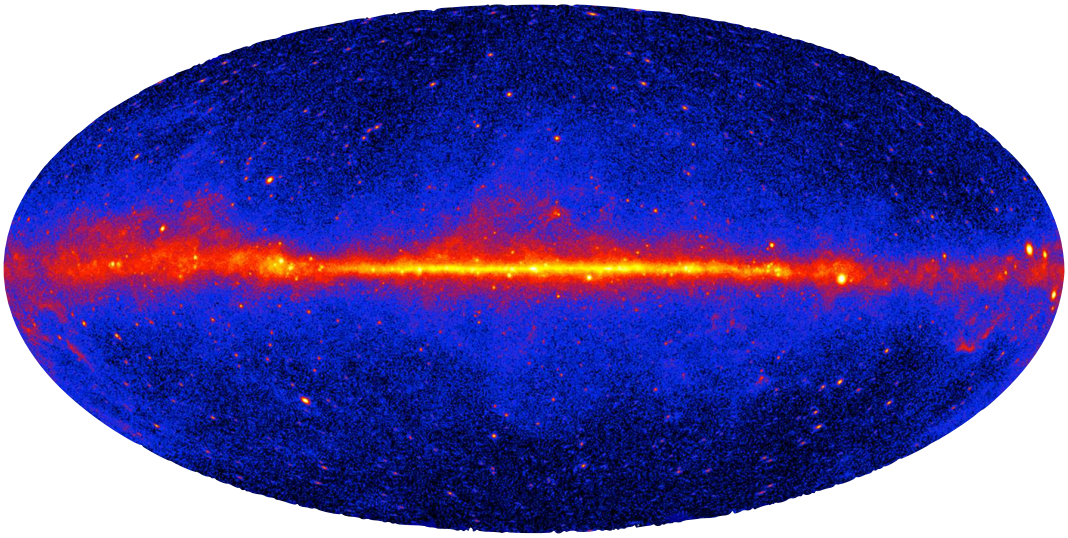
Lecture Notes on

High-Energy Astroparticle Theory

Carmelo Evoli

Gran Sasso Science Institute (GSSI), Viale F. Crispi 7, L'Aquila (Italy)
INFN Laboratori Nazionali del Gran Sasso (LNGS), Via G. Acitelli 22, Assergi (Italy)



Last updated: November 3, 2026

This document is currently under active development. The content is subject to change as revisions are ongoing. I expect to release a stable and comprehensive version no earlier than February 2027.

These lecture notes are largely based on:

- P. Blasi, *The origin of galactic cosmic rays*, A&AR, 21, 70, 2013 [1]
- D. Boncioli, *Cosmic-ray propagation in extragalactic space and secondary messengers*, Proceedings of the International School of Physics "Enrico Fermi", Course 208, Varenna 2022 [2]
- C. Evoli & U. Dupletsa, *Phenomenological models of Cosmic Ray transport in Galaxies*, Proceedings of the International School of Physics "Enrico Fermi", Course 208, Varenna 2022 [3]
- S. Gabici, C. Evoli, et al., *The origin of Galactic cosmic rays: challenges to the standard paradigm*, Int.J.Mod.Phys.D, 28, 2019 [4]
- P.D. Serpico, *Cosmic ray interactions with matter and radiation*, Proceedings of the International School of Physics "Enrico Fermi", Course 208, Varenna 2022 [5]

Additional useful readings:

- R. Blandford & D. Eichler, *Particle acceleration at astrophysical shocks: A theory of cosmic ray origin*, October 1987, 154, Physics Reports
- C.D. Dermer & G. Menon, *High energy radiation from black holes*, October 2009, Princeton University Press
- T.K. Gaisser, R. Engel & E. Resconi, *Cosmic Rays and Particle Physics*, June 2016, Cambridge University Press
- D. Overbye, *Lonely Hearts of the Cosmos: The Story of the Scientific Quest for the Secret of the Universe*, November 1999, Back Bay Books
- G.B. Rybicki & A.P. Lightman, *Radiative Processes in Astrophysics*, June 1986, Wiley-VCH
- M. Vietri, *Foundations of High-Energy Astrophysics*, April 2008, Chicago University Press

Contents

1 Introduction **1**

2 Radiative processes in astroparticle physics **2**

2.1 The multi-messenger view of astroparticle physics 2

2.2 Kinetic equation for particles with radiative losses 5

2.3 Synchrotron radiation 10

2.4 Synchrotron emission by an electron population 17

2.5 Diffuse Galactic radio emission as a probe of CR electrons 19

2.6 A cooling-break estimate for the Crab Nebula 20

2.7 Inverse Compton scattering 22

2.8 Inverse Compton emission by an electron population 31

2.9 Lepton cooling in the Milky Way 33

2.10 Pair production and γ-ray absorption 35

2.11 Optical depth and the γ-ray horizon 38

2.12 Electromagnetic cascades 41

2.13 Extragalactic cascades on the CMB and EBL 46

2.14 Nuclei–proton interactions 50

2.15 The π^0 γ-ray spectrum 57

2.16 Diffuse Galactic γ-ray emission 63

2.17 Nuclei–photon interactions 68

2.18 Neutrinos from cosmic accelerators 82

References **90**

A Appendix **96**

A.1 Formal steady-state solution with continuous energy losses 96

A.2 Motion of a charged particle in a uniform magnetic field 99

A.3 Threshold kinematics in astroparticle reactions 100

1 Introduction

ADD HERE THE COSMIC RAY PLOT.

Units. Formulas are written in Gaussian units with c, $\hbar$, and k explicit. Kinematic relations, such as thresholds and invariant masses, are often written with $c = 1$ and the metric $(+,-,-,-)$, so that $s = (P_1 + P_2)^2$ is the squared center-of-momentum energy. Particle and photon energies are quoted in eV.

Cosmology. We use a flat ΛCDM cosmology with $H_0 = 70\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$, $\Omega_m = 0.3$, and $\Omega_\Lambda = 0.7$, and neglect radiation in the expansion rate $H(z)$ of Eq. (2.208). The Hubble time is $t_\mathrm{H} = H_0^{-1} \simeq 14$ Gyr, and the Hubble radius is $c\,t_\mathrm{H} \simeq 4.3$ Gpc.

2 Radiative processes in astroparticle physics

2.1 The multi-messenger view of astroparticle physics

Astroparticle physics is intrinsically *multi-messenger*: solving the puzzles of the high-energy Universe requires piecing together complementary clues. These are carried by charged cosmic rays (CRs), photons across the electromagnetic spectrum (from radio waves to TeV–PeV γ rays), neutrinos, and, for some transient systems, gravitational waves. Each messenger probes different physics, and each observed signal contains the imprint of both its production at the source and its propagation to Earth [6–8].

Cosmic rays provide a direct sample of accelerated matter. Their energy spectrum, nuclear composition, and secondary content constrain the acceleration mechanism and the environments through which the particles have propagated. Because CRs are charged, however, Galactic and extragalactic magnetic fields deflect them and generally erase their directional association with individual sources.

Photons usually offer the best angular resolution and the largest event samples. Their spectra, light curves, and sky maps reveal source locations, variability, and spatial structure over the broad energy range summarized in Table 1. Their interpretation is not unique: leptonic and hadronic particle populations can produce similar photon spectra, while absorption and reprocessing may substantially alter the radiation before it reaches us.

Neutrinos interact so weakly that they can escape dense regions and cross cosmological distances with little attenuation. Their electrical neutrality preserves directional information. In the high-energy production channels considered here, neutrinos arise mainly from charged-pion and muon decays and trace interactions of accelerated protons or nuclei. The same small interaction probability, however, makes neutrinos difficult to detect and usually limits the available event statistics.

Gravitational waves provide a different view again: they trace the coherent bulk dynamics of compact objects and can determine the time and characteristic scale of a violent transient independently of its electromagnetic emission.

These messengers are therefore most informative when interpreted together. For example, a γ-ray spectrum alone may not identify whether electrons or nuclei power a source, whereas an associated high-energy neutrino signal provides independent evidence for hadronic interactions in that source. A common model must reproduce the spectral, spatial, and temporal information carried by all available messengers. This joint requirement helps disentangle particle acceleration and transport from radiative processing inside the source and along the line of sight. Developing the radiative framework needed for this joint interpretation is the aim of this chapter.

Why radiative processes matter. Radiative processes play two complementary roles in astroparticle physics. They are *diagnostic*, because the radiation produced by high-energy particles encodes the properties of both the particles and their environment. At the same

time, they are *dynamical*, because the interactions that generate the radiation can remove or redistribute energy and thereby reshape the particle populations we are trying to infer.

The diagnostic role follows from the fact that every emission channel couples a parent population to a target field. Synchrotron emission connects the electron spectrum to the magnetic field (Sections 2.3 and 2.4), while inverse Compton emission connects the same electrons to the ambient radiation field (Sections 2.7 and 2.8). Hadronic interactions instead link cosmic rays to gas or photon backgrounds, producing secondary γ rays and neutrinos (Sections 2.14, 2.17 and 2.18). Measurements of spectra, luminosities, and morphologies can therefore be translated into constraints on particle distributions, magnetic and radiation fields, gas densities, and source dimensions.

The dynamical role becomes important when an energy-loss or destruction time is comparable to, or shorter than, the smaller of the source age and the particle escape time. The kinetic framework developed in Section 2.2 shows how continuous losses reshape the injected spectrum and, depending on the injection history and escape, produce spectral steepening, cooling breaks, and cooling cutoffs. Synchrotron and inverse Compton losses provide the main examples for leptons (Section 2.9); photon–photon absorption and electromagnetic cascades redistribute high-energy γ-ray emission (Sections 2.12 and 2.13); and photohadronic interactions can modify both the energy and the composition of cosmic rays at the highest energies (Section 2.17).

The two roles need not be equally important in a given environment. A process may leave the parent population almost unchanged and still provide an excellent observable tracer. A classic example is the diffuse Galactic γ-ray emission. In the diffuse Galactic medium, the energy-loss time of GeV–TeV CR protons due to inelastic pp collisions typically exceeds their escape time, so these collisions make only a small correction to the proton spectrum. Yet the photons from π^0 decay produced in CR–gas collisions provide a dominant contribution to the diffuse Galactic γ-ray emission above $\sim 1\,\mathrm{GeV}$ in the plane, and the accompanying neutrinos from charged-pion decay provide an unavoidable counterpart (Fig. 1). In this way, "secondary" radiation becomes a tracer of the product of CR and gas densities, allowing us—together with local CR measurements—to map CR gradients and test propagation models across the Galaxy [9, 10].

This example, developed quantitatively in Section 2.16, illustrates the logic followed throughout this chapter: first determine how an interaction changes the parent particles, then calculate the radiation it produces, and finally connect that radiation to an observable. Keeping the diagnostic and dynamical roles together is essential for quantitative inference in high-energy astroparticle physics [5].

Band	**Wavelength** λ [cm]	**Frequency** ν [Hz]	**Energy** E [eV]
Radio	$\gtrsim 10^{2}$	$\lesssim 3\times10^{8}$	$\lesssim 10^{-6}$
Microwaves	$\sim 10^{-1}$ – 10^{2}	$\sim 3\times10^{8}$ – 3×10^{11}	$\sim 10^{-6}$ – 10^{-3}
Infrared (IR)	$\sim 8\times10^{-5}$ – 10^{-1}	$\sim 3\times10^{11}$ – 4×10^{14}	$\sim 10^{-3}$ – 1.5
Visible (optical)	$\sim 4\times10^{-5}$ – 8×10^{-5}	$\sim 4\times10^{14}$ – 7×10^{14}	~ 1.5 – 3
Ultraviolet (UV)	$\sim 10^{-6}$ – 4×10^{-5}	$\sim 7\times10^{14}$ – 3×10^{16}	~ 3 – 10^{2}
X-rays	$\sim 10^{-9}$ – 10^{-6}	$\sim 3\times10^{16}$ – 3×10^{19}	$\sim 10^{2}$ – 10^{5}
Gamma rays	$\lesssim 10^{-9}$	$\gtrsim 3\times10^{19}$	$\gtrsim 10^{5}$

Table 1: Approximate ranges of the electromagnetic spectrum in wavelength, frequency, and photon energy. Boundaries between bands are conventional and may vary slightly in the literature.

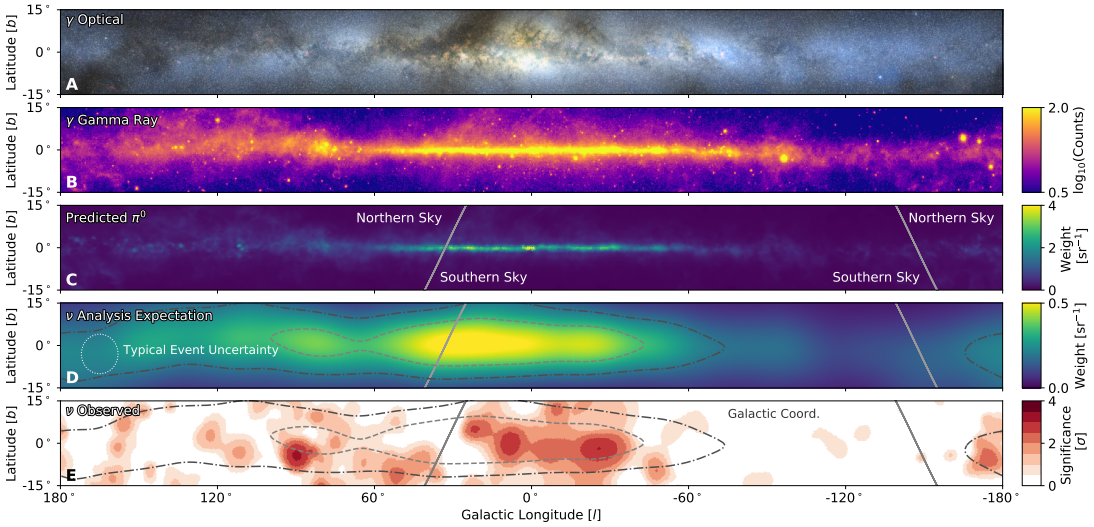


Figure 1: The Milky Way in photons and neutrinos. All panels are in Galactic coordinates ($\pm15°$ in latitude, $\pm180°$ in longitude). (A) Optical image, partly obscured by dust. (B) *Fermi*-LAT photon counts above 1 GeV in the 12-year survey. (C) Expected neutrino distribution inferred from the *Fermi*-LAT π^0 emission model. (D) The same template after accounting for IceCube's sensitivity and the angular uncertainty of cascade events. Contours enclose 20% and 50% of the predicted signal. (E) Local, pre-trial significance from an all-sky point-source search using cascade events, with the same contours as in (D). Individual hot spots did not establish point-source detections; a separate spatial-template analysis yielded 4.5σ evidence for Galactic-plane neutrino emission after accounting for trials. Adapted from [11].

2.2 Kinetic equation for particles with radiative losses

To illustrate how a consistent treatment of target fields, cooling, and escape emerges, we start from a minimal model in which particles undergo continuous energy losses. Even in this simple setting, the steady-state distribution generally differs from the injected spectrum: radiative cooling and escape imprint characteristic breaks and cutoffs that can be read as diagnostics of the acceleration region [7].

Let $N(E,t)$ denote the differential *number density* per unit energy ($\mathrm{L}^{-3}\,\mathrm{E}^{-1}$), and $Q(E,t)$ the *injection rate density* ($\mathrm{L}^{-3}\,\mathrm{T}^{-1}\,\mathrm{E}^{-1}$). Continuous losses are described by

$$\dot{E} \equiv \frac{\mathrm{d}E}{\mathrm{d}t} = -b(E)\,, \qquad b(E) > 0\,, \tag{2.1}$$

and escape is parametrized by an energy-dependent residence time $\tau_{\mathrm{esc}}(E)$.

Consider the energy bin $[E, E+\Delta E]$ at time $t+\Delta t$. The particles occupying this bin must have had energy in a neighboring interval $[E', E'+\Delta E']$ at time t, related by the backward map

$$E' = E + b(E)\,\Delta t, \qquad E' + \Delta E' = (E+\Delta E) + b(E+\Delta E)\,\Delta t. \tag{2.2}$$

Expanding $b(E+\Delta E)$ to first order in ΔE gives the Jacobian factor

$$\Delta E' = \Delta E + \frac{\partial b}{\partial E}\,\Delta E\,\Delta t. \tag{2.3}$$

Number conservation between t and $t+\Delta t$ implies

$$N(E, t+\Delta t)\,\Delta E = N(E', t)\,\Delta E'. \tag{2.4}$$

Expanding to first order in Δt,

$$N(E, t+\Delta t) = N(E,t) + \frac{\partial N}{\partial t}\,\Delta t, \qquad N(E', t) = N(E,t) + \frac{\partial N}{\partial E}\,b(E)\,\Delta t. \tag{2.5}$$

Insert Eq. (2.5) into Eq. (2.4), keep only $\mathcal{O}(\Delta t)$, and cancel the common $N\,\Delta E$:

$$\frac{\partial N}{\partial t}\,\Delta t\,\Delta E = \frac{\partial N}{\partial E}\,b\,\Delta t\,\Delta E + N\,\frac{\partial b}{\partial E}\,\Delta t\,\Delta E. \tag{2.6}$$

Dividing by $\Delta t\,\Delta E$ and taking $\Delta t \to 0$ yields

$$\frac{\partial N}{\partial t} = \frac{\partial}{\partial E}\left[b(E)\,N(E,t)\right]. \tag{2.7}$$

This is the continuity equation in energy space (without sources or sinks). The signed energy-space flux is $\dot{E}N = -bN$; equivalently, bN is the magnitude of the particle flux directed from high to low energies. The Jacobian $\Delta E'/\Delta E = 1 + (\partial_E b)\Delta t$ ensures number conservation during this compression or rarefaction in energy space.

Injection acts as the local source Q, whereas escape supplies the sink $-N/\tau_{\mathrm{esc}}$. Including both gives the full kinetic (continuity) equation,

$$\frac{\partial N}{\partial t} = \frac{\partial}{\partial E}[b(E)\,N(E,t)] - \frac{N(E,t)}{\tau_{\rm esc}(E)} + Q(E,t)\,. \tag{2.8}$$

This equation describes how a particle distribution sourced by Q evolves under radiative losses and escape. One may obtain Eq. (2.8) equivalently by integrating the phase-space transport equation over angles and volume and then projecting onto energy; the bin-mapping derivation above is, however, especially transparent for a first exposure.

Equilibrium spectrum with radiative losses. In steady state ($\partial_t N = 0$), and assuming that the distribution vanishes sufficiently fast at high energy (so that the energy-space flux $b(E)N(E) \to 0$ as $E \to \infty$), the kinetic equation Eq. (2.8) admits a closed-form solution.

If escape is neglected ($\tau_{\rm esc} \to \infty$),

$$N(E) = \frac{1}{b(E)} \int_E^\infty {\rm d}E'\, Q(E')\,, \tag{2.9}$$

whereas if escape is competitive with losses one finds

$$N(E) = \frac{1}{b(E)} \int_E^\infty {\rm d}E'\, Q(E') \exp\left[-\int_E^{E'} \frac{{\rm d}E''}{b(E'')\,\tau_{\rm esc}(E'')}\right]. \tag{2.10}$$

Derivations of Eqs. (2.9) and (2.10) are given in Section A.1.

Equations Eqs. (2.9) and (2.10) admit a simple physical interpretation. The quantity $b(E)N(E)$ is the magnitude of the downward particle flux through energy E (number per unit volume and time). In the absence of escape, this flux must balance the total injection rate into all energies $E' \geq E$, namely $\int_E^\infty Q(E')\,{\rm d}E'$, which immediately yields Eq. (2.9).

When escape is present, particles injected at E' contribute to $N(E)$ only if they remain in the system long enough to cool down to E. This is encoded by the attenuation factor

$$\exp\left[-\int_E^{E'} \frac{{\rm d}E''}{b(E'')\,\tau_{\rm esc}(E'')}\right],$$

which is precisely the survival probability against escape along the cooling trajectory from E' to E. In the limit $\tau_{\rm esc} \to \infty$ the exponential tends to unity and Eq. (2.10) reduces to the no-escape result Eq. (2.9).

In the following paragraphs we develop intuition by applying Eq. (2.9) to simple choices of $Q(E)$ (monoenergetic injection, power-law injection, finite-duration injection), highlighting how cooling produces spectral steepening, breaks, and cutoffs.

Monochromatic injection. Consider continuous, *monoenergetic* injection,

$$Q(E) = Q_0\,\delta(E - E_{\rm inj}),$$

and neglect escape. Inserting this source into Eq. (2.9) gives

$$N(E) = \frac{1}{b(E)} \int_E^\infty Q_0\,\delta(E' - E_{\rm inj})\,{\rm d}E' = \frac{Q_0}{b(E)}\,\Theta(E_{\rm inj} - E)\,, \tag{2.11}$$

i.e.

$$N(E) = \frac{Q_0}{b(E)}, \quad \text{for } E < E_{\rm inj}\,, \tag{2.12}$$

and $N(E) = 0$ for $E > E_{\rm inj}$. Thus, below $E_{\rm inj}$ the steady spectrum directly mirrors the loss law: if $b(E) \propto E^m$ with $m > 0$, then

$$N(E) \propto E^{-m}, \quad \text{for } E < E_{\rm inj}\,.$$

Spectral steepening. For a *power-law injection* $Q(E) = Q_0\,E^{-p}$ with $p > 1$, and continuous losses $b(E) = b_0\,E^m$, the no-escape steady-state solution Eq. (2.9) gives

$$N(E) = \frac{1}{b(E)}\int_E^\infty Q(E')\,{\rm d}E' \propto \frac{\int_E^\infty E'^{-p}\,{\rm d}E'}{E^m} \propto E^{-(p+m-1)}\,. \tag{2.13}$$

The equilibrium spectral index therefore differs from the injection index by $\Delta p = m - 1$: losses *steepen* the spectrum when $m > 1$, leave its slope unchanged when $m = 1$, and *harden* it when $m < 1$. For example, ionization and Coulomb losses depend only weakly on energy ($m \simeq 0$), giving $N(E) \propto E^{-(p-1)}$.

For synchrotron and inverse Compton losses in the Thomson regime (see Sections 2.3 and 2.7), one has $m = 2$: an injected spectrum with $p = 2$ yields an equilibrium spectrum $N(E) \propto E^{-3}$ in the cooling-dominated range.

The cooling break. The previous results allow us to understand the origin of the *cooling break*, one of the most informative features in the non-thermal spectra of high-energy sources: its location links source age or escape timescale to the loss rate $b(E)$, and hence to the strength of the relevant target fields.

To build intuition, first neglect radiative losses. If particles are continuously injected for a time t_*, the distribution simply accumulates:

$$N(E) \;\simeq\; Q(E)\,t_{\rm res}(E), \qquad t_{\rm res}(E) = \min\big[t_*,\,\tau_{\rm esc}(E)\big]. \tag{2.14}$$

This "reservoir" picture shows that when $t_* \ll \tau_{\rm esc}(E)$ the spectrum retains the injected slope, $N(E) \approx Q(E)\,t_*$.

Radiative losses introduce an energy-dependent residence time in energy space. When cooling is the fastest process, particles are advected through energy too quickly to accumulate and the steady state is no longer $Q \times t_*$. For a power-law injection $Q(E) \propto E^{-p}$ with $p > 1$, the loss-dominated solution is

$$N(E) = \frac{1}{b(E)}\int_E^\infty Q(E')\,{\rm d}E' = \frac{Q(E)}{p-1}\,\frac{E}{b(E)} \equiv \frac{Q(E)\,\tau_{\rm cool}(E)}{p-1}, \tag{2.15}$$

where

$$\tau_{\rm cool}(E) \equiv \frac{E}{b(E)}\,. \tag{2.16}$$

Thus $N \sim Q\,\tau_{\rm cool}$ as a scaling relation; the exact normalization retains the slope-dependent factor $(p-1)^{-1}$.

The transition between the reservoir regime and the cooling-dominated regime occurs at the *cooling break*[1] E_c, defined implicitly by

$$\tau_{\rm cool}(E_c) \simeq \min[t_*, \tau_{\rm esc}(E_c)] \,. \tag{2.17}$$

For power-law losses $b(E) = b_0 E^m$ with $m > 1$, and for an age-limited system ($t_* \ll \tau_{\rm esc}$),

$$E_c \simeq \left[\frac{1}{b_0 t_*}\right]^{1/(m-1)} \quad \Longrightarrow \quad m = 2 : \; E_c \simeq \frac{1}{b_0 t_*} \,. \tag{2.18}$$

As a result, for an injected power law $Q(E) \propto E^{-p}$, and provided the effective residence time below the break is approximately energy independent, the spectrum exhibits the canonical broken behavior,

$$N(E) \propto \begin{cases} E^{-p}, & E \ll E_c \quad \text{(age/escape limited)}, \\ E^{-(p+m-1)}, & E \gg E_c \quad \text{(loss dominated)}. \end{cases} \tag{2.19}$$

Thus the spectral index steepens by $\Delta p = m - 1$ across the break; in particular, for synchrotron/IC losses in the Thomson regime ($m = 2$) the steepening is $+1$. Locating E_c is therefore highly diagnostic, as it connects the source age (or escape time) to the target-field strengths that set $b(E)$ as shown in Eq. (2.18). If escape dominates and $\tau_{\rm esc}(E)$ is energy dependent, the low-energy branch is instead proportional to $Q(E)\tau_{\rm esc}(E)$, so its slope already differs from the injection slope.

Injection burst, cooling cutoff, and cooling spike. Consider an impulsive injection at $t = 0$ with initial spectrum $N_i(E) \equiv N(E, 0^+)$, followed by continuous radiative losses $\dot{E} = -b(E)$. A convenient starting point is the cooling-time relation along a characteristic (from Eq. (2.1)):

$$t = \int_E^{E_i} \frac{{\rm d}E'}{b(E')} \,, \tag{2.20}$$

which implicitly defines the injection energy $E_i \equiv E_i(E, t)$ of particles observed at energy E after a time t.

The spectrum at time t follows from simple number conservation under this energy-space advection,

$$N(E, t)\,{\rm d}E = N_i(E_i)\,{\rm d}E_i \quad \Longrightarrow \quad N(E, t) = N_i(E_i)\,\frac{{\rm d}E_i}{{\rm d}E} \,. \tag{2.21}$$

Differentiating Eq. (2.20) at fixed t yields the Jacobian,

$$\frac{{\rm d}E_i}{{\rm d}E} = \frac{b(E_i)}{b(E)} \,, \tag{2.22}$$

[1] The term "cooling break" is widespread but can be misleading: the particles are non-thermal and no unique temperature characterizes a power-law population. What is called "cooling" here is phase-space *advection in energy* due to continuous radiative losses, rather than thermodynamic cooling of a Maxwellian gas. A more precise term would be "radiative-loss break" (or simply "loss break"). We keep the conventional nomenclature for continuity with the literature.

so the burst solution can be written in the useful form

$$N(E,t) = \frac{b(E_i)}{b(E)}\, N_i(E_i)\,, \qquad t = \int_E^{E_i} \frac{\mathrm{d}E'}{b(E')}\,. \tag{2.23}$$

Cooling cutoff. A distinctive feature of super-linear losses ($m > 1$ for $b \propto E^m$) is that the time required to cool from arbitrarily high energy down to E,

$$t_{\rm pass}(E) \;=\; \int_E^{\infty} \frac{\mathrm{d}E'}{b(E')}\,, \tag{2.24}$$

is finite. If $t > t_{\rm pass}(E)$, then no finite E_i can satisfy Eq. (2.20), and the distribution must vanish at that energy:

$$t > t_{\rm pass}(E) \quad \Longrightarrow \quad N(E,t) = 0\,.$$

This defines a time-dependent cutoff $E_{\rm cut}(t)$ through $t_{\rm pass}(E_{\rm cut}) = t$. For $b(E) = b_0 E^m$ with $m > 1$,

$$t_{\rm pass}(E) = \frac{1}{(m-1)\, b_0\, E^{m-1}} \quad \Longrightarrow \quad E_{\rm cut}(t) = \left[(m-1)\, b_0\, t\right]^{-1/(m-1)}$$

and, in particular,

$$m = 2: \; E_{\rm cut} \;=\; \frac{1}{b_0 t}\,.$$

Under continuous injection this sharp cutoff is smoothed into a cooling break; for a single burst it remains a moving high-energy edge at $E_{\rm cut}(t)$.

Cooling spike (pile-up) for hard injection. The same mapping also explains why a pile-up may appear just below the cutoff. As $E \to E_{\rm cut}^-$, the required injection energy grows without bound ($E_i \to \infty$), and the transformation $E_i \mapsto E$ compresses a wide range of initial energies into a narrow interval near $E_{\rm cut}$. For power-law losses one may express the characteristic in terms of $E_{\rm cut}(t)$ as

$$E_i(E,t) = E\left[1 - \left(\frac{E}{E_{\rm cut}(t)}\right)^{m-1}\right]^{-1/(m-1)}, \qquad (E < E_{\rm cut})\,. \tag{2.25}$$

For an initial power law $N_i(E_i) \propto E_i^{-p}$, Eq. (2.23) then gives

$$N(E,t) \propto E^{-p}\left[1 - \left(\frac{E}{E_{\rm cut}(t)}\right)^{m-1}\right]^{-\frac{m-p}{m-1}}, \qquad (E < E_{\rm cut}). \tag{2.26}$$

Thus $N(E,t)$ diverges as $E \to E_{\rm cut}^-$ if $p < m$: sufficiently hard injection produces a cooling spike (pile-up) immediately below the cutoff. In particular, for synchrotron/IC losses in the Thomson regime ($m = 2$),

$$N(E,t) \propto E^{-p}\left(1 - \frac{E}{E_{\rm cut}}\right)^{-(2-p)},$$

so injection harder than E^{-2} ($p < 2$) yields a pile-up near $E_{\rm cut}(t)$.

In practice the spike is regulated—and typically smoothed into a finite bump—by a maximum injection energy, a finite-duration injection history, escape, and/or departures from a pure power-law loss rate (e.g. Klein–Nishina effects for IC, Section 2.7).

2.3 Synchrotron radiation

Synchrotron radiation—schematically $e + B \to e + \gamma + B$—is emitted whenever a *relativistic* charged particle is forced to curve its trajectory by a magnetic field. In the opposite, non-relativistic limit the same phenomenon is called *cyclotron emission* [6, 7].

A natural starting point is the familiar Larmor result: in c.g.s. units, a charge q undergoing an acceleration of magnitude $a = \|\mathbf{a}\|$ radiates the power

$$P_{\text{Larmor}} = \frac{2}{3}\frac{q^2 a^2}{c^3}. \tag{2.27}$$

The q^2 dependence is characteristic of emission by a single accelerated charge: at fixed acceleration the radiated field is proportional to q, and the power to its square. Further powers of the charge appear when the acceleration is itself produced by an electromagnetic force, as in synchrotron emission below or in Thomson scattering, whose cross section scales as q^4/m^2.

When the motion becomes relativistic, time dilation and relativistic beaming enhance the emission. The appropriate generalization is the *Liénard* formula [6],

$$P = \frac{2}{3}\frac{q^2\gamma^6}{c^3}\left(a^2 - (\boldsymbol{\beta}\times\mathbf{a})^2\right) = \frac{2}{3}\frac{q^2\gamma^4}{c^3}\left(a_\perp^2 + \gamma^2 a_\parallel^2\right), \tag{2.28}$$

where $\boldsymbol{\beta} = \mathbf{v}/c$, and $a_\perp$ and $a_\parallel$ are the components of $\mathbf{a}$ perpendicular and parallel to the velocity $\mathbf{v}$, respectively. This decomposition will be particularly handy for synchrotron motion, since the magnetic Lorentz force is perpendicular to $\mathbf{v}$ and thus predominantly contributes through $a_\perp$.

Motion in a magnetic field and synchrotron power for a single charge. In a purely magnetic field the Lorentz force does no work: it changes the *direction* of the momentum, but not its magnitude. As a result the particle energy is conserved and the acceleration is everywhere perpendicular to the velocity (see Section A.2). Starting from $\dot{\mathbf{p}} = \frac{q}{c}\,\mathbf{v}\times\mathbf{B}$ with $\mathbf{p} = \gamma m\mathbf{v}$, the magnitude of the *perpendicular* acceleration is

$$a_\perp = \frac{|q|\,v_\perp\,B}{\gamma\,m\,c} = \frac{|q|\,B}{\gamma\,m\,c}\,v\sin\theta, \tag{2.29}$$

where θ is the pitch angle between $\mathbf{v}$ and $\mathbf{B}$. Since $a_\parallel = 0$, inserting this expression into Eq. (2.28) immediately gives the synchrotron power emitted by a single relativistic particle:

$$P_{\text{s}} = \frac{2}{3}\frac{q^4B^2}{m^2c^3}\,\gamma^2\beta^2\sin^2\theta \xrightarrow{\beta\to 1} \frac{2}{3}\frac{q^4B^2}{m^2c^3}\,\gamma^2\sin^2\theta. \tag{2.30}$$

The fourth power of the charge has two origins: the radiation formula contributes q^2, while the magnetic acceleration itself is proportional to $|q|$.

If pitch angles are isotropically distributed, the geometric factor averages to

$$\langle\sin^2\theta\rangle = \frac{1}{4\pi}\int \mathrm{d}\Omega\,\sin^2\theta = \frac{2}{3}.$$

Averaging Eq. (2.30) and rewriting the result in terms of the magnetic energy density $U_B = B^2/(8\pi)$ and the electron Thomson cross section $\sigma_T = \frac{8\pi}{3}(e^2/m_e c^2)^2$, one obtains

$$\begin{aligned} \langle P_s \rangle &= \frac{4}{9}\frac{q^4 B^2}{m^2 c^3}\gamma^2\beta^2 \\ &= \frac{4}{3} c\,\sigma_T\, Z^4 \left(\frac{m_e}{m}\right)^2 \gamma^2\beta^2\, U_B, \qquad Z \equiv \frac{|q|}{e}. \end{aligned} \tag{2.31}$$

Two remarks are particularly useful to keep in mind. First, at *fixed particle energy* $E = \gamma mc^2$ one has $\gamma \propto 1/m$, so that $P_s \propto Z^4 m^{-4}$, a very steep mass dependence. For a proton and an electron of the same energy, for example, using the mass ratio $m_p/m_e \simeq 6\pi^5$ [12] gives $P_p/P_e = (m_e/m_p)^4 \simeq 9\times10^{-14}$. Synchrotron losses are therefore overwhelmingly more important for electrons and positrons than for nuclei. Second, the loss rate is directly proportional to the magnetic-field energy density U_B: stronger fields simply provide a larger background into which the particle can radiate.

Cooling time. For electrons ($m = m_e$) the synchrotron loss time (or cooling time) is defined as the ratio between the particle energy and its radiated power,

$$\tau_{\rm loss}(\gamma) \;\simeq\; \frac{E}{|{\rm d}E/{\rm d}t|} = \frac{\gamma m_e c^2}{\langle P_s\rangle}. \tag{2.32}$$

Using Eq. (2.31) one obtains

$$\tau_{\rm loss}(\gamma) = \frac{3}{4}\frac{m_e c}{\sigma_T}\frac{1}{\gamma\, U_B} \;\propto\; \frac{1}{E}\,. \tag{2.33}$$

In other words, *the higher the energy, the shorter the lifetime*: energetic electrons cool very efficiently. This simple scaling is one of the main reasons synchrotron-dominated sources tend to develop *steeper* spectra at high energies [7].

The Crab Nebula provides a familiar illustration: its synchrotron emission appears more compact at higher frequencies [13]. The higher-energy electrons responsible for this emission cool more rapidly and therefore radiate closer to their acceleration region.

For quick estimates it is handy to write the result in a numerical form:

$$\tau_{\rm loss}(E_e) \simeq 0.3\ {\rm Myr}\left(\frac{U_B}{\rm eV\ cm^{-3}}\right)^{-1}\left(\frac{E_e}{\rm TeV}\right)^{-1}. \tag{2.34}$$

From beaming to the single-electron spectrum. Our goal is to derive the spectrum emitted by a single electron of fixed energy. A complete calculation starts from the time profile of the radiation received by a distant observer and obtains the spectral power from the squared modulus of the Fourier amplitude of the radiated field [6]. Before presenting the results of that calculation, however, we can identify the characteristic frequency through a much simpler geometric argument. A short feature of duration $\Delta t_{\rm obs}$ in the observed waveform carries appreciable Fourier content up to frequencies of order $1/\Delta t_{\rm obs}$. Estimating

the pulse duration will therefore locate the characteristic scale of the spectrum, although it will not determine its detailed shape.

The key relativistic ingredient is *beaming*: for $\gamma \gg 1$ the radiation is confined within a narrow cone of half-opening angle $\sim 1/\gamma$ around the instantaneous velocity direction $\mathbf{v}$. As the electron gyrates, this cone sweeps across the sky. To a distant observer the signal therefore appears as a sequence of short pulses, received only when the beaming cone momentarily crosses the line of sight. The duration of these pulses is the key quantity estimated below.

Relativistic aberration. A clean way to see the origin of beaming is through relativistic aberration. Let the primed frame move with speed βc along x with respect to the unprimed (observer) frame. The Lorentz transform is

$$ct' = \gamma(ct - \beta x), \qquad x' = \gamma(x - \beta ct), \qquad y' = y\,.$$

Velocity components transform according to

$$u_x = \frac{u'_x + \beta c}{1 + \beta u'_x/c'}, \qquad u_y = \frac{u'_y}{\gamma(1 + \beta u'_x/c)}.$$

Therefore the propagation angle θ in the observer frame obeys

$$\tan\theta = \frac{u_y}{u_x} = \frac{u'_y}{\gamma(u'_x + \beta c)} = \frac{u' \sin\theta'}{\gamma(\beta c + u' \cos\theta')}.$$

For photons ($u' = c$) this reduces to the standard *aberration formula*,

$$\tan\theta = \frac{\sin\theta'}{\gamma\,(\beta + \cos\theta')}. \tag{2.35}$$

A photon emitted along the motion ($\theta' = 0$) remains forward-directed ($\theta = 0$). More interestingly, a photon emitted sideways in the emitter frame ($\theta' = \pi/2$) is seen at $\theta \simeq 1/\gamma$ when $\gamma \gg 1$ (Fig. 2). In other words, emission that is even moderately wide in the comoving frame is squeezed into a narrow forward cone in the lab frame. This is precisely the beaming effect that produces the short synchrotron pulses of a gyrating electron, and it also underlies the extreme collimation inferred in strongly relativistic outflows such as GRB jets [8].

Pulse duration and characteristic frequency. In the non-relativistic limit the emission is essentially monochromatic at the gyrofrequency $\nu_0 = \Omega_0/(2\pi) = |q|B/(2\pi mc)$ (see Table 3): the charge executes nearly uniform circular motion, the radiation pattern is broad, and there is no strong time-compression of the signal.

For relativistic particles the picture changes in a very intuitive way. Beaming squeezes the emission into a narrow cone of half-opening angle $\sim 1/\gamma$ around the instantaneous velocity. The observer therefore "sees" the electron only while this cone sweeps across the line of sight. What is intrinsically continuous emission becomes, in the observer frame, a train of short pulses. Since the frequency content is controlled by the Fourier transform of this

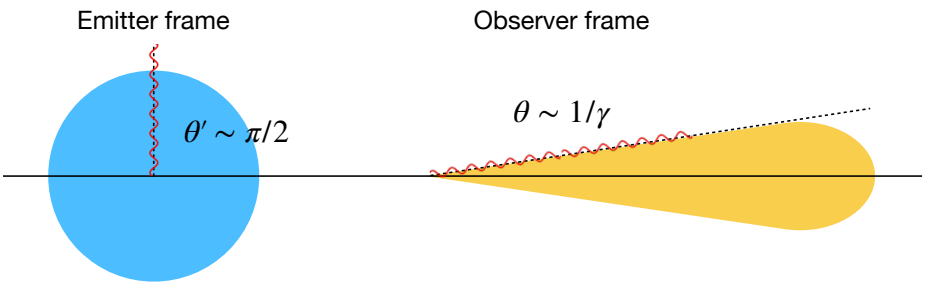


Figure 2: An isotropic pattern in the emitter frame is concentrated into a narrow cone $\sim 1/\gamma$ about the velocity in the observer frame.

time-domain signal, the highest (and most characteristic) frequencies are controlled by the *observed* pulse width $\Delta t_{\rm obs}$: roughly, $\nu_{\rm s} \sim 1/\Delta t_{\rm obs}$.

For clarity, consider an order-unity pitch angle, as depicted in Fig. 3, and let the line of sight be nearly aligned with the velocity when the beam points toward the observer. The beam sweeps across the line of sight over a small orbital phase interval

$$\Delta\phi \sim \frac{2}{\gamma}.$$

Point A marks the first position along this arc at which radiation within the cone of half-opening angle $\sim 1/\gamma$ is directed toward the observer, while B marks the last.

The corresponding *emission* duration in the lab frame is then

$$\Delta t \sim \frac{\Delta\phi}{\Omega_B} = \frac{2/\gamma}{\Omega_0/\gamma} = \frac{2}{\Omega_0},$$

where we used $\Omega_B = \Omega_0/\gamma$ with $\Omega_0 \equiv |q|B/(mc)$.

However, the *received* pulse is even shorter than Δt. Photons emitted later start closer to the observer, and therefore have a slightly smaller travel time. Denoting by δt_i the propagation time from the emission point i to the observer, one finds

$$\Delta t_{\rm obs} = (t_B + \delta t_B) - (t_A + \delta t_A) = (t_B - t_A) + (\delta t_B - \delta t_A) = \frac{\rm AB}{v} - \frac{\rm AB}{c} = \Delta t\,(1-\beta). \tag{2.36}$$

For $\gamma \gg 1$, $1-\beta \simeq 1/(2\gamma^2)$, so the pulse is compressed by roughly $2\gamma^2$:

$$\Delta t_{\rm obs} = \Delta t\,(1-\beta) \simeq \frac{\Delta t}{2\gamma^2} \quad\Longrightarrow\quad \Delta t_{\rm obs} \sim \frac{1}{\gamma^2\Omega_0}. \tag{2.37}$$

It is sometimes convenient to express the same scaling in terms of the gyroradius $r_g = v_\perp/\Omega_B$:

$$\Delta t_{\rm obs} \sim \frac{r_g}{\gamma^3 v_\perp}, \tag{2.38}$$

so for an ultra-relativistic electron ($v_\perp \simeq c$) the observed pulse width shrinks like γ^{-3} at fixed r_g.

