invariant separations. When intuition is helpful, one may employ a 1PN–consistent proxy chart $U = c^2\chi$ with $\mathbf{g} = \nabla U$; this is bookkeeping, not ontology.

3 Drag epoch as a coherence opening

3.1 Thermodynamic rates and the drag gate

Let η denote conformal time. CLASS outputs the Thomson opacity derivative $\kappa'(\eta)$, with units Mpc^{-1} in CLASS conventions ($c = 1$). Define the baryon loading

$$R(\eta) = \frac{3\rho_b}{4\rho_\gamma} = \frac{3\,\Omega_b\,a(\eta)}{4\,\Omega_\gamma}. \tag{4}$$

The physical drag rate relevant for baryon slip is

$$\Gamma_{\mathrm{drag}}(\eta) = \frac{\kappa'(\eta)}{1 + R(\eta)}. \tag{5}$$

We define a dimensionless coherence opening gate

$$G(\eta) \equiv \frac{\mathcal{H}(\eta)}{\mathcal{H}(\eta) + \Gamma_{\mathrm{drag}}(\eta)} \in [0, 1], \qquad \mathcal{H} \equiv \frac{a'}{a}. \tag{6}$$

When drag collisions dominate ($\Gamma_{\mathrm{drag}} \gg \mathcal{H}$) one has $G \ll 1$: extended comparator coherence is continuously randomized and no persistent "memory" channel is available. When expansion dominates ($\Gamma_{\mathrm{drag}} \ll \mathcal{H}$) one has $G \to 1$: the coherence channel opens.

Remark 1 (Observable proxy vs. fundamental coupling)**.** *In the* ΦBSU *interpretation, Thomson/drag physics appears here only as an* observable proxy *for a comparator decoherence rate. We do not claim that the global identity sector (A_{id}) "couples" to electromagnetism. Rather, in the standard plasma the dominant mechanism limiting comparator coherence is photon–electron drag, so κ' provides a calibrated clock for the coherence transition.*

3.2 Impulse kernel for drag–latched deposition

The v3 closure replaces post–drag accumulation by an impulse localized at the opening event. Define the (non–negative) normalized impulse kernel

$$\hat{g}(\eta) \equiv \frac{\partial_\eta G(\eta)}{\int d\eta\, \partial_\eta G(\eta)}, \qquad \int d\eta\, \hat{g}(\eta) = 1. \tag{7}$$

Because $\hat{g}$ is a derivative, its peak is expected to lie slightly after the instant $G = 1/2$ (hence slightly after the conventional z_d); this offset is an intrinsic feature of using the drag opening as an event–like write window.

4 Coherence scale and BAO–safe band–pass

Boltzmann solvers provide the diffusion (Silk) damping length $\ell_{\mathrm{D}}(\eta)$. We anchor the coherence length to drag and freeze it for the snapshot,

$$\ell_{\mathrm{coh}} \equiv \ell_{\mathrm{D}}(\eta_d), \qquad k_{\mathrm{D}} \equiv \ell_{\mathrm{coh}}^{-1}. \tag{8}$$

To suppress BAO scales while selecting interference foci around ℓ_{coh} we use a dimensionless band–pass

$$q \equiv k\,\ell_{\mathrm{coh}}, \qquad W(q) = q^2 e^{-q^2}. \tag{9}$$

This peaks at $q = 1$ ($k \simeq k_{\mathrm{D}}$) and suppresses both $q \ll 1$ (BAO/large scales) and $q \gg 1$ (deeply sub–Silk microstructure).

5 Gauge–safe focusing proxy

5.1 Bardeen/Weyl potential and a dimensionless caustic proxy

For practical coding we do not have direct access to α or ρ. In linear scalar perturbations about FLRW, the Bardeen potentials (Φ_B, Ψ_B) are gauge invariant. Define the Weyl (lensing) potential

$$\Phi_W \equiv \frac{\Phi_B + \Psi_B}{2}. \tag{10}$$

We define a dimensionless focusing/caustic readiness proxy

$$\widetilde{C}_{\rm GI}(k,\eta) \equiv \left[\frac{k^2\,\Phi_W(k,\eta)}{\mathcal{H}(\eta)^2}\right]_+, \qquad [x]_+ \equiv \max\{x,0\}. \tag{11}$$

This is already gauge invariant in linear theory (through Φ_W) and avoids the spurious gauge sensitivity of 3D divergences in fixed slicings.

