Tracy Slatyer – Lecture 2 Questions

Questions marked in green were answered during the Q&A session. No attempt was made to correct grammar/spelling issues. Where a slide number was given it is shown.

Q1: In the DM fraction of PBH plot, the mass range where PBH can make up the entire DM density (M_{PBH} ~ 10^{17} to 10^{23} g), depends strongly on the underlying mass distribution of PBHs, right? So, it is possible that for extended enough PBH mass distribution (example: for a log-normal mass distribution of very high width, for a power law mass distribution etc), there is actually no region in parameter space where PBHs can be entirety of DM. Besides, some of the evaporation constraints (example: INTEGRAL 511 keV constraints) also depend very strongly on DM density profile and for some density profile, the ""mass gap"" is actually not that large . So, the mass window where PBHs can make up the entire DM density is actually not that robust as portrayed, right? (since it depends very strongly on underlying mass distribution.)

A: It's correct that the plot I showed is assuming a delta-function PBH mass distribution, and when you allow for extended mass distributions then the constraints change. My understanding, as you say, is that typically it makes the constraints harder to evade, as a small fraction of the PBHs being in a mass range where there are sensitive constraints can exclude the whole scenario. The Carr review 2002.12778 has a section on extended mass functions (IVB).

The constraints my collaborators and I set using the INTEGRAL data (https://arxiv.org/abs/2004.00627), which to my knowledge are currently setting the lower end of the open mass range for PHBs forming 100% of DM, are almost independent of the DM density profile. The 511 keV bound depends more sensitively on the density distribution, as you say, but in terms of what PBH mass values can make up 100% of the DM, it looks like it's a difference of less than a factor of two in the minimum allowed PBH mass (looking at Fig. 1 of https://arxiv.org/pdf/1906.09994.pdf). The open(ish) window at low masses covers 5-6 orders of magnitude, so I don't think that varying the DM density profile allows you to close that window significantly.

Q2 (slide 4): Usually cosmic rays should give rise to e+ e- in equal rates right? if they come from photons pair producing. Then why e+ are sensitive probes for dark matter?

A: Most cosmic rays in our galaxy are thought to be ambient electrons/protons that are swept up in supernova shocks and accelerated to high energies -- they don't originate from pair production. Positrons are produced when protons collide with the gas, but this rate is small enough to give a subdominant contribution (which also is expected to fall with increasing energy) compared to the electron cosmic rays. Pair production in pulsar magnetospheres, and dark matter annihilation, are both examples of mechanisms that are expected to produce electrons and positrons in equal number, which is why both could in principle help explain the rise in high-energy positrons seen by AMS-02 (except that the rate of dark matter annihilation required is in conflict with other constraints).

Q3 (slide 5): If dark matter causes excess in antielectrons/anti-protons, then shouldn't it also increase the number of observed electrons/protons too, as annihilation are to pair of particle - antiparticle

A: Indeed it does! But the ambient ratios of positrons to electrons and antiprotons to protons are both much less than 1 (as the universe is matter/antimatter-asymmetric), so adding a component with a 50-50 ratio has the effect of increasing the fraction of antimatter.

Q4 (slide 5): Could we get useful limits from measuring cosmic ray protons and electrons, instead of antiprotons and positrons?

A: Such searches have a much higher background, for the reasons described above. That said, yes, you can also get constraints from these channels, it’s just hard for them to be competitive given the large background and uncertainties in that background.

Q5 (slide 5): For the heavier DM, the expected <sigma v> is greater than the thermal relic cross-section, does this imply that the backgrounds are much higher for those regions and if one has to observe an indirect signal of DM annihilation, one has to deal with backgrounds to see some signal, or I am missing some point over here.