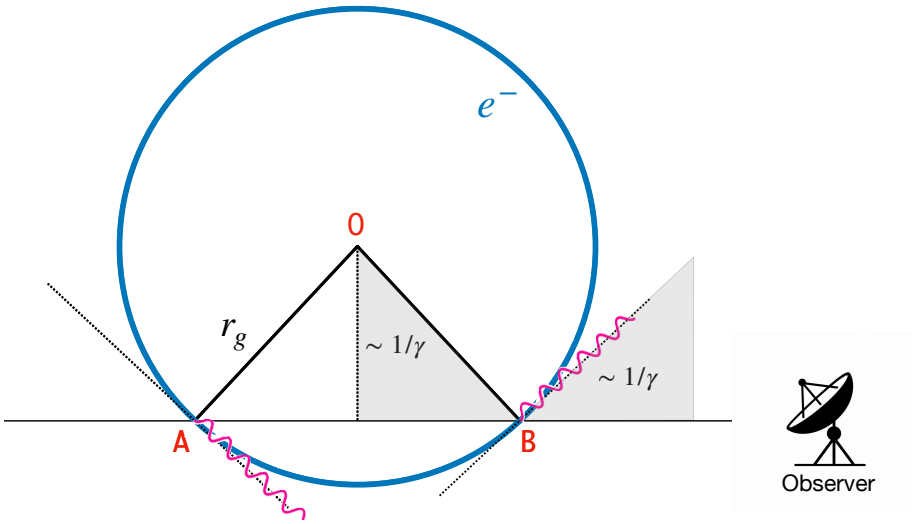


Figure 3: As the electron gyrates, its emission is confined to a cone of half-opening $\sim 1/\gamma$ about the instantaneous velocity. The observer receives radiation only while the cone sweeps across the line of sight (between points A and B), over a phase interval $\Delta\phi \sim 2/\gamma$.

The observer therefore measures a broad synchrotron spectrum with a characteristic frequency set by the inverse pulse duration. Restoring the pitch-angle dependence ($v_\perp = v\sin\theta$) suppressed in this geometric estimate gives [6]

$$\nu_s \sim \frac{1}{\Delta t_{\rm obs}} \sim \gamma^2 \nu_0 \sin\theta = \gamma^2 \frac{|q|B}{2\pi mc} \sin\theta, \tag{2.39}$$

Thus the characteristic frequency is boosted by γ^2 relative to the cyclotron frequency ν_0. For an ultra-relativistic electron, $\gamma \simeq E_e/(m_e c^2)$, and hence $\nu_s \propto B E_e^2 \sin\theta$. This relation is extremely useful diagnostically: once B is known or constrained, the frequency of an observed synchrotron photon identifies the energy of the parent electron that produced it. For a pitch angle of 90° the pulses repeat every gyroperiod $T = 2\pi/\Omega_B$. For a general pitch angle the guiding center drifts toward the observer, so successive pulses arrive separated by $T(1-\beta\cos^2\theta) \simeq T\sin^2\theta$, and the single-electron spectrum consists of harmonics of $\Omega_B/\sin^2\theta$.

Single–electron spectrum (full calculation). A rigorous derivation of the synchrotron spectrum can be carried out by treating the motion *locally* as curvature radiation: over the short visibility arc the trajectory has an (almost) constant radius of curvature, and only the *perpendicular* acceleration contributes to the radiation. For a helical orbit one may write the local curvature radius as $\rho \simeq \gamma m_e c^2/(|q|B\sin\theta)$ (equivalently, $\rho = r_g/\sin^2\theta$ if $r_g = v_\perp/\Omega_B$). This immediately sets the *characteristic* (or *critical*) frequency that sets the spectral scale:

$$\nu_c = \frac{3}{2}\gamma^2 \nu_0 \sin\theta = \frac{3}{4\pi}\frac{|q|B}{m_e c}\gamma^2 \sin\theta. \tag{2.40}$$

Numerically, for a GeV electron in a typical interstellar magnetic field the critical frequency indeed falls in the *radio* band:

$$\nu_c \simeq 16\ \mathrm{MHz} \left(\frac{E}{\mathrm{GeV}}\right)^2 \left(\frac{B}{\mu\mathrm{G}}\right) \sin\theta. \tag{2.41}$$

The remarkable outcome of the full Fourier analysis of the Liénard–Wiechert fields is that the *entire* single–electron spectrum can be written in a compact, universal form [6]:

$$\mathcal{P}_\nu(\nu,\gamma,\theta) = \frac{\sqrt{3}\,|q|^3 B}{m_e c^2}\ \sin\theta\ F\left(\frac{\nu}{\nu_c}\right), \qquad F(x) = x\int_x^\infty K_{5/3}(y)\,\mathrm{d}y. \tag{2.42}$$

The explicit $\sin\theta$ simply reminds us that synchrotron emission is powered by the *perpendicular* acceleration: no pitch angle, no synchrotron.

Here $F(x)$ is the *synchrotron kernel* (built from the modified Bessel function $K_{5/3}$), which encodes the spectral *shape*, while the prefactor sets the overall *amplitude* through B and the radiating component of the acceleration ($\propto \sin\theta$). Two useful asymptotic limits capture the familiar appearance of the spectrum,

$$F(x) \simeq \begin{cases} 2.15\, x^{1/3}, & x \equiv \nu/\nu_c \ll 1, \\ \sqrt{\dfrac{\pi x}{2}}\, e^{-x}, & x \gg 1, \end{cases}$$

so that $\mathcal{P}_\nu \propto \nu^{1/3}$ at low frequencies, with an exponential suppression above ν_c. Between these regimes $\mathcal{P}_\nu$ reaches a well–defined maximum at $\nu_{\max} \simeq 0.29\,\nu_c$ (Fig. 4); in practice, one typically works with the dimensionless ratio $x = \nu/\nu_c$, so that the spectral shape is fully specified by the kernel $F(x)$.

In many astrophysical environments the magnetic field direction varies on scales much smaller than the emitting region (turbulent or effectively random fields). In that case it is natural to average over field directions, equivalently over pitch angles. A convenient pitch–angle averaged form of Eq. (2.42) is given in [14]:

$$\left\langle \mathcal{P}_\nu(\nu,\gamma)\right\rangle_\theta = \frac{\sqrt{3}\,|q|^3 B}{m_e c^2}\ G\left(\frac{\nu}{\nu_c}\right), \qquad \nu_c \equiv \frac{3}{4\pi}\,\frac{|q|B}{m_e c}\,\gamma^2, \tag{2.43}$$

with the angle–averaged kernel

$$G(x) = x\int_x^\infty K_{5/3}(y)\,\frac{\sqrt{y^2-x^2}}{y}\,\mathrm{d}y\,, \tag{2.44}$$

and the maximum of the function $G(x)$ is shifted towards smaller values $\nu \simeq 0.23\nu_c$.

Finally, we remark that for back–of–the–envelope estimates it is often convenient to adopt the *delta approximation*, which collapses the spectrum to a single characteristic frequency:

$$\mathcal{P}_\nu(\nu) \simeq P_{\mathrm{s}}\,\delta\big(\nu-\nu_c\big), \tag{2.45}$$

with P_{s} from Eq. (2.30) and ν_c from Eq. (2.40). By construction, this approximation discards the detailed spectral shape; nevertheless, it preserves the correct scalings with B, γ, and θ, and it usually provides a reliable order–of–magnitude estimate of the emissivity.

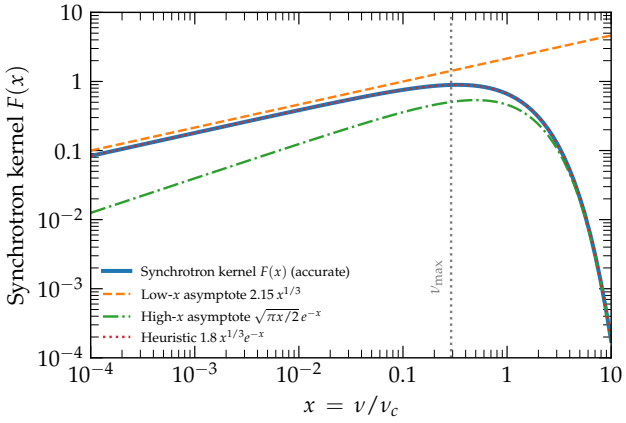


Figure 4: The synchrotron kernel $F(x)$ describing the spectrum of a single electron. The exact function $F(x) = x\int_x^\infty K_{5/3}(y)\,\mathrm{d}y$ (solid) is shown with its low- and high-x asymptotes and the heuristic approximation $1.8\,x^{1/3}e^{-x}$.

Synchrotron* burnoff *limit. A useful consequence of synchrotron scaling is that, in a "one–zone" accelerator where the same electromagnetic fields both accelerate the electron and bend its trajectory, there is a maximum synchrotron photon energy set by radiation reaction. The idea is simple: the electron cannot be energized faster than it loses energy to synchrotron emission.

Approximating the energy gain rate in an electric field of magnitude $\mathcal{E}$ as $\dot{E}_{\rm acc} \simeq e\,\mathcal{E}\,c$, and the synchrotron loss power as in Eq. (2.30)

$$P_{\rm syn} \simeq \frac{2}{3}\,\frac{e^4 B_\perp^2}{m_e^2 c^3}\,\gamma^2, \tag{2.46}$$

the balance $\dot{E}_{\rm acc} = P_{\rm syn}$ defines a radiation–reaction–limited Lorentz factor $\gamma_{\rm rad}$. Substituting $\gamma_{\rm rad}$ into the critical photon energy $\varepsilon_c \equiv h\nu_c = \frac{3}{2}\hbar\gamma^2\frac{eB_\perp}{m_e c}$, Eq. (2.40), yields the remarkable bound

$$\varepsilon_{\rm syn,max} = \frac{9}{4}\,\frac{m_e c^2}{\alpha_{\rm F}}\left(\frac{\mathcal{E}}{B_\perp}\right) \simeq 160\ \mathrm{MeV}\left(\frac{\mathcal{E}}{B_\perp}\right), \tag{2.47}$$

where $\alpha_{\rm F} = e^2/(\hbar c)$ is the fine–structure constant. This value refers to the critical energy; the peak of the single-electron spectral power lies at $0.29\,\varepsilon_c$, so quoted burnoff energies can differ by an order-unity convention. In particular, under ideal MHD conditions one expects $\mathcal{E} \lesssim B_\perp$, and therefore $\varepsilon_{\rm syn,max} \lesssim 160$ MeV in the plasma rest frame.

This scale is intriguingly close to the high–energy cutoff of the *steady* Crab Nebula synchrotron component, while the rare Crab γ-ray flares push the synchrotron emission to several hundred MeV, clearly beyond this limit [15]. Exceeding the ideal–MHD burnoff limit requires decoupling acceleration from radiation (e.g. accelerating where $B_\perp$ is small) and/or non–ideal regions with $\mathcal{E} > B_\perp$, as in magnetic reconnection [16].

2.4 Synchrotron emission by an electron population

In most astrophysical environments the radiating electrons occupy a broad range of energies rather than a single Lorentz factor. A common description of the parent distribution is a power law,

$$N_e(\gamma)\,\mathrm{d}\gamma = K\,\gamma^{-p}\,\mathrm{d}\gamma, \qquad \gamma_{\min} < \gamma < \gamma_{\max},$$

where $N_e(\gamma)$ is the differential electron number density per unit volume and Lorentz factor. We write $N_e(E)$ for the electron energy distribution denoted by $N(E)$ in Section 2.2. Number conservation under the change of variable $E = \gamma m_e c^2$ gives $N_e(\gamma)\,\mathrm{d}\gamma = N_e(E)\,\mathrm{d}E$, hence $N_e(\gamma) = m_e c^2\, N_e(E)$.

Typical non-thermal acceleration mechanisms produce slopes in the range $p \simeq 2$–3.

We denote by j_ν the angle-integrated volume emissivity (power radiated per unit volume and frequency). It follows by integrating the single–electron spectral power over the population and averaging over pitch angles:

$$j_\nu(\nu) = \int_{\gamma_{\min}}^{\gamma_{\max}} \big\langle \mathcal{P}_\nu(\nu,\gamma)\big\rangle_\theta\, N_e(\gamma)\,\mathrm{d}\gamma,$$

where $\langle\cdots\rangle_\theta$ denotes an isotropic pitch–angle average unless stated otherwise. For isotropic emission, the conventional emission coefficient per unit solid angle is $j_\nu/(4\pi)$.

With this definition, j_ν connects directly to observations: for optically thin emission, the measured specific intensity I_ν is the line-of-sight integral of $j_\nu/(4\pi)$.

To expose the dependence on B and ν, we work in the ultra-relativistic limit $\beta \simeq 1$ and approximate each single–electron spectrum by a delta function centered at its characteristic frequency,

$$\big\langle \mathcal{P}_\nu(\nu,\gamma)\big\rangle_\theta \simeq \big\langle P_{\mathrm{s}}(\gamma)\big\rangle\,\delta\big(\nu - \nu_c(\gamma)\big),$$

with $\langle P_{\mathrm{s}}(\gamma)\rangle$ from Eq. (2.31) and $\nu_c(\gamma) = \frac{3}{2}\gamma^2\nu_0$ as in Eq. (2.43). [2] Using the identity

$$\delta\big(\nu - \nu_c(\gamma)\big) = \frac{\delta(\gamma-\gamma_\nu)}{\left|\frac{\mathrm{d}\nu_c}{\mathrm{d}\gamma}\right|_{\gamma_\nu}}, \qquad \gamma_\nu \equiv \sqrt{\frac{2}{3}\frac{\nu}{\nu_0}}, \qquad \frac{\mathrm{d}\nu_c}{\mathrm{d}\gamma} = 3\nu_0\,\gamma,$$

we find

$$j_\nu(\nu) = K\int_{\gamma_{\min}}^{\gamma_{\max}} \big\langle P_{\mathrm{s}}(\gamma)\big\rangle\,\delta\big(\nu-\nu_c(\gamma)\big)\,\gamma^{-p}\,\mathrm{d}\gamma = K\,\frac{4}{9}\,\frac{c\,\sigma_T U_B}{\nu_0}\,\gamma_\nu^{1-p}. \tag{2.48}$$

Since $U_B \propto B^2$, $\nu_0 \propto B$, and $\gamma_\nu \propto (\nu/B)^{1/2}$, this immediately implies the standard optically thin synchrotron scaling [6, 7]:

$$j_\nu \;\propto\; K\,B^{\frac{p+1}{2}}\,\nu^{-\frac{p-1}{2}} \qquad \text{for} \qquad \nu_{\min} < \nu < \nu_{\max}. \tag{2.49}$$

[2] Centering the delta function on ν_c, rather than on the peak of the angle-averaged kernel at $0.23\,\nu_c$, is deliberate. For a power-law population, the electrons that dominate the emission at frequency ν have $\nu/\nu_c \simeq 0.7$, 1.0, and 1.2 for $p = 2$, 2.6, and 3, respectively: the more numerous lower-energy electrons, radiating above their own ν_c, compensate for the decline of the kernel.

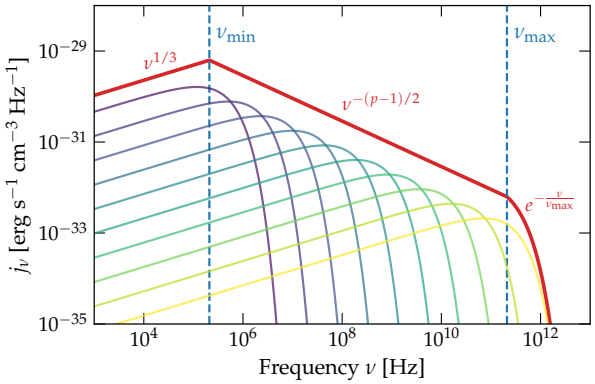


Figure 5: Synchrotron emissivity j_ν produced by a power-law electron population $N_e(\gamma) \propto \gamma^{-p}$ between $\gamma_{\rm min} = 10^2$ and $\gamma_{\rm max} = 10^5$. The dashed vertical lines mark the characteristic frequencies $\nu_{\rm min}$ and $\nu_{\rm max}$ associated with $\gamma_{\rm min}$ and $\gamma_{\rm max}$ for $B = 5\,\mu$G (see Eq. (2.50)). Between $\nu_{\rm min}$ and $\nu_{\rm max}$ the emissivity approaches the standard optically thin slope $j_\nu \propto \nu^{-(p-1)/2}$, while at $\nu \ll \nu_{\rm min}$ it tends to the low-frequency asymptote $j_\nu \propto \nu^{1/3}$. At $\nu \gtrsim \nu_{\rm max}$ the spectrum is exponentially suppressed by the single-electron kernel. The figure also shows normalized single-electron spectra for representative Lorentz factors (log-spaced between $\gamma_{\rm min}$ and $\gamma_{\rm max}$) and weighted according to $N_e(\gamma)$, illustrating how the total emissivity builds up from the superposition of individual contributions.

The corresponding validity band is set by the lowest and highest electron energies present,

$$\nu_{\rm min} \equiv \nu_c(\gamma_{\rm min}) = \frac{3}{2}\,\gamma_{\rm min}^2\,\nu_0, \qquad \nu_{\rm max} \equiv \nu_c(\gamma_{\rm max}) = \frac{3}{2}\,\gamma_{\rm max}^2\,\nu_0, \tag{2.50}$$

(up to order-unity pitch-angle factors if one keeps $\sin\theta$ explicitly).

Provided that the source remains optically thin, below the band the low-frequency tail of the single–electron kernel dominates, yielding

$$j_\nu \propto \nu^{1/3} \qquad (\nu \ll \nu_{\rm min}).$$

At sufficiently low frequencies, synchrotron self-absorption can make the source optically thick and replace this asymptote with a spectral turnover [6]. Above the band, by contrast, the emission is exponentially suppressed by the kernel of the highest-energy electrons available:

$$j_\nu \propto \exp\left(-\frac{\nu}{\nu_{\rm max}}\right) \qquad (\nu \gg \nu_{\rm max}).$$

A schematic example of the resulting broadband spectrum—and of how it emerges from summing the contributions of electrons with different Lorentz factors—is shown in Fig. 5.

From photon slope to electron slope. Introducing the synchrotron spectral index $\alpha \equiv -\mathrm{d}\ln j_\nu / \mathrm{d}\ln\nu$, the scaling in Eq. (2.49) implies the simple and widely used relation

$$\alpha = \frac{p-1}{2} \qquad \Longleftrightarrow \qquad p = 2\alpha + 1. \tag{2.51}$$

This connection is one of the main reasons synchrotron emission is such a powerful diagnostic: in the optically thin regime, a measured radio (or X-ray) slope directly constrains the underlying electron spectrum.

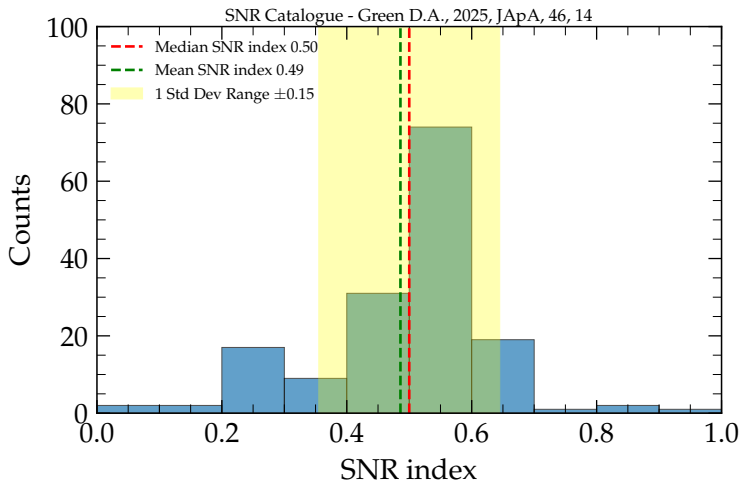


Figure 6: Distribution of the integrated radio spectral index α for Galactic SNRs in the Green catalogue [17], where the flux density scales as $S_\nu \propto \nu^{-\alpha}$. The vertical dashed lines indicate the sample median ($\alpha_{\rm med} = 0.50$) and mean ($\langle\alpha\rangle = 0.49$), while the shaded band marks the 1σ dispersion (± 0.15).

As an example, Fig. 6 shows the distribution of integrated radio spectral indices α for Galactic supernova remnants in the Green catalogue [17]. Interpreting these spectra as optically thin synchrotron emission, each measured α corresponds to a parent electron index $p = 2\alpha + 1 \sim 2$, providing a direct estimate of the emitting-electron slope across the Galactic SNR population.

Finally, note that the amplitude scales as $B^{(p+1)/2}$, rather than B^2, because changing B also shifts the characteristic frequency of each electron. Replacing the delta approximation with the full synchrotron kernel, Eq. (2.42), leaves the B- and ν-scalings unchanged and only multiplies j_ν by a smooth function of p.

2.5 Diffuse Galactic radio emission as a probe of CR electrons

After subtraction of the thermal free–free component, the diffuse Galactic synchrotron intensity (Fig. 7) is often well described over the $\sim$ 1–10 GHz band by a power law in *brightness temperature*,

$$T_b(\nu) \propto \nu^{-\beta_{\rm syn}}, \tag{2.52}$$

with a representative all-sky value $\beta_{\rm syn} \simeq 2.8$ [18]. The index varies across the sky, and several analyses find mild steepening toward higher frequencies. Here T_b is the Rayleigh–Jeans (RJ) *equivalent temperature* corresponding to the specific intensity I_ν, i.e. the temperature a blackbody would need (in the RJ limit) to reproduce the observed I_ν at frequency ν:

$$I_\nu \equiv \frac{2k_{\rm B}\nu^2}{c^2}\, T_b(\nu)\,, \tag{2.53}$$

so that $T_b \propto I_\nu\, \nu^{-2}$.

Assume optically thin synchrotron emission from a parent electron distribution $N_e(\gamma, s) = K(s)\gamma^{-p}$, as in Section 2.4. Here s is distance along the line of sight and $K(s)$ is the local

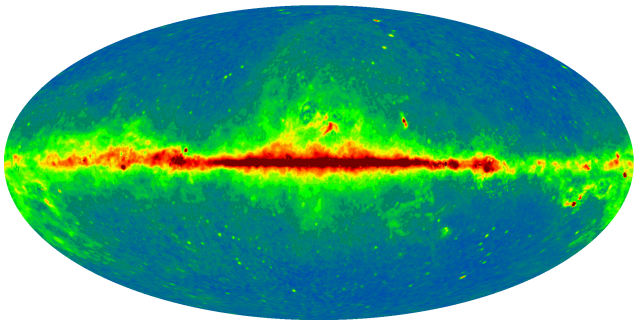


Figure 7: Synchrotron foreground in total intensity reconstructed from the three-year *WMAP* K-band (23 GHz) data using the maximum-entropy method. Adapted from [20].

normalization per unit Lorentz factor. If the slope p does not vary strongly along the line of sight, Eq. (2.49) gives $I_\nu \propto \nu^{-(p-1)/2}$. Therefore

$$T_b(\nu) \propto \nu^{-2}\, I_\nu \propto \nu^{-(p+3)/2} \quad \Longrightarrow \quad p = 2\beta_{\rm syn} - 3 \simeq 2.6\,. \tag{2.54}$$

The characteristic synchrotron frequency, Eq. (2.40), gives a direct mapping between observing band and electron energies: for a typical interstellar field $B \simeq 3\ \mu$G,

$$E_e \simeq 4.6 \left(\frac{\nu}{\rm GHz}\right)^{1/2} \left(\frac{B}{3\ \mu{\rm G}}\right)^{-1/2} {\rm GeV}\,, \tag{2.55}$$

so 1–10 GHz primarily probes electrons with $E_e \sim 5$–15 GeV (up to order-unity pitch-angle factors).

This is a particularly powerful diagnostic because it constrains the *interstellar* lepton spectrum through emission integrated along long lines of sight, and it probes an energy range in which solar modulation can significantly distort the locally measured spectrum [19]. The spectral slope gives a relatively direct estimate of p, but the intensity normalization is degenerate: schematically,

$$I_\nu \propto \nu^{-(p-1)/2} \int {\rm d}s\, K(s)\, B_\perp(s)^{(p+1)/2}.$$

Thus total intensity constrains a line-of-sight weighted combination of the normalization of the electron distribution and the magnetic field rather than either one separately.

In practice, this degeneracy is broken with additional assumptions or data. A common assumption is energy equipartition between cosmic rays and magnetic fields [21]. Alternatively, the radio data can be combined with observables that depend differently on K and B: the γ-ray emission produced by the same electrons through inverse Compton scattering and bremsstrahlung, which does not depend on B, and Faraday rotation and polarization, which trace the magnetic field [22, 23].

2.6 A cooling-break estimate for the Crab Nebula

The integrated synchrotron spectrum of the Crab Nebula provides a useful illustration of both the diagnostic power and the limitations of a cooling-break argument. The calculation

below treats the nebula as a single zone with continuous injection, negligible escape, and synchrotron losses that dominate near the break. Modern spatially resolved data instead favor a multi-zone, time-dependent system, possibly with more than one electron population, so identifying any single feature uniquely as an age-related cooling break is model dependent [24].

The Crab is a pulsar-wind nebula (PWN) associated with the supernova observed in 1054 AD and thus has an age of order $t_{\rm age} \sim 10^3\,\mathrm{yr}$. At a distance of $\simeq 2\,\mathrm{kpc}$, it spans several arcminutes in the radio and optical bands [13, 17]. Its non-thermal emission extends from radio to PeV γ rays. The low-energy hump is synchrotron radiation from relativistic $e^{\pm}$, extending into the MeV band, whereas the high-energy hump is produced predominantly by inverse Compton scattering [15, 25] (see Fig. 8).

For an illustrative one-zone estimate, consider a fiducial break in the optical/UV range,

$$\nu_{\rm br} \sim 10^{15}\,\mathrm{Hz}, \qquad \varepsilon_{\rm br} \equiv h\nu_{\rm br} \sim 4\,\mathrm{eV}\left(\frac{\nu_{\rm br}}{10^{15}\,\mathrm{Hz}}\right), \tag{2.56}$$

near the transition toward the steeper X-ray synchrotron continuum. The integrated Crab spectrum contains several changes of slope, and some multi-component models place the age-related cooling break in the infrared instead [24]. Thus the value above should be read as a fiducial scale chosen to demonstrate the inference, not as a unique decomposition of the Crab spectrum.

Recall that for continuous injection and radiative losses with $b(E) \propto E^2$, Eq. (2.19) gives a cooling-break steepening $\Delta p = 1$. The optically thin synchrotron relation $\alpha = (p-1)/2$ (cf. Eq. (2.51)) therefore implies

$$\Delta\alpha = \frac{\Delta p}{2} = 0.5\,.$$

As a consequence, if a measured feature has this change of slope, the agreement supports a cooling interpretation, although it does not by itself prove one.

In the simplest one-zone interpretation, the break marks the Lorentz factor $\gamma_{\rm br}$ at which the synchrotron cooling time equals the system age (see Eq. (2.17)):

$$\tau_{\rm syn}(\gamma_{\rm br}) \simeq t_{\rm age}. \tag{2.57}$$

Using the characteristic synchrotron frequency (Eq. (2.40)) for $\sin\theta \sim \mathcal{O}(1)$:

$$\nu_{\rm syn} \simeq \frac{3eB}{4\pi m_e c}\,\gamma^2 \simeq 4.2\times10^6\,B_{\rm G}\,\gamma^2\ \mathrm{Hz}, \tag{2.58}$$

where $B_{\rm G} \equiv B/(1\,\mathrm{G})$. Combining this with the pitch-angle-averaged synchrotron loss time in Eq. (2.33),

$$\tau_{\rm syn}(\gamma) = \frac{6\pi m_e c}{\sigma_T B^2\,\gamma}, \tag{2.59}$$

one can eliminate γ in favor of the observed break frequency and solve for the magnetic field. The result can be written as

$$\tau_{\rm syn}(\nu) \simeq \frac{6\pi m_e c}{\sigma_T}\sqrt{\frac{3e}{4\pi m_e c}}\,B^{-3/2}\,\nu^{-1/2}, \tag{2.60}$$

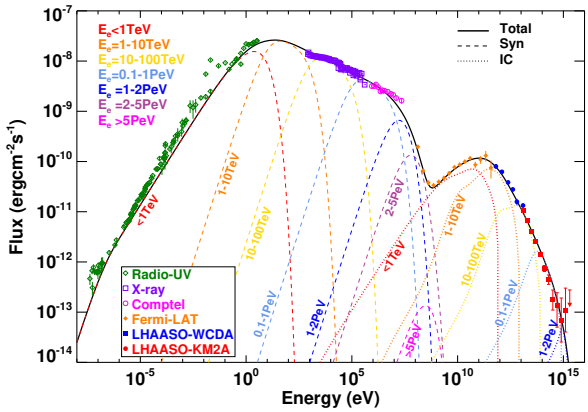


Figure 8: Model decomposition of the Crab Nebula's multi-wavelength spectrum. The curves separate the contributions from electrons in different energy intervals, from the radio synchrotron component to UHE γ rays produced by inverse Compton scattering. Adapted from [25].

and therefore, identifying $\nu = \nu_{\rm br}$ and $\tau_{\rm syn} = t_{\rm age}$,

$$B \simeq 140\,\mu{\rm G}\left(\frac{t_{\rm age}}{1\,{\rm kyr}}\right)^{-2/3}\left(\frac{\nu_{\rm br}}{10^{15}\,{\rm Hz}}\right)^{-1/3}. \tag{2.61}$$

Here $\nu_{\rm syn}$ is evaluated at $\sin\theta = 1$, whereas $\tau_{\rm syn}$ is pitch-angle averaged. Treating both consistently changes B by a few tens of percent at most, well within the accuracy of this one-zone estimate. The corresponding break Lorentz factor is then

$$\gamma_{\rm br} \simeq 1.3\times10^{6}\left(\frac{\nu_{\rm br}}{10^{15}\,{\rm Hz}}\right)^{1/2}\left(\frac{B}{140\,\mu{\rm G}}\right)^{-1/2}, \tag{2.62}$$

corresponding to an electron energy $E_{e,\rm br} \simeq 0.7\,{\rm TeV}$ (as shown also in Fig. 8). The inferred field is also close, at the expected order-unity level, to the value $B \simeq 110\,\mu$G inferred for the PeV-emitting region by LHAASO [25]. Given the multi-zone caveat, this agreement should be read only as an order-of-magnitude consistency check.

This exercise highlights the basic diagnostic power of cooling breaks: once a break frequency is identified and the source age is known, Eq. (2.61) provides an estimate of the effective magnetic field (or, conversely, a magnetic-field estimate can be turned into a radiative age). The weak dependence $B \propto \nu_{\rm br}^{-1/3}$ makes the estimate useful for order-of-magnitude work. In realistic PWNe, however, spatial variations of B, time-dependent injection, escape, additional losses, and multiple emitting zones blur a one-to-one mapping between a single $\nu_{\rm br}$ and a single cooling time.

2.7 Inverse Compton scattering

Inverse Compton (IC) scattering is the transfer of kinetic energy from a relativistic electron to a lower-energy photon through the ordinary Compton process,

$$e + \gamma \longrightarrow e' + \gamma'\,.$$

It is therefore not a different microscopic interaction. The word *inverse* describes the direction of the energy transfer in the astrophysical frame: the photon is up-scattered while the electron cools. In this sense IC is the radiative sibling of synchrotron emission. Both processes drain energy from relativistic electrons, but the target is a photon field in IC scattering and a magnetic field in synchrotron emission.

The cleanest way to understand IC scattering is to keep two frames distinct from the outset. The "lab" frame will denote the frame in which the target radiation field is isotropic, while the electron rest frame (ERF) is the natural frame for the scattering. An incident photon is Doppler-boosted into the ERF, scattered there, and boosted back to the lab. In the exact Compton process the photon loses some energy to recoil in the ERF; in the Thomson limit that loss is negligible. Nevertheless, the two Lorentz transformations can give the photon a large net energy gain in the lab, supplied by the electron.

We develop this picture in three stages. First, we recover the Thomson cross section and its angular distribution in the ERF. Next, we derive the Compton recoil kinematics and combine it with the transformations between the ERF and the lab frame. This reveals how inverse Compton scattering boosts photons and leads to the single-electron IC power in the Thomson limit. Only after establishing this result do we return to the case of appreciable recoil and examine how Klein–Nishina effects modify the cross section, the scattered-photon energy, and the electron cooling rate.

Thomson scattering (classical limit, ERF). We begin in the ERF and suppress primes throughout this paragraph. Consider a monochromatic plane wave with photon energy $\hbar\omega \ll m_e c^2$. Recoil is then negligible. We also assume that the wave is weak enough for the electron's driven motion to remain non-relativistic. In this classical picture the electric field sets the electron into oscillation, and the accelerated charge re-radiates the incident wave.

In Gaussian units, the Poynting flux of the incident wave is

$$\mathbf{S} = \frac{c}{4\pi}\,\mathbf{E}\times\mathbf{B} \quad\Longrightarrow\quad S \equiv |\mathbf{S}| = \frac{c}{4\pi}|\mathbf{E}|^2,$$

where $|\mathbf{E}| = |\mathbf{B}|$ for a plane wave in vacuum. Writing

$$\mathbf{E} = E_0\,\boldsymbol{\epsilon}\,\sin(\omega t + \phi)\,,$$

the unit vector $\boldsymbol{\epsilon}$ specifies the linear polarization and is transverse to the propagation direction.

The Lorentz force acting on the electron is

$$\mathbf{F} = q\left(\mathbf{E} + \frac{\mathbf{v}}{c}\times\mathbf{B}\right) \simeq q\mathbf{E},$$

because the magnetic force is smaller than the electric force by a factor of order $|\mathbf{v}|/c$. The corresponding acceleration is

$$\mathbf{a} = \frac{\mathbf{F}}{m_e} = \frac{q}{m_e}\mathbf{E},$$

and its squared magnitude, averaged over one wave period, is

$$\langle a^2 \rangle = \left(\frac{q}{m_e}\right)^2 \langle E^2 \rangle = \left(\frac{q}{m_e}\right)^2 \frac{E_0^2}{2} .$$

The Larmor formula Eq. (2.27) then gives the time-averaged radiated power,

$$\langle P \rangle = \frac{2}{3} \frac{q^2 \langle a^2 \rangle}{c^3} = \frac{1}{3} \frac{q^4}{m_e^2 c^3} E_0^2 .$$

The total scattering cross section is the ratio between the radiated power and the incident flux. Using the time-averaged Poynting flux

$$\langle S \rangle = \frac{c}{8\pi} E_0^2 ,$$

we obtain

$$\langle P \rangle = \sigma_{\mathrm{T}} \langle S \rangle \quad \Longrightarrow \quad \sigma_{\mathrm{T}} = \frac{\langle P \rangle}{\langle S \rangle} = \frac{8\pi}{3} \left(\frac{q^2}{m_e c^2}\right)^2 \equiv \frac{8\pi}{3} r_e^2 , \qquad r_e \equiv \frac{q^2}{m_e c^2} . \tag{2.63}$$

Here r_e is the classical electron radius, and σ_{T} is the Thomson cross section. Numerically,

$$\sigma_{\mathrm{T}} \simeq 6.652 \times 10^{-25} \ \mathrm{cm}^2 .$$

The total cross section fixes how much power is scattered, but not where that power goes. The angular dependence follows from the dipole radiation pattern. In most astrophysical applications the incident photons are treated as *unpolarized*: the incident radiation is an incoherent, equal mixture of two orthogonal transverse polarizations, and the polarization of the outgoing photon is not measured. Averaging over the incident polarization and summing over the final one gives [26]

$$\frac{\mathrm{d}\sigma_{\mathrm{T}}}{\mathrm{d}\Omega_f} = \frac{r_e^2}{2} \left(1 + \cos^2\theta\right) = \frac{3\sigma_{\mathrm{T}}}{16\pi} \left(1 + \cos^2\theta\right) ,$$

where θ is the angle between the incident and outgoing photon directions. This distribution is not isotropic: scattering is strongest in the forward and backward directions and is smaller by a factor of two at $\theta = \pi/2$. Consistently, its angular integral recovers the total Thomson cross section,

$$\int d\Omega_f \, \frac{\mathrm{d}\sigma_{\mathrm{T}}}{\mathrm{d}\Omega_f} = \sigma_{\mathrm{T}} .$$

Equally important for what follows, the pattern is invariant under $\hat{\mathbf{n}}_f \to -\hat{\mathbf{n}}_f$. Its first angular moment therefore vanishes,

$$\frac{1}{\sigma_{\mathrm{T}}} \int d\Omega_f \, \hat{\mathbf{n}}_f \frac{\mathrm{d}\sigma_{\mathrm{T}}}{\mathrm{d}\Omega_f} = 0 .$$

Thus the scattered photons have no mean outgoing momentum in the ERF, even though their angular distribution is not isotropic. This symmetry will be the key angular result in the derivation of the IC power.

Finally, consider an arbitrary photon bath rather than a single plane wave, and restore primes for ERF quantities. Let dn'_γ denote an element of the incident photon number density, resolved in energy and propagation direction. Since the electron is at rest, the corresponding scattering rate in the Thomson limit is

$$d\dot{N}'_{\rm sc} = \sigma_{\rm T} c \, dn'_\gamma \, .$$

Each scattered photon carries an outgoing energy $\varepsilon'_f \simeq \varepsilon'_i$, so the total outgoing power is

$$P'_{\rm out} = \int \varepsilon'_f \, d\dot{N}'_{\rm sc} \simeq \sigma_{\rm T} c \int \varepsilon'_i \, dn'_\gamma = \sigma_{\rm T} c \, U'_{\rm rad} \, .$$

The integral includes all incident photon energies and directions, and $U'_{\rm rad}$ is their total energy density. The directional contributions add positively even in an isotropic bath, for which the net radiation flux vanishes.

The incoming photons participating in these scatterings carry the power

$$P'_{\rm in} = \sigma_{\rm T} c U'_{\rm rad} \quad \Longrightarrow \quad P'_{\rm out} - P'_{\rm in} \simeq 0 \, .$$

To this accuracy, scattering redirects the photons without changing their total energy in the instantaneous ERF. An *anisotropic* radiation field may still transfer net momentum to the electron, whereas for an isotropic field the momentum transfers also cancel on average. The non-zero IC energy loss appears in the lab frame, where the electron is moving.

Compton scattering formula (ERF kinematics). The classical Thomson treatment assumes that the electron recoil is negligible. To quantify this assumption, we now treat the electron and photon relativistically and impose energy–momentum conservation. No scattering amplitude is needed at this stage: the energy shift is fixed entirely by kinematics.

Let P_i^μ and P_f^μ be the initial and final electron four-momenta, and K_i^μ and K_f^μ the corresponding photon four-momenta. We temporarily set $c = 1$ and use the metric $(+,-,-,-)$. Conservation of four-momentum gives

$$P_i + K_i = P_f + K_f \, .$$

The electron remains on its mass shell, $P_i^2 = P_f^2 = m_e^2$, while the photon four-momenta are null, $K_i^2 = K_f^2 = 0$. Eliminating the unobserved final electron momentum, $P_f = P_i + K_i - K_f$, therefore gives

$$\begin{aligned} P_f^2 &= (P_i + K_i - K_f)^2 \\ &= m_e^2 + 2P_i \cdot K_i - 2P_i \cdot K_f - 2K_i \cdot K_f = m_e^2 \, . \end{aligned}$$

We now evaluate these Lorentz scalars in the ERF. The initial electron is at rest, $P_i = (m_e, \mathbf{0})$. Choosing the x-axis along the incident photon direction and placing the scattered photon in the x–y plane,

$$K_i = \varepsilon'_i (1, 1, 0, 0) \, , \qquad K_f = \varepsilon'_f (1, \cos\theta, \sin\theta, 0) \, .$$

Here θ is the angle between the incident and outgoing photon directions, the same scattering angle used in the Thomson angular distribution. The required scalar products are

$$P_i \cdot K_i = m_e \varepsilon_i' , \qquad P_i \cdot K_f = m_e \varepsilon_f' , \qquad K_i \cdot K_f = \varepsilon_i' \varepsilon_f' (1 - \cos\theta) .$$

Substitution into the mass-shell condition yields

$$m_e(\varepsilon_i' - \varepsilon_f') = \varepsilon_i' \varepsilon_f' (1 - \cos\theta) .$$

Solving for the outgoing photon energy and restoring c, we obtain the Compton formula

$$\varepsilon_f' = \frac{\varepsilon_i'}{1 + \dfrac{\varepsilon_i'}{m_e c^2}(1 - \cos\theta)} . \tag{2.64}$$

The denominator is never smaller than unity. A photon scattered in the forward direction ($\theta = 0$) keeps its energy, whereas at any non-zero angle it loses energy to electron recoil. The largest loss occurs for back-scattering, $\theta = \pi$. The fractional energy transferred to the electron is

$$\frac{\varepsilon_i' - \varepsilon_f'}{\varepsilon_i'} = \frac{\dfrac{\varepsilon_i'}{m_e c^2}(1 - \cos\theta)}{1 + \dfrac{\varepsilon_i'}{m_e c^2}(1 - \cos\theta)} \simeq \frac{\varepsilon_i'}{m_e c^2}(1 - \cos\theta) ,$$

where the last expression applies when $\varepsilon_i'(1 - \cos\theta) \ll m_e c^2$. For a given scattering angle, this is the condition for negligible recoil. The Thomson regime requires recoil to remain negligible over the full angular distribution. Since $0 \le 1 - \cos\theta \le 2$, the corresponding angle-independent criterion is $\varepsilon_i' \ll m_e c^2$. The photon is then redirected in the ERF while its energy remains nearly unchanged, $\varepsilon_f' \simeq \varepsilon_i'$.

