Lecture 2 Notes (VSOP 2026): Cosmological Constraints on Dark Matter

Tracy R. Slatyer[1,*]

[1] *Center for Theoretical Physics – a Leinweber Institute, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*

I. HORIZONS AND PERTURBATIONS

This FLRW metric with dark matter/energy components seems to provide a good description of the cosmos on average, but in reality of course the cosmos is not homogeneous/isotropic – there are fluctuations in the matter and energy densities (some large, e.g. Earth!) and accompanying fluctuations in the metric which describes gravity. We can do perturbation theory in a cosmological context by considering small fluctuations around the FLRW metric. What we have done so far is zeroth-order cosmology, and this is sometimes all we need, but first-order (linear) cosmology is largely how we extracted the cosmological parameter values I cited earlier. When the perturbations become large enough, we need to resort to numerical simulations (although clever effective-field-theory methods can extend the range of perturbation theory to smaller scales). We are not going to treat cosmological perturbation theory in these lectures due to time constraints, but you should know (1) the perturbations exist and hold important information about dark matter, and (2) they can oscillate or grow over time, but their behavior is qualitatively different depending on their wavelength relative to the *horizon scale* (and also depends on the environment, e.g. whether the cosmos is matter-dominated, radiation-dominated, or dark-energy dominated).

Cosmological horizons describe the maximum distance scales over which events can be causally connected. There are a few different meanings of the horizon that differ at the $\mathcal{O}(1)$ level, but roughly speaking, this distance is set by how far light can travel in the time it takes the universe to expand by a $\mathcal{O}(1)$ factor. This expansion timescale is governed by $H = (1/a)da/dt = d\ln a/dt$ – note this parameter has units of inverse time. This suggests a characteristic (physical) scale of order cH^{-1} for the horizon distance (recalling that H has units of inverse time). The corresponding comoving scale is the physical distance divided by a/a_0, or $((a/a_0)H)^{-1}$ if we work in natural units where $c = 1$ (which we will generally do in these notes). This specific quantity is referred to as the *comoving Hubble radius*.

Two precise horizon definitions that encode causal connectedness are (1) the *particle horizon* (also sometimes called the *cosmic horizon*) and (2) the *cosmic event horizon*. The particle horizon/cosmic horizon is the maximum distance between two particles such that light emitted arbitrarily early from particle 1 could have reached particle 2 by time t, and as such:

$$d_{\mathrm{p}}(t) = c\,a(t)\int_{t_{\mathrm{i}}}^{t}\frac{dt'}{a(t')}, \tag{1}$$

where t_{i} is the earliest possible time of emission (often taken as the time when $a(t) \to 0$). The cosmic event horizon is the maximum distance such that light emitted at time t can ever reach a particle at arbitrarily late times:

$$d_{\mathrm{e}}(t) = c\,a(t)\int_{t}^{\infty}\frac{dt'}{a(t')}. \tag{2}$$

Thus these quantities describe how causally connected the universe has been in the past, and will remain in the future.

[*] tslatyer@mit.edu

A. Perturbations

It is convenient to decompose density perturbations into Fourier modes:

$$\delta(\boldsymbol{k}, t) = \int d^3x \, e^{i\boldsymbol{k}\cdot\boldsymbol{x}} \, \delta(\boldsymbol{x}, t), \qquad \delta \equiv \frac{\rho - \bar{\rho}}{\bar{\rho}}. \tag{3}$$

Here $\boldsymbol{k}$ is the spatial frequency or wavenumber associated with the mode. The wavelength of the mode is $\lambda = 2\pi/k$. Because these wavelengths stretch with the cosmic expansion, it is often better to characterize perturbations by *comoving* wavelength or, equivalently, by wavenumber (as usual the comoving wavelength is the physical wavelength rescaled by a_0/a).

B. Horizon entry and exit

Physically, we can say generally that if a mode's wavelength exceeds the horizon scale, different parts of the mode cannot communicate causally. Such a mode is said to be *frozen* or *outside the horizon.* The relevant scale for the (comoving) wavelength turns out to be the (comoving) Hubble radius. If the wavelength later becomes smaller than the Hubble radius, we say the mode has *entered the horizon.* Conversely, a mode can *exit the horizon* if its wavelength grows relative to the Hubble radius.