6 Drag–latched memory snapshot $S_d(k)$

We define the drag–latched memory transfer as

$$S_d(k) \equiv \Lambda_0 \int d\eta\; \hat{g}(\eta)\; W(k\ell_{\rm coh})\; \widetilde{C}_{\rm GI}(k,\eta)\; e^{-\Gamma_{\rm min}(\eta_d-\eta)_+}. \tag{12}$$

Here $\Gamma_{\rm min}$ is a small regulator (units Mpc^{-1}) ensuring integrability and allowing for minimal pre–drag leakage. The overall normalization Λ_0 cancels out of the rare–peak statistics below (because all spectral moments scale together), so Λ_0 can be fixed for plotting convenience.

7 From S_d to $\bar{n}$ and $b(z)$ without hand tuning

7.1 Memory field and power spectrum

Treat the drag–latched memory as a (linear) Gaussian random field at drag,

$$S(\mathbf{k}) = S_d(k)\,\mathcal{R}(\mathbf{k}), \tag{13}$$

where $\mathcal{R}$ is the primordial curvature perturbation. Then

$$P_S(k) = |S_d(k)|^2 P_{\mathcal{R}}(k), \qquad \Delta^2_S(k) = \frac{k^3}{2\pi^2} P_S(k). \tag{14}$$

All quantities are computable from the same CLASS transfer outputs used to make the diagnostic plots.

7.2 Spectral moments

Define spectral moments

$$\sigma_j^2 \equiv \int_0^\infty \frac{dk}{2\pi^2}\, k^{2+2j} P_S(k), \qquad j \in \{0,1,2\}. \tag{15}$$

The rms is $\sigma_S \equiv \sigma_0$. Define the BBKS shape parameters

$$R_* \equiv \sqrt{3}\,\frac{\sigma_1}{\sigma_2}, \qquad \gamma \equiv \frac{\sigma_1^2}{\sigma_0\sigma_2}. \tag{16}$$

7.3 One-parameter seeding: rarity ν

Seeds are defined as rare peaks of $S(\mathbf{x})$:

$$S(\mathbf{x}) \geq S_c, \qquad S_c \equiv \nu\, \sigma_S. \tag{17}$$

The dimensionless peak height ν is the *single free parameter* controlling seed rarity. The volume fraction above threshold is

$$f_{\rm seed}(\nu) = \frac{1}{2}\,\mathrm{erfc}\!\left(\frac{\nu}{\sqrt{2}}\right). \tag{18}$$

7.4 Comoving peak number density

Using Gaussian peak theory (BBKS), the differential number density of maxima is

$$\frac{dn_{\rm pk}}{d\nu} = \frac{1}{(2\pi)^2 R_*^3}\, e^{-\nu^2/2}\, G(\gamma, \gamma\nu), \tag{19}$$

where G is the BBKS peak curvature function. The cumulative comoving abundance is

$$\bar{n}(>\nu) = \int_\nu^\infty d\nu'\, \frac{dn_{\rm pk}}{d\nu'}. \tag{20}$$

Equation (20) is parameter free once ν is fixed: $\bar{n}$ is determined by the spectral moments of P_S. For interim estimates one may approximate $\bar{n} \simeq f_{\rm seed}/V_{\rm coh}$ with a coherence volume fixed by the same spectrum (see Appendix A).

7.5 Bias from the cross-spectrum (no hand-set response)

The large-scale modulation of seed abundance (and hence BAO inheritance) is implemented as a peak–background split. Rather than introduce a free response parameter, we compute the long-mode response from the *measured* cross-spectrum between S and the long-mode matter density δ in the same linear run.

For jointly Gaussian fields, the conditional mean obeys

$$\langle S(\mathbf{k}) \mid \delta(\mathbf{k})\rangle = \alpha(k)\, \delta(\mathbf{k}), \qquad \alpha(k) \equiv \frac{P_{S\delta}(k)}{P_{\delta\delta}(k)}. \tag{21}$$

In the present linear construction one can compute the same quantity directly from transfer functions, without any stochastic sampling. Because both fields are sourced by the same primordial curvature mode $\mathcal{R}(\mathbf{k})$,

$$S(\mathbf{k}) = S_d(k)\, \mathcal{R}(\mathbf{k}), \qquad \Delta_m(\mathbf{k}, z) = T_m(k, z)\, \mathcal{R}(\mathbf{k}), \tag{22}$$

where Δ_m is a gauge–invariant comoving matter density contrast (or any fixed, explicitly specified linear density variable used consistently throughout). It follows that

$$\alpha(k, z) = \frac{P_{S\Delta}(k, z)}{P_{\Delta\Delta}(k, z)} = \frac{S_d(k)}{T_m(k, z)} \quad \text{(linear transfer ratio)}. \tag{23}$$