From Compton to* inverse *Compton: the role of frames. At first sight, the word *inverse* may seem puzzling. The Compton formula was derived in the electron rest frame (ERF), where recoil can only reduce the photon energy, $\varepsilon_f' \le \varepsilon_i'$, with equality for forward scattering. The resolution is that photon energy is frame dependent. In astrophysical applications the natural lab frame is usually the one in which the target radiation is isotropic, such as the CMB rest frame, and the electron moves through it with Lorentz factor γ. The same scattering can then increase the photon energy at the expense of the electron's kinetic energy.

Let ε_i be the incident photon energy in the lab frame, and let the electron move with velocity βc. We denote by μ_i the cosine of the angle between the incident photon and the electron velocity, both measured in the lab frame. Boosting to the ERF gives

$$\varepsilon_i' = \gamma \varepsilon_i (1 - \beta \mu_i) . \tag{2.65}$$

Head-on photons have $\mu_i = -1$ and therefore receive the largest boost into the ERF, whereas photons traveling alongside the electron have $\mu_i = 1$ and are de-boosted.

After scattering, let μ'_f be the cosine of the angle between the outgoing photon and the electron velocity, measured in the ERF. The inverse boost gives

$$\varepsilon_f = \gamma\varepsilon'_f(1+\beta\mu'_f)\,. \tag{2.66}$$

Combining these two transformations with the ERF Compton formula, Eq. (2.64), yields the exact kinematic relation

$$\frac{\varepsilon_f}{\varepsilon_i} = \frac{\gamma^2(1-\beta\mu_i)(1+\beta\mu'_f)}{1+\dfrac{\gamma\varepsilon_i}{m_e c^2}(1-\beta\mu_i)(1-\cos\theta)}\,. \tag{2.67}$$

Here θ retains its previous meaning: it is the angle between the incident and outgoing photon directions in the ERF. This equation separates the two ingredients cleanly. The factors in the numerator are Doppler boosts, while the denominator accounts for electron recoil.

In the Thomson limit the recoil term is negligible, so Eq. (2.67) gives

$$\varepsilon_f \simeq \gamma^2\varepsilon_i(1-\beta\mu_i)(1+\beta\mu'_f)\,. \tag{2.68}$$

The largest gain occurs for a head-on photon, $\mu_i=-1$, that is redirected along the electron velocity, $\mu'_f=1$:

$$\varepsilon_f^{\max} \simeq \gamma^2(1+\beta)^2\varepsilon_i \simeq 4\gamma^2\varepsilon_i \qquad (\gamma\gg 1)\,.$$

This limiting geometry, however, is not that of a typical scattering. To obtain the mean photon energy and the associated electron energy-loss rate, we must combine the kinematics with the rate of encounters in the radiation field.

Single-particle power radiated in IC scattering. We retain the lab frame and notation introduced above, and assume that all relevant target photons satisfy the Thomson condition $\varepsilon'_i \ll m_e c^2$. Let $n_\gamma(\varepsilon_i)$ be the photon number density per unit energy. For an isotropic photon bath, the number density in the intervals $d\varepsilon_i$ and $d\Omega$ is

$$dn_\gamma = \frac{n_\gamma(\varepsilon_i)}{4\pi}\,d\varepsilon_i\,d\Omega\,, \qquad dU = \varepsilon_i\,dn_\gamma\,,$$

with $U_{\rm rad} = \int dU$. The rate at which the electron scatters photons from this element is

$$d\dot{N}_{\rm sc} = \sigma_{\rm T} c(1-\beta\mu_i)\,dn_\gamma\,.$$

The factor $1-\beta\mu_i$ is the relative-flux factor. It enhances head-on encounters ($\mu_i=-1$) and suppresses photons that overtake the electron from behind ($\mu_i=1$). Equivalently, it follows by computing the rate $\sigma_{\rm T} c\,dn'_\gamma$ in the ERF, where $dn'_\gamma = \gamma(1-\beta\mu_i)dn_\gamma$, and then using $dt=\gamma dt'$.

Integrating over incident directions gives the scattering rate per unit target-photon energy,

$$\begin{aligned}\frac{\mathrm{d}\dot{N}_{\rm sc}}{\mathrm{d}\varepsilon_i} &= \sigma_{\rm T} c\,n_\gamma(\varepsilon_i)\int_\Omega \frac{d\Omega}{4\pi}\,(1-\beta\mu_i)\\ &= \frac{\sigma_{\rm T} c\,n_\gamma(\varepsilon_i)}{2}\int_{-1}^{1}(1-\beta\mu_i)\,d\mu_i = \sigma_{\rm T} c\,n_\gamma(\varepsilon_i)\,,\end{aligned}$$

where the second equality follows after integration over the azimuthal angle.

We now use the photon kinematics derived above. At fixed incident direction μ_i, the symmetry of the Thomson angular distribution gives $\langle \mu_f' \rangle = 0$, so Eq. (2.68) gives a mean outgoing energy $\gamma^2 \varepsilon_i (1 - \beta \mu_i)$. Averaging over incident directions with the normalized encounter-rate weight then yields

$$\begin{aligned} \langle \varepsilon_f \rangle &= \frac{\gamma^2 \varepsilon_i}{2} \int_{-1}^{1} (1 - \beta \mu_i)^2 \, d\mu_i \\ &= \gamma^2 \left(1 + \frac{\beta^2}{3} \right) \varepsilon_i \simeq \frac{4}{3} \gamma^2 \varepsilon_i \qquad (\gamma \gg 1) \,. \end{aligned} \tag{2.69}$$

The two factors of $1 - \beta \mu_i$ in the integral come from the photon-energy boost and the encounter rate, respectively.

This result applies separately to each incident photon energy. Multiplying by the corresponding scattering rate and integrating over the target spectrum gives the total outgoing photon power,

$$\begin{aligned} P_{\rm out} &= \int_0^\infty \langle \varepsilon_f \rangle \, \frac{\mathrm{d}\dot{N}_{\rm sc}}{\mathrm{d}\varepsilon_i} \, d\varepsilon_i \\ &= \sigma_{\rm T} c \, \gamma^2 \left(1 + \frac{\beta^2}{3} \right) \int_0^\infty \varepsilon_i n_\gamma(\varepsilon_i) \, d\varepsilon_i \\ &= \sigma_{\rm T} c \, \gamma^2 \left(1 + \frac{\beta^2}{3} \right) U_{\rm rad} \,. \end{aligned}$$

The factor multiplying $\sigma_{\rm T} c$ is precisely the radiation energy density measured in the ERF, $U'_{\rm rad} = \gamma^2 (1 + \beta^2/3) U_{\rm rad}$.

The outgoing power includes the energy initially carried by the scattered photons. To obtain the power supplied by the electron, we must therefore subtract the incident photon power,

$$P_{\rm in} = \int_0^\infty \varepsilon_i \, \frac{\mathrm{d}\dot{N}_{\rm sc}}{\mathrm{d}\varepsilon_i} \, d\varepsilon_i = \sigma_{\rm T} c \, U_{\rm rad} \,.$$

Energy conservation then identifies the net gain of the radiation field with the electron energy-loss rate:

$$\begin{aligned} -\frac{\mathrm{d}E_e}{\mathrm{d}t} &= P_{\rm IC} = P_{\rm out} - P_{\rm in} \\ &= \sigma_{\rm T} c \left[\gamma^2 \left(1 + \frac{\beta^2}{3} \right) - 1 \right] U_{\rm rad} = \frac{4}{3} \gamma^2 \beta^2 \, \sigma_{\rm T} c \, U_{\rm rad} \,. \end{aligned} \tag{2.70}$$

In the last step we used $\gamma^2 - 1 = \gamma^2 \beta^2$. The limit $\beta \to 0$ provides an immediate check: a stationary electron only redirects photons in the Thomson regime, so $P_{\rm out} = P_{\rm in}$ and there is no net energy transfer.

For a relativistic electron, Eq. (2.70) has the same structure as the pitch-angle-averaged synchrotron result, with $U_{\rm B}$ replaced by $U_{\rm rad}$. It also shows that, as long as the Thomson

condition holds, the total IC power depends on the target radiation only through its total energy density, not through the detailed photon spectrum.

Klein–Nishina regime. The Thomson approximation used in Eq. (2.70) rests on one kinematic assumption: the incident photon energy must be much smaller than the electron rest energy in the ERF. A photon that is soft in the lab can violate this condition after being Doppler-boosted into the rest frame of a sufficiently energetic electron. For an incident direction μ_i, the relevant parameter is

$$x \equiv \frac{\varepsilon_i'}{m_e c^2} = \frac{\gamma \varepsilon_i}{m_e c^2}(1 - \beta \mu_i)\,. \tag{2.71}$$

Head-on encounters dominate the energy transfer and have $x \simeq 2\gamma\varepsilon_i/(m_e c^2)$ for $\gamma \gg 1$. For an isotropic target field it is conventional to introduce

$$b \equiv \frac{4\gamma \varepsilon_i}{m_e c^2}\,. \tag{2.72}$$

The numerical factor is conventional and arises naturally in the angle-averaged IC kernel. The two parameters encode the same physical transition: $b \ll 1$ is the Thomson regime, while $b \gtrsim 1$ signals the onset of Klein–Nishina (KN) effects for photons of energy ε_i. Thus, for a photon field concentrated near $\varepsilon_\star$, the transition occurs around

$$E_{e,\mathrm{KN}} \sim \frac{(m_e c^2)^2}{4\varepsilon_\star}\,. \tag{2.73}$$

For a broad photon spectrum there is no single transition energy: each part of the target field enters the KN regime at a different electron energy.

Once x is no longer small, three changes occur together. Electron recoil makes the collision inelastic in the ERF, the scattering probability becomes energy dependent and smaller than its Thomson value, and the characteristic scattered-photon energy no longer grows as $\gamma^2\varepsilon_i$. Consequently, the IC loss rate depends on the shape of the target spectrum, not only on its total energy density.

For unpolarized radiation, the differential Klein–Nishina cross section is

$$\frac{\mathrm{d}\sigma_{\mathrm{KN}}}{\mathrm{d}\Omega_f'} = \frac{r_e^2}{2}\left(\frac{\varepsilon_f'}{\varepsilon_i'}\right)^2\left(\frac{\varepsilon_f'}{\varepsilon_i'} + \frac{\varepsilon_i'}{\varepsilon_f'} - \sin^2\theta\right), \tag{2.74}$$

where $\varepsilon_f'/\varepsilon_i'$ is given by Eq. (2.64). In the limit $x \to 0$, this expression reduces to the unpolarized Thomson angular distribution. At larger x, large-angle scatterings are increasingly suppressed and the emission becomes more strongly concentrated in the forward direction.

Integrating over the final direction gives the total cross section (see e.g. [26]),

$$\sigma_{\mathrm{KN}} = \frac{3}{4}\sigma_{\mathrm{T}}\left[\frac{1+x}{x^3}\left(\frac{2x(1+x)}{1+2x} - \ln(1+2x)\right) + \frac{1}{2x}\ln(1+2x) - \frac{1+3x}{(1+2x)^2}\right]. \tag{2.75}$$

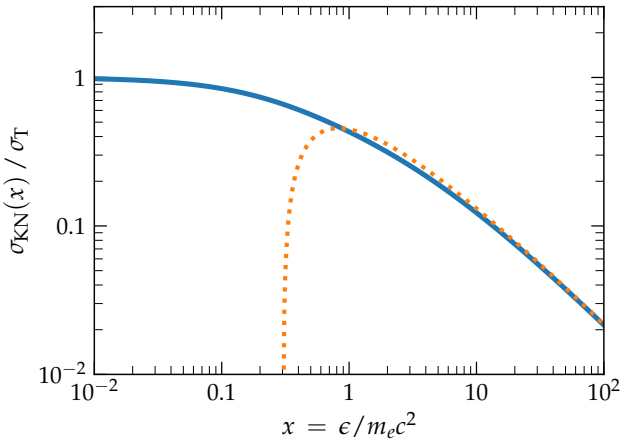


Figure 9: Klein–Nishina cross section as a function of the ERF photon energy $x = \varepsilon_i'/(m_e c^2)$, Eq. (2.75). For $x \ll 1$ the Thomson limit is recovered, while for $x \gg 1$ the cross section is strongly suppressed. The dashed orange line shows the high-energy asymptotic form, Eq. (2.77).

Expanding for soft photons ($x \ll 1$) gives

$$\sigma_{KN} \simeq \sigma_T \, (1 - 2x + \dots) \,, \tag{2.76}$$

so that $\sigma_{KN} \to \sigma_T$ as $x \to 0$. In the opposite, deep KN limit ($x \gg 1$),

$$\sigma_{KN} \simeq \frac{3}{8} \sigma_T \frac{1}{x} \left(\ln 2x + \frac{1}{2} \right) , \tag{2.77}$$

showing explicitly that the effective cross section decreases with increasing ERF photon energy, roughly as $1/x$ (up to the logarithmic correction), in agreement with Fig. 9.

The energy of the scattered photon changes at the same transition. For the geometry that maximizes the lab-frame photon energy, the exact kinematic relation Eq. (2.67) becomes, for $\gamma \gg 1$,

$$\varepsilon_f^{max} \simeq \frac{4\gamma^2 \varepsilon_i}{1 + 4\gamma \varepsilon_i / (m_e c^2)} = \frac{4\gamma^2 \varepsilon_i}{1 + b} \,. \tag{2.78}$$

This formula displays both limits:

$$\varepsilon_f^{max} \simeq \begin{cases} 4\gamma^2 \varepsilon_i, & b \ll 1, \\ \gamma m_e c^2, & b \gg 1. \end{cases} \tag{2.79}$$

The Thomson γ^2 growth therefore cannot continue indefinitely. In the deep KN regime a single scattered photon can carry an order-unity fraction of the electron energy, so IC cooling need not be a sequence of many small energy losses.

The reduction of the cooling power begins smoothly before the deep KN limit is reached. For an ultra-relativistic electron in an isotropic photon field, expansion of the full loss rate to first order beyond Thomson gives [26]

$$P_{IC} = \frac{4}{3} \gamma^2 \beta^2 \sigma_T c \, U_{rad} \left[1 - \frac{63}{10} \frac{\gamma}{m_e c^2} \frac{\langle \varepsilon_i^2 \rangle}{\langle \varepsilon_i \rangle} + \dots \right] , \tag{2.80}$$

where the moments refer to the lab-frame photon number distribution,

$$\langle \varepsilon_i^k \rangle \equiv \frac{\int_0^\infty \varepsilon_i^k n_\gamma(\varepsilon_i)\, d\varepsilon_i}{\int_0^\infty n_\gamma(\varepsilon_i)\, d\varepsilon_i}\,.$$

The expansion is valid only when the correction in brackets is much smaller than unity. It already shows why $U_{\rm rad}$ alone is no longer sufficient: higher-energy target photons are suppressed first, so two radiation fields with the same energy density but different spectra produce different IC losses.

When $b \gg 1$ across the relevant target spectrum, the extreme-KN asymptotic loss rate for an ultra-relativistic electron is [27]

$$P_{\rm IC}^{\rm KN} \simeq \frac{3}{8}\,\sigma_{\rm T} c (m_e c^2)^2 \int_0^\infty \frac{n_\gamma(\varepsilon_i)}{\varepsilon_i}\left[\ln\left(\frac{4\gamma\varepsilon_i}{m_e c^2}\right) - \frac{11}{6}\right] d\varepsilon_i\,. \tag{2.81}$$

For a monochromatic target, this power grows only logarithmically with γ, in sharp contrast with the Thomson scaling $P_{\rm IC} \propto \gamma^2$. The cooling time is therefore longer than a Thomson extrapolation would predict. Because a single collision may also remove a substantial fraction of the electron energy, a Monte Carlo approach that follows individual scatterings can be preferable to a continuous-loss treatment in the deep KN regime.

The practical consequences are closely linked: the IC spectrum turns over relative to the Thomson prediction, its highest-energy photons approach the electron energy, and the electron distribution can retain more high-energy particles because their cooling is less efficient. For broad or thermal target fields, accurate calculations use the full KN kernel or dedicated approximations rather than one universal suppression factor [28, 29].

2.8 Inverse Compton emission by an electron population

Having derived the power and characteristic photon-energy boost for a single electron, we can now ask how a population of electrons maps into an inverse Compton (IC) spectrum. In the Thomson regime, the calculation closely parallels the synchrotron case developed in Section 2.4.

Consider an isotropic power-law electron distribution,

$$n_e(\gamma)\,{\rm d}\gamma = K\,\gamma^{-p}\,{\rm d}\gamma, \qquad \gamma_{\rm min} < \gamma < \gamma_{\rm max}, \tag{2.82}$$

where $n_e(\gamma)$ is the differential number density per unit Lorentz factor. For clarity, we approximate the target radiation as a monochromatic, isotropic photon bath with lab-frame photon energy ε_i and total energy density $U_{\rm rad}$. We also assume that all relevant electrons scatter in the Thomson regime,

$$b \equiv \frac{4\gamma\varepsilon_i}{m_e c^2} \ll 1\,.$$

For an ultra-relativistic electron, the angle-averaged IC power is

$$P_{\rm IC}(\gamma) \simeq \frac{4}{3}\,\sigma_{\rm T} c\,\gamma^2 U_{\rm rad}, \tag{2.83}$$

and the mean energy of the scattered photons is

$$\langle \varepsilon_f \rangle \simeq \frac{4}{3}\gamma^2 \varepsilon_i \,. \tag{2.84}$$

These are the single-particle results of Eqs. (2.69) and (2.70). In frequency space, the corresponding characteristic IC frequency is

$$\nu_{\rm IC}(\gamma) \equiv \frac{\langle \varepsilon_f \rangle}{h} \simeq \frac{4}{3}\gamma^2 \nu_i, \qquad \nu_i \equiv \frac{\varepsilon_i}{h} \,. \tag{2.85}$$

As in the synchrotron derivation, we first replace the broad single-electron spectrum by a delta function centered on its characteristic frequency,

$$P_\nu^{\rm IC}(\gamma, \nu) \simeq P_{\rm IC}(\gamma)\, \delta\big(\nu - \nu_{\rm IC}(\gamma)\big). \tag{2.86}$$

Again, as in Section 2.4, we denote by $j_\nu^{\rm IC}$ the angle-integrated volume emissivity, namely the power radiated per unit volume and frequency. For isotropic emission, the conventional emission coefficient per unit solid angle is $j_\nu^{\rm IC}/(4\pi)$. The population emissivity is therefore

$$j_\nu^{\rm IC} = \int_{\gamma_{\rm min}}^{\gamma_{\rm max}} n_e(\gamma) P_\nu^{\rm IC}(\gamma, \nu)\, {\rm d}\gamma \,. \tag{2.87}$$

To evaluate the integral, write

$$\nu_{\rm IC}(\gamma) = A\gamma^2, \qquad A \equiv \frac{4}{3}\nu_i, \qquad \gamma_\nu = \left(\frac{\nu}{A}\right)^{1/2}.$$

The delta-function identity then gives

$$\delta\big(\nu - A\gamma^2\big) = \frac{\delta(\gamma - \gamma_\nu)}{2A\gamma_\nu} \,.$$

Substituting this result and Eq. (2.83) into Eq. (2.87), we obtain

$$j_\nu^{\rm IC} = \frac{K\sigma_{\rm T} c U_{\rm rad}}{2\nu_i}\, \gamma_\nu^{1-p} \propto K U_{\rm rad} \nu_i^{(p-3)/2} \nu^{-(p-1)/2}. \tag{2.88}$$

Thus, for a fixed seed field and with $\alpha_{\rm IC} \equiv -{\rm d}\ln j_\nu^{\rm IC}/{\rm d}\ln\nu$, the central result is

$$\boxed{\; j_\nu^{\rm IC} \propto \nu^{-(p-1)/2}, \qquad \alpha_{\rm IC} = \frac{p-1}{2}, \;} \tag{2.89}$$

valid between the characteristic frequencies associated with $\gamma_{\rm min}$ and $\gamma_{\rm max}$. The full IC kernel broadens the spectral boundaries, but it preserves this power-law index in the Thomson regime. Synchrotron and IC emission therefore inherit the same spectral slope from the same power-law electron population.

The parallel extends to the total power. For the same electron, and after averaging the synchrotron power over an isotropic pitch-angle distribution,

$$\frac{P_{\rm IC}}{\langle P_{\rm s} \rangle} = \frac{U_{\rm rad}}{U_{\rm B}} \,. \tag{2.90}$$

This ratio then applies to the optically thin bolometric luminosities of a common electron population, provided all relevant IC scatterings remain in the Thomson regime. This makes the synchrotron-to-IC ratio a powerful probe of $U_{\rm B}$ or $U_{\rm rad}$.

2.9 Lepton cooling in the Milky Way

A classic application of radiative losses is the cooling of electrons and positrons propagating through the Milky Way. Above a few tens of GeV, synchrotron emission in the Galactic magnetic field and inverse Compton (IC) scattering on the interstellar radiation field (ISRF) are typically the dominant energy-loss channels. Their competition provides a direct link between the local environment and the spectrum of Galactic cosmic-ray leptons.

If all IC scatterings occur in the Thomson regime, the synchrotron and IC powers have the same dependence on electron energy and can be added directly:

$$P_{\rm rad}^{\rm Th} = \frac{4}{3}\sigma_{\rm T} c \gamma^2 \left(U_{\rm B} + U_{\rm rad}\right), \qquad (\gamma \gg 1).$$

Using $\tau_{\rm loss} = E_e / P_{\rm rad}$, this gives

$$\tau_{\rm loss}^{\rm Th}(E_e) \simeq 0.3\ {\rm Myr} \left(\frac{U_{\rm B} + U_{\rm rad}}{1\ {\rm eV\,cm^{-3}}}\right)^{-1} \left(\frac{E_e}{1\ {\rm TeV}}\right)^{-1}. \tag{2.91}$$

Thus a TeV electron cools on a sub-Myr timescale for energy densities typical of the Solar neighborhood. The simple $\tau_{\rm loss} \propto E_e^{-1}$ scaling will provide our Thomson benchmark.

Magnetic field. The field entering the synchrotron loss rate is the total rms field sampled by the particles, including both ordered and turbulent components. A volume-averaged value $B \simeq 1$–$3\ \mu$G is a useful benchmark for Galactic propagation, although the total field in the disk can be larger and substantial spatial variations are expected. For reference, $B = 3\ \mu$G corresponds to $U_{\rm B} \simeq 0.23\ {\rm eV\,cm^{-3}}$. More generally, the magnetic energy density scales as

$$U_{\rm B} = \frac{B^2}{8\pi} \simeq 0.025\ {\rm eV\,cm^{-3}} \left(\frac{B}{1\ \mu{\rm G}}\right)^2, \tag{2.92}$$

Interstellar radiation field. The spatially uniform CMB provides an unavoidable target with $U_{\rm CMB} \simeq 0.26\ {\rm eV\,cm^{-3}}$. In the Galactic disk it is supplemented by direct stellar emission and by infrared radiation from dust. A convenient approximation represents the local ISRF as a sum of graybody components, each specified by a temperature T and an energy density $U_{{\rm rad},j}$. Typical components span the infrared ($k_{\rm B}T \sim 10^{-3}$–10^{-2} eV), optical ($k_{\rm B}T \sim 10^{-1}$ eV), and ultraviolet ($k_{\rm B}T \gtrsim 1$ eV) bands. The illustrative decomposition in Fig. 10 has a total radiation density of order $1\ {\rm eV\,cm^{-3}}$, comparable to or larger than the magnetic energy density above.

When does Klein–Nishina suppression matter? The Thomson estimate must be reconsidered when an incident photon is no longer soft in the electron rest frame. Following Eq. (2.72), define

$$b \equiv \frac{4\gamma\varepsilon_i}{m_e c^2}. \tag{2.93}$$

The crossover begins around $b \sim 1$, corresponding to

$$E_e^{\rm KN} \sim \frac{(m_e c^2)^2}{4\varepsilon_i} \simeq 65\ {\rm GeV} \left(\frac{\varepsilon_i}{1\ {\rm eV}}\right)^{-1}. \tag{2.94}$$

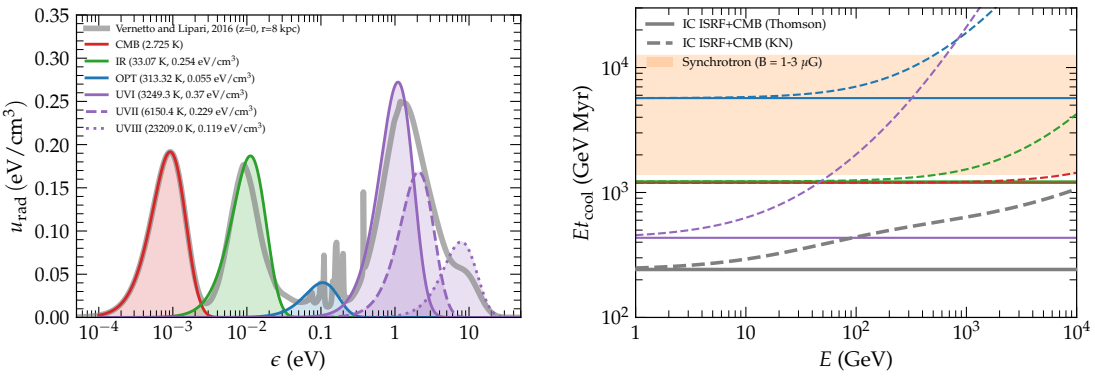


Figure 10: **Left:** Local Milky Way ISRF spectral energy density. The thick gray curve is the model of Vernetto and Lipari [30], evaluated at $z = 0$ and $r = 8$ kpc; the colored curves show an illustrative decomposition into diluted thermal components representing the CMB, dust emission, and stellar light. **Right:** Electron IC cooling on the corresponding components. Solid curves show the Thomson approximation and dashed curves include KN suppression; the colors identify the same target components as in the left panel. Thick gray curves give the total IC contribution, while the orange band shows synchrotron losses for $B = 1$–3 μG. The plotted quantity is $E_e\tau_{\rm cool}$, which is constant for the Thomson scaling $\tau_{\rm cool} \propto E_e^{-1}$; an upward turn therefore makes KN suppression directly visible.

This is not a sharp threshold: the transition is spread over energy by the angular and spectral distributions of the target photons. It nevertheless reveals the hierarchy immediately. IC losses on optical/UV photons become suppressed at electron energies of order tens of GeV, infrared targets enter the KN regime around multi-TeV energies, and the CMB, with $\varepsilon_i \simeq 6.3 \times 10^{-4}$ eV, remains close to the Thomson limit up to $E_e \sim 10^5$ GeV.

A useful KN-corrected parametrization. For quantitative work, one can retain the simplicity of the Thomson formula while accounting accurately for the smooth KN transition. The parametrization summarized below is due to Khangulyan, Aharonian, and Kelner [28]; in particular, the interpolation functions and all numerical coefficients quoted here are theirs, and are not derived in these notes.

Consider an isotropic diluted blackbody with temperature T and energy density

$$U_{\rm rad} = \kappa a T^4, \qquad 0 < \kappa \leq 1,$$

where κ is the dilution factor. Following the notation of Khangulyan et al., the IC loss power of an ultra-relativistic electron is written as

$$P_{\rm IC}(\gamma) = P_{\rm Th}(\gamma) f_{\rm KN}(u), \qquad P_{\rm Th}(\gamma) \equiv \frac{4}{3}\sigma_{\rm T} c \gamma^2 U_{\rm rad}, \qquad u \equiv \frac{4\gamma k_{\rm B} T}{m_e c^2}. \tag{2.95}$$

The dimensionless variable u is the thermal analogue of the monochromatic parameter b. The normalization is chosen so that $f_{\rm KN} \to 1$ as $u \to 0$.

For a blackbody spectral shape, the suppression factor can be written as

$$f_{\rm KN}^{\rm (BB)}(u) \equiv \frac{F_{\rm iso}(u)}{0.722\, u^2}. \tag{2.96}$$

Their compact interpolation, accurate to about 5%, is

$$F_{\rm iso}(u) \simeq G^{(0)}_{\rm iso}(u), \qquad G^{(0)}_{\rm iso}(u) = \frac{c_{\rm iso} u \ln(1 + 0.722u/c_{\rm iso})}{1 + c_{\rm iso} u / 0.822}, \qquad c_{\rm iso} = 4.62. \tag{2.97}$$

A more accurate approximation, with errors of about 1%, can be found in the same paper [28].

For an ISRF represented by several graybody components, their loss powers simply add:

$$P^{\rm tot}_{\rm IC}(\gamma) = \sum_j \frac{4}{3} \sigma_{\rm T} c \gamma^2 U_{{\rm rad},j} f_{\rm KN}(u_j), \qquad u_j = \frac{4\gamma k_{\rm B} T_j}{m_e c^2}. \tag{2.98}$$

The physical picture is now transparent. At low energy all components are in the Thomson regime, so only their summed energy density matters and $E_e \tau_{\rm cool}$ is constant. As the electron energy rises, the hottest stellar components are suppressed first; infrared and CMB photons continue to cool the electrons efficiently to higher energies. The IC loss rate therefore does not switch off at one universal energy, but changes smoothly as the relative weight of the target components evolves.

For the illustrative ISRF in Fig. 10, KN effects increase the total IC cooling time at $E_e = 1$ TeV by a factor of about 2.6 relative to a Thomson extrapolation. Even small departures from the Thomson cooling law, however, can leave measurable imprints in the propagated lepton spectrum [31]. For example, KN effects have been invoked to explain the hardening observed by AMS-02 in the cosmic-ray electron spectrum around 42 GeV [32, 33], although whether KN losses alone account for this feature depends on the adopted ISRF [29]. This example illustrates the importance of including KN effects when predicting spectral features, even when only some components of the radiation field enter the KN regime.

2.10 Pair production and γ-ray absorption

Very-high-energy γ rays emitted by distant sources (AGN, GRBs, star-forming galaxies) do not propagate through empty space. Instead, they traverse an almost ubiquitous bath of low-energy photons: the CMB and the extragalactic background light (EBL), i.e. the accumulated starlight and its dust-reprocessed infrared emission. Therefore, its spectrum has two broad peaks, in the optical at $\varepsilon \sim 1$ eV and in the infrared at $\varepsilon \sim 10^{-2}$ eV (see Fig. 11). Even though these backgrounds are extremely dilute by terrestrial standards, over cosmological distances they can efficiently absorb γ rays.

The basic reaction is photon–photon pair production,

$$\gamma + \gamma \;\rightarrow\; e^+ + e^-,$$

which removes photons from the line of sight. The observable consequence is an *attenuation* of the primary γ-ray spectrum, leading to the associated notion of a *γ-ray horizon*.

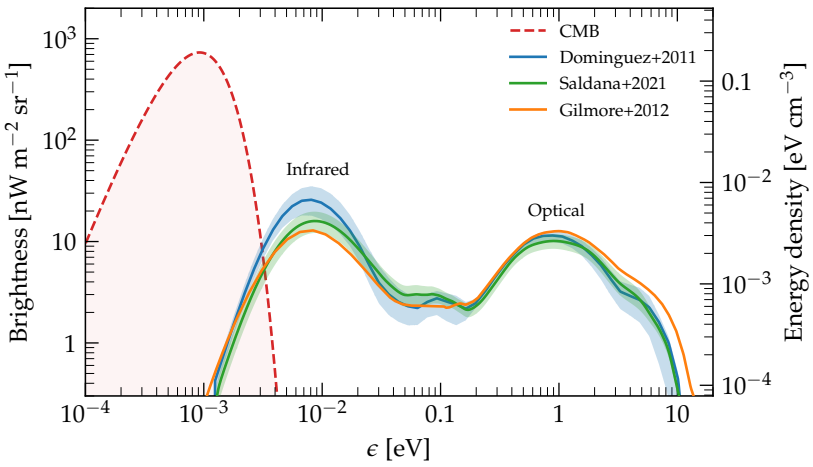


Figure 11: Photon field energy density of the CMB and the extragalactic background light (EBL) as a function of photon energy. For the EBL we show three recent models from the literature [34–36].

Threshold condition and characteristic target energies. The kinematics are conveniently expressed in terms of the Lorentz-invariant squared center-of-mass (CoM) energy. Writing four-momenta in energy units, $P_a = (E_a, c\mathbf{p}_a)$, gives

$$s = (P_a + P_b)^2 = m_a^2c^4 + m_b^2c^4 + 2\big(E_aE_b - c^2\mathbf{p}_a \cdot \mathbf{p}_b\big).$$

For two photons ($m_a = m_b = 0$, $cp = E$) of energies E_γ and ε and collision angle θ,

$$s = 2\,E_\gamma\,\varepsilon\,(1 - \cos\theta)\,.$$

Pair creation becomes possible once the CoM energy can supply the combined rest energy of the two leptons,

$$s_{\rm th} = 4m_e^2c^4 \qquad \Longrightarrow \qquad 2E_\gamma\varepsilon(1-\cos\theta) \geq 4m_e^2c^4. \tag{2.99}$$

The lowest threshold is reached in a head-on collision ($\theta = \pi$, hence $1 - \cos\theta = 2$), for which

$$E_\gamma\,\varepsilon \;\geq\; (m_ec^2)^2 \qquad \Longleftrightarrow \qquad E_\gamma \;\geq\; \frac{(m_ec^2)^2}{\varepsilon}\,. \tag{2.100}$$

Numerically,

$$E_\gamma \;\gtrsim\; 0.26\ \text{TeV}\left(\frac{\varepsilon}{\text{eV}}\right)^{-1},$$

which sets the minimum target-photon energy for a given γ ray. For instance,

- a 100 GeV photon requires $\varepsilon \geq \varepsilon_{\rm th} \simeq 2.6$ eV, in the optical/UV range;
- a 10 TeV photon requires $\varepsilon \geq \varepsilon_{\rm th} \simeq 2.6 \times 10^{-2}$ eV, in the infrared.

The cross section vanishes at threshold and peaks somewhat above it, so efficient absorption generally samples target energies a few times these minima. Other collision angles raise the threshold; a nearly isotropic radiation field therefore requires an angular average.

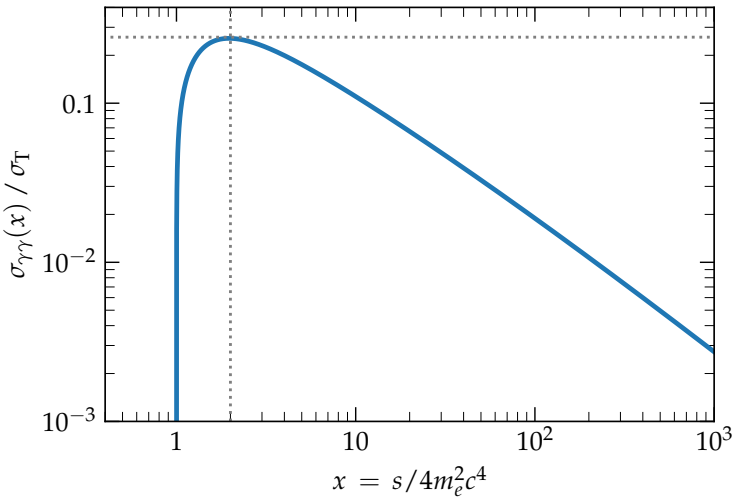


Figure 12: Photon–photon pair-production cross section in units of σ_T as a function of $x = s/(4m_e^2c^4) = \gamma^{*2}$. The cross section rises above threshold ($x = 1$), peaks at $\simeq 0.26\,\sigma_T$ for $x \simeq 2$ ($\beta^* \simeq 0.70$), and then decreases at large x.

Cross section in the center-of-mass frame. It is useful to parametrize the reaction by the dimensionless speed $\beta^* \equiv v_e^*/c$ of either lepton in the CoM frame. The star labels quantities in this frame, where the pair is produced back-to-back. Each lepton has energy

$$E_e^* = \gamma^* m_e c^2, \qquad \gamma^* = \frac{1}{\sqrt{1-\beta^{*2}}},$$

and the total CoM energy satisfies

$$\sqrt{s} = 2E_e^* = 2\gamma^* m_e c^2 \qquad \Longrightarrow \qquad \gamma^* = \frac{\sqrt{s}}{2m_e c^2}.$$

Equivalently,

$$\beta^* = \sqrt{1-\frac{1}{\gamma^{*2}}} = \sqrt{1-\frac{4m_e^2c^4}{s}} = \sqrt{1-\frac{2m_e^2c^4}{E_\gamma \varepsilon(1-\cos\theta)}},$$

which makes the threshold behavior transparent: at threshold $\beta^* \to 0$, while far above threshold $\beta^* \to 1$.

The exact cross section for unpolarized photons (the Breit–Wheeler cross section) can be written in terms of β^* [37]:

$$\sigma_{\gamma\gamma}(\beta^*) = \frac{3}{16}\,\sigma_T\,(1-\beta^{*2})\left[2\beta^*(\beta^{*2}-2) + (3-\beta^{*4})\ln\left(\frac{1+\beta^*}{1-\beta^*}\right)\right], \tag{2.101}$$

where σ_T is the Thomson cross section. Below threshold, $s < 4m_e^2c^4$, the cross section is zero.

The main limits are (see Fig. 12):

- *Near threshold* ($\beta^* \ll 1$), $\sigma_{\gamma\gamma} \to 0$ and then rises steeply as soon as pair creation becomes kinematically allowed.

- The cross section peaks at a value of order

$$\sigma_{\gamma\gamma}^{\max} \simeq 0.25\,\sigma_{\mathrm{T}},$$

  at $\beta^* \simeq 0.70$, corresponding to $s \simeq 2s_{\mathrm{th}}$ or $\sqrt{s} \simeq 1.4\sqrt{s_{\mathrm{th}}}$.

- *Far above threshold* ($s \gg 4m_e^2c^4$, $\beta^* \to 1$), a standard asymptotic form is

$$\begin{aligned}\sigma_{\gamma\gamma} &\simeq \frac{3}{8}\frac{\sigma_{\mathrm{T}}}{\gamma^{*2}}\left[\ln\left(4\gamma^{*2}\right) - 1\right] \\ &= \frac{3}{2}\,\sigma_{\mathrm{T}}\frac{m_e^2c^4}{s}\left[\ln\left(\frac{s}{m_e^2c^4}\right) - 1\right].\end{aligned} \tag{2.102}$$

Thus the high-energy cross section falls as an inverse power of s, up to a logarithmic factor. Physically, while any target photons above threshold *can* contribute, the cross section is largest when s is only modestly above $4m_e^2c^4$. For nearly head-on collisions, this means that $E_\gamma\varepsilon$ is within a factor of a few of $(m_ec^2)^2$. This makes distant TeV spectra a useful probe of the EBL: each γ-ray energy preferentially samples a range of target energies, and the attenuation encodes the intervening photon field. To connect this cross section to observations, we now integrate the interaction probability along the line of sight and over the target spectrum and directions.

2.11 Optical depth and the γ-ray horizon

Absorption along the line of sight is described by the optical depth $\tau_{\gamma\gamma}$, which determines the absorption probability,

$$P_{\mathrm{abs}}(E_\gamma) = 1 - \exp\left[-\tau_{\gamma\gamma}(E_\gamma)\right].$$

The optical depth accumulates the local absorption coefficient along the path. This coefficient is obtained by integrating over target-photon directions and energies, giving

$$\tau_{\gamma\gamma}(E_\gamma) = \int_0^R \mathrm{d}x \int \mathrm{d}\Omega\,(1-\cos\theta) \int_{\varepsilon_{\mathrm{th}}}^{\infty} \mathrm{d}\varepsilon\; n_\gamma(\varepsilon,\Omega,x)\,\sigma_{\gamma\gamma}\big(s(E_\gamma,\varepsilon,\cos\theta)\big), \tag{2.103}$$

where x parametrizes the path length and $n_\gamma(\varepsilon,\Omega,x)$ is the background-photon number density per unit energy and solid angle, allowing for a background that may be anisotropic. The flux factor $(1-\cos\theta)$ reflects the relative velocity of the two photons and the corresponding dependence of the CoM energy. The threshold energy is set locally by Eq. (2.99),

$$\varepsilon_{\mathrm{th}}(E_\gamma,\mu) = \frac{2m_e^2c^4}{E_\gamma(1-\mu)}, \qquad \mu \equiv \cos\theta. \tag{2.104}$$

For an *isotropic* background field, we define the angle-integrated differential number density as

$$n_\gamma(\varepsilon,x) \equiv \int \mathrm{d}\Omega\; n_\gamma(\varepsilon,\Omega,x).$$

Choosing the γ-ray direction as the polar axis gives $\mathrm{d}\Omega = \mathrm{d}\phi\;\mathrm{d}\mu$, with $\mu = \cos\theta$. Substituting into Eq. (2.103) and integrating over the azimuth ϕ therefore yields

$$\tau_{\gamma\gamma}(E_\gamma) = \int_0^R \mathrm{d}x \int_{-1}^{1} \mathrm{d}\mu \, \frac{1-\mu}{2} \int_{\varepsilon_{\mathrm{th}}(E_\gamma,\mu)}^{\infty} \mathrm{d}\varepsilon \, n_\gamma(\varepsilon, x) \, \sigma_{\gamma\gamma}(s), \tag{2.105}$$

where $s = 2E_\gamma\varepsilon(1-\mu)$ as before. The factor $(1-\mu)/2$ thus combines the relative-flux factor with the normalized angular average.

At each energy, the fraction of primary photons that survive propagation is $1 - P_{\mathrm{abs}}(E_\gamma)$. If $F_{\mathrm{int}}(E_\gamma)$ denotes the flux that would reach Earth in the absence of absorption, the observed primary flux is therefore

$$\begin{aligned} F_{\mathrm{obs}}(E_\gamma) &= [1 - P_{\mathrm{abs}}(E_\gamma)]\, F_{\mathrm{int}}(E_\gamma) \\ &= F_{\mathrm{int}}(E_\gamma)\, \exp\left[-\tau_{\gamma\gamma}(E_\gamma)\right]. \end{aligned}$$

The exponential factor is the survival probability of the primary photons.