For a mode with fixed comoving wavelength λ, to understand how it has evolved over history, we will thus be interested in seeing how $((a/a_0)H)^{-1}$ evolves in the expanding universe – this will tell us what the pattern of horizon entry/exit for the mode looks like. Since a_0 is constant we can just look at the scaling of $(aH)^{-1}$.

On the problem set you will explore how a evolves with t for each of dark energy domination, matter domination, and radiation domination. For dark energy domination, since ρ remains constant, we see from the Friedmann equation that H will also be constant, and then $a \propto e^{Ht}$. In this case $aH \propto He^{Ht}$ *grows exponentially* over time, so the comoving Hubble radius *falls exponentially* with time. In such a universe, all pre-existing modes eventually exit the horizon.

For matter/radiation domination, you will show that $a \propto t^k$ with $0 < k < 1$. In this case, $H = (1/a)da/dt = k/t$, so $(aH)^{-1} \propto t^{1-k}$ – thus the comoving Hubble radius *grows* with time, and more modes (with fixed comoving wavelength) enter the horizon as time goes on. Loosely speaking, a matter/radiation-dominated universe becomes more causally connected over time, whereas the opposite is true in a dark-energy-dominated universe – because we are entering a period of dark energy domination, there are galaxies that are now observable, but the light emitted from them today will never reach us.

Thus the history of a typical perturbation mode is that it exited the horizon during inflation, and then re-entered the horizon at some later time (shorter-wavelength modes earlier) during the radiation/matter-dominated epochs, at which point it began to oscillate and/or grow. We have only recently re-entered dark energy domination so only the largest-scale modes have re-exited the horizon, but we would expect more to do so over time. In general, only modes that are inside the horizon at a given time are sensitive to new physics occurring at that time. Thus often small-scale modes (which have been inside the horizon for longest) are expected to be most significantly affected by new physics; however, those modes are also usually harder to model theoretically (as perturbation theory breaks down once the perturbations grow beyond the linear regime).

II. A COSMIC TIMELINE

With all that background information, let us now summarize how we think the universe evolved, using temperature as a proxy for redshift/time.

Inflation (T = ?): At very early times, there was likely an epoch of cosmic inflation, a period of early exponential expansion. Inflation massively dilutes any existing matter/radiation and seeds nearly scale-invariant perturbations, from quantum fluctuations in the field driving inflation. These perturbation modes exit the horizon, but when they re-enter the horizon later (during matter/radiation domination) and begin to oscillate and grow, they form the seeds of structure in our cosmos.

$T \sim 100$ **GeV:** the electroweak phase transition/crossover occurs: the Higgs boson acquires a vacuum expectation value and gives masses to Standard Model particles.

$T \sim 100$ **MeV:** the QCD phase transition occurs, from quark–gluon plasma to hadronic degrees of freedom.

These epochs are very interesting, but they are not directly observable today – they could potentially yield gravitational-wave signatures that are observable in the future, and we can learn about inflation from measuring the perturbations.

$T \sim 1$ **MeV:** neutrinos decouple from the photon bath, and the temperature drops below the proton–neutron mass splitting. This is the epoch of Big Bang nucleosynthesis (BBN), it is when various elements are formed, and so late-time elemental abundances can probe its properties.

$T \sim 0.2$ **eV:** photons are no longer hot enough to reliably dissociate hydrogen. Hydrogen forms through $e^- + p \to H$ and the Universe becomes transparent to low-energy photons. This period is called the *recombination epoch* (a bit of a misnomer as the hydrogen was never previously combined!), and the *epoch of last scattering*, as it corresponds to the release of the cosmic microwave background (CMB) photons from the ionized plasma. These photons subsequently free-stream through the cosmos and provide the earliest snapshot of the Universe. Metric fluctuations born during inflation imprint themselves on these photons, allowing us to probe properties of both inflation and the plasma that was coupled to the photons.

