

PREPARED FOR SUBMISSION TO JCAP

Towards a precision calculation of N_{eff} in the Standard Model II: Neutrino decoupling in the presence of flavour oscillations and finite-temperature QED

Jack J. Bennett,^a Gilles Buldgen,^b Pablo F. de Salas,^c
 Marco Drewes,^b Stefano Gariazzo,^{d,e} Sergio Pastor,^e and
 Yvonne Y. Y. Wong^a

^aSydney Consortium for Particle Physics and Cosmology, School of Physics, The University of New South Wales, Sydney NSW 2052, Australia

^bCentre for Cosmology, Particle Physics and Phenomenology, Université catholique de Louvain, Chemin du cyclotron, 2, Louvain-la-Neuve B-1348, Belgium

^cThe Oskar Klein Centre for Cosmoparticle Physics, Department of Physics, Stockholm University, SE-106 91 Stockholm, Sweden

^dINFN, Sezione di Torino, Via P. Giuria 1, I-10125 Torino, Italy

^eInstituto de Física Corpuscular (CSIC-Universitat de València), Parc Científic UV, C/ Catedrático José Beltrán, 2, E-46980 Paterna (Valencia), Spain

E-mail: j.j.bennett@unsw.edu.au, gilles.buldgen@uclouvain.be,
pablo.fernandez@fysik.su.se, marco.drewes@uclouvain.be, gariazzo@to.infn.it,
pastor@ific.uv.es, yvonne.y.wong@unsw.edu.au

Abstract. We present in this work a new calculation of the standard-model benchmark value for the effective number of neutrinos, $N_{\text{eff}}^{\text{SM}}$, that quantifies the cosmological neutrino-to-photon energy densities. The calculation takes into account neutrino flavour oscillations, finite-temperature effects in the quantum electrodynamics plasma to $\mathcal{O}(e^3)$, where e is the elementary electric charge, and a full evaluation of the neutrino–neutrino collision integral. We provide furthermore a detailed assessment of the uncertainties in the benchmark $N_{\text{eff}}^{\text{SM}}$ value, through testing the value’s dependence on (i) optional approximate modelling of the weak collision integrals, (ii) measurement errors in the physical parameters of the weak sector, and (iii) numerical convergence, particularly in relation to momentum discretisation. Our new, recommended standard-model benchmark is $N_{\text{eff}}^{\text{SM}} = 3.0440 \pm 0.0002$, where the nominal uncertainty is attributed predominantly to errors incurred in the numerical solution procedure ($|\delta N_{\text{eff}}| \sim 10^{-4}$), augmented by measurement errors in the solar mixing angle $\sin^2 \theta_{12}$ ($|\delta N_{\text{eff}}| \sim 10^{-4}$).

Contents

1	Introduction	1
2	The system	3
2.1	Continuity equation	3
2.2	Boltzmann equation	4
3	New elements in this calculation	5
3.1	New finite-temperature QED corrections	5
3.1.1	Finite-temperature corrections to the QED equation of state	5
3.1.2	Finite-temperature QED corrections to weak scattering rates	6
3.2	Full collision integral for neutrino–neutrino scattering	8
3.3	Assessment of remaining uncertainties	8
3.3.1	Approximate treatments of the weak collision integrals	9
3.3.2	Measurement errors in the physical parameters of the neutrino sector	11
4	Numerical convergence	12
4.1	Convergence variables	12
4.1.1	Momentum discretisation	12
4.1.2	Initialisation time	14
4.2	Energy density and number density conservation tests	15
5	Physics results	17
5.1	New finite-temperature QED corrections	18
5.2	Full collision integral for neutrino–neutrino scattering	18
5.3	Assessment of remaining uncertainties	19
5.3.1	Approximate treatment of the weak collision integrals	19
5.3.2	Measurement errors in the physical parameters of the neutrino sector	19
6	Conclusions	21
A	Collision integrals	23
B	Damping approximation	25
B.1	Damping coefficients	26

1 Introduction

The effective number of neutrinos, N_{eff} , is a parameter that quantifies the ratio of cosmological energy densities in neutrino-like relics¹ to photons in the early universe shortly after neutrino decoupling (at temperatures $T \sim 1$ MeV). Within the standard model (SM) of particle physics,² its expected theoretical value, $N_{\text{eff}}^{\text{SM}}$, is 3 (for three generations of SM neutrinos),

¹In the cosmological context, a neutrino-like relic is a thermalised or partially thermalised light particle state that decouples from the primordial plasma while ultra-relativistic and remains decoupled thereafter.

²In this work, the term “standard model of particle physics” is understood to include neutrino oscillations.

plus percent-level corrections due primarily to (i) energy transport from the primordial quantum electrodynamics (QED) plasma to the neutrino sector [1–13], and (ii) finite-temperature-induced deviations of the QED plasma equation of state from an ideal gas [14–19]. A 2016, fully momentum-dependent precision transport study including neutrino flavour oscillations and finite-temperature QED corrections performed by some of us (de Salas and Pastor) put the number at $N_{\text{eff}}^{\text{SM}} = 3.045$ [20]. This number has recently been revised to $N_{\text{eff}}^{\text{SM}} = 3.044$ in references [12, 13], including a new QED correction identified in reference [19].

Interest in the N_{eff} parameter as a probe of beyond-the-standard model (BSM) light relics — and by association the precision computation of $N_{\text{eff}}^{\text{SM}}$ — has a long history. In the era of precision cosmology, the continued interest in the N_{eff} parameter stems primarily from its effects on observables such as the cosmic microwave background (CMB) anisotropies; current measurements by the ESA Planck mission already constrain the parameter to $N_{\text{eff}} = 2.99_{-0.33}^{+0.34}$ (95% C.I.) [21], severely limiting the viable parameter space of BSM states produced well after the quantum chromodynamics (QCD) phase transition ($T \sim 200$ MeV) such as light sterile neutrinos [22, 23], eV-mass axions [24], and decay products of long-lived BSM particles [25, 26]. In the near future, a conservative projection of the scientific capacity of CMB-S4 sees the 1σ sensitivity to N_{eff} improve to $\sigma(N_{\text{eff}}) \sim 0.02 \rightarrow 0.03$ [27] depending on the final configuration of the experiment, potentially probing light relic production immediately after the QCD phase transition [28, 29].

Increasing observational precision calls for a corresponding improvement in the theoretical prediction of the SM benchmark $N_{\text{eff}}^{\text{SM}}$. Recently, three of us (Gariazzo, de Salas, and Pastor) revisited the energy transport aspect of the $N_{\text{eff}}^{\text{SM}}$ computation in the presence of neutrino flavour oscillations [30], while the other four (Bennett, Buldgen, Drewes and Wong) quantified the leading and several sub-leading finite-temperature QED contributions to $N_{\text{eff}}^{\text{SM}}$ [19]. In this work, we pool our resources to supply a state-of-the-art precision computation of the SM benchmark $N_{\text{eff}}^{\text{SM}}$. Relative to the 2016 update of [20], the main improvements in this new calculation are:

1. Energy transport in the presence of flavour oscillations is now computed using the new precision neutrino decoupling code `FortEPiano` [30]³, a fully momentum-dependent decoupling code that accepts up to 3 active + 3 sterile neutrino flavours.
2. We incorporate two new, subdominant finite-temperature corrections to the QED equation of state identified in [19], notably the $\mathcal{O}(e^3)$ correction, where e is the elementary electric charge. As we shall see, this correction will account for the leading-order shift in the SM benchmark $N_{\text{eff}}^{\text{SM}}$ relative to previously reported results.
3. We include a complete numerical evaluation of the neutrino–neutrino collision integral in the presence of neutrino flavour oscillations. Of particular note are the flavour-off-diagonal entries peculiar to oscillating systems: these had previously been modelled using a damping approximation in references [20, 30].
4. We provide a detailed assessment of the uncertainties in the benchmark $N_{\text{eff}}^{\text{SM}}$ value, through testing the value’s dependence on (i) the approximate modelling of the weak collision integrals in the presence of flavour oscillations, (ii) measurement errors in the physical parameters of the weak sector, and (iii) numerical convergence.

For the reader in search of a quick number, our new SM benchmark is $N_{\text{eff}}^{\text{SM}} = 3.0440 \pm 0.0002$, where the nominal uncertainty is attributed predominantly to errors incurred in the

³This code will be made publicly available at https://bitbucket.org/ahep_cosmo/fortepiano_public.

numerical solution procedure ($|\delta N_{\text{eff}}| \sim 10^{-4}$), followed by measurement uncertainties in the solar mixing angle $\sin^2 \theta_{12}$ ($|\delta N_{\text{eff}}| \sim 10^{-4}$). This new benchmark value is identical to that obtained in reference [13] to the same number of significant digits, which had also been computed with the full neutrino–neutrino collision integral and finite-temperature QED effects to $\mathcal{O}(e^3)$. For readers interested in the details of our calculation, the rest of the paper is organised as follows. We describe the physical system, including the relevant equations of motion, in section 2, and the new elements of the present calculation in section 3. Section 4 discusses the issue of numerical convergence with respect to the initialisation and momentum discretisation procedures, in terms of controlled tests of detailed balance. Our calculations of $N_{\text{eff}}^{\text{SM}}$ under the introduction of various new elements are presented in section 5. Section 6 contains our conclusions. Two appendices document new technical details of our calculations.

2 The system

The SM benchmark effective number of neutrinos $N_{\text{eff}}^{\text{SM}}$ is defined via the ratio of the neutrino energy density ρ_ν to the photon energy density ρ_γ in the limit $T/m_e \rightarrow 0$,

$$\left. \frac{\rho_\nu}{\rho_\gamma} \right|_{T/m_e \rightarrow 0} \equiv \frac{7}{8} \left(\frac{4}{11} \right)^{4/3} N_{\text{eff}}^{\text{SM}}, \quad (2.1)$$

where T is the photon temperature and m_e is the electron mass. Its precision computation typically requires that we solve two sets of equations of motion across the relic neutrino decoupling epoch, in a time frame corresponding to $T \sim \mathcal{O}(100) \rightarrow \mathcal{O}(0.01)$ MeV. These two sets of equations are: (i) the continuity equation, which tracks the total energy density of the universe in the time frame of interest, and (ii) the generalised Boltzmann equation — also known as the quantum kinetic equation — for the density matrix of the neutrino ensemble, which follows the non-equilibrium dynamics of the ensemble in the presence of flavour oscillations and particle scattering.

2.1 Continuity equation

In a Friedmann–Lemaître–Robertson–Walker (FLRW) universe, the continuity equation reads $(d/dt)\rho_{\text{tot}} + 3H(\rho_{\text{tot}} + P_{\text{tot}}) = 0$, where ρ_{tot} and P_{tot} are, respectively, the total *physical* energy density and pressure, $H \equiv (1/a)(da/dt)$ the Hubble expansion rate, and a is the scale factor. In terms of the rescaled time, momentum, and temperature coordinates, $x \equiv am_e$, $y \equiv ap$, and $z \equiv aT_\gamma$, the continuity equation is given equivalently by (e.g., [18])

$$\frac{d}{dx} \bar{\rho}_{\text{tot}}(x, z(x)) = \frac{1}{x} [\bar{\rho}_{\text{tot}}(x, z(x)) - 3\bar{P}_{\text{tot}}(x, z(x))], \quad (2.2)$$

where $\bar{\rho}_{\text{tot}} \equiv (x/m_e)^4 \rho_{\text{tot}}$ and $\bar{P}_{\text{tot}} \equiv (x/m_e)^4 P_{\text{tot}}$ are now the *comoving* quantities. Thus, in using rescaled and comoving variables, we have factored out the effects of cosmic expansion.

For the physical system at hand, $\rho_{\text{tot}} \equiv \rho_{\text{QED}} + \rho_\nu$ and $P_{\text{tot}} \equiv P_{\text{QED}} + P_\nu$, both of which sum over the QED and the neutrino sectors.⁴ We assume the QED sector to be always in a state of thermodynamic equilibrium, so that ρ_{QED} and P_{QED} are connected by the standard thermodynamic relation $\rho_{\text{QED}} = -P_{\text{QED}} + T(\partial/\partial T)P_{\text{QED}}$.

⁴In the time frame of interest, the QED plasma is composed of photons, electrons/positrons, and muons/anti-muons, while the neutrino sector comprises three families of SM neutrinos and anti-neutrinos. We neglect the small, $\mathcal{O}(10^{-10})$ universal matter–antimatter asymmetry, so that a particle name is taken to refer to both the particle and its anti-particle.

2.2 Boltzmann equation

Schematically, the generalised Boltzmann equation for the one-particle reduced density matrix of the neutrino ensemble, $\varrho(t, p)$, in an FLRW universe is given by $\partial_t \varrho - pH \partial_p \varrho = -i[\mathbb{H}, \varrho] + \mathcal{I}[\varrho]$, where $[\mathbb{H}, \varrho] \equiv \mathbb{H}\varrho - \varrho\mathbb{H}$ denotes a commutator between the flavour oscillations Hamiltonian \mathbb{H} and ϱ , and the collision integral $\mathcal{I}[\varrho]$ encapsulates all non-unitary (scattering) effects on ϱ [31]. In terms of the rescaled coordinates, the Boltzmann equation reads (e.g., [6])

$$\frac{d\varrho(x, y)}{dx} = \frac{1}{m_e} \frac{m_e^4}{x^4} \sqrt{\frac{3m_{\text{Pl}}^2}{8\pi\bar{\rho}_{\text{tot}}}} \left(-i \left[\mathbb{H}(x, y, z(x)), \varrho(x, y) \right] + \mathcal{I}[\varrho(x, y)] \right), \quad (2.3)$$

where m_{Pl} is the Planck mass, and we note in passing that the prefactor $(m_e/x)^2 \sqrt{3m_{\text{Pl}}^2/8\pi\bar{\rho}_{\text{tot}}}$ is but the inverse Hubble expansion rate H^{-1} .

Working in the flavour basis, the oscillations Hamiltonian [31],

$$\mathbb{H}(x, y, z(x)) = \frac{x^6}{m_e^6} \frac{U\mathbb{M}U^\dagger}{2y} - 2\sqrt{2}G_F y \left(\frac{\mathbb{E}_\ell(z(x)) + \mathbb{P}_\ell(z(x))}{m_W^2} + \frac{4}{3} \frac{\mathbb{E}_\nu(x, z(x))}{m_W^2 \cos^2 \theta_W} \right), \quad (2.4)$$

comprises a vacuum and an in-medium part. The former consists of the neutrino squared-mass difference matrix $\mathbb{M} = \text{diag}(0, \Delta m_{21}^2, \Delta m_{31}^2)$ and the vacuum mixing matrix $U = R(\theta_{23})R(\theta_{13})R(\theta_{12})$ parameterised by three Euler rotation angles θ_{23} , θ_{13} , and θ_{12} . We explicitly set the Dirac CP phase, δ_{CP} , to zero in this analysis. In this way, there is no neutrino asymmetry, and we can consider the neutrino momentum distribution equal to the antineutrino one. We do note, however, that global fits are beginning to indicate a preferred δ_{CP} value in the range $[\pi, 2\pi]$ [32–35]. The CP -conserving value $\delta_{CP} = 0$ is excluded at approximately 3σ , for both neutrino mass orderings.

