International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 10 Issue: 04 | April - 2026 SJIF Rating: 8.659 ISSN: 2582-3930

Hot Dense Matter in Neutron Star Merger Remnants: Finite-Temperature Equation of State, Neutrino Transport, Bulk Viscosity, Convective Instabilities, and the Helium Spectroscopic Clock — A Unified Framework for Reading the QCD Phase Diagram from Transient Observations

Dr. Arti Pandoh Gupta[1]

[1]Government Degree College Batote, District Ramban, Jammu, India

——— ★★★ ———

Abstract – This paper constitutes the work on dense nuclear matter, advancing beyond the cold-matter equation of state (EoS), the hyperon puzzle, and the multimessenger aftermath of binary neutron star (BNS) mergers to address the most poorly understood regime of the QCD phase diagram: hot, dense, out-of-equilibrium nuclear matter at temperatures $T \approx$ 10–100 MeV and baryon densities $(1\text{–}5)\rho_0$ as realised in hypermassive neutron star (HMNS) remnants. We synthesise four interconnected theoretical frontiers that have each seen transformative advances in 2024–2025. First, we examine the three-dimensional finite-temperature EoS problem — the challenge of constructing fully tabulated $P(\rho, T, Ye)$ tables consistent with nuclear theory, heavy-ion constraints, and astrophysical observations — and evaluate the accuracy of the M^*-parametric framework for thermal extensions versus fully microscopic finite-temperature EoS tables (SFHo, DD2, BLh), demonstrating that finite-temperature EoS choices produce distinct postmerger gravitational-wave frequency evolutions at late times ($t > 50$ ms) that deviate significantly from quasi-universal relations. Second, we present a comprehensive treatment of neutrino transport physics in the merger remnant: the emergence and saturation of bulk viscosity driven by weak-interaction Urca processes, the role of trapped neutrinos in modifying flavor equilibration timescales, and the magnetic-field-induced modification of Urca rates through the Nucleon Width Approximation framework. Third, we analyse the Tayler–Spruit (TS) dynamo operating in the differentially-rotating HMNS as the dominant mechanism for field amplification beyond the Kelvin–Helmholtz phase, generating toroidal fields $B \sim 10^{16}$ G on timescales of tens of milliseconds and driving thermal instabilities and convective patterns that excite sub-dominant inertial modes with frequencies $f < 2$ kHz, potentially detectable by third-generation GW detectors. Fourth, and most originally, we develop the helium spectroscopic clock as a quantitative EoS constraint: the absence of a prominent He I λ1083.3 nm feature in AT2017gfo at 4.4 days post-merger limits the helium mass fraction in polar ejecta to $X\_He < 0.05$, which, combined with neutrino-hydrodynamic simulations, constrains the remnant lifetime to $\tau\_rem \leq$ 20–30 ms. This implies $M\_thres \leq 2.93\ M_\odot$ and, combined with causality arguments, limits the maximum NS mass to $M\_max \leq 2.3\ M_\odot$ and the radius of 1.6 $M_\odot$ NS to approximately 11–13 km. We place all four threads within a unified observational framework, demonstrating that neutrino emission, convective mode excitation, bulk viscous dissipation, and kilonova spectroscopic features constitute a quartet of orthogonal, complementary constraints on the finite-temperature EoS, each sensitive to a distinct region of the (ρ, T) plane inaccessible to cold-matter probes.

Key Words: *finite-temperature equation of state, neutron star merger remnant, neutrino transport, bulk viscosity, Urca processes, Tayler-Spruit dynamo, helium spectroscopy, AT2017gfo, convective instability, hypermassive neutron star, QCD phase diagram, kilonova.*

1. INTRODUCTION AND SERIES CONTEXT

The preceding three papers of this series established a coherent theoretical framework spanning the nuclear saturation regime (Paper I: symmetry energy, PREX-II, CREX, NICER, GW170817), the supra-nuclear cold-matter regime (Paper II: hyperon puzzle, hadron–quark phase transition, speed-of-sound diagnostics), and the multimessenger observational aftermath of BNS mergers (Paper III: postmerger gravitational waves, kilonova nucleosynthesis, r-process element identification, and third-generation detector prospects). A recurring and deliberately deferred thread across all three papers is the finite-temperature equation of state: the thermodynamic behaviour of nuclear matter at the temperatures $T \approx$ 10–100 MeV reached during and after the collision of two neutron stars. This paper addresses that thread directly and comprehensively.

The merger of two neutron stars is unique among astrophysical events in simultaneously accessing multiple sectors of the QCD phase diagram. The cold, β-equilibrated matter of the pre-merger stars — probed by NICER radii and GW tidal deformability — transitions within milliseconds of contact into shock-heated matter at temperatures far exceeding those found anywhere else in the cold universe. The HMNS remnant that forms is a transient object of extraordinary complexity: differentially rotating at several kilohertz, permeated by magnetic fields amplified to magnetar-level strengths, threaded by intense neutrino fluxes that redistribute lepton number throughout the ejecta, and simultaneously hosting nucleonic matter, hyperonic matter, and potentially a nascent quark-matter core — all at finite temperature. This is not merely a technical complication to be modelled away; the finite-temperature properties of dense matter are physically consequential observables that encode information about the strong interaction at conditions that no terrestrial experiment can reproduce.

