Final Review Recitation Handout

Fundamentals of Observational Cosmology

A 1-hour recap based on the homework from Lectures 2–16

This handout is not a point-by-point re-solution of every homework problem. Instead, it organizes the most exam-relevant chains of ideas that kept appearing across the assignments:

background geometry $\longrightarrow$ observational distances and clustering $\longrightarrow$ statistical inference $\longrightarrow$ survey modeling $\longrightarrow$ inflation / CMB / CAMB.

What the definitions are, where the approximations enter, how the formulas connect, and what the physical picture is.

Homework roadmap from Lecture 2 to Lecture 16

Lecture range	High-frequency homework themes
2–4	Cosmological redshift, conformal time, distance duality, FRW Christoffel symbols, the continuity equation, the condition for $\Phi = \Psi$, and the linear growth equation.
5–8	Supernova distance modulus, the low-redshift limit, the M–H_0 degeneracy, transverse/radial BAO, AP distortion, the Kaiser effect, and FoG.
9–10	Bayes, sample mean/variance/covariance, matrix χ^2, two-dimensional confidence regions, and the difference between a Fisher forecast and a full likelihood.
11–14	Kaiser multipoles, reconstruction, effective volume, weighted pair counts, effective redshift, the window matrix, forward modeling, and the logic of fixed-template BAO / ShapeFit / EFT modeling.
15–16	Slow-roll inflation, running, local PNG, Boltzmann codes, tight coupling, line-of-sight integration, and the CAMB workflow.

Essential formulas to memorize before the exam

Background and distances

$$1+z=\frac{a_0}{a}=\frac{1}{a},\qquad \tau(a)=\int_0^a\frac{\mathrm{d}a'}{a'^2H(a')},\qquad D_L=(1+z)^2D_A.$$

Fluids and growth

$$\dot\rho+3H(\rho+P)=0\quad\Longrightarrow\quad\rho\propto a^{-3(1+w)},$$

$$\ddot\delta_m+2H\dot\delta_m-4\pi G\bar\rho_m\,\delta_m=0,\qquad\text{and in EdS }\delta_m\propto a.$$

Supernovae and BAO

$$\mu=5\log_{10}\left(\frac{D_L}{\mathrm{Mpc}}\right)+25,\qquad D_L(z\ll1)\simeq\frac{cz}{H_0},$$

$$\Delta\theta\simeq\frac{r_s}{D_M(z)},\qquad \Delta z\simeq\frac{H(z)r_s}{c}.$$

RSD and multipoles

$$P_g^s(k,\mu)=(b+f\mu^2)^2P_m(k)=b^2(1+\beta\mu^2)^2P_m(k),\qquad \beta\equiv\frac{f}{b}.$$

$$P_0=\left(b^2+\frac{2}{3}bf+\frac{1}{5}f^2\right)P_m,\quad P_2=\left(\frac{4}{3}bf+\frac{4}{7}f^2\right)P_m,\quad P_4=\frac{8}{35}f^2P_m.$$

Statistical inference

$$\chi^2=\boldsymbol{\Delta}^T\mathbf{C}^{-1}\boldsymbol{\Delta},\qquad P(\boldsymbol\theta|D)\propto\mathcal{L}(D|\boldsymbol\theta)\,\pi(\boldsymbol\theta).$$

For a two-parameter joint confidence region, the standard numbers are

$$68\%\leftrightarrow\Delta\chi^2=2.30,\quad 95\%\leftrightarrow 6.18,\quad 99.7\%\leftrightarrow 11.83.$$

Fisher and surveys

$$F_{ij}=\frac{1}{2}\mathrm{Tr}\left(\boldsymbol\Sigma^{-1}\boldsymbol\Sigma_{,i}\boldsymbol\Sigma^{-1}\boldsymbol\Sigma_{,j}\right)+\boldsymbol\mu_{,i}^T\boldsymbol\Sigma^{-1}\boldsymbol\mu_{,j},\qquad \mathbf{F}_{\rm tot}=\mathbf{F}_1+\mathbf{F}_2.$$

$$w_{\rm FKP}(\mathbf{r})=\frac{1}{1+\bar n(\mathbf{r})P(k)},\qquad V_{\rm eff}(k)=\int\mathrm{d}^3r\left[\frac{\bar n(\mathbf{r})P(k)}{1+\bar n(\mathbf{r})P(k)}\right]^2.$$

$$P^{\rm obs}_{\ell,i}=\sum_{\ell',j}W_{ij}^{\ell\ell'}\,P^{\rm th}_{\ell'}(k_j).$$