The in-medium part of the Hamiltonian (2.4) contains a CP -symmetric (by assumption) correction to the neutrino dispersion relation [36], i.e., the ‘‘matter potential’’, proportional to the Fermi constant G_F , with m_W the W boson mass and θ_W the weak mixing angle. The terms \mathbb{E}_ℓ and \mathbb{P}_ℓ are momentum-integrals of some combinations of the charged-lepton energy $\epsilon_\ell(y) \equiv (y^2 + a^2 m_\ell^2)^{1/2}$ and (equilibrium) occupation number $f_\ell(y)$,

$$\mathbb{E}_\ell \equiv \frac{1}{\pi^2} \int dy y^2 \text{diag} \left(\epsilon_e(y) f_e(y), \epsilon_\mu(y) f_\mu(y), 0 \right)_{\text{ideal gas}} = \text{diag}(\bar{\rho}_e, \bar{\rho}_\mu, 0), \quad (2.5)$$

$$\mathbb{P}_\ell \equiv \frac{1}{3\pi^2} \int dy y^2 \text{diag} \left(\frac{y^2}{\epsilon_e(y)} f_e(y), \frac{y^2}{\epsilon_\mu(y)} f_\mu(y), 0 \right)_{\text{ideal gas}} = \text{diag}(\bar{P}_e, \bar{P}_\mu, 0), \quad (2.6)$$

which coincide respectively with the comoving energy density and pressure of the relevant charged leptons in the ideal gas limit; $\mathbb{E}_\nu \equiv \pi^{-2} \int dy y^3 \varrho(y)$ is the equivalent for an ultra-relativistic neutrino gas. Note that the $\mathbb{E}_\ell + \mathbb{P}_\ell$ term in equation (2.4) differs from its usual presentation found in, e.g., equation (2.2) of [30], which has $\mathbb{E}_\ell + \mathbb{P}_\ell$ replaced with $(4/3) \mathbb{E}_\ell$. First reported in [36], the former is in fact the more general result, while $(4/3) \mathbb{E}_\ell$ applies strictly only when the charged leptons are ultra-relativistic. In practice, however, using the incorrect expression incurs an error no larger than 10^{-5} in $N_{\text{eff}}^{\text{SM}}$.

The collision integral $\mathcal{I}[\varrho]$ incorporates in principle all weak scattering processes wherein at least one neutrino appears in either the initial or final state. As in the previous update [20], however, we account only for $2 \rightarrow 2$ processes involving (i) two neutrinos and two electrons anyway distributed in the initial and final states, and (ii) neutrino–neutrino scattering. Then,

schematically, $\mathcal{I}[\varrho] = \mathcal{I}_{\nu e}[\varrho] + \mathcal{I}_{\nu\nu}[\varrho]$ comprises 9D momentum-integrals that, at tree level, can be systematically reduced to 2D integrals of the forms

$$\mathcal{I}_{\nu e}[\varrho(y)] \propto G_F^2 \int dy_2 dy_3 \Pi_{\nu e}(y, y_2, y_3; x) F_{\nu e}(\varrho(y), \varrho(y_2), f_e(y_3), f_e(y_4)), \quad (2.7)$$

$$\mathcal{I}_{\nu\nu}[\varrho(y)] \propto G_F^2 \int dy_2 dy_3 \Pi_{\nu\nu}(y, y_2, y_3; x) F_{\nu\nu}(\varrho(y), \varrho(y_2), \varrho(y_3), \varrho(y_4)), \quad (2.8)$$

where $\Pi_{\nu e}$ and $\Pi_{\nu\nu}$ are scalar functions representing the scattering kernels, and the phase space matrices $F_{\nu e}$ and $F_{\nu\nu}$ are products of $\varrho(y)$ including quantum statistics. We defer to section 3.3.1 the discussion of the exact forms and our implementation of $\mathcal{I}[\varrho]$ in the solution procedure, but note here that this is one of the first studies in which the neutrino–neutrino scattering integral $\mathcal{I}_{\nu\nu}[\varrho(y)]$ is solved numerically “as is” without approximations.

Then, noting that $\bar{\rho}_\nu = \text{Tr}[\mathbb{E}_\nu]$ and correspondingly $\bar{P}_\nu = \bar{\rho}_\nu/3$, equations (2.2) to (2.8) form a closed system which we shall solve numerically using the precision neutrino decoupling code `FortEPiNO` [30].

3 New elements in this calculation

Aside from a new code, three aspects of our new computation of $N_{\text{eff}}^{\text{SM}}$ sets it apart from most existing calculations. Firstly, new and significant finite-temperature QED corrections to $\mathcal{O}(e^3)$ identified in [19] are incorporated into the precision calculation, including an $\mathcal{O}(e^2)$ logarithmic correction that has never been implemented in previous calculations and whose impact we demonstrate in this work to be below the current numerical error. Secondly, we solve numerically the full neutrino–neutrino collision integral — without approximation — in the presence of neutrino oscillations, as was previously attempted only in reference [13]. Thirdly, we identify and quantify the various sources of errors that can impact on the accuracy with which the benchmark $N_{\text{eff}}^{\text{SM}}$ can be computed theoretically; our treatment in this regard is similar to that of reference [13], but describes more extensively the impact of neutrino oscillation parameters beyond the constraints from oscillation experiments. We discuss these three new elements in three subsections below.

3.1 New finite-temperature QED corrections

Finite-temperature QED affects $N_{\text{eff}}^{\text{SM}}$ in two distinct ways [19]. The dominant effect is an altered equation of state for the QED plasma, through which the bulk thermodynamic quantities of the plasma — such as ρ_{QED} and P_{QED} that appear in the continuity equation (2.2) — are modified away from their ideal gas predictions in a temperature-dependent manner. A secondary effect are temperature-dependent corrections to the scattering rates of the constituent QED (quasi)particles with neutrinos; these enter at a practical level into the neutrino Boltzmann equation (2.3) through the $\mathcal{I}_{\nu e}[\varrho]$ collision integral (2.7) and, albeit minutely, through the in-medium neutrino dispersion relation in the flavour oscillation Hamiltonian (2.4).

3.1.1 Finite-temperature corrections to the QED equation of state

The finite-temperature QED partition function Z is well established in the literature (see, e.g., [37]), to a level of precision (as measured by powers in the elementary electric charge e) higher than is necessary for the present calculation [19]. The diagrammatic representation of the $\mathcal{O}(e^2)$ and $\mathcal{O}(e^3)$ corrections to $\ln Z$ used in this work are shown in figure 1. Contributions at higher orders in e have been estimated in the instantaneous decoupling limit to contribute

$$\ln Z^{(2)} + \ln Z^{(3)} = -\frac{1}{2} \left[\text{Diagram 1} \right] + \frac{1}{2} \left[\frac{1}{2} \left[\text{Diagram 2} \right] - \frac{1}{3} \left[\text{Diagram 3} \right] + \frac{1}{4} \left[\text{Diagram 4} \right] + \dots \right]$$

Figure 1: Diagrammatic expression for the $\mathcal{O}(e^2)$ and $\mathcal{O}(e^3)$ corrections to the finite-temperature QED partition function. The filled black circles represent one-particle irreducible photon self-energies at finite temperature.

at most $\delta N_{\text{eff}} \sim 4 \times 10^{-6}$ [19] — substantially below the intrinsic numerical noise of our computations (see section 4) — and are thus neglected in the benchmark calculation.

The leading-order, $\mathcal{O}(e^2)$ correction contributes $\delta N_{\text{eff}} \sim 0.01$ [14, 15, 17, 19], and is often discussed in two parts: a logarithmic contribution and a log-independent one (see equation (4.7) of [19]). The latter contribution modifies the continuity equation (2.2) as described in [18] (see also appendix B of [30]), and is now a staple ingredient in the computation of $N_{\text{eff}}^{\text{SM}}$ [6–12, 20, 30]. The former, on the other hand, has been hitherto regularly neglected, but should contribute a borderline-significant $\delta N_{\text{eff}} \sim -5 \times 10^{-5}$ [19]. In this work, we include for the first time this $\mathcal{O}(e^2)$ logarithmic contribution in the precision computation of $N_{\text{eff}}^{\text{SM}}$, through further modifications to the continuity equation (2.2) following [19].

Of greater interest, however, is the new, sub-leading $\mathcal{O}(e^3)$ correction arising from the resummation of ring diagrams depicted enclosed in square brackets in figure 1. Reference [19] estimates a sizeable contribution from this set of diagrams: $\delta N_{\text{eff}} \sim -10^{-3}$, larger than the effect of including neutrino flavour oscillations. Following the procedure laid down in [19], the said correction was recently adopted in [12, 13] in their computation of $N_{\text{eff}}^{\text{SM}}$. We likewise incorporate this $\mathcal{O}(e^3)$ correction into our present precision calculation.

3.1.2 Finite-temperature QED corrections to weak scattering rates

We distinguish four types of finite-temperature QED corrections to the 4-Fermi contact interaction that describes $2 \rightarrow 2$ neutrino–electron scattering at leading order: (a) modification to the dispersion relation, (b) vertex corrections, (c) real emission or absorption, and (d) closed fermion loops [38, 39]. These are depicted diagrammatically in figure 2.

Contributions of the type (a) amount to dressing the fermionic QED-charged propagator with a photon [40], which, in the quasiparticle approximation, can be effectively captured by way of a modified, in-medium dispersion relation in the evaluation of the relevant Feynman diagrams. In practice, diagrams of this sort are often evaluated using partially resummed propagators, wherein the (vacuum) particle mass is shifted to its thermal counterpart obtained from the self-energy computed to some fixed order in e — hence the common name “thermal mass correction”.⁵ In doing so, we actually resum infinitely many QED self-interactions, but at the same time also neglect the subdominant effects of additional quasiparticle solutions due to collective excitations in the plasma [44, 45].

In practice, implementation of the above description amounts to replacing all occurrences of m_e^2 in the 9D weak collision integral $\mathcal{I}[\rho]$ with its thermal counterpart $m_e^2 + \delta m_e^2(p, T)$, where

⁵The procedure of shifting the pole mass of the thermal propagator to account for finite-temperature corrections is generally valid at the *diagrammatic* level within the so-called quasiparticle approximation [41–43]. It is *not* a valid procedure, however, to compute finite-temperature corrections to bulk thermodynamic quantities (e.g., energy density, pressure, etc.) by replacing the vacuum particle mass with its thermal counterpart in expressions that have been established *originally for an ideal gas*. See [19] for a discussion of this issue in relation to the $N_{\text{eff}}^{\text{SM}}$ computation.

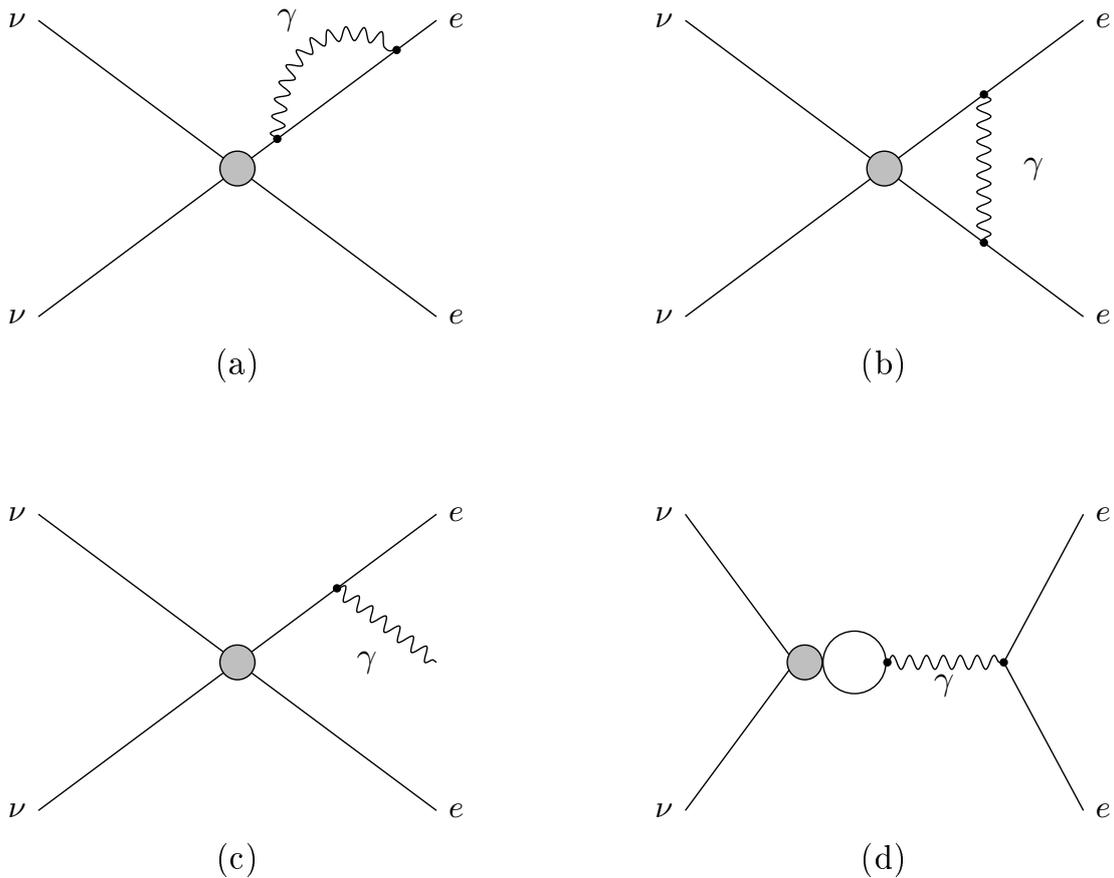


Figure 2: The only four qualitatively different finite-temperature QED corrections to the weak scattering rates.

we take $\delta m_e^2(p, T)$ to be the $\mathcal{O}(e^2)$, one-loop electron self-energy given in, e.g., equation (4.38) of [19]. This is an especially trivial task if we ignore the momentum-dependent piece in the said equation — which has been shown to contribute less than 10% to the total $\delta m_e^2(p, T)$ [18] — so that the reduction of $\mathcal{I}[\rho]$ to 2D proceeds identically as in the case with bare propagators. The same procedure can also be applied to equations (2.5) and (2.6) to approximately correct the matter potentials from the charged-lepton background.

Such a simple implementation has previously been adopted for $\mathcal{I}[\rho]$ in the precision calculations of [20, 30], as well as in the estimates of [19] which found this correction to contribute $\delta N_{\text{eff}} \sim -5 \times 10^{-5}$. For simplicity and in view of the size of the correction, we shall adopt the same approach in this work and include in addition a similarly-corrected matter potential. We stress however that the procedure is, strictly speaking, inconsistent, and leave a careful treatment of thermal-mass correction to the weak rates for future work.

Contributions of the types (b), (c) and (d) unfortunately cannot be dealt with in a similar fashion. An explicit evaluation of each individual diagram is required to assess their significance — a tedious task (see, e.g., [38, 39] for the zero-temperature case) that we also leave for a future publication.

3.2 Full collision integral for neutrino–neutrino scattering

As discussed in section 2.2, we split the weak collision integral $\mathcal{I}[\varrho]$ into two parts, $\mathcal{I}_{\nu e}[\varrho]$ and $\mathcal{I}_{\nu\nu}[\varrho]$, accounting respectively for $2 \rightarrow 2$ neutrino–electron and neutrino–neutrino scattering processes. In the flavour basis and assuming the neutrino density matrix $\varrho(y)$ to be diagonal, the *diagonal* elements of both $\mathcal{I}_{\nu e}[\varrho]$ and $\mathcal{I}_{\nu\nu}[\varrho]$, $\{\mathcal{I}[\varrho(y)]\}_{\alpha\alpha}$, where $\alpha = e, \mu, \tau$ is a flavour index, are identically the standard Boltzmann collision integral due to the said $2 \rightarrow 2$ scattering processes for the occupation number of ν_α at mode y . The *off-diagonal* elements, $\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta}$, where $\alpha \neq \beta$, are on the other hand peculiar to flavour oscillations and responsible for quantum decoherence of the neutrino ensemble.

The general, 9D forms of $\mathcal{I}_{\nu e}[\varrho]$ and $\mathcal{I}_{\nu\nu}[\varrho]$ have been determined in [31, 46]. Practical implementation, however, requires that we first reduce them to 2D integrals of the forms (2.7) and (2.8), wherein each integrand comprises a scattering kernel Π and a phase space matrix F . Many integral reduction procedures exist for the purpose (e.g., [3, 19, 22]), differing from each other only in the order in which the angular dependences of the integrand are eliminated in the intermediate steps. For $\mathcal{I}_{\nu e}[\varrho]$ we use the expressions presented in appendix A of [20], the same expressions that have been hard-coded in `FortEPiANO` for a previous study [30]. For completeness, we present them again in appendix A of this work.