Four developments from 2024–2025, reviewed comprehensively in this paper, represent transformative advances in understanding this regime. First, long-term numerical relativity simulations extending to 150 ms post-merger have demonstrated for the first time that fully tabulated finite-temperature EoS tables produce qualitatively distinct postmerger gravitational-wave frequency evolution from their parametric M* approximations, with deviations from quasi-universal relations becoming significant at late times and exciting long-lived convective inertial modes potentially within the sensitivity band of the Einstein Telescope (Giacomazzo et al. arXiv:2512.05118, 2024–2025). Second, the inclusion of microphysical bulk viscosity from Urca processes in general-relativistic merger simulations (Most et al. 2024) has demonstrated that weak-interaction-driven out-of-equilibrium dynamics imprints a measurable phase shift on the postmerger gravitational-wave signal, producing a new observable channel for constraining the direct Urca threshold density and hence the symmetry energy slope L. Third, the Tayler–Spruit (TS) dynamo has been identified as the dominant magnetic field amplification mechanism on timescales of tens of milliseconds in the HMNS core, operating in the thermally-stratified differentially-rotating remnant and generating toroidal fields that drive thermomagnetic instabilities (Barrère et al. 2024–2025). Fourth, a novel EoS constraint has emerged from kilonova spectroscopy: the non-detection of a pronounced He I λ1083.3 nm feature in AT2017gfo at 4.4 days post-merger limits the helium content of polar ejecta, constraining the remnant lifetime to at most 20–30 ms and thereby placing upper bounds on M_max and neutron star radii (Bauswein et al. arXiv:2411.03427, 2025).

Together, these four advances constitute a new observational paradigm in which the finite-temperature EoS — long treated as an uncertain background parameter in BNS merger simulations — becomes a directly measurable quantity through combined gravitational-wave, neutrino, magnetic, and spectroscopic observational channels. The aim of this paper is to synthesise these developments within a unified theoretical framework and to map, for the first time, each observable to its sensitivity domain in the (ρ, T) plane of the QCD phase diagram, providing a roadmap for the precision finite-temperature EoS measurements that third-generation GW detectors and next-generation kilonova spectroscopy will enable.

2. THE FINITE-TEMPERATURE EQUATION OF STATE: STRUCTURE, METHODS, AND CHALLENGES

2.1 Thermodynamic Structure of the Finite-Temperature EoS

The equation of state of nuclear matter at finite temperature is a function of three independent thermodynamic variables: the baryon number density ρ_B, the temperature T, and the electron fraction (charge fraction) $Y_e \equiv (n_p + n_e)/n_B$. For cold, β-equilibrated matter — the relevant limit for isolated neutron stars — Y_e is determined by the condition $\mu_n = \mu_p + \mu_e + \mu_\nu$ and the full three-dimensional EoS table $P(\rho_B, T, Y_e)$ reduces to a one-dimensional $P(\rho_B)$ curve. However, in the HMNS remnant, neither β-equilibrium nor charge neutrality at the local level is guaranteed: shock-heated regions develop large local chemical potential imbalances $\delta\mu \equiv \mu_n - \mu_p - \mu_e \neq 0$, and neutrinos are trapped in the densest, hottest regions ($T > 10$ MeV, $\rho > \rho_0$) while streaming freely in the outer layers. The full three-dimensional EoS table is therefore required for accurate merger simulations.

Constructing a physically consistent three-dimensional EoS table that satisfies all known nuclear theory and observational constraints is computationally and theoretically challenging. The leading publicly available finite-temperature EoS tables — SFHo (Steiner, Fischer, Hempel 2013), DD2 (Typel et al. 2010, extended by Hempel and Schaffner-Bielich), and BLh (Bombaci, Logoteta, Hempel 2018) — are based on relativistic mean field (RMF) models calibrated to nuclear saturation properties, binding energies, and charge radii, but extrapolated to supra-nuclear densities and finite temperatures using mean-field approximations whose validity at $T > 30$ MeV and $\rho > 3\rho_0$ is uncertain. The SFHo EoS, derived from an RMF model with a relatively soft symmetry energy ($L \approx 47$ MeV), predicts $R_{1.4} \approx 11.9$ km and M_max ≈ 2.06 M⊙, consistent with current multimessenger constraints. DD2 is stiffer ($L \approx 55$ MeV), predicting $R_{1.4} \approx 13.2$ km but still accommodating 2 M⊙ pulsars. BLh incorporates two-body and three-body nuclear forces from the Brueckner–Hartree–Fock (BHF) approach, providing a firmer microscopic basis but at the cost of less flexibility for finite-temperature extensions.

Mroczek et al. (Physical Review D, 2024) introduced a systematic finite-temperature expansion framework for the dense-matter EoS. By Taylor-expanding the pressure P(ρ, T, Y_e) in powers of both the isospin asymmetry δ = (n_n − n_p)/n_B and the thermal parameter T/μ_B around the cold β-equilibrated limit, they demonstrated that the finite-temperature EoS can be reconstructed to within 5% accuracy for μ_B ≥ 1100 MeV (corresponding to $\rho \geq \rho_0$) and $T \leq 100$ MeV. This range encompasses virtually all physically relevant conditions in the HMNS remnant ($T \leq 70$ MeV, μ_B ≥ 950 MeV). The expansion parameter (T/μ_B) < 0.1 throughout the merger-relevant domain, ensuring rapid convergence. This framework provides, for the first time, a quantitative uncertainty quantification for finite-temperature EoS extensions that does not assume a specific hadronic or quark model.

2.2 The M* Parametric Framework and Its Limitations

In merger simulations that do not have access to a fully tabulated finite-temperature EoS table for a given cold EoS model, a widely adopted approximation is the M* (effective mass) parametric framework introduced by Raithel et al. (2019). In this approach, the full finite-temperature EoS is approximated as:

$$P(n_B, T, Y_e) = P_{cold}(n_B, Y_e) + P_{thermal}(n_B, T; m^*, \gamma_{th})$$

where the thermal pressure P_thermal is parametrised by the nucleon effective mass m*(n_B) and a thermal adiabatic index γ_th. The M* model captures the leading-order effect of finite temperature on the Fermi energy of the nucleon system and reproduces the correct non-relativistic limit. Ujevic et al. (Astrophysical Journal Letters, 2024) validated this framework against full SFHo finite-temperature simulations, finding that the M* approximation accurately reproduces the temperature and thermal pressure profiles of the HMNS remnant when calibrated to the correct m* and γ_th.