In cosmological applications the same idea holds, but the path length is naturally expressed in terms of redshift z. Since $\frac{\mathrm{d}z}{\mathrm{d}t} = -(1+z)H(z)$, the proper path-length element along the photon trajectory is

$$\mathrm{d}x = c\left|\frac{\mathrm{d}t}{\mathrm{d}z}\right| \mathrm{d}z = \frac{c\ \mathrm{d}z}{(1+z)H(z)}.$$

A photon observed today with energy E_0 had energy $E_0(1+z)$ at redshift z. Using local target-photon energies and physical (proper) densities at each redshift gives

$$\begin{aligned} \tau_{\gamma\gamma}(E_0, z_s) = \int_0^{z_s} \mathrm{d}z \, \frac{c}{(1+z)H(z)} \int_{-1}^{1} \mathrm{d}\mu \, \frac{1-\mu}{2} \\ \times \int_{\varepsilon_{\mathrm{th}}(E_0,z,\mu)}^{\infty} \mathrm{d}\varepsilon \, n_\gamma(\varepsilon, z)\, \sigma_{\gamma\gamma}\Big(2E_0(1+z)\,\varepsilon\,(1-\mu)\Big), \end{aligned} \tag{2.106}$$

where z_s is the source redshift and $H(z)$ the Hubble expansion rate. The local threshold entering the energy integral is

$$\varepsilon_{\mathrm{th}}(E_0, z, \mu) = \frac{2m_e^2c^4}{E_0(1+z)(1-\mu)}\,. \tag{2.107}$$

One then defines the *γ-ray horizon* $E_{\gamma,\mathrm{hor}}(z)$ implicitly by $\tau_{\gamma\gamma}(E_{\gamma,\mathrm{hor}}, z) = 1$. The observed flux is strongly suppressed wherever $\tau_{\gamma\gamma} \gg 1$. The opacity need not increase monotonically with energy, as the CMB example below illustrates. Different EBL models predict different $\tau_{\gamma\gamma}(E_\gamma, z)$ curves; comparing these predictions to AGN spectra therefore yields direct *constraints* on the EBL, as illustrated in Fig. 13.

A back-of-the-envelope mean free path. It is instructive to estimate the scale at which $\tau_{\gamma\gamma} \sim 1$ without doing the full integral. Since τ is essentially (rate)$\times$(time), one may write

$$\tau_{\gamma\gamma} \sim \sigma_{\gamma\gamma}\, n_\gamma\, L,$$

where n_γ is a characteristic target photon density, $\sigma_{\gamma\gamma}$ an effective cross section, and L a propagation distance.

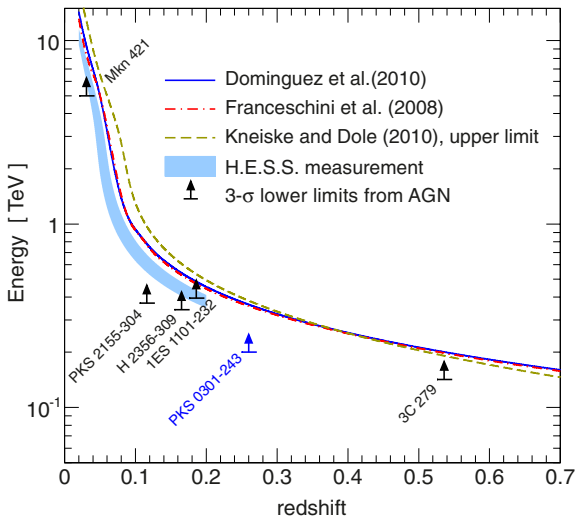


Figure 13: γ-ray horizon $E_{\gamma,\mathrm{hor}}(z)$ for different EBL models, defined by $\tau_{\gamma\gamma}(E_{\gamma,\mathrm{hor}}, z) = 1$. Lower limits derived from AGN spectra are shown for comparison. Adapted from [38].

Consider a $\sim$TeV γ ray. The head-on *threshold* target energy is $\varepsilon_{\mathrm{th}} \simeq 0.26\ \mathrm{eV}\,(E_\gamma/\mathrm{TeV})^{-1}$, i.e. infrared photons. Because $\sigma_{\gamma\gamma}$ peaks when $E_\gamma\varepsilon$ is a few times threshold, the interaction is typically most efficient for $\varepsilon \sim 0.5$–1 eV, broadly in the optical/near-IR band.

From Fig. 11, the EBL energy density in a logarithmic interval of width $\Delta \ln \varepsilon \sim 1$ around $\varepsilon \sim 1$ eV is of order

$$u_{\mathrm{EBL}} \sim 3 \times 10^{-3}\ \mathrm{eV\,cm^{-3}}.$$

Approximating the photon energies within this interval by ε gives the corresponding photon number density,

$$n_\gamma \sim \frac{u_{\mathrm{EBL}}}{\varepsilon} \sim 3 \times 10^{-3}\ \mathrm{cm^{-3}} \qquad (\varepsilon \sim 1\ \mathrm{eV}).$$

The pair-production cross section reaches

$$\sigma_{\gamma\gamma}^{\mathrm{max}} \sim 0.25\,\sigma_{\mathrm{T}} \sim 1.7 \times 10^{-25}\ \mathrm{cm^2},$$

so taking $\sigma_{\gamma\gamma} \sim \sigma_{\gamma\gamma}^{\mathrm{max}}$ as an optimistic estimate gives

$$L \sim \frac{1}{\sigma_{\gamma\gamma}\, n_\gamma} \sim 2 \times 10^{27}\ \mathrm{cm} \sim 0.6\ \mathrm{Gpc}.$$

The broad target spectrum and the angular average reduce the effective cross section below the peak value used here, shifting this estimate by factors of a few. The estimate nevertheless shows that the optical depth of TeV photons becomes order unity over distances of order a gigaparsec. For standard parameters of a flat ΛCDM cosmology, 0.6 Gpc corresponds roughly to $z \simeq 0.15$. The TeV emission of sources at $z \gg 0.1$ is therefore expected to be strongly attenuated by $\gamma\gamma \to e^+e^-$ interactions on the EBL. This energy-dependent absorption produces an apparent high-energy cutoff in the observed γ-ray spectrum.

The PeV horizon on the CMB. The same estimate applied to the CMB gives a horizon inside the Milky Way. Using the values introduced earlier, $n_{\rm CMB} \simeq 411\ \mathrm{cm}^{-3}$ and $\langle\varepsilon\rangle \simeq 6.3\times10^{-4}$ eV, the head-on threshold for photons of this mean energy lies at

$$E_\gamma \simeq \frac{m_e^2c^4}{\langle\varepsilon\rangle} \simeq 400\ \mathrm{TeV}.$$

The broad CMB spectrum smooths the onset of absorption around this scale. With $\sigma_{\gamma\gamma} \sim \sigma_{\gamma\gamma}^{\max}$ and the full CMB density,

$$L \sim \frac{1}{\sigma_{\gamma\gamma}^{\max}\, n_{\rm CMB}} \sim 1.4\times10^{22}\ \mathrm{cm} \sim 5\ \mathrm{kpc}.$$

Evaluating Eq. (2.105) with the Planck spectrum lengthens this estimate somewhat. The absorption length reaches its minimum at about 2 PeV [30],

$$\lambda_{\gamma\gamma}^{\rm CMB} \simeq 7\ \mathrm{kpc} \qquad (E_\gamma \simeq 2\ \mathrm{PeV}), \tag{2.108}$$

and grows again at higher energies, to about 11 kpc at 10 PeV and 50 kpc at 100 PeV, as the cross section decreases.

Below $\sim$ 200 TeV, CMB absorption is negligible over typical Galactic distances: the absorption length is already $\sim$ 400 kpc at 200 TeV and increases rapidly at lower energies. Between 1 and 5 PeV the optical depth to the Galactic Center (8.5 kpc) is $\tau_{\gamma\gamma} \simeq 1$. Nearby sources are only mildly absorbed: the Crab Nebula, detected up to 1.1 PeV [25], lies at 2 kpc, where $\tau_{\gamma\gamma} \simeq 0.2$ at 1 PeV. At a few PeV, sources in other galaxies are strongly attenuated: from the Large Magellanic Cloud, at 50 kpc, $\tau_{\gamma\gamma} \simeq 7$ at 2 PeV. Around these energies, γ-ray astronomy is therefore effectively Galactic. The interstellar radiation field of starlight and dust adds a smaller absorption around 100 TeV, mainly toward the inner Galaxy [30].

2.12 Electromagnetic cascades

Pair production does not destroy the energy of an absorbed γ ray; it changes the particle that carries it. The newly created electron and positron can transfer that energy back to photons by inverse Compton (IC) scattering,

$$\gamma + \gamma_{\rm b} \to e^+ + e^-, \qquad e^\pm + \gamma_{\rm b} \to e^\pm + \gamma.$$

If the IC photon is itself above the pair-production edge, it is absorbed and the cycle starts again. If it is below the edge, it escapes. The repeated sequence

$$\gamma \longrightarrow e^\pm \longrightarrow \gamma \longrightarrow \cdots$$

is an *electromagnetic cascade*.

The aim of this section is to understand two features of a fully developed cascade: the energy at which photons begin to escape, and the approximate shape of their spectrum. We first use one monochromatic background to expose the kinematics. This is deliberately a toy model. In particular, we assume a static homogeneous medium, a propagation

distance long enough that photons in the opaque range are absorbed, IC-dominated and complete electron cooling, and a characteristic-energy approximation for the IC photons. The CMB–EBL application, where two different backgrounds play different roles, follows in Section 2.13.

The distinction between a *threshold* and an *escape edge* is important. For a monochromatic background of photon energy ε_b, define

$$x \equiv \frac{E_\gamma \varepsilon_b}{m_e^2 c^4} = \frac{E_\gamma}{E_{\rm th}}, \qquad E_{\rm th} \equiv \frac{m_e^2 c^4}{\varepsilon_b}. \tag{2.109}$$

For a head-on collision, pair production is kinematically allowed when $x \geq 1$; see Eq. (2.100). This condition alone does not make the medium opaque: the cross section vanishes at threshold, and absorption also depends on the target density and propagation distance. The physical escape edge is set by $\tau_{\gamma\gamma}(E) \sim 1$. In the present sharp-threshold, optically thick approximation we identify it, to order-unity accuracy, with

$$E_c \sim E_{\rm th}. \tag{2.110}$$

Thus E_c should be read as the characteristic boundary between absorbed and escaping photons, not as a precise universal cutoff.

Three stages of the energy flow. It is helpful to follow the cascade from high to low energy before discussing its spectrum.

1. *Leading-particle stage.* Far above threshold ($x \gg 1$), pair production commonly shares the photon energy asymmetrically. If f is the fraction carried by the softer lepton, its high-energy distribution is approximately logarithmic,

$$\frac{d\sigma}{df} \propto \frac{1}{f}, \qquad f_{\rm min} = \frac{1-\sqrt{1-1/x}}{2} \simeq \frac{1}{4x}, \qquad f \leq \frac{1}{2}, \tag{2.111}$$

   where $f_{\rm min}$ is written for head-on collisions, consistent with Eq. (2.109); oblique collisions raise it by an order-unity factor. One lepton therefore retains most of the energy and, in the Klein–Nishina regime, can regenerate a photon of comparable energy. The energy then descends slowly through the leading chain, while softer secondaries feed the rest of the cascade.

2. *Multiplication stage.* As the interaction energy approaches the pair-production region, splittings become less asymmetric and the number of particles grows rapidly. Pair production and IC scattering redistribute the injected energy over many electrons, positrons, and photons. This is the stage associated with the approximately scale-invariant part of the escaping spectrum.

3. *Terminal cooling stage.* Once IC photons fall below E_c, they are no longer efficiently reabsorbed and escape. The remaining electrons continue to cool and populate progressively lower photon energies. This last stage produces the robust low-energy asymptote derived below.

The same kinematics bounds the energy at which pairs are injected. Multiplying the soft-lepton limit in Eq. (2.111) by the parent energy gives

$$E_{e,\mathrm{soft}}^{\mathrm{min}} = f_{\mathrm{min}} E_\gamma = \frac{x\left(1-\sqrt{1-1/x}\right)}{2} E_{\mathrm{th}},$$

which decreases monotonically from $E_{\mathrm{th}}/2$ at threshold to $E_{\mathrm{th}}/4$ for $x \gg 1$. Whatever the parent energy, no lepton is injected below $m_e^2c^4/(4\varepsilon_{\mathrm{b}})$. The *characteristic* injection energy of the last generation is instead set by the photons absorbed just above the edge. Near threshold the parent energy is shared approximately symmetrically, giving $E_e \sim E_c/2$. Asymmetric splittings also populate $E_c/4 \lesssim E_e \lesssim E_c/2$, and collision angles and realistic target spectra smear this boundary further. Thus $E_c/2$ is a characteristic lower injection scale, not a strict boundary. Since the IC photon energy scales as E_e^2 in the Thomson regime, the corresponding spectral break is spread over up to a factor of ~ 4 in photon energy.

In the Thomson limit, an electron of energy E_e emits IC photons with characteristic energy

$$E_\gamma(E_e) \simeq \frac{4}{3}\left(\frac{E_e}{m_e c^2}\right)^2 \varepsilon_{\mathrm{IC}}, \tag{2.112}$$

where $\varepsilon_{\mathrm{IC}}$ is the characteristic energy of the photons that dominate IC cooling. We therefore define

$$E_X \equiv E_\gamma\left(E_e = \frac{E_c}{2}\right). \tag{2.113}$$

Photons below E_X are made by electrons cooling below the pair-injection scale; photons between E_X and E_c retain information about the multiplication stage.

For the one-background toy model, $\varepsilon_{\mathrm{IC}} = \varepsilon_{\mathrm{b}}$ and $E_c \sim m_e^2c^4/\varepsilon_{\mathrm{b}}$, so, formally,

$$E_X \sim \frac{E_c}{3}. \tag{2.114}$$

The putative interval $E_X < E_\gamma < E_c$ then spans less than one half-decade, and Eq. (2.112) does not really apply at its lower boundary. An electron of energy $E_c/2$ has $\gamma = m_e c^2/(2\varepsilon_{\mathrm{b}})$, so the Klein–Nishina parameter of Eq. (2.72) is

$$b = \frac{4\gamma\varepsilon_{\mathrm{b}}}{m_e c^2} = 2,$$

and the Thomson estimate would hand the photon $E_X/E_e = 2/3$ of the electron energy. More generally, every pair-produced lepton has $E_e \geq m_e^2c^4/(4\varepsilon_{\mathrm{b}})$ and hence $b \geq 1$: in a single field, freshly injected pairs never scatter in the Thomson regime on the photons that created them. Their energy losses are discrete rather than continuous, and the mapping $E_\gamma \propto E_e^2$ holds only for electrons that have already cooled well below the injection scale. A single monochromatic field therefore teaches the spectral logic but does not produce a broad intermediate power law. The CMB–EBL hierarchy will instead give $E_X \ll E_c$, with the last pairs scattering deep in the Thomson regime.

Mapping the electron flow into escaping photons. Let $Q_e(E_e)dE_e$ be the total number of electrons and positrons injected by pair production in the interval dE_e, integrated over the history of the cascade and per injected primary. The number that subsequently cool through energy E_e is

$$\Phi_e(E_e) \equiv \int_{E_e}^{\infty} Q_e(E'_e)\, dE'_e. \tag{2.115}$$

This definition makes the bookkeeping transparent: every particle injected above E_e crosses that energy once if cooling is monotonic and complete.

While the particles cross an interval dE_e, they radiate an energy of order dE_e. In the characteristic-energy approximation the corresponding number of photons is

$$dN_\gamma \simeq \frac{dE_e}{E_\gamma(E_e)}. \tag{2.116}$$

If $Y_\gamma(E_\gamma) \equiv dN_\gamma/dE_\gamma$ denotes the escaping photon yield per primary, then

$$Y_\gamma(E_\gamma) \simeq \Phi_e(E_e)\frac{dE_e}{dE_\gamma}\frac{1}{E_\gamma}, \qquad E_e = E_e(E_\gamma). \tag{2.117}$$

In the Thomson regime, $E_\gamma \propto E_e^2$, and hence

$$Y_\gamma(E_\gamma) \propto \Phi_e(E_e(E_\gamma))\, E_\gamma^{-3/2}. \tag{2.118}$$

This relation isolates the well-controlled IC mapping from the less certain electron flow during multiplication.

Low-energy branch: a cooling result. For $E_e \lesssim E_c/2$, direct pair injection is small in the sharp-edge approximation, and it vanishes identically below $E_{\rm th}/4$; we therefore set $Q_e(E_e) \simeq 0$. Equation (2.115) then gives

$$\Phi_e(E_e) \simeq \Phi_0, \qquad E_e \lesssim \frac{E_c}{2}. \tag{2.119}$$

Every injected lepton crosses every lower energy once. Substitution into Eq. (2.118) yields

$$Y_\gamma(E_\gamma) \propto E_\gamma^{-3/2}, \qquad E_\gamma \lesssim E_X. \tag{2.120}$$

Within the assumptions of complete Thomson cooling and the characteristic-energy approximation, this is a genuine analytic result. The full IC kernel smooths the boundaries but preserves the low-energy $E_\gamma^{-3/2}$ asymptote.

Intermediate branch: a plausibility argument. The next step does *not* have the same status. For $E_e \gtrsim E_c/2$, fresh pairs are injected while the cascade is multiplying, and Φ_e is determined by the coupled pair-production and IC transport equations. Energy conservation alone does not fix its energy dependence.

A useful Heitler-like closure is to suppose that a developed multiplication stage is approximately self-similar. In this picture the cascade places comparable injected energy in each logarithmic electron interval,

$$E_e^2 Q_e(E_e) \simeq \text{const}. \tag{2.121}$$

If this scaling extends over a sufficiently broad interval, then $Q_e \propto E_e^{-2}$ and

$$\Phi_e(E_e) = \int_{E_e}^{\infty} Q_e(E'_e)\, dE'_e \propto E_e^{-1}. \tag{2.122}$$

Using $E_e \propto E_\gamma^{1/2}$ in Eq. (2.118) gives, conditionally,

$$Y_\gamma(E_\gamma) \propto E_\gamma^{-2}, \qquad E_X \lesssim E_\gamma \lesssim E_c. \tag{2.123}$$

Equivalently, $E_\gamma^2 Y_\gamma$ is constant: the escaping photons carry comparable energy per logarithmic interval.

This chain of reasoning makes the assumption visible. Approximate scale invariance makes Eq. (2.121) plausible; it does not derive it. Solving the full kinetic problem a posteriori validates an almost flat $E_\gamma^2 Y_\gamma$ segment, but numerical calculations typically find a slightly harder spectrum, $Y_\gamma \propto E_\gamma^{-1.9}$, together with smooth rather than sharp transitions [39]. We retain E^{-2} as a transparent analytic benchmark, not as an exact prediction.

The standard analytic template. Collecting the robust cooling asymptote and the simulation-supported multiplication benchmark gives

$$Y_\gamma(E_\gamma) = \frac{K}{E_X} \begin{cases} \left(\dfrac{E_\gamma}{E_X}\right)^{-3/2}, & E_\gamma \le E_X, \\ \left(\dfrac{E_\gamma}{E_X}\right)^{-2}, & E_X < E_\gamma \le E_c, \\ 0, & E_\gamma > E_c. \end{cases} \tag{2.124}$$

The number spectrum steepens from $E^{-3/2}$ to approximately E^{-2} at E_X and is suppressed at the escape edge E_c. Its weak dependence on the energy and identity of a sufficiently energetic primary is the origin of the term *universal cascade spectrum* [40, 41]. Here "universal" refers to this limiting, fully developed shape; it does not make the approximations above exact.

If the primary deposits an energy E_0 into the IC–pair-production channel, energy conservation fixes the normalization:

$$E_0 = \int dE_\gamma\, E_\gamma Y_\gamma(E_\gamma) = K E_X \left[2 + \ln\left(\frac{E_c}{E_X}\right)\right], \tag{2.125}$$

and therefore

$$K = \frac{E_0}{E_X\,[2 + \ln(E_c/E_X)]}. \tag{2.126}$$

The same normalization follows from the electron flow. In the Thomson regime $E_e = (E_c/2)(E_\gamma/E_X)^{1/2}$, and Eq. (2.117) with $\Phi_e = \Phi_0$ reproduces the low-energy branch with $K = \Phi_0 E_c/(4E_X)$. With $\Phi_e = \Phi_0 E_c/(2E_e)$ above $E_c/2$, the energy radiated by leptons whose IC photons escape, $\int \Phi_e\, dE_e$ up to $E_e(E_c)$, reproduces Eq. (2.125). The closure of Eq. (2.121) and the template are therefore mutually consistent.

This assumes that competing losses are negligible. Broad radiation fields, finite propagation distances, cosmological redshift, Klein–Nishina effects, and magnetic deflections or synchrotron losses all shift or smooth the template. We return to these limitations after applying it to the CMB and EBL in the next section.

2.13 Extragalactic cascades on the CMB and EBL

We now apply the cascade picture to intergalactic space. The essential new ingredient is that pair production and IC cooling need not be controlled by the same radiation field. The extragalactic background light (EBL) contains optical and infrared photons, whereas the cosmic microwave background (CMB) contains much softer but far more abundant photons. Schematically,

$$\varepsilon_{\rm CMB} \ll \varepsilon_{\rm EBL}, \qquad U_{\rm CMB} \gg U_{\rm EBL},$$

where U_i denotes radiation energy density.

These inequalities create a useful division of labour over the energy range relevant to the last cascade generations:

- the EBL sets the *last pair-production edge*, because its larger photon energy allows lower-energy γ rays to produce pairs;
- the CMB usually dominates *electron cooling*, because its energy density is larger and IC scattering remains in the Thomson regime to higher electron energies than it does on the EBL.

This is not a general rule based only on the largest target-photon energy or number density. The absorption rate contains the full integral over $n(\varepsilon)\sigma_{\gamma\gamma}$ and the path length, while the Thomson IC loss rate is proportional to $U_{\rm rad}$ and is reduced by Klein–Nishina effects. The statements above summarize which terms dominate for the observed CMB and EBL in the final stages of a long extragalactic cascade.

In a monochromatic representation of each background, define the pair-production threshold scales

$$E_{c,i} \equiv \frac{m_e^2 c^4}{\varepsilon_i}, \qquad i \in \{\mathrm{CMB}, \mathrm{EBL}\}. \tag{2.127}$$

Because $\varepsilon_{\rm CMB} \ll \varepsilon_{\rm EBL}$,

$$E_{c,\rm EBL} \ll E_{c,\rm CMB}.$$

The resulting energy flow can be summarized as

$$\begin{aligned} E_\gamma \gtrsim E_{c,\rm CMB} &: \text{high-energy leading-particle cycling,} \\ E_{c,\rm EBL} \lesssim E_\gamma \lesssim E_{c,\rm CMB} &: \text{cascade multiplication, chiefly EBL PP and CMB IC,} \\ E_\gamma \lesssim E_{c,\rm EBL} &: \text{terminal cooling and escaping photons.} \end{aligned}$$

This hierarchy supplies the broad multiplication range that was absent in the single-field toy model.

The two observable energy scales. Let "PP" denote the photons that control the final pair-production edge and "IC" those that control electron cooling. In the sharp-threshold approximation,

$$E_c \simeq E_{c,\rm PP} \equiv \frac{m_e^2 c^4}{\varepsilon_{\rm PP}}. \tag{2.128}$$

As in Section 2.12, the last near-threshold conversion produces electrons and positrons with characteristic energy

$$E_{e,\mathrm{last}} \sim \frac{E_c}{2}. \tag{2.129}$$

Their IC photons mark the transition between the terminal-cooling and multiplication parts of the escaping spectrum:

$$E_X \equiv E_\gamma\left(E_e = \frac{E_c}{2}\right). \tag{2.130}$$

If the last pairs scatter in the Thomson regime on photons of energy $\varepsilon_{\mathrm{IC}}$, Eq. (2.112) gives

$$\begin{aligned} E_X &\simeq \frac{4}{3}\left(\frac{E_c}{2m_e c^2}\right)^2 \varepsilon_{\mathrm{IC}} \\ &\simeq \frac{1}{3} E_c \frac{\varepsilon_{\mathrm{IC}}}{\varepsilon_{\mathrm{PP}}}. \end{aligned} \tag{2.131}$$

For extragalactic cascades, PP on the EBL and IC cooling on the CMB therefore give

$$E_c \simeq \frac{m_e^2 c^4}{\varepsilon_{\mathrm{EBL}}}, \qquad E_X \simeq \frac{1}{3} E_c \frac{\varepsilon_{\mathrm{CMB}}}{\varepsilon_{\mathrm{EBL}}} \ll E_c. \tag{2.132}$$

For the representative values $\varepsilon_{\mathrm{EBL}} = 0.1$ eV and $\varepsilon_{\mathrm{CMB}} = 6.3 \times 10^{-4}$ eV,

$$E_c \simeq 2.6\ \mathrm{TeV}, \qquad E_X \simeq 5.5\ \mathrm{GeV}, \qquad \frac{E_c}{E_X} \simeq 4.8 \times 10^2. \tag{2.133}$$

The IC scattering that defines E_X is safely Thomson in this example. The small ratio $\varepsilon_{\mathrm{CMB}}/\varepsilon_{\mathrm{EBL}}$ has opened more than two decades between E_X and E_c, in contrast to $E_X \sim E_c/3$ for the one-field example.

The quoted numbers are illustrative. The EBL is broad rather than monochromatic, and the physical escape energy is determined by $\tau_{\gamma\gamma}(E, z) \sim 1$, not by kinematics alone. Source distance, redshift, and the evolving EBL therefore shift and smooth both characteristic features.

From the spectral template to a diffuse intensity. We now use the analytic template of Eq. (2.124). Here $N_\gamma(E)$ denotes the present-day differential number density of escaping cascade photons, rather than the per-primary yield used in the previous section:

$$N_\gamma(E) = \frac{K}{E_X} \begin{cases} \left(\dfrac{E}{E_X}\right)^{-3/2}, & E \le E_X, \\ \left(\dfrac{E}{E_X}\right)^{-2}, & E_X < E \le E_c, \\ 0, & E > E_c. \end{cases} \tag{2.134}$$

Its normalization is fixed by the present cascade radiation density,

$$\begin{aligned} \omega_{\mathrm{cas}} &\equiv \int dE\, E\, N_\gamma(E) \\ &= K E_X \left[2 + \ln\left(\frac{E_c}{E_X}\right)\right], \end{aligned} \tag{2.135}$$

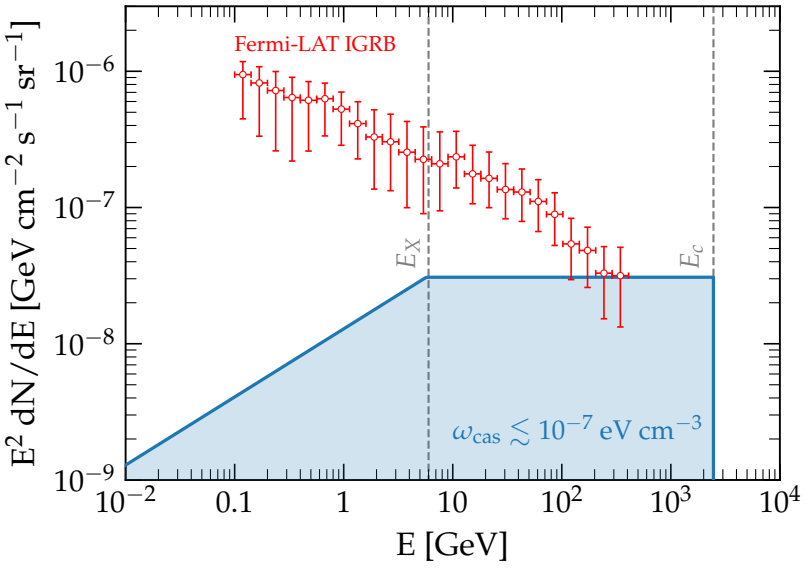


Figure 14: Analytic two-field cascade template for IC cooling on the CMB ($\varepsilon_{\rm IC} = 6.3 \times 10^{-4}$ eV) and pair production on a representative EBL component ($\varepsilon_{\rm PP} = 0.1$ eV), normalized to $\omega_{\rm cas} = 10^{-7}$ eV cm^{-3}. The points show the *Fermi*-LAT IGRB measurement [42]. Dashed lines mark the characteristic E_X and E_c of Eq. (2.133).

or

$$K = \frac{\omega_{\rm cas}}{E_X\,[2 + \ln(E_c/E_X)]}. \tag{2.136}$$

For a homogeneous isotropic photon population, number density and differential intensity are related by

$$I_\gamma(E) = \frac{c}{4\pi} N_\gamma(E). \tag{2.137}$$

On the benchmark E^{-2} branch,

$$E^2 I_\gamma(E) = \frac{c}{4\pi} K E_X = \frac{c}{4\pi} \frac{\omega_{\rm cas}}{2 + \ln(E_c/E_X)}. \tag{2.138}$$

Thus the approximately flat part of an $E^2 I_\gamma$ plot measures the cascade energy per logarithmic photon interval. Detailed kinetic calculations give a slope closer to $E^{-1.9}$, for which $E^2 I_\gamma \propto E^{0.1}$ rather than being exactly constant [39]. Equation (2.138) should consequently be used as an analytic estimate.

The cascade bound. The isotropic γ-ray background (IGRB) measured by *Fermi*-LAT provides a calorimetric limit on electromagnetic energy injected at much higher energies. At every energy the cascade contribution must not exceed the measured background. If an energy E is chosen on the intermediate branch, inversion of Eq. (2.138) gives the simple bound

$$\omega_{\rm cas} \lesssim \frac{4\pi}{c} \left[E^2 I_\gamma(E)\right]_{\rm IGRB} \left[2 + \ln\left(\frac{E_c}{E_X}\right)\right]. \tag{2.139}$$

At the high-energy end of the *Fermi*-LAT spectrum, $E^2 I_\gamma$ is a few $\times 10^{-8}\,\mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}$. Together with $E_c/E_X \sim 10^2$–10^3, this gives the useful scale

$$\omega_{\mathrm{cas}} \lesssim \mathcal{O}(10^{-7})\ \mathrm{eV\,cm^{-3}}. \tag{2.140}$$

The estimate is intentionally conservative because it assigns the entire IGRB at the chosen energy to the cascade. Precision bounds compare a numerically propagated spectrum with all energy bins and depend on the source-redshift distribution, EBL model, and resolved-source contribution; values from a few 10^{-8} to several 10^{-7} $\mathrm{eV\,cm^{-3}}$ are possible under different assumptions [39, 42].

This is the *cascade bound*. Interactions of ultrahigh-energy cosmic rays with background photons inject neutral-pion photons, Bethe–Heitler pairs, and electrons from charged-pion decays; the photopion channel simultaneously produces neutrinos. The electromagnetic secondaries cascade into the *Fermi*-LAT band even when their injection energy is far higher. Consequently, the observed IGRB limits the electromagnetic energy associated with proton-dominated UHECR models and, indirectly, their cosmogenic-neutrino flux. Translating the present density ω_{cas} into a limit on source power requires a cosmological propagation calculation; it is not fixed by the static estimate alone.

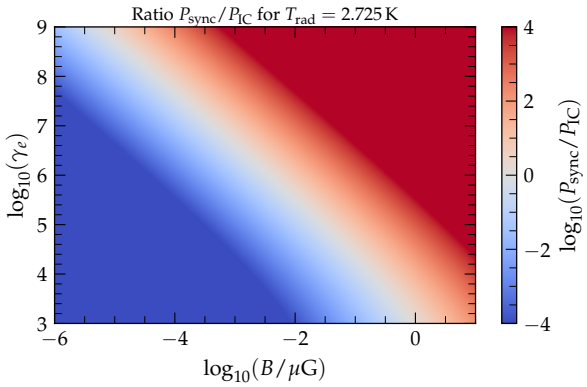


Figure 15: Ratio of synchrotron to IC power for cascade pairs in a CMB-dominated target field, including Klein–Nishina suppression. The contour $P_{\mathrm{syn}} = P_{\mathrm{IC}}$ separates IC-dominated cascade regeneration from synchrotron-dominated reprocessing.

When the standard cascade is modified. The template assumes that pairs return most of their energy to IC photons. Magnetic fields introduce synchrotron losses. A convenient comparison is

$$\frac{P_{\mathrm{syn}}}{P_{\mathrm{IC}}} = \frac{U_B}{U_{\mathrm{rad}} F_{\mathrm{KN}}(\gamma)}, \qquad U_B = \frac{B^2}{8\pi}, \tag{2.141}$$

where $F_{\mathrm{KN}} \leq 1$ describes Klein–Nishina suppression of IC cooling. In the Thomson regime on the CMB, equality of the magnetic and radiation energy densities corresponds to

$$B_{\mathrm{cr}} = \sqrt{8\pi U_{\mathrm{CMB}}} \simeq 3.2\ \mu\mathrm{G}\,(1+z)^2. \tag{2.142}$$

For $B \gtrsim B_{\rm cr}$, synchrotron emission diverts energy from the usual GeV–TeV cascade. At very high electron energies $F_{\rm KN} \ll 1$, so this can occur at fields well below $B_{\rm cr}$. Even much weaker intergalactic fields can deflect the pairs, altering the angular extent and arrival time of a point-source cascade without dominating its energy losses.

Finally, real cosmological cascades are not described by one fixed pair of energies. The EBL has a broad evolving spectrum, all photon energies redshift, nearby cascades may be underdeveloped, and the injection history matters. These effects replace the sharp breaks by smooth transitions and reduce "universality" to an approximate or weak statement. The analytic template nevertheless remains valuable because it identifies the energy flow and gives a transparent first estimate before a numerical transport calculation.

2.14 Nuclei–proton interactions

We now turn from leptonic radiation to hadronic interactions. Relativistic cosmic-ray protons and nuclei collide with ambient gas, composed mainly of hydrogen and helium. These collisions redistribute the projectile energy among secondary hadrons and thereby connect cosmic rays to γ rays, neutrinos, and secondary charged particles.

These collisions are governed by the strong interaction, which is short-ranged and acts on nuclear length scales. A first estimate of their cross section is therefore obtained by treating a nucleus as an extended object with the empirical radius

$$R_N \simeq r_0\, A^{1/3}, \qquad r_0 \simeq 1.2\ \text{fm}, \tag{2.143}$$

where A is the mass number [43].

In a simple black-disk picture, the corresponding geometrical cross section is

$$\sigma_{\rm geom} \simeq \pi R_N^2. \tag{2.144}$$

Substituting the above scaling gives

$$\sigma_{\rm geom} \simeq \pi (r_0 A^{1/3})^2 \simeq 4.5\, A^{2/3}\ \text{fm}^2 \simeq 45\, A^{2/3}\ \text{mb}, \tag{2.145}$$

where $1\ \text{fm}^2 = 10\ \text{mb}$. This estimate is not a precision formula: a nucleus–nucleus collision depends on both radii, and measured inelastic cross sections retain an energy dependence. It nevertheless captures the useful large-A scaling $\sigma \propto A^{2/3}$ [44].

The interaction time is the mean time between collisions in a medium with target number density $n_{\rm target}$. For a relativistic projectile,

$$\tau_{\rm int} \simeq \frac{1}{n_{\rm target}\, \sigma_{\rm geom}\, c}\,. \tag{2.146}$$

For a cosmic-ray nucleus propagating through hydrogen of number density $n_{\rm H}$, Eq. (2.145) gives

$$\tau_{\rm int} \simeq 25\, A^{-2/3} \left(\frac{n_{\rm H}}{\text{cm}^{-3}}\right)^{-1} \text{Myr}. \tag{2.147}$$

Two simple but important messages follow immediately:

- in dense environments (e.g. molecular clouds, starburst galaxies) the hadronic interaction time can become comparable to, or shorter than, competing timescales such as escape;
- since $\sigma \propto A^{2/3}$, heavier nuclei have shorter interaction times, $\tau_{\rm int} \propto A^{-2/3}$.

This is an interaction time, not yet an energy-loss time. The latter also depends on the mean fractional energy loss per collision and is derived below for proton–proton interactions.

Unlike the electromagnetic processes discussed earlier, the hadronic cross sections and secondary-particle yields considered here cannot generally be predicted from first principles with the required accuracy. Both fragmentation and pion production therefore rely on phenomenological models and parametrizations constrained by laboratory and accelerator measurements.

Fragmentation cross sections. When a cosmic-ray nucleus interacts inelastically at GeV/n energies, the projectile itself almost always breaks up into lighter nuclei, often together with newly produced particles such as pions. This process is called *fragmentation* (or *spallation*). To a good approximation the fragment keeps the velocity of its parent, and hence its kinetic energy per nucleon. This is known as the *straight-ahead* approximation and is well justified above a few GeV/n [45].

Within this approximation, fragmentation is described by partial cross sections $\sigma_{i\to j}(T)$ for producing a fragment j from a projectile i, which depend only on the kinetic energy per nucleon T of the projectile.

These are *inclusive* cross sections: they count the production of fragment j regardless of what else emerges from the collision. *Exclusive* cross sections, instead, refer to a fully specified final state. The ratio of an inclusive cross section to the total inelastic cross section, $P_{i\to j} = \sigma_{i\to j}/\sigma_i$, is the mean number of fragments j produced per inelastic collision.

Moreover, unstable fragments whose decay times in the propagation frame are much shorter than the transport time can be treated as decaying immediately after production. For example, decay times of a kyr or less are short compared with Myr for Galactic transport or Gyr for cosmological propagation. These isotopes are called *ghost nuclei* and need not be propagated explicitly. Their production cross sections, weighted by the relevant decay branching ratios, are added to the direct production cross sections of their daughters, giving *cumulative* cross sections [45, 46].

For example, for $^{12}\mathrm{C} + p$ at 3.66 GeV/n, removing a single nucleon gives $\sigma_{^{12}\mathrm{C}\to{}^{11}\mathrm{C}} \simeq 29$ mb and $\sigma_{^{12}\mathrm{C}\to{}^{11}\mathrm{B}} \simeq 28$ mb, and removing two gives $\sigma_{^{12}\mathrm{C}\to{}^{10}\mathrm{B}} \simeq 12$ mb [47]. Light fragments have large partial cross sections too, $\sigma_{^{12}\mathrm{C}\to{}^{4}\mathrm{He}} \simeq 160$ mb and $\sigma_{^{12}\mathrm{C}\to{}^{1}\mathrm{H}} \simeq 140$ mb. The total inelastic cross section is $\sigma_{\mathrm{C}} \sim 250$ mb [45], close to the geometric estimate of Eq. (2.145). The partial cross sections summed over all fragments exceed it, because a single collision produces several fragments. These values are direct cross sections, however, and ghost nuclei can contribute as much as direct production. In this example, the $^{11}\mathrm{C}$ fragment decays to $^{11}\mathrm{B}$ by β^+ decay with $t_{1/2} \simeq 20$ min, so the cumulative cross section for producing $^{11}\mathrm{B}$ is $\sigma_{^{12}\mathrm{C}\to{}^{11}\mathrm{B}} + \sigma_{^{12}\mathrm{C}\to{}^{11}\mathrm{C}} \simeq 57$ mb, about twice the direct value.

Fragmentation depends on the detailed structure of the projectile, through isospin and odd–even effects and the Coulomb barrier that hinders proton emission, so each reaction behaves somewhat differently. Nuclear codes reach 5–10% at best, for a few channels in narrow energy ranges, and those channels are not necessarily the ones that matter most for cosmic rays [45, 46]. In practice, measurements must be interpolated.

Various parametrizations have been proposed in the literature, based on observed regularities and fits to the available data. The semi-empirical formulae of Silberberg and Tsao exploit regularities in the mass difference between projectile and fragment and in the neutron-to-proton ratio of the fragment [48–50]. The empirical formulae of Webber and coworkers factorize the cross section into a charge-dependent term, an isotopic distribution, and an energy dependence fitted at one energy and extrapolated to the others [51, 52]. Propagation codes such as GALPROP combine these formulae with nuclear codes and renormalize them, reaction by reaction, to whatever data exist [53, 54].

The main limitation of these approaches, however, is that the data themselves are sparse. Most measurements lie below a few tens of MeV/n, or between a few hundred MeV/n and a couple of GeV/n with only one or a few points per channel. Results from different groups can differ by significant factors, and many reactions relevant to astrophysics have never been measured. Above a few GeV/n, nearly all parametrizations assume energy-independent cross sections, although the nucleon–nucleon inelastic cross section is known to rise with energy. The resulting uncertainties range from 5–10% for the best-measured proton-induced channels to 20–50% or more for others [45, 46]. By contrast, current experiments such as AMS-02 measure cosmic-ray fluxes to 1–3%.

This disparity has direct consequences for one of the pillars of cosmic-ray physics: the reconstruction of the Galactic journey of cosmic rays from their nuclear composition, as we will see in **??**. The grammage, residence time, and diffusion coefficient inferred from these abundances are only as accurate as the nuclear inputs used to interpret them. As already anticipated in 2002 [45], production cross sections have thus become a limiting factor in the interpretation of secondary-to-primary ratios. Fully exploiting the promise of precise cosmic-ray observations therefore demands a corresponding advance in nuclear physics, making new cross-section measurements—particularly between 1 and 100 GeV/n—a central priority [55].

But hadronic collisions do more than break nuclei apart: they also create pions, whose decays turn cosmic rays into γ rays and neutrinos alike—a textbook case of *multi-messenger astrophysics*.

Pion production in proton–proton collisions. The simplest nucleus–nucleus interaction is a proton–proton (pp) collision. It is also one of the most important for multi-messenger astrophysics: once the pion-production threshold is exceeded, part of the available center-of-mass energy goes into creating new hadrons, the lightest of which are pions. Their decays provide a direct link between charged cosmic rays, γ rays, and neutrinos.