On the other hand, the $\mathcal{I}_{\nu\nu}[\varrho]$ collision integral (2.8) is highly nonlinear: its phase space matrix $F_{\nu\nu}$ couples $\varrho(y)$ at four different modes. For this reason, with the exception of the recent study of [13] which does consider the integral in full, $\mathcal{I}_{\nu\nu}[\varrho]$ has thus far mostly been treated only in an approximate fashion. The 2016 update of [20], for example, replaced its diagonal entries with the standard Boltzmann collision integrals for the occupation numbers — a procedure equivalent to assuming a flavour-diagonal $\varrho(y)$ in the evaluation of $F_{\nu\nu}$ — while modelling the off-diagonal elements using a damping approximation (to be discussed in subsection 3.3.1). Reference [30] likewise adopted the off-diagonal damping approximation, but opted to neglect the diagonal entries. Meanwhile, the recent study of [12], while treating the diagonal entries in the manner of [20], did not model the off-diagonal ones at all.

In the present work, we incorporate the full $\mathcal{I}_{\nu\nu}[\varrho]$ collision integral “as is” in the precision computation of the benchmark $N_{\text{eff}}^{\text{SM}}$. The complete expression of its 2D reduced form can be found in appendix A. Because it couples $\varrho(y)$ at four different modes y , numerical evaluation of the phase space matrix $F_{\nu\nu}$ always requires a minimum of one interpolation of $\varrho(y)$ in y -space in order to enforce energy conservation amongst the four neutrino states participating in a $\nu\nu \rightarrow \nu\nu$ collision process. Linear interpolation between two nodes closest to the desired y -value suffices for the off-diagonal entries $\{\varrho(y)\}_{\alpha\beta}$, $\alpha \neq \beta$. For the diagonal entries, $\{\varrho(y)\}_{\alpha\alpha}$, we find that interpolating the normalised $\{\varrho(y)\}_{\alpha\alpha}/f_{\text{eq}}(y)$, where f_{eq} is the relativistic Fermi–Dirac distribution, generally offers better stability than interpolating the bare $\{\varrho(y)\}_{\alpha\alpha}$. The most critical numerical issue in this regard, however, is the choice of momentum discretisation scheme. We shall discuss this in detail in section 4.

3.3 Assessment of remaining uncertainties

Aside from the $\mathcal{O}(e^4)$ and higher finite-temperature corrections to the QED equation of state that have been previously quantified in [19], uncertainties in the SM benchmark $N_{\text{eff}}^{\text{SM}}$ can arise also from our treatment of the weak sector. To this end, we identify three classes of errors: (i) optional physical approximations provided in `FortEPiANO` to stabilise and expedite the evaluation of the weak collision integrals, (ii) measurement errors in the physical parameters of the neutrino sector, and (iii) numerical non-convergence originating in the discretisation and initialisation procedures of the solution scheme. The last class of errors, class (iii), is

inherent in all numerical approximations to solutions of differential equations, and warrants a separate discussion in section 4. Error sources (i) and (ii) are physical in origin, which we describe in two subsections below.

3.3.1 Approximate treatments of the weak collision integrals

Since tracking the decoupling of the neutrino sector from the QED plasma is the main objective of our calculation, the minimum set-up of `FortEPiANO` always includes a full and non-negotiable numerical evaluation of the diagonal components of the neutrino–electron collision integral $\mathcal{I}_{\nu e}[\varrho]$. Beyond this bare minimum requirement, however, several approximations and/or alternative implementations of the remaining $\mathcal{I}[\varrho]$ terms may be considered to facilitate their computation. Reference [30], for example, assumed $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$ and a damping approximation for all off-diagonal entries in what we shall call the **minimum set-up**.

Diagonal neutrino–neutrino collision integral. There are good reasons to think that the diagonal entries of the neutrino–neutrino collision integral, $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$, may be effectively dispensable. Phenomenologically, the $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ terms serve to transport energy between different neutrino flavours, a role that may conceivably be fulfilled to a good extent by large-mixing flavour oscillations in those studies that account for the latter. In view that $\mathcal{I}_{\nu\nu}[\varrho]$ is highly nonlinear and its full evaluation comes with substantial numerical uncertainties (see section 5.2), alternative implementations of $\mathcal{I}_{\nu\nu}[\varrho]$ may be desirable for a numerically stable outcome, *provided* of course that accuracy is not compromised.

In this work, we test three different implementations of $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$:

1. $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$, as in [30]: the minimum set-up uses this approximation;
2. As in [12, 20], we evaluate $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ assuming a **diagonal** $\varrho(y)$; and
3. A **full** evaluation of $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ without assumptions;

and we reiterate that the last named option is itself a novel element of the present work. See section 3.2.

Off-diagonal damping approximation. If the deviation of the neutrino ensemble from equilibrium is minimal, the off-diagonal entries of the weak collision integral, $\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta}$, where $\alpha \neq \beta$, can be systematically manipulated into the form

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta} = -\{D(x, y, z)\}_{\alpha\beta}\{\varrho(y)\}_{\alpha\beta}, \quad (3.1)$$

where $\{D(x, y, z)\}_{\alpha\beta} = \{D_{\nu e}(x, y, z)\}_{\alpha\beta} + \{D_{\nu\nu}(x, y, z)\}_{\alpha\beta}$ are flavour-dependent damping coefficients. Equation (3.1) is the so-called “damping approximation”. At the practical level, the approximation effectively decouples the evolution of $\{\varrho(y)\}_{\alpha\beta}$ not only from other momentum modes $y' \neq y$ but also from other entries of $\varrho(y)$ at the same mode y , and thus, where it is valid, offers an extremely efficient computational pathway.⁶

In the context of precision $N_{\text{eff}}^{\text{SM}}$ computation and related topics (e.g., sterile neutrino thermalisation), there are historically two understandings of the damping approximation (3.1) and hence two disparate sets of damping coefficients in the existing literature derived under different assumptions:

⁶There is in principle also a similar damping approximation for the diagonal entries of $\mathcal{I}[\varrho]$. See appendix B. We do not however use it in our computation of $N_{\text{eff}}^{\text{SM}}$.

1. At each mode y , standard linear response instructs us to equate all occurrences of $\varrho(y')$ in the phase space matrix F — *except* for $y' = y$ — to their equilibrium expectations, i.e., flavour-diagonal and $\{\varrho(y)\}_{\alpha\alpha} = f_{\text{eq}}(y)$, where $f_{\text{eq}}(y)$ is some equilibrium distribution. The procedure immediately renders $\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta}$ into the form (3.1), where any remaining integration can be immediately performed to yield the damping coefficients.

This approach was previously used in [47] to compute $\{D(x, y, z)\}_{\alpha\beta}$ in the $m_e = 0$ (i.e., $x = 0$) limit assuming Maxwell–Boltzmann statistics. In this work, we generalise these $x = 0$ calculations to Fermi–Dirac statistics including Pauli blocking, and find

$$\{D_{\nu\nu}(y, z)\}_{\alpha\beta} = \mathcal{D}(y, z) \equiv \frac{2 G_F^2 y z^4}{(2\pi)^3} d(y/z), \quad (3.2)$$

$$\{D_{\nu e}(y, z)\}_{\alpha\beta} = \frac{1}{8} [(2 \sin^2 \theta_W \pm 1)_\alpha^2 + (2 \sin^2 \theta_W \pm 1)_\beta^2 + 8 \sin^4 \theta_W] \mathcal{D}(y, z), \quad (3.3)$$

where the notation $(2 \sin^2 \theta_W \pm 1)_\alpha$ indicates that the “+” sign applies to $\alpha = e$ and “−” to $\alpha = \mu, \tau$. The function $d(s = y/z)$ is a number of order 100 shown in figure (B.10). Its exact form — expressed as a double momentum integral — is given in equation (B.9); for computational ease, however, we also supply a fitting function $d_{\text{fit}}(s)$ in equation (B.11), which reproduces $d(s)$ in the interval $s \in [10^{-4}, 10^3]$ to better than 0.25% accuracy. Details of the calculation can be found in appendix B.

2. Reference [46] proposed to simplify $\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta}$ under the ansatz

$$\varrho(y) = \frac{f_{\text{eq}}(y)}{f_{\text{eq}}(\langle y \rangle)} \varrho(\langle y \rangle), \quad (3.4)$$

such that $\varrho(y)$ at all modes y evolve *in phase*, where $\langle y \rangle$ denotes a representative momentum. Upon integration in y , the procedure yields a thermally-averaged $\langle \{\mathcal{I}[\varrho(y)]\}_{\alpha\beta} \rangle$, expressed in the damping form (3.1) in terms of $\varrho(\langle y \rangle)$ and a set of thermally-averaged damping coefficients $\langle \{D(x, y, z)\}_{\alpha\beta} \rangle$. In introducing the ansatz (3.4), the primary motivation of [46] was to reduce the generalised Boltzmann equation (2.3) to a single set of “quantum rate equations” evaluated at the representative momentum $\langle y \rangle$. We note however that some subsequent works (e.g., [48]) invoked a scaling argument to obtain y -dependent damping coefficients such as $\{D(x, y, z)\}_{\alpha\beta} = [y/\langle y \rangle] \langle \{D(x, y, z)\}_{\alpha\beta} \rangle$.

Physically, the ansatz (3.4) effectively removes those parts of the collision integral $\mathcal{I}[\varrho]$ associated with flavour-blind elastic scattering. The y -scaled versions of the damping coefficients derived under this scheme, particularly for $\mathcal{I}_{\nu\nu}[\varrho]$, have historically been employed in the precision $N_{\text{eff}}^{\text{SM}}$ calculations of [20, 30]. We do not consider them in this work, however, as our goals differ substantively from the original intentions of [46] that motivated the ansatz (3.4).⁷

In the present study, we test two different implementations of $\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta}$: (i) a full, real-time numerical evaluation of all entries, and (ii) $\{\mathcal{I}_{\nu e}[\varrho(y)]\}_{\alpha\beta}$ and $\{\mathcal{I}_{\nu\nu}[\varrho(y)]\}_{\alpha\beta}$ computed under the damping approximation (3.1), with damping coefficients (3.2) and (3.3) obtained from linear response. The latter implementation corresponds to the minimum set-up.

⁷One could also interpret the ansatz (3.4) physically as an assumption of a time-scale hierarchy between flavour-blind elastic scattering and other (inelastic and flavour-dependent) processes, whereby the former are taken to be always fast enough to set all $\varrho(y)$ in phase, while the latter determine the slower aspects of $\varrho(y)$ ’s evolution. For ultrarelativistic neutrinos that only have weak interactions and hence only one collision time-scale, the assumption of such a hierarchy is unfortunately ill-conceived. The “in phase” assumption is likewise not borne out by our numerical results. These constitute more reasons to reject the ansatz (3.4).

	Parameter [Units]	Value $\pm 1\sigma$ uncertainty	Reference
QED	$\alpha/10^{-3}$	$7.2973525693 \pm 0.0000000011$	[49]
	m_e [MeV]	$0.51099895000 \pm 0.00000000015$	[49]
Weak	$\sin^2 \theta_W$	0.23871 ± 0.00009	[50, 51]
	g_L	0.727	[50, 51]
	\tilde{g}_L	-0.273	[50, 51]
	g_R	0.233	[50, 51]
	G_F [10^{-5}GeV^{-2}]	1.1663787 ± 0.0000006	[52]
	m_W [GeV]	80.379 ± 0.012	[52]
Neutrino	$\sin^2 \theta_{12}/10^{-1}$	3.18 ± 0.16	[33]
	$\sin^2 \theta_{13}/10^{-2}$	$2.200^{+0.069}_{-0.062}$	[33]
	$\sin^2 \theta_{23}/10^{-1}$	5.74 ± 0.14	[33]
	Δm_{21}^2 [10^{-5}eV^2]	$7.50^{+0.22}_{-0.20}$	[33]
	Δm_{31}^2 [10^{-3}eV^2] (NO)	$2.55^{+0.02}_{-0.03}$	[33]
	Δm_{31}^2 [10^{-3}eV^2] (IO)	$-2.45^{+0.03}_{-0.02}$	[33]

Table 1: Central values and 1σ uncertainties of the physical constants used in this work. The neutrino parameter values have been derived in the global fit [33] assuming a normal mass ordering (NO), with the exception of the last entry which assumes an inverted mass ordering (IO). The weak mixing angle value quoted here corresponds to $\sin^2 \theta_W(0)_{\overline{\text{MS}}}$ in the modified minimal subtraction ($\overline{\text{MS}}$) scheme, established via SM renormalisation group running from the measured value of $\sin^2 \theta_W(m_Z)_{\overline{\text{MS}}} = 0.23124 \pm 0.00006$ at the Z -pole [50, 51].

3.3.2 Measurement errors in the physical parameters of the neutrino sector

The fine structure constant $\alpha \equiv e^2/4\pi$ and the electron mass m_e have both been experimentally determined to nine significant digits [49] (see table 1). In the context of computing finite-temperature corrections to $N_{\text{eff}}^{\text{SM}}$, these parameters are essentially infinitely well known.

Likewise, the small uncertainties in the weak sector constants — G_F , $\sin^2 \theta_W$, m_W , g_L , \tilde{g}_L , and g_R — are not expected to impact on $N_{\text{eff}}^{\text{SM}}$ by more than $|\delta N_{\text{eff}}| \sim 10^{-5}$, the intrinsic numerical noise of the computation (see section 4). To see this, note that as a rule of thumb, shifting the neutrino decoupling temperature T_d by a fractional $\sim 0.1\%$ induces a change of $|\delta N_{\text{eff}}| \lesssim 10^{-4}$ [19]. Then, by simple power counting, varying G_F even by as much as 3σ away from its central values (i.e., relative change of $\sim \pm 10^{-4}\%$) can only have a negligible effect on $N_{\text{eff}}^{\text{SM}}$. The impact of varying the W -boson mass m_W must be similarly imperceptible, since in our parameterisation m_W enters the picture only through flavour oscillations, which are themselves known to be a subdominant, at most $|\delta N_{\text{eff}}| \sim 0.001$ effect. For $\sin^2 \theta_W$, we have explicitly tested that substituting $\sin^2 \theta_W(m_Z)_{\overline{\text{MS}}}$ for $\sin^2 \theta_W(0)_{\overline{\text{MS}}}$ (see caption of table 1 for definitions) shifts $N_{\text{eff}}^{\text{SM}}$ by a minute $|\delta N_{\text{eff}}| \sim 4 \times 10^{-5}$.

Physical parameters of the neutrino sector are, on the other hand, generally far less well measured. The most poorly-determined parameter from the global fit of [33] for example, the solar neutrino mixing angle $\sin^2 \theta_{12}$, has a 1σ -uncertainty of about 5%, as shown in table 1. This is not to mention the as-yet-undetermined sign of Δm_{31}^2 , i.e., a normal or inverted ordering of the neutrino masses, although global fits tend to favour $\Delta m_{31}^2 > 0$ (normal ordering) marginally [32–35, 53, 54].

In this work, we compute our SM benchmark $N_{\text{eff}}^{\text{SM}}$ using the central measured values of the physical parameters given in table 1, assuming a normal ordering of the neutrino masses

($m_3 > m_2 > m_1$). To explore the dependence of $N_{\text{eff}}^{\text{SM}}$ on the parameters of the neutrino sector, we vary the neutrino mass splittings and vacuum mixing angles one at a time by well beyond 5σ away on both sides from their central values. Such a large range of variations more than covers the small differences in the numbers obtained from the three independent global fits of [33–35]. For completeness, we repeat the exercise for an inverted mass ordering ($m_2 > m_1 > m_3$), assuming the same parameter values and uncertainties given in table 1.⁸

4 Numerical convergence

We compute $N_{\text{eff}}^{\text{SM}}$ using the neutrino decoupling code `FortEPiaNO`, which employs the `DLSODA` routine from the `ODEPACK` Fortran package [55, 56] to propagate the Boltzmann–continuity system of equations. Before presenting our results in section 5, we perform first in this section a series of numerical convergence tests to ensure that the settings in `FortEPiaNO` satisfy detailed balance at a level adequate for the computation of $N_{\text{eff}}^{\text{SM}}$ to the desired precision.

To begin with, we note that for an absolute and relative tolerance in `DLSODA` fixed at 10^{-7} , the final outcome N_{eff} may differ by $|\delta N_{\text{eff}}| \sim 10^{-5}$ depending on our choices of compiler (`Intel Fortran` or `GFortran`) and even computer on which to run the code. This latter number therefore forms our estimate of the intrinsic numerical noise of the $N_{\text{eff}}^{\text{SM}}$ calculation, and sets a baseline for the assessment of numerical convergence.