Inflation / CAMB

$$r\simeq\frac{8}{N_\star},\qquad n_s\simeq1-\frac{2}{N_\star}\qquad (V=\tfrac{1}{2}m^2\phi^2),$$

$$\Delta b_{\rm PNG}\propto\frac{f_{\rm NL}^{\rm loc}}{k^2},\qquad \Theta_\ell(k,\tau_0)=\int_0^{\tau_0}\mathrm{d}\tau\,S_T(k,\tau)\,j_\ell[k(\tau_0-\tau)].$$

Module A: background geometry and linear growth (Lectures 2–4)

A1. The three things most likely to be tested in background geometry

(i) Redshift is not just ordinary special-relativistic Doppler shift.
In an FRW universe, the more fundamental relation is

$$1 + z = \frac{\lambda_{\rm obs}}{\lambda_{\rm em}} = \frac{a_0}{a_{\rm em}}.$$

For a very nearby source one can approximate the effect as Doppler, but on cosmological scales the natural interpretation is that the wavelength is stretched by the expansion of spacetime.

(ii) Conformal time is the time variable that makes many evolution equations cleaner.
From

$$\mathrm{d}\tau = \frac{\mathrm{d}t}{a}$$

we immediately get

$$\mathcal{H} \equiv \frac{a'}{a} = aH, \qquad \frac{\mathrm{d}}{\mathrm{d}\tau} = a\frac{\mathrm{d}}{\mathrm{d}t}.$$

(iii) Distance duality is one of the best derivation questions for an exam.
It connects the luminosity distance and the angular-diameter distance:

$$D_L = (1+z)^2 D_A.$$

A2. Representative problem: derive $D_L = (1+z)^2 D_A$ in one chain

Step 1: define the luminosity distance using the flux.

$$F = \frac{L}{4\pi D_L^2}.$$

When photons arrive at the observer, **the energy of each photon** is reduced by a factor $(1+z)^{-1}$ and **the arrival rate** is also reduced by $(1+z)^{-1}$. Therefore

$$F = \frac{L}{4\pi D_M^2 (1+z)^2}.$$

Hence

$$D_L = (1+z) D_M.$$

Step 2: define the angular-diameter distance from angular size.
If the physical transverse size of the source at emission is ℓ and the observed angle is θ, then

$$D_A = \frac{\ell}{\theta}.$$

FRW geometry gives

$$\ell = a_{\rm em} D_M \theta = \frac{D_M}{1+z}\theta,$$

so

$$D_A = \frac{D_M}{1+z}.$$

Step 3: combine the two results

$$\boxed{D_L = (1+z)^2 D_A.}$$

Common mistakes
1. Be explicit about whether you mean D_M, D_A, or D_L; do not mix up the transverse comoving distance and the angular-diameter distance.
2. State clearly where the two factors of $(1+z)^{-1}$ in the luminosity distance come from: **energy redshift** and **time dilation**.
3. This relation relies on light propagating along null geodesics, photon-number conservation, and an ordinary metric theory.

A3. Representative problem: why the growth equation has exactly that structure

A classic homework derivation is to connect the continuity, Euler, and Poisson equations:

$$\delta_m' = -\theta_m + 3\Phi', \qquad \theta_m' + \mathcal{H}\theta_m = k^2\Psi.$$

On sub-horizon scales, the potentials vary slowly, so one often neglects $3\Phi'$, and in the absence of anisotropic stress one has $\Phi = \Psi$. Using the Fourier-space Poisson equation

$$-k^2\Psi = 4\pi G a^2 \bar{\rho}_m \, \delta_m,$$

we obtain

$$\delta_m'' + \mathcal{H}\delta_m' - 4\pi G a^2 \bar{\rho}_m \, \delta_m = 0.$$

Changing back from conformal time to cosmic time,

$$\boxed{\ddot{\delta}_m + 2H\dot{\delta}_m - 4\pi G \bar{\rho}_m \, \delta_m = 0.}$$

In an Einstein–de Sitter universe, substitute

$$\delta_m \propto a^n$$

and one finds the growing mode

$$\boxed{\delta_m \propto a.}$$

The exam-level summary of this section is really just three lines:
(1) $\Phi = \Psi$ only holds when there is **no anisotropic stress**;
(2) the three terms in the growth equation correspond to inertia, Hubble friction, and gravitational clustering;
(3) in EdS, the growing mode is $\delta \propto a$, which is the standard benchmark for quick estimates.