$T \sim 2 \times 10^{-3} - 0.2$eV**:** this period is often called the cosmic dark ages. The fluctuations in the matter density grow nonlinearly and seed cosmic structure formation, yielding a filamentary cosmic web of mostly dark matter. Regions of high dark-matter density attract ordinary matter via gravity; stars begin to form in these gas clumps, leading to galaxies. At the end of this period, early stars heat the gas and ionize the hydrogen, marking cosmic dawn and reionization.

III. OUTLINE OF CONSTRAINTS ON DARK MATTER PROPERTIES FROM COSMOLOGY

We have powerful constraints on dark-matter properties from the epoch of BBN onward. In this section we will discuss some of those general bounds, which constrain the space of possibilities for all dark matter models.

A. N_{eff} bounds

BBN observables are the abundances of light nuclei, especially helium and deuterium. The helium case is the simplest: it depends on how many neutrons are bound into helium nuclei before they decay to protons. This is sensitive to the age of the Universe when the temperature drops below the neutron–proton mass splitting, which is controlled by the Hubble expansion rate $H(z)$, with $H^2 \propto \rho$ as we discussed above.

For this reason, BBN (and later the CMB) is sensitive to the total energy density of the universe and hence to the number of species populating the radiation bath. This is usually quoted as a bound on the number of effective relativistic degrees of freedom, in terms of an effective number of neutrino species. Extra radiation would increase the energy density of the radiation field and hence modify the Universe's expansion rate, which needs to be compared to the rate of nuclear processes to calculate the nucleus yields from BBN.

At late times one usually writes

$$\rho_{\text{rad}} = \rho_\gamma \left[1 + \frac{7}{8}\left(\frac{4}{11}\right)^{4/3} N_{\text{eff}}\right], \tag{4}$$

where N_{eff} is the number of effective neutrino species, which we typically break into a SM contribution (close to 3) and an exotic contribution ($N_{\text{eff}} = 3 + \Delta N_\nu$). The reason for the extra prefactors is that the Standard Model neutrinos have a lower temperature than the photons, by a factor of $\approx \left(\frac{11}{4}\right)^{1/3}$, and the energy density of a radiation species in equilibrium scales as the 4th power of its temperature (note this is consistent with ρ redshifting as $(1+z)^4$ and temperature redshifting as $T \propto (1+z)$). Qualitatively, this is because after the neutrinos' interactions with the photon bath become inefficient, electrons and positrons annihilate and heat up the photons, but not the neutrinos). I will not go through this calculation in detail but it follows from the conservation of entropy in the expanding universe.

A careful calculation obtains $N_{\rm eff}^{\rm SM} = 3.045$ for the SM contribution, and the current observational limits are $N_{\rm eff}^{\rm Planck\ 2018} = 2.99 \pm 0.17$ from the CMB [1], and $N_{\rm eff}^{\rm BBN\ only} \simeq (2.93\text{–}2.96) \pm 0.3$ from BBN [2], with the exact central value depending on the code used to calculate the nuclear yields. While the BBN limits have a modestly larger error bar, they constrain the expansion history at a much earlier time / higher temperature, which can be useful.

In particular, BBN limits directly constrain the possibility that dark matter was still relativistic (and hence radiation-like) at $T \sim 1$ MeV. If dark matter, or a new particle to which it couples, is lighter than a few MeV and has a comparable temperature to Standard Model particles, it will in general violate these bounds; lighter DM can evade these limits if its temperature is significantly colder than the Standard Model (implying upper limits on its interaction rate), since the contribution to the energy density scales as T^4. Even if the DM is cold enough to evade the bounds, if it annihilates or decays into Standard Model particles, that can differentially heat the photons and neutrinos, which again modifies $N_{\rm eff}$ (e.g. [3]).

B. The CMB

Like BBM, the CMB sets limits on the energy content in relativistic species, but beyond that, it gives our best measurement of the total dark-matter density (to percent-level accuracy), and tightly constrains deviations of the dark-matter equation of state around recombination. For example, no more than a few percent of dark matter can decay between the CMB epoch and today [4]. The CMB also sets stringent limits on couplings between dark matter and ordinary matter or radiation — dark matter must remain "dark" in the sense that it is only very weakly coupled to the Standard Model. We will return to these bounds when we discuss specific models.