4.1 Convergence variables

Two more settings at our disposal in the execution of `FortEPiaNO` may impact on the numerical convergence of the final SM benchmark $N_{\text{eff}}^{\text{SM}}$: (i) our choice of momentum binning, both in terms of the binning method and the number of bins used, and (ii) the initialisation time x_{in} . We elaborate on our choices in these regards below.

4.1.1 Momentum discretisation

To numerically solve the Boltzmann equation (2.3), `FortEPiaNO` discretises the neutrino momentum distribution on a y -grid. Several discretisation schemes and quadrature methods exist for the purpose. The original version of `FortEPiaNO` employed in reference [30] to study light sterile neutrino thermalisation uses the method of Gauss–Laguerre (GL) quadrature, which proves to be very efficient at sampling spectral distortions at low momenta, reducing significantly discretisation-related numerical errors at a smaller computational cost.

In this work, in addition to GL quadrature, we explore also the method of Newton–Cotes quadrature. We briefly describe these two quadrature methods below and comment on the convergence properties of N_{eff} under these different discretisation schemes.

Gauss–Laguerre (GL) quadrature The method of generalised GL quadrature estimates an integral of the type $\int_0^\infty dy y^n e^{-y} f(y)$ using N_y weighted nodes on the y -grid corresponding to the first N_y roots, y_i , of the Laguerre polynomial of some order N . Observe that we do not use all N roots, both because of computational ease and because we expect the higher y -nodes’ contributions to N_{eff} to be exponentially suppressed. Rather, given a user-specified y_{max} , `FortEPiaNO` automatically selects only those nodes whose roots satisfy $0 < y_i < y_{\text{max}}$, with y_{N_y} sitting just below y_{max} and y_{N_y+1} sitting just above.

⁸The global best-fit values vary slightly between our choice of neutrino mass ordering [33–35], although the variations are statistically insignificant ($\lesssim 1\sigma$).

N_y grid	$\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$	N_y	$N_{\text{eff}}^{\text{SM}}$ (no osc)	$N_{\text{eff}}^{\text{SM}}$ (NO)
GL	$\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$	40	3.0426	3.0436
		60	3.0426	3.0435
		80	3.0425	3.0435
	Diagonal ϱ	40	3.0434	3.0442
		60	3.0433	3.0441
		80	3.0433	3.0441
	Full	40	3.0434	3.0439
		60	3.0433	3.0439
		80	3.0433	3.0439
NC	$\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$	40	3.0428	3.0438
		60	3.0426	3.0436
		80	3.0426	3.0436
	Diagonal ϱ	40	3.0436	3.0444
		60	3.0434	3.0442
		80	3.0433	3.0442
	Full	40	3.0436	3.0441
		60	3.0434	3.0439
		80	3.0433	3.0439

Table 2: Convergence of $N_{\text{eff}}^{\text{SM}}$ under variations of the momentum discretisation scheme (Gauss–Laguerre vs Newton–Cotes) and the number of y -nodes N_y , for three different implementations of the diagonal entries of the neutrino–neutrino collision integral, $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$. See section 3.3.1 for an explanation of the implementations. We report results with no neutrino oscillations (no osc) and with oscillations assuming a normal mass ordering (NO) and the central neutrino mixing parameter values of table 1.

This particular method of y_i selection also means that the actual value of the highest y always fluctuates below the nominal y_{max} , at a distance depending on both our choices of y_{max} and N_y . Using the default setting of $y_{\text{max}} = 20$, table 2 shows the outcome N_{eff} under different set-ups *vis à vis* the neutrino–neutrino collision integral with respect to variations of N_y in the range $N_y \in [40, 80]$. The bottom right panel of figure 3 shows in addition variations with respect to our choice of y_{max} in the range $y_{\text{max}} \in [20, 40]$ for the full calculation. Clearly, provided we do not decrease significantly the ratio y_{max}/N_y , variations in N_y and y_{max} within even fairly large ranges do not generate differences larger than $|\delta N_{\text{eff}}| \sim 10^{-4}$, indicating that numerical convergence can be achieved at the $|\delta N_{\text{eff}}| \sim 10^{-4}$ level or better using GL quadrature.

Newton–Cotes (NC) quadrature The NC quadrature method can be applied to any momentum grid, and we use for our benchmark $N_{\text{eff}}^{\text{SM}}$ calculation 60 linearly-spaced nodes between $y_{\text{min}} = 0.01$ and $y_{\text{max}} = 20$ under the minimum set-up, and 80 linearly-spaced nodes between $y_{\text{min}} = 0.01$ and $y_{\text{max}} = 30$ in a full calculation including the complete neutrino–neutrino collision. As with the GL method, table 2 and the top and bottom left panels of figure 3 demonstrate that NC also performs at $|\delta N_{\text{eff}}| \sim \mathcal{O}(10^{-4})$ level or better with respect to variations in N_y and y_{max} , again provided we do not decrease y_{max}/N_y substantially. In terms of variations in y_{min} , figure 3 shows our N_{eff} estimates to be remarkably stable, shifting

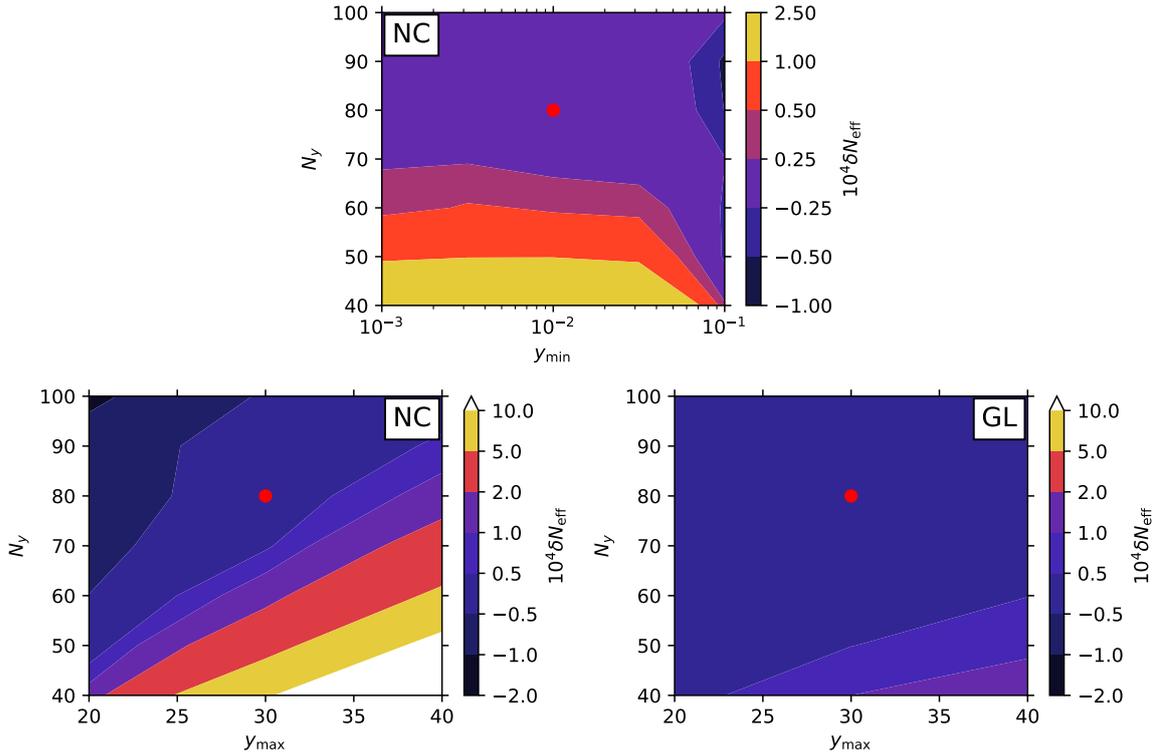


Figure 3: Change in $N_{\text{eff}}^{\text{SM}}$ in a full calculation including the complete neutrino–neutrino collision integral, with respect to variations in the number of y -nodes and their range. *Top:* Variations in N_y and y_{min} using the Newton–Cotes quadrature method. *Bottom left:* Variations in N_y and y_{max} using the Newton–Cotes quadrature method. *Bottom right:* Variations in N_y and y_{max} using the Gauss–Laguerre quadrature method. The setting used in our benchmark full calculation including the complete neutrino–neutrino collision integral ($y_{\text{min}} = 0.01$ where applicable, $y_{\text{max}} = 30$, and $N_y = 80$) is marked by a red dot on each plot.

no more than $|\delta N_{\text{eff}}| \sim 2 \times 10^{-5}$ for a wide range of y_{min} values.⁹

Note that figure 3 is also useful for the purpose of finding the most economical momentum-grid configuration in order to achieve a required precision. The execution time of `FortEPiANO` does not scale linearly with N_y : while the number of y -dependent differential equations that need to be solved does grow with N_y , the 2D integrals for the collision terms are computed on a $N_y \times N_y$ grid. Thus, being able to reduce the required N_y to achieve a target precision can translate into a significant reduction in computation time.

4.1.2 Initialisation time

The neutrino density matrix $\varrho(x, y)$ is initialised at x_{in} with vanishing off-diagonal entries and diagonal entries equal to the ideal gas equilibrium occupation numbers at the rescaled temperature z_{in} , where z_{in} is found by evolving the continuity equation (2.2) from an even earlier time, $x_0 \equiv (1/10)(m_e/m_\mu)$, with $z_0 \equiv a_0 \times 10 m_\mu = 1$, assuming equilibrium in the

⁹Note that these conclusions do not apply to y_{min} above 0.1 or y_{max} below 20, because the momentum range $0.1 \lesssim y \lesssim 20$ is the minimum that we need to sample of the neutrino momentum distribution, in order to capture 99.99% of the (equilibrium) neutrino number density and 99.9995% of the energy density.

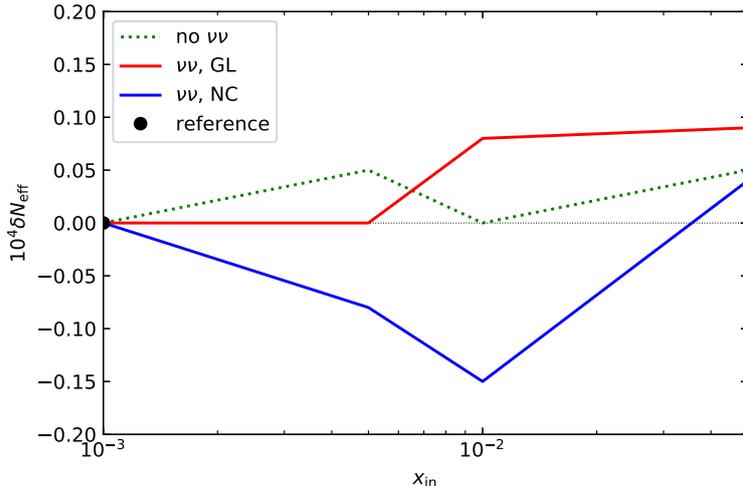


Figure 4: Change in $N_{\text{eff}}^{\text{SM}}$ with respect to variations in the initialisation time x_{in} . The default choice is $x_{\text{in}} = 0.001$. The solid red and blue lines denote the full calculation including the full neutrino–neutrino collision term using respectively the GL and the NC methods, while the green dotted line represents the minimum set-up computed with NC.

neutrino sector. Our default choice of $x_{\text{in}} = 0.001$ corresponds to $z_{\text{in}} - 1 = 2.9 \times 10^{-4}$ or $T_{\text{in}} \simeq 511$ MeV, a temperature at which the neutrino collision rate significantly dominates over both the flavour oscillation and the Hubble expansion rate (see, e.g., figure I of [57]), so that the settings for $\varrho(x_{\text{in}}, y)$ hold to an excellent approximation.

To establish convergence with respect to this initialisation procedure, figure 4 shows the changes in $N_{\text{eff}}^{\text{SM}}$ within the minimum set-up (dotted green) computed with NC and the full calculation including full neutrino–neutrino collisions using the GL (solid red) and the NC (solid blue) methods, as we systematically lower x_{in} over a range $x_{\text{in}} \in [0.001, 0.05]$. The latter number corresponds to $z_{\text{in}} = 1.098$ ¹⁰ or, equivalently, $T_{\text{in}} \simeq 10$ MeV, technically still considerably above the nominal neutrino decoupling temperature, $T_d \sim \mathcal{O}(1)$ MeV. Evidently, for the entire range of x_{in} tested, deviations from the default calculation are always below $|\delta N_{\text{eff}}| \sim 2 \times 10^{-5}$ and hence consistent with numerical noise. We therefore conclude that in the context of computing $N_{\text{eff}}^{\text{SM}}$, the choice of initialisation time x_{in} even as late as $x_{\text{in}} = 0.05$ is not a limiting factor. This reaffirms the findings of [20, 30], which recommended $x_{\text{in}} = 0.05$ for three-flavour decoupling calculations.

4.2 Energy density and number density conservation tests

As a final test of the convergence properties of FortEPiANO, we devise several controlled tests to quantify the extent to which FortEPiANO conserves neutrino number and energy under variations in the momentum discretisation described in section 4.1.1. Specifically, we solve the Boltzmann–continuity system under the following conditions for the collision integrals:

1. **No collisions**, i.e., $\mathcal{I}_{\nu e} = \mathcal{I}_{\nu\nu} = 0$.

¹⁰This number comes from our choice of normalisation of z and the fact that we take into account the presence of muons at high temperatures. When muons become non-relativistic, they transfer their entropy to the rest of the plasma, with the result that the photon temperature grows by a factor $(57/43)^{1/3} \simeq 1.098$.

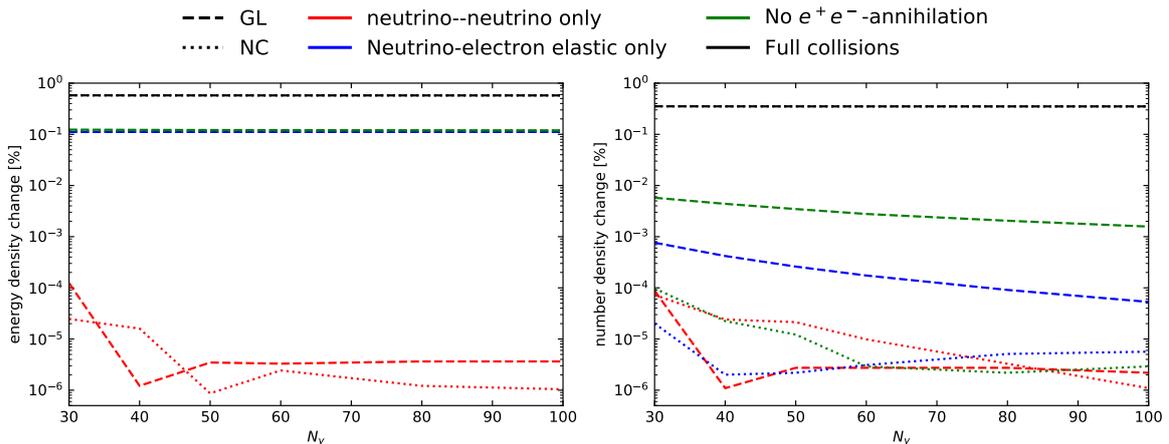


Figure 5: *Left:* Fractional change in percent in the comoving energy density relative to the initial state at $x_{\text{in}} = 0.001$, under variations of the quadrature method (Gauss–Laguerre vs Newton–Cotes), the number of y -nodes N_y , and modelling of the collision integral $\mathcal{I}[\rho]$. *Right:* Fractional change in percent in the comoving number density.

2. **Neutrino–neutrino collisions only**, i.e., we set $\mathcal{I}_{\nu e} = 0$, leaving only $\mathcal{I}_{\nu\nu}$ operative.
3. **Neutrino–electron elastic scattering only**, i.e., we set $\mathcal{I}_{\nu\nu} = 0$, as well as $\mathcal{I}_{\nu e}^{\text{ann}} = 0$ in equation (A.1).
4. **No e^+e^- -annihilation**, i.e., we set $\mathcal{I}_{\nu e}^{\text{ann}} = 0$ in equation (A.1), while keeping all other collision terms operative.
5. **Full collisions**, i.e., we take into account all the contributions to $\mathcal{I}_{\nu e}$ and $\mathcal{I}_{\nu\nu}$.

