Precision Cosmology With Large Scale Structure
David Weinberg, Ohio State University
ICTP Cosmology Summer School 2015
Lecture 3: Observational Prospects

I have cut this lecture back to be mostly about BAO because I still have holdover topics from previous lectures to cover.

Where are we now?

CMB

Planck: All sky (40,000 deg^2), near cosmic-variance limited for temperature anisotropy, noise limited for polarization.

ACT and SPT: Higher resolution experiments over smaller areas (hundreds or thousands of deg^2), frequency coverage for Sunyaev-Zeldovich as well as primary anisotropy, polarization.

Smaller area high-sensitivity experiments for B-mode polarization (e.g., BICEP).

Type Ia Supernovae

Data sets are compilations of surveys of different redshift ranges, local to $z \approx 1.2$.

"Compilation" requires enormous effort to bring different data sets to common calibration and analysis.

Union 2.1, broad compilation.

JLA – local samples, SDSS-II from $z = 0.1 - 0.4$, SNLS (Canada-France-Hawaii Telescope) from $z = 0.4 - 1$, HST at $z > 1$.

Approaching 1000 SNe (740 for JLA).

Optical imaging/weak lensing

SDSS – 10^4 deg^2, 5 bands (*ugriz*), depth $r \approx 22$, weak lensing source density approximately $1\,\mathrm{arcmin}^{-2}$; SDSS "Stripe 82" went 2 magnitudes deeper over 200 deg^2

Pan-STARRS – 3×10^4 deg^2, depth comparable to SDSS

CFHTLens – 150 deg^2,

Dark Energy Survey (DES) is ongoing, discussed at the end.

Redshift surveys

SDSS-I/II: 10^6 galaxies, broadly selected, median $z = 0.1$; 10^5 luminous red galaxies (LRGs) sparsely sampling structure to $z = 0.45$

SDSS-III BOSS: 1.5×10^6 luminous galaxies, $z = 0.2 - 0.7$, $7\times$ effective volume of SDSS-I/II LRGs; spectra of 160,000 quasars at $z = 2 - 4$ for Lyα forest analysis

Clusters

Variety of cluster catalogs selected by different methods.

Largest are optical cluster catalogs from SDSS, with similar methods now being applied to DES.

Variety of X-ray selected catalogs with tradeoffs between area and depth.

SZ cluster catalogs from Planck (all sky, shallow), SPT and ACT (smaller area, deeper).

BAO surveys and analysis

Reading: *OPCA* ch. 4; Aubourg et al. 2015, arXiv:1411.1074

BAO surveys are interesting in their own right, and they illustrate points relevant to other observational probes.

Basic BAO methodology

Pressure waves that propagate in the pre-recombination universe imprint a characteristic scale on matter clustering, which at late times gets imprinted on the clustering of galaxies or other matter tracers.

This signal appears as damped oscillations in the power spectrum or a single peak in the correlation function.

The scale is determined by the sound horizon at recombination,

$$r_s = \int_0^{t_*} \frac{c_s(t)}{a(t)} dt, \quad c_s(t) = \frac{c}{\sqrt{3}} \left[1 + \frac{3\rho_b(z)}{4\rho_\gamma(z)}\right]^{-1/2}$$

With Planck CMB determinations of $\Omega_b h^2$ and $\Omega_c h^2$, comoving size is $r_s = 147.49$ Mpc, uncertainty of 0.4%.

By measuring this scale in transverse and line-of-sight directions in structure at redshift z, infer $D_A(z)$ and $H(z)$.

In practice, fit a calculated model to observed $P(k,\mu)$ or $\xi(s,\mu)$, but r_s captures scaling with cosmological parameters.

Compared to supernovae, BAO have several distinctive features as a probe of expansion history:

• Because r_s can be computed from first principles, the BAO measurement is in absolute units, while SNe measure distance ratios. (If we knew H_0 perfectly, then we could calibrate SNe in absolute units from low redshift systems in Hubble flow.)

