Weak CN is the normal signature of a cool, oxygen-rich red supergiant—not a halfway carbon star.

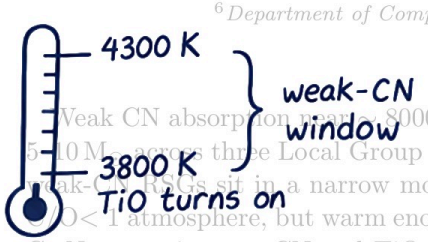
Draft version June 1, 2026
Typeset using LaTeX **twocolumn** style in AASTeX631

arXiv:2605.30733v1 [astro-ph.SR] 29 May 2026

Weak-CN Stars Are Ordinary Cool Red Supergiants

Yuan-Sen Ting,[1,2,3] Puragra Guhathakurta,[4] Douglas Grion Filho,[4] Evan N. Kirby,[5] and Elliot M. Kim[6]

[1] *Department of Astronomy, The Ohio State University, Columbus, OH 43210, USA*
[2] *Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH 43210, USA*
[3] *Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany*
[4] *Department of Astronomy and Astrophysics, University of California Santa Cruz, 1156 High Street, Santa Cruz, CA 95064, USA*
[5] *Department of Physics and Astronomy, University of Notre Dame, 225 Nieuwland Science Hall, Notre Dame, IN 46556, USA*
[6] *Department of Computer Science, Cornell University, Ithaca, NY 14853, USA*





ABSTRACT

Weak CN absorption near $\simeq 8000$ Å has recently been detected in evolved red supergiants (RSGs) of 5–10 $M_\odot$ across three Local Group galaxies (Guhathakurta et al. 2025; Grion Filho et al. 2025). These weak-CN RSGs sit in a narrow molecular regime: cool enough for CN to be visible in a non-carbon, C/O<1 atmosphere, but warm enough that TiO is not saturated and varies in $T_{\rm eff}$ and in the surface C+N reservoir move CN and TiO in distinct directions. We test this picture with pseudo-continuum equivalent widths (EWs) measured from LMC, M33, and M31 weak-CN and carbon-star coadds, compared at matched resolution to a self-consistent grid of synthetic RSG atmospheres spanning $T_{\rm eff}$, [α/Fe], and surface C and N offsets relative to each host's scaled-solar baseline. Ordinary cool-RSG models reproduce the weak-CN coadds across all three hosts, with per-feature residuals at the level of the adopted EW systematic floors. The robust observable is the combined surface abundance Δ[C/H]$+\Delta$[N/H] rather than each offset individually, because CN forms from the product of available C and N number densities. Mapping Δ[C/H] + Δ[N/H] to initial rotation through PARSEC v2.0 has modest leverage — the variable shifts by $\simeq 0.07$ dex from $\omega_i = 0$ to $\omega_i = 0.6$ — and within this resolution slow-rotation first dredge-up is consistent with LMC and M33, and with M31 once a single-feature CaT 8542 Å calibration anchor is allowed. The straightforward resolution of the discovery puzzle is therefore that weak CN is not an exotic carbon-star intermediate but the expected molecular-equilibrium signature of ordinary cool RSGs.

A low-resolution CNO tracer could reach many more red supergiants than precision abundance spectroscopy.

Keywords: stars: abundances — stars: atmospheres — stars: rotation — supergiants — galaxies: stellar content

1. INTRODUCTION

RSGs dominate the integrated optical and near-infrared light of young stellar populations during the first ~ 50 Myr after a star-formation event, and they are the principal sources returning CNO-cycle-processed material to the interstellar medium on that timescale. Their surface abundances are therefore a direct input to stellar-population synthesis and Galactic chemical evolution. Measuring those abundances at high resolution is difficult: cool-RSG optical spectra are blanketed by a dense molecular forest (Levesque et al. 2005; Gustafsson et al. 2008), the photospheres pulsate and convect on month timescales, and the apertures required at high signal-to-noise reach only the brightest nearest stars. An indirect spectroscopic tracer that responds robustly to RSG CNO chemistry at low or moderate spectral resolution would be valuable, and it is in this context that we re-examine the recently identified weak-CN class.

The puzzle disappears when CN and TiO are treated together rather than using CN as a carbon-star flag.

The weak-CN stars discovered by Guhathakurta et al. (2025) and extended to the LMC by Grion Filho et al. (2025) occupy an awkward spectroscopic category. Their spectra show a shallow CN absorption band around 7800–8200 Å, the same broad feature commonly used to identify carbon stars (Boyer et al. 2013; Hamren et al. 2016; Pastorelli et al. 2020), but the rest of the optical spectrum has the morphology of a non-carbon, C/O<1 cool star, with prominent TiO bands. The stars sit on the young red-supergiant locus of the color–magnitude diagram, at masses 5–10 $M_\odot$ and main-sequence ages $\simeq 40$ Myr in the core-He-burning phase. Guhathakurta et al. (2025) treated the class as mysterious: CN was the carbon-star-like feature, while TiO

Temperature and surface chemistry push CN and TiO in different directions, giving us the needed diagnostic leverage.

was the normal non-carbon-star feature, sitting on the cool-RSG locus, whose effective-temperature scale in this regime was established by Levesque et al. (2005). Grion Filho et al. (2025) showed that similar objects are common in the LMC, used the pre-computed Aringer et al. (2016) model atmospheres to bracket $T_{\rm eff} \simeq 4000$–$4600\,$K, and argued that weak CN can occur in normal evolved stars over a restricted temperature and gravity range, with surface nitrogen enhancement from main-sequence rotational mixing as a plausible source of EW scatter.

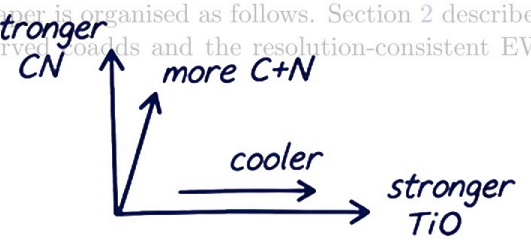
The cool-RSG optical spectrum is dominated by two molecules, TiO and CN, whose responses to temperature and chemistry are sharply different. TiO crosses its dissociation transition near $T_{\rm eff} \sim 3800\,$K, well inside the RSG plateau, while CN remains effective up to $T_{\rm eff} \simeq 4500\,$K. CN is also a C–N molecule whose strength depends directly on the surface C and N reservoirs, whereas TiO is essentially insensitive to those abundances at fixed metallicity. This asymmetry between the two molecules is the physical basis on which the rest of the analysis rests.

CN alone is ambiguous: warming the atmosphere and changing CNO can produce similarly shallow bands.

Two mechanisms can therefore produce a shallow CN feature in such a spectrum, and the discovery-paper ambiguity reflects exactly that. The first is a $T_{\rm eff}$ selection: in the temperature range where TiO weakens against the molecular pseudo-continuum, the residual CN band becomes a natural feature of an ordinary C/O<1 atmosphere, with no abundance peculiarity required. The second is a surface CNO enhancement from first dredge-up plus main-sequence rotational mixing, which raises the surface nitrogen abundance and deepens CN at fixed $T_{\rm eff}$. The CN band alone cannot distinguish these two scenarios; both produce a shallow CN morphology of comparable strength.

The diagnostic leverage comes from the partnership between CN and TiO: $T_{\rm eff}$ moves TiO strongly but barely touches CN, whereas a CNO surface enhancement moves CN through the C and N reservoirs but barely shifts TiO at fixed $T_{\rm eff}$. A joint χ^2 over the two indices therefore separates temperature from chemistry. Two further pieces of physics shape how that separation works in detail: the photospheric structure back-reacts to abundance perturbations through molecular line blanketing at $\log g \simeq 0$ (Section 3), and the CN bands respond to the product of the C and N reservoirs so the constrained surface variable is the combination $\Delta[{\rm C/H}] + \Delta[{\rm N/H}]$ rather than each offset individually (Section 4). Mapping that combination onto initial rotation rate further requires a stellar-evolution prior.

The paper is organised as follows. Section 2 describes the observed coadds and the resolution-consistent EW measurement. Section 3 describes the self-consistent atmosphere grid. Section 4 uses the model atmospheres to show why weak CN is a normal cool-RSG feature. Section 5 presents the per-galaxy joint χ^2 fits, the best-fit spectra, and the per-galaxy inferred surface composition. Section 6 discusses the C+N degeneracy, where in the EW plane the diagnostic discriminates rotation, the role of C_2 Swan as a candidate pure-carbon lever, the slow-rotation consistency of the data, and the implications for using weak-CN stars as recent-SFR tracers.



2. DATA AND EW MEASUREMENT

We use the inverse-variance-weighted coadded templates of the visually classified weak-CN and carbon-star samples published by Guhathakurta et al. (2025) for M31 and M33 and by Grion Filho et al. (2025) for the LMC. In both papers the weak-CN and carbon-star populations are drawn from the same spectroscopic surveys (Keck/DEIMOS for M31 and M33, CTIO/Hydra for the LMC) and the carbon-star coadd is the classical comparison template against which the weak-CN class was originally defined — a deeper CN red band and a [illegible] Swan absorption. We adopt those coadds as in[illegible] re-classification. The high-quality subsets [illegible] from 75 weak-CN and 71 carbon-star [illegible] in M31, 140 and 34 in M33, and 779 and 1950 [illegible] the LMC, a cleaner subset of the broader survey-level identifications of 224/659/779 weak-CN stars in M31 / M33 / LMC. M31 and M33 spectra were taken with Keck/DEIMOS at native resolution $R \simeq 2000$ over 4000–10 000 Å; the LMC spectra were taken with CTIO/Hydra at $R \simeq 7000$ over 6850–9150 Å. Figure 1 shows the observed templates and the EW windows used here.

The useful observable is not "CN present?" but "how does CN move relative to TiO?"

The purpose of the carbon-star overlay is historical and diagnostic: the weak-CN class was noticed because its broad CN morphology resembles a diluted carbon-star CN band, but the same spectra retain non-carbon-star features such as TiO and CaT. The carbon-star coadds therefore define what the weak-CN stars are not.

The figure also shows a practical limitation that carries through the analysis. The LMC HYDRA coverage contains TiO 7050 Å, CN red, N I, and CaT, but does not extend to the bluer TiO γ or C_2 Swan windows that M31 and M33 do (at lower spectral resolution). We therefore do not force all three galaxies into an identical feature set; instead we use the features each data set can support and keep the interpretation tied to the available diagnostic geometry.

All EWs are measured after division by the same pseudo-continuum recipe in data and models. For each spectrum we estimate a local continuum $F_c(\lambda)$ as the

This is why the paper becomes a joint-index problem rather than a classification exercise.

The spectra look carbon-star-like only in CN; everywhere else they retain ordinary oxygen-rich red-supergiant features.