At the level of representative channels, one may write schematically

$$p + p \;\to\; \begin{cases} p + p + \pi^0, \\ p + n + \pi^+, \\ p + p + \pi^+ + \pi^-, \\ \;\vdots \end{cases} \tag{2.148}$$

with multi-pion final states (and heavier mesons) becoming increasingly important as the collision energy grows. Pions are short-lived and decay essentially at the production site, feeding the electromagnetic and neutrino channels.

Neutral pions decay electromagnetically,

$$\pi^0 \to \gamma + \gamma, \tag{2.149}$$

while charged pions decay weakly via muons,

$$\pi^+ \to \mu^+ + \nu_\mu, \qquad \mu^+ \to e^+ + \bar{\nu}_\mu + \nu_e, \tag{2.150}$$

$$\pi^- \to \mu^- + \bar{\nu}_\mu, \qquad \mu^- \to e^- + \nu_\mu + \bar{\nu}_e. \tag{2.151}$$

Hence each $\pi^\pm$ produces three neutrinos or antineutrinos along its decay chain, whereas each π^0 produces two photons. At the source the flavor ratio from an uncooled pion–muon chain is approximately $\nu_e : \nu_\mu : \nu_\tau = 1 : 2 : 0$, counting neutrinos and antineutrinos. The simultaneous injection of γ rays and neutrinos with related energy budgets is therefore a characteristic hadronic signature. Propagation and cooling inside the source can, however, modify the observable relation. Neutrino oscillations during propagation to Earth also change the flavor ratio.

Pion production turns on only once the collision energy exceeds the kinematic threshold. For a projectile proton of total lab-frame energy E_p scattering on a proton at rest, the Mandelstam invariant is $s = 2m_p^2c^4 + 2m_pc^2E_p$. Threshold corresponds to producing the lightest allowed final state at rest in the center-of-mass frame, i.e. $\sqrt{s_{\rm th}} = (2m_p + m_{\pi^0})c^2$ (see Section A.3), which gives

$$E_p^{\rm th} = \frac{(2m_p + m_{\pi^0})^2 - 2m_p^2}{2m_p}\, c^2 \simeq 1.22\ \text{GeV}, \tag{2.152}$$

or, equivalently, a *kinetic* threshold

$$T_p^{\rm th} = E_p^{\rm th} - m_pc^2 \simeq 280\ \text{MeV}. \tag{2.153}$$

The charged channel $pp \to pn\pi^+$ has a slightly higher threshold, $T_p^{\rm th} \simeq 290$ MeV. Thus mildly relativistic protons already suffice to initiate pion production in ambient hydrogen.

Cross-sections for proton–proton collisions. The *inelastic* cross section $\sigma_{\rm inel}$ for pp collisions quantifies the probability that an interaction produces new hadrons (most importantly pions), as opposed to purely elastic scattering where the protons emerge intact.

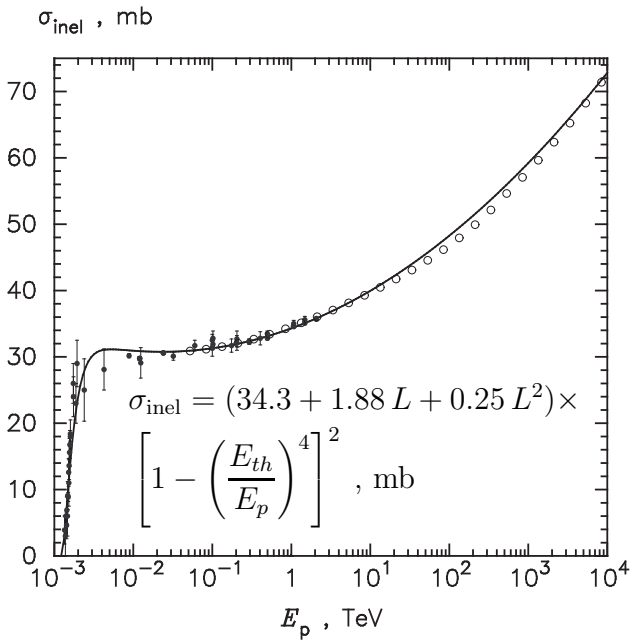


Figure 16: Inelastic proton–proton cross section as a function of the projectile proton energy. Adapted from [56].

In an inelastic event, a non-negligible fraction of the projectile energy is transferred to secondary particles, thereby powering γ-ray and neutrino emission through pion decays.

Kinematically, inelasticity turns on rapidly once the pion-production threshold is crossed ($T_p^{\rm th} \simeq 280$ MeV, or equivalently $E_p^{\rm th} \simeq 1.22$ GeV for a target proton at rest). Above a few GeV the cross section is of order ~ 30–40 mb, and at higher energies it increases only slowly, reaching ~ 60 mb by PeV projectile energies (see Fig. 16). This mild rise reflects the well-known logarithmic growth of hadronic cross sections at high center-of-mass energy.

For practical calculations it is convenient to adopt a smooth analytic fit. A widely used parametrization is [56]

$$\sigma_{\rm inel}(E_p) = \left(34.3 + 1.88\,L + 0.25\,L^2\right)\left[1 - \left(\frac{E_{\rm th}}{E_p}\right)^4\right]^2 \text{ mb}, \qquad L \equiv \ln\!\left(\frac{E_p}{1\ \text{TeV}}\right), \quad (2.154)$$

where E_p is the total projectile energy in the lab frame and $E_{\rm th} \simeq 1.22 \times 10^{-3}$ TeV is the π^0 production threshold energy. The fit applies for $E_p \geq E_{\rm th}$; below threshold, $\sigma_{\rm inel} = 0$. The threshold factor enforces $\sigma_{\rm inel} \to 0$ as $E_p \to E_{\rm th}$, while for $E_p \gg E_{\rm th}$ it tends to unity, leaving the characteristic slow $\ln^2(E_p)$ rise at very high energies. This fit will be used below when estimating hadronic yields.

For orientation, take a representative inelastic proton–proton cross section in the GeV–TeV range, $\sigma_{pp,\rm inel} \simeq 30$ mb. The corresponding *interaction* (collision) timescale is

$$\tau_{\rm int}^{pp} = \frac{1}{n_{\rm gas}\, c\, \sigma_{pp,\rm inel}} \simeq 3.5 \times 10^7\ \text{yr} \left(\frac{n_{\rm gas}}{1\ \text{cm}^{-3}}\right)^{-1} \left(\frac{\sigma_{pp,\rm inel}}{30\ \text{mb}}\right)^{-1}. \quad (2.155)$$

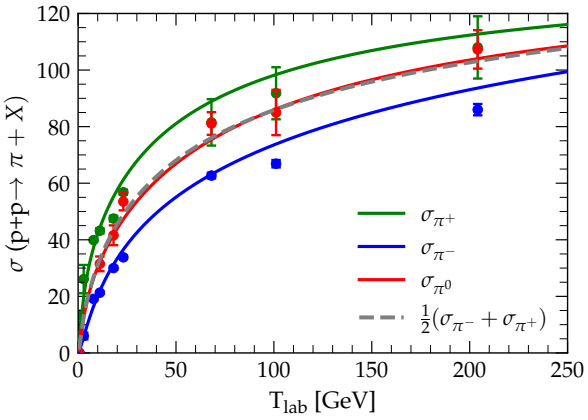


Figure 17: Total inclusive (multiplicity–weighted) cross sections for pion production in proton–proton collisions, using the parametrizations of Norbury [58].

An inelastic collision does not remove the entire proton energy. If $\kappa_{\rm inel}$ is the mean fractional loss per collision, with $\kappa_{\rm inel} \sim 0.5$ a useful high-energy benchmark, then

$$\left|\frac{dE}{dt}\right| \simeq \kappa_{\rm inel}\,\frac{E}{\tau_{\rm int}^{pp}} = \kappa_{\rm inel}\, n_{\rm gas}\, c\, \sigma_{pp,\rm inel}\, E. \tag{2.156}$$

The corresponding loss time is therefore

$$\begin{aligned}\tau_{\rm cool}^{pp} \equiv \frac{E}{|dE/dt|} &\simeq \frac{1}{\kappa_{\rm inel}\, n_{\rm gas}\, c\, \sigma_{pp,\rm inel}} = \frac{\tau_{\rm int}^{pp}}{\kappa_{\rm inel}} \\ &\simeq 7.0\times10^{7}\ {\rm yr}\left(\frac{n_{\rm gas}}{1\ {\rm cm}^{-3}}\right)^{-1}\left(\frac{\sigma_{pp,\rm inel}}{30\ {\rm mb}}\right)^{-1}\left(\frac{\kappa_{\rm inel}}{0.5}\right)^{-1}.\end{aligned} \tag{2.157}$$

This timescale is crucial in γ-ray astrophysics: for optically thin pp emission, the scaling $W_p \propto L_\gamma \tau_{\rm cool}^{pp}$ connects the *measured* γ-ray luminosity L_γ to the energy W_p stored in relativistic protons [57].

Inclusive cross sections for pion production. In astrophysical applications we are rarely interested in a single exclusive channel such as $p + p \to p + n + \pi^+$. Rather, we want the *overall yield* of pions produced in pp interactions, summed over all accompanying secondaries. This motivates the use of *inclusive* pion-production cross sections, conventionally denoted $\sigma_{pp\to\pi X}$, where X stands for an arbitrary hadronic final state.

Concretely, $\sigma_{pp\to\pi^a X}$ is obtained by summing (and integrating over phase space) all processes of the form

$$p + p \to \pi^a + X, \qquad a \in \{+, 0, -\}.$$

In this "inclusive" sense the cross section effectively counts pion *multiplicity*: at high energy, where multi-pion production is common, $\sigma_{pp\to\pi X}$ should be viewed as $\sigma_{\rm inel} \times \langle n_\pi \rangle$ rather than as the probability to produce at least one pion.[3]

[3] Below the two-pion threshold, the pion multiplicity is one, so the inclusive cross section for a given pion charge equals the sum over its allowed exclusive channels.

A convenient set of empirical fits for the total inclusive cross sections as functions of the projectile proton kinetic energy T_p, measured in GeV in the lab frame, is provided in [58]:

$$\sigma_{pp\to\pi^+X}(T_p) = \left(0.00717 + 0.0652\,\frac{\ln T_p}{T_p} + \frac{0.162}{T_p^2}\right)^{-1} \text{mb}, \tag{2.158}$$

$$\sigma_{pp\to\pi^-X}(T_p) = \left(0.00456 + \frac{0.0846}{T_p^{1/2}} + \frac{0.577}{T_p^{3/2}}\right)^{-1} \text{mb}, \tag{2.159}$$

$$\sigma_{pp\to\pi^0X}(T_p) = \left(0.007 + 0.1\,\frac{\ln T_p}{T_p} + \frac{0.3}{T_p^2}\right)^{-1} \text{mb}, \tag{2.160}$$

where the numerical value of T_p is understood in GeV and the resulting cross sections are in mb. These compact fits describe total, multiplicity-weighted yields, rather than exclusive cross sections. They must be set to zero below the corresponding production thresholds; the expressions themselves do not enforce this cutoff.

The approximate production thresholds imposed on these fits are:

- π^+ production turns on at $T_p \simeq 289$ MeV.
- π^- production has a substantially higher threshold, $T_p \simeq 600$ MeV, because charge conservation forbids single-π^- production in pp collisions; the lightest channels producing a π^- necessarily involve at least one additional pion (e.g. π^+) in the final state.
- π^0 production turns on at $T_p \simeq 280$ MeV.

As a consequence, $\sigma_{pp\to\pi^+X}$ switches on earlier and remains larger than $\sigma_{pp\to\pi^-X}$ over a broad energy range, a charge asymmetry that is clearly visible in Fig. 17.

At sufficiently high energies, where isospin symmetry becomes a useful guide and the pion multiplicities are large, one often finds the approximate relation

$$\sigma_{pp\to\pi^0X} \simeq \frac{1}{2}\left(\sigma_{pp\to\pi^+X} + \sigma_{pp\to\pi^-X}\right), \tag{2.161}$$

reflecting the tendency of the strong interaction to populate charged and neutral pion states at comparable rates.

The π^+/π^- asymmetry has direct astrophysical consequences for *secondary* leptons and neutrinos. Charged-pion decays feed $e^\pm$ and (anti)neutrinos, and in ordinary matter targets the dominance of π^+ production implies a larger *secondary* positron yield than electron yield. This charge asymmetry is therefore a useful diagnostic when interpreting cosmic-ray secondaries and diffuse multi-messenger emission in hadronic scenarios [59, 60].

Pion emissivity. The *pion emissivity* $q_\pi(E_\pi)$ is the production rate of pions per unit volume, time, and energy. It provides the bridge between a parent cosmic-ray proton population and the secondary γ rays (from π^0 decay) and neutrinos (from $\pi^\pm$ decay) produced in hadronic interactions [56].

Consider a homogeneous hydrogen target of number density n_H. For one specified pion charge state, the general production rate is

$$q_\pi(E_\pi) = c\, n_\mathrm{H} \int_{E_p^\mathrm{th}}^{\infty} \mathrm{d}E_p\, N_p(E_p)\, \frac{\mathrm{d}\sigma_{pp\to\pi X}(E_p)}{\mathrm{d}E_\pi}\,, \tag{2.162}$$

where $N_p(E_p)$ is the differential proton number density and $\mathrm{d}\sigma_{pp\to\pi X}/\mathrm{d}E_\pi$ is the (inclusive) differential cross section for producing a pion with energy E_π. For back-of-the-envelope estimates, the delta-functional approximation replaces the broad secondary distribution by one effective pion carrying a fixed fraction of the projectile energy,

$$E_\pi \simeq K_\pi\, E_p, \tag{2.163}$$

so that the differential cross section is approximated as

$$\frac{\mathrm{d}\sigma_{pp\to\pi X}}{\mathrm{d}E_\pi} \simeq \sigma_\mathrm{inel}(E_p)\, \delta\big(E_\pi - K_\pi E_p\big). \tag{2.164}$$

Here K_π is an effective energy fraction for the selected pion species. A common pedagogical choice is

$$K_\pi \simeq 0.17, \tag{2.165}$$

motivated by an overall inelasticity of order 0.5 and an approximately equal division of the pion energy budget among the three charge states. This approximation is intended to reproduce the characteristic energy mapping, not the pion multiplicity or the threshold-region spectrum.

Substituting Eq. (2.164) into Eq. (2.162) and using $\delta(E_\pi - K_\pi E_p) = (1/K_\pi)\,\delta(E_p - E_\pi/K_\pi)$ yields immediately

$$q_\pi(E_\pi) \simeq \frac{c\, n_\mathrm{H}}{K_\pi}\, \sigma_\mathrm{inel}\!\left(\frac{E_\pi}{K_\pi}\right)\, N_p\!\left(\frac{E_\pi}{K_\pi}\right). \tag{2.166}$$

Two features are noteworthy:

- typical pion energies track the parent proton energies as $E_\pi \sim K_\pi E_p$.
- if $N_p(E_p) \propto E_p^{-\alpha_p}$, then (up to the slow energy dependence of σ_inel) one has $q_\pi(E_\pi) \propto E_\pi^{-\alpha_p}$.

2.15 The π^0 γ-ray spectrum

Neutral-pion decay provides a direct link between hadronic interactions and observable γ rays:

$$\pi^0 \to \gamma + \gamma, \tag{2.167}$$

with an extremely short lifetime, $\tau_{\pi^0} \simeq 8.4 \times 10^{-17}\,\mathrm{s}$, so that the decay effectively occurs at the production site.

For the decay kinematics below, we use units with $c = 1$, so masses are expressed in energy units.

In the pion rest frame (primes), momentum conservation enforces back-to-back photons, and energy conservation fixes their energies to

$$E'_\gamma = \frac{m_\pi}{2}, \tag{2.168}$$

where $m_\pi \simeq 135\ \mathrm{MeV}$ is the neutral-pion mass.

In the lab frame, a pion with speed $\beta_\pi \equiv v_\pi/c$ and Lorentz factor γ_π boosts the photon energy according to the standard Lorentz transformation,

$$E_\gamma = \frac{m_\pi}{2}\,\gamma_\pi\,(1 + \beta_\pi \cos\theta'), \tag{2.169}$$

where θ' is the emission angle of the photon in the pion rest frame measured with respect to the pion boost direction. Varying $\cos\theta' \in [-1, 1]$ gives the kinematic endpoints

$$E_\gamma^{\mathrm{min/max}} = \frac{m_\pi}{2}\,\gamma_\pi\,(1 \mp \beta_\pi) = \frac{E_\pi \mp p_\pi}{2}, \qquad p_\pi = \sqrt{E_\pi^2 - m_\pi^2}. \tag{2.170}$$

Two limiting cases are worth noting:

- Non-relativistic pions ($\beta_\pi \ll 1$). One has $E_\gamma^{\mathrm{min}} \simeq E_\gamma^{\mathrm{max}} \simeq m_\pi/2$, and the box collapses to a monochromatic two-photon line.
- Ultra-relativistic pions ($\gamma_\pi \gg 1$). The upper endpoint approaches the pion energy, $E_\gamma^{\mathrm{max}} \simeq E_\pi$, while the lower endpoint becomes very small, $E_\gamma^{\mathrm{min}} \simeq m_\pi^2/(4E_\pi) \ll E_\pi$.

Because π^0 is spinless, the decay is isotropic in its rest frame. For *one* of the two photons[4],

$$\frac{\mathrm{d}N_\gamma}{\mathrm{d}\Omega'} = \frac{1}{4\pi} \quad \Longrightarrow \quad \mathrm{d}N_\gamma = \frac{1}{2}\,\mathrm{d}\big(\cos\theta'\big)\,. \tag{2.171}$$

Differentiating Eq. (2.169) gives

$$\mathrm{d}E_\gamma = \frac{m_\pi}{2}\,\gamma_\pi\beta_\pi\,\mathrm{d}\big(\cos\theta'\big)\,, \tag{2.172}$$

so that the energy distribution for *one photon* is flat between the endpoints:

$$\left.\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma}\right|_{\text{one photon}} = \frac{1}{m_\pi\gamma_\pi\beta_\pi} = \frac{1}{\sqrt{E_\pi^2 - m_\pi^2}}. \tag{2.173}$$

Including both photons from the decay simply doubles this result. Thus, for a pion of fixed energy E_π,

[4] Writing $\mu' = \cos\theta'$, the angular measure after azimuthal integration is $2\pi\,\mathrm{d}\mu'$, with $\mu' \in [-1, 1]$.

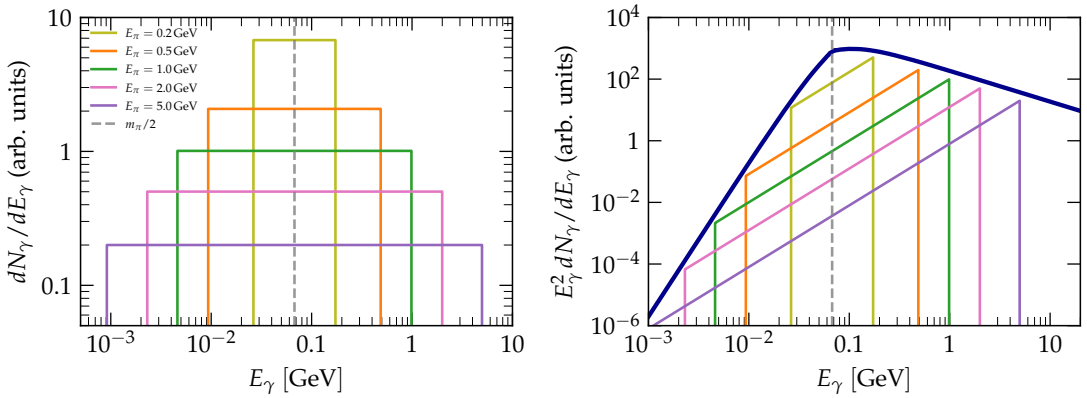


Figure 18: Illustration of how the pion-decay feature arises. *Left:* Box-like γ-ray spectra produced by individual π^0 mesons with different energies E_π. Each box is flat in E_γ between $E_\gamma^{\min}$ and $E_\gamma^{\max}$, yet all share the same logarithmic midpoint at $E_\gamma = m_\pi/2$ (vertical dashed line). *Right:* Superposition of many such boxes for a broad pion distribution. The photon-number spectrum retains a turnover tied to $m_\pi/2$; multiplying by powers of E_γ shifts the apparent peak but not the underlying kinematic scale.

$$\left.\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma}\right|_{\pi^0\to 2\gamma} = \frac{2}{m_\pi \gamma_\pi \beta_\pi} = \frac{2}{\sqrt{E_\pi^2 - m_\pi^2}}, \qquad E_\gamma \in [E_\gamma^{\min}, E_\gamma^{\max}], \tag{2.174}$$

and $\mathrm{d}N_\gamma/\mathrm{d}E_\gamma = 0$ outside this interval. Thus a pion of fixed nonzero momentum produces a box-like spectrum in E_γ. The height is independent of E_γ, and its integral over the interval is two. In the limit $p_\pi \to 0$, the width vanishes and the normalized box tends to $2\delta(E_\gamma - m_\pi/2)$.

To characterize the box on a logarithmic energy axis, it is natural to consider its *geometric* midpoint,

$$E_{\gamma,\mathrm{mid}} \equiv \sqrt{E_{\gamma,\min} E_{\gamma,\max}} \qquad \Longleftrightarrow \qquad \log E_{\gamma,\mathrm{mid}} = \frac{1}{2}\left(\log E_{\gamma,\min} + \log E_{\gamma,\max}\right). \tag{2.175}$$

Using Eq. (2.170) we find

$$E_{\gamma,\min} E_{\gamma,\max} = \left(\frac{m_\pi}{2}\gamma_\pi\right)^2 (1-\beta_\pi^2) = \left(\frac{m_\pi}{2}\right)^2, \tag{2.176}$$

since $\gamma_\pi^2(1-\beta_\pi^2) = 1$. Therefore

$$E_{\gamma,\mathrm{mid}} = \sqrt{E_{\gamma,\min} E_{\gamma,\max}} = \frac{m_\pi}{2} \simeq 67.5\ \mathrm{MeV}, \tag{2.177}$$

independent of the pion energy.

In realistic sources, neutral pions are produced with a broad distribution of energies. The γ-ray output is therefore a superposition of many such box-like spectra, one for each pion energy, all sharing the same logarithmic midpoint at $E_\gamma = m_\pi/2$. As shown explicitly below, the resulting photon-number spectrum is invariant under $E_\gamma \to m_\pi^2/(4E_\gamma)$, so it is

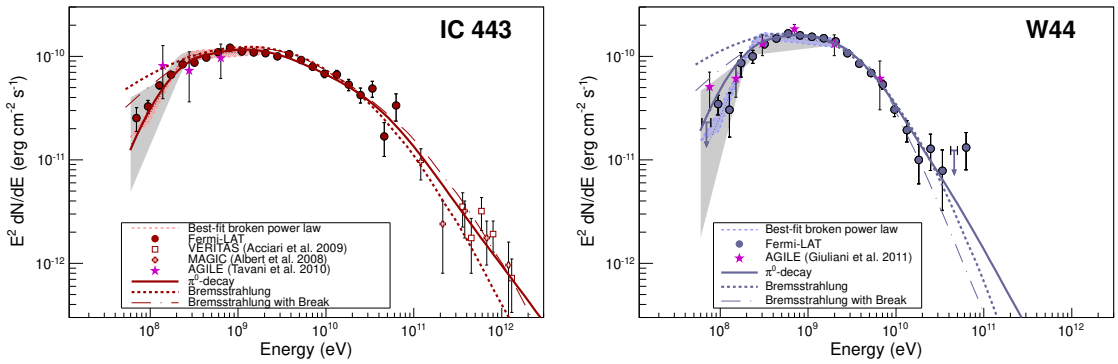


Figure 19: Evidence for a pion bump in two middle-aged supernova remnants interacting with molecular clouds, as observed by *Fermi*-LAT in 2013 (IC 443 and W44). Adapted from [61].

symmetric in $\log E_\gamma$ about $m_\pi/2 \simeq 67.5\,\mathrm{MeV}$. This kinematic turnover is the characteristic *pion bump* (Fig. 18). In an energy-weighted representation such as $E_\gamma^2 q_\gamma$, the factor E_γ^2 breaks that symmetry and moves the visible maximum to higher energy.

Whereas the normalization and high-energy tail reflect the parent proton population and the production cross section, the low-energy turnover has a robust kinematic scale set by the pion mass. That is why a deviation from a pure power law around this characteristic energy is an intriguing hint of a hadronic origin.

A celebrated example is provided by *Fermi*-LAT observations [61], which revealed a low-energy spectral feature consistent with π^0 decay in the γ-ray emission of two middle-aged supernova remnants interacting with molecular clouds (Fig. 19). This detection offered some of the clearest evidence that, at least in these systems, the γ-ray emission is powered by hadronic cosmic rays accelerated in the remnant and interacting with dense target material.

Gamma-ray emissivity. The *γ-ray emissivity* $q_\gamma(E_\gamma)$ is the production rate of photons of energy E_γ per unit time, volume, and energy. For hadronic γ rays from neutral-pion decay, q_γ follows by folding the π^0 emissivity with the decay spectrum: for each pion energy E_π we add the contribution of the two photons produced in the decay and then sum over all pion energies that can kinematically yield E_γ.

Using the single-pion decay spectrum derived above, Eq. (2.174), we can write

$$q_\gamma(E_\gamma) = \int_{E_{\pi,\min}(E_\gamma)}^{\infty} \mathrm{d}E_\pi \; q_{\pi^0}(E_\pi) \left(\frac{\mathrm{d}N_\gamma}{\mathrm{d}E_\gamma}\right)_{\pi^0\to 2\gamma}, \tag{2.178}$$

where $E_{\pi,\min}(E_\gamma)$ is the minimum pion energy required to produce a photon of energy E_γ. To determine $E_{\pi,\min}$, we recall that for a pion of fixed lab-frame energy E_π the two photon energies lie in the interval $[E_{\gamma,\min}, E_{\gamma,\max}]$, with endpoints

$$E_{\gamma,\min\,/\,\max} = \frac{E_\pi \mp p_\pi}{2}, \qquad p_\pi = \sqrt{E_\pi^2 - m_\pi^2}. \tag{2.179}$$

From these expressions one immediately obtains two useful identities:

$$E_{\gamma,\min} + E_{\gamma,\max} = E_\pi, \qquad E_{\gamma,\min}\, E_{\gamma,\max} = \frac{E_\pi^2 - p_\pi^2}{4} = \frac{m_\pi^2}{4}. \tag{2.180}$$

Eliminating $E_{\gamma,\min}$ in favor of $E_{\gamma,\max}$ using $E_{\gamma,\min} = m_\pi^2/(4E_{\gamma,\max})$ gives

$$E_\pi = E_{\gamma,\min} + E_{\gamma,\max} = E_{\gamma,\max} + \frac{m_\pi^2}{4E_{\gamma,\max}}. \tag{2.181}$$

Now fix an observed photon energy E_γ. The minimum parent energy is reached when that photon lies at one endpoint of the decay box. For $E_\gamma \geq m_\pi/2$ it is the forward-emission endpoint, $E_{\gamma,\max} = E_\gamma$; for $E_\gamma \leq m_\pi/2$ it is the backward-emission endpoint, $E_{\gamma,\min} = E_\gamma$. Either branch gives

$$E_{\pi,\min}(E_\gamma) = E_\gamma + \frac{m_\pi^2}{4E_\gamma}. \tag{2.182}$$

This expression is unchanged by $E_\gamma \leftrightarrow m_\pi^2/(4E_\gamma)$, which proves the logarithmic symmetry of the exact photon-number emissivity in Eq. (2.183).

Using the box-like photon spectrum from a single π^0 decay, Eq. (2.174), the γ-ray emissivity becomes

$$q_\gamma(E_\gamma) = 2\int_{E_{\pi,\min}(E_\gamma)}^{\infty} \mathrm{d}E_\pi\, q_{\pi^0}(E_\pi)\, \frac{1}{\sqrt{E_\pi^2 - m_\pi^2}} = 2\int_{E_\gamma + \frac{m_\pi^2}{4E_\gamma}}^{\infty} \mathrm{d}E_\pi\, \frac{q_{\pi^0}(E_\pi)}{\sqrt{E_\pi^2 - m_\pi^2}}\,. \tag{2.183}$$

In the ultra-relativistic regime, $E_\gamma \gg m_\pi$ and $E_\pi \gg m_\pi$, one may approximate $E_{\pi,\min} \simeq E_\gamma$ and $\sqrt{E_\pi^2 - m_\pi^2} \simeq E_\pi$, yielding the simpler scaling form

$$q_\gamma(E_\gamma) \simeq 2\int_{E_\gamma}^{\infty} \frac{\mathrm{d}E_\pi}{E_\pi}\, q_{\pi^0}(E_\pi). \tag{2.184}$$

At high energies the γ-ray spectrum closely tracks the pion spectrum. For a power-law pion emissivity, $q_{\pi^0}(E_\pi) \propto E_\pi^{-\alpha_\pi}$, one finds

$$q_\gamma(E_\gamma) \propto E_\gamma^{-\alpha_\pi}. \tag{2.185}$$

Since in the same regime the pion emissivity largely mirrors the parent proton spectrum, $N_p(E_p) \propto E_p^{-\alpha_p}$, it follows that hadronic gamma rays inherit, to a very good approximation, the proton spectral index:

$$q_\gamma(E_\gamma) \propto E_\gamma^{-\alpha_p}. \tag{2.186}$$

Finally, we can estimate a characteristic photon energy using the effective pion energy fraction K_π introduced in Eq. (2.163). The mean photon energy from a pion of energy E_π is $E_\pi/2$, thus

$$E_\gamma \sim K_\gamma\, E_p, \qquad K_\gamma \simeq \frac{K_\pi}{2} \sim 10\%. \tag{2.187}$$

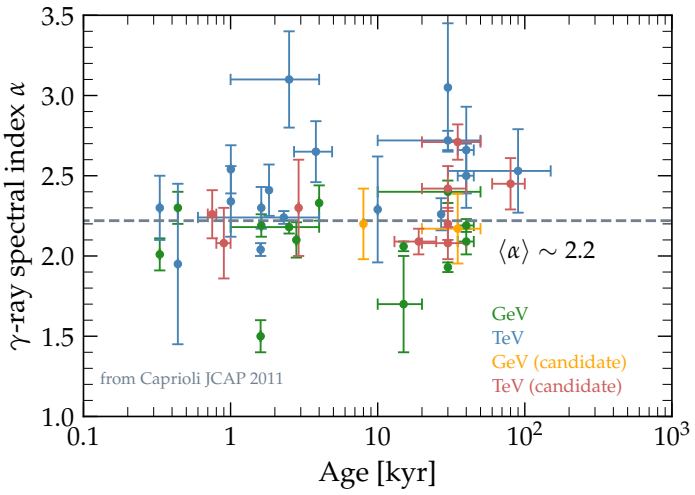


Figure 20: Observed GeV/TeV γ-ray spectral indices of supernova remnants of different ages. Adapted from [62].

This behavior is in sharp contrast with inverse Compton (IC) scattering by relativistic electrons. For an electron spectrum $N_e(E_e) \propto E_e^{-p}$, the IC γ-ray spectrum in the Thomson regime typically scales as

$$q_\gamma^{\mathrm{IC}}(E_\gamma) \propto E_\gamma^{-(p+1)/2},$$

i.e. with a different relation between the photon index and the index of the parent particles (see Section 2.8).

Supernova remnants provide a relevant example. The sample summarized in Fig. 20 contains many γ-ray–bright remnants with photon indices steeper than 2, typically $\Gamma_\gamma \simeq 2.2$–3. Within a hadronic interpretation, such steep γ-ray spectra directly point to parent hadrons with comparably steep distributions, $N_p(E) \propto E^{-\alpha_p}$ with $\alpha_p \gtrsim 2$, since $q_\gamma(E_\gamma) \propto E_\gamma^{-\alpha_p}$ to good approximation. By contrast, IC emission from an electron spectrum $N_e(E) \propto E^{-p}$ is expected to be harder; in the Thomson regime one finds approximately $q_\gamma(E_\gamma) \propto E_\gamma^{-1.5}$ for $p = 2$. The steep photon spectra in this sample are therefore compatible with hadronic parent spectra of comparable index, but are not by themselves evidence for pion production.

Although the two mechanisms relate the parent-particle and photon indices differently, this difference alone does not conclusively distinguish leptonic from hadronic emission: proton and electron injection indices may differ, radiative cooling can steepen the electron distribution, and Klein–Nishina effects can modify the IC photon index. Spectral information is therefore most useful when combined with spatial correlations between the γ-ray emission and target gas, constraints on the target radiation fields, and the pion-decay turnover discussed before [63].

2.16 Diffuse Galactic γ-ray emission

We now use the pion-decay formalism developed in Section 2.15 to address a simple observational question: if the cosmic rays measured near Earth illuminate the gas of the Galactic disk, what diffuse γ-ray intensity should they produce?

The physical chain is simple. Cosmic-ray protons and nuclei collide with interstellar gas, share part of their energy with neutral pions, and the pions decay into photons. Integrating the resulting emissivity along a line of sight produces the gas-correlated component of the diffuse sky mapped by *Fermi*-LAT. Inverse Compton scattering, bremsstrahlung, unresolved sources, and nearly isotropic backgrounds also contribute. Here we isolate the hadronic component and ask whether its expected normalization is comparable to the observed diffuse intensity. This order-of-magnitude calculation cannot by itself determine the precise balance among the different components, but it can tell us whether pion decay is naturally important or necessarily subdominant.

The calculation is simple enough to perform by hand, yet transparent enough to reveal the essential scaling and to provide a useful benchmark for more detailed numerical models.

Pion production in interstellar gas. We move beyond the delta-functional approximation by retaining the full distribution of pion energies. For cosmic-ray protons interacting with a pure hydrogen target, the pion emissivity is

$$q_\pi(E_\pi) = c\, n_{\rm H} \int_{E_p^{\rm th}}^{\infty} {\rm d}E_p\; N_p(E_p)\, \frac{{\rm d}\sigma_{pp\to\pi X}}{{\rm d}E_\pi}(E_p, E_\pi), \tag{2.188}$$

Here π denotes a neutral pion, N_p is the differential proton density, and ${\rm d}\sigma_{pp\to\pi X}/{\rm d}E_\pi$ is the inclusive differential cross section. Nuclear projectiles and helium in the gas can be included through additional terms or an approximate nuclear enhancement factor.

Energy conservation guarantees that the differential cross section vanishes for $E_p < E_\pi$. We may therefore replace the threshold by E_π as the formal lower limit, keeping the cross section zero below the physical production threshold:

$$q_\pi(E_\pi) = c\, n_{\rm H} \int_{E_\pi}^{\infty} {\rm d}E_p\; N_p(E_p)\, \frac{{\rm d}\sigma_{pp\to\pi X}}{{\rm d}E_\pi}(E_p, E_\pi). \tag{2.189}$$

Cosmic-ray measurements are usually reported as a differential intensity rather than a number density. For an isotropic, ultra-relativistic proton population,

$$I_p(E_p) = \frac{c}{4\pi}\, N_p(E_p), \tag{2.190}$$

so that

$$q_\pi(E_\pi) = 4\pi\, n_{\rm H} \int_{E_\pi}^{\infty} {\rm d}E_p\; I_p(E_p)\, \frac{{\rm d}\sigma_{pp\to\pi X}}{{\rm d}E_\pi}(E_p, E_\pi). \tag{2.191}$$

This is the natural starting point for the calculation: I_p supplies the projectiles, while the differential cross section describes how each collision distributes energy among the pions.

The high-energy scaling kernel. It is useful to separate the probability of an inelastic collision from the pion yield per collision:

$$\frac{\mathrm{d}\sigma}{\mathrm{d}E_\pi}(E_p, E_\pi) = \sigma_{\mathrm{inel}}(E_p)\,\frac{f(E_p, E_\pi)}{E_\pi}, \tag{2.192}$$

where f is the dimensionless inclusive yield. Defining the pion energy fraction

$$x \equiv \frac{E_\pi}{E_p} \in [0,1], \tag{2.193}$$

the yield per inelastic event can be written in either variable:

$$f(E_p, E_\pi) = E_\pi \frac{\mathrm{d}N_\pi}{\mathrm{d}E_\pi} = x\frac{\mathrm{d}N_\pi}{\mathrm{d}x}.$$

At this stage the separation simply defines f. The approximation enters in the high-energy regime, where σ_{inel} changes slowly and approximate scaling suggests that the shape of the yield depends mainly on x, rather than independently on E_p and E_π. We then write

$$f(E_p, E_\pi) \simeq f(x).$$

Rather than assume a delta-functional energy partition, here we calibrate this kernel using a hadronic Monte Carlo event generator, which simulates individual collisions and the secondary particles they produce. Specifically, we use SIBYLL 2.3e [64] to generate 10^5 inelastic fixed-target pp collisions at $E_p = 1$ PeV. At this high projectile energy, the simulated pion yield provides a useful approximation to the asymptotic shape.

An endpoint-suppressed log-parabola provides a compact fit to the simulated neutral-pion distribution,

$$f_{\mathrm{Sib}}(x) = A\exp\left[-a\ln^2\left(\frac{x}{x_0}\right)\right](1-x)^k, \qquad \begin{array}{ll} A = 1.97, & x_0 = 2.80\times10^{-4}, \\ a = 0.0411, & k = 2.17. \end{array} \tag{2.194}$$

Here x_0 sets the broad turnover, the logarithmic term describes the two sides of the distribution, and $(1-x)^k$ enforces the high-energy endpoint. The coefficients follow from minimizing residuals in $\log f$ over $10^{-5} \le x \le 0.8$. As Fig. 21 shows, the rms residual is 0.016 dex, corresponding to a typical multiplicative deviation of about 4%.

Equation (2.194) is then a phenomenological representation of SIBYLL at this energy. Its purpose here is to replace the delta-functional approximation by a realistic broad yield while remaining simple enough to integrate.

For an order-of-magnitude estimate, we freeze the slowly varying inelastic cross section at

$$\sigma_{\mathrm{inel}}(E_p) \simeq \sigma_0 \simeq 30\ \mathrm{mb}, \tag{2.195}$$

so that $\mathrm{d}\sigma/\mathrm{d}E_\pi \simeq \sigma_0 f(x)/E_\pi$. Over a limited energy interval, we also approximate the proton intensity by a power law,

$$I_p(E_p) = I_{10}\left(\frac{E_p}{E_{10}}\right)^{-\alpha_p}, \tag{2.196}$$

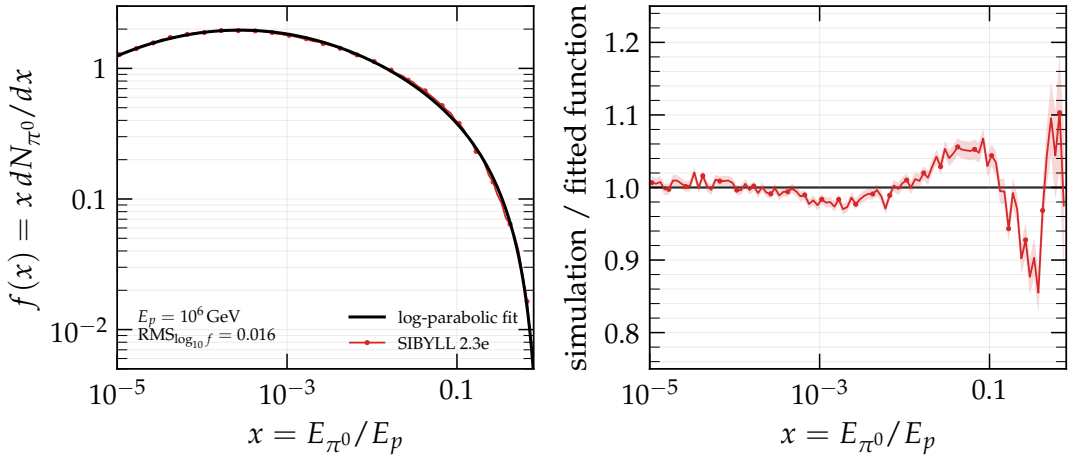


Figure 21: Inclusive neutral-pion scaling function in fixed-target pp collisions at $E_p = 10^6$ GeV. The simulated distribution uses SIBYLL 2.3e [64] with 10^5 inelastic events. Points show $x\,\mathrm{d}N_{\pi^0}/\mathrm{d}x$, and the shaded band is the Poisson uncertainty of each bin. The black curve is the fitted function in Eq. (2.194), obtained using only bins with at least 30 pions; the right panel shows the simulation-to-fit ratio.

where $E_{10} = 10$ GeV. The normalization I_{10} denotes $I_p(E_{10})$.

Inserting Eqs. (2.192) and (2.196) into Eq. (2.191) gives

$$q_\pi(E_\pi) \simeq 4\pi\, n_{\mathrm{H}}\, I_{10}\, \sigma_0 \int_{E_\pi}^{\infty} \mathrm{d}E_p \, \left(\frac{E_p}{E_{10}}\right)^{-\alpha_p} \frac{f(E_\pi/E_p)}{E_\pi}. \tag{2.197}$$

The integral becomes particularly transparent after changing variables to $x = E_\pi/E_p$, for which $E_p = E_\pi/x$ and $\mathrm{d}E_p = -E_\pi x^{-2}\mathrm{d}x$. The result is

$$q_\pi(E_\pi) = 4\pi\, n_{\mathrm{H}}\, I_{10}\, \sigma_0 \left(\frac{E_\pi}{E_{10}}\right)^{-\alpha_p} \underbrace{\int_0^1 \mathrm{d}x\, x^{\alpha_p - 2}\, f(x)}_{Y(\alpha_p)}. \tag{2.198}$$

The dimensionless factor

$$Y(\alpha_p) \equiv \int_0^1 \mathrm{d}x\, x^{\alpha_p - 2}\, f(x) \tag{2.199}$$

is the *spectrum-weighted yield*, or spectrum-weighted moment. It packages the secondary-energy distribution into one number for a given proton index. For $\alpha_p = 2.7$, the fitted kernel gives $Y(2.7) \simeq 3.18 \times 10^{-2}$.