A: The constraints being above the thermal relic line means we don’t have the ability to exclude thermal relic cross sections at those high masses. This is not because the backgrounds are higher at high energies (the opposite is usually true), but because the signal gets smaller at high energies. To understand why, remember there is a 1/m_DM^2 in the equation (from yesterday) for the annihilation flux observed at telescopes - we constrain the DM mass density from astrophysical observations, not its number density, and so heavier DM has a lower number density. The annihilation rate scales as the square of the number density, and hence as 1/m_DM^2.

Q6 (slide 6): Did they actively design the spectrometer on Voyager to measure cosmic rays? If not, how did they adapt it to do so?

A: Yes, it was designed to measure cosmic rays, my understanding is they wanted to see how the flux changed as one moved through and out of the solar system. It was a very good decision! There’s some more information here: https://nssdc.gsfc.nasa.gov/nmc/experiment/display.action?id=1977-084A-08

Q7 (slide 4): Minimal Model and data are in very consistent for the highest energy points. Is there any concern about possibly over-fitting?

A: The minimal model for the positrons is just a smooth fitting function, to be clear, not a physical model - it's a power law + a power law with an exponential cutoff. The number of free parameters here is pretty tiny compared to the number of data points, so I don't think the issue is overfitting in that sense. However, I do think the fact that the model goes through ~all the data points (the chi^2 per degree of freedom is smaller than 1) reflects that there are almost certainly bin-to-bin correlations in the AMS-02 data, and so treating the data points and error bars as entirely uncorrelated leads to fits that look better than they should. This is especially important when trying to assess the significance of possible signals in the antiproton channel (the positron excess is so huge that it isn't as much of a concern there, but this matters a lot for signals that are on the borderline of being statistically significant, as in the case of antiprotons). For example, see these papers: https://arxiv.org/abs/1903.01472 , https://arxiv.org/abs/1906.07119 (note they get very different answers for the final significance!)

Q8 (slide 7): When did people start realising that Voyager could be used for dark matter bounds?

A: Good question! The first I learned of it was Boudaud's paper in 2016 (https://arxiv.org/abs/1612.07698). This popular-press article (https://www.forbes.com/sites/brucedorminey/2018/07/10/nasas-voyager-1-spacecraft-opens-door-on-new-way-to-look-for-dark-matter) quotes Alan Cummings, who was on the Voyager science team starting in 1973, as saying (paraphrasing) he had never dreamed that Voyager would set constraints on dark matter. Of course, when it launched, we were only just starting to learn about dark matter from galactic rotation curves. I don't know if people speculated about Voyager's utility for dark matter searches in the years between its launch and 2016 - I'll leave that literature search as an exercise for the reader.

Q9 (slide 7): Why do the cross section limits for DM decay and DM annihilation into e+e- display spiky features between 0.01 GeV and 0.05 GeV and around 10 GeV?

A: The low-energy ones are based on Voyager data, and checking the data plot (e.g. Fig. 2 in https://arxiv.org/pdf/1612.07698.pdf), it appears there are just four Voyager data points. Dark matter annihilating to electrons produces quite a peaked electron spectrum (it's a delta-function in energy, E=m_DM, before propagation), and so I can imagine that in this case the signal is essentially just a spike in a single bin. As the DM mass changes, the relevant bin will change, and when the number of bins is small this can induce features in the constraint. I would guess this is what's going on, at least for the low-energy data - the Voyager data points are indeed at 0.01-0.05 GeV, so this is exactly the mass range where I'd expect the signal to correspond to a sharp peak falling almost entirely within a single bin. (For higher DM masses, you will never see the spike with Voyager data, only its low-energy tail, which spans all the bins - and for lower DM masses, the signal is entirely outside the Voyager data, so there's no constraint.) Something similar may be happening in the AMS-02 10 GeV data, but there are many more data points in that case so it's not so easy to diagnose.

Q10 (slide 11): Where does the value of 1 GeV come from?

A: The mass of a hydrogen atom is about 1 GeV. Thus there's about 5 GeV of energy stored in dark matter mass for each hydrogen atom (since dark matter has about five times the total mass of visible matter, which is mostly hydrogen).

