Solutions: Astrophysical Probes of Fundamental Physics

Based on TASI Lectures, Parts I & II

Problem 1: Volume vs. Surface Emission

(a) The BSM particle ϕ is produced throughout the stellar volume and free-streams out, so the total BSM luminosity is simply the production rate integrated over the volume:

$$L_\phi = \epsilon_\phi \cdot \frac{4}{3}\pi R_\odot^3. \tag{1}$$

(b) The BSM luminosity from the red giant core is

$$L_\phi(\mathrm{RG}) = \epsilon_\phi(T_c) \cdot \frac{4}{3}\pi r_c^3, \tag{2}$$

while that from the Sun is

$$L_\phi(\odot) = \epsilon_\phi(T_\odot) \cdot \frac{4}{3}\pi R_\odot^3. \tag{3}$$

Their ratio is

$$\frac{L_\phi(\mathrm{RG})}{L_\phi(\odot)} = \left(\frac{r_c}{R_\odot}\right)^3 \left(\frac{T_c}{T_\odot}\right)^n = (10^{-2})^3 \cdot (10)^n = 10^{n-6}. \tag{4}$$

Setting this equal to unity gives $\boxed{n = 6}$.

(c) Production rates for BSM particles typically scale as steep power laws of temperature (e.g., T^4 or higher, depending on the process). Since $n = 6$ is a fairly modest power law by particle physics standards, many BSM production mechanisms will easily exceed this threshold. The steep temperature dependence more than compensates for the factor of $\sim 10^{-6}$ in volume, making the hot, compact cores of red giants, horizontal branch stars, and supernovae significantly more luminous in BSM particles than the entire Sun. This is why evolved stars—not the Sun—typically set the strongest energy-loss constraints.

Problem 2: White Dwarf Cooling and New Physics

(a) The number of WDs observed at a given luminosity is proportional to the time spent at that luminosity:

$$\frac{dN}{dL} \propto \left|\frac{dt}{dL}\right|. \tag{5}$$

The cooling rate is determined by the total power output: $dE/dt = -(L_\gamma + L_\phi)$. If a BSM channel L_ϕ dominates at high temperatures (high luminosities), it increases the total cooling rate in that regime. The WD passes through the high-luminosity phase faster, spending less time there. This means $|dt/dL|$ decreases at high L, producing a **deficit of WDs at the bright (hot) end** of the luminosity function relative to the SM-only prediction.

This is precisely the logic used to interpret the observed WD luminosity function: at the hot end, SM plasmon decay to neutrinos already accelerates cooling relative to surface photon emission alone, and the good agreement with data constrains any *additional* BSM cooling that would further deplete the hot WD population.

(b) A steeply temperature-dependent production rate ($\epsilon_\phi \propto T^8$) concentrates its effect on a specific part of the luminosity function. It dominates at high temperatures and becomes negligible at low temperatures, creating a *localized* deficit—a sharp feature that stands out against the smooth SM prediction. Such a feature is easy to distinguish from astrophysical uncertainties in the cooling models, which tend to affect the luminosity function more broadly.

In contrast, a temperature-independent production rate adds a constant L_ϕ to the cooling at all temperatures. At high luminosities where $L_\gamma \gg L_\phi$, the BSM contribution is a negligible correction. It only becomes important at low luminosities where $L_\gamma \lesssim L_\phi$, and even there, its effect is a gradual, featureless acceleration of cooling. This diffuse modification is much harder to distinguish from uncertainties in the baseline WD cooling physics (e.g., crystallization, convective coupling, chemical stratification), all of which also affect the faint end.

The temperature-dependent particle therefore produces a more distinctive and constrainable signature, because it creates a sharp, localized deviation in a part of the luminosity function where the SM prediction is well understood.

Problem 3: How Stable is "Stable"?

(a) The survival probability for a single particle is $e^{-t_0/\tau}$, so the fraction that has decayed is

$$f_{\rm dec} = 1 - e^{-t_0/\tau} \approx \frac{t_0}{\tau} \tag{6}$$

to leading order, valid for $\tau \gg t_0$.

(b) Each DM decay converts one particle's rest mass energy into visible products. Thus the fraction of the total DM rest mass energy deposited into the gas is $f_{\rm dec} \approx t_0/\tau$.

Setting this equal to $f \sim 10^{-9}$ (the fraction needed to fully reionize the universe):

$$\frac{t_0}{\tau} \sim 10^{-9} \quad \Longrightarrow \quad \tau \sim \frac{t_0}{10^{-9}} \sim \frac{10^{18}\text{ s}}{10^{-9}} \sim 10^{27}\text{ s}. \tag{7}$$

This is already within an order of magnitude of the quoted bound $\tau \gtrsim 10^{26}$ s, showing that the sensitivity ultimately comes from the enormous energy budget of DM relative to the small ionization energy of hydrogen. Even a tiny fraction of DM decaying releases enough energy to have dramatic cosmological consequences.