However, long-term simulations extending to 150 ms post-merger (Giacomazzo et al. arXiv:2512.05118, December 2024) revealed a critical limitation of the M* framework. At early times ($t < 30$ ms post-merger), M* and tabulated EoS simulations agree well in the dominant f_2 postmerger GW frequency. At later times ($t > 50$ ms), distinct differences emerge in the GW frequency evolution: tabulated EoS simulations develop a systematic frequency drift driven by compositional and thermal gradients in the remnant that the M* approximation, with its simplified thermal structure, cannot reproduce. These late-time deviations are associated with the excitation of inertial oscillation modes — rotationally driven oscillations with frequencies $f < 2$ kHz — through convective patterns that develop from Ledoux and Solberg–Høiland instabilities in the thermally stratified differentially-rotating remnant. The physical process is as follows: differential rotation creates a radial temperature gradient that, combined with the entropy stratification from shock heating, generates a non-zero Brunt–Väisälä frequency N in the remnant core; when the Solberg–Høiland criterion for convective instability is satisfied, sustained convective patterns develop that excite inertial modes with sub-f_2 frequencies. The tabulated EoS preserves this thermal stratification consistently, while M* parametric models, which do not self-consistently track the entropy profile, suppress the instability.

Table –1: Properties of Leading Finite-Temperature EoS Tables and Their Postmerger GW Predictions

Model	L (MeV)	$R_{1.4}$ (km)	M_max (M⊙)	f_2 at 40ms (kHz)	Late-drift	Inertial modes
SFHo	47.1	11.9	2.06	~3.0	Moderate	Yes (tabulated)
DD2	55.0	13.2	2.42	~2.6	Strong	Yes (tabulated)
BLh (BHF)	61.0	12.5	2.10	~2.8	Moderate	Weaker
M* approx.	Varies	Cold EoS	Cold EoS	Agrees early	Suppressed	Not captured

3. NEUTRINO TRANSPORT AND BULK VISCOSITY IN THE MERGER REMNANT

3.1 Neutrino Emission: Luminosities, Flavour Hierarchy, and EoS Coupling

Neutrinos are the dominant cooling agent in the HMNS remnant, carrying away thermal energy at a luminosity of several $\times 10^{53}$ erg s^{-1} during the first few tens of milliseconds post-merger. The systematic study by Cusinato et al. (European Physical Journal A, 2022, extended in 2024) across a large sample of numerical relativity simulations covering a broad range of EoS models established several robust features of the neutrino emission. The total neutrino luminosity decreases with increasing reduced tidal deformability $\tilde{\Lambda}$ of the binary, because a stiffer EoS produces less violent merger dynamics and correspondingly milder shock heating. Peak luminosities can reach twice the time-averaged value. Electron antineutrino luminosities dominate over electron neutrino luminosities by a factor of 2–3 in the first few milliseconds, reflecting the greater abundance of positrons that drive $\bar{\nu}\_e$ emission through positron capture on neutrons.

A critical observational consequence of these results is that the neutrino luminosity and its time evolution encode information about the tidal deformability $\tilde{\Lambda}$ of the binary, and therefore about the cold EoS. A measurement of the neutrino burst from a nearby ($\leq$ 5 Mpc) BNS merger with a large neutrino detector (such as the Hyper-Kamiokande or DUNE) could in principle constrain $\tilde{\Lambda}$ independently of gravitational-wave measurements, providing a third complementary EoS probe from a single merger event. The mean neutrino energies — approximately 10 MeV for ν_e, 12 MeV for $\bar{\nu}\_e$, and 15–20 MeV for heavy-flavour neutrinos — reflect the temperature at the respective neutrino spheres and are nearly independent of the binary parameters, providing a robust calibration point for neutrino transport models.

3.2 Bulk Viscosity from Urca Processes: Emergence and EoS Consequences

In isolated cold neutron stars, β-equilibrium is maintained by the modified Urca processes ($n + n \rightarrow n + p + e^- + \bar{\nu}\_e$ and its reverse), with rates that are suppressed by the available phase space near the Fermi surface. In the HMNS remnant, the merger-induced density oscillations drive the matter out of β-equilibrium on a dynamical timescale, creating a non-zero chemical potential imbalance $\delta\mu \equiv \mu\_n - \mu\_p - \mu\_e$. The relaxation of δμ back to zero through Urca processes drives a net energy dissipation that constitutes the bulk viscosity of the nuclear matter.

The significance of this process was first quantitatively established in general-relativistic simulations by Most and Raithel (2021, 2024), who incorporated direct and modified Urca processes (in the neutrino-transparent regime) into the hydrodynamic evolution equations as a source term for the bulk scalar pressure correction Π:

$$d\Pi/dt + \Pi/\tau_0 = -\zeta/\tau_0 \times \nabla\cdot v$$

where ζ(ω, T, ρ) is the frequency-dependent bulk viscosity coefficient, τ_0 is the Urca relaxation timescale, and v is the fluid velocity. Their key finding is that bulk viscosity from weak interactions produces a measurable phase shift in the postmerger gravitational-wave signal, shifting the dominant oscillation frequency and modifying the long-term phase evolution in a way that depends on the direct Urca threshold density n_DU. This threshold — the density above which the direct Urca process $n \rightarrow p + e + \bar{\nu}\_e$ is kinematically allowed without requiring a spectator nucleon — is directly controlled by the proton fraction at supra-nuclear densities and hence by the symmetry energy slope L. EoS models with high L reach the direct Urca threshold at lower densities, enabling faster β-equilibration and consequently larger bulk viscous dissipation in the remnant.

The 2025 extension of this work by Tambe, Chatterjee, Alford, and Haber incorporated the effect of the amplified magnetic field on Urca rates through the Nucleon Width Approximation framework (Physical Review C, 111, 035809, 2025). The magnetic field B ~ 10^{14}–10^{15} G in the postmerger remnant quantizes the available electron phase space through Landau levels, modifying both the direct Urca rate and the modified Urca cross-section. Their central result is that strong magnetic fields shift the effective flavor equilibration condition, modifying the bulk viscous dissipation rate by up to 30–50% for the strongest physically plausible field configurations. This represents a non-negligible systematic uncertainty in postmerger EoS inference from bulk viscosity signatures that must be accounted for in third-generation detector analysis pipelines.

