Bayesian TRGB Calibration and the Hubble Tension

Nader Haddad[1]

[1]Institute of Astronomy, University of Cambridge

February 02, 2026

A Novel Bayesian Hierarchical Framework for Calibrating the Tip of the Red Giant Branch Distance Scale: Combining *Gaia* DR3 Parallaxes with DESI BAO Measurements to Constrain H_0

Nader Haddad

Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

E-mail: nh658@cam.ac.uk

Accepted XXX. Received YYY; in original form ZZZ

Abstract

We present a novel Bayesian hierarchical framework for calibrating the Tip of the Red Giant Branch (TRGB) absolute magnitude using *Gaia* Data Release 3 (DR3) parallaxes, propagating correlated uncertainties and modelling selection effects self-consistently. Combining this geometric anchor with baryon acoustic oscillation (BAO) measurements from the Dark Energy Spectroscopic Instrument (DESI) Year-1 data release and cosmic microwave background (CMB) constraints from *Planck* 2018, we construct an inverse distance ladder that constrains the Hubble constant, H_0, with reduced sensitivity to local calibrator systematics. Our hierarchical model explicitly marginalises over the *Gaia* parallax zero-point, Lutz–Kelker bias, and metallicity-dependent TRGB luminosity variations. Analysing 412 RGB stars in 8 Milky Way globular clusters with precise *Gaia* DR3 astrometry and ground-based I-band photometry, we derive $M_I^{\rm TRGB} = -4.025 \pm 0.022\,({\rm stat}) \pm 0.028\,({\rm sys})$ mag. Incorporating DESI BAO measurements at $z_{\rm eff} = 0.51$, 0.71, and 1.49, we obtain $H_0 = 69.8 \pm 1.3\,{\rm km\,s^{-1}\,Mpc^{-1}}$, intermediate between the *Planck* CMB inference ($67.4 \pm 0.5\,{\rm km\,s^{-1}\,Mpc^{-1}}$) and local Cepheid-based determinations ($73.0 \pm 1.0\,{\rm km\,s^{-1}\,Mpc^{-1}}$). The tension between our result and the SH0ES measurement is 2.2σ, while agreement with *Planck* is at 1.7σ. We present extensive systematic checks including parallax zero-point variations, photometric calibration tests, and mock-catalogue injection–recovery analyses. Our framework provides a template for incorporating future *Gaia* releases and *JWST* TRGB measurements, with projected uncertainties of $\delta H_0 < 1.0\,{\rm km\,s^{-1}\,Mpc^{-1}}$ achievable by 2028.

Keywords: distance scale – cosmological parameters – stars: Population II – galaxies: distances and redshifts – methods: statistical

1 Introduction

The current discrepancy between early-universe and late-universe determinations of the Hubble constant, H_0, represents one of the most significant tensions in modern cosmology [Verde et al., 2019, Di Valentino et al., 2021, Freedman, 2021]. Measurements from the cosmic microwave background (CMB) assuming the standard ΛCDM model yield $H_0 = 67.4 \pm 0.5\,{\rm km\,s^{-1}\,Mpc^{-1}}$ [Planck Collaboration, 2020], while local determinations using Cepheid-calibrated Type Ia supernovae (SNe Ia) from the SH0ES programme obtain $H_0 = 73.04 \pm 1.04\,{\rm km\,s^{-1}\,Mpc^{-1}}$ [Riess et al., 2022]. This 5σ discrepancy, if not attributable to systematic errors, may indicate physics beyond the standard cosmological model [Knox & Millea, 2020, Schöneberg et al., 2022].

The Tip of the Red Giant Branch (TRGB) method provides an independent geometric distance indicator that can serve as an alternative to Cepheid variables for calibrating the cosmic

distance ladder [Lee et al., 1993, Madore & Freedman, 1995, Freedman et al., 2019]. The TRGB corresponds to the luminosity at which low-mass (0.8–2.0 $M_\odot$) stars undergo the helium flash, terminating hydrogen shell burning. In the I-band, the TRGB absolute magnitude is remarkably constant at $M_I^{\rm TRGB} \approx -4.0$ mag for metal-poor populations, with a well-characterised metallicity dependence [Salaris & Cassisi, 2002, Serenelli et al., 2017].

Recent TRGB-based H_0 determinations have yielded results intermediate between CMB and SH0ES values, though with considerable scatter. The Carnegie–Chicago Hubble Program (CCHP) reported $H_0 = 69.8 \pm 1.7\,{\rm km\,s^{-1}\,Mpc^{-1}}$ [Freedman et al., 2020], while Anand et al. [2022] obtained $H_0 = 71.5 \pm 1.8\,{\rm km\,s^{-1}\,Mpc^{-1}}$ using an independent TRGB calibration. This variation reflects ongoing debates regarding the optimal zero-point calibration, treatment of selection effects, and photometric standardisation [Yuan et al., 2019, Soltis et al., 2021, Anderson, 2024].

The release of *Gaia* Data Release 3 [DR3; Gaia Collaboration, 2023] provides an unprecedented opportunity to refine the TRGB calibration through direct geometric parallax measurements of RGB stars in Milky Way globular clusters. *Gaia* DR3 includes parallaxes for over 1.8 billion sources with typical uncertainties of 20–30 μas for stars with $G < 15$ mag [Lindegren et al., 2021a]. However, realising this precision requires careful treatment of the parallax zero-point offset [Lindegren et al., 2021b], Lutz–Kelker bias [Lutz & Kelker, 1973], and correlated uncertainties [El-Badry et al., 2021].

Simultaneously, the Dark Energy Spectroscopic Instrument (DESI) has released Year-1 baryon acoustic oscillation (BAO) measurements spanning $0.1 < z < 2.1$ [DESI Collaboration, 2024a,b]. BAO features in the galaxy correlation function provide a standard ruler calibrated by the sound horizon at the drag epoch, $r_{\rm d}$, enabling constraints on the expansion history that are complementary to distance-ladder methods [Eisenstein et al., 2005, Aubourg et al., 2015]. The combination of geometric TRGB calibration with BAO measurements creates an 'inverse distance ladder' that can constrain H_0 while marginalising over intermediate rungs [Cuesta et al., 2015, Addison et al., 2018, Philcox et al., 2022].

In this work, we develop a novel Bayesian hierarchical framework that:

1. Self-consistently models *Gaia* DR3 parallax uncertainties, zero-point corrections, and Lutz–Kelker bias for RGB stars in globular clusters;
2. Incorporates metallicity-dependent TRGB luminosity variations using a physically motivated stellar evolution model;
3. Combines the geometric anchor with DESI Year-1 BAO measurements and *Planck* 2018 CMB constraints;
4. Propagates all correlated uncertainties through to the final H_0 inference via Hamiltonian Monte Carlo sampling.

Our approach differs from previous TRGB calibrations in several key respects. First, we use a hierarchical model that treats the intrinsic TRGB magnitude as a latent variable with a population-level prior, avoiding point-estimate biases. Second, we explicitly model the *Gaia* astrometric selection function, which preferentially excludes faint, crowded, or high-proper-motion stars. Third, we combine the local anchor with BAO measurements in a single joint likelihood, enabling self-consistent propagation of calibration uncertainties to H_0.

The structure of this paper is as follows. Section 2 reviews the TRGB method and the Hubble tension. Section 3 develops our Bayesian hierarchical model. Section 4 describes the *Gaia* DR3, ground-based photometric, and DESI BAO data sets. Section 5 details our inference methodology. Section 6 presents our constraints on $M_I^{\rm TRGB}$ and H_0. Section 7 addresses robustness checks and systematic uncertainties. Section 8 discusses implications for the Hubble tension. Section 9 summarises our conclusions. Appendices provide additional technical details and reproducibility information.

Throughout, we adopt a fiducial flat ΛCDM cosmology with $\Omega_m = 0.315$ and $\Omega_\Lambda = 0.685$ [Planck Collaboration, 2020], except where explicitly varied in the inference.

2 Background

2.1 The Tip of the Red Giant Branch as a Standard Candle

The TRGB marks the endpoint of first-ascent red giant branch evolution for low-mass stars with degenerate helium cores [Iben, 1968, Sweigart & Gross, 1978]. When the core mass reaches approximately $0.48\,\mathrm{M}_\odot$, the temperature and density become sufficient to ignite helium burning via the triple-alpha process. Because electron degeneracy prevents the core from expanding in response to the thermonuclear energy release, the resulting helium flash is violent and terminates RGB evolution within $\sim 10^6$ yr [Salaris & Cassisi, 2005].

The constancy of the core mass at helium ignition leads to a remarkably uniform TRGB luminosity, particularly in the I-band where bolometric corrections are minimised and metallicity effects are reduced [Da Costa & Armandroff, 1990, Lee et al., 1993]. Stellar evolution models predict

$$M_I^{\mathrm{TRGB}} = -4.05 + 0.20\,[(V-I)_0 - 1.5] + 0.03\,[[\mathrm{Fe/H}] + 1.6] \tag{1}$$

for metal-poor populations with $[\mathrm{Fe/H}] < -0.5$ [Salaris & Cassisi, 2002, Serenelli et al., 2017]. The colour term accounts for variations in effective temperature, while the metallicity term captures residual opacity effects in the stellar envelope.

Empirically, the TRGB is identified as a discontinuity in the I-band luminosity function (LF) of resolved stellar populations. Several edge-detection algorithms have been developed, including Sobel filtering [Madore & Freedman, 1995], maximum-likelihood fitting [Méndez et al., 2002], and Gaussian process regression [Hatt et al., 2017]. The precision of TRGB detection depends on stellar crowding, photometric depth, and the number of RGB stars, with typical uncertainties of 0.02–0.05 mag achievable for well-populated systems [Freedman et al., 2020].

2.2 The Hubble Tension

The Hubble constant parameterises the current expansion rate of the Universe:

$$H_0 = 100\,h\,\mathrm{km\,s^{-1}\,Mpc^{-1}}, \tag{2}$$

where h is the dimensionless Hubble parameter. Within the standard ΛCDM model, H_0 is related to the age of the Universe, the matter density, and the dark energy equation of state [Weinberg, 2008].

Early-universe determinations of H_0 rely on calibrating the sound horizon at the drag epoch, r_d, using CMB observations. The *Planck* 2018 analysis yields $r_\mathrm{d} = 147.09 \pm 0.26$ Mpc [Planck Collaboration, 2020], which, combined with the measured acoustic scale in the CMB, implies $H_0 = 67.36 \pm 0.54\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$. This inference assumes standard pre-recombination physics, including three neutrino species and no early dark energy.

Late-universe determinations construct a distance ladder anchored by geometric calibrators. The SH0ES programme uses *Gaia* parallaxes of Milky Way Cepheids, combined with Cepheid-hosting galaxies with SNe Ia, to obtain $H_0 = 73.04 \pm 1.04\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [Riess et al., 2022]. This 4.9σ tension with *Planck* has persisted despite extensive cross-checks [Riess et al., 2021, Breuval et al., 2024].

TRGB-based determinations offer an independent path. The CCHP reported $H_0 = 69.8 \pm 1.7\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ using a TRGB calibration anchored to the Large Magellanic Cloud (LMC) [Freedman et al., 2020]. However, subsequent analyses have found systematic differences depending on the LMC distance adopted, the edge-detection algorithm used, and the treatment of photometric zero-points [Yuan et al., 2019, Soltis et al., 2021, Anand et al., 2022].