Module B: SN, BAO, AP, and RSD (Lectures 5–8)

B1. Supernovae: why they measure shape well but do not set the absolute scale by themselves

In the low-redshift limit,

$$D_L(z) \simeq \frac{cz}{H_0}, \qquad \mu = 5\log_{10}\left(\frac{D_L}{\text{Mpc}}\right) + 25.$$

Substituting $D_L = (c/H_0)d_L(z; \Omega_\mathrm{m}, \Omega_\Lambda, \cdots)$ into the apparent magnitude

$$m(z) = M + \mu(z)$$

gives

$$m(z) = M - 5\log_{10} H_0 + (\text{terms that depend only on redshift and shape parameters}).$$

Therefore supernova data naturally constrain a combination

$$\mathcal{M} \equiv M - 5\log_{10} H_0 + \text{constant},$$

not M alone and not H_0 alone. This is exactly the M–H_0 **degeneracy** emphasized throughout the homework.

B2. The two BAO observables: transverse and radial

The most useful approximations from the homework are

$$\Delta\theta \simeq \frac{r_s}{D_M(z)}, \qquad \Delta z \simeq \frac{H(z) r_s}{c}.$$

For example, at $z = 0.8$ take

$$r_s = 147\,\mathrm{Mpc}, \qquad D_M = 2800\,\mathrm{Mpc}, \qquad H(z) = 109\,\mathrm{km\,s^{-1}\,Mpc^{-1}},$$

then

$$\Delta\theta \simeq \frac{147}{2800} \simeq 0.0525 \text{ rad} \simeq 3.0^\circ,$$
$$\Delta z \simeq \frac{109 \times 147}{2.998 \times 10^5} \simeq 0.053.$$

Physical picture:
BAO is a **standard ruler**. The transverse mode gives D_M/r_s, while the radial mode gives $H(z)r_s$.
That is why BAO is simultaneously sensitive to distance and to the expansion rate.

B3. What do BAO, AP, and RSD each really measure?

Effect	Physical origin	Main sensitivity
BAO	The standard ruler left by primordial photon–baryon sound waves	Geometry: $D_M(z)$ and $H(z)$
AP	Using the wrong fiducial cosmology when converting angles and redshifts to distances	**Pure geometric** distortion
RSD	Peculiar velocities along the line of sight modify the observed redshift	**Growth** and dynamics, usually written in terms of f or $f\sigma_8$

A minimal AP notation is

$$s_\perp^{\mathrm{fid}} = \alpha_\perp s_\perp, \qquad s_\parallel^{\mathrm{fid}} = \alpha_\parallel s_\parallel.$$

If $\alpha_\parallel < 1$, the fiducial cosmology is underestimating the radial scale.

B4. The single most important RSD formula for exams: Kaiser

In linear theory,

$$P_g^s(k,\mu) = \left(b + f\mu^2\right)^2 P_m(k) = b^2(1+\beta\mu^2)^2 P_m(k), \qquad \beta = \frac{f}{b}.$$

Expanding into Legendre multipoles gives

$$P_0(k) = \left(b^2 + \frac{2}{3}bf + \frac{1}{5}f^2\right) P_{\rm m}(k),$$

$$P_2(k) = \left(\frac{4}{3}bf + \frac{4}{7}f^2\right) P_{\rm m}(k),$$

$$P_4(k) = \frac{8}{35}f^2 P_{\rm m}(k).$$

If only the monopole coefficient is required, the most common derivation is

$$P_0(k) = \frac{1}{2}\int_{-1}^{1} P_g^s(k,\mu)\,\mathrm{d}\mu = b^2\left(1 + \frac{2\beta}{3} + \frac{\beta^2}{5}\right) P_{\rm m}(k).$$

Common RSD pitfalls
1. Kaiser describes **large-scale linear** coherent infall; FoG describes **small-scale nonlinear** random motions inside halos.
2. The anisotropy from AP is not the same as the anisotropy from RSD; the first is a wrong geometric mapping, the second is real velocity information.
3. In an exam answer, it is very helpful to mention both pictures explicitly: **large-scale squashing** and **small-scale elongation**.