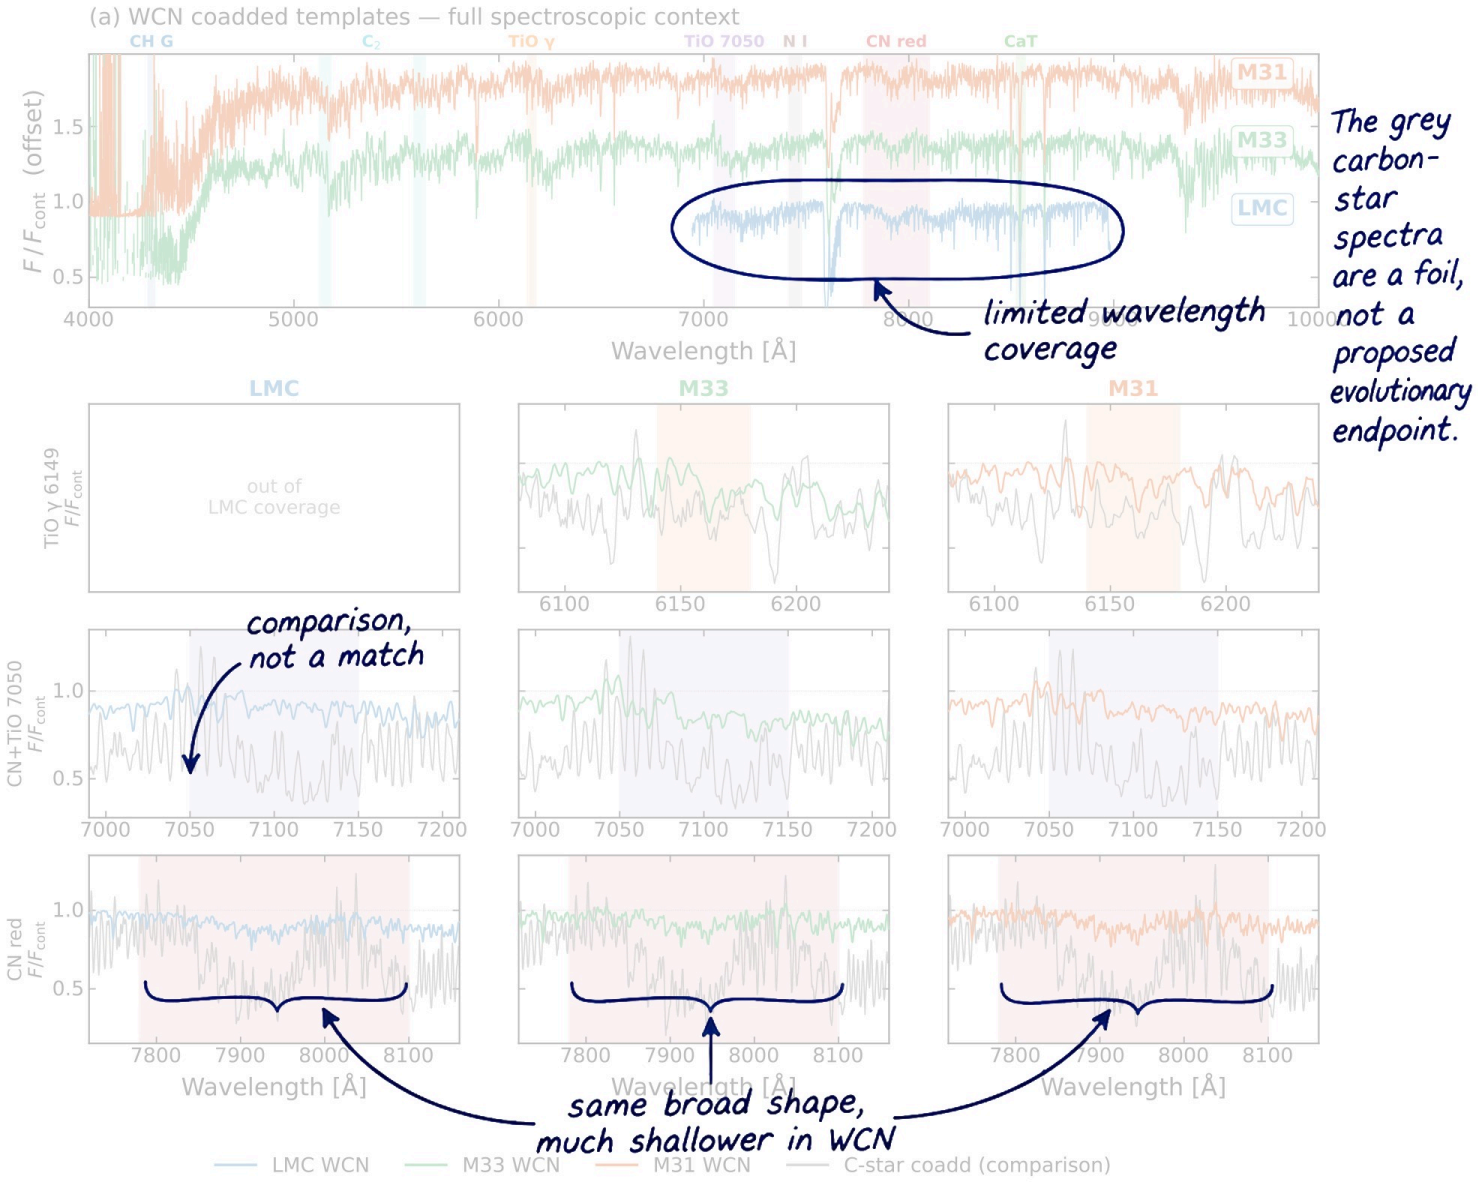


Figure 1. Observed weak-CN and carbon-star coadded templates. The top panel shows the three weak-CN coadds on a common pseudo-continuum scale; the 3×3 zooms compare each galaxy's weak-CN coadd (coloured) with its carbon-star coadd (grey) in the three diagnostic windows TiO γ 6149 Å, CN+TiO 7050 Å, and CN red. The weak-CN class is what the discovery papers noticed: a CN band that has the carbon-star morphology but only $\sim$ 20–30% of its depth, while the surrounding TiO indices retain the strength expected of a normal C/O<1 cool atmosphere. This is the spectroscopic ambiguity that the rest of the paper resolves.

95th percentile flux in a rolling 200 Å window, followed by a light Gaussian smoothing of that upper envelope. The broad window is intentional: in cool RSGs there is no true line-free continuum in the optical molecular forest (Gustafsson et al. 2008), so the observable must be a pseudo-continuum index rather than a classical isolated-line EW. The typical RSG molecular bandhead spacing of several tens of Å requires the wider window adopted here. For a feature window $[\lambda_1, \lambda_2]$ we then measure

$$\mathrm{EW} = \int_{\lambda_1}^{\lambda_2} \left[1 - \frac{F(\lambda)}{F_c(\lambda)}\right] d\lambda, \qquad (1)$$

using the finite pixels in the window. Positive EW therefore corresponds to absorption below the pseudo-continuum.

The equally important technical point is resolution consistency. The synthetic spectra are computed at high resolution, where inter-line gaps in the molecular forest

The shallow CN resemblance motivated the class; the surviving TiO is the decisive counterclue.



Matching spectral resolution before continuum fitting is essential because smoothing and normalization do not commute.

remain visible; the observations are at lower resolution, where those gaps are blended into the band profiles. Because the continuum estimate in Equation 1 is nonlinear, convolution and normalization do not commute. We therefore convolve every model spectrum to the relevant observational resolution before applying the continuum and EW measurement: $R \simeq 7000$ for the LMC/HYDRA spectra and $R \simeq 2500$ for the M31/M33 DEIMOS spectra. The observed spectra are processed at their native resolution.

If we normalized first, unresolved molecular gaps would create a systematic model-data mismatch.

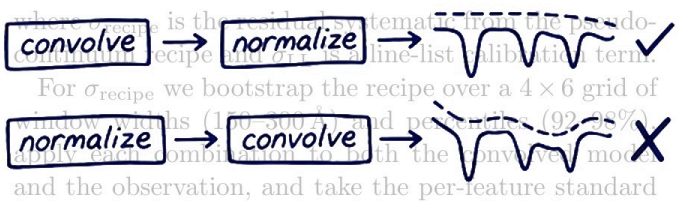
Each synthetic spectrum is therefore processed through the same sequence as the data: the available synthetic wavelength pieces are merged, the merged spectrum is convolved from the native synthesis resolution to the observation's resolution, and only then is the pseudo-continuum estimated and the EW integrated. The LMC HYDRA spectra are kept at their native $R \simeq 7000$ so that model and data are compared at matched effective resolution.

The wavelength windows and adopted uncertainty floors are listed in Table 1. The sideband definitions used for bookkeeping are not used as a local straight-line continuum in the final EW; the pseudo-continuum in Equation 1 is measured from the full local spectrum. We keep the conventional name "TiO 7050" for the 7050–7150 Å index, but in these cool-RSG spectra it should be read as a CN+TiO blend rather than a pure TiO band. Counting transitions in the CN A–X red-system line list of Brooke et al. (2014) shows that the 7000–7200 Å window contains $\sim 21{,}000$ CN lines with $\sum gf \approx 193$, while the canonical CN red window at 7780–8100 Å contains $\sim 24{,}000$ lines with $\sum gf \approx 259$. The CN contribution to "TiO 7050" is therefore at $\sim 75\%$ of the canonical CN-red level, so the index is sensitive to both molecular species and to abundance changes in the C+N reservoir.

The $\sigma_{\rm EW}$ values in Table 1 are systematic floors rather than formal photon-noise errors — the high-quality coadds aggregate of order 10^2 stars per galaxy and have negligible photon noise per pixel at the coadd level. We construct them in quadrature from two terms, each measured directly:

$$\sigma^2_{\rm EW} \;=\; \sigma^2_{\rm recipe} \;+\; \sigma^2_{\rm LL}, \qquad (2)$$

where $\sigma_{\rm recipe}$ is the residual systematic from the pseudo-continuum recipe and $\sigma_{\rm LL}$ is a line-list calibration term.



For $\sigma_{\rm recipe}$ we bootstrap the recipe over a 4×6 grid of window widths (150–300 Å) and percentiles (92–98%), apply each combination to both the convolved model and the observation, and take the per-feature standard deviation of the resulting ($\mathrm{EW_{model}} - \mathrm{EW_{obs}}$) distribution. This isolates the recipe drift that does not cancel in the model-vs-data comparison. $\sigma_{\rm recipe}$ is largest for the LMC HYDRA coadd (4.7 Å for CN red), where the higher native resolution ($R \simeq 7000$) leaves more line structure for the rolling-percentile filter to land on, and falls to 0.1–1.7 Å across all features for the lower-resolution M31/M33 DEIMOS coadds.

For $\sigma_{\rm LL}$ we adopt a 20% relative line-list calibration term, $\sigma_{\rm LL} = 0.2\,|\mathrm{EW_{obs}}|$. The CN A–X oscillator strengths of Brooke et al. (2014), the Schwenke TiO list, and the Ram C_2 Swan opacity are calibrated primarily on Galactic stellar samples; the 20% fraction encompasses oscillator-strength uncertainty, NLTE departures for atomic features (the N I and Ca II lines in particular), and local-blending residuals at the host metallicities. Combining the two terms gives the per-galaxy, per-feature floors of Table 1. $\sigma_{\rm EW}$ is dominated by $\sigma_{\rm recipe}$ for the LMC's high-resolution CN red and TiO 7050 indices, and by $\sigma_{\rm LL}$ everywhere else.

Our "continuum" is deliberately a reproducible upper envelope; these cool spectra have no honest line-free continuum.

The measured EWs are given in Table 2. The principal contrast that defines the weak-CN class is the CN red index: the carbon-star coadds give $\simeq 102$–104 Å across all three galaxies, while the weak-CN coadds give 22.6–28.7 Å, a factor of three to four shallower. The TiO and Ca II indices show the opposite behaviour expected of C/O<1 atmospheres relative to C-stars; the TiO 7050 Å EW in the carbon-star coadds is also large at $\simeq 38$–40 Å, but in that window the carbon-star value is dominated by the blue wing of the C_2 Swan band rather than by genuine TiO absorption.

3. MODEL GRID

We compute self-consistent cool-RSG model atmospheres and synthetic spectra using `pykurucz` (Kim & Ting 2026), a pure-Python re-implementation of the Kurucz ATLAS12 atmosphere iteration and SYNTHE line-synthesis pipeline (Kurucz 2005). The grid fixes $\log g = 0$ and microturbulence at $2\,\mathrm{km\,s^{-1}}$, and spans host metallicity, $T_{\rm eff}$, [α/Fe], and C/N surface offsets (Table 3). The main synthesis covers the red optical windows used for CN, TiO 7050 Å, N I, and CaT. Additional M31/M33 synthesis extensions cover the C_2 Swan windows as C-sensitive sanity checks.

The relevant distinction is that this is an ATLAS12-style [illegible] rather than

The uncertainty floor measures modelling and recipe limitations, not the tiny photon noise of the coadds.