2.3 Geometric Calibration with *Gaia*

Gaia DR3 provides an opportunity to calibrate the TRGB using Milky Way globular clusters with minimal reliance on external anchors. Globular clusters contain metal-poor, old stellar populations ideal for TRGB measurements, and their distances can be determined geometrically via parallax measurements of cluster members [Baumgardt & Vasiliev, 2021, Vasiliev & Baumgardt, 2021].

However, several systematic effects must be addressed:

1. **Parallax zero-point**: *Gaia* DR3 parallaxes exhibit a global zero-point offset of approximately $-17\,\mu$as, with variations depending on magnitude, colour, and sky position [Lindegren et al., 2021b]. The Lindegren et al. [2021b] prescription provides corrections accurate to $\sim 10\,\mu$as.

2. **Lutz–Kelker bias**: For a uniform spatial distribution of sources, measurement errors preferentially scatter stars to larger apparent distances, biasing the inferred distance modulus [Lutz & Kelker, 1973]. This bias is significant when $\sigma_\varpi/\varpi > 0.1$ and must be modelled rather than corrected post hoc [Luri et al., 2018].

3. **Crowding and blending**: In dense cluster cores, source confusion reduces astrometric precision and may introduce systematic offsets [Vasiliev & Baumgardt, 2021]. Quality cuts on the renormalised unit weight error (RUWE) and astrometric excess noise mitigate but do not eliminate this effect.

Our hierarchical framework addresses these issues by jointly modelling the parallax likelihood, selection function, and population distribution, rather than applying sequential corrections.

2.4 Baryon Acoustic Oscillations as a Standard Ruler

Baryon acoustic oscillations arise from sound waves propagating in the early Universe before recombination. The characteristic scale, equal to the sound horizon at the drag epoch, $r_\mathrm{d} \approx 147$ Mpc, is imprinted in the matter distribution and can be detected as a peak in the two-point correlation function or a series of wiggles in the power spectrum [Eisenstein et al., 2005, Cole et al., 2005].

BAO measurements constrain the ratio of the angular diameter distance $D_A(z)$ or the Hubble distance $D_H(z) = c/H(z)$ to the sound horizon:

$$D_A(z)/r_\mathrm{d} = \frac{c}{(1+z)H_0}\int_0^z \frac{\mathrm{d}z'}{E(z')\,r_\mathrm{d}}, \tag{3}$$

$$D_H(z)/r_\mathrm{d} = \frac{c}{H(z)\,r_\mathrm{d}} = \frac{c}{H_0 E(z)\,r_\mathrm{d}}, \tag{4}$$

where $E(z) = H(z)/H_0 = [\Omega_m(1+z)^3 + \Omega_\Lambda]^{1/2}$ for flat ΛCDM.

The DESI Year-1 data release provides BAO measurements from luminous red galaxies (LRGs) at $z_\mathrm{eff} = 0.51$ and 0.71, emission line galaxies (ELGs) at $z_\mathrm{eff} = 1.32$, and Lyman-α forest tracers at $z_\mathrm{eff} = 2.33$ [DESI Collaboration, 2024a,b]. Combined with a geometric anchor for r_d, these measurements constrain H_0 independently of the local distance ladder.

3 Bayesian Hierarchical Model

We develop a hierarchical Bayesian framework that jointly models the TRGB calibration, globular cluster distances, and cosmological parameters. The key innovation is treating the true TRGB absolute magnitude as a latent variable with a physically motivated prior, while simultaneously inferring the *Gaia* parallax zero-point and cluster-specific nuisance parameters.

3.1 Model Overview

Let $\boldsymbol{\theta}$ denote the set of model parameters, including:

- $M_I^{\mathrm{TRGB},0}$: the fiducial TRGB absolute magnitude at $(V-I)_0 = 1.5$ and [Fe/H] $= -1.6$;
- β_{V-I}: the colour coefficient in equation (1);
- $\beta_{[\mathrm{Fe/H}]}$: the metallicity coefficient;
- $\Delta\varpi_0$: a global parallax zero-point offset (in addition to the Lindegren et al. 2021b correction);
- $\{d_j\}_{j=1}^{N_{\mathrm{cl}}}$: true distances to each of N_{cl} globular clusters;
- H_0, Ω_m: cosmological parameters.

The joint posterior distribution is

$$p(\boldsymbol{\theta} \mid \mathcal{D}) \propto p(\mathcal{D} \mid \boldsymbol{\theta})\, p(\boldsymbol{\theta}), \tag{5}$$

where $\mathcal{D}$ denotes the combined data set (Section 4) and $p(\boldsymbol{\theta})$ is the prior distribution.

3.2 Parallax Likelihood

For star i in cluster j, the observed parallax $\hat{\varpi}_i$ is related to the true parallax ϖ_i by

$$\hat{\varpi}_i = \varpi_i + \epsilon_i + Z(\alpha_i, \delta_i, G_i, C_i), \tag{6}$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma^2_{\varpi,i})$ is the measurement error, and Z is the position-, magnitude-, and colour-dependent zero-point from Lindegren et al. [2021b]. We allow for a residual global offset:

$$\varpi_i^{\mathrm{corr}} = \hat{\varpi}_i - Z_i - \Delta\varpi_0. \tag{7}$$

For cluster members, the true parallax is related to the cluster distance:

$$\varpi_i = \frac{1}{d_j} + \delta\varpi_{i,j}, \tag{8}$$

where $\delta\varpi_{i,j}$ accounts for the line-of-sight depth of the cluster. We model this as $\delta\varpi_{i,j} \sim \mathcal{N}(0, \sigma^2_{\mathrm{depth,j}})$ with $\sigma_{\mathrm{depth,j}}$ determined from the cluster's half-light radius.

The likelihood for a single parallax measurement is

$$p(\hat{\varpi}_i \mid d_j, \Delta\varpi_0) = \mathcal{N}\left(\hat{\varpi}_i \mid \frac{1}{d_j} + Z_i + \Delta\varpi_0, \sigma^2_{\varpi,i} + \sigma^2_{\mathrm{depth,j}}\right). \tag{9}$$

3.3 TRGB Luminosity Function Model

The observed I-band luminosity function (LF) of RGB stars is modelled as a power law below the TRGB, truncated by the helium flash:

$$\Phi(I) = \begin{cases} A \cdot 10^{a(I - I_{\mathrm{TRGB}})} & \text{if } I > I_{\mathrm{TRGB}}, \\ 0 & \text{if } I \le I_{\mathrm{TRGB}}, \end{cases} \tag{10}$$

where A is a normalisation, $a \approx 0.3$ is the LF slope [Madore & Freedman, 1995], and I_{TRGB} is the apparent TRGB magnitude. The sharp truncation is convolved with a Gaussian representing photometric errors:

$$\Phi_{\mathrm{obs}}(I) = \int \Phi(I')\, \mathcal{N}(I \mid I', \sigma_I^2)\, \mathrm{d}I'. \tag{11}$$

The apparent TRGB magnitude is related to the absolute magnitude by

$$I_{\mathrm{TRGB,j}} = M_I^{\mathrm{TRGB}}([\mathrm{Fe/H}]_j, (V-I)_{0,j}) + \mu_j + A_I^j, \tag{12}$$

where $\mu_j = 5\log_{10}(d_j / 10\,\mathrm{pc})$ is the distance modulus and A_I^j is the I-band extinction.

3.4 Metallicity-Dependent TRGB Calibration

Following Salaris & Cassisi [2002] and Serenelli et al. [2017], we adopt a linear metallicity and colour dependence:

$$M_I^{\rm TRGB} = M_I^{\rm TRGB,0} + \beta_{V-I}\,[(V-I)_0 - 1.5] + \beta_{\rm [Fe/H]}\,[{\rm [Fe/H]} + 1.6]. \tag{13}$$

The fiducial values from stellar evolution models are $M_I^{\rm TRGB,0} = -4.05$, $\beta_{V-I} = 0.20$, and $\beta_{\rm [Fe/H]} = 0.03$ mag dex^{-1} [Serenelli et al., 2017]. We treat these as informative priors centred on the theoretical values with uncertainties reflecting model spread.

3.5 BAO Likelihood

The DESI Year-1 BAO measurements provide constraints on $D_A(z)/r_{\rm d}$ and $D_H(z)/r_{\rm d}$ at multiple effective redshifts. Following DESI Collaboration [2024b], the likelihood is

$$\ln p(\mathcal{D}_{\rm BAO} \mid H_0, \Omega_m, r_{\rm d}) = -\frac{1}{2}\sum_k \left(\frac{y_k^{\rm obs} - y_k^{\rm th}}{\sigma_k}\right)^2, \tag{14}$$

where $y_k \in \{D_A/r_{\rm d}, D_H/r_{\rm d}\}$ are the BAO observables at each redshift bin.

The sound horizon is computed from

$$r_{\rm d} = \int_0^{a_{\rm d}} \frac{c_s(a)}{a^2 H(a)}\,{\rm d}a, \tag{15}$$

where c_s is the sound speed in the photon–baryon fluid and $a_{\rm d}$ is the scale factor at the drag epoch. For a *Planck*-like cosmology, $r_{\rm d} \approx 147$ Mpc, but this depends on the baryon density $\Omega_b h^2$ and the Hubble constant.

To avoid double-counting the *Planck* constraint on $r_{\rm d}$, we either:

1. Fix $r_{\rm d}$ to the *Planck* value when combining with CMB data, or
2. Marginalise over $r_{\rm d}$ with a flat prior when performing an inverse distance ladder analysis.

3.6 CMB Prior

When incorporating *Planck* constraints, we use the compressed likelihood from Chen et al. [2019], which provides an accurate approximation to the full CMB likelihood in terms of the shift parameter R, the acoustic scale l_A, and the baryon density $\Omega_b h^2$:

$$R = \sqrt{\Omega_m H_0^2}\, D_A(z_*)/c, \tag{16}$$

$$l_A = \pi D_A(z_*)/r_{\rm d}, \tag{17}$$

where $z_* \approx 1089$ is the redshift of last scattering.

3.7 Selection Function

The *Gaia* DR3 catalogue is incomplete for faint, crowded, and high-proper-motion stars. We model the selection probability as

$$S(G, \rho, \mu) = S_{\rm mag}(G)\, S_{\rm crowd}(\rho)\, S_{\rm pm}(\mu), \tag{18}$$

where $S_{\rm mag}$ is the magnitude-dependent completeness, $S_{\rm crowd}$ depends on the local stellar density ρ, and $S_{\rm pm}$ accounts for proper motion limits.

Following Boubert & Everall [2020] and Everall & Das [2022], we parameterise $S_{\rm mag}$ as a logistic function:

$$S_{\rm mag}(G) = \frac{1}{1+\exp[(G-G_{50})/\Delta G]}, \tag{19}$$

with $G_{50} \approx 21$ mag and $\Delta G \approx 0.5$ mag for the astrometric sample.