Module C: Bayes, χ^2, covariance, and Fisher (Lectures 6, 9, 10)

C1. The truly essential formulas in statistical inference are not many

Bayes' theorem

$$P(\boldsymbol{\theta}|D) = \frac{\mathcal{L}(D|\boldsymbol{\theta})\,\pi(\boldsymbol{\theta})}{Z}.$$

Independent Gaussian errors

$$\chi^2 = \sum_i \frac{(d_i - m_i)^2}{\sigma_i^2}.$$

Correlated data

$$\chi^2 = \boldsymbol{\Delta}^T \mathbf{C}^{-1} \boldsymbol{\Delta}, \qquad \boldsymbol{\Delta} = d - m.$$

Joint confidence regions in a two-parameter plane

$$68\% \leftrightarrow \Delta\chi^2 = 2.30, \qquad 95\% \leftrightarrow 6.18.$$

C2. Why is two-dimensional 68% not $\Delta\chi^2 = 1$?

This is one of the favorite conceptual questions from Lectures 6 and 10. The reason is not "because someone chose it", but because a **one-dimensional interval probability** and a **two-dimensional enclosed probability** are not the same thing. For two parameters, $\Delta\chi^2$ follows a chi-square distribution with 2 degrees of freedom, so the contour enclosing 68% of the **joint probability** corresponds to

$$\boxed{\Delta\chi^2 = 2.30}$$

rather than 1. The value 1 is only the familiar 1σ number for a one-dimensional Gaussian.

C3. What do the two Fisher terms mean physically?

The general Gaussian Fisher matrix is

$$F_{ij} = \frac{1}{2}\mathrm{Tr}\Big(\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_{,i}\boldsymbol{\Sigma}^{-1}\boldsymbol{\Sigma}_{,j}\Big) + \boldsymbol{\mu}_{,i}^T\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}_{,j}.$$

The two terms mean:

- **The first term: information from parameter dependence of the covariance itself.** For observables such as CMB or weak-lensing power spectra, a lot of information can live in how the variance/covariance changes.
- **The second term: information from parameter dependence of the mean signal.** For observables such as supernova or BAO distance points, the main effect is often that the theoretical mean moves in data space.

If the covariance is approximately fixed, one often keeps only the second term.

C4. Why is Fisher used so often?

Because two homework conclusions are extraordinarily useful:

1. **The parameter covariance is approximately**

$$\mathbf{C}_\theta \approx \mathbf{F}^{-1}.$$

2. **The Fisher matrices of independent data sets can be added directly**

$$\boxed{\mathbf{F}_{\mathrm{tot}} = \mathbf{F}_1 + \mathbf{F}_2.}$$

So Fisher is ideal for survey design and for building intuition about which data set breaks which degeneracy.

C5. But when does Fisher fail?

The two most typical cases are:

- **When a parameter lies close to a physical boundary.** For example $\sum m_\nu \geq 0$ or $r \geq 0$. The true posterior is truncated by the boundary, whereas Fisher gives a symmetric Gaussian.

- **When the posterior is strongly curved or non-Gaussian.** For example a banana-shaped degeneracy. Fisher only captures the local tangent ellipse.

A one-line distinction:
A **Fisher forecast** asks "how well could the experiment do near a chosen fiducial point?"
A **full likelihood inference** asks "where in parameter space do the actual data really prefer?"
The first is fast; the second is the real answer.

Module D: practical survey analysis (Lectures 11–14)

D1. Why does the effective volume appear?

The FKP weight is

$$w_{\mathrm{FKP}}(\mathbf{r}) = \frac{1}{1+\bar{n}(\mathbf{r})P(k)},$$

and the corresponding effective volume is

$$\boxed{V_{\mathrm{eff}}(k) = \int \mathrm{d}^3r \left[\frac{\bar{n}(\mathbf{r})P(k)}{1+\bar{n}(\mathbf{r})P(k)}\right]^2}.$$

The most important thing to remember about this formula is the two limits:

- If $\bar{n}P \gg 1$, then
$$\frac{\bar{n}P}{1+\bar{n}P} \to 1,$$
so
$$V_{\mathrm{eff}} \to V_{\mathrm{survey}}.$$
This means the sample is very dense, and the error budget is dominated not by shot noise but by cosmic variance.
- If $\bar{n}P \ll 1$, then
$$\frac{\bar{n}P}{1+\bar{n}P} \approx \bar{n}P,$$
so
$$V_{\mathrm{eff}} \propto \int (\bar{n}P)^2 \mathrm{d}^3r,$$
which is much smaller than the geometric volume: shot noise is severe.