For the large–scale response we take α_L as the low–k limit of $\alpha(k, z_f)$ at formation $z_f \simeq z_d$ (implemented numerically as an average over a fixed low–k band, e.g. $k < k_{\rm BAO}/10$). Because the band–pass filter satisfies $W(q) \propto q^2$ as $q \to 0$, the transfer ratio is generically suppressed on super–BAO scales; in applications to BAO/RSD it is therefore natural to retain the full scale dependence $\alpha(k)$ and interpret Eqs. (23)–(25) as defining a *scale–dependent* response bias. Then a long overdensity δ_L shifts the effective peak height by

$$\nu_{\rm eff} = \frac{S_c - \alpha_L \delta_L}{\sigma_S} = \nu - \frac{\alpha_L}{\sigma_S}\, \delta_L. \tag{24}$$

The Lagrangian bias is therefore

$$b_L(\nu) = \frac{1}{\bar{n}} \frac{\partial \bar{n}}{\partial \delta_L} = -\frac{\alpha_L}{\sigma_S} \frac{\partial \ln \bar{n}}{\partial \nu}, \tag{25}$$

and the Eulerian bias at formation is

$$b_E(\nu) = 1 + b_L(\nu). \tag{26}$$

If seeds form at $z_f \simeq z_d$ and thereafter evolve passively, a standard passive evolution model gives

$$b(z) = 1 + [b_E(\nu; z_f) - 1] \frac{D(z_f)}{D(z)}, \tag{27}$$

where $D(z)$ is the linear growth factor of the baseline cosmology. Equation (27) provides the immediate bridge to RSD tests.

7.6 Normalization invariance (removing Λ_0)

The overall amplitude Λ_0 in Eq. (12) is *not* an additional physical parameter. If $S_d \to c\, S_d$ for any constant c, then $P_S \to c^2 P_S$ and $\sigma_j \to c\sigma_j$, but the BBKS shape parameters R_* and γ in Eq. (16) are unchanged. The threshold definition $S_c \equiv \nu\sigma_S$ in Eq. (17) rescales in the same way, and the response ratio α_L/σ_S in Eq. (25) is invariant. Consequently, $\{\bar{n}(\nu), b(z;\nu)\}$ depend only on the *shape* of $S_d(k)$, not its arbitrary normalization. In implementation one may therefore fix Λ_0 by a convenient plotting normalization (e.g. $\sigma_S \equiv 1$) without affecting any published predictions.

7.7 Illustrative ν scan (production BBKS abundance)

For orientation, the current "production" reference run (Planck–2018 baseline; drag–latched deposition; BBKS moments and BBKS peak abundance) gives the following mapping between rarity and comoving abundance. For the bias column we list the standard rare–tail closed form $b_E^{(\mathrm{rt})} \simeq 1 + \nu - 1/\nu$ as a compact *prior* for tracer strength; the fully scale-dependent response bias is computed from Eq. (23)–(25).

ν	$\bar{n}(>\nu)$ [Mpc^{-3}]	$1/\bar{n}$ [Mpc^3]	$b_E^{(\mathrm{rt})}$	heuristic label
2.0	4.73×10^{-6}	2.11×10^{5}	2.50	group/cluster-like
3.0	1.07×10^{-6}	9.36×10^{5}	3.67	rare massive seeds
4.0	6.42×10^{-8}	1.56×10^{7}	4.75	SMBH-scale rarity
5.0	1.17×10^{-9}	8.52×10^{8}	5.80	extreme rarity

These values are deterministic once the memory spectrum is fixed; there are no hand-set abundances.

8 Numerical results: acceptance checks and figures

8.1 BAO/CMB hygiene criteria

The closure is acceptable only if standard photon–baryon acoustic physics remains intact. The working acceptance checks are:

1. **(A) Sound horizon unchanged.** The CLASS value r_d matches the Planck–2018 baseline to numerical tolerance.

2. **(B) Scale separation.** $S_d(k)$ peaks near $k \sim k_{\mathrm{D}}$ and satisfies $k_{\mathrm{peak}} \gg k_{\mathrm{BAO}} \equiv 2\pi/r_d$.

3. **(C) BAO inheritance (tracer, not driver).** BAO wiggle *phase* is unchanged in $P(k)/P_{\rm smooth}(k)$, and the configuration–space BAO bump positions agree once shot noise is removed.