Q11 (slide 11): If dark matter annihilates into hidden sectors, then there will be almost no limits from the ionisation of hydrogen atom?

A: If the hidden-sector particles decay back to the Standard Model (except neutrinos) on timescales shorter than a few hundred thousand years, then the CMB constraints will be very similar to the case of direct annihilation. If they decay with a lifetime longer than that, then the constraints will be modified more significantly. If they don't decay at all and are just stable dark radiation, then indeed these ionization constraints won't apply - instead you would look at CMB constraints on the number of light effective degrees of freedom, plus structure formation and the matter power spectrum.

Q12 (slide 11): Does this late time annihilations, also affect the formation of heavier nuclei like Helium etc

A: Good question! Energy injection during the cosmic dark ages can ionize helium *atoms*, and we take that into account in our code. But if you're asking about the formation of helium *nuclei*, then that's a question about Big Bang nucleosynthesis - and indeed, dark matter annihilation can affect nucleosynthesis too, e.g. by producing energetic photons that dissociate the newly-formed nucleons. Here's a nice review, although it's getting a bit old now: https://arxiv.org/abs/0906.2087 . The constraints are typically a bit weaker than those from the CMB / ionization, but they're quite independent.

Q13 (slide 12): what happens at even higher DM masses, i.e. $m \sim 10^6$ GeV - 10^{10} GeV?

A: That's a great question. I don't think it's been calculated, but I have a physical argument that the slope of the line shouldn't change much. The slope you see at high masses is just $\langle \sigma v \rangle \propto m_{DM}$. The reason that's true is that the total energy injected by annihilation scales as $m_{DM} \times$ rate = $m_{DM} (n^2 \langle \sigma v \rangle/2) = \frac{1}{2} \langle \sigma v \rangle \rho_{DM}^2/m_{DM}$, so if m_{DM} increases by a factor of 10 (say) then you need to increase $\langle \sigma v \rangle$ by a factor of 10 as well to get the same energy injection. Now, there is an energy-dependent (and particle-species-dependent) efficiency factor which controls how efficiently this total injected energy gets converted into ionization. This efficiency factor basically has two parts, (1) how much power goes into electromagnetically-interacting particles like photons and electrons, as opposed to neutrinos (for high enough DM masses, this is typically roughly constant as a function of DM mass, for a fixed SM final state), (2) what fraction of the energy carried by electromagnetically-interacting particles goes into ionization.

At those very high masses, electromagnetically-interacting particles injected by annihilation/decay have enough energy that they immediately enter a pair production cascade by scattering off CMB photons (injected photons will scatter against a CMB photon and pair-produce; injected electrons and positrons will inverse Compton scatter to transfer their energy into high-energy photons, which then pair-produce). This happens very quickly, and the effect of the cascade is essentially to translate all the energy of the injected particles into photons just below the threshold to efficiently pair-produce on the CMB. At redshift around 1000 when the CMB is being formed, that threshold is about 100 GeV, and then it goes up linearly with decreasing redshift (so it's 1 TeV at z=100). So more or less, I expect electromagnetically-interacting particles injected at energies above 100 GeV - 1 TeV to get promptly converted into 100 GeV - 1 TeV photons, and consequently to ionize the universe with a similar efficiency as 100 GeV - 1 TeV photons. Thus the efficiency shouldn't change very much with DM mass once the mass is above about 1 TeV, and this means we just get the simple relation that the constrained annihilation cross section <sigma v> \propto m_DM, and that this relationship should continue to high DM masses.

Q14 (slide 13): In presence of sommerfeld enhancements, will these bounds become even stringent?

A: Yes, Sommerfeld enhancement boosts the annihilation rate at low velocities, and so CMB constraints are very powerful for constraining Sommerfeld-enhanced models.

Q15 (slide 13): What are the limits on very light dark matter candidates such as Axion?