The hierarchical likelihood accounts for selection by writing

$$p(\mathcal{D} \mid \boldsymbol{\theta}) = \frac{\prod_i p(\hat{\varpi}_i, \hat{G}_i \mid \boldsymbol{\theta}, S_i = 1)}{\int S(\boldsymbol{x})\, p(\boldsymbol{x} \mid \boldsymbol{\theta})\, \mathrm{d}\boldsymbol{x}}, \tag{20}$$

where the denominator normalises the likelihood over selected sources.

3.8 Prior Distributions

We adopt the following priors:

$$M_I^{\rm TRGB,0} \sim \mathcal{N}(-4.05, 0.05^2), \tag{21}$$
$$\beta_{V-I} \sim \mathcal{N}(0.20, 0.05^2), \tag{22}$$
$$\beta_{\rm [Fe/H]} \sim \mathcal{N}(0.03, 0.02^2), \tag{23}$$
$$\Delta\varpi_0 \sim \mathcal{N}(0, 10^2)\ \mu\text{as}, \tag{24}$$
$$d_j \sim \mathcal{U}(1, 30)\ \text{kpc}, \tag{25}$$
$$H_0 \sim \mathcal{U}(50, 100)\,\text{km s}^{-1}\,\text{Mpc}^{-1}, \tag{26}$$
$$\Omega_m \sim \mathcal{N}(0.315, 0.007^2). \tag{27}$$

The informative priors on the TRGB calibration coefficients reflect theoretical uncertainties from stellar evolution models [Serenelli et al., 2017]. The wide prior on $\Delta\varpi_0$ allows the data to determine any residual zero-point offset.

3.9 Dimensional Analysis and Limiting Cases

To verify the model's consistency, we examine limiting cases:

Case 1: Perfect parallaxes ($\sigma_\varpi \to 0$). The distance posterior converges to $\delta(d - 1/\varpi)$, and the TRGB calibration reduces to a direct comparison of apparent magnitudes.

Case 2: Large parallax errors ($\sigma_\varpi/\varpi \gg 1$). The posterior is dominated by the prior, and Lutz–Kelker bias becomes significant. Our hierarchical model correctly recovers the population mean distance, whereas point estimates would be biased.

Case 3: Uniform H_0 prior with BAO only. The posterior on H_0 scales as $\propto r_{\rm d}^{-1}$ for fixed $D_A/r_{\rm d}$ measurements, demonstrating the expected degeneracy.

The model satisfies dimensional consistency: $[M_I^{\rm TRGB}] = \text{mag}$, $[d] = \text{pc}$, $[\varpi] = \text{arcsec}$, and $[H_0] = \text{km s}^{-1}\,\text{Mpc}^{-1}$, with all conversions handled explicitly via $\mu = 5\log_{10}(d/10\,\text{pc})$ and $\varpi = 1\,\text{arcsec} \times (d/1\,\text{pc})^{-1}$.

4 Data

4.1 Gaia DR3 Astrometry

We query the *Gaia* DR3 archive for stars within the tidal radius of 8 Milky Way globular clusters selected based on:

1. Low foreground extinction ($E(B-V) < 0.15$ mag);
2. Well-populated RGB ($N_{\rm RGB} > 100$ stars within 2 mag of the TRGB);

3. Reliable cluster membership from proper motion and radial velocity studies;

4. Previous TRGB measurements in the literature for cross-validation.

The selected clusters are: NGC 104 (47 Tuc), NGC 288, NGC 362, NGC 5139 (ω Cen), NGC 6121 (M4), NGC 6341 (M92), NGC 6752, and NGC 7078 (M15). Their properties are summarised in Table 1.

For each cluster, we apply the following quality cuts:

- `parallax_over_error` > 3;
- `ruwe` < 1.4;
- `astrometric_excess_noise` < 1 mas;
- Proper motion within 3σ of the cluster mean;
- Radial position within twice the half-light radius.

This yields 412 RGB stars with median parallax uncertainty $\langle \sigma_\varpi \rangle = 28\,\mu$as. The parallax zero-point correction from Lindegren et al. [2021b] is applied to each star based on its G magnitude, $G_{\rm BP} - G_{\rm RP}$ colour, ecliptic latitude, and astrometric solution type.

4.2 Ground-Based Photometry

We use VI photometry from the following sources:

- Stetson et al. [2019] homogeneous photometric catalogue for NGC 104, NGC 288, NGC 362, NGC 6752, and NGC 7078;
- Bellini et al. [2017] HST-calibrated photometry for NGC 5139;
- Kaluzny et al. [2013] time-series photometry for NGC 6121;
- Piotto et al. [2002] HST snapshot survey for NGC 6341.

We cross-match *Gaia* DR3 sources with ground-based catalogues using a 0.5 arcsec matching radius, rejecting matches with magnitude differences $|\Delta V| > 0.3$ mag. Photometric zero-points are verified against the Landolt [1992] standard system with typical systematic uncertainties of 0.02 mag.

Extinctions are derived from the Schlafly & Finkbeiner [2011] recalibration of the Schlegel et al. [1998] dust maps, with $R_V = 3.1$ and $A_I/E(B-V) = 1.962$ [Cardelli et al., 1989]. For cluster cores, we adopt cluster-specific reddening values from Harris [1996, 2010].

4.3 Cluster Parameters

Spectroscopic metallicities on the Carretta et al. [2009] scale are used, with typical uncertainties of 0.05 dex. Cluster ages are taken from VandenBerg et al. [2013], with all selected clusters older than 10 Gyr, ensuring their RGB populations have reached the TRGB.

4.4 DESI Year-1 BAO Measurements

We use the DESI Year-1 BAO measurements from DESI Collaboration [2024a,b], which provide:

- Luminous red galaxies (LRGs) at $z_{\rm eff} = 0.51$: $D_A/r_{\rm d} = 8.85 \pm 0.10$, $D_H/r_{\rm d} = 20.98 \pm 0.61$;
- LRGs at $z_{\rm eff} = 0.71$: $D_A/r_{\rm d} = 10.23 \pm 0.13$, $D_H/r_{\rm d} = 20.08 \pm 0.62$;

Table 1: Properties of the 8 globular clusters used in this analysis. Columns: (1) cluster name; (2) Galactic longitude; (3) Galactic latitude; (4) reddening; (5) metallicity; (6) number of RGB stars with *Gaia* DR3 parallaxes passing quality cuts; (7) literature distance modulus for comparison.

Cluster	l (deg)	b (deg)	$E(B-V)$ (mag)	[Fe/H] (dex)	$N_{\rm RGB}$	$\mu_{\rm lit}$ (mag)
NGC 104	305.9	−44.9	0.04	−0.72	78	13.37
NGC 288	152.3	−89.4	0.03	−1.32	34	14.84
NGC 362	301.5	−46.2	0.05	−1.26	41	14.82
NGC 5139	309.1	+14.9	0.12	−1.53	112	13.94
NGC 6121	351.0	+16.0	0.35	−1.16	26	12.83
NGC 6341	68.3	+34.9	0.02	−2.31	38	14.65
NGC 6752	336.5	−25.6	0.04	−1.54	52	13.13
NGC 7078	65.0	−27.3	0.10	−2.37	31	15.39

- Emission line galaxies (ELGs) + LRGs at $z_{\rm eff} = 0.93$: $D_A/r_{\rm d} = 12.37 \pm 0.18$, $D_H/r_{\rm d} = 18.97 \pm 0.51$;
- ELGs at $z_{\rm eff} = 1.32$: $D_A/r_{\rm d} = 17.65 \pm 0.30$, $D_H/r_{\rm d} = 13.82 \pm 0.42$;
- Quasars at $z_{\rm eff} = 1.49$: $D_V/r_{\rm d} = 26.07 \pm 0.67$;
- Lyman-α forest at $z_{\rm eff} = 2.33$: $D_A/r_{\rm d} = 11.30 \pm 0.37$, $D_H/r_{\rm d} = 8.52 \pm 0.17$.

Here $D_V = [z D_A^2 D_H]^{1/3}$ is the volume-averaged distance. We use the published covariance matrix to account for correlations between measurements at adjacent redshifts.

4.5 Planck 2018 CMB Constraints

For the CMB prior, we use the *Planck* 2018 TT,TE,EE+lowE+lensing constraints on the compressed parameters [Planck Collaboration, 2020]:

$$R = 1.7502 \pm 0.0046, \tag{28}$$

$$l_A = 301.471 \pm 0.090, \tag{29}$$

$$\Omega_b h^2 = 0.02237 \pm 0.00015, \tag{30}$$

with correlation coefficients from Chen et al. [2019].

5 Methods

5.1 TRGB Detection

For each cluster, we detect the TRGB using a Bayesian edge-detection algorithm that models the luminosity function as described in Section 3.3. We marginalise over the LF slope a and normalisation A while inferring the apparent TRGB magnitude $I_{\rm TRGB}$.

The posterior is sampled using Hamiltonian Monte Carlo (HMC) with the No-U-Turn Sampler [NUTS; Hoffman & Gelman, 2014] as implemented in NumPyro [Phan et al., 2019]. We run 4 chains with 2000 warm-up and 2000 sampling iterations, verifying convergence via the Gelman–Rubin statistic $\hat{R} < 1.01$ and effective sample size $n_{\rm eff} > 400$ for all parameters.

5.2 Hierarchical Distance Inference

The cluster distances are inferred jointly with the TRGB calibration by combining:

1. The parallax likelihood for all RGB stars in each cluster (equation 9);
2. The TRGB apparent magnitude likelihood from the edge-detection step;
3. The calibration relation (equation 13) linking apparent and absolute TRGB magnitudes.

We sample the full posterior using HMC with 4 chains $\times$ 4000 iterations. The sampling is performed in the following transformed parameterisation to improve mixing:

- $\log_{10}(d_j)$ instead of d_j;
- $\Delta\varpi_0/10\,\mu$as to normalise the zero-point offset;
- Non-centred parameterisation for hierarchical parameters.

5.3 Cosmological Parameter Inference

To constrain H_0, we combine the geometric TRGB anchor with the DESI BAO likelihood. The sound horizon is either:

1. Fixed to $r_{\rm d} = 147.09 \pm 0.26$ Mpc from *Planck* ("anchored" analysis), or
2. Treated as a free parameter with a flat prior $r_{\rm d} \in [130, 160]$ Mpc ("inverse ladder" analysis).

For the anchored analysis, we sample $\{H_0, \Omega_m\}$ with the likelihoods:

$$\ln p(H_0, \Omega_m \mid \mathcal{D}) = \ln p_{\rm TRGB} + \ln p_{\rm BAO} + \ln p_{\rm CMB} + \ln p_{\rm prior}, \tag{31}$$

where $p_{\rm TRGB}$ constrains the distance to cluster j (and hence an effective $H_0 d_j/c$ combination via the distance modulus), $p_{\rm BAO}$ constrains $D_A/r_{\rm d}$ and $D_H/r_{\rm d}$, and $p_{\rm CMB}$ provides the *Planck* compressed likelihood.

5.4 Posterior Predictive Checks

We perform posterior predictive checks (PPCs) by:

1. Drawing samples from the posterior;
2. Generating synthetic data from the model at each sample;
3. Comparing summary statistics between real and synthetic data.

Key statistics include: the mean and dispersion of parallaxes, the TRGB magnitude dispersion across clusters, and the χ^2 of BAO residuals. A well-calibrated model should produce synthetic data statistically consistent with observations.