Theoretically, we expect the comoving neutrino number density to be conserved in only the first four cases, while only the first two scenarios conserve the comoving energy density as well. Numerical detailed balance can be deemed well satisfied if the fractional number and, where appropriate, energy density losses in cases 1-4, can be kept substantially below the target physical changes in the system represented by case 5.

Figure 5 shows the fractional non-conservation in the comoving neutrino number and energy densities, defined respectively as the relative changes in the final-state comoving number and energy densities relative to the initial state. Evidently, where both number and energy conservation are expected (i.e., no collisions, and neutrino–neutrino collisions only), FortEPiaNO is able to replicate the outcome at $10^{-5}\%$ or better using both the Gauss–Laguerre and the Newton–Cotes quadrature schemes. Number conservation also holds at the same level for both methods when the energy-changing neutrino–electron collisions are switched on.

Interestingly, divergences between the Gauss–Laguerre and the Newton–Cotes scheme begin to appear when neutrino–electron interactions are turned on. Here, we see in figure 5 that while NC continues to incur non-conservation in the number density at the $< 10^{-5}\%$ level, the amount of non-conservation is three orders of magnitude as large under the GL scheme and does not converge under variation of N_y . Nonetheless, at $\sim 10^{-3}\%$ the degree of non-conservation of number density from numerical artefacts is still substantially below our accuracy goal. We therefore do not deem the numerical breakdown of detailed balance to be

	$N_{\text{eff}}^{\text{SM}}$ (no osc)	$N_{\text{eff}}^{\text{SM}}$ (NO)	$N_{\text{eff}}^{\text{SM}}$ (IO)
Benchmark A — $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$			
Assuming: <ul style="list-style-type: none"> • $(2)\overline{\ln} + (2)\ln + (3)+$ type (a) weak rates • Damping for $\{\mathcal{I}_{\nu e}[\varrho]\}_{\alpha\beta}$ • $N_y = 60$, $y_{\text{max}} = 20$, NC linearly spaced y_i 	3.04263	3.04360	3.04361
Alternative estimates			
Momentum grid			
$N_y = 40$, $y_{\text{max}} = 20$, GL spacing of y_i nodes	3.04261	3.04355	3.04360
Integrals for off-diagonal $\{\mathcal{I}_{\nu e}[\varrho]\}_{\alpha\beta}$			
$N_y = 60$, $y_{\text{max}} = 20$, NC linearly spaced y_i	3.04261	3.04357	3.04362
$N_y = 40$, $y_{\text{max}} = 20$, GL spacing of y_i	3.04261	3.04357	3.04364
Finite-temperature QED corrections			
$(2)\overline{\ln}$	3.04361	3.04458	
$(2)\overline{\ln} + (2)\ln$	3.04358	3.04452	
$(2)\overline{\ln} + (3)$	3.04264	3.04361	
$(2)\overline{\ln} + (2)\ln + (3)$	3.04263	3.04360	

Table 3: Summary of our $N_{\text{eff}}^{\text{SM}}$ estimates within the minimum set-up (i.e., no diagonal neutrino–neutrino collision integral), computed for different settings of neutrino flavour oscillations and finite-temperature QED corrections. The labels $(2)\overline{\ln}$, $(2)\ln$, and (3) denote respectively the $\mathcal{O}(e^2)$ logarithmic, $\mathcal{O}(e^2)$ log-independent, and $\mathcal{O}(e^3)$ finite-temperature corrections to the QED equation of state.

of concern in our calculation, and estimate the numerical error incurred in the final $N_{\text{eff}}^{\text{SM}}$ to be $|\delta N_{\text{eff}}| \sim 10^{-4}$, based on the convergence tests reported in table 2 and figure 3.

5 Physics results

Having established that it is possible to achieve $|\delta N_{\text{eff}}| \sim 10^{-4}$ numerical convergence, we are finally in a position to report our physics results. First and foremost, the new SM benchmark effective number of neutrinos has been determined to be

$$N_{\text{eff}}^{\text{SM}} = 3.0440 \pm 0.0002, \quad (5.1)$$

under the conditions of (i) a normal neutrino mass ordering, (ii) finite-temperature QED corrections to the QED equation of state to $\mathcal{O}(e^3)$ and to the weak rates to $\mathcal{O}(e^2)$ of the type (a) (see section 3.1.2 and figure 2 for the classification), and (iii) a full numerical treatment of the complete neutrino–electron and neutrino–neutrino collision integrals. The quoted uncertainty is dominated by errors incurred in the numerical solution procedure ($|\delta N_{\text{eff}}| \sim 10^{-4}$), followed by measurement uncertainties in the solar mixing angle $\sin^2 \theta_{12}$ within its experimentally allowed 3σ -range ($|\delta N_{\text{eff}}| \sim 10^{-4}$).

Table 3 summarises the contributions of various elements to $N_{\text{eff}}^{\text{SM}}$ computed within the minimum set-up, while table 4 reports those calculations that incorporate diagonal elements of the neutrino–neutrino collision integral, $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$, implemented in different ways. We shall discuss these results in more detail below. Note however in table 3 that switching the neutrino

	$N_{\text{eff}}^{\text{SM}}$ (no osc)	$N_{\text{eff}}^{\text{SM}}$ (NO)	$N_{\text{eff}}^{\text{SM}}$ (IO)
Benchmark B — $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} \neq 0$			
Assuming: <ul style="list-style-type: none"> • (2)ln + (2)ln + (3)+ type (a) weak rates • Full $\mathcal{I}_{\nu e}[\varrho]$ and $\mathcal{I}_{\nu\nu}[\varrho]$ • $N_y = 80$, $y_{\text{max}} = 30$, NC linearly spaced y_i 	3.04341	3.04398	3.04399
Alternative estimates			
Momentum grid			
$N_y = 80$, $y_{\text{max}} = 30$, GL spacing of y_i	3.04334	3.04392	3.04392
$N_y = 80$, $y_{\text{max}} = 20$, NC linearly spaced y_i	3.04334	3.04389	3.04391
$N_y = 80$, $y_{\text{max}} = 20$, GL spacing of y_i	3.04334	3.04386	3.04393
Off-diagonal collision terms			
Damping terms, NC quadrature		3.04408	
Damping terms, GL quadrature		3.04399	
Neutrino–neutrino collision integral - $y_{\text{max}} = 20$			
Diagonal ϱ	3.04333	3.04416	
Full ϱ , interpolate ϱ /FD only in diagonal	3.04334	3.04389	
Full ϱ , interpolate ϱ /FD also in off-diagonal	3.04334	3.04389	

Table 4: Summary of our $N_{\text{eff}}^{\text{SM}}$ estimates, obtained for various implementations of the diagonal neutrino–neutrino collision integral. We report results with no neutrino oscillations (no osc), with oscillations assuming a normal mass ordering (NO) or an inverted mass ordering (IO) and the central neutrino mixing parameter values of table 1.

mass ordering from normal to inverted generates but a negligible $|\delta N_{\text{eff}}| \sim 10^{-5}$ difference in $N_{\text{eff}}^{\text{SM}}$. Henceforth, we shall concern ourselves mainly with a normal mass ordering.

5.1 New finite-temperature QED corrections

Relative to the 2016 calculations of [20], table 3 shows that the main modification to $N_{\text{eff}}^{\text{SM}}$ originates in the $\mathcal{O}(e^3)$ finite-temperature correction to the QED equation of state. This correction generates a change of $\delta N_{\text{eff}} \sim -0.001$ under the Benchmark A settings, bringing the 2016 value of 3.045 [20] down to 3.0440, consistent with the estimate of [19] within the instantaneous decoupling approximation.

We further test the impact of various combinations of finite-temperature QED corrections on $N_{\text{eff}}^{\text{SM}}$, also shown in table 3. These numbers are again largely consistent with the instantaneous-decoupling estimates of [19], and reaffirm the conclusions of [19] that the $\mathcal{O}(e^2)$ logarithmic correction to the QED equation of state ($\delta N_{\text{eff}} \sim -4 \times 10^{-5}$) and the weak rate corrections of type (a) ($|\delta N_{\text{eff}}| \lesssim 10^{-4}$) are strictly not necessary to achieve a prediction of $N_{\text{eff}}^{\text{SM}}$ to four-digit significance.

5.2 Full collision integral for neutrino–neutrino scattering

Evaluating the full neutrino–neutrino collision integral in the computation of $N_{\text{eff}}^{\text{SM}}$, especially its diagonal entries $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$, raises the number by $\delta N_{\text{eff}} \simeq 4 \times 10^{-4}$ to $N_{\text{eff}}^{\text{SM}} = 3.0440$. This is the Benchmark B value reported in table 4 including neutrino oscillations, and also our final recommended SM benchmark presented in equation (5.1). The choice between using the

full neutrino density matrix $\varrho(y)$ (“full”) or only its diagonal elements (“diagonal ϱ ”) in the evaluation of the diagonal entries $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ impacts on the final outcome at a comparable, $|\delta N_{\text{eff}}| \simeq 2 \times 10^{-4}$ level.¹¹ We therefore recommend that a full evaluation of $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ be employed whenever neutrino oscillations are switched on.

Another point of note in table 4 is that including nonzero $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ terms in the calculation diminishes the net contribution of neutrino oscillations to the final outcome $N_{\text{eff}}^{\text{SM}}$: Within the minimum set-up (i.e., $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$), oscillations add $\delta N_{\text{eff}} \sim 0.001$ to the no-oscillation estimate. With $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} \neq 0$ this contribution consistently reduces to $\delta N_{\text{eff}} \sim 5 \times 10^{-4}$. This is a physically sensible result: allowing neutrinos to scatter amongst themselves reduces the reliance of the system on oscillations as a means of energy transport between different flavours. From the perspective of computing $N_{\text{eff}}^{\text{SM}}$ to four-digit significance, this result indicates that while flavour oscillations are not as crucial an ingredient as once envisaged, they are nonetheless of sufficient importance to include in a calculation to avoid rounding errors.

5.3 Assessment of remaining uncertainties

We have seen at the end of section 4.2 that the uncertainty incurred in $N_{\text{eff}}^{\text{SM}}$ from our numerical solution procedure is estimated to be $|\delta N_{\text{eff}}| \sim 10^{-4}$. The next finite-temperature correction to the QED equation of state is expected to enter at $\delta N_{\text{eff}} \sim 4 \times 10^{-6}$ [19]. Finite-temperature QED corrections to the weak rates of the types (b), (c) and (d), on the other hand, have yet to be determined, although one might reasonably guess that their contributions to $N_{\text{eff}}^{\text{SM}}$ are similar in magnitude to that of the type (a), which sits at just below $\delta N_{\text{eff}} \sim 10^{-4}$ according to table 3. Other sources of uncertainties investigated and their impact on $N_{\text{eff}}^{\text{SM}}$ are presented in tables 3 and 4 as well as figure 6, on which we elaborate below.

5.3.1 Approximate treatment of the weak collision integrals

Approximate treatments of the diagonal neutrino–neutrino collision integral and their impact on $N_{\text{eff}}^{\text{SM}}$ as detailed in table 4 have already been discussed in section 5.2. To paraphrase the conclusion, setting $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$ generally underestimates $N_{\text{eff}}^{\text{SM}}$ by about $\delta N_{\text{eff}} \simeq 4 \times 10^{-4}$, while approximating the full density matrix ϱ with only its diagonal entries in the evaluation of $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha}$ (the “diagonal ϱ ” approximation) generates a comparable difference of $|\delta N_{\text{eff}}| \simeq 2 \times 10^{-4}$ and should hence be avoided when neutrino oscillations are switched on.

Of lesser impact is the exact modelling of the *off-diagonal* entries of the neutrino–electron and neutrino–neutrino collision integrals, $\{\mathcal{I}_{\nu e}[\varrho]\}_{\alpha\beta}$ and $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\beta}$, when neutrino oscillations are operative (these terms are irrelevant in the no-oscillation case). While this is one of the first works to use a full evaluation of $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\beta}$ in the computation of $N_{\text{eff}}^{\text{SM}}$, we find that substituting $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\beta}$ and $\{\mathcal{I}_{\nu e}[\varrho]\}_{\alpha\beta}$ with the damping approximation (3.1) generates a difference of $|\delta N_{\text{eff}}| \lesssim 10^{-4}$. We therefore conclude that the off-diagonal damping approximation may be a reasonable compromise if computing resources need to be diverted elsewhere.

5.3.2 Measurement errors in the physical parameters of the neutrino sector

Figure 6 shows changes in the predicted $N_{\text{eff}}^{\text{SM}}$ with respect to variations in the neutrino mass splittings and mixing parameters used in the calculation, for both a normal and an inverted

¹¹The diagonal ϱ approximation combined with no flavour oscillations is equivalent to the scenario considered in references [7, 8, 11]. The N_{eff} values reported in these works are consistently larger than our no oscillations Benchmark B of 3.04341 by about 0.001, because they did not include the $\mathcal{O}(e^3)$ finite-temperature correction to the QED equation of state in their computation.

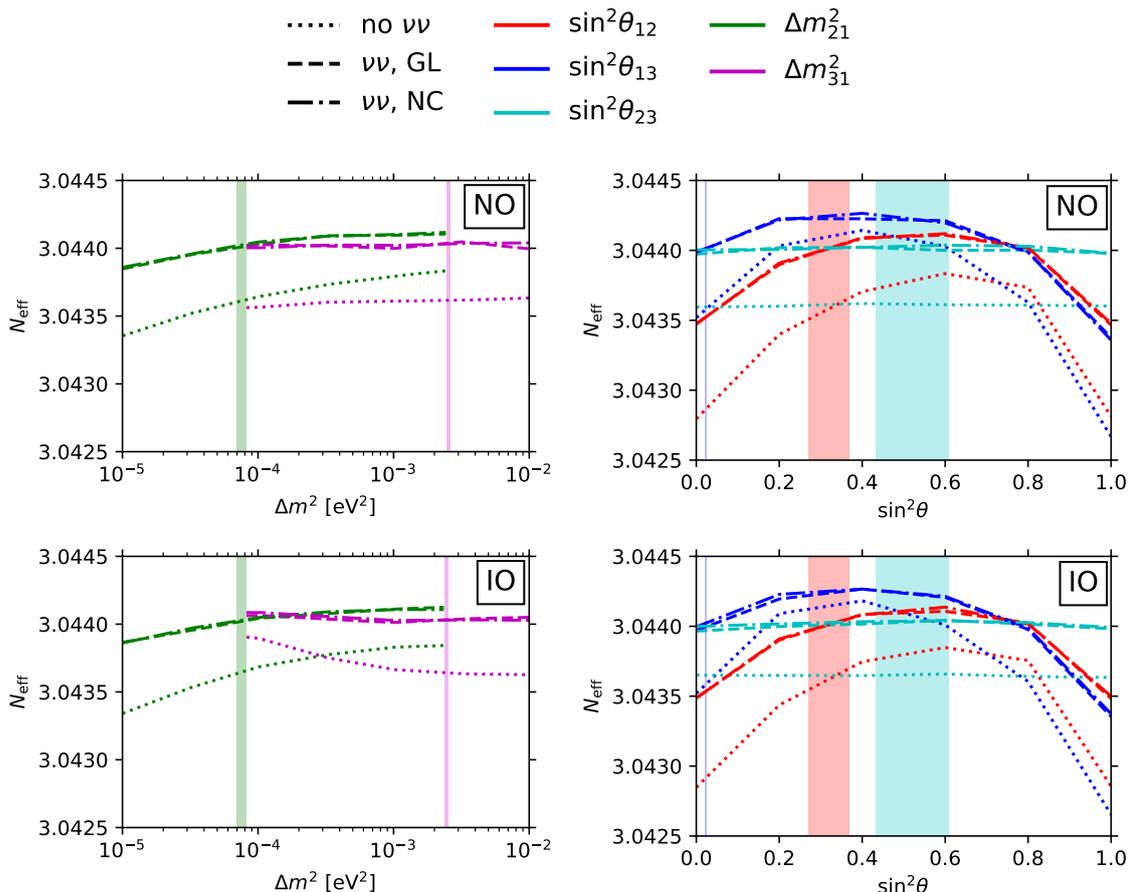


Figure 6: Change in $N_{\text{eff}}^{\text{SM}}$ under variations in the physical parameters of the neutrino sector, assuming a normal mass ordering (top row) and an inverted mass ordering (bottom row). We consider variations in the mass splittings Δm_{21}^2 and Δm_{31}^2 , assuming at all times $\Delta m_{21}^2 < |\Delta m_{31}^2|$ (left column), and in the mixing angles $\sin^2 \theta_{12}$, $\sin^2 \theta_{13}$, and $\sin^2 \theta_{23}$ (right column). We show results for both the minimum set-up (dotted lines) and the full calculation (dot-dash lines). Coloured bands represent the current 3σ -range of each parameter given in table 1 [33].

neutrino mass ordering. We consider these variations within the minimum set-up where $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} = 0$, as well as in the context of a full, $\{\mathcal{I}_{\nu\nu}[\varrho]\}_{\alpha\alpha} \neq 0$ calculation.