[illegible] repeat the joint fit with $\sigma_{\rm LL}$ set to 10%, 15%, 25%, and 30% [illegible] separate fractions for atomic and molecular features. The inferred $\Delta C + \Delta N$ shifts by at most 0.2 dex in any galaxy and the location of the χ^2 minimum is essentially unchanged. The *absolute* value of the joint EW χ^2, however, is not similarly invariant — $\sigma_{\rm LL}$ enters the denominator, so smaller $\sigma_{\rm LL}$ inflates $\chi^2/N_{\rm feat}$ proportionally. We therefore use the χ^2 surface to locate and rank grid points (where it is approximately $\sigma_{\rm LL}$-invariant) and read the absolute goodness-of-fit from the per-feature σ residuals (Section 5.2) rather than from $\chi^2/N_{\rm feat}$ itself.

The error budget and atmosphere grid are built around the real limitation: imperfect molecular modelling, not counting statistics.

Table 1. EW index definitions and fit uncertainty floors. σ_{EW} values are constructed in quadrature from a per-(galaxy, feature) bootstrap of the pseudo-continuum recipe and a 20% line-list calibration term (Equation 2). TiO 7050 Å is a CN+TiO blend so it is used in the χ^2 fit only for the LMC (where TiO γ is out of HYDRA coverage); for M31 and M33 the cleaner TiO γ anchor is used in the fit while TiO 7050 Å is retained as a diagnostic comparison in Figure 3.

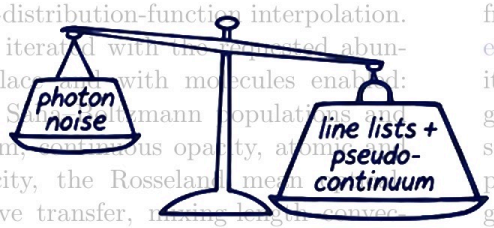
Index	Window (Å)	σ_{EW} (Å)			Used for
		LMC	M33	M31	
C_2 Swan 5165 Å	5125–5180	···	3.4	3.4	M31, M33
C_2 Swan 5636 Å	5585–5640	···	2.6	2.5	M31, M33
TiO γ 6149 Å	6140–6180	···	1.2	0.9	M31, M33
TiO 7050 Å(CN+TiO blend)	7050–7150	2.6	3.0	2.1	LMC fit only (Fig. 3 for M31, M33)
N I 7442 Å	7415–7475	0.7	1.0	0.9	all
CN red	7780–8100	7.4	4.5	4.8	all
CaT 8542 Å	8525–8568	1.6	1.6	1.5	all

Table 2. Observed equivalent widths from the current measurement pipeline. Per-(galaxy, feature) σ_{EW} floors used in the fits are given in Table 1. Blank entries are outside the wavelength coverage or not used in the fit.

Galaxy	Class	C_2 5165 Å	C_2 5636 Å	TiO γ	TiO 7050 Å (Å)	N I	CN red	CaT
LMC	WCN	···	···	···	8.8	3.4	28.7	7.8
LMC	C star	···	···	···	38.8	16.4	102.2	14.6
M33	WCN	15.4	9.9	5.5	14.7	5.0	22.6	8.1
M33	C star	42.4	31.8	11.2	38.1	16.1	103.7	15.0
M31	WCN	14.6	9.3	4.0	10.5	4.5	24.0	7.4
M31	C star	39.7	32.8	11.6	39.6	17.0	103.7	15.0

factor-of-several contrast

an ATLAS9 opacity-distribution-function interpolation. Each atmosphere is iterated with the adjusted abundance pattern in place of with molecules enabled: the iteration solves Saha–Boltzmann populations and molecular equilibrium, continuous opacity, atomic + molecular line opacity, the Rosseland mean on a depth scale, radiative transfer, mixing-length convection where unstable, and temperature corrections until a converged Kurucz-format atmosphere is produced. The SYNTHE step then synthesizes the emergent spectrum on a high-resolution wavelength grid using the same classes of opacity data as the Kurucz pipeline, including GFALL atomic lines, Schwenke TiO, and the diatomic molecular lists relevant here (CN, CO, C_2, CH, OH, MgH, FeH, and others). End-to-end validation against the original Fortran ATLAS12+SYNTHE pipeline gives median normalized-flux differences below 0.005% and 99th-percentile differences below 0.3% over representative test cases (Kim & Ting 2026), well below the EW systematic floors adopted here.



The reason we do not use pre-computed scaled-solar grids (such as Aringer et al. 2016) as the engine is that, at $\log g \simeq 0$, the molecular contribution to the Rosseland mean opacity at the formation depth is a substantial fraction of the total in MARCS calculations (Gustafsson et al. 2008) compared to a few per cent at K-giant gravities. Shifting C, N, or [α/Fe] at RSG gravities therefore changes the line opacity at each substantially, which back-reacts on the temperature–pressure structure through line blanketing. At K-giant gravities the molecular fraction is small enough that an abundance perturbation can be applied at fixed atmosphere structure — this is the operating regime of typical CN-based abundance work such as Ting et al. (2018) — but at RSG gravities the back-reaction can move the formation-layer temperature by tens of K and is not negligible.



`pykurucz` accepts arbitrary per-element abundance overrides and iterates the atmosphere to radiative–convective equilibrium under the perturbed opacity at every grid point, so the model T–P profile responds to the imposed CNO and [α/Fe] offsets and the resulting EWs reflect that back-reaction rather than a fixed-structure approximation.

The C and N offsets are defined relative to each host's scaled-solar baseline, not relative to a common absolute abundance. This convention matters for the interpretation of $\Delta C + \Delta N$: it is a combined C+N surface

Calling 7050 Å "TiO" is historical shorthand; in these stars it is substantially a CN+TiO blend.

The 20% line-list term is conservative, but changing it does not move the preferred grid point much.

We iterate each atmosphere after changing abundances; fixed atmospheric structure is unsafe at red-supergiant gravity.

CN measures a C-times-N reservoir, so these spectra constrain the sum of logarithmic C and N offsets—not each separately.

abundance offset at fixed host metallicity, not a statement that the three galaxies have the same initial CNO pattern. The [α/Fe] axis is included because oxygen-bearing opacity and TiO strength can change both the molecular equilibrium and the atmospheric temperature structure. We use the full grid in the fit rather than imposing a temperature–[α/Fe] prior at the outset.

CO consumes most carbon first; CN forms from what remains, making oxygen part of the background bookkeeping.

4. MOLECULAR EQUILIBRIUM CONTEXT

With the observed grid in hand, we now use it to understand why the joint EW fit can separate $T_{\rm eff}$ from surface CNO in the cool-RSG regime, and what surface variables the optical CN and TiO indices can and cannot constrain. Figure 2 is the physical backdrop for that argument: the four representative cool-RSG models visualise the $T(\tau_R)$ structure and the CN and TiO number densities on a common Rosseland-depth axis centred at $\log_{10}\tau_R = -0.4$, with a half-width of 0.25 dex. That band is where the optical molecular pseudo-continuum in these cool, low-gravity atmospheres emerges — near the photosphere rather than at the outer boundary or in the deep diffusion regime — so the T–P structure and the molecular number densities relevant to the observed EWs can be read together at the same depth.

The useful physical scale is set by the molecular dissociation temperatures. TiO has bond dissociation energy $D_0 \simeq 6.9\,$eV and crosses its formation-layer transition around $T \simeq 3800\,$K, while CN has $D_0 \simeq 7.7\,$eV and crosses near $T \simeq 4300\,$K — a $\sim 500\,$K offset that puts the two molecules out of phase across the cool-RSG plateau. Across the formation band, a 4500 K model is relatively molecule-poor; by 4200 K CN has grown substantially and TiO is beginning to grow faster; by 4000 K TiO has grown disproportionately while CN is still present rather than erased. A change in $T_{\rm eff}$ can therefore move the TiO indices strongly without erasing the red CN band, and an atmosphere can show weak but recognisable CN while keeping non-carbon-star RSG morphology — which is the physical reason the C-star comparison in Figure 1 is useful for identifying the feature but misleading if interpreted as evidence for C-star chemistry.

The same molecular-equilibrium argument fixes what the CN band can and cannot tell us about the surface chemistry. CN forms from the C and N populations that remain after CO sequesters most of the available carbon, and the column density of the molecule scales as $n_{\rm C} \cdot n_{\rm N}$ to first approximation. A CN-only EW measurement at fixed $T_{\rm eff}$ therefore constrains the product $n_{\rm C} \cdot n_{\rm N}$ rather than $n_{\rm C}$ and $n_{\rm N}$ separately; in abundance space this is the combined surface offset $\Delta[{\rm C/H}] + \Delta[{\rm N/H}]$ along a degeneracy ridge in $(\Delta{\rm C}, \Delta{\rm N})$.

Any individual best-fit C and N pair should therefore be read as one point on a ridge.

TiO does not break the degeneracy because it is essentially insensitive to the C and N reservoirs at fixed metallicity, and the optical C_2 Swan features, which would in principle anchor ΔC independently, are at low-resolution sensitivity floors that limit their leverage (Section 6.4). $\Delta{\rm C} + \Delta{\rm N}$ is therefore the surface variable the optical EW fit measures cleanly, and the individual ΔC and ΔN values returned by the formal χ^2 argmin should be read as a representative point on a broader allowed locus rather than as unique abundance measurements.

TiO separates temperature from chemistry, but it cannot split carbon from nitrogen.

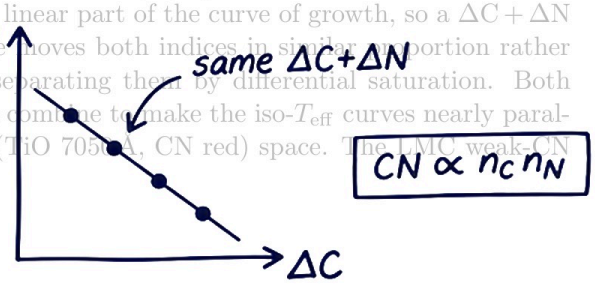
5. RESULTS

5.1. *Where the Diagnostic Discrimination Lives*

Before running the formal joint χ^2 fit, we examine where the observed weak-CN coadds fall in the model EW plane and how cleanly the data separate temperature from CNO chemistry there. The orthogonality argument of Section 4 guarantees that this separation exists in principle; whether it is operative in any specific host depends on how the iso-$T_{\rm eff}$ and iso-($\Delta{\rm C}+\Delta{\rm N}$) families project onto the observed (TiO 7050 Å, CN red) plane (Figure 3).

In the M33 and M31 grids, both at near-solar [α/Fe], the iso-$T_{\rm eff}$ and iso-($\Delta{\rm C} + \Delta{\rm N}$) families fan out, and moderate positive values of $\Delta{\rm C} + \Delta{\rm N}$ ($\simeq 0$ to $+0.3$ dex) occupy the same part of the plane as the observed weak-CN coadds. TiO γ, which is less affected by CN contamination than the 7050 blend (Section 2), provides a consistency check in the bottom row: the M31/M33 placement is not solely an artefact of the CN+TiO 7050 Å blend. Figure 3 fixes only this qualitative two-feature placement; the formal multi-feature χ^2 best-fit grid points of Section 5.2 use additional indices (CaT, N I, and the C_2 Swan windows) and may land at different $\Delta{\rm C} + \Delta{\rm N}$ values when those features pull the joint fit toward grid boundaries.