• When the precision is high enough to measure transverse and line-of-sight scales separately, BAO measure $H(z)$ as well as $D_A(z)$. Otherwise, measure angle average, yielding $D_V(z) = [D_A(z)]^{2/3}[cz/H(z)]^{1/3}$.

• The achievable precision of BAO increases with redshift because there is more volume at high redshift.

• The CMB measures the angular scale of the *same* standard ruler at $z = 1100$, with almost perfect precision.

• Even the most powerful BAO surveys seem likely to be limited by statistical errors rather than systematic errors.

Sampling and volume: sample variance vs. shot noise

Roughly speaking, precision of BAO distance measurement is set by ability to centroid the bump in the correlation function.

Width of this bump in linear theory is ~ 10 Mpc at 150 Mpc, or $\sim 7\%$, and it can be centroided to a fraction of that.

For a survey that fully mapped 3-d structure over the whole sky, error in correlation function or power spectrum is set by cosmic variance, limit of finite number of modes in observable universe.

The cosmic variance limit on BAO precision in the redshift shell $z = 1 - 1.3$ is about 0.2% in $D_A(z)$ and 0.1% in $H(z)$; at $z = 0.5 - 0.7$ the corresponding numbers are 0.35% and 0.65%.

If the survey instead covers a fraction $f_{\rm sky}$, the error is larger by $f_{\rm sky}^{-1/2}$. The error caused by finite survey volume is usually referred to as sample variance.

If the galaxy distribution is sparse, then there is an additional error contribution from shot noise — poor measurement of the structure that is there, rather than limited amount of structure to measure.

If one estimates the power spectrum $P(k)$ from a redshift survey of volume V that contains $N_k = 4\pi k^2 dk(V/16\pi^3)$ Fourier modes, the fractional error is

$$\frac{\sigma_P}{P} = N_k^{-1/2} \left(1 + \frac{1}{nP}\right) \propto V^{-1/2} \left(1 + \frac{1}{nP}\right).$$

Here n is mean number density of the tracer galaxies and P is the power spectrum *of the galaxies*.
For $nP \gg 1$, the error is dominated by the finite volume.
For $nP < 1$, the precision is degraded by shot noise, and this degradation becomes large for $nP \ll 1$. Really one would like to have $nP \geq 2$.
The relevant scale for BAO is roughly $k = 0.2h\,\mathrm{Mpc}^{-1}$. The number density required to get $nP = 1$ is about

$$n \approx 4 \times 10^{-4} h^3\,\mathrm{Mpc}^{-3} \sigma_{8,g}^{-2}.$$

More strongly biased galaxies have higher $\sigma_{8,g}^2$, so a lower number density is required.
Thus, surveys of highly luminous galaxies are an efficient way to measure BAO, as they are the brightest objects at a given redshift and they are strongly biased.
The SDSS-III BOSS survey has a comoving space density $n \approx 3 \times 10^{-4} h^3\,\mathrm{Mpc}^{-3}$ and a galaxy bias $b_g = 2$, so its BAO errors are limited mainly by sample variance.
There is more volume at higher redshift, so the sample variance errors for BAO are smaller. However, you have to measure more galaxy redshifts to map this larger volume, and the galaxies are fainter, so it's harder to get high n.
If you want to measure structure on smaller scales (higher k), you generally need higher n to avoid shot noise because $P(k)$ is a decreasing function of k.

Reconstruction

Although the 150 Mpc scale of the BAO is well into the linear regime, the $\sim$ 10 Mpc scale of the *width* of the BAO peak is mildly non-linear.
Non-linear evolution broadens the BAO peak, which one can think of as mass "diffusing" out of the BAO shell.
A broader peak is harder to centroid, so the precision of the measurement is degraded.
(One can also think in Fourier space: non-linear evolution suppresses the oscillations at higher k, so one has weaker features to measure.)
Reconstruction uses the Zeldovich approximation (or some related method) to "move structure back" to its original position, sharpening the BAO peak.
Even fairly simple reconstruction schemes seem to work quite well, significantly improving the precision of BAO measurements (by factor $1.5 - 2$).