3.3 Composition g-Modes and Bulk Viscous Damping

A particularly elegant manifestation of bulk viscosity in warm neutron stars is the resonant damping of composition g-modes — buoyancy-driven oscillations whose restoring force is the gradient of the electron fraction Y_e in the stellar interior. In warm stars ($T \approx 1–5$ MeV), the Urca relaxation timescale τ_{Urca} becomes comparable to the characteristic period of composition g-modes (typically 0.1–1 s), producing a resonance condition in which the mode's driving and damping rates are comparable. Rau et al. (2023, extended 2024) introduced the complex-valued dynamic sound speed:

$$c_s^2(\omega, T) = c_{s,ad}^2 + i\omega\zeta(\omega)/(\varepsilon + P)$$

where $c_{s,ad}^2$ is the adiabatic sound speed, $\zeta(\omega)$ the frequency-dependent bulk viscosity, and ε, P the energy density and pressure. Using this frequency-dependent generalisation, they computed the complex frequencies of composition g-modes as a function of temperature for three representative EoS models, demonstrating that bulk viscous damping completely suppresses composition g-modes at $T \geq 3$ MeV and can significantly modify the mode spectrum even at $T \approx 1$ MeV. This suppression is directly observable as a deficit of low-frequency ($f < 1$ kHz) substructure in the postmerger GW spectrum for remnants that remain long-lived enough to cool below the relevant temperature threshold, providing a time-domain EoS diagnostic that is entirely complementary to the dominant f_2 oscillation.

4. MAGNETIC FIELD AMPLIFICATION: BEYOND THE KELVIN-HELMHOLTZ PHASE

4.1 The Tayler–Spruit Dynamo in Differentially-Rotating Remnants

In our preceding paper (Paper III), we discussed the amplification of magnetic fields during the merger through the Kelvin–Helmholtz instability (KHI) at the shear interface, which operates on millisecond timescales and amplifies the pre-merger field from $B \sim 10^{12}$ G to $B \sim 10^{14}$ G. However, after the KHI saturates at approximately 5–10 ms post-merger, a question arises: what mechanism sustains and potentially further amplifies the magnetic field over the longer timescale of the remnant lifetime (10–100 ms)?

The answer, established by recent 3D GRMHD simulations (Barrère et al. 2024–2025), is the Tayler–Spruit (TS) dynamo. This mechanism operates specifically in differentially-rotating, stably-stratified media — precisely the conditions prevailing in the HMNS core at radii $r < 10$ km, where the density profile is Ledoux-stable but the differential rotation profile shows a strong radial shear with $\Omega \sim 4000–7000\ \mathrm{s}^{-1}$. The TS dynamo mechanism involves two coupled instabilities: the winding of a seed poloidal field by differential rotation to produce a toroidal field (the Ω-effect), and the subsequent Tayler (Pitts–Tayler) kink instability of the toroidal field that regenerates a poloidal component. The cycle sustains itself as long as differential rotation is maintained.

Barrère et al. computed the Brunt–Väisälä frequency in the TS-unstable region of the HMNS core from the 3D GRMHD simulation of Kiuchi et al. (2024), finding $N \approx 4970\ \mathrm{s}^{-1}$ at a radius of $R_{TS} \approx 7$ km where $\rho \approx 3.7 \times 10^{14}\ \mathrm{g\ cm}^{-3}$. The TS dynamo generates a predominantly toroidal field growing at a rate $\sigma_{TS} \approx \omega_A$ (the Alfvén frequency), amplifying the post-KHI field from $B \sim 10^{14}$ G toward $B \sim 10^{16}$ G on a timescale of $\tau_{TS} \sim 10–50$ ms. The resulting strong toroidal field drives thermomagnetic instabilities through ambipolar diffusion that create localised regions of non-uniform entropy, coupling the magnetic and thermal structure of the remnant in a way that the purely hydrodynamic M* approximation cannot capture.

4.2 Convective Instabilities, Inertial Modes, and Sub-f_2 Gravitational-Wave Emission

The thermomagnetic coupling driven by the TS dynamo has a directly observable consequence: it establishes a sustained pattern of convective cells in the HMNS core that persist well beyond 100 ms post-merger. These convective patterns excite inertial oscillation modes — rotationally-modified internal waves whose restoring force is the Coriolis effect — at frequencies substantially below the dominant f_2 quadrupolar mode. The mode frequencies depend on the local rotation rate Ω, the Brunt–Väisälä frequency N, and the density stratification, all of which are sensitive to the finite-temperature EoS through the entropy profile.

Long-term simulations by Giacomazzo et al. (arXiv:2512.05118, December 2024) extending to 150 ms post-merger with both tabulated SFHo and DD2 EoS tables found that these inertial modes appear consistently in both EoS models, with frequencies $f_{iner} \approx 1–2$ kHz, but that their precise frequencies and growth rates differ between EoS models at a level of 100–200 Hz. Crucially, these modes are excited exclusively in the fully tabulated EoS simulations; the M* approximation, which does not self-consistently propagate the entropy stratification, fails to excite them at comparable amplitude. The energy carried in these inertial modes, while subdominant (approximately 5–15% of the total GW energy compared to f_2), is potentially within the sensitivity range of the Einstein Telescope for merger events at distances ≤ 20 Mpc, making them a new EoS-sensitive gravitational-wave observable available only to third-generation detectors.