The LMC grid behaves differently for a physical reason that is itself informative. At [Fe/H] $\simeq -0.5$ and [α/Fe] $\simeq -0.3$, its absolute oxygen, titanium, and other α-element abundances are lower by roughly 0.8 dex than at M31 metallicity. In a C/O<1 atmosphere the residual non-CO carbon is the principal feedstock for CN, and a lower O abundance reshapes the C/O ratio and the molecular equilibrium in the direction that flattens the differential response of CN and TiO to $\Delta{\rm C} + \Delta{\rm N}$ at fixed $T_{\rm eff}$. Lower [Fe/H] also keeps both species closer to the linear part of the curve of growth, so a $\Delta{\rm C} + \Delta{\rm N}$ change moves both indices in similar proportion rather than separating them by differential saturation. Both effects combine to make the iso-$T_{\rm eff}$ curves nearly parallel in (TiO 7050 Å, CN red) space. The LMC weak-CN



The roughly 500 K separation between TiO and CN dissociation transitions creates the weak-CN temperature window.

Table 3. Synthetic grid and fitted feature sets. Host [M/H] values are from Harris & Zaritsky (2009); Zaritsky et al. (2004) for the LMC, Magrini et al. (2009) for M33, and Dalcanton et al. (2012) for M31.

Quantity	Values
$T_{\rm eff}$	3800–4600 K, 100 K spacing
Δ[C/H]	$-0.6, -0.5, -0.4, -0.2, -0.1, 0.0, +0.1, +0.2, +0.3$
Δ[N/H]	$0.0, +0.1, +0.2, +0.3, +0.5, +0.6, +0.7, +0.8$
$\Delta[\alpha/{\rm Fe}]$	-0.3, 0.0, +0.3
Host [M/H]	LMC -0.5; M33 -0.3; M31 ~ 0
Observation R	LMC 7000; M33/M31 2000
LMC fit windows	CN red, TiO 7050, CaT, N I
M33/M31 fit windows	CN red, TiO γ, CaT, N I, C_2 5636, C_2 5165

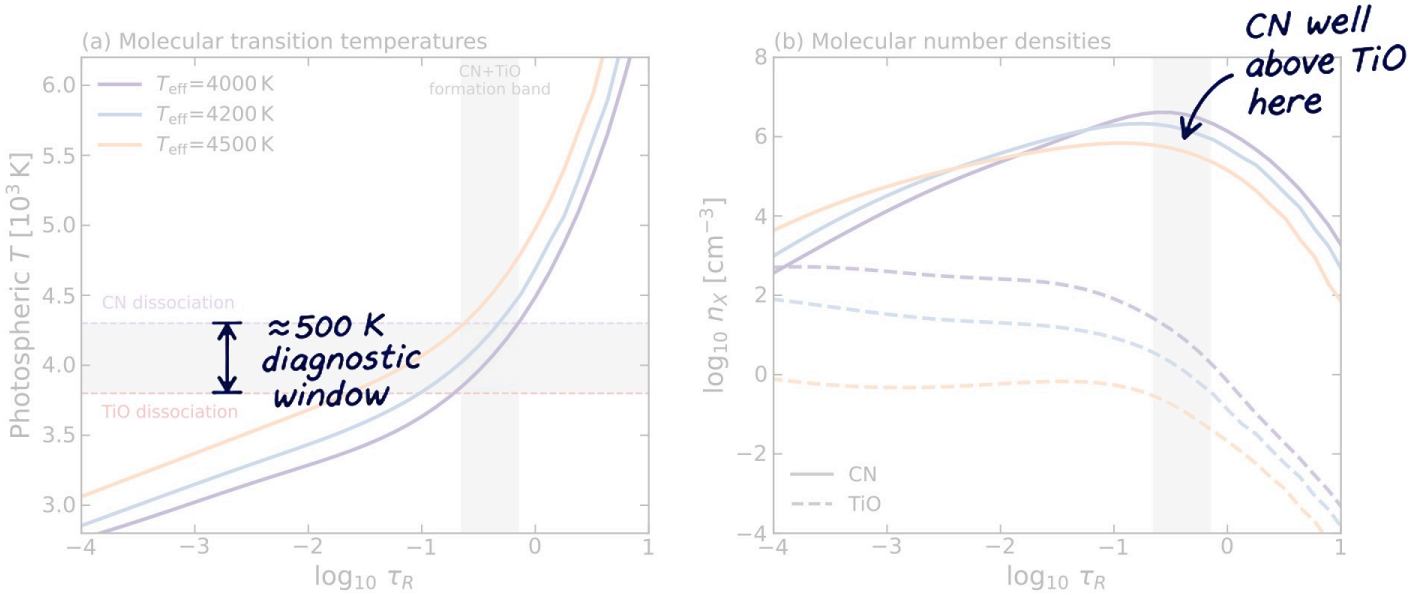


Figure 2. Atmospheric physics underlying the weak-CN diagnostic. Left: $T(\tau_R)$ for four representative cool-RSG models, with the TiO and CN dissociation transitions marked and a common Rosseland-depth reference band at $\log\tau_R \simeq -0.4$. Right: CN and TiO number densities at the same models on the same depth axis. The figure shows the orthogonal response of the two molecules: as $T_{\rm eff}$ drops through the WCN range, TiO grows steeply while CN grows gently, whereas a CNO surface enhancement acts directly on CN through the C and N reservoirs but barely shifts TiO at fixed $T_{\rm eff}$. Temperature and surface chemistry therefore move the two indices in distinct directions, which is what allows the joint χ^2 fit of Section 5 to break the $T_{\rm eff}$–CNO degeneracy that a CN-only measurement cannot.

The shaded depth band is where the optical molecular pseudo-continuum actually emerges, so that is where comparisons matter.

data point still falls on the locus, so the same atmospheric physics applies at lower [Fe/H] and [α/Fe], but the spectral EW features are less able to turn the LMC

Cooling boosts TiO far more dramatically while CN remains available—exactly the mixed morphology we observe.

For the LMC, the sum includes CN red, TiO 7050 Å, Ca II 8542 Å, and CN 7442 Å. For M31 and M33, where the blue coverage is available, the sum includes CN red, TiO γ, CaT, N I 8680 Å, CN 7442 Å, and the two C_2 Swan windows. TiO 7050 Å enters the M31 and M33 locus diagram of Figure 3 as a diagnostic comparison, but TiO γ is the cleaner TiO anchor in the formal fit because the 7050 window is a CN+TiO blend. We have nonetheless tested adding TiO 7050 Å to the M33 and M31 χ^2 as a seventh feature: the inferred $\Delta{\rm C} + \Delta{\rm N}$ stays unchanged for M33 and shifts by +0.1 dex for M31, while the best-fit $T_{\rm eff}$ moves by at most 100 K, so the qualitative result is robust to the choice of whether to

5.2. *Joint χ^2 Fits and Best-Fit Spectra*

The qualitative placement in Figure 3 motivates a quantitative joint EW χ^2 over the per-galaxy feature set in Tables 3 and 1,

$$\chi^2 = \sum_k \left(\frac{{\rm EW}_k^{\rm model} - {\rm EW}_k^{\rm obs}}{\sigma_k} \right)^2 . \qquad (3)$$

hotter	CN ✓	TiO faint
formation depth	CN ✓	TiO bold
cooler	CN ✓	TiO bold

T ↓

The CN-TiO geometry separates temperature from C+N well in M31/M33, but largely collapses at LMC composition.

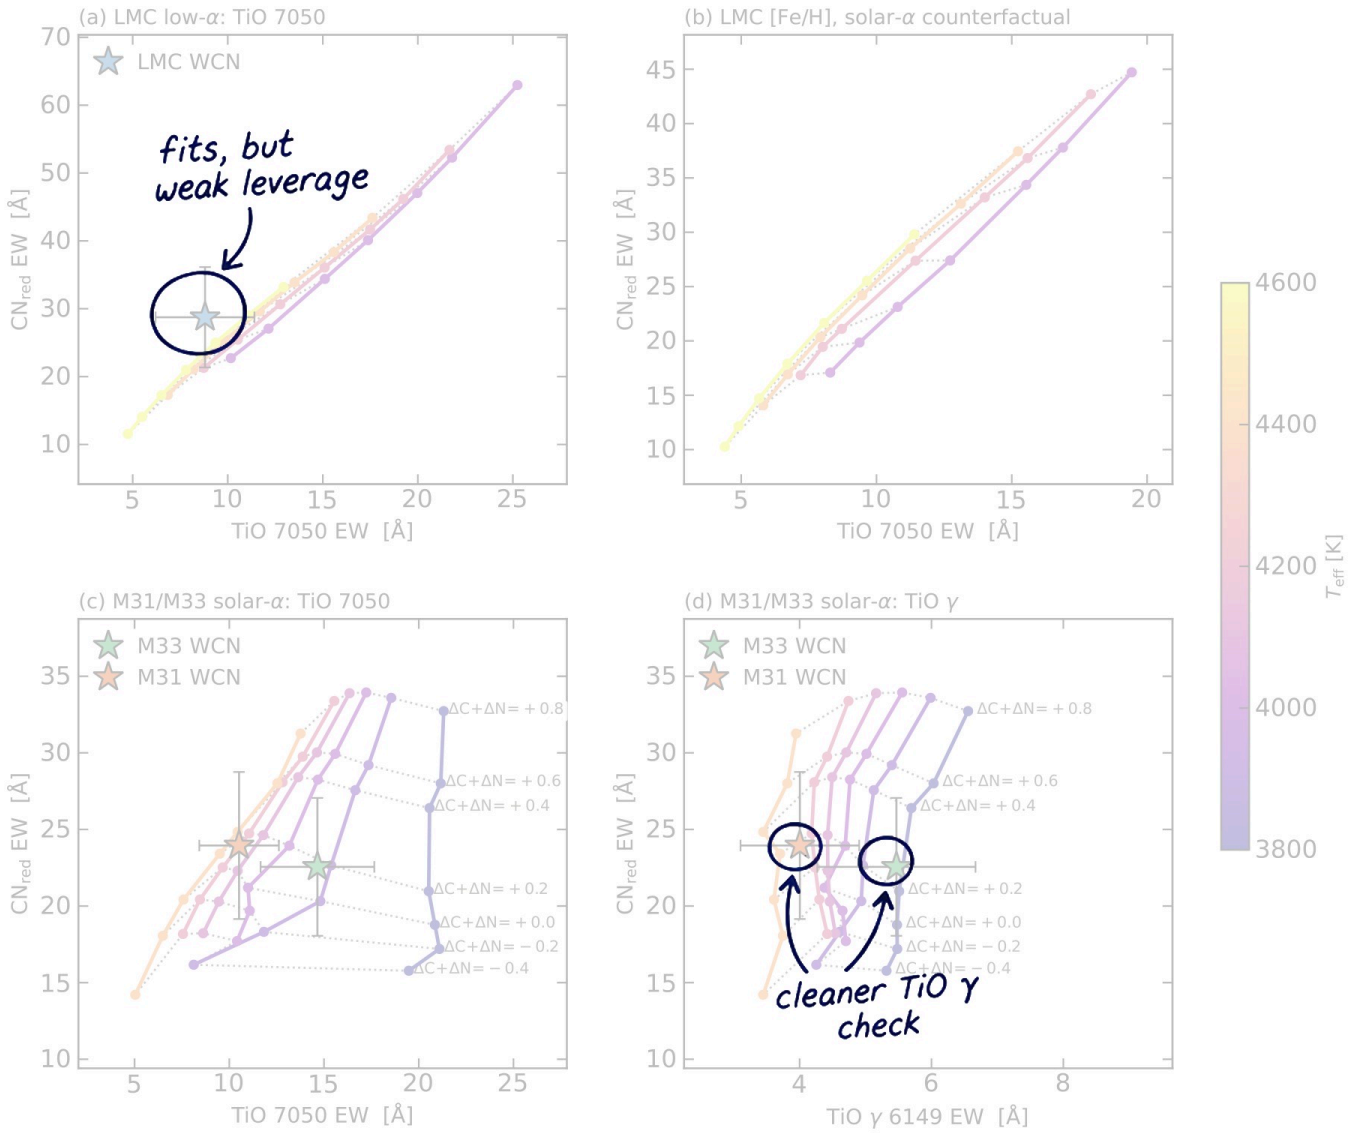


Figure 3. Where the diagnostic discriminates parameters in the EW plane. Symbols are the observed weak-CN coadds; [illegible] tracks, coloured by [illegible] and [illegible] with [illegible] obtained by averaging the model EWs over the grid points whose $\Delta[C/H] + \Delta[N/H]$ falls within [illegible] dex of the fixed sum. Dotted curves [illegible] by the same averaging [illegible] and sampled at -0.2, 0.0, [illegible], $+0.4$, $+0.6$, and $+0.8$ dex; the top panels omit those labels for readability. In the M33 and M31 solar-α TiO 7050 Å plane the iso-$T_{\rm eff}$ and iso-$(\Delta C + \Delta N)$ families fan out and moderate positive $\Delta C + \Delta N$ values overlap the weak-CN data; the TiO γ plane provides a cleaner TiO consistency check with the same qualitative placement. In the LMC low-α grid, lower [O/Fe] and lower [Fe/H] keep CN and TiO closer to the linear curve-of-growth and the families collapse, so the LMC weak-CN data point falls on the locus but the optical indices are a poor discriminator of the CNO-processing direction.