4. **(D) Drag localization.** The impulse kernel $\hat{g}$ is narrow and concentrated near η_d.

8.2 Production figure (BBKS moments, transfer ratio, BAO acceptance)

Figure 1 shows the current "production" diagnostic summary generated by the pipeline summarized in Appendix B. It demonstrates in a single view: (i) a narrow drag impulse kernel $\hat{g}$ and drag-latched seeding, (ii) a memory spectrum peaked near the Silk/diffusion scale and well separated from BAO, (iii) BBKS spectral moments and the resulting peak abundance $\bar{n}(\nu)$, (iv) the scale-dependent transfer ratio $\alpha(k)$, and (v) BAO inheritance checks in both Fourier and configuration space.

For completeness we include the v2 BAO inheritance check as an appendix figure (Fig. 3); it illustrates that once shot noise is removed from $\xi(r)$, a biased tracer and the seed tracer share the same BAO peak position as the matter field (bias rescales amplitude, not the peak location).

9 Outlook: from drag seeds to supermassive BECOs (Tardis–buoyancy)

The present note stops at a *linear* seeding map. The next required programme step is nonlinear growth to supermassive compact centres. Two modelling points are essential for consistency with ΦBSU:

1. **"Gravity as coordinates" does not remove *invariants*.** In ΦBSU the metric is a chart representation, but the buoyancy gradient $a_\mu = -\partial_\mu \ln \rho$ is a local invariant field. Therefore any "Tardis" interior picture must be formulated in terms of coordinate-invariant observables: redshift factors, tidal invariants, null congruence focusing (Raychaudhuri/Sachs), and invariant volumes computed from the operational metric. Large interior volume is not a gauge artefact if it appears in such invariants.

2. **Tardis–buoyancy as a saturation/response problem.** A practical modelling route is to posit a ΦBSU-consistent saturation of focusing in the ρ-sector: as $|\nabla \ln \rho|$ grows, caustic focusing is cut off and an effective stiff/ultradense "lattice" phase develops. One may implement this as a nonlinear closure for a_μ or for a tidal invariant $T_{\mu\nu} = \nabla_\mu a_\nu$, calibrated so that the exterior is BH-mimicking while the interior admits an enlarged proper volume ("Tardis") and a fluid outer layer.

This growth stage must be confronted with SMBH demographics, early galaxy assembly, and (eventually) GW ringdown constraints; it is logically downstream from the BAO/CMB-safe seeding map provided here.

A Coherence volume approximation (interim)

For quick estimates before implementing the full BBKS curvature function $G(\gamma, \gamma\nu)$, one may approximate the coherence cell size by the spectrum-defined scale

$$R_{\rm coh} \equiv \frac{\sigma_0}{\sigma_1}, \qquad V_{\rm coh} \equiv (2\pi)^{3/2} R_{\rm coh}^3, \tag{28}$$

and use $\bar{n} \simeq f_{\rm seed}/V_{\rm coh}$. This introduces *no* new free scale: $R_{\rm coh}$ is fixed by P_S. The main text recommends switching to the BBKS formula (20) for the publication pipeline.

B Pipeline unit audit and code artefacts

For transparency and reproducibility, the numerical prototype was built as a thin wrapper around CLASS/`classy` outputs (thermodynamics, background, scalar perturbations), followed by post-processing steps that implement the filters and rare-peak mapping described in the main text.

Script artefacts. The working directory contains small, single-file scripts that mirror the equation chain in the main text:

- `beco_v2.py`: drag-latched deposition, peak-background split modulation, and diagnostic plots.
- `beco_v2_bao_diagnostics.py`: configuration-space BAO peak localisation with explicit shot-noise removal and stability checks.
- `beco_pipeline.py`: a compact "unit-audit" pipeline that reproduces the gate/rate/coherence hierarchy and verifies dimensional consistency.
- `beco_production.py`: "production" pipeline that computes BBKS spectral moments $(\sigma_0, \sigma_1, \sigma_2, R_*, \gamma)$, BBKS peak abundance $\bar{n}(\nu)$, transfer ratio $\alpha(k)$, and BAO acceptance diagnostics; writes `beco_production.png` and `beco_summary.json`.

All scripts assume CLASS conventions ($c = 1$, conformal time in Mpc, k in Mpc^{-1}, $\mathcal{H}$ in Mpc^{-1}) and require a local `classy` installation.