Table –2: Finite-Temperature EoS Observables and Their Sensitivity Domain in the (ρ, T) Plane

Observable	Density range	Temp. range	Primary EoS sensitivity	Detection channel
f_2 GW frequency	$2–3\rho_0$	10–50 MeV	Cold EoS + thermal pressure	ET/CE (kHz)
Late f_2 drift	$2–3\rho_0$	30–60 MeV	Finite-T entropy stratification	ET/CE (t > 50 ms)
Inertial modes	$1.5–2.5\rho_0$	20–50 MeV	Brunt-Väisälä freq.; TS dynamo	ET/CE (f < 2 kHz)
Bulk viscosity GW phase	$1–3\rho_0$	10–50 MeV	Urca threshold; L parameter	ET/CE (phase)
Neutrino luminosity	$1–3\rho_0$	10–70 MeV	Tidal deformability $\tilde{\Lambda}$; shock heat	HyperK/DUNE (≤ 5 Mpc)
Helium spectral clock	Ejecta ($\leq\rho_0$)	5–30 MeV	Remnant lifetime; M_thres	VLT/JWST/ELT

5. THE HELIUM SPECTROSCOPIC CLOCK: A NEW EoS CONSTRAINT FROM KILONOVA SPECTRA

5.1 Theoretical Basis: Helium Production in Neutrino-Driven Winds

When the HMNS remnant survives for more than a few milliseconds without collapsing to a black hole, the intense neutrino flux irradiating the remnant's accretion disk and dynamical ejecta raises the electron fraction Y_e of the polar wind component through the reaction $\nu_e + n \rightarrow p + e^-$. High-Y_e ejecta (Y_e ≳ 0.4) undergo alpha-rich freeze-out, synthesising significant quantities of helium-4 (^{4}He) as a byproduct. The mass fraction of helium produced in the polar wind is a steeply increasing function of the remnant lifetime τ_rem: a long-lived remnant (100–200 ms) irradiates a large polar wind mass and produces a substantial helium abundance X_He ≳ 0.1, while a short-lived remnant (τ_rem ≤ 20 ms) irradiates comparatively little material and produces X_He < 0.05 in the polar ejecta.

Bauswein, Darbha, Hotokezaka, Kasen, and collaborators (arXiv:2411.03427, November 2024; published Physical Review D, 2025) exploited this connection to derive a novel EoS constraint from the AT2017gfo kilonova spectrum. The He I λ1083.3 nm line — a near-infrared transition from the metastable $2s^3S$ state of neutral helium — produces a characteristic P Cygni absorption feature in the 800–1200 nm range when helium is present at mass fractions X_He ≳ 0.05 in the line-forming region of the kilonova ejecta. Non-local thermodynamic equilibrium (NLTE) population modelling of helium under the physical conditions of the AT2017gfo ejecta at 4.4 days post-merger shows that the metastable $2s^3S$ state is significantly populated by high-energy photons from radioactive r-process decay chains and by the UV-optical radiation field of the hot kilonova continuum.

5.2 The AT2017gfo Constraint and Its EoS Interpretation

The AT2017gfo VLT X-shooter spectrum at 4.4 days post-merger shows an absorption feature near 1 μm that has been attributed primarily to the Sr II near-infrared triplet (Watson et al. 2019). Bauswein et al. (2024–2025) analysed the constraint that the helium mass fraction must place on this feature: they demonstrated that a helium mass fraction X_He ≳ 0.05 in the polar ejecta would produce a He I P Cygni feature at the correct wavelength with sufficient equivalent width to be detectable in the observed spectrum, but that the observed spectrum is inconsistent with such a prominent helium contribution.

This upper limit X_He < 0.05 translates, through neutrino-hydrodynamic simulations of merger remnants with multiple EoS models, into an upper limit on the remnant lifetime: τ_rem ≤ 20–30 ms for the GW170817 system with M_tot ≈

2.74 M⊙. Because the remnant lifetime is an EoS-dependent quantity — stiffer EoS models support longer-lived remnants before collapse, while softer EoS models collapse faster — this lifetime upper limit translates directly into an EoS constraint. Specifically, it implies that the threshold binary mass for prompt (or very rapid) gravitational collapse M_thres satisfies M_thres ≤ 2.93 M⊙. Using the empirical relations between M_thres, M_max, and the stellar radius R_max (established by Bauswein et al. 2017 and refined subsequently), this bound implies:

$$M\_max \leq 2.3\,M\odot \quad \text{and} \quad R(1.6\,M\odot) \approx 12 \pm 1\text{ km} \quad (\text{for } M\_max \approx 2.0\text{–}2.15\,M\odot)$$

This constraint rules out a significant subset of current EoS models — specifically those that predict simultaneously large M_max and large $R_{1.4}$, which would produce long remnant lifetimes inconsistent with the helium abundance limit. Combined with the lower limit on $R_{1.4}$ from the argument that AT2017gfo's brightness rules out prompt collapse ($R \gtrsim 10.5$ km), the helium spectroscopic clock provides a two-sided constraint on the neutron star radius.

5.3 The Helium vs. Strontium Debate and Resolution

An important complication in interpreting the 1 μm P Cygni feature of AT2017gfo is the competing attribution to He I λ1083.3 nm versus Sr II. Sneppen, Damgaard, Watson, Collins, Shingles, and Sim (Astronomy & Astrophysics, December 2024) performed a comprehensive collisional-radiative NLTE modelling of helium in the AT2017gfo ejecta spanning epochs from 0.92 to 4.4 days post-merger. Their central conclusion is definitive: the 1 μm P Cygni feature in AT2017gfo is inconsistent with a He I interpretation in both its emergence time and spectro-temporal evolution.

Specifically, the helium feature predicted by their NLTE model fails on two independent grounds. First, the earliest detection of the 1 μm feature at 1.17 days post-merger (SALT spectrum) shows an absorption minimum at ~800 nm that is inconsistent with the wavelength expected for a He I P Cygni feature expanding at the velocities inferred from the Sr II identification. Second, the temporal evolution of the feature equivalent width and absorption minimum wavelength is inconsistent with the NLTE helium model, which predicts a different time profile for the metastable $2s^3S$ population as the ejecta cools. The Sr II interpretation, by contrast, correctly predicts the emergence time, absorption wavelength, and temporal evolution in LTE models.

The resolution reconciles both studies: the Sr II identification is secure and robust; the absence of a He I feature is therefore a true constraint on the helium abundance (X_He < 0.05), and the helium spectroscopic clock of Bauswein et al. is valid precisely because the observed feature is definitively attributed to Sr II rather than to helium. These two contemporaneous papers are thus complementary rather than contradictory: one establishes what the 1 μm feature is (Sr II); the other establishes what it cannot be (He I at the level that would be produced by a long-lived remnant).