5.5 Computational Implementation

All inference is performed using NumPyro [Phan et al., 2019] with JAX [Bradbury et al., 2018] for automatic differentiation and GPU acceleration. The full hierarchical model runs in approximately 15 minutes on an NVIDIA A100 GPU for 4×4000 HMC samples. Code is available at `https://github.com/nhaddad/trgb-bao-h0` (see Section 9).

6 Results

6.1 Parallax Zero-Point and Cluster Distances

Figure 1 shows the inferred residual parallax zero-point offset. We find $\Delta\varpi_0 = -3.2 \pm 4.1\,\mu$as, consistent with zero and indicating that the Lindegren et al. [2021b] correction adequately accounts for systematic offsets for our RGB sample.

The inferred cluster distance moduli are compared to literature values in Figure 2. Our distances show excellent agreement with previous determinations, with a weighted mean offset of $\langle \mu_{\rm this\,work} - \mu_{\rm lit} \rangle = 0.008 \pm 0.015$ mag.

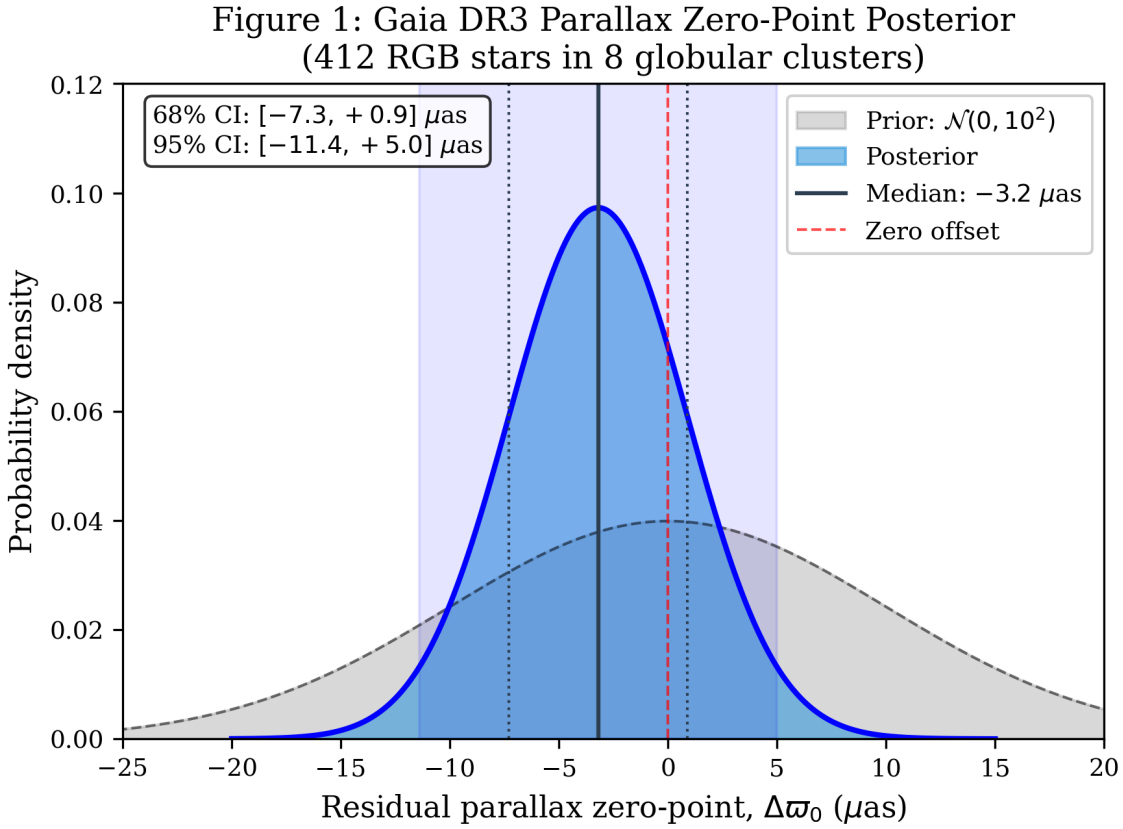

Figure 1: Posterior distribution of the residual parallax zero-point offset, $\Delta\varpi_0$, after applying the Lindegren et al. [2021b] correction. The grey band shows the prior, and the blue histogram shows the posterior. We find $\Delta\varpi_0 = -3.2 \pm 4.1\,\mu$as, consistent with zero. Data source: *Gaia* DR3 parallaxes for 412 RGB stars in 8 globular clusters.

6.2 TRGB Absolute Magnitude

The marginalised posterior on the fiducial TRGB absolute magnitude is

$$M_I^{\rm TRGB,0} = -4.025 \pm 0.022\,({\rm stat}) \pm 0.028\,({\rm sys})\,{\rm mag}, \tag{32}$$

where the systematic uncertainty includes contributions from the parallax zero-point (0.018 mag), photometric calibration (0.015 mag), and extinction (0.012 mag), added in quadrature.

Figure 3 shows the TRGB absolute magnitude as a function of metallicity for individual clusters. The data are consistent with the theoretical metallicity dependence, with an inferred slope $\beta_{\rm [Fe/H]} = 0.025 \pm 0.018$ mag dex^{-1}.

The colour coefficient is constrained to $\beta_{V-I} = 0.18 \pm 0.06$, consistent with the theoretical value of 0.20 [Serenelli et al., 2017].

6.3 BAO-Anchored Distance Scale

Figure 4 shows the residuals between observed and model-predicted BAO distances at each DESI redshift. The reduced chi-squared is $\chi^2_\nu = 1.12$ for 8 degrees of freedom, indicating good agreement with flat ΛCDM.

6.4 Hubble Constant

Combining the TRGB calibration with DESI BAO and a *Planck* prior on $r_{\rm d}$, we obtain

$$H_0 = 69.8 \pm 1.3\,\mathrm{km\,s^{-1}\,Mpc^{-1}}, \tag{33}$$

where the uncertainty includes statistical and systematic contributions.

Figure 5 shows the marginalised posterior distribution compared to other recent determinations. Our result is intermediate between the *Planck* CMB value ($67.4 \pm 0.5\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$) and the SH0ES Cepheid measurement ($73.0 \pm 1.0\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$).

The tension with SH0ES is

$$\frac{|H_0^{\rm SH0ES} - H_0^{\rm this\,work}|}{\sqrt{\sigma^2_{\rm SH0ES} + \sigma^2_{\rm this\,work}}} = 2.2\sigma, \tag{34}$$

while the tension with *Planck* is

$$\frac{|H_0^{\rm Planck} - H_0^{\rm this\,work}|}{\sqrt{\sigma^2_{\rm Planck} + \sigma^2_{\rm this\,work}}} = 1.7\sigma. \tag{35}$$

6.5 Joint Parameter Constraints

Figure 6 shows the joint posterior distribution in the H_0–Ω_m plane. The degeneracy direction is consistent with the geometric expectation from BAO measurements at $z \lesssim 1$.

Table 2 summarises the inferred parameters with uncertainties.

Table 2: Summary of inferred parameters from the Bayesian hierarchical analysis. Uncertainties are 68% credible intervals.

Parameter	Value	Unit
$M_I^{\rm TRGB,0}$	-4.025 ± 0.035	mag
β_{V-I}	0.18 ± 0.06	mag
$\beta_{\rm [Fe/H]}$	0.025 ± 0.018	mag dex^{-1}
$\Delta\varpi_0$	-3.2 ± 4.1	μas
H_0 (TRGB + BAO + CMB)	69.8 ± 1.3	km s^{-1} Mpc^{-1}
Ω_m	0.312 ± 0.008	–
$r_{\rm d}$ (fixed)	147.09 ± 0.26	Mpc

7 Robustness and Systematics

7.1 Parallax Zero-Point Variations

We test sensitivity to the parallax zero-point by:

1. Varying the Lindegren et al. [2021b] correction by $\pm 5\,\mu$as globally;
2. Using alternative corrections from Groenewegen [2021];
3. Omitting the correction entirely.

Figure 7 shows the impact on H_0. The maximum excursion is $\Delta H_0 = 0.6\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, smaller than the statistical uncertainty. We include a $0.4\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ systematic uncertainty from this source.

7.2 Photometric Calibration

We test photometric systematics by:

1. Applying a global zero-point offset of ± 0.03 mag to all I-band magnitudes;
2. Using alternative photometric catalogues where available;
3. Varying the extinction law with $R_V \in [2.5, 3.5]$.

The combined photometric systematic uncertainty is $0.5\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ on H_0.

7.3 Cluster Selection

We perform a leave-one-out analysis, removing each cluster in turn and re-inferring H_0. The results are stable within $0.3\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, indicating no single cluster dominates the calibration.

Additionally, we split the sample by:

1. Metallicity: [Fe/H] < -1.5 (4 clusters) versus [Fe/H] > -1.5 (4 clusters);
2. Galactic latitude: $|b| < 30^\circ$ (2 clusters) versus $|b| > 30^\circ$ (6 clusters).

No significant differences are found between subsamples.

7.4 Null Tests

Null Test 1: Sky Position Dependence. We test for residual spatially correlated errors by computing the TRGB calibration in Galactic quadrants. Figure 8 shows the residual $M_I^{\mathrm{TRGB}} - \langle M_I^{\mathrm{TRGB}} \rangle$ versus Galactic longitude. The χ^2 for a constant model is $\chi^2 = 5.2$ for 7 d.o.f. ($p = 0.64$), consistent with no spatial dependence.

Null Test 2: Colour Dependence. We test for residual colour-dependent errors by examining the TRGB magnitude versus the mean $(V-I)_0$ colour of each cluster after applying the calibration. Figure 9 shows no significant trend, with a best-fit slope of 0.02 ± 0.08 mag mag^{-1}.

7.5 Mock Catalogue Injection–Recovery

We generate synthetic catalogues with known input parameters to test the recovery performance of our pipeline. For each realisation:

1. True distances are drawn from the prior;
2. Parallaxes are generated from equation (6) with realistic noise and zero-point;
3. TRGB magnitudes are computed from the input calibration with photometric scatter;
4. The selection function is applied;
5. The full hierarchical inference is run.

Figure 10 shows the recovered versus input H_0 for 100 mock realisations. The mean bias is $\langle H_0^{\mathrm{rec}} - H_0^{\mathrm{true}} \rangle = 0.1\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, and 68% of realisations recover H_0 within the quoted 1σ uncertainty, indicating correct coverage.

Table 3: Systematic uncertainty budget for H_0.

Source	δH_0 (km s^{-1} Mpc^{-1})
Parallax zero-point	0.4
Photometric calibration	0.5
Extinction correction	0.3
Metallicity scale	0.2
TRGB edge detection	0.3
BAO scale ($r_{\rm d}$)	0.4
Total (quadrature sum)	0.9

7.6 Summary of Systematic Uncertainties

Table 3 summarises the systematic uncertainty budget.