Unit-audit figure. Figure 2 illustrates the key intermediate quantities as functions of redshift for representative modes, and serves as a regression test for the end-to-end pipeline (rates, diffusion scale, focusing proxy, filtered focusing, and memory depth). The seeding probability in panel (f) is a *diagnostic* nonlinearity (logistic excursion) and is not used for the publishable ν-peak mapping in Sec. 20–27.

C v2 BAO inheritance check (prototype figure)

References

[1] C.-P. Ma and E. Bertschinger, *Cosmological perturbation theory in the synchronous and conformal Newtonian gauges*, Astrophys. J. **455** (1995) 7.

[2] W. Hu and N. Sugiyama, *Small Scale Cosmological Perturbations: an Analytic Approach*, Astrophys. J. **444** (1995) 489.

[3] Planck Collaboration, *Planck 2018 results. VI. Cosmological parameters*, Astron. Astrophys. **641** (2020) A6.

[4] J. M. Bardeen, J. R. Bond, N. Kaiser and A. S. Szalay, *The statistics of peaks of Gaussian random fields*, Astrophys. J. **304** (1986) 15.

[5] E. Säkkinen, ΦBSU — *Buoyant Separverse Unification — Field Theory, Part I* (updated 2026).

[6] E. Säkkinen, *A New 4-Dimensional Cosmological Principle and Buoyant Spacetime* (2024).

[7] E. Säkkinen, *A New Reflection-Graded Hypersymmetry on* $M_4 \times K_2$ (updated 2026).