First of all, we note that the minimum set-up is, in comparison with the full calculation, visibly more sensitive to variations in the neutrino parameters. Ramping up $\sin^2 \theta_{12}$ from zero to its best-fit value of 0.32, for example, generates $\delta N_{\text{eff}} \simeq 0.001$ in the minimal set-up and $\delta N_{\text{eff}} \simeq 0.0005$ in the full calculation. This result is consistent with our earlier observation in section 5.2 that including flavour oscillations in the calculation has a stronger impact on the outcome $N_{\text{eff}}^{\text{SM}}$ within the minimum set-up than within the full calculation.

The actual variations in $N_{\text{eff}}^{\text{SM}}$ with respect to changes in the parameter values — especially if we restrict our attention to the 3σ regions (shaded regions) — are however generally quite small. Variations with respect to the neutrino mass splittings, notwithstanding the relatively large, percent-level 1σ -uncertainties in Δm_{21}^2 and Δm_{31}^2 , are practically indiscernible in

the 3σ region. Physically, such insensitivity of $N_{\text{eff}}^{\text{SM}}$ to Δm_{21}^2 and Δm_{31}^2 reflects the fact that the typical oscillation frequency, $\Delta m_{ij}^2/2\langle p \rangle$, in all oscillation channels is much larger than the Hubble expansion rate across the neutrino decoupling epoch (see, e.g., figure I of [57]). This means that the oscillation probability averaged over a Hubble time is effectively dependent only on the mixing angle; the exact mass splittings largely drop out of the picture.

For the neutrino mixing angles, we likewise find no discernible dependence of $N_{\text{eff}}^{\text{SM}}$ on $\sin^2 \theta_{23}$ across the whole parameter range tested ($\sin^2 \theta_{23} \in [0, 1]$), clearly because this oscillation channel merely swaps the ν_μ and the ν_τ populations. Except for a largely inconsequential dependence on the muon energy density in the matter potential part of the flavour oscillations Hamiltonian (2.4),¹² these two populations are, for the $N_{\text{eff}}^{\text{SM}}$ calculation, essentially identical. Changing the rate at which they are swapped therefore has no real impact on the outcome.

The only instances in which we find a reasonably strong, $\delta N_{\text{eff}} \sim \mathcal{O}(10^{-4})$ response of $N_{\text{eff}}^{\text{SM}}$ to parameter variations are the cases of $\sin^2 \theta_{12}$ and $\sin^2 \theta_{13}$, where tuning these parameters towards $\sin^2 \theta_{ij} \simeq 0.5$ tends to increase $N_{\text{eff}}^{\text{SM}}$. These mixing angles control the oscillations between the ν_e and the $\nu_{\mu,\tau}$ populations, the former of which has the important feature that it decouples from the QED plasma last because of the charged-current $\nu_e e$ -interactions unique to it. Then, a larger mixing angle enhances the energy transfer rate between ν_e and $\nu_{\mu,\tau}$, and keeps the latter populations effectively coupled to the QED plasma for a longer time to partake in the entropy transfer from e^+e^- -annihilation.

For a $\pm 3\sigma$ ($\sim 20\%$) variation in $\sin^2 \theta_{12}$ from its central value, figure 6 shows that $N_{\text{eff}}^{\text{SM}}$ can change by as much as $\delta N_{\text{eff}} \sim \pm 10^{-4}$, consistent with expectations (i.e., $\sim 20\%$ of the change due to flavour oscillations $\delta N_{\text{eff}} \sim 5 \times 10^{-4}$). For similar relative variations in $\sin^2 \theta_{13}$, the change in $N_{\text{eff}}^{\text{SM}}$ is an order of magnitude smaller. But this is merely because the central value of $\sin^2 \theta_{13}$ itself is an order of magnitude smaller than $\sin^2 \theta_{12}$ to begin with, and hence plays only a subdominant role to the latter in terms of facilitating energy transfer between the ν_e and the $\nu_{\mu,\tau}$ populations. Had it been possible to increase $\sin^2 \theta_{13}$ to ~ 0.5 , figure 6 shows that $N_{\text{eff}}^{\text{SM}}$ would have been enhanced by $\delta N_{\text{eff}} \sim 2 - 3 \times 10^{-4}$.

6 Conclusions

We have updated in this work the standard-model benchmark value for the effective number of neutrinos, $N_{\text{eff}}^{\text{SM}}$, that quantifies the cosmological neutrino-to-photon energy densities, and estimated its uncertainties. Our recommended value of $N_{\text{eff}}^{\text{SM}} = 3.0440 \pm 0.0002$ has been established through a careful tracking of relic neutrino decoupling in the presence of neutrino flavour oscillations (assuming a normal neutrino mass ordering), as well as finite-temperature effects in the QED plasma. The error estimate takes into account numerical uncertainty at the level of $|\delta N_{\text{eff}}| \sim 10^{-4}$ due to the discretisation of the momentum grid, plus a physical error of $|\delta N_{\text{eff}}| \sim 10^{-4}$ arising from the current 3σ measurement uncertainty in the neutrino oscillation parameters. Our central $N_{\text{eff}}^{\text{SM}}$ value is in perfect agreement with the calculation of reference [13] which incorporated the same physics; our nominal theoretical uncertainty is somewhat larger than that quoted in [13], however, owing to different accounts of the physical error from the oscillation parameters.

Relative to the 2016 calculation [20] which has a nominal uncertainty of ± 0.001 , the benchmark has shifted by $\delta N_{\text{eff}} \simeq -0.001$ because of a hitherto neglected $\mathcal{O}(e^3)$ finite-

¹²The $\nu_\mu \leftrightarrow \nu_\tau$ vacuum oscillation frequency typically supersedes the Hubble expansion rate at $T \sim 20$ MeV, by which time the muon energy density has become strongly Boltzmann-suppressed. For this reason, the muon energy density that appears in the matter potential is largely inconsequential for the computation of N_{eff} .

Standard-model corrections to $N_{\text{eff}}^{\text{SM}}$	Leading-digit contribution
m_e/T_d correction	+0.04
$\mathcal{O}(e^2)$ FTQED correction to the QED EoS	+0.01
Non-instantaneous decoupling+spectral distortion	-0.005
$\mathcal{O}(e^3)$ FTQED correction to the QED EoS	-0.001
Flavour oscillations	+0.0005
Type (a) FTQED corrections to the weak rates	$\lesssim 10^{-4}$

Table 5: Leading-digit contributions from various SM corrections, in order of importance, accounted for in this work that make up the final $N_{\text{eff}}^{\text{SM}} - 3$. The largest, m_e/T_d correction is identically the Rel correction discussed in [19], while the non-instantaneous decoupling+spectral distortion correction is defined relative to an estimate of $N_{\text{eff}}^{\text{SM}}$ in the limit of instantaneous decoupling assuming $T_d = 1.3453$ MeV [19].

temperature correction to the QED equation state. We note that the oft-quoted 2005 number $N_{\text{eff}}^{\text{SM}} = 3.046$ [6] also misses the $\mathcal{O}(e^3)$ correction and has a nominal uncertainty of ± 0.002 [6]. A leading-digit breakdown of the various SM effects that contribute to $N_{\text{eff}}^{\text{SM}}$'s deviation from 3 is presented in table 5.

The uncertainty in the benchmark $N_{\text{eff}}^{\text{SM}}$ we quote in equation (5.1) is a conservative sum of the two broad classes of errors examined in this work: numerical convergence of the solution procedure, notably momentum discretisation, and measurement uncertainties in the physical parameters of the neutrino sector. It is chiefly dominated by the former, and is further augmented by contributions from measurement errors in the solar neutrino mixing angle $\sin^2 \theta_{12}$. Other uncertainties, including higher-order finite-temperature QED corrections, measurement errors in the other oscillation parameters, transients, etc., all fall below the intrinsic numerical noise of FortEPiANO, which, in the context of computing $N_{\text{eff}}^{\text{SM}}$, is in the ballpark of $|\delta N_{\text{eff}}| \sim 10^{-5}$.

Relative to the nominal uncertainty of $\delta N_{\text{eff}} \sim \pm 2 \times 10^{-4}$ given in equation (5.1), we believe we have exhausted all possible effects within the standard model of particle physics that would change $N_{\text{eff}}^{\text{SM}}$ appreciably. An uncertainty of this magnitude also more than suffices to minimise the total error budget in the inference of cosmological parameters from the forthcoming generation of cosmological observations [27]. Nevertheless, even if only for completeness, estimates of the types (b), (c) and (d) finite-temperature QED corrections to the weak rates remain on the table, and there is certainly scope for beating down the uncertainty in $N_{\text{eff}}^{\text{SM}}$ even further. In the latter regard, the measurement of $\sin^2 \theta_{12}$, for example, is expected to be improved by the next generation of neutrino oscillations experiment to a better-than-1% determination [58] (see also [59]). A dedicated investigation of the stability of the weak collision integral evaluation could potentially also eliminate a large and currently dominating chunk in the nominal uncertainty in $N_{\text{eff}}^{\text{SM}}$. We leave these for future work.

Acknowledgments

JJB and Y³W are supported in part by the Australian Government through the Australian Research Council's Discovery Project (project DP170102382) and Future Fellowship (project FT180100031) funding schemes. GB acknowledges the support of the National Fund for Scientific Research (F.R.S.- FNRS Belgium) through a FRIA grant. MaD and Y³W acknowledge support from the ASEM-DUO fellowship programme of the Belgian Académie de recherche

et d'enseignement supérieur (ARES). PFdS acknowledges support by the Vetenskapsrådet (Swedish Research Council) through contract No. 638-2013-8993 and the Oskar Klein Centre for Cosmoparticle Physics. SG and SP are supported by the Spanish grants FPA2017-85216-P (AEI/FEDER, UE), PROMETEO/2018/165 (Generalitat Valenciana) and the Red Consolider MultiDark FPA2017-90566-REDC. SG acknowledges financial support by the ‘‘Juan de la Cierva-Incorporación’’ program (IJC2018-036458-I) of the Spanish MICINN until September 2020, from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 754496 (project FELLINI) starting from October 2020, and thanks the Institute for Nuclear Theory at the University of Washington for its hospitality and the Department of Energy for partial support during the preparation of this work. We thank Julien Froustey and Oleksandr Tomalak for pointing out an error in the original manuscript.

A Collision integrals

Accounting only for $2 \rightarrow 2$ neutrino–electron and neutrino–neutrino collision processes, the weak collision integral $\mathcal{I}[\varrho(x, y)]$ splits naturally into two parts, $\mathcal{I}[\varrho(x, y)] = \mathcal{I}_{\nu e}[\varrho(x, y)] + \mathcal{I}_{\nu\nu}[\varrho(y)]$. At tree level, their general forms can be found in, e.g., [31]. Here, we give their 2D-reduced forms under the assumptions of (i) spatial homogeneity and isotropy, and (ii) CP -symmetry. The integral reduction follows the procedure of [3]. We note however other reduction methods exist [19, 22], which yield formally different but numerically identical results.

Neutrino–electron. The neutrino–electron collision integral $\mathcal{I}_{\nu e}[\varrho(x, y)]$ splits further into a scattering and an annihilation part,

$$\mathcal{I}_{\nu e}[\varrho(x, y)] = \frac{G_F^2}{(2\pi)^3 y^2} \{I_{\nu e}^{\text{sc}}[\varrho(x, y)] + I_{\nu e}^{\text{ann}}[\varrho(x, y)]\}, \quad (\text{A.1})$$

with

$$\begin{aligned} I_{\nu e}^{\text{sc}} &= \int dy_2 dy_3 \frac{y_2}{\epsilon_2} \\ &\times \left\{ (\Pi_{2a}^{\text{sc}}(y, y_2) + \Pi_{2b}^{\text{sc}}(y, y_4)) \left[F_{\text{sc}}^{LL}(\varrho^{(1)}, f_e^{(2)}, \varrho^{(3)}, f_e^{(4)}) + F_{\text{sc}}^{RR}(\varrho^{(1)}, f_e^{(2)}, \varrho^{(3)}, f_e^{(4)}) \right] \right. \\ &\quad \left. - 2x^2 \Pi_1^{\text{sc}}(y, y_3) \left[F_{\text{sc}}^{RL}(\varrho^{(1)}, f_e^{(2)}, \varrho^{(3)}, f_e^{(4)}) + F_{\text{sc}}^{LR}(\varrho^{(1)}, f_e^{(2)}, \varrho^{(3)}, f_e^{(4)}) \right] \right\}, \end{aligned} \quad (\text{A.2})$$

$$\begin{aligned} I_{\nu e}^{\text{ann}} &= \int dy_2 dy_3 \frac{y_3}{\epsilon_3} \\ &\times \left\{ \Pi_{2b}^{\text{ann}}(y, y_4) F_{\text{ann}}^{LL}(\varrho^{(1)}, \varrho^{(2)}, f_e^{(3)}, f_e^{(4)}) + \Pi_{2a}^{\text{ann}}(y, y_3) F_{\text{ann}}^{RR}(\varrho^{(1)}, \varrho^{(2)}, f_e^{(3)}, f_e^{(4)}) \right. \\ &\quad \left. + x^2 \Pi_1^{\text{ann}}(y, y_2) \left[F_{\text{ann}}^{RL}(\varrho^{(1)}, \varrho^{(2)}, f_e^{(3)}, f_e^{(4)}) + F_{\text{ann}}^{LR}(\varrho^{(1)}, \varrho^{(2)}, f_e^{(3)}, f_e^{(4)}) \right] \right\}, \end{aligned} \quad (\text{A.3})$$

and $\epsilon_i^2 = x^2 + y_i^2$ is a rescaled energy. The scattering kernels are given by

$$\begin{aligned} \Pi_1^{\text{sc}}(y, y_3) &= y y_3 D_1 + D_2(y, y_3, y_2, y_4), \\ \Pi_1^{\text{ann}}(y, y_2) &= y y_2 D_1 - D_2(y, y_2, y_3, y_4), \\ \Pi_{2a}^{\text{sc}}(y, y_2)/2 &= y \epsilon_2 y_3 \epsilon_4 D_1 + D_3 - y \epsilon_2 D_2(y_3, y_4, y, y_2) - y_3 \epsilon_4 D_2(y, y_2, y_3, y_4), \\ \Pi_{2b}^{\text{sc}}(y, y_4)/2 &= y \epsilon_2 y_3 \epsilon_4 D_1 + D_3 + \epsilon_2 y_3 D_2(y, y_4, y_2, y_3) + y \epsilon_4 D_2(y_2, y_3, y, y_4), \\ \Pi_{2a}^{\text{ann}}(y, y_3)/2 &= y y_2 \epsilon_3 \epsilon_4 D_1 + D_3 + y \epsilon_3 D_2(y_2, y_4, y, y_3) + y_2 \epsilon_4 D_2(y, y_3, y_2, y_4), \\ \Pi_{2b}^{\text{ann}}(y, y_4)/2 &= y y_2 \epsilon_3 \epsilon_4 D_1 + D_3 + y_2 \epsilon_3 D_2(y, y_4, y_2, y_3) + y \epsilon_4 D_2(y_2, y_3, y, y_4), \end{aligned} \quad (\text{A.4})$$

where the functions D_i are defined as follows [3]:

$$\begin{aligned}
D_1(a, b, c, d) &= \frac{16}{\pi} \int_0^\infty \frac{d\lambda}{\lambda^2} \prod_{i=a,b,c,d} \sin(\lambda i), \\
D_2(a, b, c, d) &= -\frac{16}{\pi} \int_0^\infty \frac{d\lambda}{\lambda^4} \prod_{i=a,b} [\lambda i \cos(\lambda i) - \sin(\lambda i)] \prod_{j=c,d} \sin(\lambda j), \\
D_3(a, b, c, d) &= \frac{16}{\pi} \int_0^\infty \frac{d\lambda}{\lambda^6} \prod_{i=a,b,c,d} [\lambda i \cos(\lambda i) - \sin(\lambda i)].
\end{aligned} \tag{A.5}$$

All three integrals can be evaluated analytically. See, e.g., [60] for the complete expressions.