6. MAPPING THE QCD PHASE DIAGRAM WITH FINITE-TEMPERATURE MERGER OBSERVABLES

One of the most powerful conceptual contributions of this paper is the mapping of each observable discussed above to a specific region of the temperature-density (T–n_B) plane of the QCD phase diagram. This mapping provides observers and theorists with a precise statement of which aspect of the finite-temperature EoS each measurement constrains, enabling systematic planning of observational programs and theoretical calculations.

At baryon densities $n\_B \leq \rho_0$ and temperatures $T \leq 5$ MeV, the relevant physical processes are the neutrino-driven wind dynamics and the synthesis of helium and light r-process elements. The helium spectroscopic clock is sensitive to this regime: it constrains the neutrino luminosity integrated over the first τ_rem ≈ 20–30 ms, which in turn constrains the thermal pressure at $n\_B \approx 0.5\text{–}1.0\ \rho_0$ and $T \approx 5\text{–}10$ MeV in the neutrino-emitting surface layers of the remnant. This is the most accessible region of the finite-temperature QCD phase diagram because it overlaps with conditions reachable in terrestrial nuclear experiments (heavy-ion collisions at HADES/GSI energies probe $T \approx 50\text{–}80$ MeV, $n\_B \approx 2\text{–}3\ \rho_0$, but with isospin-symmetric conditions), and with the validity range of chiral effective field theory ($\rho \leq \rho_0$, $T \leq 30$ MeV).

At $n\_B \approx 1\text{–}3\ \rho_0$ and $T \approx 10\text{–}50$ MeV, the dominant processes are Urca equilibration and bulk viscous dissipation. The bulk viscosity phase shift in postmerger GW signals is most sensitive to the Urca threshold density n_DU, which is set by the proton fraction at $n\_B \approx 1.5\text{–}2.5\ \rho_0$ — directly controlled by the symmetry energy slope L. This regime is inaccessible to laboratory experiments (no existing facility reaches these combined densities and temperatures simultaneously in neutron-rich matter) but can be constrained by the combination of the symmetry energy slope from nuclear structure experiments and the bulk viscosity signature from ET/CE postmerger observations.

At n_B ≈ 2–5 ρ_0 and T ≈ 30–100 MeV, the conditions in the HMNS core approach the hadron–quark crossover boundary. The dominant f_2 postmerger GW frequency is most sensitive to this regime: it reflects the stiffness of matter at the central densities of the remnant, which are typically 2–3 ρ_0 for the dominant mass configurations at merger and up to 4–5 ρ_0 in the innermost core. The hot quark matter EoS in this regime, studied by Issifu et al. (European Physical Journal A, August 2025) using the density-dependent quark mass model (DDQM) extended to finite temperature through a lattice QCD-motivated approach, shows that at fixed entropy per baryon s/n_B, the merger remnant becomes more massive and larger as s/n_B increases, with the neutrino abundance also increasing strongly. This entropy-driven expansion of the remnant at fixed baryon mass is a direct signature of the quark matter EoS that differentiates it from hadronic EoS models and could in principle be identified through the GW frequency evolution during the first 10 ms post-merger.

The finite-temperature expansion framework of Mroczek et al. (Physical Review D, 2024) places all four observable regimes on a common theoretical footing. Their demonstration that the full P(ρ, T, Y_e) table can be reconstructed to within 5% from cold-matter constraints plus a small number of finite-temperature parameters (the effective mass m* and heat capacity at fixed density) provides the mathematical basis for a unified inference framework in which all four observable channels — helium spectroscopy, neutrino luminosity, bulk viscosity GW phase, and late-time inertial modes — simultaneously constrain the same underlying finite-temperature EoS table.

7. UNIFIED BAYESIAN INFERENCE FRAMEWORK FOR THE FINITE-TEMPERATURE EoS

The mapping of observables to the (ρ, T) plane established in the previous section suggests a natural Bayesian inference framework for simultaneously constraining the finite-temperature EoS from multiple observable channels. We outline this framework here, noting that its full implementation requires sophisticated multi-messenger analysis pipelines that represent a major theoretical and computational challenge for the community.

The fundamental object to be inferred is the three-dimensional EoS table P(ρ, T, Y_e), parametrised through the finite-temperature expansion as a cold EoS posterior (constrained by the cold-matter observables of Papers I–III: PREX-II, CREX, NICER, GW170817 tidal deformability) plus a finite-temperature extension parametrised by the effective mass function m*(ρ) and the dimensionless heat capacity ratio c_V/c_V^{ideal}. The likelihood function is the product of four independent likelihood terms:

$$L\_total = L\_He \times L\_\nu \times L\_\zeta \times L\_inert$$

where L_He is the likelihood of the helium abundance upper limit from AT2017gfo spectroscopy, L_ν is the likelihood of the neutrino luminosity time-profile from a hypothetical nearby merger detection, L_ζ is the likelihood of the bulk viscosity GW phase shift from ET/CE postmerger observations, and L_inert is the likelihood of the inertial mode frequency from late-time GW spectroscopy. Each likelihood term constrains a distinct sector of the finite-temperature EoS:

L_He constrains primarily the entropy density at n_B ≈ 0.5ρ_0, T ≈ 5–15 MeV, through the remnant lifetime. L_ν constrains the shock-heating efficiency and neutrino sphere temperature, sensitive to P(n_B, T) at intermediate densities and temperatures. L_ζ constrains the Urca threshold density n_DU and hence the symmetry energy through the bulk viscosity coefficient ζ(n_DU, T). L_inert constrains the Brunt–Väisälä frequency profile N(ρ, T) and hence the entropy gradient in the HMNS core.

The complementarity of these four constraints is maximum: they access different density-temperature regions and different microphysical processes, so that their intersection places constraints on the finite-temperature EoS that no single channel could achieve alone. The full implementation of this framework requires: (i) a suite of long-term (t > 100 ms) GRMHD merger simulations with fully tabulated finite-temperature EoS tables spanning the prior parameter space; (ii) coupled neutrino transport at the M1 or two-moment level with Doppler corrections at all orders in v/c; (iii) NLTE kilonova spectral synthesis with complete atomic data for helium and light r-process elements; and (iv) GW parameter estimation codes that include bulk viscosity phase evolution and inertial mode templates.