D2. Why is reconstruction written as $\delta_{\mathrm{rec}} = \delta_d - \delta_s$?

The four-step logic of reconstruction is ideal for a short-answer question:

1. Use the smoothed density field to estimate the large-scale displacement;
2. move the galaxies backward to obtain the displaced field, δ_d;
3. move the random catalog using the same mask and selection to obtain the shifted field, δ_s;
4. subtract them:
$$\delta_{\mathrm{rec}} = \delta_d - \delta_s.$$

The point of the subtraction is to remove the smooth survey geometry and mean-density background, leaving only the reconstructed overdensity.

Rec-Iso vs. Rec-Sym:

- **Rec-Iso** tends to remove large-scale Kaiser anisotropy as well, so the post-reconstruction quadrupole is often small; this is convenient for BAO-only analyses.
- **Rec-Sym** preserves more RSD information, so the post-reconstruction quadrupole need not be small; this is useful when modeling monopole and quadrupole together.

D3. The window matrix is the master key of the second half of the course

Once the continuous convolution is discretized, it becomes

$$\boxed{P^{\rm obs}_{\ell,i} = \sum_{\ell',j} W^{\ell\ell'}_{ij} P^{\rm th}_{\ell'}(k_j).}$$

This has two major consequences:

1. **The survey window mixes multipoles.**
   Even if the data vector contains only P_0 and P_2, the theoretical forward model often still needs P_4, because after convolution the higher multipoles leak into the lower observed ones.
2. **The theory grid should be finer than the data grid.**
   Convolution, interpolation, and binning all mix neighboring wavenumbers. If the theory grid is too coarse, the convolution itself becomes contaminated by discretization error.

D4. One-sentence summary of Lecture 14: modeling is a trade-off between information and robustness

Modeling route	One-line interpretation
Fixed-template BAO	Few parameters and high robustness; mainly extracts geometric information from the BAO peak.
ShapeFit	Uses more broadband shape information than BAO-only, while keeping a comparatively compact parametrization.
Direct EFT full modeling	More theory freedom, access to higher k and more information, but much stronger demands on counterterms, bias control, and systematics.

The key message to remind students of in a recitation is this:
What is measured is not a "clean theory", but rather **theory + survey window + estimator + reconstruction convention + covariance**.
The homework in the second half of the course was really training this forward-model mindset.

Module E: inflation, PNG, CMB, and CAMB (Lectures 15–16)

E1. The slow-roll inflation result most likely to appear directly on an exam

For the quadratic potential

$$V(\phi) = \frac{1}{2}m^2\phi^2,$$

the homework asked students to derive

$$r \simeq \frac{8}{N_\star}, \qquad n_s \simeq 1 - \frac{2}{N_\star}.$$

Thus

$$N_\star = 50: \quad r \simeq 0.16, \quad n_s \simeq 0.96,$$
$$N_\star = 60: \quad r \simeq 0.133, \quad n_s \simeq 0.967.$$

Physical conclusion:
The value of n_s is reasonable, but the predicted r is rather large, so simple large-field monomial models are less favored than plateau-type single-field slow-roll models.

E2. Why is Lyα more useful for running than for amplitude?

The point of this homework question is not to memorize a slogan, but to understand the phrase **lever arm in scale**.
A_s mainly controls the overall normalization, while the running α_s controls how the spectrum bends as a function of $\ln k$. Lyα extends the constraint from CMB scales to much smaller scales, so it is much more sensitive to *how the spectrum bends*, not just to yet another overall amplitude measurement.

E3. Why does local PNG prefer large scales and highly biased tracers?

The key scaling in the homework is

$$\Delta b_{\rm PNG} \propto \frac{f_{\rm NL}^{\rm loc}}{k^2}.$$

If $k_2 = 2k_1$, then

$$\frac{\Delta b(k_2)}{\Delta b(k_1)} = \frac{1}{4}.$$

So doubling the wavenumber leaves only one quarter of the signal. That is why local PNG is all about **using the smallest possible** k. In addition, highly biased tracers often bring an extra enhancement proportional to $(b-1)$, which is why QSO-like samples are especially valuable.

E4. CAMB is really just a very clear pipeline

The four CAMB homework questions can be compressed into one picture:

set the background cosmology and primordial spectrum
$\rightarrow$ choose initial conditions
$\rightarrow$ evolve the Einstein–Boltzmann equations
$\rightarrow$ construct source functions
$\rightarrow$ project to C_ℓ and $P(k)$.