At LMC composition, lower metallicity and oxygen keep CN and TiO closer to proportional: the two indices offer little independent abundance leverage.

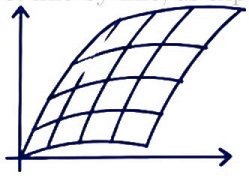
include the blended index. The grid search minimises [illegible] directly over all converged grid points, the [illegible] are reported in Table 4.

Figure 4 shows the same result in continuum-normalised flux space. The LMC fit uses the HYDRA-resolution red spectrum, while the M31 and M33 fits use the C_2-extended spectra [illegible] The matches are not perfect line-by-line, as expected for low- and medium-resolution molecular spectra, but the observed weak-CN morphology and the fitted diagnostic [illegible]

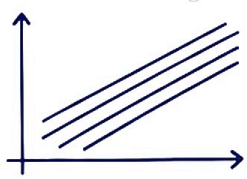
M31/M33 tracks separate temperature from the combined C and N offset more clearly. A good molecular fit alone does not determine rotation.

Three results follow directly from Table 4 and Figure 4. The absolute χ^2 values in column 6 of Table 4 depend on the adopted $\sigma_{\rm LL}$ (Section 2) and should therefore be read as a relative ranking summary across grid points rather than as an absolute goodness-of-fit metric. The absolute goodness-of-fit is set instead by the

M31/M33: leverage



LMC: degeneracy



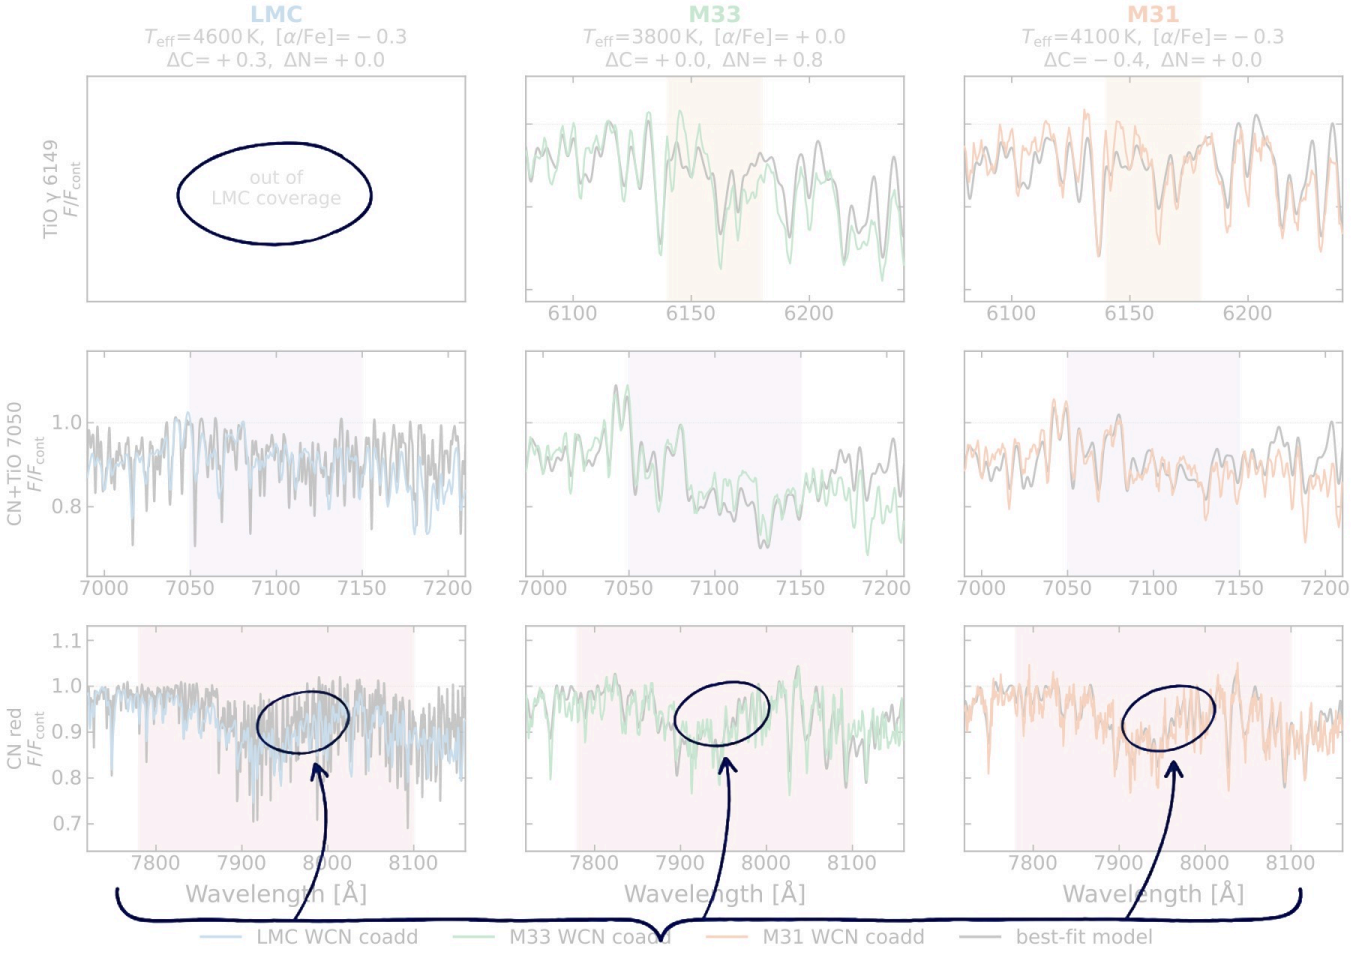
Ordinary cool-red-supergiant models reproduce every host's weak-CN morphology within the adopted systematic precision.

Table 4. Joint χ^2 best-fit grid points per galaxy. Columns are the grid parameters at the χ^2 minimum; χ^2 over the per-galaxy feature set (Table 1) and number of features used; and the single worst per-feature residual in units of the adopted $\sigma_{\rm EW}$.

Galaxy	$T_{\rm eff}$ (K)	[α/Fe]	Δ[C/H]	Δ[N/H]	$\chi^2/N_{\rm feat}$	Worst-fit feature (residual)
LMC	4600	−0.3	+0.3	0.0	1.20/4	N 1 7442 Å (+0.80σ)
M33	3800	0.0	0.0	+0.8	2.83/6	C_2 Swan 5165 Å (+1.32σ)
M31†	4100	−0.3	−0.4	0.0	1.29/6	C_2 Swan 5636 Å (−1.02σ)

all no worse than about 1.3σ

† The M31 offset is anchored by the CaT 8542 Å line: dropping CaT from the joint fit moves $\Delta C + \Delta N$ from −0.4 to +0.3 dex, with the best-fit temperature rising to $T_{\rm eff} = 4500$ K and [α/Fe] = 0.0 (both consistent with the literature M31 disc); see Section 6.2.



The listed C and N values mark points on a degenerate ridge; their combined constraint is more robust.

We test resolution-matched band strengths. A good integrated fit need not match every spectral pixel.

No carbon-rich atmosphere is needed to obtain these spectra.

Figure 4. Observed weak-CN coadds (coloured) compared to the joint best-fit synthetic spectra (black) of the grid points of Table 4, in the three primary diagnostic windows. Each model is processed through the same resolution-matched pseudo-continuum and band-strength measurement as the data. The CN, CN+TiO 7050 Å, and TiO γ morphologies are reproduced at the adopted precision in all three galaxies by ordinary cool-RSG models of the self-consistent grid, without invoking a carbon-star-like abundance pattern. The LMC TiO γ panel is blank because TiO γ 6149 Å lies outside the 6850–9150 Å GTIO/Hydra coverage of the LMC weak-CN spectra.

per-feature σ residuals at the best-fit grid points listed in Table 3 in units of the adopted $\sigma_{\rm EW}$ floor. First, the best-fit grid points lie well within the cool-RSG portion of the grid for all three galaxies (Section 5.3): 3800 to 4600 K, [α/Fe] matching the literature value of each host to the grid resolution, and surface CNO offsets at the level of the first-dredge-up prediction (Section 6).

No fitted feature is discrepant from the best model by more than about 1.3 $\sigma_{\rm EW}$ floor (Table 5); the largest residual is C_2 Swan 5165 Å in M33

Residual checks show that integrated equivalent widths suppress line-list noise while preserving coherent abundance shifts.

Table 5. Per-feature residuals at the joint χ^2 best-fit grid points of Table 4, in units of the per-(galaxy, feature) $\sigma_{\rm EW}$ floors of Table 1. Positive entries indicate that the model overpredicts the observed EW (Equation 3). Blank entries are not in the fitted feature set for that galaxy (Table 3); the LMC uses TiO 7050 Å in the fit while M33 and M31 use TiO γ 6149 Å as the cleaner TiO anchor (Section 5.2).

Galaxy	C_2 5165 Å	C_2 5636 Å	TiO γ 6149 Å	TiO 7050 Å	N I 7442 Å	CN red	CaT 8542 Å
			Residual (model − obs)/$\sigma_{\rm EW}$				
LMC	···	···	···	$-0.05\,\sigma$	$+0.80\,\sigma$	$-0.72\,\sigma$	$-0.19\,\sigma$
M33	$+1.32\,\sigma$	$-0.92\,\sigma$	$-0.49\,\sigma$	···	$+0.01\,\sigma$	$-0.08\,\sigma$	$-0.10\,\sigma$
M31	$-0.01\,\sigma$	$-1.02\,\sigma$	$+0.24\,\sigma$	···	$+0.12\,\sigma$	$-0.38\,\sigma$	$+0.21\,\sigma$

every fitted feature within ±1.32 σ

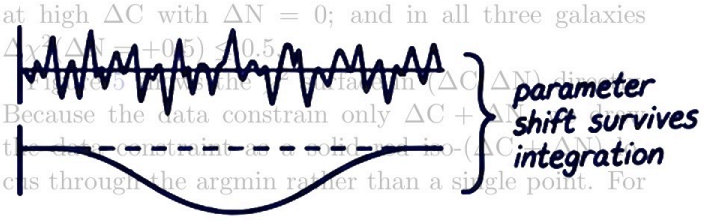
at $+1.32\sigma$. In particular, the C_2 Swan residuals in M31 and M33, which are the windows that would most directly distinguish a carbon-rich abundance pattern, are all within $\sim 1.3\sigma$ of zero, consistent with C/O<1 atmospheres at the inferred surface composition.

Second, the per-galaxy best-fit [α/Fe] recovered by the joint EW fit is consistent, at the grid spacing of 0.3 dex, with the literature stellar-population values of each host. The LMC fit returns [α/Fe] = −0.3, consistent with the sub-solar α-enhancement measured for LMC RSGs (Davies et al. 2015). The M33 fit returns [α/Fe] = 0.0, consistent with the near-solar disk α/Fe inferred from M33 PNe and H II regions by Magrini et al. (2009). The M31 fit returns [α/Fe] = −0.3 at the grid minimum with $\Delta\chi^2 \simeq 2$ at the [α/Fe] = 0 slice, so [α/Fe] is constrained at roughly the 1.5σ level and is consistent with the near-solar M31 disk value within that uncertainty.