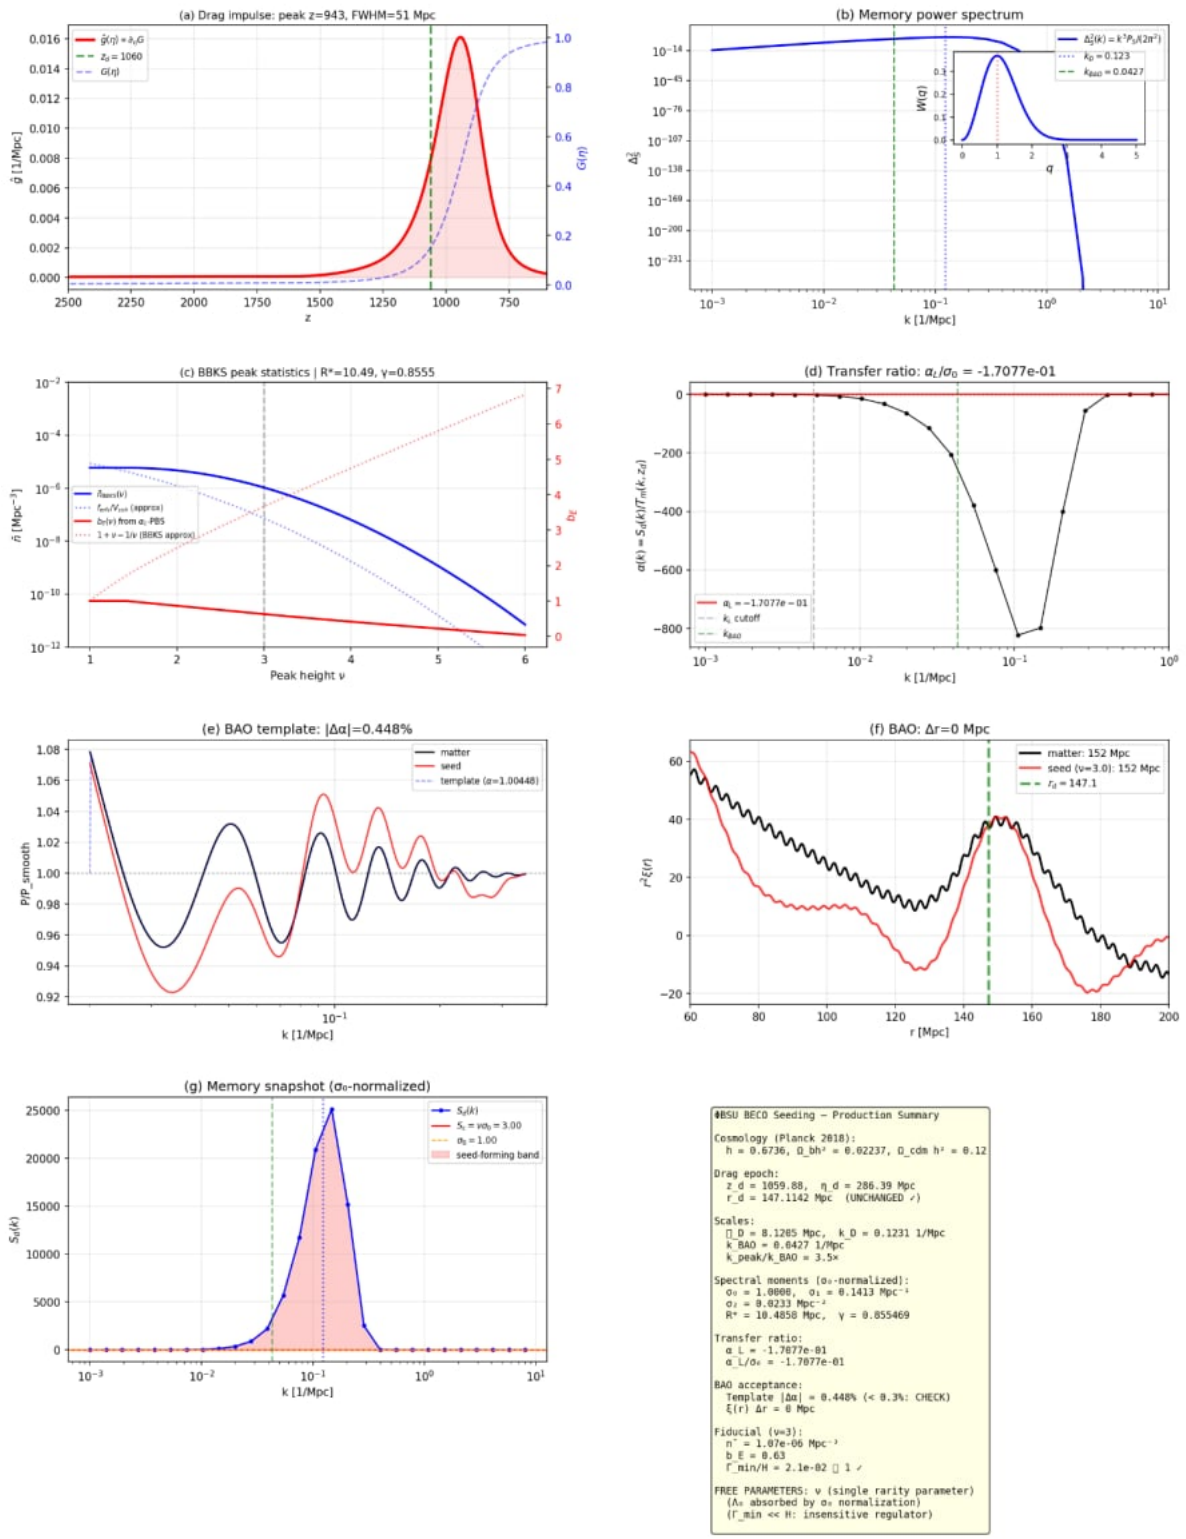

Figure 1: ΦBSU BECO seeding "production" summary figure (CLASS baseline; σ_0-normalized). Panel (a): drag impulse kernel $\hat{g}(\eta) \propto \partial_\eta G(\eta)$ (peak near $z \simeq 943$, FWHM $\simeq 51$ Mpc in conformal distance). Panel (b): dimensionless memory power Δ_S^2 with the fixed band-pass window $W(q)$. Panel (c): BBKS peak abundance $\bar{n}(>\nu)$ (blue) and response-bias diagnostics. Panel (d): transfer ratio $\alpha(k)$ (Eq. (23)); its low-k suppression reflects the built-in $W \propto q^2$ behaviour. Panels (e)–(f): BAO "hygiene" checks: wiggle-phase alignment and BAO bump position in $\xi(r)$ (shot-noise removed). Panel (g): the drag-latched memory snapshot $S_d(k)$ after σ_0 normalization; the shaded band indicates the seed-forming k-range for a representative ν. The inset summary lists the key numerical invariants of the run.