Lastly, the phase space matrices F_{sc}^{ab} and F_{ann}^{ab} , where $a, b = R, L$, are given by

$$\begin{aligned}
F_{\text{sc}}^{ab} &\equiv f_e^{(4)}(\mathbb{1} - f_e^{(2)})G^a \varrho^{(3)} G^b (\mathbb{1} - \varrho^{(1)}) - f_e^{(2)}(\mathbb{1} - f_e^{(4)})\varrho^{(1)} G^b (\mathbb{1} - \varrho^{(3)})G^a + \text{h.c.} \\
F_{\text{ann}}^{ab} &\equiv f_e^{(3)} f_e^{(4)} G^a (\mathbb{1} - \varrho^{(2)})G^b (\mathbb{1} - \varrho^{(1)}) - (\mathbb{1} - f_e^{(3)})(\mathbb{1} - f_e^{(4)})G^a \varrho^{(2)} G^b \varrho^{(1)} + \text{h.c.},
\end{aligned} \tag{A.6}$$

with $\varrho^{(i)} = \varrho(y_i)$, $f_e^{(i)} = f_e(y_i) \times \mathbb{1}$, and

$$\begin{aligned}
G^L &= \text{diag}(g_L, \tilde{g}_L, \tilde{g}_L), \\
G^R &= \text{diag}(g_R, g_R, g_R)
\end{aligned} \tag{A.7}$$

are the weak coupling matrices for left- and right-handed particles, conventionally taken to be $g_L = \sin^2 \theta_W + 1/2$, $\tilde{g}_L = \sin^2 \theta_W - 1/2$, and $g_R = \sin^2 \theta_W$. In our calculation of $N_{\text{eff}}^{\text{SM}}$, we use the values given in table 1 at zero-momentum transfer [50, 51].

Neutrino–neutrino. The neutrino–neutrino collision integral likewise splits into a scattering and an annihilation part,

$$\mathcal{I}_{\nu\nu}[\varrho(x, y)] = \frac{G_F^2}{(2\pi)^3 y^2} \{ I_{\nu\nu}^{\text{sc}}[\varrho(x, y)] + I_{\nu\nu}^{\text{pair}}[\varrho(x, y)] \}, \tag{A.8}$$

with

$$I_{\nu\nu}^{\text{sc}} = \frac{1}{4} \int dy_2 dy_3 \Pi_{2a}^{\text{sc}}(y, y_2) F_{\text{sc}}^{\nu\nu}(\varrho^{(1)}, \varrho^{(2)}, \varrho^{(3)}, \varrho^{(4)}), \tag{A.9}$$

$$I_{\nu\nu}^{\text{pair}} = \frac{1}{4} \int dy_2 dy_3 \Pi_{2b}^{\text{sc}}(y, y_4) F_{\text{pair}}^{\nu\nu}(\varrho^{(1)}, \varrho^{(2)}, \varrho^{(3)}, \varrho^{(4)}). \tag{A.10}$$

The scattering kernels are identically those given above in equation (A.4), and the phase space matrices are

$$\begin{aligned}
F_{\text{sc}}^{\nu\nu} &\equiv (\mathbb{1} - \varrho^{(1)})\varrho^{(3)} \left[(\mathbb{1} - \varrho^{(2)})\varrho^{(4)} + \text{tr}(\dots) \right] - \varrho^{(1)}(\mathbb{1} - \varrho^{(3)}) \left[\varrho^{(2)}(\mathbb{1} - \varrho^{(4)}) + \text{tr}(\dots) \right] + \text{h.c.}, \\
F_{\text{pair}}^{\nu\nu} &\equiv (\mathbb{1} - \varrho^{(1)})(\mathbb{1} - \varrho^{(2)}) \left[\varrho^{(4)}\varrho^{(3)} + \text{tr}(\dots) \right] - \varrho^{(1)}\varrho^{(2)} \left[(\mathbb{1} - \varrho^{(4)})(\mathbb{1} - \varrho^{(3)}) + \text{tr}(\dots) \right] \\
&\quad + (\mathbb{1} - \varrho^{(1)})\varrho^{(3)} \left[\varrho^{(4)}(\mathbb{1} - \varrho^{(2)}) + \text{tr}(\dots) \right] - \varrho^{(1)}(\mathbb{1} - \varrho^{(3)}) \left[(\mathbb{1} - \varrho^{(4)})\varrho^{(2)} + \text{tr}(\dots) \right] + \text{h.c.},
\end{aligned} \tag{A.11}$$

where the notation $\text{tr}(\dots)$ denotes the trace of the preceding term.

B Damping approximation

Observe that all collision integrals (A.2), (A.3), (A.9), and (A.10) come in the form

$$\mathcal{I}[\varrho(y)] = (\mathbb{1} - \varrho(y))\Gamma_y^< - \varrho(y)\Gamma_y^> + \text{h.c.}, \quad (\text{B.1})$$

where $\Gamma_y^<$ and $\Gamma_y^>$ are the gain and loss terms respectively, and it is understood that they are 3×3 complex matrices. Consider first an off-diagonal element, $\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta}$, with $\alpha \neq \beta$. Defining $\Gamma_y \equiv \Gamma_y^< + \Gamma_y^>$ and writing out equation (B.1) explicitly in index notation, we find

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta} = -\{\varrho(y)\}_{\alpha\beta} \left[\{\Gamma_y\}_{\beta\beta} + \{\Gamma_y^*\}_{\alpha\alpha} \right] - \sum_{\gamma \neq \beta} \{\varrho(y)\}_{\alpha\gamma} \{\Gamma_y\}_{\gamma\beta} - \sum_{\gamma \neq \alpha} \{\varrho(y)\}_{\gamma\beta} \{\Gamma_y^*\}_{\gamma\alpha}, \quad (\text{B.2})$$

where we have used the fact that $\{\varrho(y)\}_{\alpha\beta} = \{\varrho(y)^*\}_{\beta\alpha}$, and summation over γ is implied. Observe that the two sums contain only off-diagonal entries of Γ_y .

Suppose now Γ_y is diagonal. This might be the case if, for example, the density matrices $\varrho(y_2), \varrho(y_3), \varrho(y_4)$ that constitute the integrand of Γ_y are all diagonal, or if their off-diagonal entries all oscillate with different phases so that $\{\Gamma_y\}_{\alpha\beta}$ integrates to zero. Then, equation (B.2) immediately simplifies to

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta} \simeq -\{D(y)\}_{\alpha\beta} \{\varrho(y)\}_{\alpha\beta}, \quad (\text{B.3})$$

which is of a damping form, with damping coefficient $\{D(y)\}_{\alpha\beta} \equiv \{\Gamma_y\}_{\beta\beta} + \{\Gamma_y^*\}_{\alpha\alpha}$. This is the origin of the ‘‘off-diagonal damping approximation’’ (3.1).

The exercise can be repeated also for a diagonal entry of $\mathcal{I}[\varrho(y)]$. Writing out equation (B.1) in index form, we find

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\alpha} = 2\{\Gamma_y^<\}_{\alpha\alpha} - 2\{\varrho(y)\}_{\alpha\alpha} \{\Gamma_y\}_{\alpha\alpha} - \sum_{\gamma \neq \alpha} \left[\{\varrho(y)\}_{\alpha\gamma} \{\Gamma_y\}_{\gamma\alpha} + \{\varrho(y)^*\}_{\alpha\gamma} \{\Gamma_y^*\}_{\gamma\alpha} \right], \quad (\text{B.4})$$

where we have used the fact that $\{\varrho(y)\}_{\alpha\alpha}$ and $\{\Gamma_y\}_{\alpha\alpha}$ are real. Assuming again that all off-diagonal entries of Γ_y are negligible, equation (B.4) simplifies to

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\alpha} \simeq 2\{\Gamma_y^<\}_{\alpha\alpha} - 2\{\varrho(y)\}_{\alpha\alpha} \{\Gamma_y\}_{\alpha\alpha}. \quad (\text{B.5})$$

Lastly, if all particle species — besides the one at the mode y represented by $\varrho(y)$ — are in a state of thermal equilibrium, then detailed balance requires that $\Gamma_y^> = e^{\epsilon_y/z} \Gamma_y^<$, with a rescaled energy ϵ_y associated with the mode y , and a rescaled temperature z . Consequently, equation (B.5) can be recast in the form

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\alpha} \simeq -\{R(y)\}_{\alpha} \left[\{\varrho(y)\}_{\alpha\alpha} - f_{\text{eq}}(y) \right], \quad (\text{B.6})$$

where $\{R(y)\}_{\alpha} \equiv 2\{\Gamma_y\}_{\alpha\alpha}$ is often called a ‘‘repopulation coefficient’’ [2, 46, 47], and for $\epsilon_y = y$ (valid for nearly massless neutrinos), $f_{\text{eq}}(y)$ is simply the relativistic Fermi–Dirac distribution. Equation (B.6), then, is the ‘‘diagonal damping approximation’’.

Comparing the definitions of the diagonal and off-diagonal damping coefficients, $\{R(y)\}_{\alpha}$ and $\{D(y)\}_{\alpha\beta}$, we identify the relation

$$\{D(y)\}_{\alpha\beta} = \frac{1}{2} \left[\{R(y)\}_{\alpha} + \{R(y)\}_{\beta} \right], \quad (\text{B.7})$$

which is well known especially in the context of α and β representing respectively an active and a sterile neutrino flavour. Then, combining equations (B.3) and (B.5), we can now write down a “general damping approximation”,

$$\{\mathcal{I}[\varrho(y)]\}_{\alpha\beta} = -\{D(y)\}_{\alpha\beta} \left[\{\varrho(y)\}_{\alpha\beta} - \delta_{\alpha\beta} f_{\text{eq}}(y) \right], \quad (\text{B.8})$$

for both diagonal $\alpha = \beta$ and off-diagonal $\alpha \neq \beta$ entries of a collision integral.

Before proceeding to derive the repopulation and damping coefficients, note that while for completeness we have given the general damping approximation (B.8), in practical implementations the diagonal version typically does not have sufficient accuracy relative to our goals, owing to the difficulty in knowing the correct form of $f_{\text{eq}}(y)$ to which the neutrino ensemble should tend. Therefore, the minimum set-up of our calculation always solves the diagonal entries of the neutrino–electron collision integral $\mathcal{I}_{\nu e}[\varrho(y)]$ in full. Note however that the approximate damping scheme proposed in reference [22] may circumvent this problem.

B.1 Damping coefficients

Neutrino–neutrino. To evaluate the repopulation and hence damping coefficients corresponding to the neutrino–neutrino collision integral $\mathcal{I}_{\nu\nu}[\varrho(y)]$, we first assume that $\varrho(y') = f_{\text{eq}}(y')\mathbb{1}$ for all $y' \neq y$, where $f_{\text{eq}}(y')$ is the relativistic Fermi–Dirac distribution of some temperature which, for simplicity, we shall take to be the QED plasma temperature z .¹³ Then, the repopulation coefficient due to neutrino–neutrino collisions, $\{R_{\nu\nu}(y)\}_{\alpha} \equiv 2\{\Gamma_y^{\nu\nu}\}_{\alpha\alpha}$, can be constructed from the collision integral (A.8) to give

$$\begin{aligned} \{R_{\nu\nu}(y)\}_{\alpha} &= \frac{2G_F^2}{(2\pi)^3 y^2} \int dy_2 dy_3 \left[\Pi_{2a}^{\text{sc}}(y, y_2) + 2\Pi_{2b}^{\text{sc}}(y, y_4) \right] \\ &\quad \times \left([1 - f_{\text{eq}}(y_2)] f_{\text{eq}}(y_3) f_{\text{eq}}(y_4) + f_{\text{eq}}(y_2) [1 - f_{\text{eq}}(y_3)] [1 - f_{\text{eq}}(y_4)] \right) \\ &\equiv \mathcal{D}(y, z), \end{aligned} \quad (\text{B.9})$$

which is *flavour-blind* as expected (because in thermal equilibrium, there are equally many neutrinos and antineutrinos in all flavours).

For relativistic Fermi–Dirac distributions, the function $\mathcal{D}(y, z)$ evaluates to

$$\mathcal{D}(y, z) = \frac{2G_F^2 y z^4}{(2\pi)^3} d(y/z), \quad (\text{B.10})$$

which is predominantly linear in y , with a residual y -dependence encapsulated in the function $d(s = y/z)$, shown in figure 7 as a function of $s = y/z$. For computational ease, $d(s)$ can be fitted in the interval $s \in [10^{-4}, 10^3]$ to better than 0.25% accuracy by the curve

$$d_{\text{fit}}(s) = d_0 e^{-1.01s} + d_{\infty} (1 - e^{-0.01s}) + (e^{-0.01s} - e^{-1.01s}) \left[\frac{a_0 + a_1 \ln(s) + a_2 \ln^2(s)}{1 + b_1 \ln(s) + b_2 \ln^2(s)} \right], \quad (\text{B.11})$$

where $d_0 = 129.875$ and $d_{\infty} = 100.999$ are the asymptotic values of the function as $s \rightarrow 0$ and $s \rightarrow \infty$ respectively, and the fitting coefficients are $a_0 = 90.7332$, $a_1 = -48.4473$, $a_2 = 20.1219$, $b_1 = -0.529157$, and $b_2 = 0.20649$.

¹³It may be desirable to use the neutrino temperature estimated during run time instead of the QED plasma temperature. In practice, however, we find no significant difference in the outcome N_{eff} between these choices.

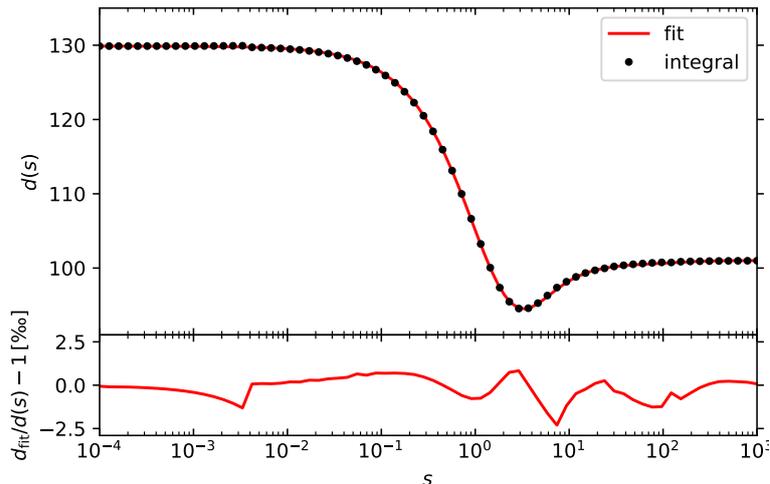


Figure 7: The function $d(s)$, as defined in equation (B.10). The dots denote actual numerical evaluations of the integral (B.9), while the line represents the fitting function $d_{\text{fit}}(s)$ of equation (B.11). The bottom panel shows the fractional difference between the two.