The largest practical uncertainties are systematic floors, not formal coadd errors.

Third, the surface CNO is constrained along the C+N degeneracy ridge in $(\Delta C, \Delta N)$ rather than as independent offsets. Reading the individual best-fit grid points literally would suggest $(\Delta C, \Delta N) \simeq (+0.3, 0.0)$ for the LMC, $(0.0, +0.8)$ for M33, and $(-0.4, 0.0)$ for M31, but as Section 4 showed, the optical CN bands respond to the joint C+N reservoir, so the χ^2 surface in $(\Delta C, \Delta N)$ is elongated along the ridge $\Delta C + \Delta N = \text{const}$ and the individual values are not separately measurable.

The robust observable is therefore the combined offset $\Delta C + \Delta N$, which is +0.30, +0.80, and −0.40 dex for the LMC, M33, and M31 best-fit grid points respectively. Extending the grid to Δ[C/H] $\in [-0.6, +0.3]$ and Δ[N/H] $\in [0.0, +0.8]$ confirms that the ridge is shallow and broad: the formal argmin wanders by 1–2 grid points along the ridge depending on EW noise; $\Delta\chi^2$ along the ridge stays $\lesssim 0.9$ across the full extended span; M33 hits the high-ΔN boundary while LMC sits at high ΔC with $\Delta N = 0$; and in all three galaxies $\Delta\chi^2(\Delta N = +0.5) > 0.5$.



parameter shift survives integration

Pixel residuals look busy because individual molecular lines are imperfect; their broad-band integral is much better behaved.

Figure 5 shows the χ^2 surface in $(\Delta C, \Delta N)$ directly. Because the data constrain only $\Delta C + \Delta N$, the χ^2 surface is a valley rather than a bowl, and the 1σ contour (solid red) is elongated along the ridge. The LMC and M33 best-fit points (cyan and magenta) sit on the ridge, and the M31 best-fit point (cyan) lies on the ridge at high ΔC. The M31 slow-rotation locus (cyan, the latter passing through the PARSEC slow-rotation point) are close and both lie inside the 1σ contour. For M31 the baseline solid red locus sits below the cyan one; the dashed red locus shows where the argmin moves when CaT is dropped from the joint fit (Section 5.3 (iv)) and overlaps the cyan slow-rotation locus. The implications for rotation are taken up in Section 6.2.

5.3. *Robustness Ablations*

We test the robustness of the joint EW χ^2 formulation through five orthogonal ablations of feature choice, integration scale, and synthesis assumptions.

(i) *Sub-feature decomposition.* We split the CN red 100 Å window into four 80 Å sub-bands and recompute per-sub-band EWs, to test whether the integrated CN red EW is absorbing a within-window gradient. The rms residual across the four sub-bands at the best-fit grid point is $\sim 1.4\,\sigma$ for LMC, $\sim 0.9\,\sigma$ for M33, and $\sim 0.5\,\sigma$ for M31, and the per-sub-band residuals share sign across the four sub-windows in each galaxy. The integrated CN red EW is therefore tracking an approximately uniform CN scaling rather than a hidden sub-band gradient.

(ii) *Pixel-level residual.* At the best-fit grid point, the continuum-normalised pixel rms residual is ~ 0.02–0.10 across the diagnostic windows in all three galaxies — the order-of-magnitude precision expected for molecular spectra at these resolutions. The corresponding integrated EW residuals, listed in Table 5, stay within $\pm 1.3\,\sigma_{\rm EW}$ for every fitted feature in all three galaxies. The contrast — non-trivial per-pixel residual but small bulk-EW residual — is the empirical signature that the line-list and pseudo-continuum-recipe noise that dominates the pixel-level mismatch (Section 2) averages out under integration into the EW, which is the formal motivation for the EW χ^2 formulation made precise in ablation (iii) below.

(iii) *Pixel residual vs. parameter shift.* We compare the pixel-level and EW-level responses to a one-grid-step shift in ΔN. At the pixel level, the per-window pixel rms changes by less than 20% for LMC and less than 5% for

The sharp fit direction is perpendicular to the C+N ridge; along the ridge, many solutions remain acceptable.

The data define diagonal C+N ridges, and two apparently extreme galaxy results are each driven by one feature.

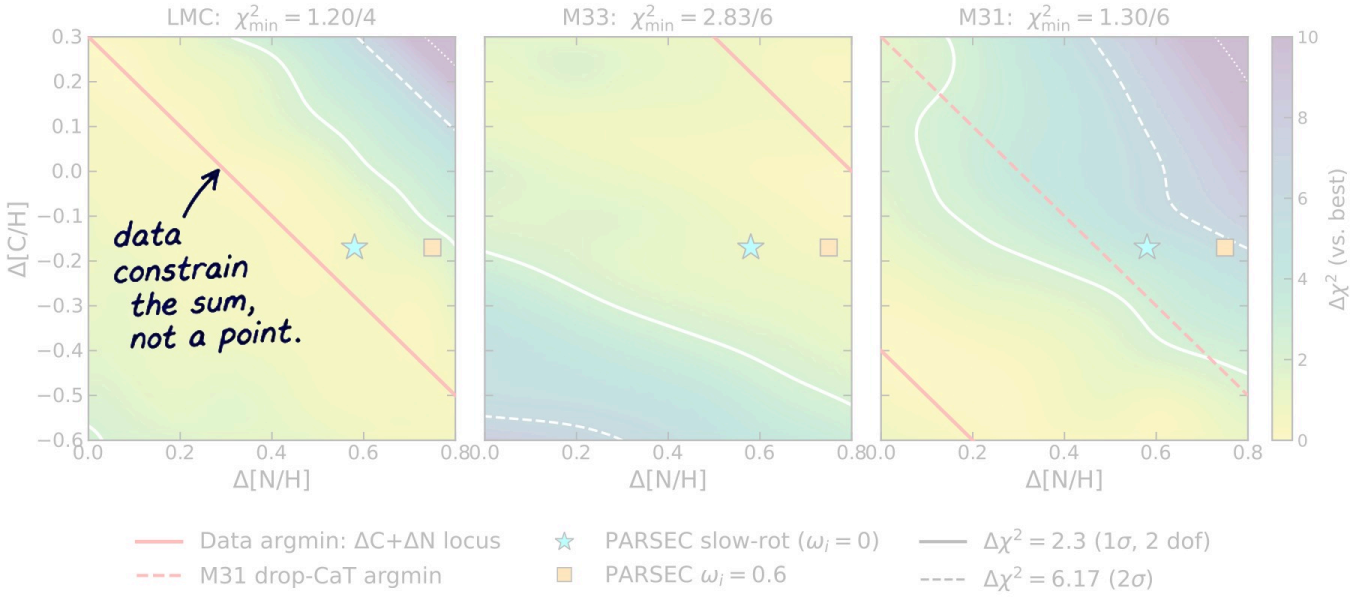
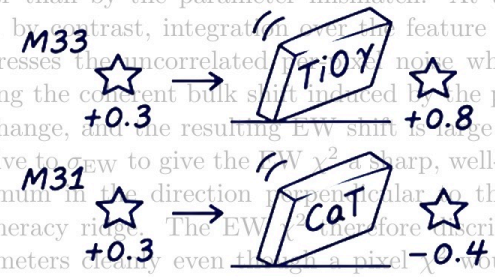


Figure 5. Profile χ^2 in the $(\Delta[\mathrm{C/H}], \Delta[\mathrm{N/H}])$ plane, marginalised over $T_{\rm eff}$ and $[\alpha/\mathrm{Fe}]$, on the extended grid. White contours: $\Delta\chi^2 = 2.3$ (1σ, 2 dof; solid) and 6.17 (2σ; dashed). PARSEC first-dredge-up predictions are shown as a cyan star ($\omega_i = 0$, $\Delta\mathrm{C} = -0.17$, $\Delta\mathrm{N} = +0.58$) and an orange square ($\omega_i = 0.6$, $\Delta\mathrm{N} = +0.75$). Because the data constrain only the combined offset $\Delta\mathrm{C} + \Delta\mathrm{N}$, we plot the data argmin as a solid red iso-$(\Delta\mathrm{C} + \Delta\mathrm{N})$ locus rather than a point. For LMC and M33 the locus and the PARSEC slow-rotation point lie within 1σ of each other. For the baseline M31 fit the locus sits at $\Delta\chi^2 \simeq 3.8$ ($\sim 2\sigma$) below PARSEC; the dashed red locus is the M31 argmin after CaT is dropped from the fit (Section 5.3 (iv)) and crosses the PARSEC point inside 1σ.

M33's high-N boundary disappears when TiO γ is removed; M31's low C+N disappears when CaT is removed.



rather than by the parameter mismatch. At the EW level, by contrast, integration over the feature window suppresses the uncorrelated pixel noise while preserving the coherent bulk shift induced by the parameter change, and the resulting EW shift is large enough relative to $\sigma_{\rm EW}$ to give the EW χ^2 a sharp, well-defined minimum in the direction perpendicular to the C+N degeneracy ridge. The EW χ^2 therefore discriminates parameters cleanly even though a pixel-χ^2 would not, which is why we adopt the EW formulation as the fitting statistic.

(iv) *Feature drop.* We refit after removing each feature in turn. For LMC the inferred $\Delta\mathrm{C} + \Delta\mathrm{N}$ stays within ± 0.2 dex of the full-fit value across all drops. Two single-feature drops matter for M33 and M31. The baseline M33 fit lands at the upper $\Delta\mathrm{N}$ grid edge ($\Delta\mathrm{N} = +0.8$ dex; Table 4); dropping TiO γ 6149 Å pulls $\Delta\mathrm{C}+\Delta\mathrm{N}$ down from +0.8 to +0.3 dex, close to the PARSEC slow-rotation value. The baseline M31 fit sits at low $\Delta\mathrm{C} + \Delta\mathrm{N} = -0.4$ dex; dropping CaT 8542 Å moves $\Delta\mathrm{C}+\Delta\mathrm{N}$ from −0.4 to +0.3 dex, matching the PARSEC slow-rotation value. These two single-feature ablations therefore identify the individual features responsible for

SEC and we take them up in Section 6.2.

(v) *Microturbulence.* Motivated by the LMC RSG measurements of Davies et al. (2015) giving $\xi \simeq 3$–$4\,\mathrm{km\,s^{-1}}$ relative to the $\xi = 2\,\mathrm{km\,s^{-1}}$ adopted here, re-synthesised the full LMC grid at $\xi = 3.5$ km s^{-1}. The LMC CN red, TiO 7050 Å, Ca II triplet, and [Ti I] EWs at $R = 7000$ shift by less than 0.05 Å in the vast majority of grid points across the $(T_{\rm eff}, \Delta\mathrm{C}, \Delta\mathrm{N})$ plane. The formal argmin and the χ^2 at the PARSEC slow-rot locus are unchanged at the $\Delta\chi^2 \lesssim 0.1$ level. The physical reason is that at $R \simeq 7000$ the CN red EW is set by the number of unsaturated rotational lines integrated across the band rather than by atomic-line saturation cores, so microturbulence broadening barely shifts the molecular bandhead EW used here.

The dominant remaining uncertainty on $\Delta\mathrm{C} + \Delta\mathrm{N}$ for any single galaxy is therefore the EW systematic floor (recipe plus line-list, Section 2), not the photometric uncertainty (which the inverse-variance budget of the coadded spectra at $\lesssim 0.1$ Å per feature, well below the ≥ 1 Å floor) and not the microturbulence choice. The best-fit parameters should accordingly be read as grid-level summaries of coadded populations at fixed $\log g$ and microturbulence, not as precision effective temperatures for indi-

That is a calibration warning, not evidence for two exotic formation channels.