This is the reward for the change of variables: all details of the broad pion distribution have collapsed into $Y(\alpha_p)$. Within the scaling approximation, Eq. (2.198) then shows immediately that the pion emissivity inherits the proton power-law index. The energy dependence of σ_{inel} and scaling violations produce only mild corrections to this simple conclusion.

From pions to photons. The final local step is the neutral-pion decay derived in the previous section. Substituting Eq. (2.198) into Eq. (2.183), and taking the high-energy limit $E_\gamma \gg m_\pi$, gives

$$q_\gamma(E_\gamma) \simeq 8\pi\, n_{\rm H}\, \sigma_0\, Y(\alpha_p)\, I_{10} \int_{E_\gamma}^{\infty} \frac{{\rm d}E_\pi}{E_\pi} \left(\frac{E_\pi}{E_{10}}\right)^{-\alpha_p}. \tag{2.200}$$

The remaining integral is elementary, yielding

$$q_\gamma(E_\gamma) = \frac{8\pi}{\alpha_p}\, n_{\rm H}\, \sigma_0\, Y(\alpha_p)\, I_{10} \left(\frac{E_\gamma}{E_{10}}\right)^{-\alpha_p}. \tag{2.201}$$

Equivalently, in the commonly used $E^2 q(E)$ representation,

$$E_\gamma^2\, q_\gamma(E_\gamma) = \frac{8\pi}{\alpha_p}\, n_{\rm H}\, \sigma_0\, Y(\alpha_p)\, E_{10}^2\, I_{10} \left(\frac{E_\gamma}{E_{10}}\right)^{2-\alpha_p}. \tag{2.202}$$

Two conclusions are already visible. First, within the high-energy and scaling approximations, the hadronic γ-ray spectrum inherits the proton index, $q_\gamma \propto E_\gamma^{-\alpha_p}$. Second, its normalization is the product of four physical ingredients: the gas density, the inelastic cross section, the spectrum-weighted pion yield, and the cosmic-ray intensity.

From volume emissivity to sky intensity. A γ-ray telescope measures intensity on the sky rather than local volume emissivity. Along a direction $\hat{\boldsymbol{n}}$, the two are related by

$$I_\gamma(E_\gamma, \hat{\boldsymbol{n}}) = \frac{1}{4\pi} \int_{\rm los} {\rm d}s\; q_\gamma(E_\gamma, \boldsymbol{r}). \tag{2.203}$$

Introducing the hydrogen column density

$$N_{\rm H} \equiv \int_{\rm los} {\rm d}s\; n_{\rm H}, \tag{2.204}$$

and assuming that the proton spectrum is approximately uniform along the line of sight, Eq. (2.202) becomes

$$E_\gamma^2\, I_\gamma(E_\gamma) = \frac{2}{\alpha_p}\, N_{\rm H}\, \sigma_0\, Y(\alpha_p)\, E_{10}^2\, I_p(E_{10}) \left(\frac{E_\gamma}{E_{10}}\right)^{2-\alpha_p}. \tag{2.205}$$

This is the central observational relation. The volume density $n_{\rm H}$ has become the measurable gas column $N_{\rm H}$, and $N_{\rm H}\sigma_0$ plays the role of an effective interaction depth. The γ-ray sky therefore traces the cosmic-ray population weighted by the amount of target material along each line of sight.

A Galactic benchmark. For a representative line of sight through the inner Galactic plane, take $E_\gamma = 1$ GeV and adopt

- a characteristic hydrogen column density $N_{\rm H} \simeq 10^{22}\ {\rm cm}^{-2}$ [65];
- the locally measured proton intensity, $E_{10}^2 I_p(E_{10}) \simeq 0.2\ {\rm GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}$ [66];

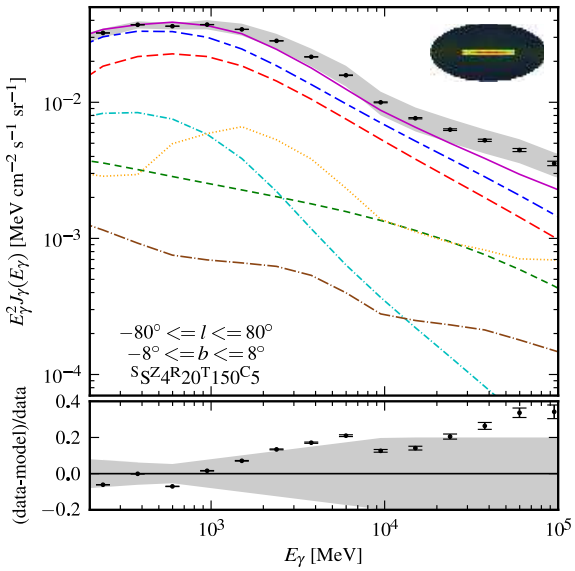


Figure 22: Diffuse γ-ray spectrum from the inner Galaxy measured by *Fermi*-LAT. Data points show the extracted spectrum; error bars are statistical, and the gray band represents the systematic uncertainty associated with the LAT effective area. The inset shows the photon counts in the region of interest. Representative GALPROP curves separate the principal emission components: hadronic π^0 decay, inverse Compton scattering, and bremsstrahlung. Adapted from [68].

- a spectral index $\alpha_p \simeq 2.7$, for which the spectrum-weighted yield is $Y(\alpha_p) \simeq 3.2 \times 10^{-2}$.

The calculation so far includes only cosmic-ray protons colliding with interstellar hydrogen. Helium and heavier nuclei, both in the cosmic rays and in the interstellar gas, contribute additional pion production. Their contribution is usually accounted for by a *nuclear enhancement factor* multiplying the proton–proton emissivity. Adopting $\sigma_0 \simeq 30$ mb and an enhancement factor of 1.8 [67], and multiplying the right-hand side of Eq. (2.205) by this factor, we obtain

$$E_\gamma^2 I_\gamma(E_\gamma = 1\ \mathrm{GeV}) \simeq 10^{-5}\ \mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}. \tag{2.206}$$

This back-of-the-envelope estimate already gives an intensity comparable to that measured along the Galactic plane near 1 GeV (Fig. 22). A quantitative comparison with the data requires more detailed modelling of the spatial distributions of cosmic rays and atomic and molecular gas, together with the contributions from inverse Compton scattering, bremsstrahlung, unresolved sources, and instrumental backgrounds [10].

The significance of this estimate also lies in its scaling: I_γ is proportional to the gas column and to the cosmic-ray intensity sampled by that gas. The hadronic γ-ray emission therefore traces both the cosmic-ray distribution and the interstellar gas along the line of sight. In practice, much of its small-scale structure reflects variations in the gas column, since the cosmic-ray population is comparatively smooth on these scales.

This is the modern version of the basic argument made by Hayakawa [69, 70]: a measured cosmic-ray flux, folded with a gas column and a hadronic yield, predicts a diffuse γ-ray

signal on the scale actually observed. The agreement does not prove that the cosmic-ray density is perfectly uniform throughout the Galaxy. It does show, however, that the GeV–TeV particles measured near the Sun are not merely a local population: cosmic rays permeate the Galactic disk and form a large-scale *cosmic-ray sea*. How diffusion and confinement establish that population is the subject of **??**.

2.17 Nuclei–photon interactions

Energetic sources are often embedded in intense radiation fields: radiation from accretion disks and broad-line regions in AGN, prompt and afterglow emission in GRBs, or starlight and dust emission in star-forming environments. High-energy cosmic-ray protons and nuclei can interact with these photons, losing energy, producing secondary particles, or changing species.

The relevant channels are, schematically,

- *photopion production*, $N + \gamma \to N + \pi$, where $N = p, n$ denotes a nucleon; for nuclei the process is often described per nucleon;
- *photopair production* (Bethe–Heitler), $N + \gamma \to N + e^+e^-$, with $N = p$ here; the extension to nuclei will be motivated later;
- *photodisintegration* of nuclei, $A + \gamma \to A' + N$, and more general fragmentation channels.

Although characteristic cross sections for photopion and photopair production on protons are smaller than inelastic pp cross sections, the photon density near a compact source can far exceed the gas density. Photonuclear interactions can then dominate the energy losses and secondary production, shaping the spectrum and composition of the escaping cosmic rays and the accompanying γ-ray and neutrino signals (see Section 2.18).

A typical example is the propagation of ultra-high-energy cosmic rays (UHECRs, roughly $E \gtrsim 10^{18}$ eV) through intergalactic space. The Universe is permeated by the CMB, with present-day photon number density

$$n_{\mathrm{CMB},0} \simeq 411 \ \mathrm{cm}^{-3},$$

compared with the mean baryon number density today,

$$n_{b,0} \sim \text{few} \times 10^{-7} \ \mathrm{cm}^{-3}.$$

This large target density, together with the boost of photon energies into the projectile rest frame, makes interactions with diffuse radiation backgrounds the dominant non-adiabatic propagation effect at the highest energies.

Recall that, in an expanding Universe, these interactions must compete with *adiabatic* losses due to cosmological redshift. For an ultra-relativistic particle, the energy-loss rate and instantaneous loss timescale are:

$$\left(\frac{dE}{dt}\right)_{\mathrm{ad}} = -H(z)\,E, \qquad \tau_{\mathrm{ad}}(z) \equiv \frac{E}{|dE/dt|_{\mathrm{ad}}} = \frac{1}{H(z)}. \tag{2.207}$$

For a flat ΛCDM cosmology, neglecting the radiation contribution to the expansion rate, the Hubble–Lemaître rate is

$$H(z) = H_0\sqrt{\Omega_m(1+z)^3+\Omega_\Lambda}, \tag{2.208}$$

with $\Omega_m + \Omega_\Lambda = 1$; representative values are $\Omega_m \simeq 0.3$ and $\Omega_\Lambda \simeq 0.7$.

Taking $H_0 \simeq 70\ \mathrm{km\,s^{-1}\,Mpc^{-1}}$ gives a present-day loss timescale $\tau_{\mathrm{ad},0} = H_0^{-1} \simeq 14$ Gyr and a corresponding loss length

$$\lambda_{\mathrm{ad},0} \equiv c\,\tau_{\mathrm{ad},0} = \frac{c}{H_0} \simeq 4.3\ \mathrm{Gpc}. \tag{2.209}$$

As we will detail later, the photonuclear processes introduced here can have interaction or energy-loss times comparable to, or shorter than, the propagation and adiabatic timescales, thereby shaping the propagated UHECR spectrum and composition. For protons, familiar examples are the *GZK suppression* from photopion production [71, 72] and the Bethe–Heitler *pair-production dip* [73].

To quantify these effects, we first derive the interaction and energy-loss rates in a general photon field. We then apply the results to each channel, discussing rest-frame thresholds, cross sections, inelasticities, and characteristic propagation lengths.

Interaction rates and energy-loss lengths. For a relativistic nucleus (or proton) with Lorentz factor Γ moving through a photon bath, the interaction rate is obtained from "flux × cross section", integrated over the target photon phase space. The relevant kinematic invariant is

$$s = m_A^2c^4 + 2\,E_A\,\varepsilon\,(1-\beta\cos\theta) = m_A^2c^4 + 2\,m_Ac^2\,\varepsilon'\,,$$

where m_A is the nuclear mass, $E_A = \Gamma m_A c^2$ its lab-frame energy, and θ the angle between the projectile and photon momenta. The photon energy ε in the lab frame is related to its energy in the nucleus rest frame by

$$\varepsilon' = \Gamma\,\varepsilon\,(1-\beta\cos\theta).$$

The cross section can therefore be tabulated as $\sigma(s)$, or equivalently as $\sigma(\varepsilon')$, with the threshold enforced by its vanishing for $\varepsilon' < \varepsilon'_{\mathrm{th}}$.

Let $\mathrm{d}n_\gamma/(\mathrm{d}\varepsilon\,\mathrm{d}\Omega)$ be the differential photon density per unit energy and solid angle. The interaction rate per projectile is

$$\frac{\mathrm{d}N_{\mathrm{int}}}{\mathrm{d}t} = c\int \mathrm{d}\varepsilon \int \mathrm{d}\Omega\,(1-\beta\cos\theta)\,\frac{\mathrm{d}n_\gamma}{\mathrm{d}\varepsilon\,\mathrm{d}\Omega}\,\sigma\big(s(\varepsilon,\theta)\big)\,. \tag{2.210}$$

The factor $c(1-\beta\cos\theta)$ is the relative flux of the projectile and photon beams in the lab frame.

Isotropic photon field. If the background is isotropic, $\mathrm{d}n_\gamma/(\mathrm{d}\varepsilon\,\mathrm{d}\Omega) = n_\gamma(\varepsilon)/(4\pi)$, and with $\mu \equiv \cos\theta$ one finds

$$\frac{\mathrm{d}N_{\mathrm{int}}}{\mathrm{d}t} = c\int \mathrm{d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{2}\int_{-1}^{1}\mathrm{d}\mu\,(1-\beta\mu)\,\sigma(\varepsilon')\,, \qquad \varepsilon' = \Gamma\varepsilon(1-\beta\mu). \tag{2.211}$$

Now change integration variable from μ to ε' at fixed ε. Since

$$\varepsilon' = \Gamma\varepsilon(1-\beta\mu) \quad \Longrightarrow \quad \mathrm{d}\mu = -\frac{\mathrm{d}\varepsilon'}{\Gamma\varepsilon\beta}, \qquad (1-\beta\mu) = \frac{\varepsilon'}{\Gamma\varepsilon},$$

the angular integral becomes

$$\int_{-1}^{1} \mathrm{d}\mu\,(1-\beta\mu)\,\sigma(\varepsilon') = \frac{1}{\Gamma^2\varepsilon^2\,\beta}\int_{\varepsilon'_{\min}}^{\varepsilon'_{\max}} \mathrm{d}\varepsilon'\,\varepsilon'\,\sigma(\varepsilon')\,, \tag{2.212}$$

with limits set by $\mu \in [-1,1]$:

$$\varepsilon'_{\min} = \Gamma\varepsilon(1-\beta), \qquad \varepsilon'_{\max} = \Gamma\varepsilon(1+\beta).$$

Inserting Eq. (2.212) into Eq. (2.211) yields

$$\frac{\mathrm{d}N_{\mathrm{int}}}{\mathrm{d}t} = \frac{c}{2\Gamma^2\beta}\int_0^\infty \mathrm{d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{\varepsilon^2}\int_{\varepsilon'_{\min}}^{\varepsilon'_{\max}} \mathrm{d}\varepsilon'\,\varepsilon'\,\sigma(\varepsilon')\,, \tag{2.213}$$

Here the cross section vanishes below threshold. Equivalently, the lower limit can be replaced by $\max(\varepsilon'_{\min}, \varepsilon'_{\mathrm{th}})$, with the inner integral set to zero if this exceeds $\varepsilon'_{\max}$.

Ultra-relativistic limit. For $\Gamma \gg 1$ one has $\beta \simeq 1$, $\varepsilon'_{\min} \simeq \varepsilon/(2\Gamma)$ and $\varepsilon'_{\max} \simeq 2\Gamma\varepsilon$. For soft target photons, $\varepsilon/(2\Gamma)$ is negligible compared with the rest-frame threshold, so this lower rest-frame energy limit can be set to zero. Swapping the order of integration then gives the standard compact form

$$\frac{\mathrm{d}N_{\mathrm{int}}}{\mathrm{d}t} \simeq \frac{c}{2\Gamma^2}\int_{\varepsilon'_{\mathrm{th}}}^\infty \mathrm{d}\varepsilon'\,\sigma(\varepsilon')\,\varepsilon'\int_{\varepsilon'/(2\Gamma)}^\infty \mathrm{d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{\varepsilon^2}\,. \tag{2.214}$$

From rates to interaction and energy-loss lengths. The interaction time and mean free path are

$$\tau_{\mathrm{int}} = \left(\frac{\mathrm{d}N_{\mathrm{int}}}{\mathrm{d}t}\right)^{-1}, \qquad \lambda_{\mathrm{int}} \simeq c\,\tau_{\mathrm{int}} \quad (\beta \simeq 1).$$

Let $\kappa(\varepsilon')$ be the mean fraction of the projectile energy removed per interaction, called the *inelasticity*. The mean energy-loss rate is obtained by $\sigma \to \kappa\,\sigma$ in Eq. (2.214):

$$-\frac{1}{E}\frac{\mathrm{d}E}{\mathrm{d}t} \simeq \frac{c}{2\Gamma^2}\int_{\varepsilon'_{\mathrm{th}}}^\infty \mathrm{d}\varepsilon'\,\kappa(\varepsilon')\,\sigma(\varepsilon')\,\varepsilon'\int_{\varepsilon'/(2\Gamma)}^\infty \mathrm{d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{\varepsilon^2}\,, \tag{2.215}$$

The corresponding loss time and loss length are

$$\tau_{\mathrm{loss}} \equiv \frac{E}{|\mathrm{d}E/\mathrm{d}t|}, \qquad \ell_{\mathrm{loss}} \simeq c\,\tau_{\mathrm{loss}}.$$

Photopion production

The leading inelastic channel for sufficiently energetic protons in a radiation field is *photopion production*, in which a target photon excites the proton to a baryonic resonance that subsequently decays with pion emission. Close to threshold the interaction is dominated by the Δ-resonance,

$$p + \gamma \to \Delta^+ \to \begin{cases} p + \pi^0, \\ n + \pi^+, \end{cases} \tag{2.216}$$

with an intrinsic lifetime $\sim 10^{-23}\,\mathrm{s}$. Neutral pions promptly decay into $\gamma\gamma$, while charged pions decay through $\pi^\pm \to \mu^\pm \nu_\mu$ followed by $\mu^\pm \to e^\pm \nu_e \nu_\mu$, producing high-energy neutrinos and leptons (following the same chain as in Section 2.14).

Thresholds. The kinematics are conveniently expressed in terms of the Mandelstam invariant

$$s = \left(P_p + P_\gamma\right)^2 = m_p^2 c^4 + 2\,E_p\,\varepsilon\,(1 - \beta_p \cos\theta)\,,$$

where ε is the target-photon energy in the lab frame, θ is the collision angle, and $\beta_p = v_p/c$ (so $\beta_p \simeq 1$ for ultra-relativistic protons). Neglecting the small proton–neutron mass difference, single-pion production becomes possible when $s \geq (m_p + m_\pi)^2 c^4$, giving the threshold proton energy (cf. Eq. (A.28))

$$E^\pi_{p,\mathrm{th}} = \frac{[(m_p + m_\pi)^2 - m_p^2]c^4}{2\,\varepsilon\,(1 - \beta_p \cos\theta)} = \frac{(m_\pi^2 + 2 m_p m_\pi)c^4}{2\,\varepsilon\,(1 - \beta_p \cos\theta)}. \tag{2.217}$$

For head-on collisions, $1 - \beta_p \cos\theta \simeq 1 + \beta_p \simeq 2$, this can be written numerically as

$$E^\pi_{p,\mathrm{th}} \simeq 1.1 \times 10^{20}\,\mathrm{eV} \left(\frac{\varepsilon}{6.3 \times 10^{-4}\,\mathrm{eV}}\right)^{-1}. \tag{2.218}$$

This corresponds to a head-on threshold of order 100 EeV for interactions with "typical" CMB photons. In practice, the onset of efficient losses occurs somewhat lower because interactions sample the higher-energy tail of the CMB. The precise onset follows from the convolution of the photon spectrum with the energy-dependent cross section.

Inelasticity. The inelasticity $\kappa_{p\gamma}$ is the fraction of the incoming proton energy lost per interaction. Strictly at threshold, the outgoing proton and pion are both at rest in the CoM frame; therefore, boosting to the lab frame gives

$$E^{\mathrm{lab}}_{p,\mathrm{out}} = \gamma_{\mathrm{CoM}} m_p c^2, \qquad E^{\mathrm{lab}}_{\pi} = \gamma_{\mathrm{CoM}} m_\pi c^2,$$

where γ_{CoM} is the Lorentz factor of the CoM frame relative to the lab.

Energy conservation gives $E^{\mathrm{lab}}_{p,\mathrm{in}} + \varepsilon = E^{\mathrm{lab}}_{p,\mathrm{out}} + E^{\mathrm{lab}}_{\pi}$. Neglecting the incident photon energy ε compared with $E^{\mathrm{lab}}_{p,\mathrm{in}}$, we obtain

$$\kappa^{(\mathrm{th})}_{p\gamma} \equiv \frac{E^{\mathrm{lab}}_{p,\mathrm{in}} - E^{\mathrm{lab}}_{p,\mathrm{out}}}{E^{\mathrm{lab}}_{p,\mathrm{in}}} \simeq \frac{E^{\mathrm{lab}}_{\pi}}{E^{\mathrm{lab}}_{p,\mathrm{out}} + E^{\mathrm{lab}}_{\pi}} = \frac{m_\pi}{m_p + m_\pi} \simeq 0.13\,. \tag{2.219}$$

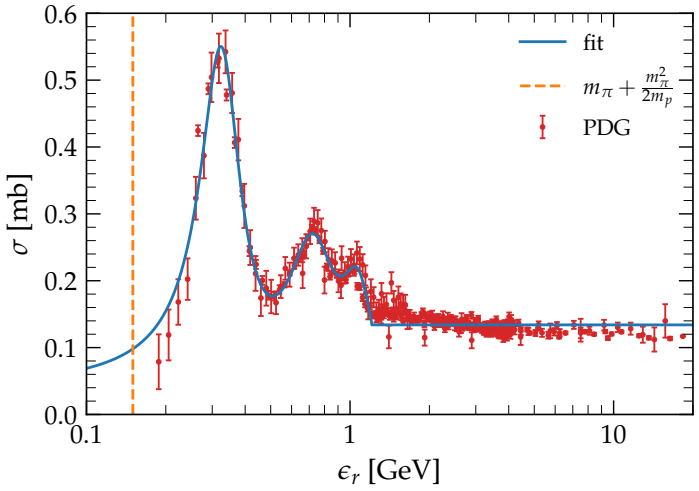


Figure 23: Photopion production cross section $\sigma_{p\gamma}$ as a function of the photon energy in the proton rest frame, ε'. Points show experimental measurements compiled by the PDG [75], while the solid curve is the analytic fit with parameters reported in Table 2. The dashed vertical line marks the kinematic threshold for pion production.

The common boost factor cancels in this ratio.

At higher center-of-mass energies—in particular around the Δ peak and in the multi-pion regime—the produced pions carry additional kinetic energy, and the typical inelasticity rises to $\kappa_{p\gamma} \sim 0.2$–$0.5$ [74]. Consequently, only a handful of interactions are sufficient for a proton to lose a substantial fraction of its initial energy in a dense photon field.

Cross sections. The photopion cross section is most conveniently expressed as a function of the photon energy in the proton rest frame, ε' (see Fig. 23). The Lorentz transformation gives

$$\varepsilon' = \Gamma_p \, \varepsilon \, (1 - \beta_p \cos\theta), \tag{2.220}$$

This variable is directly related to the invariant s through

$$s = m_p^2 c^4 + 2 m_p c^2 \, \varepsilon' \quad \Longrightarrow \quad \varepsilon' = \frac{s - m_p^2 c^4}{2 m_p c^2}. \tag{2.221}$$

Phenomenologically, one often decomposes the total photohadronic cross section as

$$\sigma_{p\gamma}(\varepsilon') = \sigma_{\rm res}(\varepsilon') + \sigma_{\rm dir}(\varepsilon') + \sigma_{\rm mp}(\varepsilon') + \sigma_{\rm diff}(\varepsilon'), \tag{2.222}$$

corresponding to (i) baryonic *resonances*, (ii) *direct* (non-resonant) meson production, (iii) *multi-pion* production at high energies, and (iv) a *diffractive* component.

- *Resonance region.* The dominant feature close to threshold is the $\Delta^+(1232)$ resonance, with $m_\Delta \simeq 1.232\,\mathrm{GeV}$ and width $w_\Delta \simeq 0.115\,\mathrm{GeV}$. Its decay fixes the relative weights of neutral and charged pion channels,

$$p\gamma \to \Delta^+ \to \begin{cases} p\,\pi^0 & (2/3) \\ n\,\pi^+ & (1/3) \end{cases} \tag{2.223}$$

  Within this Δ-resonance approximation, π^0 production is therefore about twice as frequent as π^+ production.

  At somewhat higher ε', additional N^* and Δ^* resonances contribute; for many applications their detailed structure can be absorbed into an effective "resonance bump" on top of a smoother continuum.

- *Direct meson production.* Non-resonant production includes the single-pion channel $p\gamma \to n\pi^+$ and, at higher energies, channels such as $p\gamma \to \Delta^{++}\pi^-$ and $p\gamma \to \Delta^0\pi^+$. The subsequent decay of the outgoing Δ produces an additional pion [74].

- *Multi-pion channel.* At sufficiently large ε' the available phase space opens up and multi-pion production becomes important,

$$p + \gamma \to p + a\,\pi^0 + b\,(\pi^+ + \pi^-)\,, \tag{2.224}$$

  with pion multiplicities a, b increasing with energy. This regime is also the one in which the mean inelasticity rises substantially compared to the resonance region.

Analytic approximation. For many back-of-the-envelope estimates it is sufficient to replace the detailed photopion physics by a small number of *effective* channels that reproduce the gross features of $\sigma_{p\gamma}(\varepsilon')$ and the mean inelasticity:

(i) *Resonance region* ($\varepsilon' \lesssim 1\,\mathrm{GeV}$). The cross section is modeled as the sum of three Breit–Wigner (Lorentzian) resonances, dominated by the $\Delta(1232)$. For the first resonance, we use a mean inelasticity $\kappa_{p\gamma}^{(1\pi)} \simeq 0.2$ and the branching ratio $p\pi^0 : n\pi^+ = 2:1$. The higher resonances are included here only to approximate the total cross section. Each resonance is parametrized as

$$\sigma_{\mathrm{res}}(\varepsilon') = \frac{\mathcal{A}}{\pi}\,\frac{w/2}{(\varepsilon' - \varepsilon'_{\mathrm{r}})^2 + (w/2)^2}\,, \tag{2.225}$$

where $\mathcal{A}$ is the integrated strength of the Lorentzian over the full real axis, $\varepsilon'_{\mathrm{r}}$ the resonance energy, and w its width. The physical cross section is set to zero below the pion-production threshold. The fit parameters for the three resonances are reported in Table 2, and the corresponding fit is shown in Fig. 23.

(ii) *Multi-pion continuum* ($\varepsilon' \gtrsim 1$ GeV). At higher energies the cross section is approximated as a constant,

$$\sigma_{\mathrm{mp}} \simeq 0.13\ \mathrm{mb},$$

with a larger mean inelasticity $\kappa_{p\gamma}^{(\mathrm{mp})} \simeq 0.5$. The transferred energy is often assumed to be shared approximately equally among π^0, π^+, and π^-.

In this simplified picture, each interaction transfers a substantial fraction of the proton energy into pions. Hence, in a sufficiently intense photon field only a few interactions are needed to strongly degrade the proton energy and to efficiently produce secondary γ rays and neutrinos.

Because the energy loss occurs in relatively *large, discrete* steps, fluctuations can be important: the number of interactions undergone, and the stochastic partition of energy among secondaries, can noticeably affect the propagated spectrum on finite path lengths. This is one of the motivations for Monte Carlo treatments of UHECR propagation, which explicitly sample interaction points and final-state kinematics rather than relying solely on a continuous-loss approximation [76, 77].

i	$\mathcal{A}_i$ (mb GeV)	$\varepsilon'_{r,i}$ (GeV)	w_i (GeV)	Peak value (mb)
1	0.109	0.323	0.139	0.50
2	0.127	0.713	0.370	0.22
3	0.100	1.069	0.378	0.17

Table 2: Best-fit Breit–Wigner resonance parameters (Fig. 23).

Interaction length and the GZK suppression. A useful order-of-magnitude estimate of the photopion *interaction length* follows from

$$\tau_{p\gamma}^{-1} \sim c\, n_\gamma\, \sigma_{p\gamma}, \qquad \lambda_{p\gamma} \simeq c\, \tau_{p\gamma} \sim \frac{1}{n_\gamma\, \sigma_{p\gamma}},$$

where we have implicitly absorbed the angular factor $\langle 1 - \cos\theta \rangle = \mathcal{O}(1)$ appropriate to an approximately isotropic photon bath. For the CMB, $n_{\rm CMB} \simeq 411\ {\rm cm}^{-3}$, and using as a reference cross section the peak value of the Δ-resonance one finds

$$\lambda_{p\gamma}^{\pi} \sim \frac{1}{n_{\rm CMB}\, \sigma_{p\gamma}^{(\pi)}} \simeq 2\ {\rm Mpc} \left(\frac{\sigma_{p\gamma}^{\pi}}{0.5\,{\rm mb}} \right)^{-1}, \qquad (\Delta \text{ region}). \tag{2.226}$$

This estimate uses the full CMB density as though every photon sampled the peak cross section. Near threshold only the high-energy tail and a restricted range of angles contribute, so the exact convolution in Eq. (2.214) gives a longer, energy-dependent interaction length.

What ultimately matters for the attenuation of UHECRs is the *energy-loss length*,

$$\ell_{p\gamma}^{\pi} \equiv \left| \frac{1}{E} \frac{dE}{dx} \right|^{-1} \sim \frac{\lambda_{p\gamma}^{\pi}}{\kappa_{p\gamma}^{\pi}} \sim \mathcal{O}(10)\,{\rm Mpc}, \tag{2.227}$$

assuming $\kappa_{p\gamma}^{\pi} \sim 0.2$ (single-pion production in the Δ region). At higher energies multi-pion production increases the mean inelasticity, but the cross section is smaller; to zeroth order their impact on $\ell_{p\gamma}$ is comparable because $\sigma_{p\gamma}^{(\rm mp)} \kappa_{p\gamma}^{(\rm mp)} \sim \sigma_{p\gamma}^{(1\pi)} \kappa_{p\gamma}^{(1\pi)}$.

This simple estimate already captures the key point: above the photopion threshold on the CMB ($E_p \gtrsim$ few $\times 10^{19}$ eV, allowing for the high-energy tail of the Planck spectrum), UHE protons cannot propagate over cosmological distances without substantial energy degradation. The resulting propagation-induced suppression of the proton flux was identified independently by Greisen and by Zatsepin & Kuzmin and is commonly called the GZK suppression [71, 72].

More recent measurements of the cosmic-ray spectrum and composition above 10^{18} eV have challenged this proton-dominated picture. The composition is inferred from the measured distributions of $X_{\max}$, the atmospheric depth at which an air shower reaches its maximum number of particles. Joint fits to these distributions and the spectrum favor a mixed nuclear composition. In this interpretation, the observed high-energy suppression reflects the combined effects of a maximum rigidity attainable at the sources and photodisintegration during propagation, rather than the proton GZK effect alone [78, 79].

Photopair production (Bethe–Heitler)

Below the photopion threshold, the dominant electromagnetic loss channel for UHE protons (and nuclei) in a photon field is *photopair production* (Bethe–Heitler),

$$p + \gamma \to p + e^+ + e^-, \tag{2.228}$$

and, more generally for a nucleus of charge Z and mass number A,

$$A + \gamma \to A + e^+ + e^-. \tag{2.229}$$

In contrast to photopion production—which typically removes a sizable fraction of the projectile energy in a single interaction—Bethe–Heitler pair production has a very small inelasticity, so energy is drained through many interactions and can often be treated as a continuous loss term.

Thresholds. As for photopion production, it is convenient to work with the photon energy in the projectile rest frame,

$$\varepsilon' = \Gamma_p \, \varepsilon \, (1 - \beta_p \cos\theta), \tag{2.230}$$

For a proton, $\Gamma_p = E_p/(m_p c^2)$ and $\beta_p = v_p/c$; ε is the lab-frame photon energy and θ the collision angle. Pair production becomes kinematically allowed when $s = m_p^2 c^4 + 2E_p\varepsilon(1 - \beta_p \cos\theta)$ is equal to $s_{\rm th} = (m_p + 2m_e)^2 c^4$, and one finds the lab-frame threshold (cf. Eq. (A.28)):

$$E^{ee}_{p,\rm th} = \frac{4m_e^2 c^4 + 4 m_e m_p c^4}{2\varepsilon(1 - \beta_p \cos\theta)} \simeq \frac{2 m_e m_p \, c^4}{\varepsilon \, (1 - \beta_p \cos\theta)}. \tag{2.231}$$

For head-on collisions, $1 - \beta_p \cos\theta \simeq 2$, this reduces to

$$E^{ee}_{p,\rm th} \simeq \frac{m_e m_p \, c^4}{\varepsilon} \simeq 7.6 \times 10^{17}\,\text{eV} \left(\frac{\varepsilon}{6.3 \times 10^{-4}\,\text{eV}} \right)^{-1}. \tag{2.232}$$

Thus, for the CMB, Bethe–Heitler losses become relevant already at $E_p \sim 10^{18}$ eV, well below the onset of photopion production. For nuclei the threshold scales as $E_{A,\rm th} \simeq A \, E_{p,\rm th}$, i.e. the threshold *per nucleon* is approximately unchanged.

Inelasticity. As for photopion production, the threshold products are at rest in the CoM frame and acquire the same boost into the lab frame. The fraction of the incoming proton energy transferred to the pair is therefore set by the pair-to-total mass ratio:

$$\kappa_{p\gamma}^{ee} \equiv \frac{\Delta E}{E} \simeq \frac{2m_e}{m_p + 2m_e} \simeq \frac{2m_e}{m_p} \simeq 10^{-3}. \tag{2.233}$$

This threshold value provides an order-of-magnitude estimate away from threshold, up to factors associated with the kinetic energy of the pair and angular effects.

Because $\kappa_{p\gamma}^{ee} \ll 1$, many interactions are required to significantly degrade the projectile energy, motivating a continuous-loss treatment. For nuclei, the same argument gives

$$\kappa_{A\gamma}^{ee} \sim \frac{2m_e}{Am_p}, \tag{2.234}$$

so the inelasticity decreases with increasing mass number.

Cross section. At $\varepsilon' \gg m_e c^2$, the total Bethe–Heitler cross section admits a compact high-energy approximation (for a point-like charge) [80],

$$\sigma_{\rm BH}(\varepsilon') \simeq \frac{7}{6\pi}\,\alpha_{\rm F}\,\sigma_{\rm T}\,Z^2\left[\ln\!\left(\frac{2\varepsilon'}{m_e c^2}\right) - \frac{109}{42}\right], \qquad (\varepsilon' \gg m_e c^2), \tag{2.235}$$

where $\sigma_{\rm T}$ is the Thomson cross section and $\alpha_{\rm F}$ the fine-structure constant. The logarithmic dependence implies that $\sigma_{\rm BH}$ varies only slowly with energy; for protons ($Z = 1$) it is typically of order $\sigma_{\rm BH} \sim 10^{-27}\,\mathrm{cm}^2$ (mb scale) in the range relevant for UHECR propagation, and it scales approximately as Z^2 for nuclei.

For a head-on collision, Eq. (2.230) gives $\varepsilon' \simeq 2\Gamma_p \varepsilon$, with $\Gamma_p = E_p/(m_p c^2)$. Thus, for a typical CMB photon,

$$\varepsilon' \simeq 1.3\ \mathrm{MeV}\left(\frac{E_p}{10^{18}\ \mathrm{eV}}\right)\left(\frac{\varepsilon}{6.3\times 10^{-4}\ \mathrm{eV}}\right).$$

At $E_p = 10^{18}$ eV, this is only slightly above the pair-production threshold $\varepsilon'_{\rm th} \simeq 2m_e c^2 \simeq 1.02$ MeV; the high-energy approximation in Eq. (2.235) is therefore not fully justified at this energy. Nevertheless, we retain $\sigma_{\rm BH} \sim 1$ mb as a rough order-of-magnitude benchmark for the estimates below.

Interaction length and the pair-production dip. A crude mean free path estimate follows from

$$\lambda_{p\gamma}^{ee} \sim \frac{1}{n_\gamma\,\sigma_{\rm BH}} \simeq \mathrm{Mpc}\left(\frac{\sigma_{\rm BH}}{1\,\mathrm{mb}}\right)^{-1}.$$

However, what is relevant for spectral attenuation is the *energy-loss length*,

$$\ell_{p\gamma}^{ee} \equiv \left|\frac{1}{E}\frac{dE}{dx}\right|^{-1} \sim \frac{\lambda_{p\gamma}^{ee}}{\kappa_{p\gamma}^{ee}} \sim \frac{\mathrm{Mpc}}{10^{-3}} \sim 10^3\ \mathrm{Mpc} \sim \mathcal{O}(\mathrm{Gpc}), \tag{2.236}$$

with a non-trivial energy dependence arising from the convolution over the CMB spectrum and the logarithmic growth of $\sigma_{\rm BH}(\varepsilon')$. Thus, for $E_p \gtrsim 10^{18}$ eV, Bethe–Heitler losses on

the CMB can be competitive with adiabatic losses, while at higher energies photopion production takes over (see Eq. (2.227)).

This intermediate-energy loss channel imprints a characteristic, fairly broad *dip* in the propagated *proton* spectrum. In the so-called *dip model*, the observed "ankle" feature in the UHECR spectrum is interpreted primarily as this Bethe–Heitler pair-production dip of extragalactic protons propagating on the CMB, rather than as a purely astrophysical transition feature; the detailed shape then depends on the injection slope, source evolution, and the cosmological propagation history through $\ell^{ee}(E,z)$ [73].

The pure-proton dip model no longer provides a satisfactory description of high-precision measurements of the all-particle UHECR spectrum. Joint fits to the spectrum and composition instead favor models with a mixed nuclear composition, in which the ankle is associated with a transition between source populations or source components [78, 79].

For nuclei, a simple scaling applies *at fixed Lorentz factor*, so that the photon energies in the projectile rest frame are the same. To leading order, the total cross section scales as $\sigma^{ee}_{A\gamma} \propto Z^2$, while the inelasticity scales as $\kappa^{ee}_{A\gamma} \propto 1/A$. Thus

$$\ell^{ee}_{A\gamma}(\Gamma) \simeq \frac{A}{Z^2}\,\ell^{ee}_{p\gamma}(\Gamma).$$

For most stable nuclei, $A \simeq 2Z$, giving $A/Z^2 \simeq 4/A$: heavier nuclei have shorter pair-production loss lengths at the same Γ. This compares a nucleus of total energy $E_A \simeq AE_p$ with a proton of energy E_p, rather than particles of the same total energy.

Photodisintegration

In addition to photopair production and photopion production, an essential process for UHECR nuclei propagating through ambient photon fields is *photodisintegration*,

$$(A,Z) + \gamma \;\rightarrow\; (A-i,\,Z-j) \;+\; \text{fragments}\,, \tag{2.237}$$

in which the nucleus loses one or several nucleons ($i = 1, 2, \dots$) and, more generally, may also eject light clusters (e.g. α particles). Bethe–Heitler losses leave the nuclear species intact and gradually reduce its Lorentz factor. Photodisintegration instead changes the species: to a good approximation the daughter nucleus and emitted fragments inherit the Lorentz factor of the parent, while its mass number decreases.

As in other photo-nuclear processes it is convenient to work with the photon energy ε' in the nucleus rest frame.

Kinematic thresholds. For a nucleus of Lorentz factor Γ interacting with a background photon of lab-frame energy ε at angle θ (between photon and nucleus momenta),

$$\varepsilon' = \Gamma\,\varepsilon\,(1-\beta\cos\theta) \simeq \Gamma\,\varepsilon\,(1-\cos\theta), \tag{2.238}$$

so the *minimum-threshold* configuration is head-on ($\theta = \pi$), for which $\varepsilon' \simeq 2\Gamma\varepsilon$.

A channel with rest-frame threshold $\varepsilon'_{\rm th}$ therefore turns on when

$$\Gamma \gtrsim \Gamma_{\rm th} \simeq \frac{\varepsilon'_{\rm th}}{2\varepsilon} \quad\Longrightarrow\quad \frac{E_{\rm th}}{A} \simeq \Gamma_{\rm th}\, m_N c^2 \simeq m_N c^2\,\frac{\varepsilon'_{\rm th}}{2\varepsilon}, \tag{2.239}$$

so thresholds are most transparently expressed in terms of Γ or, equivalently, the energy per nucleon E/A.

For most nuclei the relevant $\varepsilon'_{\rm th}$ is set by nuclear binding and collective excitations: single-nucleon emission typically requires $\varepsilon'_{\rm th} \sim \mathcal{O}(10)\,\mathrm{MeV}$, so different species have broadly similar *Lorentz-factor* thresholds. Notable exceptions include very weakly bound nuclei (e.g. ^{9}Be, with a neutron separation energy of order MeV) and tightly bound light nuclei (e.g. ^{4}He, whose dominant breakup channels have higher effective thresholds, $\sim 20\,\mathrm{MeV}$).

From Eq. (2.239), one obtains the convenient scaling

$$E_{\rm th} \approx A\,4.7\times10^{15}\,\mathrm{eV}\left(\frac{\varepsilon'_{\rm th}}{10\ \mathrm{MeV}}\right)\left(\frac{\varepsilon}{\mathrm{eV}}\right)^{-1}. \tag{2.240}$$

This immediately explains why the EBL is crucial for nuclei. For CMB photons with $\varepsilon_{\rm CMB} \simeq 6.3\times10^{-4}\,\mathrm{eV}$,

$$\frac{E_{\rm th}}{A} \sim 7.5\times10^{18}\,\mathrm{eV}\left(\frac{\varepsilon'_{\rm th}}{10\ \mathrm{MeV}}\right) \quad\Longrightarrow\quad E_{\rm th} \sim 4\times10^{20}\,\mathrm{eV}\ \ (A=56), \tag{2.241}$$

so photodisintegration on the CMB becomes relevant only at very high energies. In contrast, for EBL photons with $\varepsilon \sim 10^{-2}$–1 eV, the threshold is lower by factors of approximately 15–1500. Photodisintegration of intermediate and heavy nuclei can therefore begin at $E_{\rm th} \sim 10^{17}$–10^{19} eV, well below the CMB threshold.