Fitting data

Once you have measured the correlation function or power spectrum, you need to infer the BAO scale by fitting a model to the data.
One thing this requires is having a covariance matrix of your measurement errors.
This is typically done by creating large numbers (hundreds, preferably > 1000) of mock catalogs that have realistic clustering properties and measuring the correlation function for these large number of mock catalogs to get

$$C_{ij} = \langle (\xi_i - \bar{\xi}_i)(\xi_j - \bar{\xi}_j \rangle.$$

Creating and analyzing the mock catalogs can be a major computational expense.
You then need to fit a model to the data. One choice is a complete linear or non-linear prediction of the correlation function.
However, an important virtue of BAO is that it uses a localized feature whose location should be insensitive to complications like scale-dependent galaxy bias or smooth errors in the galaxy selection function.

One therefore fits a model that adds nuisance parameters to describe broad-band distortions of $\xi(s, \mu)$ and allows the amplitude of the peak to be free. You fit for the location of the peak marginalizing over these nuisance parameters.

Nuisance parameters degrade your statistical errors but immunize you against answers biased by systematic effects. There is some science and some art to deciding what kinds of nuisance parameters to add so that you remove biases but don't unnecesarily degrade your signal.

BAO systematics

Potential systematics in BAO are:

(a) astrophysical: galaxy bias changes the BAO peak location or messes up your model fit in some other way

(b) observational: problems with your observations (e.g., varying galaxy selection with varying seeing) cause you to measure the wrong correlation function or power spectrum

(c) cosmological: you don't calculate the right r_s

Regarding (a), tests on matter and halos in N-body simulations indicate that shifts in the BAO peak are a fraction of a percent, so correcting for these shifts with moderate accuracy will reduce any biases to 0.1% or less.

Furthermore, reconstruction appears to remove these shifts entirely.

Assessing (b) is very data set dependent. Most systematics are unlikely to introduce a localized feature like the BAO peak. They can affect broad-band power, which is an important reason for the marginalization described above.

Rather exhaustive tests in BOSS show that all the observational systematics we can identify have small impact relative to the errors.

This is true for both the galaxies and the Lyα forest, but we are more confident for the galaxies because we have a lot more experience there.

Regarding (c), an example of physics that could change r_s is the addition of a neutrino species or early dark energy. However, these would be interesting discoveries in themselves. One just has to be aware that there are some standard cosmological assumptions going into the r_s calculation, and violations of those assumptions could change the answer.

BAO with the Lyman-alpha forest

Lyα forest absorption provides a tracer of structure at $z = 2 - 4$. Each quasar sightline probes many independent structures, so this is an efficient way to map structure.

Measuring BAO requires many sightlines and a high enough density to measure many cross-correlations for each sightline.

BOSS did this with 160,000 quasars over 10,000 deg^2.

Most of the usual principles of BAO measurement apply, but with some twists.

The analog of shot noise is the photon noise in the quasar spectra, which cause the flux you measure to differ from the smooth flux.

There is an additional contribution to noise from the fact that your quasar sightlines do not fully sample the structure in the volume you are mapping.

In BOSS, these two error contributions are comparable, and they are both much larger than the cosmic variance error. To reach the cosmic variance limit, we would have to have a much higher quasar density and higher S/N spectra.

A good discussion of these S/N issues is McQuinn & White 2011 (MNRAS 415, 2257).

The biggest observational systematics have to do with determining the quasar continuum, as absorption is measured with respect to this continuum. You don't have to do it right on a quasar-by-quasar basis, but you need to avoid systematics that bias the BAO measurement.

BAO with 21cm intensity mapping

With a giant radio interferometer like SKA, it is possible to do a redshift survey of $\sim 10^9$ galaxies using 21cm redshifts.

But resolving all the galaxies is overkill for BAO, for which one only needs to resolve ∼Mpc scales.

The idea of radio intensity mapping is to use a low resolution radio interferometer to map the overall HI intensity field, which is a kind of smoothed map of the galaxy density field.

This should have a BAO signal just like the galaxy distribution.

The principal challenges are contamination by terrestrial sources and foreground emission. It is not yet clear whether these are easy or hard to overcome.