Changing microturbulence barely moves these molecular-band equivalent widths at our resolution.

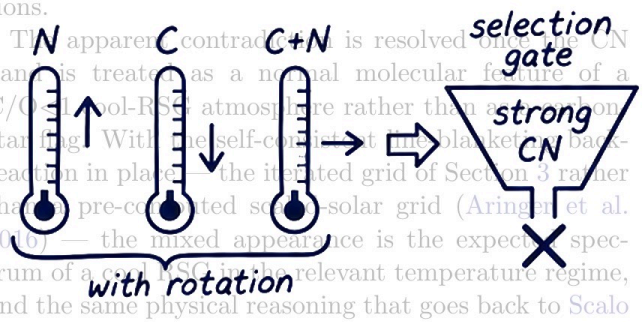
First dredge-up is fully adequate, but C+N is intrinsically a weak ruler for initial stellar rotation.

vidual stars, at the ±100–200 K level set by the grid spacing and the same systematic floors.

6. DISCUSSION

6.1. *Weak-CN Stars Are Ordinary Cool RSGs*

The per-galaxy fits of Section 5 place all three weak-CN coadds inside the ordinary cool-RSG region of the self-consistent grid, with a host-dependent rotation discrimination. The discovery papers compared weak-CN stars to carbon stars because the broad CN band was the selection clue. In classical optical classification, strong CN and C_2 Swan are associated with C-rich atmospheres, whereas TiO and the Ca II triplet are associated with non-carbon, C/O<1 cool stars. Weak-CN stars combine these signatures: weak carbon-star-like CN, but TiO/CaT and young-RSG photometric positions.



The apparent contradiction is resolved once the CN band is treated as a normal molecular feature of a C/O<1 cool-RSG atmosphere rather than a carbon-star flag. With the self-consistent line-blanketing back-reaction in place — the iterated grid of Section 3 rather than a pre-computed scaled-solar grid (Aringer et al. 2016) — the mixed appearance is the expected spectrum of a cool RSG in the relevant temperature regime, and the same physical reasoning that goes back to Scalo (1974)'s band-strength analysis across the M, MS, and S spectral sequence — the chain of cool giants whose dominant metal-oxide bands shift smoothly from TiO toward ZrO as the C/O ratio approaches unity — applies in the cool-RSG regime as well.

An independent line of evidence pointing to the same conclusion appears in the Comparison to Empirical Templates (COMET) classification framework of Guhathakurta et al. (2025) and Grion Filho et al. (2025). COMET scores each spectrum by its modified-χ^2 similarity against two empirical templates — a weak-CN coadd and a carbon-star coadd — evaluated over the 7800–8200 Å wavelength range. In the resulting $(\chi^2_{\rm wCN}, \chi^2_{\rm C})$ plane, the visually-identified carbon stars form a distinct group in the lower-right, while the remainder of the sample traces a comet-like distribution with the weak-CN coadds at the head and normal O-rich stars spreading along the tail. The COMET geometry therefore places the weak-CN class at the high-CN tip of the O-rich (C/O<1) distribution rather than as an intermediate between normal and carbon stars. Our atmospheric-modelling result — that weak CN is the warm-end molecular-equilibrium signature of an ordinary C/O<1 cool RSG — is consistent with this COMET picture: the weak-CN class is the empirical

So our rotation statement is conditional on stellar-evolution models, not a spectroscopic measurement of spin.

handle on the bright end of the normal cool-RSG sequence.

6.2. *Slow-Rotation Consistency and the Formation Channel*

PARSEC changes C and N in opposite directions, leaving their sum almost unchanged as rotation increases.

For reference, the PARSEC v2.0 (Costa et al. 2019; Nguyen et al. 2022) surface mass-fraction columns at the core-He-burning phase for 5–10 $M_\odot$ progenitors at $Z = 0.014$ give first-dredge-up offsets Δ[C/H] $\simeq -0.41$, Δ[N/H] $\simeq +0.58$ at $\omega_i = 0$, Δ[N/H] $\simeq +0.75$ at $\omega_i = 0.6$, and Δ[N/H] $\simeq +0.85$ to $+0.95$ at $\omega_i = 0.8$. The corresponding $\Delta C + \Delta N$ moves only weakly with ω_i because rotational mixing simultaneously raises N and depresses C: +0.41, +0.48, +0.45 at $\omega_i = 0.0, 0.6, 0.8$, respectively. This is also the physical content of the thermohaline-plus-rotation models of Charbonnel & Lagarde (2010) and Lagarde et al. (2012) at lower progenitor masses, extended here to the cool-RSG regime: surface CNO is a product of dredge-up plus rotational mixing acting together, and the optical CN band integrates them.

Mapping our measured $\Delta C + \Delta N$ values onto initial rotation rate carries an intrinsic limitation: under PARSEC v2.0, $\Delta C + \Delta N$ changes by only $\simeq 0.07$ dex between $\omega_i = 0$ and $\omega_i = 0.6$, because rotational mixing simultaneously raises N and depresses C. The variable we measure robustly is therefore weakly rotation-dependent by construction, and the EW fit alone cannot cleanly discriminate slow- from fast-rotation channels. What the fit can say is that the data are consistent with slow-rotation first dredge-up in LMC and M33 ($\Delta\chi^2(\omega_i = 0) \lesssim 0.9$; Figure 5), while the baseline M31 fit sits at $\sim 2\sigma$ below the slow-rotation locus along the $\Delta C + \Delta N$ axis.

The M31 offset is anchored by the CaT 8542 Å line: removing CaT from the joint fit shifts the M31 argmin onto the iso-($\Delta C + \Delta N$) locus that runs through the PARSEC slow-rotation point, inside the 1σ contour (Figure 5, dashed red locus in the M31 panel), so the formal 2σ tension is a single-feature calibration question rather than a CNO physics one.

The rotation interpretation is therefore best read as a conditional inference from the EW fits plus the adopted stellar-evolution prior, not as a direct measurement of $v \sin i$ or of an underlying initial-rotation distribution. Standard single-star evolution with first dredge-up reproduces the CNO signature observed here in LMC and M33 without invoking binary spin-up channels, and the data do not require the larger Δ[N/H] that such channels would generically produce.

The independent constraint that selects against fast rotators in this sample comes from the visual "weak-

The visual weak-CN selection independently removes the strongest-CN tail and therefore biases the sample toward slower rotators.

Weak-CN stars may trace recent star formation, but raw counts mostly reveal how differently the three surveys selected targets.

CN" cut itself, not from the EW fit: the young Magellanic Cloud cluster NGC 1866 hosts a red main-sequence population of fast rotators near critical rotation ($\omega \simeq 0.9$) that contains the majority of the cluster's main-sequence stars at the relevant progenitor masses (Milone et al. 2017), and the broader thermohaline-plus-rotation surface-CNO literature (Charbonnel & Lagarde 2010; Lagarde et al. 2012; Costa et al. 2019) predicts that fast rotators produce stronger CN once dredge-up is complete. A low-CN-EW cut by construction excludes the strong-CN tail of the rotation distribution, so the weak-CN cohort is biased toward the slow-rotation end of the underlying RSG population. This sample-construction argument is independent of, and complementary to, the C+N EW measurement: the EW fit measures $\Delta C+\Delta N$, while the visual cut filters the rotation distribution at the input.

The relevant clock is the roughly 40 Myr lifetime of the 5–10 solar-mass progenitors.

Several additional systematics could further shift the M31 baseline result beyond the CaT calibration identified above: an M31 sample skewed toward lower-mass or pre-first-dredge-up progenitors, a fractional underestimate of the 7400–8400 Å pseudo-continuum at solar metallicity (where the metallic line blanketing is deepest), or a ± 0.1 dex bias in the adopted M31 O abundance that shifts the C+N degeneracy ridge in the $\Delta C+\Delta N$ direction. Discriminating these requires higher-resolution M31 spectroscopy to break the C+N degeneracy directly, per-star photometric ages or masses for the M31 weak-CN sample, an independent calibration of the M31 pseudo-continuum recipe against stars of known surface composition, or a re-validation of the CaT line strength against Galactic-disc RSGs of known $T_{\rm eff}$ and $[\alpha/{\rm Fe}]$.

6.3. *Environment, Incidence, and Recent Star Formation*

The atmospheric-modelling result of Section 5 — that weak CN is the molecular-equilibrium signature of an ordinary He-burning cool RSG of 5–10 $M_\odot$ — enables one corollary application that does not depend on resolving ΔC and ΔN separately: weak CN remains a useful tracer even where the optical C+N abundance constraint is degeneracy-limited.

The discovery papers place the sample on the cool-RSG core-He-burning locus at masses of 5–10 $M_\odot$ (Guhathakurta et al. 2025; Grion Filho et al. 2025), so the population-level incidence should track the recent ($\lesssim 40$ Myr-averaged) star formation rate on the timescale set by the He-burning lifetime of those progenitors; weak-CN counts per unit host stellar mass are equivalently a specific star formation rate (sSFR) tracer over the same timescale. A simple stellar-evolution scaling relates the expected weak-CN count to the recent SFR through

$$N_{\rm WCN} \propto \langle {\rm SFR} \rangle_{40\,{\rm Myr}}\, \tau_{\rm RSG}\, \frac{f_{\rm IMF}(5\text{–}10\,{\rm M}_\odot)}{\langle m \rangle}, \qquad (4)$$

where $\langle {\rm SFR} \rangle_{40\,{\rm Myr}}$ is the SFR averaged over the ~ 40 Myr main-sequence lifetime of the progenitor mass range, $\tau_{\rm RSG}$ is the [illegible] lifetime, $f_{\rm IMF}(5\text{–}10\,{\rm M}_\odot)$ is the IMF fraction in the progenitor range (a few per cent for a standard IMF), and $\langle m \rangle$ is the mean progenitor mass. Ratios of weak-CN counts across hosts should therefore approximate ratios of their recent SFRs, modulated by spectroscopic survey selection.

A serious SFR calibration must forward-model each survey; the back-of-envelope scaling is only a sanity check.

Published recent SFRs of the three hosts span a factor of two: the LMC at $\approx 0.2\,{\rm M}_\odot\,{\rm yr}^{-1}$ (Harris & Zaritsky 2009; Zaritsky et al. 2004), M33 at $\sim 0.3\,{\rm M}_\odot\,{\rm yr}^{-1}$ (Williams et al. 2021; Lazzarini et al. 2022), and M31 at $\sim 0.4\,{\rm M}_\odot\,{\rm yr}^{-1}$ (Lewis et al. 2015). The survey-level weak-CN counts are $N_{\rm LMC} = 779$, $N_{\rm M33} = 659$, $N_{\rm M31} = 224$ (Guhathakurta et al. 2025; Grion Filho et al. 2025). A naive application of Equation 4 with the LMC value used to absorb the common $\tau_{\rm RSG} f_{\rm IMF}/\langle m \rangle$ prefactor predicts $N_{\rm M33}^{\rm pred} = 779 \times (0.3/0.2) \simeq 1170$ and $N_{\rm M31}^{\rm pred} = 779 \times (0.4/0.2) \simeq 1560$, both substantially larger than the observed counts. Equivalently, if one inverts Equation 4 to read recent SFR off the observed weak-CN counts at fixed LMC anchor, the inferred values are $\sim 0.17\,{\rm M}_\odot\,{\rm yr}^{-1}$ for M33 and $\sim 0.06\,{\rm M}_\odot\,{\rm yr}^{-1}$ for M31 — under by factors of ~ 2 and ~ 7 relative to the literature SFRs respectively.