8. DISCUSSION: THE HOT QCD PHASE DIAGRAM AND TERRESTRIAL CONNECTIONS

The measurements and theoretical developments reviewed in this paper collectively define a new window into the QCD phase diagram at finite baryon density and temperature — a region of fundamental importance for particle physics and cosmology but notoriously inaccessible to first-principles QCD calculations due to the sign problem at finite baryon chemical potential. The conditions in the HMNS remnant ($T \approx 10$–100 MeV, $n\_B \approx 1$–$5\ \rho_0$, $Y\_e \approx 0.05$–0.4) span a trajectory in the QCD phase diagram that is qualitatively different from that accessed by heavy-ion collisions at RHIC or LHC ($T \approx 150$–500 MeV, $n\_B \approx 0$, symmetric matter) or by terrestrial neutron-rich matter experiments ($T \approx 0$, $n\_B \approx 0.5$–$2\ \rho_0$).

The HADES experiment at GSI and its successor at FAIR (CBM) are the closest terrestrial probes of the merger-relevant regime, accessing $T \approx 50$–80 MeV and $n\_B \approx 2$–$3\ \rho_0$ in heavy-ion collisions at beam energies of a few GeV per nucleon. The pion multiplicity and strangeness production ratios measured at HADES provide constraints on the thermal pressure at these conditions that complement the astrophysical constraints from merger remnants. A quantitative comparison between HADES measurements and the thermal EoS inferred from merger GW observations would, for the first time, test whether the nuclear EoS extrapolated from heavy-ion collisions in symmetric matter is consistent with the neutron-rich, neutrino-loaded conditions of the HMNS remnant.

Similarly, the MUSES (Modular Unified Solver of the Equation of State) collaboration (Kumar et al., Living Reviews in Relativity, 27, 3, 2024) has developed a comprehensive theoretical framework for constructing finite-temperature EoS models consistently constrained by lattice QCD at high temperature, perturbative QCD at high density, chiral EFT at low density, heavy-ion collision data at intermediate conditions, and astrophysical observations at ultra-high density. The MUSES framework explicitly addresses the thermodynamic consistency conditions — positivity of specific heat, Le Chatelier's principle, causality — that must be satisfied by any physical EoS table across the full (ρ, T) plane, and provides a publicly available EoS solver that takes nuclear matter observables as inputs and generates GRMHD-compatible tabulated EoS output. This tool, now validated against multiple existing EoS tables, will be essential for implementing the unified Bayesian inference framework outlined above.

9. CONCLUSIONS

This paper has synthesised four interconnected theoretical and observational frontiers in the physics of hot dense matter in binary neutron star merger remnants, providing for the first time a unified mapping from observables to their sensitivity domains in the finite-temperature QCD phase diagram. The principal conclusions of this work are as follows:

1. Fully tabulated finite-temperature EoS tables (SFHo, DD2, BLh) produce qualitatively distinct postmerger gravitational-wave signatures from M* parametric approximations at late times ($t > 50$ ms post-merger). Specifically, convective instabilities driven by Ledoux and Solberg–Høiland criteria in the thermally stratified HMNS core excite inertial oscillation modes at $f \approx 1$–2 kHz that are not captured by M* simulations. These modes are potentially detectable by the Einstein Telescope for merger events at distances ≤ 20 Mpc and constitute a new finite-temperature EoS diagnostic.

2. Weak-interaction-driven bulk viscosity from direct and modified Urca processes introduces a measurable phase shift in the postmerger gravitational-wave signal. The phase shift magnitude depends on the direct Urca threshold density n_DU, which is set by the symmetry energy slope L. Strong magnetic fields ($B \sim 10^{14}$–10^{15} G) modify Urca rates by 30–50% through Landau quantization of electron phase space, constituting a systematic uncertainty that must be incorporated in future EoS inference pipelines.

3. The Tayler–Spruit dynamo, operating on timescales of tens of milliseconds in the differentially-rotating HMNS core, amplifies the post-KHI magnetic field to $B \sim 10^{16}$ G in a predominantly toroidal configuration. This drives thermomagnetic instabilities and convective patterns that couple the magnetic and thermal structure of the remnant and sustain long-lived inertial modes. The TS dynamo is an intrinsically finite-temperature effect, operating specifically in the stably-stratified, thermally-structured remnant.

4. The helium spectroscopic clock, derived from the non-detection of He I λ1083.3 nm in AT2017gfo at 4.4 days post-merger, constrains the polar ejecta helium mass fraction to $X\_He < 0.05$ and hence the remnant lifetime to $\tau\_rem \leq 20$–

30 ms. Combined with EoS-dependent collapse threshold relations and causality arguments, this limits M_max ≤ 2.3 M⊙ and constrains R(1.6 M⊙) ≈ 12 ± 1 km. The simultaneous determination by Sneppen et al. (2024) that the 1 μm feature is definitively attributed to Sr II rather than He I validates and strengthens this constraint.

5. Neutrino emission from the HMNS remnant encodes the tidal deformability $\tilde{\Lambda}$ of the binary through its control over shock-heating intensity: stiffer EoS systems produce systematically lower neutrino luminosities. A future neutrino burst detection from a nearby BNS merger with Hyper-Kamiokande or DUNE would provide an independent EoS constraint from a third particle species — complementing gravitational-wave and photon observations.

6. A unified Bayesian inference framework combining the helium spectroscopic constraint, neutrino luminosity, bulk viscosity GW phase, and inertial mode frequency simultaneously constrains the full three-dimensional P(ρ, T, Y_e) EoS table. The four constraints are maximally complementary, accessing different density-temperature sectors and different microphysical processes. Implementation of this framework requires long-term GRMHD simulations with tabulated EoS, coupled M1 neutrino transport with full Doppler corrections, and NLTE kilonova spectral synthesis — representing the frontier of computational nuclear astrophysics.