The single formula students should remember most clearly is the line-of-sight expression:

$$\Theta_\ell(k, \tau_0) = \int_0^{\tau_0} \mathrm{d}\tau \, S_T(k,\tau)\, j_\ell[k(\tau_0 - \tau)].$$

Its meaning is:

- **Cosmology dependence** is largely compressed into the source function $S_T(k,\tau)$;
- **geometric projection** is largely encoded by the spherical-Bessel kernel j_ℓ.

This is exactly why the line-of-sight method is much faster than evolving every ℓ mode all the way to today.

E5. Why does tight coupling imply $\theta_b \simeq \theta_\gamma$?

Before recombination, Thomson scattering is extremely rapid, and the collision term contains

$$\dot\kappa(\theta_\gamma - \theta_b).$$

If $\dot\kappa$ is very large, then any mismatch between θ_b and θ_γ is immediately damped away, so

$$\theta_b - \theta_\gamma = \mathcal{O}(\dot\kappa^{-1}),$$

and at leading order we obtain

$$\boxed{\theta_b \simeq \theta_\gamma.}$$

Last 5 minutes: exam answer template and checklist

F1. A template for long derivation questions

When students see a derivation problem, encourage them to write in the following order:

1. **Write the definition first.** For example, start with the definition of D_L, D_A, χ^2, or $P^s(k,\mu)$.
2. **State the approximations next.** For example: sub-horizon limit, linear theory, no anisotropic stress, independent Gaussian errors, or fixed covariance.
3. **Keep the dependence visible.** Do not plug in numbers too early; first show the dependence on $H(z)$, $D_M(z)$, f, b, or $\bar nP$.
4. **End with one sentence of physical interpretation.** For example: "this is a purely geometric effect", "this quantity is growth-sensitive", "cosmic variance dominates here", or "this is the result of tight coupling".

F2. Ten self-check questions before the exam

1. Can you write down $D_L = (1+z)^2 D_A$ in 30 seconds?
2. Can you explain why $\Phi = \Psi$ requires vanishing anisotropic stress?
3. Can you derive the growth equation from δ' and θ'?
4. Can you explain why supernovae do not determine M and H_0 separately?
5. Can you state clearly what BAO, AP, and RSD each measure?
6. Can you write down the Kaiser formulas for P_0, P_2, and P_4?
7. Can you write the matrix χ^2 for correlated data?

8. Can you explain why two-dimensional 68% corresponds to $\Delta\chi^2 = 2.30$?

9. Can you state the two limiting cases of V_{eff}?

10. Can you explain in one sentence why the CAMB line-of-sight method is fast?

F3. A minimal blackboard outline for the instructor

If you only want to write the bare minimum on the board, a good order is:

1. $1 + z = 1/a, \; \tau = \int \mathrm{d}a/(a^2 H), \; D_L = (1+z)^2 D_A$
2. $\ddot{\delta} + 2H\dot{\delta} - 4\pi G\rho\,\delta = 0, \; \delta_{\text{EdS}} \propto a$
3. $\mu = 5\log_{10}(D_L/\text{Mpc}) + 25, \; \mathcal{M} = M - 5\log_{10} H_0 + \text{const}$
4. $\Delta\theta = r_s/D_M, \; \Delta z = H r_s/c, \; P_g^s = (b + f\mu^2)^2 P_m$
5. $\chi^2 = \boldsymbol{\Delta}^T \mathbf{C}^{-1} \boldsymbol{\Delta}, \; \Delta\chi^2_{68\%,2D} = 2.30, \; \mathbf{F}_{\text{tot}} = \mathbf{F}_1 + \mathbf{F}_2$
6. $V_{\text{eff}} = \int [\bar{n}P/(1+\bar{n}P)]^2 \mathrm{d}^3 r, \; P^{\text{obs}} = \mathbf{W} P^{\text{th}}, \; r \simeq 8/N_\star, \; \Theta_\ell = \int S_T j_\ell$

A final sentence for the students:
The genuinely hard part of the second half of this course is not the number of formulas. It is that a single observable usually mixes **geometry, growth, statistics, window effects, and modeling conventions** at the same time.
As long as you can separate, in each problem, **which quantity is a definition, which step is an approximation, and which combination is actually constrained by the data**, the final exam will look much less messy.