Ultralight axions decay to SM particles with a very long lifetime, because they are so light and weakly-coupled to the SM - they would also only decay (or annihilate) to very low-energy photons or neutrinos. This means they won't ionize the gas, and are generally very difficult to detect through annihilation or decay, unless there's some enhancement to these processes. A more promising channel in most models is to look for their *oscillation* to/from Standard Model photons.

Q16 (slide 14): Why does 1e-9 of the DM getting annihilated correspond to a 1e-6 distortion in the CMB spectrum?

A: The energy density of radiation in the universe today is roughly 1/3000 that of matter - at earlier times, it was larger relative to the matter density. Even today, if we converted 10^(-9) of the energy density stored in matter into radiation, it would only be a O(10^(-6)) distortion to the CMB (3000*10^(-9) = 3*10^(-6)). At any earlier time, the effect would be even smaller. In reality, simple models of DM annihilation that aren't yet excluded by other searches tend to predict a distortion to the CMB of order 10^(-9)-10^(-10).

Q17 (slide 15): Why is 21cm emission so special? Are there another spectra which also can be used? To be specific can we also use transitions like higher n to the ground state?

A: The higher-n transitions to the ground state are the Lyman-series transitions, like the Lyman-alpha line I also discussed. These are great for measuring the temperature at late times, but when the hydrogen in the universe is close to 100% neutral (as is true before reionization), anything energetic enough to excite hydrogen from the ground state gets completely absorbed and there's very little signal (this is called the Gunn-Peterson trough). The 21cm transition is much weaker (smaller rate), and so we can hope to see 21cm emission even from epochs when the universe was full of neutral hydrogen gas. Certainly, though, other emission lines could be used as well - here is a review of some strategies: https://arxiv.org/abs/1709.09066

Q18: Why is there skepticism on the EDGES result?

A: That's a whole talk in itself, but the gist is (1) this is a really hard measurement, this is the first claim of a detection, and it hasn't been confirmed by any other experiments, (2) the EDGES result is extremely surprising under standard cosmology.

They found a very deep 21cm absorption trough at z\sim 17, implying either: there was a lot more 21cm radiation illuminating the gas at that time than we previously thought; the 100-million-year-old universe was very cold, colder than we thought was possible even setting the heating from stars to *zero*; or we just have cosmology very wrong for this epoch (which would be a bit surprising as we have measurements of the CMB and large-scale structure which seem to work perfectly). If the measurement turns out to be correct, it's absolutely the discovery of the decade.

Q19: If dark matter dominantly talks with species in the dark sector, is there any experimental way still to determine its nature apart from the gravitational ones

A: See the answer to Q11 above - if the dark-sector particles talk to the SM, there are many ways to search for them. If it's a fully secluded dark sector, however, then we would mostly be restricted to cosmological and gravitational probes, e.g. measurements of the number of light degrees of freedom through their effects on the expansion history, and the distribution of dark matter throughout the cosmos. Note these probes can also constrain interactions between dark-sector particles, so they're not testing solely gravitational *interactions*, even if the searches themselves rely on the gravitational effects of dark matter.

Q20 (slide 24): Could the 7 keV sterile neutrino also explain the LSND anomaly?

A: Unfortunately, as far as I know, the LSND anomaly prefers an eV-scale sterile neutrino - much lighter than 7 keV.

Q21 (slide 25): This is a bit of an unrelated question, but one that I never had the opportunity to ask. In Cuoco's 2017 paper, that figure's (the second one in your slide) caption says it is showing the best fit to the antiproton/proton ratio to the AMS data, but the units are for a flux. I assume it's a typo and it is the antiproton flux not the ratio, but maybe I am wrong.

A: Good catch, I hadn't noticed that before! I think the typo is in the y-axis label and this is actually the antiproton ratio, based on the values on the y-axis being around 10^(-5)-10^(-4) - compare to the raw AMS-02 data in their 2016 paper:
https://journals.aps.org/prl/abstract/10.1103/PhysRevLett.117.091103