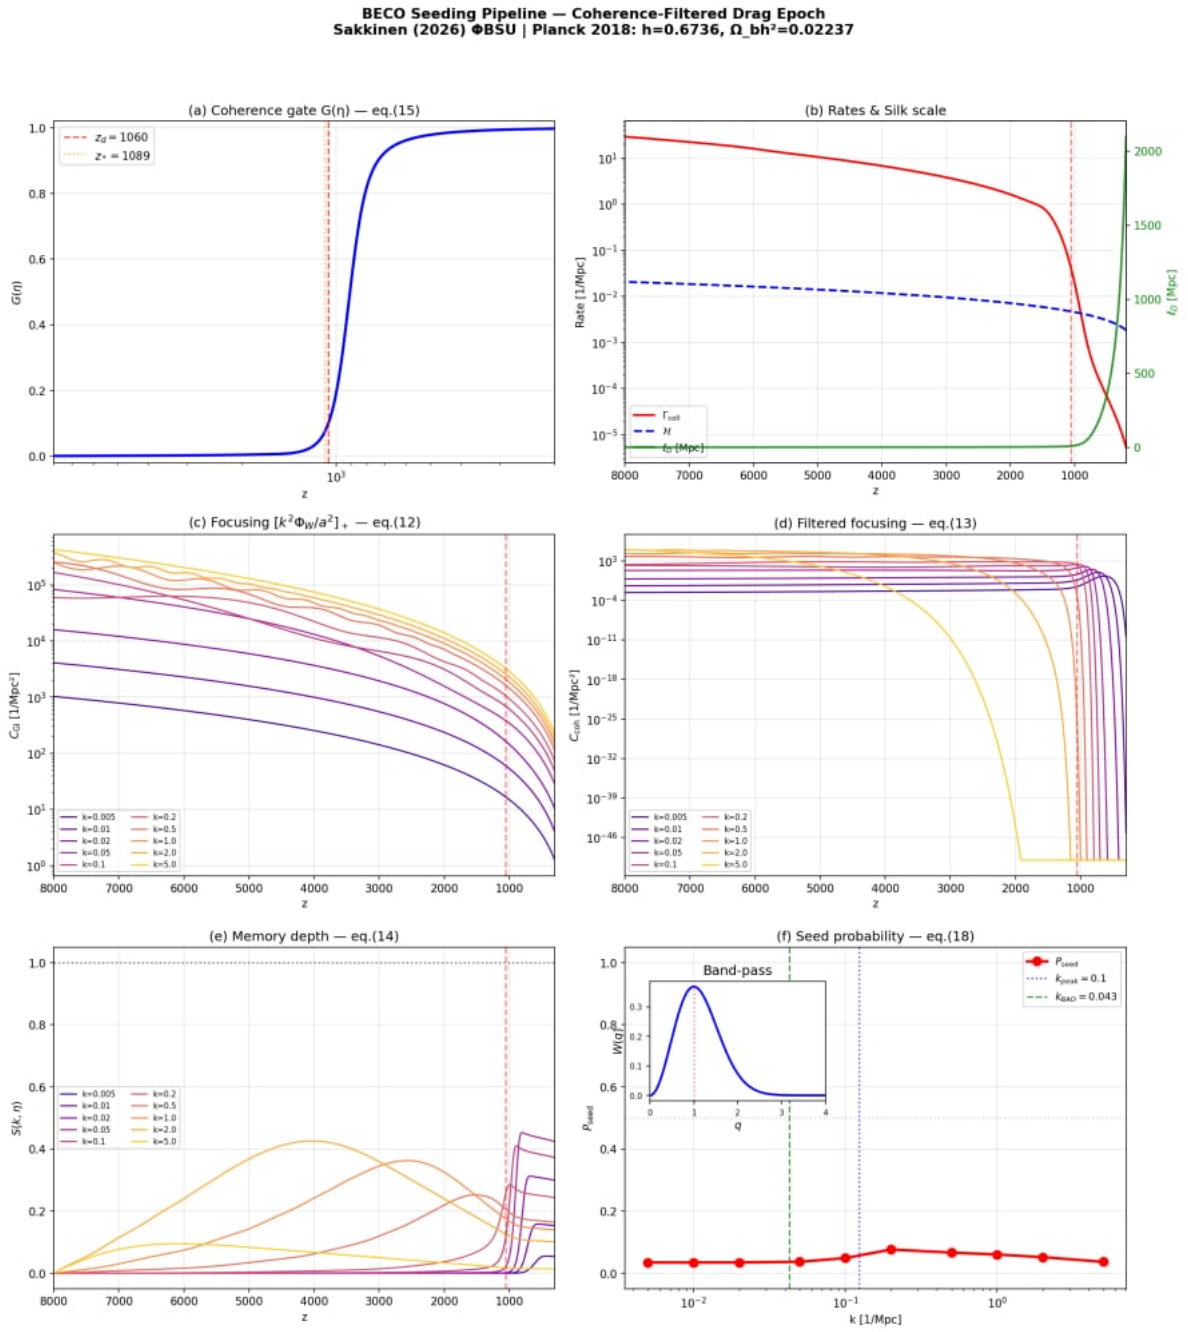

Figure 2: "Unit-audit" pipeline figure generated by `beco_pipeline.py`. Panels show: (a) the coherence gate $G(\eta)$, (b) the collision/expansion rates and the diffusion length ℓ_D, (c) the Weyl-potential-based focusing proxy, (d) the band-pass filtered focusing $C_{\rm coh}$, (e) the memory depth evolution $S(k,\eta)$, and (f) a diagnostic