This four-paper series has traced the problem of dense nuclear matter from its most precisely measured properties near saturation through the exotic-matter regime of neutron star cores, the rich multimessenger aftermath of binary mergers, and now into the finite-temperature interior of the hypermassive remnant. The common thread is the nuclear force — specifically, how its isospin-asymmetric component organises matter under pressures and temperatures that no laboratory can reproduce. We are at the extraordinary juncture where spectroscopic observations of stellar remnants 40 megaparsecs away constrain the thermal pressure of nuclear matter at temperatures of 10 million degrees, where gravitational waves carry the phase imprint of microsecond weak-interaction processes, and where the absence of a spectral line in a dying star's optical glow limits the maximum mass of all neutron stars in the universe. The next decade, with its third-generation gravitational-wave detectors, space telescopes, and refined nuclear theory, will complete this program.

ACKNOWLEDGEMENT

The author gratefully acknowledges the support of Government Degree College Batote, District Ramban, Jammu, India. The author thanks the international communities of nuclear astrophysics, gravitational-wave astronomy, neutrino physics, and kilonova spectroscopy for the foundational works that make this synthesis possible. Special acknowledgement is due to the MUSES Collaboration, the LIGO-Virgo-KAGRA Collaboration, the NICER Science Team, and the VLT X-shooter spectroscopy team whose data and theoretical frameworks underpin every section of this paper.

REFERENCES

1. Giacomazzo, B. et al.: On the Treatment of Thermal Effects in the Equation of State on Neutron Star Merger Remnants. arXiv:2512.05118 (December 2024 / 2025).

2. Mroczek, D., Yao, N., Zine, K., Noronha-Hostler, J., Dexheimer, V., Haber, A., Most, E.: Finite-Temperature Expansion of the Dense-Matter Equation of State. Physical Review D (2024). arXiv:2404.01658.

3. Ujevic, M., Somasundaram, R., Dietrich, T., Margueron, J., Tews, I.: Validity of a Finite-Temperature Framework for Neutron Star Mergers. Astrophysical Journal Letters (February 2024).

4. Cusinato, M., Guercilena, F.M., Perego, A. et al.: Neutrino Emission from Binary Neutron Star Mergers: Characterising Light Curves and Mean Energies. European Physical Journal A, 58, 99 (2022).

5. Most, E.R., Raithel, C.A.: Emergence of Microphysical Bulk Viscosity in Binary Neutron Star Post-Merger Dynamics. arXiv:2207.00442; Physical Review D (2024).

6. Tambe, P., Chatterjee, D., Alford, M., Haber, A.: Effect of Magnetic Fields on Urca Rates in Neutron Star Mergers. Physical Review C, 111, 035809 (2025). arXiv:2409.09423.

7. Alford, M.G., Haber, A., Zhang, Z.: Beyond Modified Urca: The Nucleon Width Approximation for Flavor-Changing Processes in Dense Matter. Physical Review C, 110, L052801 (2024).

8. Rau, P.B., Schwenk, A.: Composition g-Modes and Bulk Viscous Damping in Warm Neutron Stars. Physical Review D (2023–2024).

9. Barrère, P. et al.: Tayler–Spruit Dynamo in Binary Neutron Star Merger Remnants. HAL:hal-04843892 (2024–2025); Astronomy & Astrophysics (2025).

10. Kiuchi, K. et al.: High-Resolution GRMHD Simulations of Binary Neutron Star Mergers. Physical Review D (2024).

11. Bauswein, A. et al.: Helium as an Indicator of the Neutron-Star Merger Remnant Lifetime and its Potential for Equation of State Constraints. Physical Review D (2025). arXiv:2411.03427.

12. Sneppen, A., Damgaard, R., Watson, D., Collins, C.E., Shingles, L., Sim, S.A.: Helium Features are Inconsistent with the Spectral Evolution of the Kilonova AT2017gfo. Astronomy & Astrophysics, 692, A134 (2024).

13. Issifu, A., da Silva, F.M., Santos, L.C.N., Menezes, D.P., Frederico, T.: Hot Quark Matter and Merger Remnants. European Physical Journal A (August 2025). DOI: 10.1140/epja/s10050-025-01654-x.

14. Kumar, R., Dexheimer, V., Jahan, J., Noronha, J., Noronha-Hostler, C. et al. (MUSES Collaboration): Theoretical and Experimental Constraints for the Equation of State of Dense and Hot Matter. Living Reviews in Relativity, 27, 3 (2024).

15. Breschi, M. et al.: Constraining the Equation of State in Neutron-Star Cores via the Long-Ringdown Signal. Nature Communications, 16, 1136 (2025).

16. Watson, D. et al.: Identification of Strontium in the Merger of Two Neutron Stars. Nature, 574, 497–500 (2019).

17. Abbott, R. et al.: QCD Constraints on Isospin-Dense Matter and the Nuclear Equation of State. Physical Review Letters, 134, 011903 (2025).

18. Choudhury, D. et al.: A NICER View of PSR J0437−4715. Astrophysical Journal Letters, 971, L20 (2024).

19. Sorensen, A. et al.: Dense Nuclear Matter Equation of State from Heavy-Ion Collisions. Progress in Particle and Nuclear Physics, 134, 104080 (2024).

20. Reed, B.T., Fattoyev, F.J., Horowitz, C.J., Piekarewicz, J.: Density Dependence of the Symmetry Energy in the Post-PREX-CREX Era. Physical Review C, 109, 035803 (2024).

BIOGRAPHY

[Photo]	Dr. Arti Pandoh Gupta is the Incharge Principal at Government Degree College Batote, District Ramban, Jammu, India. She holds a Doctorate in Theoretical Nuclear Physics from the University of Jammu. She has published four papers in peer-reviewed international journals and has presented her research at national and international conferences. Her research interests span nuclear structure theory, nuclear symmetry energy, finite-temperature dense matter, neutron star merger physics, kilonova spectroscopy, and multimessenger nuclear astrophysics. This paper is the fourth in her series on dense nuclear matter connecting laboratory nuclear physics to compact star astrophysics and gravitational-wave science.