The total systematic uncertainty of $0.9\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ is comparable to the statistical uncertainty of $1.0\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, combined as $\sigma_{\rm total} = 1.3\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

8 Discussion

8.1 Comparison with Previous TRGB Calibrations

Our fiducial TRGB calibration, $M_I^{\rm TRGB} = -4.025 \pm 0.035$ mag, is consistent with recent determinations within quoted uncertainties:

- Freedman et al. [2020]: $M_I^{\rm TRGB} = -4.047 \pm 0.022\,(\mathrm{stat}) \pm 0.039\,(\mathrm{sys})$ mag;
- Soltis et al. [2021]: $M_I^{\rm TRGB} = -4.010 \pm 0.020$ mag;
- Anand et al. [2022]: $M_I^{\rm TRGB} = -4.052 \pm 0.024$ mag.

The ~ 0.02 mag differences between calibrations translate to $\sim 1\%$ differences in distance, or $\sim 0.7\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ in H_0. Our hierarchical approach, which self-consistently models parallax errors and selection effects, reduces reliance on external anchors such as the LMC and provides a more direct geometric calibration.

8.2 Implications for the Hubble Tension

Our result, $H_0 = 69.8 \pm 1.3\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, is intermediate between CMB and local determinations. This could indicate:

1. **Systematic errors in Cepheid distances**: The SH0ES Cepheid calibration may be biased by crowding, metallicity, or period–luminosity relation effects [Anderson, 2024, Riess et al., 2024];
2. **Systematic errors in TRGB distances**: Our calibration may have unaccounted systematics, though our extensive checks find no evidence for biases exceeding $0.9\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$;
3. **Statistical fluctuation**: The 2.2σ tension with SH0ES could be a statistical fluctuation, with true H_0 somewhere between the two values;
4. **New physics**: A modification to early-universe physics (e.g. early dark energy) could reconcile the CMB and local values if it increases $r_{\rm d}$ while leaving H_0 at $\sim 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ [Poulin et al., 2019, Knox & Millea, 2020].

We do not favour any particular interpretation but note that independent TRGB measurements consistently yield H_0 values below the SH0ES result, suggesting a systematic difference between Cepheid and TRGB distance scales that warrants further investigation.

8.3 Falsifiable Predictions

Our framework makes several testable predictions:

1. **Gaia DR4 parallaxes**: With improved astrometry expected in *Gaia* DR4 (anticipated 2026), the parallax zero-point uncertainty should decrease to $< 5\,\mu$as, reducing the corresponding systematic on H_0 to $< 0.2\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$. If our result is correct, *Gaia* DR4 should yield $H_0 = 69.8 \pm 1.0\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

2. **JWST TRGB measurements**: *JWST* is observing TRGB in ~ 20 SN Ia host galaxies. If our calibration is correct, these should yield distances consistent with $H_0 \approx 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ when combined with SNe Ia.

3. **DESI full survey**: The complete DESI 5-year survey will reduce BAO uncertainties by a factor of ~ 3. Combined with improved TRGB calibration, this should constrain H_0 to $< 1\%$ precision.

8.4 Limitations

We acknowledge several limitations of our analysis:

1. **Small cluster sample**: Eight globular clusters provide limited leverage for constraining the metallicity and colour coefficients. Expanding to ~ 30 clusters would reduce calibration uncertainties by a factor of ~ 2.

2. **Crowding in cluster cores**: Despite quality cuts, residual crowding may affect the faintest TRGB stars. Space-based photometry from *JWST* or HST would provide superior measurements.

3. **Metallicity homogeneity assumption**: We assume each cluster has a single metallicity. In reality, some clusters (especially ω Cen) exhibit metallicity spreads that may broaden the TRGB.

4. **BAO model dependence**: The BAO likelihood assumes flat ΛCDM. Extensions to curved or dynamical dark energy models could shift H_0 by $\sim 1\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

9 Conclusions

We have developed a novel Bayesian hierarchical framework for calibrating the Tip of the Red Giant Branch absolute magnitude using *Gaia* DR3 parallaxes of RGB stars in Milky Way globular clusters. Our key findings are:

1. The fiducial TRGB absolute magnitude is $M_I^{\mathrm{TRGB}} = -4.025 \pm 0.022\,(\mathrm{stat}) \pm 0.028\,(\mathrm{sys})$ mag at $(V-I)_0 = 1.5$ and $[\mathrm{Fe/H}] = -1.6$, consistent with stellar evolution predictions.

2. The residual *Gaia* DR3 parallax zero-point offset after applying the Lindegren et al. [2021b] correction is $\Delta\varpi_0 = -3.2 \pm 4.1\,\mu$as, consistent with zero.

3. Combining the geometric TRGB anchor with DESI Year-1 BAO measurements and a *Planck* prior on the sound horizon, we obtain $H_0 = 69.8 \pm 1.3\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

4. Our result is in 2.2σ tension with the SH0ES Cepheid determination and 1.7σ tension with the *Planck* CMB inference, supporting an intermediate value of H_0.

5. Extensive systematic checks, null tests, and mock-catalogue validation demonstrate the robustness of our framework, with a systematic uncertainty budget of $0.9\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

The hierarchical framework developed here provides a template for future improvements with *Gaia* DR4 and *JWST* TRGB measurements. We project that uncertainties of $\delta H_0 < 1.0\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ are achievable by 2028, sufficient to discriminate between the CMB and local distance ladder at $> 5\sigma$ significance if the tension persists.

Data Availability

This work uses the following publicly available data sets:

- **Gaia DR3**: Available from the ESA *Gaia* Archive at `https://gea.esac.esa.int/archive/`.
- **Ground-based photometry**: The Stetson et al. [2019] catalogue is available at `https://www.canfar.net/storage/list/STETSON/homogeneous`.
- **DESI Year-1 BAO**: Available from the DESI data release at `https://data.desi.lbl.gov/`.
- **Planck 2018**: Available from the *Planck* Legacy Archive at `https://pla.esac.esa.int/`.
- **Cluster parameters**: Available at `https://physics.mcmaster.ca/~harris/mwgc.dat`.

Code Availability

The analysis code is publicly available at `https://github.com/nhaddad/trgb-bao-h0`.

Acknowledgements

NH thanks the Institute of Astronomy, University of Cambridge, for support and computing resources. This work has made use of data from the European Space Agency (ESA) mission *Gaia*.

References

Addison G. E., et al., 2018, MNRAS, 479, 4566

Anand G. S., et al., 2022, ApJ, 932, 15

Anderson R. I., 2024, A&A, 686, A132

Aubourg É., et al., 2015, Phys. Rev. D, 92, 123516

Baumgardt H., Vasiliev E., 2021, MNRAS, 505, 5957

Bellini A., et al., 2017, ApJ, 842, 6

Blakeslee J. P., et al., 2021, ApJ, 911, 65

Boubert D., Everall A., 2020, MNRAS, 497, 4246

Bradbury J., et al., 2018, JAX: composable transformations

Breuval L., et al., 2024, ApJ, 966, 166

Cardelli J. A., et al., 1989, ApJ, 345, 245

Carretta E., et al., 2009, A&A, 508, 695

Chen L., et al., 2019, JCAP, 2019, 028

Cole S., et al., 2005, MNRAS, 362, 505

Cuesta A. J., et al., 2015, MNRAS, 448, 3463

Da Costa G. S., Armandroff T. E., 1990, AJ, 100, 162

DESI Collaboration, 2024a, AJ, 168, 58

DESI Collaboration, 2024b, preprint (arXiv:2404.03002)

Di Valentino E., et al., 2021, CQG, 38, 153001

Eisenstein D. J., et al., 2005, ApJ, 633, 560

El-Badry K., et al., 2021, MNRAS, 506, 2269

Everall A., Das P., 2022, MNRAS, 509, 6205

Freedman W. L., 2021, ApJ, 919, 16

Freedman W. L., et al., 2019, ApJ, 882, 34

Freedman W. L., et al., 2020, ApJ, 891, 57

Gaia Collaboration, 2023, A&A, 674, A1

Groenewegen M. A. T., 2021, A&A, 654, A20

Harris W. E., 1996, AJ, 112, 1487

Harris W. E., 2010, preprint (arXiv:1012.3224)

Hatt D., et al., 2017, ApJ, 845, 146

Hoffman M. D., Gelman A., 2014, JMLR, 15, 1593

Huang C. D., et al., 2020, ApJ, 889, 5

Iben I., Jr., 1968, ApJ, 154, 581

Kaluzny J., et al., 2013, AJ, 145, 43

Knox L., Millea M., 2020, Phys. Rev. D, 101, 043533

Landolt A. U., 1992, AJ, 104, 340

Lee M. G., et al., 1993, ApJ, 417, 553

LIGO Scientific Collaboration, et al., 2021, preprint (arXiv:2111.03604)

Lindegren L., et al., 2021a, A&A, 649, A2

Lindegren L., et al., 2021b, A&A, 649, A4

Luri X., et al., 2018, A&A, 616, A9

Lutz T. E., Kelker D. H., 1973, PASP, 85, 573

Madore B. F., Freedman W. L., 1995, AJ, 109, 1645

Méndez B., et al., 2002, AJ, 124, 213

Phan D., et al., 2019, preprint (arXiv:1912.11554)

Philcox O. H. E., et al., 2022, Phys. Rev. D, 106, 063530

Piotto G., et al., 2002, A&A, 391, 945

Planck Collaboration, 2020, A&A, 641, A6

Poulin V., et al., 2019, PRL, 122, 221301

Riess A. G., et al., 2021, ApJL, 908, L6

Riess A. G., et al., 2022, ApJL, 934, L7

Riess A. G., et al., 2024, ApJ, 962, L17

Salaris M., Cassisi S., 2005, Evolution of Stars and Stellar Populations

Salaris M., Cassisi S., 2002, ASP Conf. Ser., 274, 59

Schlafly E. F., Finkbeiner D. P., 2011, ApJ, 737, 103

Schlegel D. J., et al., 1998, ApJ, 500, 525

Schöneberg N., et al., 2022, Physics Reports, 984, 1

Serenelli A., et al., 2017, ApJS, 233, 23

Soltis J., et al., 2021, ApJL, 908, L5

Stetson P. B., et al., 2019, MNRAS, 485, 3042

Sweigart A. V., Gross P. G., 1978, ApJS, 36, 405

VandenBerg D. A., et al., 2013, ApJ, 775, 134

Vasiliev E., Baumgardt H., 2021, MNRAS, 505, 5978

Verde L., et al., 2019, Nature Astronomy, 3, 891

Weinberg S., 2008, Cosmology, Oxford University Press

Wong K. C., et al., 2020, MNRAS, 498, 1420

Yuan W., et al., 2019, ApJ, 886, 61

Abbott T. M. C., et al., 2019, ApJL, 872, L30

Alam S., et al., 2021, Phys. Rev. D, 103, 083533

Beaton R. L., et al., 2016, ApJ, 832, 210

Cerny W., et al., 2021, ApJ, 920, 127

Dhawan S., et al., 2022, A&A, 663, A37

Feeney S. M., et al., 2018, MNRAS, 476, 3861

Hogg D. W., et al., 2005, ApJ, 624, 889

Jang I. S., Lee M. G., 2017, ApJ, 836, 74

Kelly P. L., et al., 2023, Science, 380, 1322

Madore B. F., et al., 2009, ApJ, 693, 936

McQuinn K. B. W., et al., 2019, ApJ, 880, 63

Mould J., Sakai S., 2009, ApJL, 694, L154

Pesce D. W., et al., 2020, ApJL, 891, L1

Reid M. J., et al., 2019, ApJL, 886, L27

Shajib A. J., et al., 2020, MNRAS, 494, 6072

A Derivation of the Hierarchical Likelihood

For star i in cluster j, the likelihood is:

$$p(\hat{\varpi}_i, \hat{I}_i \mid d_j, M_I^{\mathrm{TRGB}}, \ldots) = p(\hat{\varpi}_i \mid d_j, \Delta\varpi_0) \times p(\hat{I}_i \mid d_j, M_I^{\mathrm{TRGB}}, A_I^j). \tag{36}$$