seed probability at drag. The main text replaces this diagnostic excursion by the rare-peak mapping with $S_c = \nu\sigma_S$.

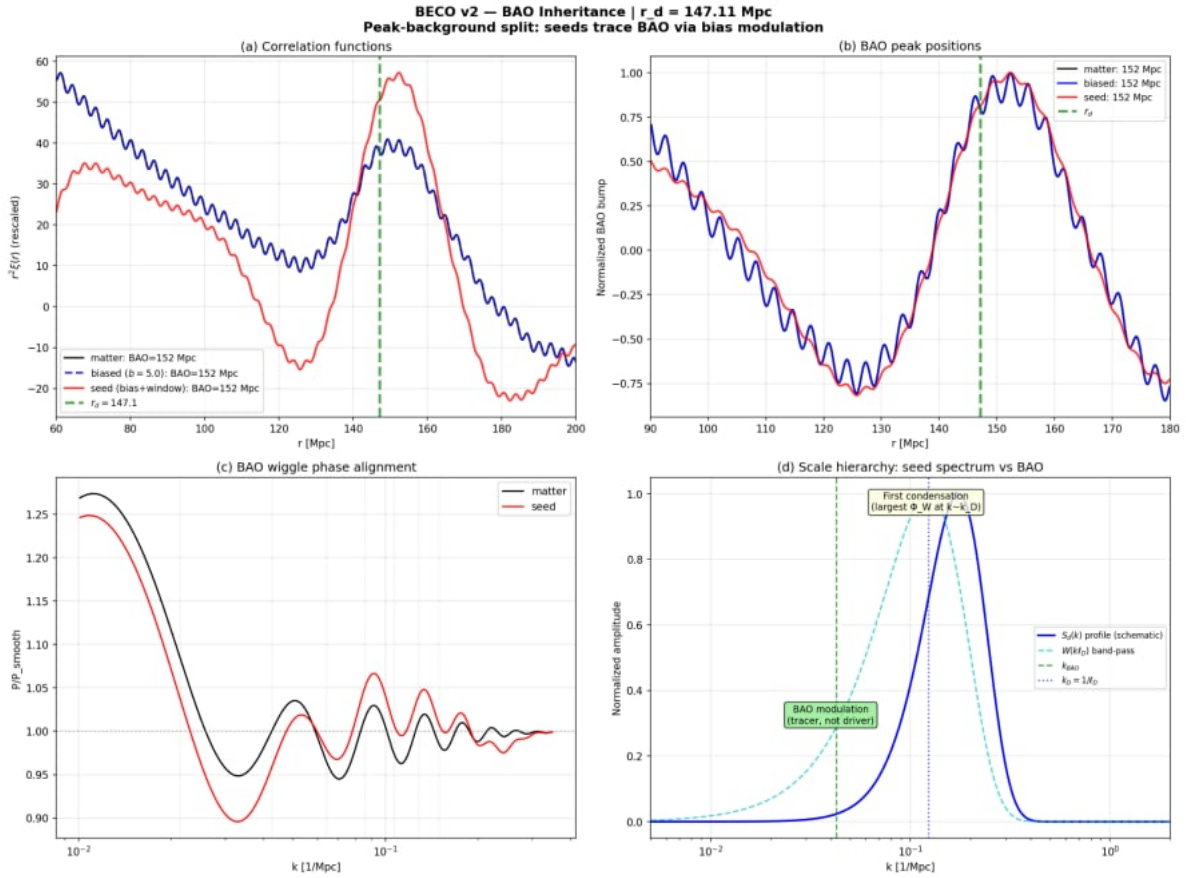

Figure 3: BAO inheritance check for the v2 prototype. The biased tracer and seed peaks align with the matter BAO peak once shot noise is removed from $\xi(r)$ (a constant $1/\bar{n}$ contributes only to $\xi(0)$).