(c) If Planck can detect a $\sim 1\%$ change in the ionization fraction, the detectable threshold drops to $f \sim 10^{-2} \times 10^{-9} = 10^{-11}$. This gives

$$\tau \sim \frac{10^{18}\text{ s}}{10^{-11}} \sim 10^{29}\text{ s}. \tag{8}$$

This is even *longer* than the actual quoted bound of $\sim 10^{26}$ s. The discrepancy tells us that additional factors reduce the sensitivity relative to our idealized estimate: not all decay energy goes into ionizing hydrogen (some goes into heating, some into photons below the ionization threshold, some is redshifted away), so there is an efficiency factor $f_{\text{eff}} < 1$ that can be as small as $\sim 10^{-1}$–10^{-2} depending on the DM mass and decay channel. Additionally, the timing of energy injection matters: energy deposited well before or well after recombination has a different impact on the CMB than energy deposited near $z \sim 1000$. The quoted bound of 10^{26} s represents the *weakest* constraint across all masses—for specific masses and channels, constraints can be significantly stronger.

Problem 4: Choosing Your Target

(a) Target B has a J-factor that is 10^3 times larger than Target A, so the expected annihilation signal from Target B is about 1000 times stronger. However, Target B also has an astrophysical gamma-ray background ~ 100 times the expected DM signal. To extract a DM signal from Target B, you would need to model and subtract this background to better than $\sim 1\%$ accuracy, which is extremely challenging in practice. Mismodeling the background can produce spurious excesses or artificially strong limits. Target A, despite its $1000\times$ smaller signal, has backgrounds well below the instrumental sensitivity, so any detected flux can be attributed to DM with high confidence. This is why, in practice, dwarf spheroidal galaxies are the preferred targets for setting *robust* constraints on DM annihilation: the systematic uncertainties are much smaller, even if the raw signal is weaker.

(b) For decay, the signal scales with the D-factor. Target B has $D \sim 10^{18.5}$ vs. $D \sim 10^{18}$ for Target A—only a factor of ~ 3 difference. Now the signal advantage of the galaxy cluster is marginal, while it still carries the same factor of $\sim 100\times$ background penalty. The trade-off clearly favors Target A. More broadly, because the D-factor involves only one power of ρ (rather than ρ^2), it is much less sensitive to the peak density in the target. Dense environments like galaxy clusters do not "win" as dramatically for decay searches as they do for annihilation. This is why the notes point out that even the "blank sky" (the local DM halo) has a D-factor comparable to that of the Perseus cluster.

(c) The J-factor is $\int \rho^2 \, dr \, d\Omega$. If the DM density field has fluctuations (substructure), then by the inequality $\langle \rho^2 \rangle \geq \langle \rho \rangle^2$ (a consequence of the variance being non-negative: $\mathrm{Var}(\rho) = \langle \rho^2 \rangle - \langle \rho \rangle^2 \geq 0$), the true J-factor computed from the clumpy field exceeds what one would estimate from a smooth profile with the same total mass. This is the "boost factor." In contrast, the D-factor is $\int \rho \, dr \, d\Omega$, which is linear in ρ and therefore depends only on the total integrated mass along the line of sight. Redistributing that mass into clumps does not change the integral—$\langle \rho \rangle$ is the same regardless of how the mass is distributed.

Problem 5: The p-Wave Penalty

(a) Using the non-relativistic equipartition relation $\frac{1}{2} m_\chi v^2 = \frac{3}{2} T_f$ with $T_f = m_\chi / 20$:

$$v^2 = \frac{3T_f}{m_\chi} = \frac{3}{20} = 0.15. \tag{9}$$

So $v \sim 0.4\, c$ and $v^2 \sim 0.1$ at freeze-out.

(b) In a present-day galaxy with $v \sim 10^{-3}\, c$:

$$v_{\text{today}}^2 \sim 10^{-6}. \tag{10}$$

(c) The ratio of the present-day to freeze-out annihilation rates for a p-wave model is

$$\frac{\langle \sigma v \rangle_{\text{today}}}{\langle \sigma v \rangle_{\text{f.o.}}} = \frac{v_{\text{today}}^2}{v_{\text{f.o.}}^2} \sim \frac{10^{-6}}{10^{-1}} = 10^{-5}. \tag{11}$$

The present-day annihilation rate is suppressed by five orders of magnitude relative to the freeze-out rate. Since the freeze-out cross section for p-wave DM is similar to the s-wave thermal relic value ($\sim 3 \times 10^{-26}$ cm^3/s, up to the mild $v^2 \sim 0.1$ correction), the effective present-day cross section is $\sim 3 \times 10^{-31}$ cm^3/s. This is far below the sensitivity of current gamma-ray telescopes, making p-wave DM essentially invisible to standard indirect detection methods—even though it freezes out with nearly the same cross section as s-wave DM. This is a concrete example of why a null result from indirect detection does not rule out thermal DM in general; it specifically constrains s-wave annihilation.