Inelasticity. A key point is that photodisintegration changes A but leaves the Lorentz factor Γ of the nuclear remnant essentially unchanged. In the initial nucleus rest frame, the kinetic energy imparted to the nuclear center of mass by photon absorption is small:

$$E_{\rm rec} \sim \frac{(\varepsilon')^2}{2\,A\,m_N c^2} \ll m_N c^2. \tag{2.242}$$

In the GDR regime, the fragment kinetic energies are also small compared with their rest energies. The leading fragment therefore retains nearly the parent's Lorentz factor. Boosting back to the lab frame,

$$E \simeq \Gamma\,A\,m_N c^2 \quad\Longrightarrow\quad \frac{\Delta E}{E} \simeq \frac{\Delta A}{A}, \tag{2.243}$$

i.e. the energy per nucleon is approximately conserved,

$$\frac{E}{A} \simeq \Gamma\,m_N c^2 \qquad \text{(nearly constant)}. \tag{2.244}$$

For the emission of i nucleons in a single interaction, the effective inelasticity is therefore of order $\kappa_A \simeq i/A$ (typically small for heavy nuclei), even though the *composition* can change rapidly through a sequence of such events. Here ΔE refers to the energy removed from the leading nuclear fragment. The emitted nucleons retain approximately their share of the parent energy, so this is not an equivalent loss from the total cosmic-ray energy budget.

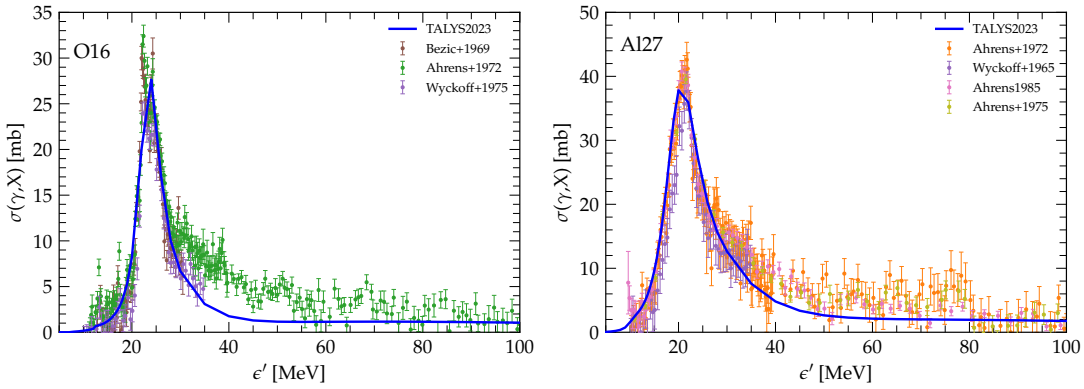


Figure 24: Photodisintegration cross section measurements from the EXFOR database (points) compared with TALYS theoretical evaluations (solid line) for two representative nuclei: ^{16}O (left) and ^{27}Al (right). Adapted from [81].

Cross sections. The dominant breakup channels depend on the photon energy in the nuclear rest frame, ε' (see Fig. 24). A schematic division is:

- *Giant dipole resonance (GDR):* for $\varepsilon' \lesssim 30$ MeV the photoabsorption strength is dominated by the collective electric dipole ($E1$) mode. For most nuclei the cross section peaks at $\varepsilon' \sim 15$–25 MeV, and de-excitation proceeds mainly through the emission of one or two nucleons, $(A,Z)+\gamma \to (A-1,Z)+n$ and/or $(A,Z)+\gamma \to (A-1,Z-1)+p$, with neutron emission typically favored in heavy nuclei. The effective onset is controlled by the lowest particle-emission threshold, typically close to the one-neutron separation energy $S_n \sim 8$ MeV.

- *Multi-nucleon emission / quasi-deuteron (QD) region:* for $30\ \text{MeV} \lesssim \varepsilon' \lesssim 150$ MeV the wavelength of the photon is comparable to nuclear length scales and absorption on correlated nucleon pairs becomes efficient (the "quasi-deuteron" picture). In this regime the total cross section is only weakly energy dependent compared to the sharp GDR peak, and the typical outcome is the ejection of a nucleon pair (e.g. np), often followed by additional nucleons from the ensuing intra-nuclear cascade.

- *Baryonic resonances and photopion production:* at still higher ε', photons excite baryonic resonances on individual nucleons. For a free nucleon, the pion-production threshold is

$$\varepsilon'_{\pi,\text{th}} \simeq m_\pi c^2 \left(1 + \frac{m_\pi}{2m_N}\right) \approx 145\ \text{MeV}, \tag{2.245}$$

  (with the Δ-resonance region at a few hundred MeV). Using Eq. (2.239), photopion production on the CMB requires

$$\frac{E}{A} \gtrsim m_N c^2 \frac{\varepsilon'_{\pi,\text{th}}}{2\varepsilon_{\text{CMB}}} \sim 10^{20}\ \text{eV}, \tag{2.246}$$

  i.e. it turns on only at the very highest energies.

  In nuclei, produced pions may escape (yielding high-energy secondaries via π decay) or be reabsorbed; in the latter case, the deposited energy excites the nucleus and

typically leads to the emission of several nucleons. For propagation through the EBL/CMB, the high threshold generally makes photopion channels subdominant to GDR photodisintegration, except at the extreme end of the spectrum.

Giant dipole resonance (GDR). As discussed above, the GDR produces a broad photoabsorption peak in the nucleus rest frame and is typically the *dominant* contribution to photodisintegration in UHECR propagation through the EBL/CMB. It is therefore useful to make explicit how the GDR is implemented in practical calculations of interaction and attenuation length scales.

The GDR peak energy is typically $\varepsilon'_{\rm GDR} \sim 10$–$20$ MeV (with only a weak dependence on A), and the characteristic width is $\Delta_{\rm GDR} \sim 3$–10 MeV [82, 83]. A convenient normalization is provided by the Thomas–Reiche–Kuhn (TRK) sum rule for the integrated electric-dipole strength. Approximating this strength as concentrated in the GDR gives

$$\Sigma_{\rm TRK} \simeq \int_0^\infty {\rm d}\varepsilon'\, \sigma_{\rm GDR}(\varepsilon') \simeq \frac{2\pi^2 e^2 \hbar}{m_N c}\,\frac{NZ}{A} \approx 60\,\frac{NZ}{A}\ {\rm mb\,MeV}, \tag{2.247}$$

where $N = A - Z$ is the neutron number and m_N the nucleon mass. This is an approximate normalization of the GDR contribution, rather than an integral of all photonuclear channels at arbitrarily high energies [82, 83].

A simple δ-function approximation preserves this integrated strength:

$$\sigma_{\rm GDR}(\varepsilon') \simeq \Sigma_{\rm TRK}\,\delta\big(\varepsilon' - \varepsilon'_{\rm GDR}\big)\,. \tag{2.248}$$

It is often sufficient for back-of-the-envelope estimates.

More refined parametrizations replace Eq. (2.248) by a Lorentzian (or a sum of two Lorentzians for deformed nuclei) and treat the single- and double-nucleon channels explicitly in the GDR region, matched to a roughly flat cross section between $\varepsilon'_1 \simeq 30$ MeV and $\varepsilon'_{\rm max} \simeq 150$ MeV to account for multi-nucleon emission [82, 83].

Interaction lengths with a δ-function GDR. For an isotropic photon background with differential number density $n_\gamma(\varepsilon)$ (per unit energy and volume in the lab frame), the photodisintegration rate of a nucleus of Lorentz factor Γ can be written as (cf. Eq. (2.214))

$$\tau_{A\gamma}^{-1}(\Gamma) = \frac{c}{2\Gamma^2}\int_{\varepsilon'_{\rm th}}^\infty {\rm d}\varepsilon'\,\sigma_A(\varepsilon')\,\varepsilon' \int_{\varepsilon'/(2\Gamma)}^\infty {\rm d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{\varepsilon^2}\,, \tag{2.249}$$

Here $\sigma_A(\varepsilon')$ is the total rest-frame breakup cross section, summed over all channels. Retaining only the GDR and using Eq. (2.248) collapses the ε' integral:

$$\tau_{A\gamma}^{-1}(\Gamma) \simeq \frac{c}{2\Gamma^2}\,\Sigma_{\rm TRK}\,\varepsilon'_{\rm GDR}\int_{\varepsilon'_{\rm GDR}/(2\Gamma)}^\infty {\rm d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{\varepsilon^2}\,. \tag{2.250}$$

The lower limit,

$$\varepsilon_0 \equiv \frac{\varepsilon'_{\rm GDR}}{2\Gamma}, \tag{2.251}$$

is the minimum lab-frame photon energy that can excite the resonance in a head-on collision.

For a blackbody at temperature T, such as the CMB,

$$n_\gamma(\varepsilon) = \frac{8\pi}{(hc)^3}\,\frac{\varepsilon^2}{e^{\varepsilon/(k_{\rm B}T)}-1},$$

so the remaining integral is analytic:

$$\int_{\varepsilon_0}^{\infty} \mathrm{d}\varepsilon\,\frac{n_\gamma(\varepsilon)}{\varepsilon^2} = \frac{8\pi k_{\rm B}T}{(hc)^3}\,\ln\!\left[\frac{1}{1-e^{-\varepsilon_0/(k_{\rm B}T)}}\right]. \tag{2.252}$$

Finally, the interaction length is

$$\lambda_{A\gamma}(\Gamma) \equiv c\,\tau_{A\gamma}(\Gamma). \tag{2.253}$$

Since photodisintegration approximately conserves Γ, one has $E \simeq \Gamma A m_N c^2$ and thus $\Delta E/E \simeq \Delta A/A$ (cf. Eq. (2.243)). If the dominant channel ejects on average i nucleons per interaction, then

$$\left|\frac{1}{E}\frac{\mathrm{d}E}{\mathrm{d}t}\right| \simeq \left|\frac{1}{A}\frac{\mathrm{d}A}{\mathrm{d}t}\right| \simeq \frac{i}{A}\,\tau_{A\gamma}^{-1}(\Gamma), \tag{2.254}$$

so the corresponding effective energy-loss length satisfies

$$\ell_{A\gamma}(E) \equiv \left|\frac{E}{\mathrm{d}E/\mathrm{d}x}\right| \simeq \frac{A}{i}\,\lambda_{A\gamma}(E). \tag{2.255}$$

Thus $\ell_{A\gamma}$ is typically much larger than the mean free path. The distinction is physical: $\lambda_{A\gamma}$ controls how quickly the identity of the nucleus changes, whereas $\ell_{A\gamma}$ tracks the energy retained by the leading fragment.

As an example of a numerical evaluation, we consider iron nuclei interacting with the CMB. We take $A = 56$, $Z = 26$, $\varepsilon'_{\rm GDR} \simeq 20$ MeV and $\Sigma_{\rm TRK} \simeq 60(NZ/A)\,\mathrm{mb\,MeV} \approx 8.4\times10^2\ \mathrm{mb\,MeV}$, with $T_{\rm CMB} = 2.725$ K. Using Eqs. (2.252) and (2.255), we obtain the following order-of-magnitude estimates:

- $E = 3\times10^{20}$ eV:

$$\lambda_{A\gamma}^{(\rm CMB)} \sim 0.7\ \mathrm{Mpc}, \quad\Longrightarrow\quad \ell_{A\gamma} \sim \frac{A}{i}\lambda_{A\gamma} \sim 40\ \mathrm{Mpc} \quad (i \simeq 1).$$

- $E = 10^{21}$ eV:

$$\lambda_{A\gamma}^{(\rm CMB)} \sim 0.04\ \mathrm{Mpc}, \quad\Longrightarrow\quad \ell_{A\gamma} \sim \frac{A}{i}\lambda_{A\gamma} \sim 2\ \mathrm{Mpc} \quad (i \simeq 1).$$

These estimates are meant to capture the scaling and characteristic magnitudes. At the highest energies, these short attenuation lengths indicate that iron nuclei cannot retain their energy and identity over cosmological distances. A quantitative horizon requires following the full nuclear cascade as the energy and mass number evolve. Realistic calculations replace the δ-approximation by the finite GDR width, include the growth of multi-nucleon channels (so $i > 1$ at higher ε'), and adopt detailed photon backgrounds.

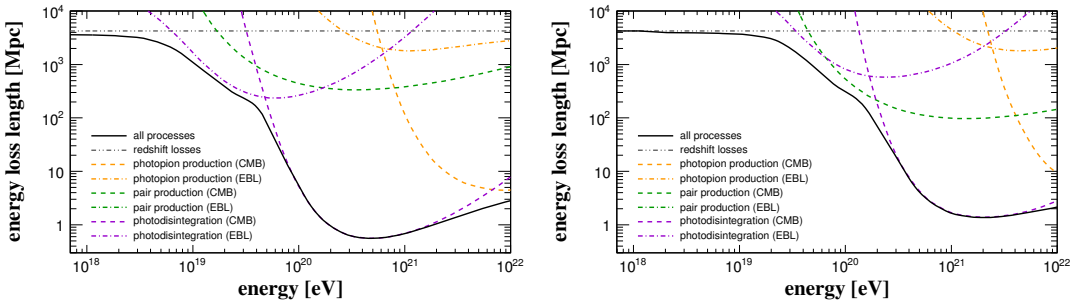


Figure 25: Energy-loss lengths at $z = 0$ for ^{14}N (left) and ^{56}Fe (right). Colors denote the interaction channel and line styles the target photon field, as indicated in the legend. The thick black curve shows the total loss length including all processes. Curves computed with CRPropa 3 using TALYS-based photodisintegration tables [84] and the EBL model of Gilmore *et al.* [35]. Adapted from [86].

More recent approaches use photodisintegration cross sections from modern nuclear-model calculations (e.g. TALYS-based tables [84]) coupled to extended reaction networks that track many isotopes (including unstable isobars) and their decay channels, yielding a more faithful description of the nuclear cascade during propagation. Figure 25 shows the corresponding loss lengths for representative nuclei and illustrates that, despite its simplicity, the δ-function GDR estimate captures the correct *scale* and the rapid energy dependence of the attenuation length.

Finally, significant uncertainties remain in photonuclear cross sections away from well-measured nuclides and in poorly constrained multi-particle emission channels. These propagate into non-negligible systematics in UHECR propagation predictions, and hence in inferences on source spectra and composition from UHECR data [85].

2.18 Neutrinos from cosmic accelerators

High-energy neutrinos are a *golden channel* for identifying hadronic cosmic-ray sources. As already discussed in the context of hadronic γ-ray emission in Section 2.14, neutral and charged pion decays connect γ rays and neutrinos to the same parent population of accelerated protons and nuclei. Gamma rays can also be produced by relativistic electrons; high-energy neutrinos provide a direct signature of hadronic interactions with gas or radiation.

Their weak interactions also allow neutrinos to escape from dense source interiors where γ rays are absorbed, while their electrical neutrality preserves information about the source direction. They can therefore probe otherwise hidden acceleration regions, including the immediate surroundings of the supermassive black holes powering AGN. IceCube's evidence for neutrino emission from the obscured active galaxy NGC 1068 illustrates this possibility [87].

This high-energy observational window opened only recently, with IceCube's discovery of an astrophysical neutrino flux in 2013 [88], and is now being expanded by KM3NeT in the Mediterranean [89]. Building on the production mechanisms introduced above, we now

ask what neutrino energies and flavors to expect, how their emission relates to γ rays, and what the observed flux reveals about the power and efficiency of cosmic accelerators.

Characteristic energies and flavors. Near the photopion threshold, the Δ^+ resonance gives

$$p + \gamma \to \Delta^+ \to \begin{cases} p + \pi^0 & (2/3), \\ n + \pi^+ & (1/3), \end{cases} \tag{2.256}$$

as discussed in Section 2.17. The energy fraction carried by the pion follows directly from two-body kinematics. In the Δ-resonance rest frame,

$$E_\pi^* = \frac{(m_\Delta^2 - m_N^2 + m_\pi^2)c^2}{2m_\Delta}.$$

If the resonance has Lorentz factor γ_Δ, the pion energy in the source frame is

$$E_\pi = \gamma_\Delta \left(E_\pi^* + \beta_\Delta c p_\pi^* \cos\theta^*\right).$$

Averaging over the decay angle gives $\langle E_\pi \rangle = \gamma_\Delta E_\pi^*$, while $E_\Delta = \gamma_\Delta m_\Delta c^2 \simeq E_p$ because the target photon energy is negligible compared with E_p. Hence

$$\left\langle \frac{E_\pi}{E_p} \right\rangle \simeq \frac{E_\pi^*}{m_\Delta c^2} = \frac{m_\Delta^2 - m_N^2 + m_\pi^2}{2m_\Delta^2} \simeq 0.22\,. \tag{2.257}$$

The neutrino energy can be estimated in the same way. In the charged pion rest frame, two-body decay gives

$$E_{\nu_\mu}^{\text{prompt}*} = \frac{(m_\pi^2 - m_\mu^2)c^2}{2m_\pi}, \qquad E_\mu^* = \frac{(m_\pi^2 + m_\mu^2)c^2}{2m_\pi}.$$

After averaging over the pion-frame decay angle, the prompt neutrino therefore carries the fraction

$$\frac{\langle E_{\nu_\mu}^{\text{prompt}} \rangle}{E_\pi} = \frac{1 - r_\pi}{2} \simeq 0.21, \qquad r_\pi \equiv \frac{m_\mu^2}{m_\pi^2}.$$

The muon carries the complementary fraction $(1 + r_\pi)/2 \simeq 0.79$. If its energy is shared approximately equally among the electron and the two neutrinos in the three-body decay, each secondary neutrino carries

$$\frac{\langle E_\nu^{(\mu)} \rangle}{E_\pi} \sim \frac{1}{3}\frac{1 + r_\pi}{2} \simeq 0.26.$$

Thus all three neutrinos may be assigned the representative energy $E_\pi/4$. Similarly, $\pi^0 \to 2\gamma$ gives the average photon energy $E_\pi/2$ [90]. The characteristic secondary energies are therefore

$$E_\gamma \simeq \frac{E_\pi}{2} \simeq 0.1 E_p, \qquad (\pi^0 \to 2\gamma), \tag{2.258}$$

$$E_\nu \simeq \frac{E_\pi}{4} \simeq 0.05 E_p, \qquad (\pi^\pm \to \mu^\pm \to e^\pm). \tag{2.259}$$

These are delta-functional estimates in the source frame. The exact decay spectra are broad, and multi-pion production broadens the mapping from E_p to secondary energies further.

An uncooled pion–muon chain produces the approximate source flavor ratio

$$\nu_e : \nu_\mu : \nu_\tau \simeq 1 : 2 : 0, \tag{2.260}$$

where neutrinos and antineutrinos are combined. Oscillations over astrophysical baselines redistribute this flux to a ratio close to $1 : 1 : 1$ at Earth. Strong synchrotron or interaction losses of pions and muons can change both the energy mapping and the flavor ratio [91].

Neutron beta decay provides an additional, lower-energy channel,

$$n \to p + e^- + \bar{\nu}_e. \tag{2.261}$$

In the neutron rest frame the available kinetic energy is only $Q \simeq 0.782\ \mathrm{MeV}$, so a boosted decay typically gives

$$\frac{E_{\bar{\nu}_e}}{E_n} \sim \frac{\text{few} \times 10^{-1}\ \mathrm{MeV}}{m_n c^2} \sim \text{few} \times 10^{-4}. \tag{2.262}$$

At a fixed parent-nucleon energy these neutrinos are therefore about two orders of magnitude less energetic than those from the pion–muon chain.

Neutrino and γ-ray production rates. Let

$$Q_i(E) \equiv \frac{\mathrm{d}N_i}{\mathrm{d}E\,\mathrm{d}t} \tag{2.263}$$

be the source production rate per unit energy of species i, and define $Q_{\pi^\pm} = Q_{\pi^+} + Q_{\pi^-}$. Here Q_{ν_α} includes both ν_α and $\bar{\nu}_\alpha$, as does the intensity I_{ν_α} below. In the delta-functional decay approximation, each charged pion produces three neutrinos near $E_\nu = E_\pi/4$. Including the Jacobian $\mathrm{d}E_\pi/\mathrm{d}E_\nu = 4$,

$$\sum_\alpha Q_{\nu_\alpha}(E_\nu) \simeq 12\, Q_{\pi^\pm}(4E_\nu), \tag{2.264}$$

or, in terms of power per logarithmic energy interval,

$$\sum_\alpha E_\nu^2 Q_{\nu_\alpha}(E_\nu) \simeq \frac{3}{4}\left[E_\pi^2 Q_{\pi^\pm}(E_\pi)\right]_{E_\pi = 4E_\nu}. \tag{2.265}$$

The factor $3/4$ is the fraction of charged-pion energy carried by the three neutrinos in this approximation.

Similarly, two photons are produced near $E_\gamma = E_\pi/2$. The Jacobian now gives

$$E_\gamma^2 Q_\gamma(E_\gamma) \simeq \left[E_\pi^2 Q_{\pi^0}(E_\pi)\right]_{E_\pi = 2E_\gamma}. \tag{2.266}$$

Introduce the charged-to-neutral pion ratio, evaluated at the same pion energy,

$$R_\pi(E_\pi) \equiv \frac{Q_{\pi^\pm}(E_\pi)}{Q_{\pi^0}(E_\pi)}. \tag{2.267}$$

The Δ-resonance gives $R_\pi \simeq 1/2$, whereas high-energy pp collisions give $R_\pi \simeq 2$; multi-pion photohadronic interactions interpolate between these simple limits. Combining the two decay channels gives the source relation

$$\frac{1}{3}\sum_\alpha E_\nu^2 Q_{\nu_\alpha}(E_\nu) \simeq \frac{R_\pi(4E_\nu)}{4}\left[E_\gamma^2 Q_\gamma(E_\gamma)\right]_{E_\gamma=2E_\nu}. \tag{2.268}$$

The left-hand side is the flavor-averaged rate; after oscillations it is also a useful approximation for each flavor at Earth.

This *γ-ray–neutrino multi-messenger relation* provides a quick, order-of-magnitude estimate of the pionic signal in one channel from observations in the other [91]. We stress, however, that it applies at production. In general, internal absorption and electromagnetic cascades can strongly modify the escaping photon spectrum without changing the neutrino output.

Diffuse neutrino intensity. For one source, define the all-flavor neutrino luminosity per logarithmic energy interval by

$$\mathcal{L}_\nu(E_\nu) \equiv \sum_\alpha E_\nu^2 Q_{\nu_\alpha}(E_\nu). \tag{2.269}$$

For a population with local comoving density n_s, the corresponding local luminosity density is

$$\mathcal{E}_{\nu,0}(E_\nu) \equiv E_\nu^2 \frac{\mathrm{d}Q_\nu}{\mathrm{d}V} \simeq n_s \mathcal{L}_\nu(E_\nu). \tag{2.270}$$

Here $Q_\nu = \sum_\alpha Q_{\nu_\alpha}$, and $\mathrm{d}Q_\nu/\mathrm{d}V$ is the production rate per unit energy, time, and comoving volume, summed over the source population.

This is a differential luminosity per logarithmic interval, not a bolometric luminosity.

The cosmological line-of-sight integral can be summarized as

$$\sum_\alpha E_\nu^2 I_{\nu_\alpha}(E_\nu) \simeq \frac{c}{4\pi}\, t_{\rm H}\, \xi_z\, \mathcal{E}_{\nu,0}(E_\nu) = \frac{\xi_z R_{\rm H}}{4\pi}\, n_s \mathcal{L}_\nu(E_\nu), \tag{2.271}$$

where $t_{\rm H} = H_0^{-1}$, $R_{\rm H} = c t_{\rm H} \simeq 4.3\ \mathrm{Gpc}$, and the dimensionless factor ξ_z contains the source evolution, redshift losses, and spectral-energy mapping. Its value is of order unity but is population and spectrum dependent.

Around $E_\nu \sim 100\ \mathrm{TeV}$, a representative all-flavor IceCube intensity is

$$\sum_\alpha E_\nu^2 I_{\nu_\alpha}^{\rm obs} \sim \text{few} \times 10^{-8}\ \mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}, \tag{2.272}$$

with an energy dependence that is not exactly E_ν^{-2} [91, 92]. Equation (2.271) then implies the benchmark

$$n_s \mathcal{L}_\nu \sim 10^{36}\, \xi_z^{-1}\ \mathrm{erg\,Mpc^{-3}\,s^{-1}} \sim 3 \times 10^{43}\, \xi_z^{-1}\ \mathrm{erg\,Mpc^{-3}\,yr^{-1}} \tag{2.273}$$

per logarithmic energy interval. This luminosity-density requirement is compared with representative candidate source populations in Fig. 26.

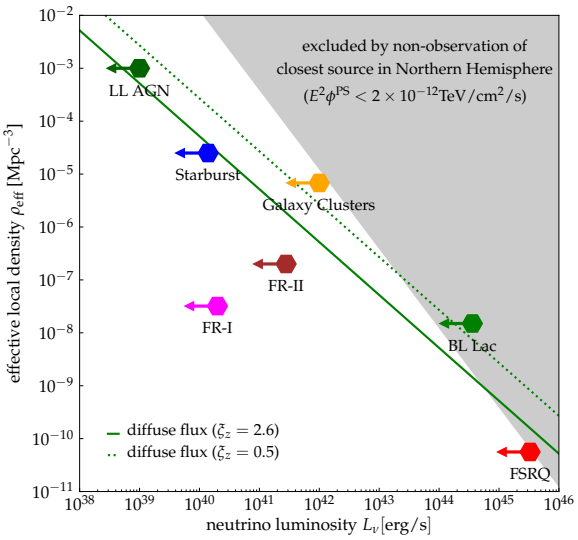


Figure 26: Neutrino luminosity-density requirements inferred from IceCube compared with representative candidate source populations. Adapted from [91].

Cosmic-ray power and pion-production efficiency. The observed UHECR intensity provides a complementary estimate of the power supplied by extragalactic accelerators. To construct a proton benchmark, we first assign this power to protons. This is a simplifying assumption, rather than a measurement of a pure-proton flux; the mixed composition discussed in Section 2.17 will enter below. Around the EeV to tens-of-EeV range, a representative intensity is

$$E_p^2 I_p^{\rm UHE} \sim 10^{-7}\ {\rm GeV\,cm^{-2}\,s^{-1}\,sr^{-1}} \tag{2.274}$$

[93]. After accounting for cosmological propagation, the proton benchmark corresponds to a local source power density of order

$$\mathcal{E}_{p,0}(E_p) \equiv E_p^2 \frac{{\rm d}Q_p}{{\rm d}V} \sim 5 \times 10^{36}\ {\rm erg\,Mpc^{-3}\,s^{-1}}, \tag{2.275}$$

per logarithmic interval, with an order-unity dependence on source evolution, composition, and propagation losses.

Let $f_\pi(E_p)$ be the fraction of the injected proton power that is converted into all pion charge states before escape. Thus $0 \le f_\pi \le 1$; for a thin target, $f_\pi \simeq \kappa t_{\rm res}/\tau_{\rm int}$, where $t_{\rm res}$ is the residence time, whereas a hadronic calorimeter approaches $f_\pi \to 1$. The fraction of pion power entering charged pions is $R_\pi/(1+R_\pi)$. Using $E_\pi \simeq 0.2 E_p$,

$$E_\pi^2 \frac{{\rm d}Q_{\pi^\pm}}{{\rm d}V} \simeq \frac{R_\pi(E_\pi)}{1+R_\pi(E_\pi)} \left[f_\pi(E_p)\, E_p^2 \frac{{\rm d}Q_p}{{\rm d}V} \right]_{E_p \simeq 5E_\pi}. \tag{2.276}$$

Combining this result with Eqs. (2.265) and (2.271) gives the flavor-averaged diffuse intensity

$$\frac{1}{3}\sum_\alpha E_\nu^2 I_{\nu_\alpha} \simeq \frac{c\, t_{\rm H} \xi_z}{4\pi} \frac{1}{4} \frac{R_\pi}{1+R_\pi} \left[f_\pi(E_p)\, E_p^2 \frac{{\rm d}Q_p}{{\rm d}V} \right]_{E_p \simeq 20E_\nu}, \tag{2.277}$$

where R_π is evaluated at $E_\pi \simeq 4E_\nu$. As in Eq. (2.271), the cosmological energy mapping is included in ξ_z.

Evaluating Eq. (2.277) requires the proton power at $E_p \simeq 20E_\nu$. At the IceCube energies, $E_\nu \sim 100$ TeV, this means $E_p \sim 2$ PeV. The UHECR power density of Eq. (2.275) is instead inferred from particles at $E_p \sim 10^{18}$–$10^{19.5}$ eV, well above the parent energies relevant to IceCube. The two energy ranges can be connected only through an assumption on the source spectrum. The Waxman–Bahcall argument assumes an injection spectrum $Q_p(E_p) \propto E_p^{-2}$. The power per logarithmic interval, $E_p^2\,\mathrm{d}Q_p/\mathrm{d}V$, is then the same in every decade, and the UHECR value applies unchanged several decades lower in energy. A softer spectrum puts more power at PeV energies than the UHECR normalization implies, and the neutrino flux can then exceed the estimate below.

Under this assumption, Eq. (2.275) and $R_\pi \simeq 1/2$–2 give the benchmark

$$\frac{1}{3}\sum_\alpha E_\nu^2 I_{\nu_\alpha} \sim (1\text{–}2)\times 10^{-8}\, f_\pi \xi_z \ \mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}. \tag{2.278}$$

This is the scale associated with the Waxman–Bahcall argument [94].

For a fixed injected proton power, the largest neutrino output is reached when $f_\pi \to 1$: essentially all that power at the relevant energies is converted into pions before escape. This energy-budget estimate is therefore also called the *calorimetric limit*, illustrated by the green dashed curve in Fig. 27 [91].

Real sources need not reach the calorimetric limit: if the gas or photon density is low, or if protons escape rapidly, their residence time can be shorter than the pion-production energy-loss time, giving $f_\pi < 1$. In particular, sources of the observed UHECRs must allow a substantial component to escape before losing most of its energy. They may nevertheless be calorimetric at lower energies, where magnetic confinement can be more effective.

If the escaping UHECRs are nuclei, an additional constraint comes from their survival against photodisintegration rather than photopion production. Requiring that no more than about half of the energy in the leading nuclear component be removed before escape limits the photon density and, indirectly, the photopion efficiency. For an E^{-2} injection spectrum, and using the narrow-resonance approximation for a power-law photon field, this gives the all-flavor nucleus-survival benchmark [90, 95]

$$\sum_\alpha E_\nu^2 I_{\nu_\alpha} \lesssim 2\times 10^{-9}\,\xi_z \left(\frac{A}{56}\right)^{-0.21} \mathrm{GeV\,cm^{-2}\,s^{-1}\,sr^{-1}}, \qquad A > 4.$$

This all-flavor Murase–Beacom benchmark is about an order of magnitude below the corresponding all-flavor proton scale: nuclear survival restricts the photon density before photopion production becomes efficient.

The connection to the observed UHECR power assumes that the neutrino sources also supply a substantial escaping cosmic-ray component. A *hidden* population may instead be efficient at producing neutrinos while contributing little to the observed UHECR flux; its injected power is then not fixed by the UHECR normalization. Likewise, γ rays produced

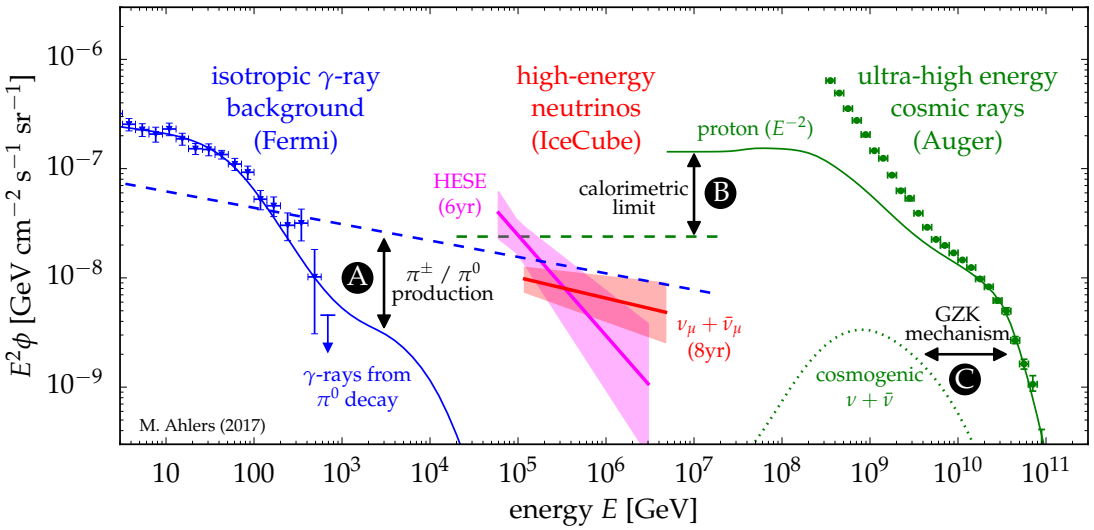


Figure 27: Multi-messenger panorama of cosmic rays, γ rays, and neutrinos. Adapted from [91].

together with the neutrinos may be absorbed inside the source and reprocessed to lower energies. NGC 1068, introduced at the start of this section, is a prominent candidate for such an obscured accelerator [87].

Why neutrinos probe cosmological distances. A neutrino emitted at redshift z_s with energy $E_{\rm em}$ is observed with energy $E_0 = E_{\rm em}/(1+z_s)$. Its survival probability is $P_{\rm surv} = \exp(-\tau_\nu)$, where the optical depth accumulated along the expanding line of sight is

$$\tau_\nu(E_0, z_s) = \int_0^{z_s} \frac{c\,\mathrm{d}z}{(1+z)H(z)} \sum_j n_j(z)\,\big\langle (1-\beta_j\cos\theta)\,\sigma_{\nu j}(s)\big\rangle\,. \tag{2.279}$$

Here j runs over the cosmological targets, principally relic neutrinos and baryons, with proper number density n_j. The average is over target momenta and collision angles; $\beta_j = v_j/c$, and θ is the angle between the neutrino and target momenta. The factor $1-\beta_j\cos\theta$ accounts for their relative flux, and s is evaluated at the local neutrino energy $E_\nu = E_0(1+z)$. For a freely expanding background, $n_j(z) = n_{j,0}(1+z)^3$.

Well below the weak-boson mass scale, $s \ll m_W^2 c^4$, the four-fermion estimate is $\sigma_{\nu\bar{\nu}} \sim G_{\rm F}^2 s(\hbar c)^2/\pi$, where $G_{\rm F}$ has its usual units of inverse energy squared. For relativistic relic-neutrino targets, $s \sim E_\nu k_{\rm B} T_\nu$. Since both E_ν and T_ν scale as $1+z$, the local interaction rate scales, between channel thresholds, as

$$\Gamma_{\nu\bar{\nu}} \sim c\,n_\nu \langle\sigma_{\nu\bar{\nu}}\rangle \propto (1+z)^5. \tag{2.280}$$

A useful low-redshift scale in this approximation is the mean free time

$$\Gamma_{\nu\bar{\nu},0}^{-1} \sim 10^{27}\ \mathrm{yr}\, N_{\rm ch}^{-1} \left(\frac{E_0}{\mathrm{TeV}}\right)^{-1}, \tag{2.281}$$

where $N_{\rm ch}$ is an effective channel factor that includes weak-coupling weights and kinematic thresholds [96]. This is enormously longer than the Hubble time, even at TeV–PeV energies; scattering on the mean cosmological baryon density is negligible as well. When relic

neutrinos are non-relativistic, $k_{\mathrm{B}}T_\nu$ in the estimate of s is replaced parametrically by $m_\nu c^2$, shortening the mean free time but leaving $\tau_\nu \ll 1$ for sources at astrophysical redshifts. Neutrinos can therefore reach us from cosmological distances: expansion redshifts their energies and oscillations mix their flavors, but neither process removes them. At much higher energies or sufficiently early epochs, interactions with the relic-neutrino background can instead impose a finite horizon, particularly near resonant annihilation.

References

[1] P. Blasi, "The Origin of Galactic Cosmic Rays," *Astron. Astrophys. Rev.* **21** (2013) 70, arXiv:1311.7346 [astro-ph.HE].

[2] D. Boncioli, "Cosmic-ray propagation in extragalactic space and secondary messengers," *Proc. Int. Sch. Phys. Fermi* **208** (2024) 315–351, arXiv:2309.12743 [astro-ph.HE].

[3] C. Evoli and U. Dupletsa, "Phenomenological models of cosmic ray transport in galaxies," *Proc. Int. Sch. Phys. Fermi* **208** (2024) 257–313, arXiv:2309.00298 [astro-ph.HE].

[4] S. Gabici, C. Evoli, D. Gaggero, P. Lipari, P. Mertsch, E. Orlando, A. Strong, and A. Vittino, "The origin of Galactic cosmic rays: challenges to the standard paradigm," *Int. J. Mod. Phys. D* **28** no. 15, (2019) 1930022, arXiv:1903.11584 [astro-ph.HE].

[5] P. D. Serpico, "Cosmic ray interactions with matter and radiation," *Proc. Int. Sch. Phys. Fermi* **208** (2024) 1–40, arXiv:2308.04361 [astro-ph.HE].

[6] G. B. Rybicki and A. P. Lightman, *Radiative Processes in Astrophysics*. Wiley, New York, 1979.

[7] M. S. Longair, *High Energy Astrophysics*. Cambridge University Press, Cambridge, 2011.

[8] G. Ghisellini, *Radiative Processes in High Energy Astrophysics*, vol. 873. 2013. arXiv:1202.5949 [astro-ph.HE].

[9] R. Yang, F. Aharonian, and C. Evoli, "Radial distribution of the diffuse γ-ray emissivity in the Galactic disk," *Phys. Rev. D* **93** no. 12, (2016) 123007, arXiv:1602.04710 [astro-ph.HE].

[10] L. Tibaldo, D. Gaggero, and P. Martin, "Gamma rays as probes of cosmic-ray propagation and interactions in galaxies," *Universe* **7** no. 5, (2021) 141, arXiv:2103.16423 [astro-ph.HE].

[11] **IceCube** Collaboration, R. Abbasi *et al.*, "Observation of high-energy neutrinos from the Galactic plane," *Science* **380** no. 6652, (2023) 1338–1343, arXiv:2307.04427 [astro-ph.HE].

[12] F. Lenz, "The Ratio of Proton and Electron Masses," *Phys. Rev.* **82** no. 4, (1951) 554–554.

[13] K. K. Madsen *et al.*, "Broadband X-ray Imaging and Spectroscopy of the Crab Nebula and Pulsar with NuSTAR," *Astrophys. J.* **801** no. 1, (2015) 66, arXiv:1502.07765 [astro-ph.HE].

[14] F. Aharonian and A. M. Taylor, "Limitations on the Photo-disintegration Process as a Source of VHE Photons," *Astropart. Phys.* **34** (2010) 258–266, arXiv:1005.3230 [astro-ph.HE].

[15] R. Zanin, "Crab Observational Status: Nebulae, Pulsations, and Flares," *Astrophys. Space Sci. Libr.* **446** (2017) 101–133.

[16] B. Cerutti, G. R. Werner, D. A. Uzdensky, and M. C. Begelman, "Simulations of particle acceleration beyond the classical synchrotron burnoff limit in magnetic reconnection: An explanation of the Crab flares," *Astrophys. J.* **770** (2013) 147, arXiv:1302.6247 [astro-ph.HE].

[17] D. A. Green, "A Revised Galactic Supernova Remnant Catalogue," *Bull. Astron. Soc. India* **37** (2009) 45, arXiv:0905.3699 [astro-ph.HE].

[18] P. Platania, M. Bensadoun, M. Bersanelli, G. De Amici, A. Kogut, S. Levin, D. Maino, and G. F. Smoot, "A determination of the spectral index of galactic synchrotron emission in the 1-10 GHz range," *Astrophys. J.* **505** (1998) 473, arXiv:astro-ph/9707252.

[19] A. W. Strong, E. Orlando, and T. R. Jaffe, "The interstellar cosmic-ray electron spectrum from synchrotron radiation and direct measurements," *Astron. Astrophys.* **534** (2011) A54, arXiv:1108.4822 [astro-ph.HE].

[20] **WMAP** Collaboration, G. Hinshaw *et al.*, "Three-year Wilkinson Microwave Anisotropy Probe (WMAP) observations: temperature analysis," *Astrophys. J. Suppl.* **170** (2007) 288, arXiv:astro-ph/0603451.

[21] R. Beck and M. Krause, "Revised equipartition & minimum energy formula for magnetic field strength estimates from radio synchrotron observations," *Astron. Nachr.* **326** (2005) 414–427, arXiv:astro-ph/0507367.

[22] M. Unger and G. R. Farrar, "The Coherent Magnetic Field of the Milky Way," *Astrophys. J.* **970** no. 1, (2024) 95, arXiv:2311.12120 [astro-ph.GA].

[23] M. Unger and G. R. Farrar, "The Random Magnetic Field of the Milky Way," 8, 2026.

[24] M. Lyutikov, T. Temim, S. Komissarov, P. Slane, L. Sironi, and L. Comisso, "Interpreting Crab Nebula's synchrotron spectrum: two acceleration mechanisms," *Mon. Not. Roy. Astron. Soc.* **489** no. 2, (2019) 2403–2416, arXiv:1811.01767 [astro-ph.HE].

[25] **LHAASO** Collaboration, Z. Cao *et al.*, "Peta–electron volt gamma-ray emission from the Crab Nebula," *Science* **373** no. 6553, (2021) 425–430, arXiv:2111.06545 [astro-ph.HE].

[26] G. R. Blumenthal and R. J. Gould, "Bremsstrahlung, synchrotron radiation, and compton scattering of high-energy electrons traversing dilute gases," *Rev. Mod. Phys.* **42** (1970) 237–270.

[27] D. Khangulyan, F. Aharonian, and A. M. Taylor, "On the Properties of Inverse Compton Spectra Generated by Upscattering a Power-law Distribution of Target Photons," *Astrophys. J.* **954** no. 2, (2023) 186, arXiv:2307.12467 [astro-ph.HE].