Radio intensity mapping could prove to be an inexpensive way to measure BAO over large sky areas from $z = 0.5 - 3$.

Some principles of survey design

Think about what you're trying to do. What are your basic observables? What physical parameters do you want to measure?

One can use forecast parameter errors to define figures of merit and use them to optimize instrument or survey design.

However, it is often equally valuable to focus on the precision of measuring the basic observable (e.g., $D_A(z)$ or $f(z)\sigma_8(z)$ instead of w or γ).

Think about what are likely to be the limiting factors on your measurements:

• statistical errors from finite survey volume

• statistical errors from shot noise or something analogous

• systematic errors in the observations (e.g., photo-z errors or shear measurement biases for weak lensing)

• theoretical modeling errors (e.g., redshift-space distortion modeling, baryonic effects on weak lensing signals)

Then think about what you can do to reduce or mitigate those limiting factors, especially if finite survey volume is not the main one.

One can frequently treat systematic biases by parameterizing the effect and marginalizing (e.g., uncertainties in photo-z calibration or shear calibration). The tighter the prior you have on these nuisance parameters, the less statistical penalty you pay.

At an instrument level, the basic trades usually have to do with

• telescope aperture vs. field of view

• pixel size: image sampling vs. field of view

For spectroscopic surveys, there are significantly different considerations depending whether one is using fibers, slits, or slitless spectroscopy.

For a specified instrument design, the basic trade in survey design is usually depth vs. area — long exposures or short exposures?

Details of target selection, redshift sampling, etc. may also matter.

What's on the horizon for dark energy experiments?

Brief Review: Weinberg et al. 2013, arXiv:1309.5380

Dark Energy Survey – 5000 deg^2, 2 mag deeper than SDSS, $\sim 30\times$ CFHTLens. Weak lensing, clusters, SNe, photometric BAO.

Hyper-Suprime Cam – roughly like DES, but different in detail.

eBOSS – extend BOSS to $z = 1 - 2$ BAO measurements with quasars

HETDEX – BAO at $z = 3$ with Lyα emission line galaxies

eROSITA – all sky X-ray cluster catalog, 10^5 clusters

DESI – redshift survey of $\sim$ 30M galaxies for BAO and RSD, roughly $z = 0 - 1.4$, plus denser and larger Lyα forest survey at $z = 2 - 4$

Subaru PFS – roughly like DESI, probably smaller area and higher redshift

LSST – optical imaging of 20,000 deg^2, weak lensing source density of 30 arcmin^{-2} (?)

Euclid – optical imaging of 15,000 deg^2, similar source density to LSST; slitless spectroscopic survey for BAO measurements at $z = 0.7 - 2$, large area but low density

WFIRST – infrared imaging of 2,200 deg^2, somewhat higher source density than LSST or Euclid; slitless spectroscopic survey for BAO measurements at $z = 1 - 2$, densely sampled (typically $nP > 2$). Designed with control of weak lensing systematics as a paramount consideration.

Where might we be in 2030?

Roughly speaking, current measurement precision is $\sim 1 - 2\%$ for expansion history, $\sim 5 - 10\%$ for growth.

Goal of LSST/Euclid/WFIRST is to get to $\sim 0.1 - 0.3\%$ precision for both, with greater redshift range, better control of systematics, more cross-checks across methods.

Lots of discovery space.

If current anomalies are real, they will be well confirmed by DES/HSC.

Alternatively, they could go away, but convincing new anomalies could come out of the greater precision of LSST/Euclid/WFIRST.

These might point us in a clear direction, or we might just have an anomaly without knowing what it means.

Something might come from a surprising direction: time variation of α, fifth-force experiments, etc.

Alternatively, everything could remain consistent with ΛCDM. It's not clear there are feasible ways to gain another factor of 10 in precision, but there are some possibilities, and good ideas could come along between now and then.

We should have a cosmological measurement of neutrino mass.

There will also be big advances in CMB experiments, perhaps including detections of B-mode polarization from gravity waves or evidence for primordial non-Gaussianity.

These will better constrain inflation models or alternatives, maybe point to some surprising understanding of cosmic acceleration.