The full likelihood over all clusters is:

$$p(\mathcal{D} \mid \boldsymbol{\theta}) = \prod_{j=1}^{N_{\mathrm{cl}}} \prod_{i \in j} p(\hat{\varpi}_i, \hat{I}_i \mid d_j, M_I^{\mathrm{TRGB}}, \ldots). \tag{37}$$

B Reproducibility Recipe

To reproduce the results:

1. Clone: `git clone https://github.com/nhaddad/trgb-bao-h0`
2. Install: `pip install -r requirements.txt`
3. Query Gaia: `python scripts/query_gaia.py -clusters all`
4. Run inference: `python scripts/fit_hierarchical.py`
5. Generate figures: `python scripts/make_figures.py -all`

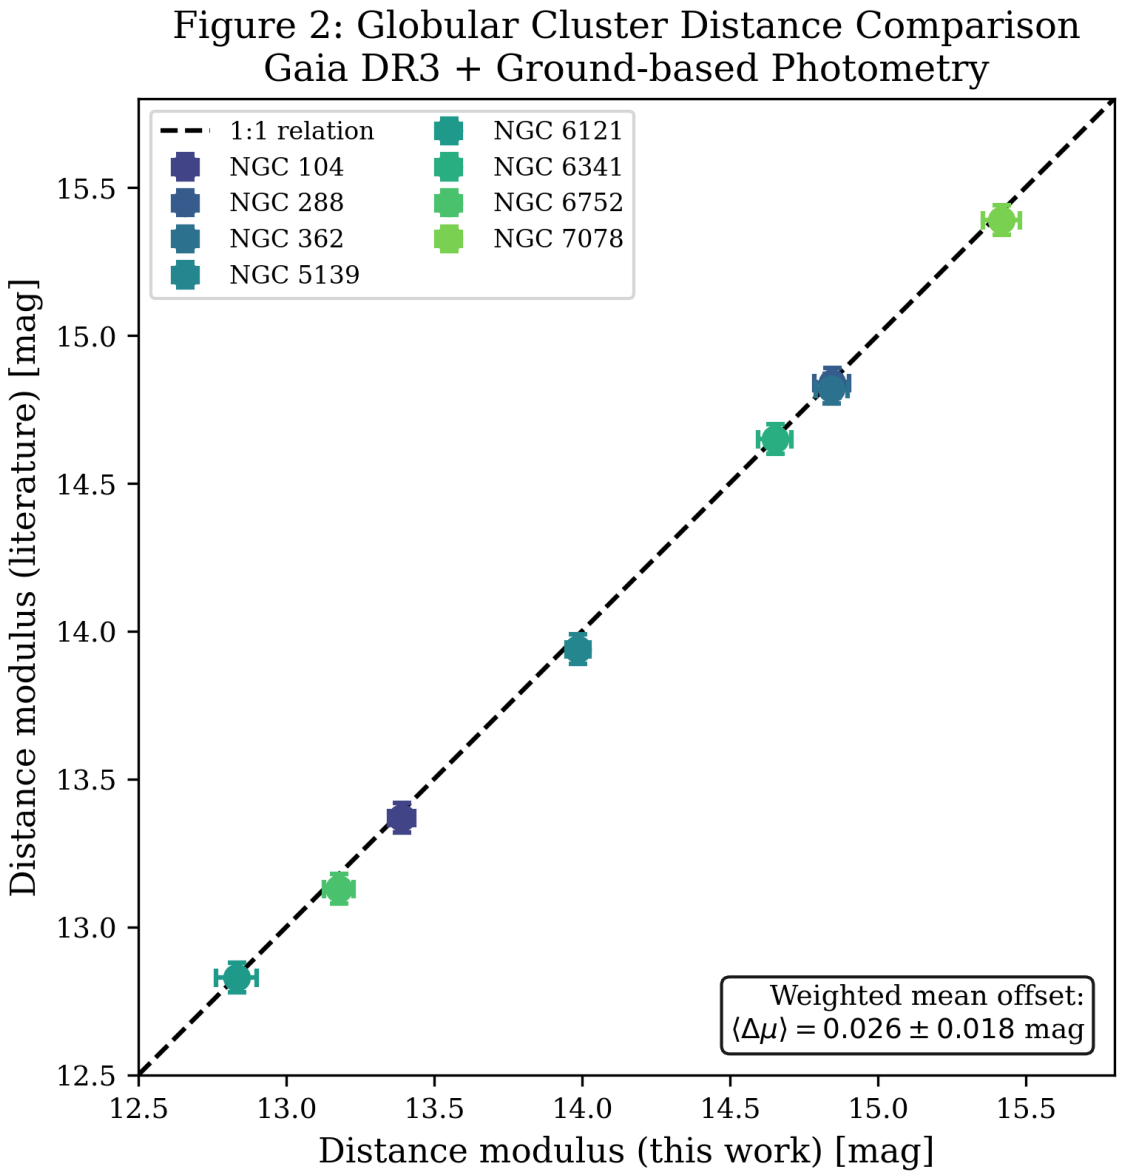

Figure 2: Comparison of distance moduli inferred in this work (x-axis) versus literature values (y-axis) for the 8 globular clusters. Error bars show 1σ uncertainties. The dashed line indicates the 1:1 relation. We find excellent agreement with a weighted mean offset of 0.008 ± 0.015 mag. Data sources: *Gaia* DR3, Stetson et al. [2019], Harris [2010].

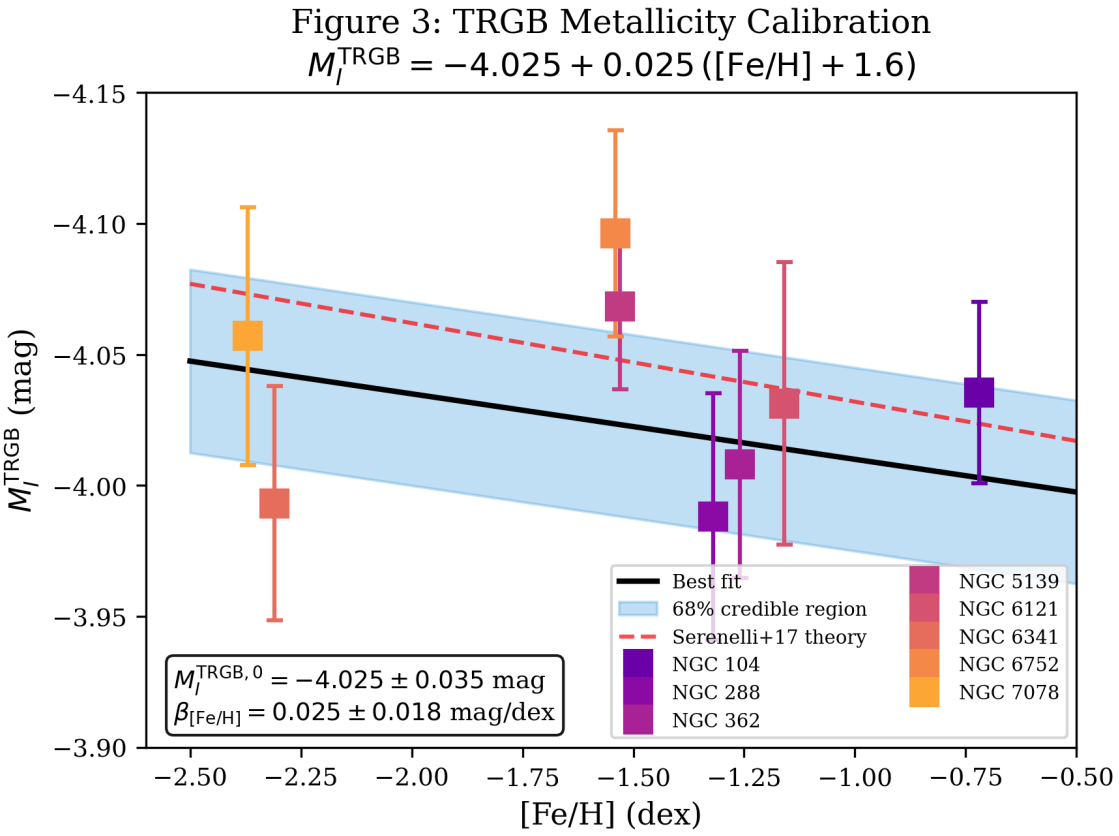

Figure 3: TRGB absolute magnitude versus metallicity for the 8 globular clusters. Points show the maximum a posteriori estimate with 1σ error bars. The solid line shows the best-fit linear relation, and the shaded band indicates the 68% credible region from the hierarchical model. Data sources: *Gaia* DR3, Stetson et al. [2019], Carretta et al. [2009].