In terms of perturbations, the CMB constrains only relatively large-scale modes (long wavelengths). Smaller scales and later times are accessed through large-scale structure observations.

C. Structure formation

After photons decouple from the ionized plasma, matter perturbations continue to grow under gravity and eventually collapse into virialized structures. Most dark matter must be cold: truly hot, relativistic dark matter cannot exceed roughly $\sim 1\%$ of the total [5]. Dark matter that is too fast-moving erases structure on small scales. A similar suppression of structure on small scales occurs if dark matter couples to radiation, with a scale set by the free-streaming length of the radiation, or if the dark sector introduces a new physical scale, such as the de Broglie wavelength of the dark matter particle. Only modes that are inside the horizon — and therefore evolving rather than frozen — are affected by these mechanisms. Redshift-dependent effects therefore induce a scale dependence, with differing effects between modes that were or were not inside the horizon at that redshift.

It is often convenient to work with the power spectrum

$$P(k) = \left\langle |\delta(\boldsymbol{k})|^2 \right\rangle , \tag{5}$$

which describes the power in density fluctuations at a given scale. The effects described above can be viewed as modifications to the power spectrum on small scales.

1. *Limits on dark matter free-streaming*

Our current observational bounds on the free-streaming scale $\lambda_{\rm fs}$ (which can be translated into bounds on the dark matter velocity or interaction rate) include:

- Milky Way satellites: $\lambda_{\rm fs} \lesssim 10\,h^{-1}\,{\rm kpc}$ [6], where $h \equiv \frac{H_0}{100\ {\rm km\,s^{-1}\,Mpc^{-1}}} \approx 0.7$.
- Strong gravitational lensing: $\lambda_{\rm fs} \lesssim 6$–$7\,{\rm kpc}$ [7].
- Lyman-α forest, i.e. clustering of gas at $z \sim 2$–6: $\lambda_{\rm fs} \lesssim 10\,{\rm kpc}$ [8].
- Looking for disruptions to stellar streams in the Milky Way gives comparable sensitivity [9].

A kiloparsec (kpc) is an astrophysical unit, 1 parsec is about 3.3 light years. For calibration, the distance from the Earth to the center of the Milky Way is about 8.5 kpc.

These searches probe the tiniest dark-matter halos that can form stars, and are now pushing into the regime of "truly dark" halos with no bound stars. There is a nice review article on how different dark matter models can modify the power spectrum by Bechtol et al [10].

2. *Limits on ultra-light ("fuzzy") dark matter*

The same observations can be used to constrain dark matter whose de Broglie wavelength ($\lambda = 2\pi/p$) is comparable to these free-streaming scales. For calibration, for dark matter with a velocity of $10^{-3} - 10^{-4}c$, which is typical of present-day galaxies, a mass of 10^{-21} eV corresponds to a de Broglie wavelength of $\sim 2-20$ kpc. Thus in this mass range, quantum/wave effects should be visible on the scale of (dwarf) galaxies. On one hand this is an extremely tiny mass (we expect at least one neutrino to have a mass around 0.1 eV), but on the other hand it is the most robust lower limit we have on the dark matter mass scale. Analyses like those described above for free-streaming set limits around $m_{\rm DM} \gtrsim 2 - 3 \times 10^{-21}$ eV [6, 9].

However, in this specific case (ultralight dark matter) modern analyses claim to do quite a bit better, attaining limits of $m_{\rm DM} \gtrsim 10^{-19} - 10^{-18}$ eV. For example, [11] claims bounds of a few $\times 10^{-19}$ eV, based on detailed modeling of wave interference effects in dwarf galaxies and how they would affect the stars in those galaxies (they actually claim a stronger limit, around 8×10^{-18} eV, but it relies on a particular system being an accurately-measured galaxy, which is a much more tenuous claim).