[28] D. Khangulyan, F. A. Aharonian, and S. R. Kelner, "Simple Analytical Approximations for Treatment of Inverse Compton Scattering of Relativistic Electrons in the Blackbody Radiation Field," *Astrophys. J.* **783** (2014) 100, arXiv:1310.7971 [astro-ph.HE].

[29] K. Fang, X.-J. Bi, S.-J. Lin, and Q. Yuan, "Klein–Nishina Effect and the Cosmic Ray Electron Spectrum," *Chin. Phys. Lett.* **38** no. 3, (2021) 039801, arXiv:2007.15601 [astro-ph.HE].

[30] S. Vernetto and P. Lipari, "Absorption of very high energy gamma rays in the Milky Way," *Phys. Rev. D* **94** no. 6, (2016) 063009, arXiv:1608.01587 [astro-ph.HE].

[31] L. Stawarz, V. Petrosian, and R. D. Blandford, "On the Energy Spectra of GeV/TeV Cosmic Ray Leptons," *Astrophys. J.* **710** (2010) 236–247, arXiv:0908.1094 [astro-ph.GA].

[32] **AMS** Collaboration, M. Aguilar *et al.*, "Towards Understanding the Origin of Cosmic-Ray Electrons," *Phys. Rev. Lett.* **122** no. 10, (2019) 101101.

[33] C. Evoli, P. Blasi, E. Amato, and R. Aloisio, "Signature of Energy Losses on the Cosmic Ray Electron Spectrum," *Phys. Rev. Lett.* **125** no. 5, (2020) 051101, arXiv:2007.01302 [astro-ph.HE]. [Erratum: Phys.Rev.Lett. 126, 249901 (2021)].

[34] A. Dominguez *et al.*, "Extragalactic Background Light Inferred from AEGIS Galaxy SED-type Fractions," *Mon. Not. Roy. Astron. Soc.* **410** (2011) 2556, arXiv:1007.1459 [astro-ph.CO].

[35] R. C. Gilmore, R. S. Somerville, J. R. Primack, and A. Dominguez, "Semi-analytic modeling of the EBL and consequences for extragalactic gamma-ray spectra," *Mon. Not. Roy. Astron. Soc.* **422** (2012) 3189, arXiv:1104.0671 [astro-ph.CO].

[36] A. Saldana-Lopez, A. Domínguez, P. G. Pérez-González, J. Finke, M. Ajello, J. R. Primack, V. S. Paliya, and A. Desai, "An observational determination of the evolving extragalactic background light from the multiwavelength HST/CANDELS survey in the Fermi and CTA era," *Mon. Not. Roy. Astron. Soc.* **507** no. 4, (2021) 5144–5160, arXiv:2012.03035 [astro-ph.CO].

[37] R. J. Gould and G. P. Schreder, "Pair Production in Photon-Photon Collisions," *Phys. Rev.* **155** (1967) 1404–1407.

[38] **H.E.S.S.** Collaboration, A. Abramowski *et al.*, "Measurement of the extragalactic background light imprint on the spectra of the brightest blazars observed with H.E.S.S," *Astron. Astrophys.* **550** (2013) A4, arXiv:1212.3409 [astro-ph.HE].

[39] V. Berezinsky and O. Kalashev, "High energy electromagnetic cascades in extragalactic space: physics and features," *Phys. Rev. D* **94** no. 2, (2016) 023007, arXiv:1603.03989 [astro-ph.HE].

[40] V. S. Berezinskii, S. V. Bulanov, V. A. Dogiel, and V. S. Ptuskin, *Astrophysics of cosmic rays*. North-Holland, Amsterdam, 1990.

[41] A. Capanema, *Cascaded gamma-ray Counterpart of the IceCube Neutrinos*. PhD thesis, Pontifícia Universidade Católica do Rio de Janeiro, 2020.

[42] **Fermi-LAT** Collaboration, M. Ackermann *et al.*, "The spectrum of isotropic diffuse gamma-ray emission between 100 MeV and 820 GeV," *Astrophys. J.* **799** (2015) 86, arXiv:1410.3696 [astro-ph.HE].

[43] S. S. M. Wong, *Introductory Nuclear Physics*. John Wiley & Sons, New York, 2 ed., 1998.

[44] J. R. Letaw, R. Silberberg, and C. H. Tsao, "Proton-nucleus total inelastic cross sections: an empirical formula for $E > 10$ MeV," *Astrophys. J. Suppl.* **51** (Mar., 1983) 271.

[45] D. Maurin, R. Taillet, F. Donato, P. Salati, A. Barrau, and G. Boudoul, "Galactic cosmic ray nuclei as a tool for astroparticle physics," 12, 2002.

[46] Y. Genolini, D. Maurin, I. V. Moskalenko, and M. Unger, "Current status and desired precision of the isotopic production cross sections relevant to astrophysics of cosmic rays: Li, Be, B, C, and N," *Phys. Rev. C* **98** no. 3, (2018) 034611, arXiv:1803.04686 [astro-ph.HE].

[47] A. Korejwo, T. Dzikowski, M. Giller, J. Wdowczyk, V. V. Perelygin, and A. V. Zarubin, "The measurement of isotopic cross-sections of C-12 beam fragmentation on liquid hydrogen at 3.66-GeV/nucleon," *J. Phys. G* **26** (2000) 1171–1186.

[48] R. Silberberg and C. H. Tsao, "Partial Cross-Sections in High-Energy Nuclear Reactions, and Astrophysical Applications. I. Targets With z <= 28.," *Astrophys. J. Suppl.* **25** (1973) 315–333.

[49] R. Silberberg and C. H. Tsao, "Partial Cross-Sections in High-Energy Nuclear Reactions, and Astrophysical Applications. II. Targets Heavier than Nickel," *Astrophys. J. Suppl.* **25** (1973) 335–367.

[50] R. Silberberg, C. H. Tsao, and A. F. Barghouty, "Updated Partial Cross Sections of Proton-Nucleus Reactions," *Astrophys. J.* **501** no. 2, (1998) 911–919.

[51] W. R. Webber, J. C. Kish, and D. A. Schrier, "Formula for calculating partial cross-sections for nuclear reactions of nuclei with E >= 200-MeV/nucleon in hydrogen targets," *Phys. Rev. C* **41** (1990) 566–571.

[52] W. R. Webber, A. Soutoul, J. C. Kish, and J. M. Rockstroh, "Updated Formula for Calculating Partial Cross Sections for Nuclear Reactions of Nuclei with $Z \leq 28$ and $E > 150$ MeV Nucleon^{-1} in Hydrogen Targets," *Astrophys. J. Suppl.* **144** no. 1, (Jan., 2003) 153–167.

[53] A. W. Strong, I. V. Moskalenko, and V. S. Ptuskin, "Cosmic-ray propagation and interactions in the Galaxy," *Ann. Rev. Nucl. Part. Sci.* **57** (2007) 285–327, arXiv:astro-ph/0701517.

[54] I. V. Moskalenko, A. W. Strong, and S. G. Mashnik, "Propagation of cosmic rays: Nuclear physics in cosmic-ray studies," *AIP Conf. Proc.* **769** no. 1, (2005) 1612–1617, arXiv:astro-ph/0411400.

[55] D. Maurin *et al.*, "Precision cross-sections for advancing cosmic-ray physics and other applications: A comprehensive programme for the next decade," *Phys. Rept.* **1161** (2026) 1–81, arXiv:2503.16173 [astro-ph.HE].

[56] S. R. Kelner, F. A. Aharonian, and V. V. Bugayov, "Energy spectra of gamma-rays, electrons and neutrinos produced at proton-proton interactions in the very high energy regime," *Phys. Rev. D* **74** (2006) 034018, arXiv:astro-ph/0606058. [Erratum: Phys.Rev.D 79, 039901 (2009)].

[57] F. A. Aharonian, "High-energy cosmic rays and gamma radiation," *Proc. Int. Sch. Phys. Fermi* **208** (2024) 391–425.

[58] J. W. Norbury, "Total cross section parameterizations for pion production in nucleon-nucleon collisions," *Nucl. Instrum. Meth. B* **267** (2009) 1209–1212.

[59] I. V. Moskalenko and A. W. Strong, "Production and propagation of cosmic ray positrons and electrons," *Astrophys. J.* **493** (1998) 694–707, arXiv:astro-ph/9710124.

[60] P. Lipari, "Charged pion chain decay and the cosmic ray positron flux," 10, 2025.

[61] **Fermi-LAT** Collaboration, M. Ackermann *et al.*, "Detection of the Characteristic Pion-Decay Signature in Supernova Remnants," *Science* **339** (2013) 807, arXiv:1302.3307 [astro-ph.HE].

[62] D. Caprioli, "Understanding hadronic gamma-ray emission from supernova remnants," *JCAP* **05** (2011) 026, arXiv:1103.2624 [astro-ph.HE].

[63] D. Caprioli, "Particle acceleration at shocks: An introduction," *Proc. Int. Sch. Phys. Fermi* **208** (2024) 143–181, arXiv:2307.00284 [astro-ph.HE].

[64] F. Riehn, R. Engel, A. Fedynitch, T. K. Gaisser, and T. Stanev, "Hadronic interaction model Sibyll 2.3d and extensive air showers," *Phys. Rev. D* **102** no. 6, (2020) 063002, arXiv:1912.03300 [hep-ph].

[65] K. M. Ferriere, "The interstellar environment of our galaxy," *Rev. Mod. Phys.* **73** (2001) 1031–1066, arXiv:astro-ph/0106359.

[66] **PAMELA** Collaboration, O. Adriani *et al.*, "PAMELA Measurements of Cosmic-ray Proton and Helium Spectra," *Science* **332** (2011) 69–72, arXiv:1103.4055 [astro-ph.HE].

[67] M. Mori, "Nuclear enhancement factor in calculation of Galactic diffuse gamma-rays: a new estimate with DPMJET-3," *Astropart. Phys.* **31** (2009) 341–343, arXiv:0903.3260 [astro-ph.HE].

[68] **Fermi-LAT** Collaboration, M. Ackermann *et al.*, "Fermi-LAT Observations of the Diffuse Gamma-Ray Emission: Implications for Cosmic Rays and the Interstellar Medium," *Astrophys. J.* **750** (2012) 3, arXiv:1202.4039 [astro-ph.HE].

[69] S. Hayakawa, "Propagation of the Cosmic Radiation through Intersteller Space," *Progress of Theoretical Physics* **8** no. 5, (Nov., 1952) 571–572.

[70] S. Hayakawa and S. Kobayashi, "Propagation of the Cosmic Radiation through Interstellar Space," *Journal of Geomagnetism and Geoelectricity* **5** no. 3, (Jan., 1953) 83–93.

[71] K. Greisen, "End to the cosmic ray spectrum?," *Phys. Rev. Lett.* **16** (1966) 748–750.

[72] G. T. Zatsepin and V. A. Kuzmin, "Upper limit of the spectrum of cosmic rays," *JETP Lett.* **4** (1966) 78–80.

[73] V. Berezinsky, A. Z. Gazizov, and S. I. Grigorieva, "Dip in UHECR spectrum as signature of proton interaction with CMB," *Phys. Lett. B* **612** (2005) 147–153, arXiv:astro-ph/0502550.

[74] A. Mucke, R. Engel, J. P. Rachen, R. J. Protheroe, and T. Stanev, "SOPHIA: Monte Carlo simulations of photohadronic processes in astrophysics," *Comput. Phys. Commun.* **124** (2000) 290–314, arXiv:astro-ph/9903478.

[75] **Particle Data Group** Collaboration, S. Navas *et al.*, "Review of particle physics," *Phys. Rev. D* **110** no. 3, (2024) 030001.

[76] **CRPropa** Collaboration, E. Armengaud, G. Sigl, T. Beau, and F. Miniati, "CRPropa: a numerical tool for the propagation of uhe cosmic rays, gamma-rays and neutrinos," *Astropart. Phys.* **28** (2007) 463–471, arXiv:astro-ph/0603675.

[77] R. Aloisio, D. Boncioli, A. F. Grillo, S. Petrera, and F. Salamida, "SimProp: a Simulation Code for Ultra High Energy Cosmic Ray Propagation," *JCAP* **10** (2012) 007, arXiv:1204.2970 [astro-ph.HE].

[78] **Pierre Auger** Collaboration, A. Aab *et al.*, "Combined fit of spectrum and composition data as measured by the Pierre Auger Observatory," *JCAP* **04** (2017) 038, arXiv:1612.07155 [astro-ph.HE]. [Erratum: JCAP 03, E02 (2018)].

[79] **Pierre Auger** Collaboration, D. Boncioli, "A new view of UHECRs with the Pierre Auger Observatory," *PoS* **UHECR2024** (2025) 027, arXiv:2509.15862 [astro-ph.HE].

[80] H. Bethe and W. Heitler, "On the Stopping of fast particles and on the creation of positive electrons," *Proc. Roy. Soc. Lond. A* **146** (1934) 83–112.

[81] L. A. Dourado, D. Boncioli, A. di Matteo, C. Evoli, S. Petrera, and F. Salamida, "Towards the implementation of the photodisintegration interaction in SimProp Sirente," *PoS* **UHECR2024** (2025) 067.

[82] J. L. Puget, F. W. Stecker, and J. H. Bredekamp, "Photonuclear Interactions of Ultrahigh-Energy Cosmic Rays and their Astrophysical Consequences," *Astrophys. J.* **205** (1976) 638–654.

[83] E. Khan, S. Goriely, D. Allard, E. Parizot, T. Suomijarvi, A. J. Koning, S. Hilaire, and M. C. Duijvestijn, "Photodisintegration of ultra-high-energy cosmic rays revisited," *Astropart. Phys.* **23** (2005) 191–201, arXiv:astro-ph/0412109.

[84] A. Koning, S. Hilaire, and S. Goriely, "TALYS: modeling of nuclear reactions," *Eur. Phys. J. A* **59** no. 6, (2023) 131. [Erratum: Eur.Phys.J.A 59, 146 (2023)].

[85] E. Kido, T. Inakura, M. Kimura, N. Kobayashi, S. Nagataki, N. Shimizu, A. Tamii, and Y. Utsuno, "Evaluations of uncertainties in simulations of propagation of ultrahigh-energy cosmic-ray nuclei derived from microscopic nuclear models," *Astropart. Phys.* **152** (2023) 102866, arXiv:2206.03447 [astro-ph.HE].

[86] R. Alves Batista, D. Boncioli, A. di Matteo, A. van Vliet, and D. Walz, "Effects of uncertainties in simulations of extragalactic UHECR propagation, using CRPropa and SimProp," *JCAP* **10** (2015) 063, arXiv:1508.01824 [astro-ph.HE].

[87] **IceCube** Collaboration, R. Abbasi *et al.*, "Evidence for neutrino emission from the nearby active galaxy NGC 1068," *Science* **378** no. 6619, (2022) 538–543, arXiv:2211.09972 [astro-ph.HE].

[88] **IceCube** Collaboration, M. G. Aartsen *et al.*, "Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector," *Science* **342** (2013) 1242856, arXiv:1311.5238 [astro-ph.HE].

[89] **KM3NeT** Collaboration, S. Aiello *et al.*, "Observation of an ultra-high-energy cosmic neutrino with KM3NeT," *Nature* **638** no. 8050, (2025) 376–382. [Erratum: Nature 640, E3 (2025)].

[90] M. S. Muzio, "Ultrahigh energy cosmic rays and neutrino flux models," *Eur. Phys. J. ST* **234** no. 16, (2025) 4939–4949, arXiv:2502.11834 [astro-ph.HE].

[91] M. Ahlers and F. Halzen, "Opening a New Window onto the Universe with IceCube," *Prog. Part. Nucl. Phys.* **102** (2018) 73–88, arXiv:1805.11112 [astro-ph.HE].

[92] **IceCube** Collaboration, M. G. Aartsen *et al.*, "Observation of High-Energy Astrophysical Neutrinos in Three Years of IceCube Data," *Phys. Rev. Lett.* **113** (2014) 101101, arXiv:1405.5303 [astro-ph.HE].

[93] **Pierre Auger** Collaboration, P. Abreu *et al.*, "The energy spectrum of cosmic rays beyond the turn-down around 10^{17} eV as measured with the surface detector of the Pierre Auger Observatory," *Eur. Phys. J. C* **81** no. 11, (2021) 966, arXiv:2109.13400 [astro-ph.HE].

[94] E. Waxman and J. N. Bahcall, "High-energy neutrinos from astrophysical sources: An Upper bound," *Phys. Rev. D* **59** (1999) 023002, arXiv:hep-ph/9807282.

[95] K. Murase and J. F. Beacom, "Neutrino Background Flux from Sources of Ultrahigh-Energy Cosmic-Ray Nuclei," *Phys. Rev. D* **81** (2010) 123001, arXiv:1003.4959 [astro-ph.HE].

[96] P. Gondolo, G. Gelmini, and S. Sarkar, "Cosmic neutrinos from unstable relic particles," *Nucl. Phys. B* **392** (1993) 111–136, arXiv:hep-ph/9209236.

A Appendix

A.1 Formal steady-state solution with continuous energy losses

We derive the steady-state solutions quoted in Section 2.2. The kinetic equation is

$$\frac{\partial N}{\partial t} = \frac{\partial}{\partial E}[b(E)\,N(E,t)] - \frac{N(E,t)}{\tau_{\rm esc}(E)} + Q(E,t), \tag{A.1}$$

where N is the differential number density, Q is the injection rate density, $b(E) > 0$ is defined by $\dot{E} = -b(E)$, and $\tau_{\rm esc}(E) > 0$ is the escape time. We assume throughout that these functions are sufficiently regular for the operations below and that a steady state exists.

The sign convention is worth making explicit. The signed flux in energy space is $\dot{E}N = -bN$, so Eq. (A.1) is equivalent to

$$\frac{\partial N}{\partial t} + \frac{\partial}{\partial E}\left(\dot{E}\,N\right) = -\frac{N}{\tau_{\rm esc}} + Q.$$

Thus losses transport particles toward lower energies, escape removes particles locally, and injection adds them.

Steady state without escape. Set $\partial_t N = 0$ and take $\tau_{\rm esc} \to \infty$. Defining the magnitude of the downward energy-space flux,

$$Y(E) \equiv b(E)N(E), \tag{A.2}$$

the kinetic equation becomes

$$\frac{\mathrm{d}Y}{\mathrm{d}E} = -Q(E). \tag{A.3}$$

It is useful to work first on a finite interval $E \le E' \le E_{\rm u}$, where $E_{\rm u}$ lies above the highest energy at which particles enter the system. Integrating Eq. (A.3) gives

$$Y(E_{\rm u}) - Y(E) = -\int_E^{E_{\rm u}} Q(E')\,\mathrm{d}E', \tag{A.4}$$

or

$$Y(E) = Y(E_{\rm u}) + \int_E^{E_{\rm u}} Q(E')\,\mathrm{d}E'. \tag{A.5}$$

The boundary value $Y(E_{\rm u})$ would represent a downward particle flux entering the interval from still higher energies. For a closed upper boundary there is no such incoming flux, and therefore

$$Y(E_{\rm u}) = b(E_{\rm u})N(E_{\rm u}) = 0. \tag{A.6}$$

If the source extends without a sharp upper cutoff, we take $E_{\rm u} \to \infty$, provided the integral converges and $\lim_{E\to\infty} b(E)N(E) = 0$. The solution is then

$$b(E)N(E) = \int_E^{\infty} Q(E')\,\mathrm{d}E', \qquad N(E) = \frac{1}{b(E)}\int_E^{\infty} Q(E')\,\mathrm{d}E'. \tag{A.7}$$

The result has a direct conservation-law interpretation: the downward flux through E equals the total rate at which particles are injected above E.

Steady state with escape. Retaining escape and again setting $\partial_t N = 0$ gives

$$\frac{\mathrm{d}}{\mathrm{d}E}[b(E)N(E)] - \frac{N(E)}{\tau_{\mathrm{esc}}(E)} = -Q(E). \tag{A.8}$$

In terms of $Y = bN$, this is the linear first-order equation

$$\frac{\mathrm{d}Y}{\mathrm{d}E} - a(E)Y(E) = -Q(E), \qquad a(E) \equiv \frac{1}{b(E)\tau_{\mathrm{esc}}(E)}. \tag{A.9}$$

The dimensions provide a quick check: $[Y] = \mathrm{L}^{-3}\mathrm{T}^{-1}$ and $[a] = \mathrm{E}^{-1}$, so every term in Eq. (A.9) has units $\mathrm{L}^{-3}\mathrm{T}^{-1}\mathrm{E}^{-1}$.

Choose an arbitrary reference energy E_0 and define the integrating factor

$$\mathcal{I}(E) \equiv \exp\left[-\int_{E_0}^{E} a(\epsilon)\,\mathrm{d}\epsilon\right]. \tag{A.10}$$

Because $\mathcal{I}' = -a\mathcal{I}$, multiplying Eq. (A.9) by $\mathcal{I}$ gives

$$\frac{\mathrm{d}}{\mathrm{d}E}[\mathcal{I}(E)Y(E)] = -\mathcal{I}(E)Q(E). \tag{A.11}$$

Integrating from E to E_{u} and dividing by $\mathcal{I}(E)$ yields the general finite-boundary solution,

$$\begin{aligned} Y(E) = Y(E_{\mathrm{u}})\exp\left[-\int_E^{E_{\mathrm{u}}} \frac{\mathrm{d}\epsilon}{b(\epsilon)\tau_{\mathrm{esc}}(\epsilon)}\right] \\ + \int_E^{E_{\mathrm{u}}} \mathrm{d}E'\,Q(E')\exp\left[-\int_E^{E'} \frac{\mathrm{d}E''}{b(E'')\tau_{\mathrm{esc}}(E'')}\right]. \end{aligned} \tag{A.12}$$

The arbitrary reference energy E_0 has canceled, as it must.

For a closed upper boundary, $Y(E_{\mathrm{u}}) = 0$. Taking $E_{\mathrm{u}} \to \infty$ under the same convergence assumptions as above gives

$$N(E) = \frac{1}{b(E)}\int_E^{\infty} \mathrm{d}E'\,Q(E')\,\exp\left[-\int_E^{E'} \frac{\mathrm{d}E''}{b(E'')\tau_{\mathrm{esc}}(E'')}\right]. \tag{A.13}$$

Interpretation as a survival probability. A particle cooling from E' to $E < E'$ follows

$$\mathrm{d}t = -\frac{\mathrm{d}E}{b(E)}.$$

The integrated escape hazard accumulated along that trajectory is

$$\int \frac{\mathrm{d}t}{\tau_{\mathrm{esc}}} = \int_E^{E'} \frac{\mathrm{d}E''}{b(E'')\tau_{\mathrm{esc}}(E'')}. \tag{A.14}$$

Consequently,

$$\mathcal{S}(E \leftarrow E') \equiv \exp\left[-\int_E^{E'} \frac{\mathrm{d}E''}{b(E'')\tau_{\mathrm{esc}}(E'')}\right] \tag{A.15}$$

is precisely the probability that a particle injected at E' remains in the system while cooling to E. The solution can therefore be written in Green-function form,

$$N(E) = \int_0^{\infty} G(E,E')Q(E')\,\mathrm{d}E', \qquad G(E,E') = \frac{\Theta(E'-E)}{b(E)}\mathcal{S}(E \leftarrow E'). \tag{A.16}$$

Direct verification. Let

$$Y(E) = \int_E^\infty Q(E')\mathcal{S}(E \leftarrow E')\,\mathrm{d}E'.$$

Since $\mathcal{S}(E \leftarrow E) = 1$ and

$$\frac{\partial}{\partial E}\mathcal{S}(E \leftarrow E') = a(E)\mathcal{S}(E \leftarrow E'),$$

Leibniz's rule gives

$$\frac{\mathrm{d}Y}{\mathrm{d}E} = -Q(E) + a(E)Y(E). \tag{A.17}$$

Hence $Y' - aY = -Q$, which is exactly Eq. (A.9). This check fixes both the sign and the limits in the exponential.

Useful limits. The formal solution reproduces the expected physical regimes:

- If $\tau_{\mathrm{esc}} \to \infty$, then $\mathcal{S} \to 1$, and Eq. (A.13) reduces to the no-escape result Eq. (A.7).
- For monoenergetic injection $Q(E) = Q_0\delta(E - E_{\mathrm{inj}})$,

$$N(E) = \frac{Q_0}{b(E)}\mathcal{S}(E \leftarrow E_{\mathrm{inj}})\Theta(E_{\mathrm{inj}} - E). \tag{A.18}$$

  Escape therefore attenuates the no-escape cooling trail by its survival probability.
- If escape is much faster than cooling and Q, b, and τ_{esc} vary slowly over $\Delta E \sim b(E)\tau_{\mathrm{esc}}(E)$, the integral is local and gives

$$N(E) \simeq Q(E)\tau_{\mathrm{esc}}(E). \tag{A.19}$$

  Conversely, when losses dominate, the survival factor is nearly unity over the energies that contribute appreciably, and the no-escape solution applies.

For example, if $Q(E) = Q_0E^{-p}$ with $p > 1$ and $b(E) = b_0E^m$, the loss-dominated solution is

$$N(E) = \frac{Q_0}{b_0(p-1)}E^{-(p+m-1)}, \tag{A.20}$$

recovering the steepening derived in Eq. (2.13). Competition with escape produces a transition near

$$t_{\mathrm{cool}}(E) \equiv \frac{E}{b(E)} \sim \tau_{\mathrm{esc}}(E).$$

For a source of finite age, the analogous age-limited break occurs when $t_{\mathrm{cool}} \sim t_*$. That case is not a strict steady state below the break and requires either the time-dependent solution or the reservoir argument developed in Section 2.2.

A.2 Motion of a charged particle in a uniform magnetic field

We work in c.g.s. Gaussian units. A particle of charge q and rest mass m moves with velocity $\mathbf{v}$ in a constant magnetic field $\mathbf{B}$. We denote $\beta \equiv v/c$, $\gamma = (1-\beta^2)^{-1/2}$, and $\hat{\mathbf{b}} \equiv \mathbf{B}/B$ (the unit vector along the field).

The starting point is the Lorentz force,

$$\dot{\mathbf{p}} \equiv \frac{\mathrm{d}}{\mathrm{d}t}(\gamma m \mathbf{v}) = q\left(\mathbf{E} + \frac{\mathbf{v}}{c} \times \mathbf{B}\right) .$$

Throughout this appendix we set $\mathbf{E} = 0$. This is a good idealization in many astrophysical plasmas: in the fluid rest frame the electric field vanishes because of the extremely high conductivity (ideal MHD).

A first and very useful observation is that a magnetic field does no work:

$$\frac{\mathrm{d}E}{\mathrm{d}t} = \mathbf{v} \times \dot{\mathbf{p}} = q\,\mathbf{v} \times \mathbf{E} = 0 \quad \Longrightarrow \quad E = \gamma m c^2 = \text{const}\,.$$

Hence, both γ and the speed $|\mathbf{v}|$ are constant in time. The acceleration is therefore purely geometric: it can only *change the direction* of $\mathbf{v}$, not its magnitude.

With $\mathbf{E} = 0$ the equation $\gamma m\,\mathbf{a} = \frac{q}{c}\,\mathbf{v} \times \mathbf{B}$ immediately gives

$$\mathbf{v} \cdot \mathbf{a} = \frac{q}{\gamma m c}\mathbf{v} \cdot (\mathbf{v} \times \mathbf{B}) = 0\,,$$

so $\mathbf{a} \perp \mathbf{v}$: the acceleration is centripetal (curvature) rather than tangential (speed change).

To describe the motion, decompose the velocity into components parallel and perpendicular to $\mathbf{B}$,

$$\mathbf{v} = v_\parallel \hat{\mathbf{b}} + \mathbf{v}_\perp\,.$$

Projecting the Lorentz equation along $\mathbf{B}$ yields

$$\frac{\mathrm{d}}{\mathrm{d}t}(\gamma m v_\parallel) = \frac{q}{c}\,\mathbf{B} \times (\mathbf{v} \times \mathbf{B}) = 0 \quad \Longrightarrow \quad p_\parallel = \gamma m v_\parallel = \text{const}\,.$$

Because $p = \gamma m v$ is also constant (the speed is fixed), the *perpendicular* component is constant as well:

$$p_\perp = \sqrt{p^2 - p_\parallel^2} = p\,\sqrt{1-\mu^2} = \text{const}\,, \quad \mu \equiv \frac{p_\parallel}{p} = \cos\theta\,.$$

In other words, the *pitch angle* θ between the momentum and the field line is conserved.

In the plane orthogonal to $\mathbf{B}$ we have

$$\gamma m\,\frac{\mathrm{d}\mathbf{v}_\perp}{\mathrm{d}t} = \frac{q}{c}\,\mathbf{v}_\perp \times \mathbf{B}\,, \quad \frac{\mathrm{d}v_\parallel}{\mathrm{d}t} = 0\,,$$

i.e. uniform circular motion plus uniform drift along $\mathbf{B}$. Choosing coordinates with $\mathbf{B} = B\,\hat{\mathbf{z}}$ and introducing an initial phase ϕ_0, an explicit solution is

$$\begin{aligned}
v_x(t) &= -\,v_\perp \sin(\Omega_B t + \phi_0)\,, & x(t) &= x_g + r_g \cos(\Omega_B t + \phi_0)\,,\\
v_y(t) &= \quad\; v_\perp \cos(\Omega_B t + \phi_0)\,, & y(t) &= y_g + r_g \sin(\Omega_B t + \phi_0)\,,\\
v_z(t) &= v_\parallel = \beta c\,\mu\,, & z(t) &= z_g + v_\parallel t\,,
\end{aligned}$$

where (x_g, y_g, z_g) are the guiding–center coordinates, and $v_\perp = \Omega_B r_g$ ties the speed to the radius as in any circular motion. The path is a *helix* of fixed radius and constant pitch.

These results introduce two natural scales. The *gyro* (relativistic) angular frequency and the gyroradius are

$$\Omega_B = \frac{qB}{\gamma mc} = \frac{\Omega_0}{\gamma}, \quad r_g = \frac{v_\perp}{\Omega_B} = \frac{\gamma mc\, v_\perp}{qB} = \frac{p_\perp c}{qB}. \tag{A.21}$$

Here

$$\Omega_0 \equiv \frac{qB}{mc}$$

is the *cyclotron* (non-relativistic) angular frequency; dividing by 2π gives ordinary frequencies ν_0 and $\nu_B = \nu_0/\gamma$. The factor of γ has a clear interpretation: relativity *slows* the rotation as seen in the lab frame (time dilation), so the frequency drops by $1/\gamma$ while the radius grows by γ at fixed $v_\perp$.

Cosmic particles are often relativistic ($v_\perp \simeq c$), for which

$$r_g \simeq \frac{\gamma mc^2}{qB} \approx \frac{E}{qB}.$$

Magnetic rigidity. It is often convenient to define the *magnetic rigidity* $R \equiv pc/|q|$, so that

$$r_g = \frac{R}{B}.$$

For ultra-relativistic particles ($p \simeq E/c$) and charge number $Z = |q|/e$,

$$r_g \simeq 1.08 \text{ pc} \left(\frac{E}{\text{PeV}}\right)\left(\frac{Z\,B}{\mu\text{G}}\right)^{-1},$$

a handy back-of-the-envelope rule for Galactic and extragalactic settings.

For quick reference, Table 3 recaps the symbols and their relativistic/non-relativistic forms, and Table 4 provides order-of-magnitude scales relevant to cosmic-ray physics.

A.3 Threshold kinematics in astroparticle reactions

Many reactions relevant to high-energy astrophysics become kinematically allowed only once the available center-of-mass (CoM) energy exceeds a threshold: $\gamma\gamma$ pair production, photo-pion production, photodisintegration, hadronic pion production, and so on. It is therefore useful to collect a few compact formulas that translate a CoM threshold into a corresponding threshold energy in the laboratory (LAB) frame.

For simplicity, throughout this section we set $c = 1$.

General invariant threshold condition. Consider a generic reaction

$$a + b \longrightarrow 1 + 2 + \cdots + n,$$

Quantity	**Non-relativistic**	**Relativistic**
Angular frequency	$\Omega_0 = \dfrac{qB}{mc}$	$\Omega_B = \dfrac{qB}{\gamma mc} = \dfrac{\Omega_0}{\gamma}$
Ordinary frequency	$\nu_0 = \dfrac{\Omega_0}{2\pi} = \dfrac{1}{2\pi}\dfrac{qB}{mc}$	$\nu_B = \dfrac{\Omega_B}{2\pi} = \dfrac{\nu_0}{\gamma}$
Gyroradius	$r_0 = \dfrac{mc\,v_\perp}{qB}$	$r_g = \dfrac{\gamma mc\,v_\perp}{qB} = \gamma r_0$

Table 3: *Notation used in these notes.* Here $v_\perp$ is the speed perpendicular to $\mathbf{B}$; $p_\perp = \gamma m v_\perp$. For $v_\perp \simeq c$, $r_g \simeq \gamma mc^2/(qB)$.

Quantity	**Galactic CR**	**UHECR**
	$B \sim 1\ \mu$G, $E \sim 10$ GeV	$B \sim 1$ nG, $E \sim 1$ EeV
Angular freq. Ω_B	$\sim 10^{-3}$ s^{-1}	$\sim 10^{-14}$ s^{-1}
Ordinary freq. ν_B	$\sim 10^{-4}$ Hz	$\sim 10^{-15}$ Hz
Gyroradius r_g	$\sim 3\times10^{13}$ cm $\approx$ 10^{-5} pc	$\sim 3\times10^{24}$ cm $\approx$ Mpc

Table 4: *Representative scales* for a proton ($Z = 1$) with $v_\perp \simeq c$. For general Z, $r_g \propto 1/Z$ and $\Omega_B, \nu_B \propto Z$.

where a and b are the incoming particles and the final-state particles $1,\ldots,n$ have rest masses $m_1,\ldots,m_n$. Denote by P_a^μ and P_b^μ the incoming 4-momenta, and by P_i^μ the outgoing ones.

The key quantity is the Lorentz-invariant Mandelstam variable

$$s \equiv (P_a + P_b)^2.$$

We adopt the metric signature $(+,-,-,-)$ and $P^\mu = (E, \mathbf{p})$, so that $P^2 = m^2$. In the CoM frame, $\mathbf{p}_a + \mathbf{p}_b = 0$ and therefore

$$s = E_{\rm cm}^2,$$

with $E_{\rm cm}$ the total CoM energy.

At threshold the CoM energy is minimal, i.e. the final state carries no kinetic energy in the CoM frame (all outgoing particles have vanishing three-momenta). Hence

$$E_{\rm cm}^{(\rm thr)} = \sum_{i=1}^{n} m_i,$$

and the invariant threshold condition is

$$s_{\rm thr} = \left(\sum_{i=1}^{n} m_i\right)^2. \tag{A.22}$$

In practice, threshold energies are obtained as follows:

1. Compute s_{thr} from the final-state masses via Eq. (A.22).
2. Express s in the chosen LAB frame in terms of the incoming energies, momenta, and collision angle.
3. Set $s = s_{\text{thr}}$ and solve for the desired threshold energy.

General $2 \to n$ ***threshold for a fixed target.*** A common configuration in cosmic-ray physics is a projectile a impinging on a stationary target b in the LAB frame (e.g. $pp \to pp\pi$ processes).

Let m_a and m_b be the corresponding rest masses, and take particle b at rest:

$$P_a = (E_a, \mathbf{p}_a), \qquad P_b = (m_b, \mathbf{0}).$$

Then

$$s = (P_a + P_b)^2 = m_a^2 + m_b^2 + 2m_b\, E_a. \tag{A.23}$$

Setting $s = s_{\text{thr}}$ gives the LAB threshold *total* energy of the projectile:

$$E_a^{(\text{thr})} = \frac{s_{\text{thr}} - m_a^2 - m_b^2}{2m_b}. \tag{A.24}$$

This is the master formula for fixed-target thresholds.

Example: pion production. For proton–proton collisions producing a single neutral pion,

$$p + p \to p + p + \pi^0,$$

one has $m_a = m_b = m_p$ and

$$s_{\text{thr}} = (2m_p + m_{\pi^0})^2.$$

Inserting into Eq. (A.24) gives the threshold *total* energy of the projectile proton,

$$E_p^{(\text{thr})} = \frac{(2m_p + m_{\pi^0})^2 - 2m_p^2}{2m_p} = m_p + 2m_{\pi^0} + \frac{m_{\pi^0}^2}{2m_p} \simeq 1.22\ \text{GeV}. \tag{A.25}$$

General $2 \to n$ ***threshold for relativistic particles.*** In many radiative reactions the incoming particles are ultra-relativistic (e.g. $\gamma\gamma \to e^+e^-$ or $p\gamma$ processes), so that the available CoM energy depends sensitively on the collision geometry, i.e. on the angle θ between their momenta.

For two particles with 4-momenta P_a and P_b, masses m_a and m_b, energies E_a and E_b, and relative angle θ between $\mathbf{p}_a$ and $\mathbf{p}_b$, the invariant can be written as

$$s = m_a^2 + m_b^2 + 2E_aE_b\,(1 - \beta_a\beta_b \cos\theta)\,, \tag{A.26}$$

where $\beta_{a,b} \equiv v_{a,b}$. This form makes the angular dependence explicit:

- *Head-on*: $\theta = \pi$ $(\cos\theta = -1)$,

- *Parallel*: $\theta = 0$ ($\cos\theta = +1$).

For ultra-relativistic particles ($\beta_a \simeq \beta_b \simeq 1$),

$$s \simeq m_a^2 + m_b^2 + 2E_a E_b\,(1 - \cos\theta). \tag{A.27}$$

Since $1 - \cos\theta$ is maximal for $\theta = \pi$, head-on encounters maximize s at fixed E_a and E_b and are therefore the most effective for reaching threshold.

Setting $s = s_{\rm thr}$ and solving for (say) E_a yields the corresponding threshold condition,

$$E_a^{(\rm thr)} = \frac{s_{\rm thr} - m_a^2 - m_b^2}{2E_b\,(1 - \beta_a\beta_b\cos\theta)} \simeq \frac{s_{\rm thr} - m_a^2 - m_b^2}{2E_b(1 - \cos\theta)}, \tag{A.28}$$

where the last expression applies in the ultra-relativistic limit. For head-on collisions ($\theta = \pi$) this further reduces to

$$E_a^{(\rm thr)} \simeq \frac{s_{\rm thr} - m_a^2 - m_b^2}{4E_b}. \tag{A.29}$$

Example: $\gamma\gamma$ pair production. Consider electron–positron pair production via photon–photon annihilation,

$$\gamma + \gamma \to e^+ + e^-,$$

where the incoming photons have energies E_γ and ε, and collide at an angle θ. Since $m_a = m_b = 0$, the threshold invariant is

$$s_{\rm thr} = (2m_e)^2 = 4m_e^2.$$

Using the ultra-relativistic expression Eq. (A.28) with $m_a = m_b = 0$,

$$E_\gamma\,\varepsilon\,(1 - \cos\theta) \geq 2m_e^2. \tag{A.30}$$

For head-on collisions ($\theta = \pi$, so $1 - \cos\theta = 2$) this reduces to

$$E_\gamma\,\varepsilon \geq m_e^2, \tag{A.31}$$

or, equivalently,

$$E_\gamma^{(\rm thr)} = \frac{m_e^2}{\varepsilon} \simeq 0.26\ \text{TeV}\,\left(\frac{\varepsilon}{1\ \text{eV}}\right)^{-1}. \tag{A.32}$$

This relation underlies the connection between TeV γ rays and optical/IR background photons discussed in the main text.

Collider vs. fixed-target. It is instructive to compare fixed-target and collider configurations for symmetric $a + a$ reactions. Consider two identical particles of mass m and energy E colliding head-on. In the CoM frame their 4-momenta are

$$P_1 = (E, \mathbf{p}), \qquad P_2 = (E, -\mathbf{p}),$$

so that

$$s = (P_1 + P_2)^2 = (2E)^2 = 4E^2. \tag{A.33}$$

To reach a given threshold s_{thr}, a collider therefore requires a beam energy

$$E_{\text{beam}}^{(\text{collider})} = \frac{1}{2}\sqrt{s_{\text{thr}}}. \tag{A.34}$$

By contrast, for a fixed-target experiment with the same projectile and target mass m, Eq. (A.24) gives

$$E_{\text{proj}}^{(\text{thr})} = \frac{s_{\text{thr}} - 2m^2}{2m} \simeq \frac{s_{\text{thr}}}{2m} \qquad (s_{\text{thr}} \gg m^2). \tag{A.35}$$

Thus, at large s_{thr} the fixed-target threshold grows linearly with s_{thr}, whereas in a collider it grows only as $\sqrt{s_{\text{thr}}}$. This is why colliders are vastly more efficient than fixed-target setups at reaching very high CoM energies.

Cosmic rays, however, are observed at extreme *LAB* energies, and produce new particles when they collide with nuclei (or protons) approximately at rest in the atmosphere. A natural question is therefore: *What cosmic-ray energy E_{CR} is required to reach a CoM energy comparable to that of the LHC in a fixed-target collision on a proton at rest?*

For a cosmic-ray proton of energy E_{CR} colliding with a proton at rest, Eq. (A.23) yields

$$s = 2m_p^2 + 2m_p E_{\text{CR}} \simeq 2m_p E_{\text{CR}} \qquad (E_{\text{CR}} \gg m_p), \tag{A.36}$$

so that

$$E_{\text{CR}} \simeq \frac{s}{2m_p} = \frac{(\sqrt{s})^2}{2m_p}. \tag{A.37}$$

Taking $\sqrt{s} \equiv \sqrt{s_{\text{LHC}}} \simeq 13.6$ TeV and $m_p \simeq 0.94$ GeV, one finds

$$E_{\text{CR}} \simeq \frac{(1.36 \times 10^4 \text{ GeV})^2}{2 \times 0.94 \text{ GeV}} \simeq 10^{17} \text{ eV}.$$

Thus, only cosmic rays with energies $\gtrsim 10^{17}$ eV colliding on a proton at rest reach CoM energies comparable to those of the LHC.