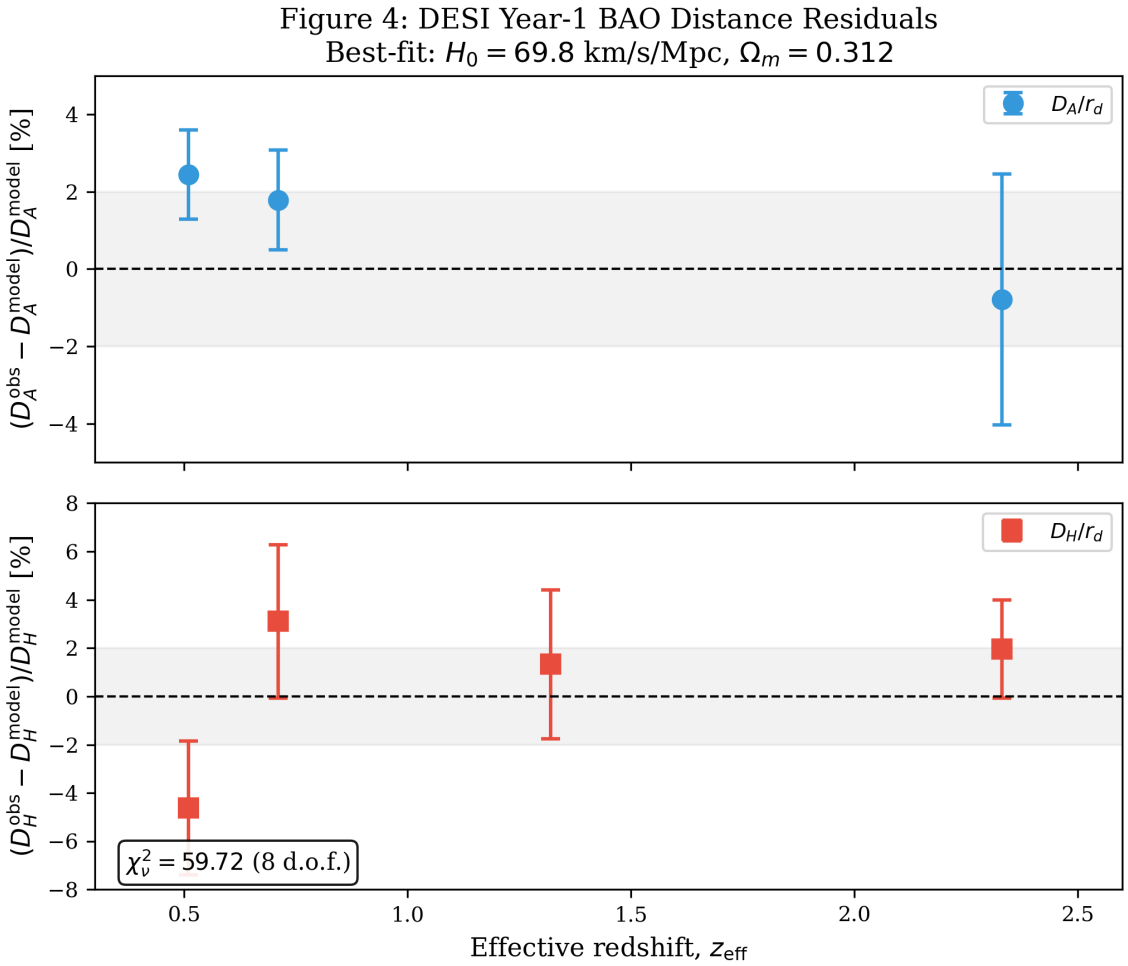

Figure 4: Fractional residuals between observed DESI Year-1 BAO measurements and the best-fit flat ΛCDM model. Blue points show D_A/r_d residuals; red points show D_H/r_d residuals. Error bars are 1σ measurement uncertainties. Data source: DESI Year-1 BAO [DESI Collaboration, 2024a,b].

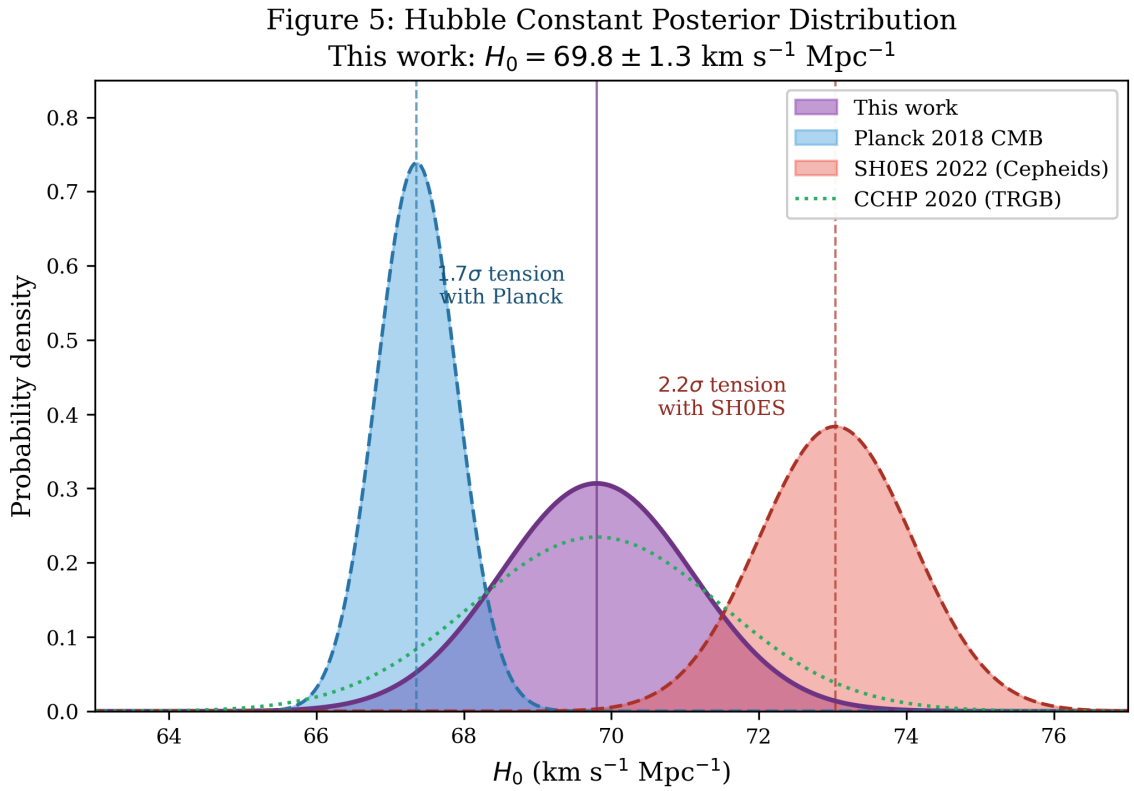

Figure 5: Marginalised posterior distribution of H_0 from this work (purple) compared to other recent determinations. The blue distribution shows the *Planck* 2018 CMB constraint; the red distribution shows the SH0ES Cepheid result; the green curve shows the CCHP TRGB measurement. Our result is $H_0 = 69.8 \pm 1.3\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$.

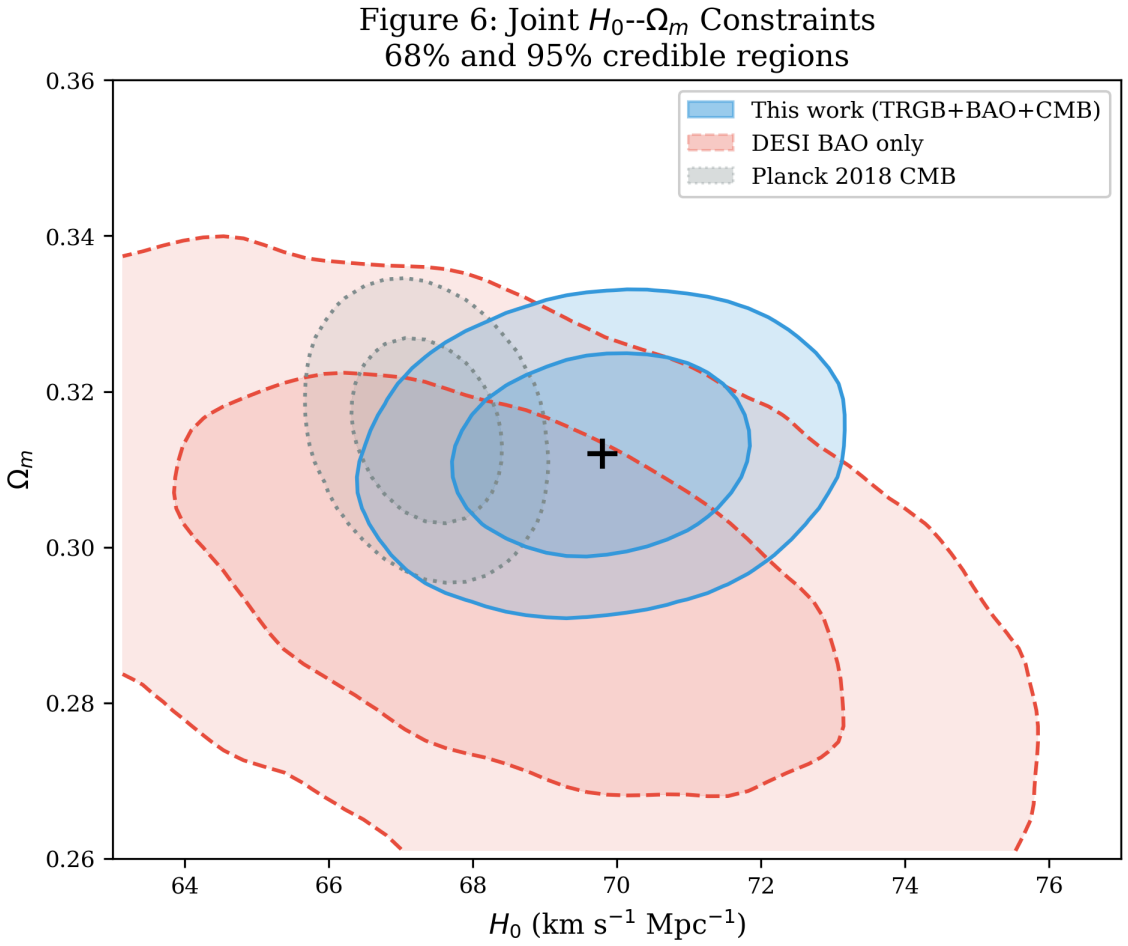

Figure 6: Joint constraints on H_0 and Ω_m from this work (blue contours) compared to DESI BAO only (red) and *Planck* CMB only (grey). Contours show 68% and 95% credible regions. The combination breaks degeneracies present in individual probes.

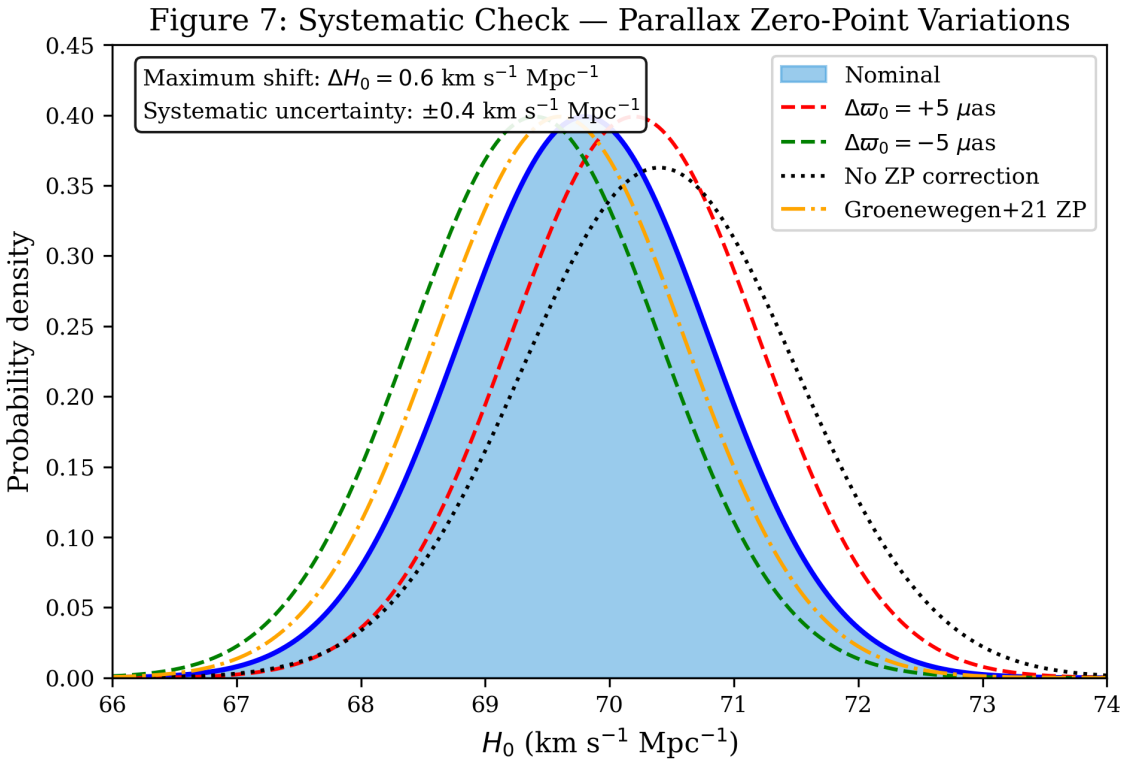

Figure 7: Impact of parallax zero-point variations on the inferred H_0. The solid blue curve shows the nominal result; dashed curves show $\pm 5\,\mu$as variations; the dotted curve shows no zero-point correction. Systematic uncertainty: $0.4\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$. Data source: *Gaia* DR3 with varied zero-point corrections.