3. *The Tremaine-Gunn bound*

For fermionic dark matter, we can do quite a bit better than the mass bounds above, because of constraints on how tightly the dark matter can be packed. This is the Tremaine-Gunn bound. The Pauli exclusion principle mandates that for identical fermions, the phase space density satisfies:

$$\frac{n}{p^3} \sim \frac{\rho}{m_\chi (m_\chi v)^3} \lesssim 2 \tag{6}$$

which implies

$$m_\chi \gtrsim \left(\frac{\rho}{v^3}\right)^{1/4}. \tag{7}$$

The smallest galaxies we observe have (very roughly) $v \sim 10^{-5}$, and dark matter mass densities around 1 GeV/cm^3. In such systems, this constraint gives:

$$m_\chi \sim \left(\frac{1\ {\rm GeV/cm^3}}{(10^{-5})^3}\right)^{1/4} \sim 10^{-6.5}\ {\rm GeV} \sim 300\ {\rm eV}. \tag{8}$$

Thus, particle dark matter below this approximate scale must be bosonic. A more careful calculation [12] obtained a lower mass bound of 410 eV (for fermionic DM) based on three separate dwarf spheroidal galaxies.

In the next lectures, we will explore classes of models that survive these constraints and how they can be tested experimentally.

[1] N. Aghanim *et al.* (Planck), "Planck 2018 results. VI. Cosmological parameters," Astron. Astrophys. **641**, A6 (2020), [Erratum: Astron.Astrophys. 652, C4 (2021)], arXiv:1807.06209 [astro-ph.CO].

[2] Cara Giovanetti, Mariangela Lisanti, Hongwan Liu, Siddharth Mishra-Sharma, and Joshua T. Ruderman, "Cosmological parameter estimation with a joint-likelihood analysis of the cosmic microwave background and big bang nucleosynthesis," Phys. Rev. D **112**, 063530 (2025), arXiv:2408.14531 [astro-ph.CO].

[3] Nashwan Sabti, James Alvey, Miguel Escudero, Malcolm Fairbairn, and Diego Blas, "Refined Bounds on MeV-scale Thermal Dark Sectors from BBN and the CMB," JCAP **01**, 004 (2020), arXiv:1910.01649 [hep-ph].

[4] Théo Simon, Guillermo Franco Abellán, Peizhi Du, Vivian Poulin, and Yuhsin Tsai, "Constraining decaying dark matter with BOSS data and the effective field theory of large-scale structures," Phys. Rev. D **106**, 023516 (2022), arXiv:2203.07440 [astro-ph.CO].

[5] Maria Archidiacono, Steen Hannestad, Alessandro Mirizzi, Georg Raffelt, and Yvonne Y. Y. Wong, "Axion hot dark matter bounds after Planck," JCAP **10**, 020 (2013), arXiv:1307.0615 [astro-ph.CO].

[6] E. O. Nadler *et al.* (DES), "Milky Way Satellite Census. III. Constraints on Dark Matter Properties from Observations of Milky Way Satellite Galaxies," Phys. Rev. Lett. **126**, 091101 (2021), arXiv:2008.00022 [astro-ph.CO].

[7] D. Gilman *et al.*, "The free-streaming length of dark matter from JWST observations of 28 strong gravitational lenses," (2026), arXiv:2606.05277 [astro-ph.CO].

[8] Vid Iršič *et al.*, "Unveiling dark matter free streaming at the smallest scales with the high redshift Lyman-alpha forest," Phys. Rev. D **109**, 043511 (2024), arXiv:2309.04533 [astro-ph.CO].

[9] Nilanjan Banik, Jo Bovy, Gianfranco Bertone, Denis Erkal, and T. J. L. de Boer, "Novel constraints on the particle nature of dark matter from stellar streams," JCAP **10**, 043 (2021), arXiv:1911.02663 [astro-ph.GA].

[10] Keith Bechtol *et al.*, "Snowmass2021 Cosmic Frontier White Paper: Dark Matter Physics from Halo Measurements," in *Snowmass 2021* (2022) arXiv:2203.07354 [hep-ph].

[11] Simon May, Neal Dalal, and Andrey Kravtsov, "Updated bounds on ultra-light dark matter from the tiniest galaxies," (2025), arXiv:2509.02781 [astro-ph.CO].

[12] Alexey Boyarsky, Oleg Ruchayskiy, and Dmytro Iakubovskyi, "A Lower bound on the mass of Dark Matter particles," JCAP **03**, 005 (2009), arXiv:0808.3902 [hep-ph].