Neutrino–electron. The repopulation coefficients corresponding to the neutrino–electron collision integral $\mathcal{I}_{\nu e}[\varrho(y)]$, $\{R_{\nu e}(y)\}_\alpha \equiv 2\{\Gamma_y^{\nu e}\}_{\alpha\alpha}$, can be established similarly under the assumption of $\varrho(y') = f_{\text{eq}}(y')\mathbb{1}$ for all $y' \neq y$, where, again, $f_{\text{eq}}(y')$ is the relativistic Fermi–Dirac distribution with the QED plasma temperature z . We likewise assume the electron phase space distribution to be given by the same relativistic Fermi–Dirac form, i.e., $f_e(y) = f_{\text{eq}}(y)$. Then, using equation (A.1) to construct $\{\Gamma_y^{\nu e}\}_{\alpha\alpha}$, we find

$$\begin{aligned} \{R_{\nu e}(y)\}_\alpha &= \frac{G_F^2}{2(2\pi)^3 y^2} [(2\sin^2\theta_W \pm 1)_\alpha^2 + 4\sin^4\theta_W] \int dy_2 dy_3 [\Pi_{2a}^{\text{sc}}(y, y_2) + 2\Pi_{2b}^{\text{sc}}(y, y_4)] \\ &\quad \times \left([1 - f_{\text{eq}}(y_2)]f_{\text{eq}}(y_3)f_{\text{eq}}(y_4) + f_{\text{eq}}(y_2)[1 - f_{\text{eq}}(y_3)][1 - f_{\text{eq}}(y_4)] \right) \\ &= \frac{1}{4} [(2\sin^2\theta_W \pm 1)_\alpha^2 + 4\sin^4\theta_W] \mathcal{D}(y, z), \end{aligned} \tag{B.12}$$

where in the prefactor $(2\sin^2\theta_W \pm 1)_\alpha$ the plus sign “+” is understood to apply to $\alpha = e$ and “−” to $\alpha = \mu, \tau$, and $\mathcal{D}(y, z)$ is the same function given in equation (B.10).

It is of course also possible to retain a finite x , i.e., a finite electron mass m_e , in the computation of the repopulation and hence damping coefficients. However, the final outcome will have a more complicated time dependence than that contained in $\mathcal{D}(y, z)$, and for simplicity we have opted to omit this additional dependence.

References

- [1] S. Dodelson and M. S. Turner, “Nonequilibrium neutrino statistical mechanics in the expanding universe,” *Phys. Rev. D* **46** (1992) 3372–3387.
- [2] S. Hannestad and J. Madsen, “Neutrino decoupling in the early universe,” *Phys. Rev. D* **52** (1995) 1764–1769, [arXiv:astro-ph/9506015](#) [astro-ph].

- [3] A. D. Dolgov, S. H. Hansen, and D. V. Semikoz, “Nonequilibrium corrections to the spectra of massless neutrinos in the early universe,” *Nucl. Phys. B* **503** (1997) 426–444.
- [4] A. D. Dolgov, S. H. Hansen, and D. V. Semikoz, “Nonequilibrium corrections to the spectra of massless neutrinos in the early universe: Addendum,” *Nucl. Phys. B* **543** (1999) 269–274.
- [5] S. Esposito, G. Miele, S. Pastor, M. Peloso, and O. Pisanti, “Nonequilibrium spectra of degenerate relic neutrinos,” *Nucl. Phys. B* **590** (2000) 539–561.
- [6] G. Mangano, G. Miele, S. Pastor, T. Pinto, O. Pisanti, and P. D. Serpico, “Relic neutrino decoupling including flavor oscillations,” *Nucl. Phys. B* **729** (2005) 221–234, [arXiv:hep-ph/0506164 \[hep-ph\]](#).
- [7] J. Birrell, C.-T. Yang, and J. Rafelski, “Relic Neutrino Freeze-out: Dependence on Natural Constants,” *Nucl. Phys. B* **890** (2014) 481–517, [arXiv:1406.1759 \[nucl-th\]](#).
- [8] E. Grohs, G. M. Fuller, C. T. Kishimoto, M. W. Paris, and A. Vlasenko, “Neutrino energy transport in weak decoupling and big bang nucleosynthesis,” *Phys. Rev. D* **93** (2016) 083522, [arXiv:1512.02205 \[astro-ph.CO\]](#).
- [9] M. Escudero, “Neutrino Decoupling Beyond the Standard Model: CMB constraints on the Dark Matter mass with a fast and precise N_{eff} evaluation,” *JCAP* **02** (2019) 007, [arXiv:1812.05605 \[hep-ph\]](#).
- [10] M. Escudero Abenza, “Precision Early Universe Thermodynamics made simple: N_{eff} and Neutrino Decoupling in the Standard Model and beyond,” *JCAP* **05** (2020) 048, [arXiv:2001.04466 \[hep-ph\]](#).
- [11] J. Froustey and C. Pitrou, “Incomplete neutrino decoupling effect on big bang nucleosynthesis,” *Phys. Rev. D* **101** (2020) 043524, [arXiv:1912.09378 \[astro-ph.CO\]](#).
- [12] K. Akita and M. Yamaguchi, “A precision calculation of relic neutrino decoupling,” *JCAP* **08** (2020) 012, [arXiv:2005.07047 \[hep-ph\]](#).
- [13] J. Froustey, C. Pitrou, and M. C. Volpe, “Neutrino decoupling including flavour oscillations and primordial nucleosynthesis,” *JCAP* **12** (2020) 015, [arXiv:2008.01074 \[hep-ph\]](#).
- [14] D. A. Dicus, E. W. Kolb, A. M. Gleeson, E. C. G. Sudarshan, V. L. Teplitz, and M. S. Turner, “Primordial Nucleosynthesis Including Radiative, Coulomb, and Finite Temperature Corrections to Weak Rates,” *Phys. Rev. D* **26** (1982) 2694.
- [15] A. Heckler, “Astrophysical applications of quantum corrections to the equation of state of a plasma,” *Phys. Rev. D* **49** (1994) 611–617.
- [16] N. Fornengo, C. Kim, and J. Song, “Finite temperature effects on the neutrino decoupling in the early universe,” *Phys. Rev. D* **56** (1997) 5123–5134, [hep-ph/9702324](#).
- [17] R. E. Lopez and M. S. Turner, “An Accurate Calculation of the Big Bang Prediction for the Abundance of Primordial Helium,” *Phys. Rev. D* **59** (1999) 103502, [arXiv:astro-ph/9807279 \[astro-ph\]](#).
- [18] G. Mangano, G. Miele, S. Pastor, and M. Peloso, “A Precision calculation of the effective number of cosmological neutrinos,” *Phys. Lett. B* **534** (2002) 8–16, [arXiv:astro-ph/0111408 \[astro-ph\]](#).
- [19] J. J. Bennett, G. Buldgen, M. Drewes, and Y. Y. Wong, “Towards a precision calculation of the effective number of neutrinos N_{eff} in the Standard Model I: The QED equation of state,” *JCAP* **03** (2020) 003, [arXiv:1911.04504 \[hep-ph\]](#).
- [20] P. F. de Salas and S. Pastor, “Relic neutrino decoupling with flavour oscillations revisited,” *JCAP* **07** (2016) 051, [arXiv:1606.06986 \[hep-ph\]](#).
- [21] **Planck** Collaboration, N. Aghanim *et al.*, “Planck 2018 results. VI. Cosmological parameters,” *Astron. Astrophys.* **641** (2020) A6, [arXiv:1807.06209 \[astro-ph.CO\]](#).

- [22] S. Hannestad, R. S. Hansen, T. Tram, and Y. Y. Y. Wong, “Active-sterile neutrino oscillations in the early Universe with full collision terms,” *JCAP* **08** (2015) 019, [arXiv:1506.05266 \[hep-ph\]](#).
- [23] S. Hagstotz *et al.*, “Bounds on light sterile neutrino mass and mixing from cosmology and laboratory searches,” [arXiv:2003.02289 \[astro-ph.CO\]](#).
- [24] M. Archidiacono, S. Hannestad, A. Mirizzi, G. Raffelt, and Y. Y. Y. Wong, “Axion hot dark matter bounds after Planck,” *JCAP* **10** (2013) 020, [arXiv:1307.0615 \[astro-ph.CO\]](#).
- [25] J. Hasenkamp and J. Kersten, “Dark radiation from particle decay: cosmological constraints and opportunities,” *JCAP* **08** (2013) 024, [arXiv:1212.4160 \[hep-ph\]](#).
- [26] P. Di Bari, S. F. King, and A. Merle, “Dark Radiation or Warm Dark Matter from long lived particle decays in the light of Planck,” *Phys. Lett. B* **724** (2013) 77–83, [arXiv:1303.6267 \[hep-ph\]](#).
- [27] **CMB-S4** Collaboration, K. N. Abazajian *et al.*, “CMB-S4 Science Book, First Edition,” [arXiv:1610.02743 \[astro-ph.CO\]](#).
- [28] M. Archidiacono, T. Basse, J. Hamann, S. Hannestad, G. Raffelt, and Y. Y. Y. Wong, “Future cosmological sensitivity for hot dark matter axions,” *JCAP* **05** (2015) 050, [arXiv:1502.03325 \[astro-ph.CO\]](#).
- [29] K. N. Abazajian and J. Heeck, “Observing Dirac neutrinos in the cosmic microwave background,” *Phys. Rev. D* **100** (2019) 075027, [arXiv:1908.03286 \[hep-ph\]](#).
- [30] S. Gariazzo, P. F. de Salas, and S. Pastor, “Thermalisation of sterile neutrinos in the early Universe in the 3+1 scheme with full mixing matrix,” *JCAP* **07** (2019) 014, [arXiv:1905.11290 \[astro-ph.CO\]](#).
- [31] G. Sigl and G. Raffelt, “General kinetic description of relativistic mixed neutrinos,” *Nucl. Phys. B* **406** (1993) 423–451.
- [32] P. F. de Salas, D. V. Forero, C. A. Ternes, M. Tórtola, and J. W. F. Valle, “Status of neutrino oscillations 2018: 3 σ hint for normal mass ordering and improved CP sensitivity,” *Phys. Lett. B* **782** (2018) 633–640, [arXiv:1708.01186 \[hep-ph\]](#).
- [33] P. F. de Salas *et al.*, “2020 global reassessment of the neutrino oscillation picture,” *JHEP* **02** (2021) 071, [arXiv:2006.11237 \[hep-ph\]](#).
- [34] I. Esteban, M. Gonzalez-Garcia, M. Maltoni, T. Schwetz, and A. Zhou, “The fate of hints: updated global analysis of three-flavor neutrino oscillations,” *JHEP* **09** (2020) 178, [arXiv:2007.14792 \[hep-ph\]](#).
- [35] F. Capozzi, E. Di Valentino, E. Lisi, A. Marrone, A. Melchiorri, and A. Palazzo, “Addendum to: Global constraints on absolute neutrino masses and their ordering,” *Phys. Rev. D* **101** (2020) 116013, [arXiv:2003.08511 \[hep-ph\]](#).
- [36] D. Nötzold and G. Raffelt, “Neutrino Dispersion at Finite Temperature and Density,” *Nucl. Phys. B* **307** (1988) 924–936.
- [37] J. I. Kapusta and C. Gale, *Finite-temperature field theory: Principles and applications*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2011.
- [38] O. Tomalak and R. Hill, “Theory of elastic neutrino-electron scattering,” *Phys. Rev. D* **101** (2020) 033006, [arXiv:1907.03379 \[hep-ph\]](#).
- [39] R. J. Hill and O. Tomalak, “On the effective theory of neutrino-electron and neutrino-quark interactions,” *Phys. Lett. B* **805** (2020) 135466, [arXiv:1911.01493 \[hep-ph\]](#).
- [40] V. P. Silin, “On the electromagnetic properties of a relativistic plasma,” *Sov. Phys. JETP* **11** (1960) 1136–1140.

- [41] P. B. Arnold, G. D. Moore, and L. G. Yaffe, “Effective kinetic theory for high temperature gauge theories,” *JHEP* **01** (2003) 030, [arXiv:hep-ph/0209353 \[hep-ph\]](#).
- [42] A. Anisimov, W. Buchmuller, M. Drewes, and S. Mendizabal, “Nonequilibrium Dynamics of Scalar Fields in a Thermal Bath,” *Annals Phys.* **324** (2009) 1234–1260, [arXiv:0812.1934 \[hep-th\]](#).
- [43] M. Drewes, “On the Role of Quasiparticles and thermal Masses in Nonequilibrium Processes in a Plasma,” [arXiv:1012.5380 \[hep-th\]](#).
- [44] V. V. Klimov, “Collective excitations in a hot quark-gluon plasma,” *Sov. Phys. JETP* **55** (1982) 199–204. [*Zh. Eksp. Teor. Fiz.* **82** (1982) 336–345].
- [45] H. A. Weldon, “Effective Fermion Masses of Order gT in High Temperature Gauge Theories with Exact Chiral Invariance,” *Phys. Rev. D* **26** (1982) 2789.
- [46] B. H. McKellar and M. J. Thomson, “Oscillating doublet neutrinos in the early universe,” *Phys. Rev. D* **49** (1994) 2710–2728.
- [47] N. F. Bell, R. R. Volkas, and Y. Y. Y. Wong, “Relic neutrino asymmetry evolution from first principles,” *Phys. Rev. D* **59** (1999) 113001, [arXiv:hep-ph/9809363 \[hep-ph\]](#).
- [48] K. Kainulainen and A. Sorri, “Oscillation induced neutrino asymmetry growth in the early universe,” *JHEP* **02** (2002) 020, [arXiv:hep-ph/0112158](#).
- [49] E. Tiesinga, P. J. Mohr, D. B. Newell, and B. N. Taylor, “The 2018 CODATA Recommended Values of the Fundamental Physical Constants.” <https://physics.nist.gov/constants>, 2018.
- [50] K. Kumar, S. Mantry, W. Marciano, and P. Souder, “Low Energy Measurements of the Weak Mixing Angle,” *Ann. Rev. Nucl. Part. Sci.* **63** (2013) 237–267, [arXiv:1302.6263 \[hep-ex\]](#).
- [51] J. Erler and S. Su, “The Weak Neutral Current,” *Prog. Part. Nucl. Phys.* **71** (2013) 119–149, [arXiv:1303.5522 \[hep-ph\]](#).
- [52] **Particle Data Group** Collaboration, P. Zyla *et al.*, “Review of Particle Physics,” *PTEP* **2020** (2020) 083C01.
- [53] I. Esteban, M. Gonzalez-Garcia, A. Hernandez-Cabezudo, M. Maltoni, and T. Schwetz, “Global analysis of three-flavour neutrino oscillations: synergies and tensions in the determination of θ_{23} , δ_{CP} , and the mass ordering,” *JHEP* **01** (2019) 106, [arXiv:1811.05487 \[hep-ph\]](#).
- [54] P. F. de Salas, S. Gariazzo, O. Mena, C. A. Ternes, and M. Tórtola, “Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects,” *Front. Astron. Space Sci.* **5** (2018) 36, [arXiv:1806.11051 \[hep-ph\]](#).
- [55] A. C. Hindmarsh, “ODEPACK, a Systematized Collection of ODE Solvers,” in *Scientific Computing*, R. Stepleman, ed. Elsevier, 1983.
- [56] K. Radhakrishnan and A. C. Hindmarsh, “Description and Use of LSODE, the Livermore Solver for Ordinary Differential Equations,” Tech. Rep. UCRL-ID-113855, Lawrence Livermore National Laboratory, 1993.
- [57] C. Lunardini and A. Yu. Smirnov, “High-energy neutrino conversion and the lepton asymmetry in the universe,” *Phys. Rev. D* **64** (2001) 073006, [arXiv:hep-ph/0012056 \[hep-ph\]](#).
- [58] **JUNO** Collaboration, F. An *et al.*, “Neutrino Physics with JUNO,” *J. Phys. G* **43** (2016) 030401, [arXiv:1507.05613 \[physics.ins-det\]](#).
- [59] S. A. R. Ellis, K. J. Kelly, and S. W. Li, “Current and Future Neutrino Oscillation Constraints on Leptonic Unitarity,” *JHEP* **12** (2020) 068, [arXiv:2008.01088 \[hep-ph\]](#).
- [60] D. N. Blaschke and V. Cirigliano, “Neutrino Quantum Kinetic Equations: The Collision Term,” *Phys. Rev. D* **94** (2016) 033009, [arXiv:1605.09383 \[hep-ph\]](#).