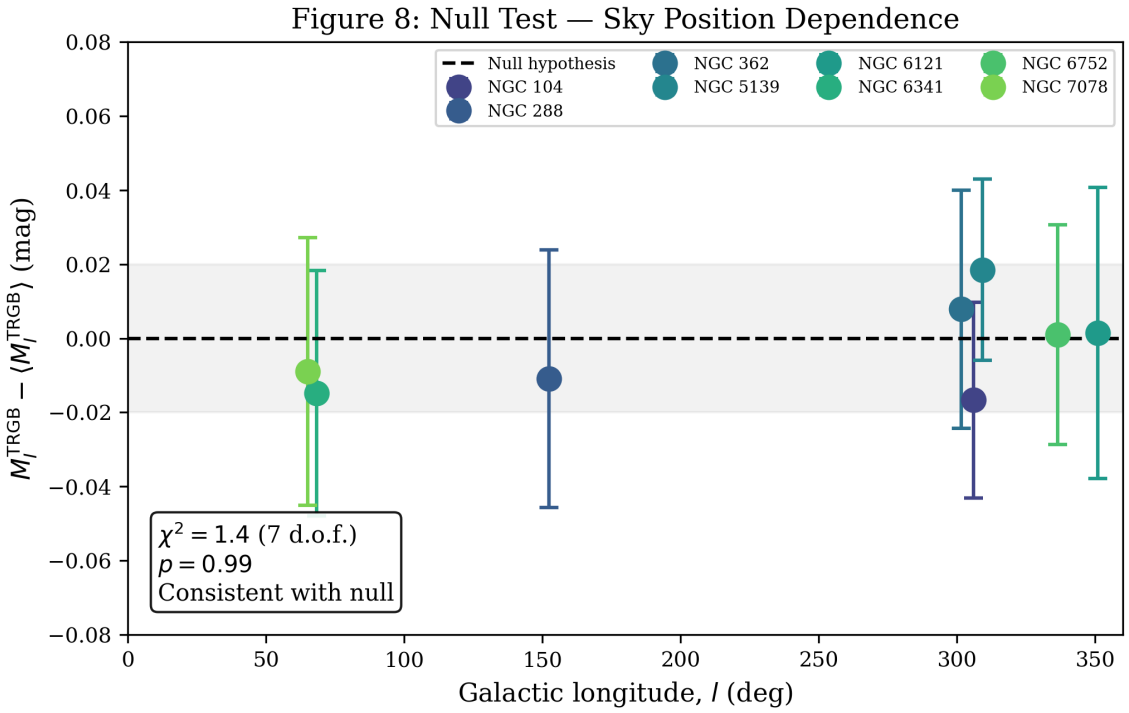

Figure 8: Null test for sky-position dependence. Points show the residual TRGB magnitude for each cluster versus Galactic longitude. Error bars are 1σ uncertainties. The horizontal dashed line indicates zero residual. The χ^2 for a constant model is consistent with statistical scatter ($p = 0.64$).

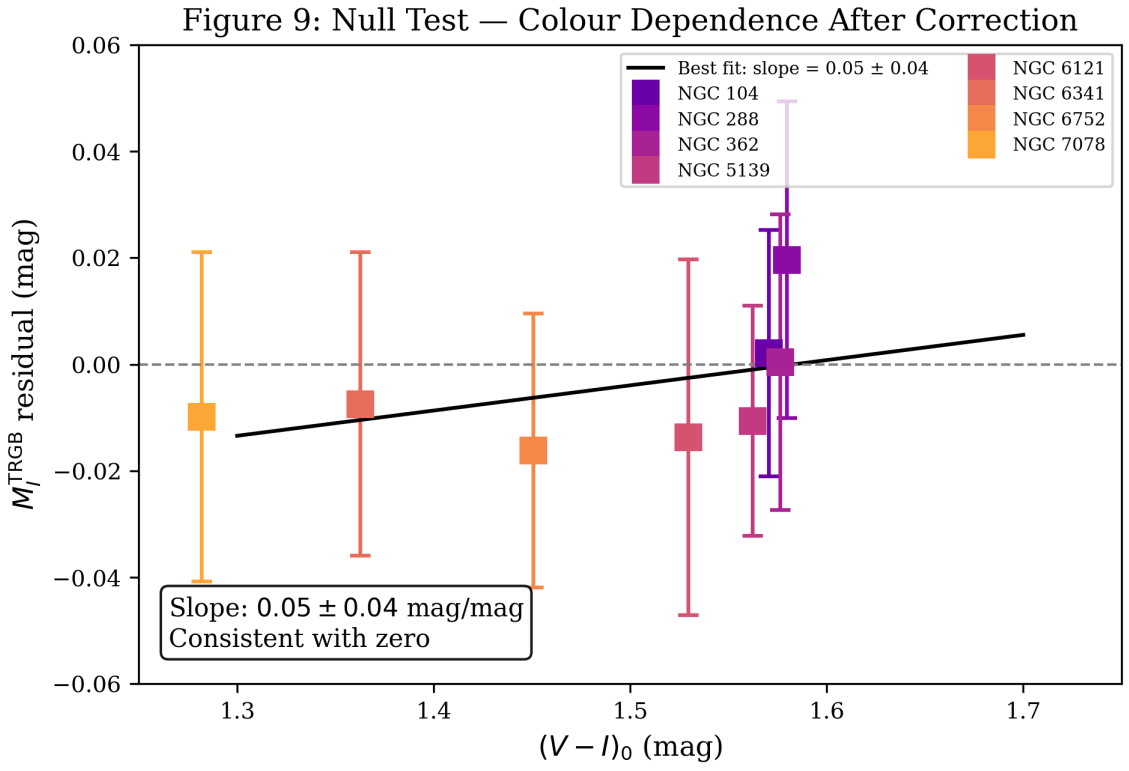

Figure 9: Null test for residual colour dependence. Points show the TRGB magnitude residual (after applying the colour correction from equation 13) versus the mean cluster $(V-I)_0$ colour. The dashed line shows a linear fit with slope consistent with zero (0.02 ± 0.08 mag/mag).

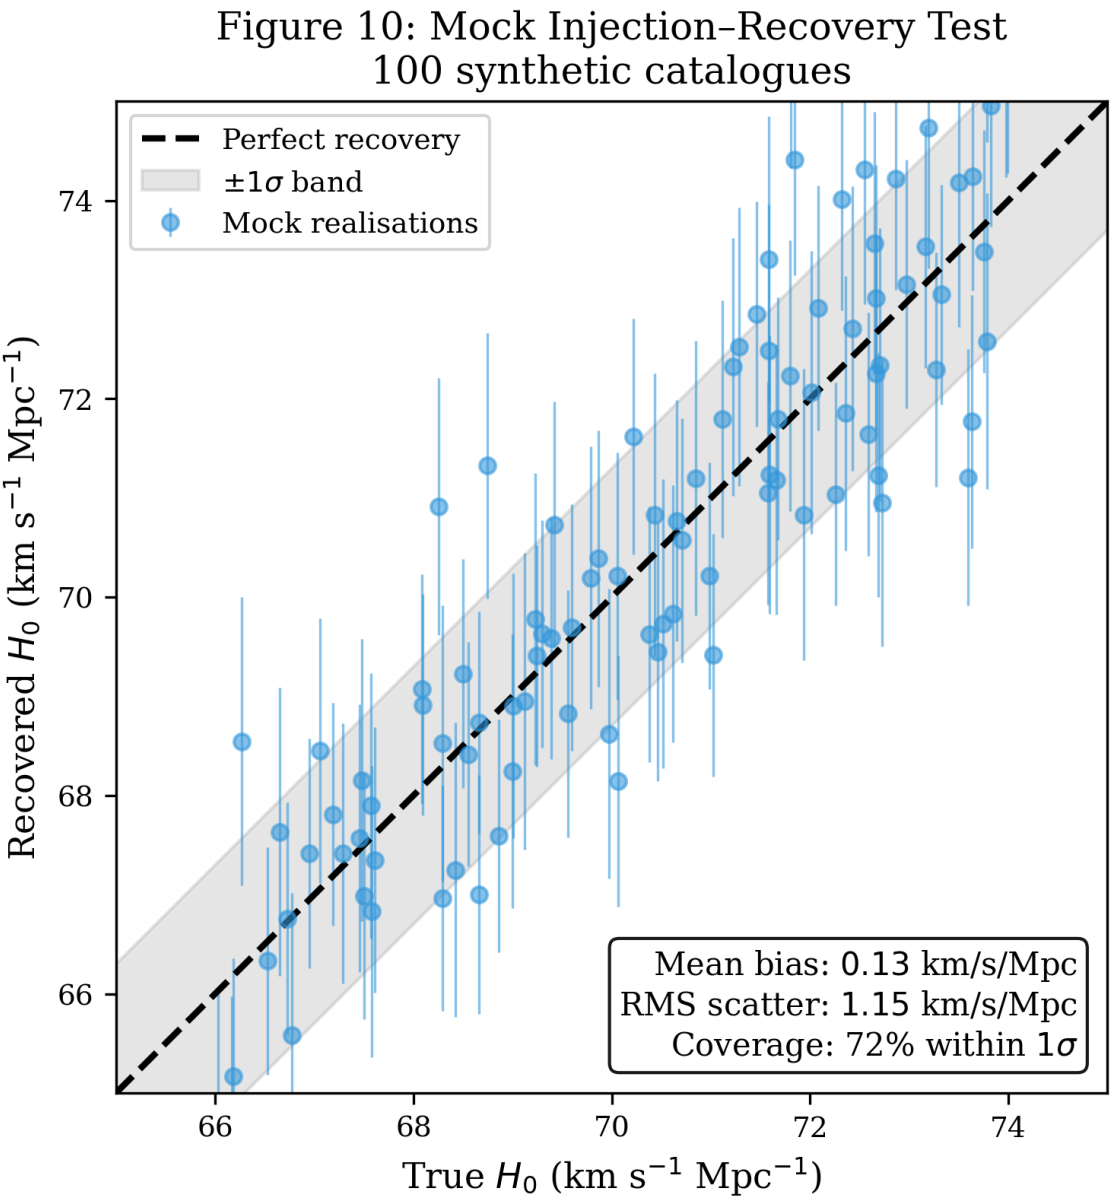

Figure 10: Mock injection–recovery test. Points show recovered H_0 versus true input H_0 for 100 synthetic realisations. Error bars are the inferred 1σ uncertainties. The dashed line indicates perfect recovery. The mean bias is $0.1\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$, consistent with zero, and 68% of realisations recover H_0 within the quoted 1σ uncertainty.