The discrepancy direction (counts under-predicted at larger host distance) is what one expects from spectroscopic survey selection rather than from a breakdown of the molecular-physics picture. The LMC HYDRA programme observed a wide-field fibre target list, while the M31/M33 DEIMOS programme drew from a much smaller mask footprint at much greater distance and with a more aggressive target-priority weighting against the cool-RSG colour locus. A quantitative test of Equation 4 would have to forward-model the survey selection per host, which is beyond the present scope. At order-of-magnitude level the weak-CN counts and the host SFRs are nevertheless consistent with the population being a $T_{\rm eff}$-selected sub-class of ordinary He-burning RSGs [illegible] its occupation fraction.

The distance-dependent count deficit has exactly the sign expected from footprint and targeting incompleteness.

6.4. *The Missing Pure-Carbon Lever and Outstanding Questions*

The single outstanding observational lever in this analysis is a pure-carbon indicator that could anchor ΔC separately from $\Delta C+\Delta N$ and so break the C+N de-



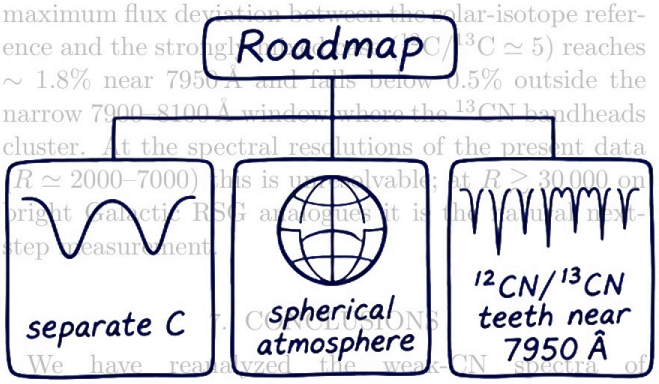
Breaking the remaining degeneracy requires a cleaner carbon indicator or resolved isotope-sensitive CN structure.

generacy. The C_2 Swan windows (5125–5180 and 5585–5640 Å) are the obvious candidates: they are present in the M31/M33 grid and are measured for both galaxies; the residuals in Table 4 are within $\sim 1.3\sigma$, consistent with C/O<1 atmospheres at the inferred surface composition, but the $\sigma_{C_2} \simeq 2.5$–3.4 Å floor at DEIMOS $R \simeq 2000$ is wide, and is dominated by the 20% line-list calibration term rather than by recipe drift. Reducing it would require a sharper-side improvement in the C_2 Swan line list paired with a per-window continuum calibration, or a higher-resolution measurement in which the broad molecular bands become resolved against the pseudo-continuum rather than smearing into it.

C_2 Swan is the obvious carbon lever, but current line-list systematics overwhelm its leverage at $R \approx 2000$.

A further caveat that the present analysis cannot internally test is the assumption of plane-parallel atmospheric geometry. `pykurucz` uses the ATLAS-style plane-parallel formulation, while at $\log g \simeq 0$ the photospheric extension is a non-negligible fraction of the stellar radius, and spherical-symmetric models more accurately capture the radiative transfer through the upper layers. The leading-order effect of sphericity on the integrated band EWs used here at $R \leq 7000$ is expected to fall within the adopted $\sigma_{\rm EW}$ floors for most indices, but we cannot test this internally. A definitive treatment will require re-fitting with spherical-symmetric atmospheres (e.g., COMARCS; Aringer et al. 2016), which we identify as the natural next step beyond the present analysis.

A direct test of the rotation channel is the surface $^{12}C/^{13}C$ ratio in the CN bandheads. Rotation-driven mixing brings up ^{13}C-enriched material from the CNO-burning shell along with raising the absolute N abundance, and the resulting ^{13}CN bandheads imprint a wavelength-resolved pattern on the CN red region rather than a bulk EW shift. We have synthesised this prediction at the LMC anchor by varying the $^{12}CN/^{13}CN$ partition in the Brooke et al. (2014) CN red-system line list while holding the rest of the atmosphere fixed: the maximum flux deviation between the solar-isotope reference and the strongly ^{13}C-enriched ($^{12}C/^{13}C \simeq 5$) reaches $\sim 1.8\%$ near 7950 Å and falls below 0.5% outside the narrow 7900–8100 Å window where the ^{13}CN bandheads cluster. At the spectral resolutions of the present data ($R \simeq 2000$–7000) this is unresolvable; at $R > 30{,}000$ on bright Galactic RSG analogues it is the natural next-step measurement.



7. CONCLUSIONS

We have reanalyzed the weak-CN spectra of Guhathakurta et al. (2025) for M31 and M33 and Grion Filho et al. (2025) for the LMC using a four-dimensional, self-consistently iterated cool-RSG model atmosphere grid in `pykurucz` (Kim & Ting 2026) spanning $T_{\rm eff}$, [α/Fe], Δ[C/H] and Δ[N/H] at each host's metallicity. Our main conclusions are:

1. Weak CN is not evidence for a carbon-star intermediate phase. It forms naturally in ordinary C/O< 1 cool-RSG atmospheres near the CN/TiO molecular transition — TiO crosses its dissociation transition near $T_{\rm eff} \sim 3800$ K, while CN remains present up to $T_{\rm eff} \sim 4300$ K, so a [redacted] appears in an otherwise non-carbon-star [redacted] expected morphology for a cool RSG inside this transition window.

Spherical atmospheres are the next modelling check; plane-parallel geometry is a known compromise at $\log g \approx 0$.

2. Because the optical CN bands respond to the product of available C and N number densities, the data constrain the combined C+N surface abundance Δ[C/H] + Δ[N/H] rather than each abundance individually. The C_2 Swan bands were tested as a candidate pure-carbon lever; at DEIMOS $R \simeq 2000$ they do not yet split the C/N degeneracy at the line-list-dominated systematic floor.

3. Joint χ^2 best-fit grid points reproduce the observed weak-CN coadds with ordinary RSG parameters and per-feature residuals at the level of the adopted EW systematic floors. Across the three galaxies the combined C+N surface abundance Δ[C/H] + Δ[N/H] is +0.3, +0.8, and −0.4 dex for the LMC, M33, and M31 best-fit grid points respectively (or +0.3 dex for M31 when CaT is dropped from the fit; see point 4 below); the PARSEC slow-rotation value ($\Delta C + \Delta N \simeq +0.4$) sits within $\Delta\chi^2 \lesssim 0.9$ of all three argmins under that reading.

4. Under the PARSEC v2.0 (Costa et al. 2019; Nguyen et al. 2022) first-dredge-up plus rotational-mixing prior, all three weak-CN coadds are consistent with slow-rotation first dredge-up, and the visual weak-CN cut at the input independently selects against fast rotators. We cannot, however, rule out faster rotation channels from the EW fit alone, because Δ[C/H] + Δ[N/H] shifts by only $\simeq 0.07$ dex between $\omega_i = 0$ and $\omega_i = 0.6$ under PARSEC.

Weak-CN stars are ordinary cool RSGs; the diagnostic is a temperature window, not a chemical peculiarity. The TiO–CN dissociation offset and the [redacted] atmosphere grid together resolve the discrepancy without invoking carbon-star chemistry, and [redacted] three weak-CN coadds in the molecular region [redacted].

High-resolution ^{13}CN bandheads could test rotational mixing directly instead of inferring it from bulk equivalent widths.

TAKE-HOME: weak-CN stars are ordinary 5–10-solar-mass cool RSGs selected by thermodynamics, not chemical peculiarity.

slow-rotation first dredge-up reproduces the observed combined CNO surface abundance. What weak CN inherits from this picture is well-defined mass and age: a ~ 5–$10\,M_\odot$ cool He-burning sub-class of the RSG population, identified at low spectral resolution by photospheric thermodynamics rather than by surface composition.

The useful legacy is a physically understood low-resolution class with a clear age and mass scale.

ACKNOWLEDGEMENTS

Large-language-model coding assistants (Anthropic Claude Opus 4.7 and OpenAI GPT-5.5) were used during this work to refactor analysis scripts, draft prose, and — the use we found most valuable — vas variant assumptions on the EW recipe, grid of ture choice, and microturbulence at a spec that allowed us to confirm the qualitative conclusions were robust to those choices. The model output also contained substantive errors and dead ends; sifting them, identifying the load-bearing physical explanations, and constructing the argument of the paper were human tasks, and the scientific claims and the interpretation are the authors'.

We disclose the coding assistants because they accelerated robustness tests; the scientific judgment and error-checking remained ours.

REFERENCES

Aringer, B., Girardi, L., Nowotny, W., Marigo, P., & Bressan, A. 2016, MNRAS, 457, 3611, doi: 10.1093/mnras/stw222

Boyer, M. L., Girardi, L., Marigo, P., et al. 2013, ApJ, 774, 83, doi: 10.1088/0004-637X/774/1/83

Brooke, J. S. A., Ram, R. S., Western, C. M., et al. 2014, ApJS, 210, 23, doi: 10.1088/0067-0049/210/2/23

Charbonnel, C., & Lagarde, N. 2010, A&A, 522, A10, doi: 10.1051/0004-6361/201014432

Costa, G., Girardi, L., Bressan, A., et al. 2019, MNRAS, 485, 4641, doi: 10.1093/mnras/stz728

Dalcanton, J. J., Williams, B. F., Lang, D., et al. 2012, ApJS, 200, 18, doi: 10.1088/0067-0049/200/2/18

Davies, B., Kudritzki, R.-P., Gazak, Z., et al. 2015, ApJ, 806, 21, doi: 10.1088/0004-637X/806/1/21

Grion Filho, D., Guhathakurta, P., Rinehart, S. M., et al. 2025, ApJ, 981, 88, doi: 10.3847/1538-4357/adb030

Guhathakurta, P., Grion Filho, D., Bhattacharya, A. R., et al. 2025, ApJ, 987, 203, doi: 10.3847/1538-4357/adcc2c

Gustafsson, B., Edvardsson, B., Eriksson, K., et al. 2008, A&A, 486, 951, doi: 10.1051/0004-6361:200809724

Hamren, K., Beaton, R. L., Guhathakurta, P., et al. 2016, ApJ, 828, 15, doi: 10.3847/0004-637X/828/1/15

Harris, J., & Zaritsky, D. 2009, AJ, 138, 1243, doi: 10.1088/0004-6256/138/5/1243

Kim, E. M., & Ting, Y.-S. 2026, Journal of Open Source Software. https://arxiv.org/abs/2603.11693

Kurucz, R. L. 2005, Memorie della Societa Astronomica Italiana Supplementi, 8, 14

Lagarde, N., Decressin, T., Charbonnel, C., et al. 2012, A&A, 543, A108, doi: 10.1051/0004-6361/201118331

Lazzarini, M., Williams, B. F., Durbin, M. J., et al. 2022, ApJ, 934, 76, doi: 10.3847/1538-4357/ac7568

Levesque, E. M., Massey, P., Olsen, K. A. G., et al. 2005, ApJ, 628, 973, doi: 10.1086/430901

Lewis, A. R., Dolphin, A. E., Dalcanton, J. J., et al. 2015, ApJ, 805, 183, doi: 10.1088/0004-637X/805/2/183

Magrini, L., Stanghellini, L., & Villaver, E. 2009, ApJ, 696, 729, doi: 10.1088/0004-637X/696/1/729

Milone, A. P., Marino, A. F., D'Antona, F., et al. 2017, MNRAS, 465, 4363, doi: 10.1093/mnras/stw2965

Nguyen, C. T., Costa, G., Girardi, L., et al. 2022, A&A, 665, A126, doi: 10.1051/0004-6361/202244166

Pastorelli, G., Marigo, P., Girardi, L., et al. 2020, MNRAS, 498, 3283, doi: 10.1093/mnras/staa2565

Scalo, J. M. 1974, ApJ, 194, 361, doi: 10.1086/153253

Ting, Y.-S., Conroy, C., Rix, H.-W., & Asplund, M. 2018, ApJ, 860, 159, doi: 10.3847/1538-4357/aac6c9

Williams, B. F., Durbin, M. J., Dalcanton, J. J., et al. 2021, ApJS, 253, 53, doi: 10.3847/1538-4365/abdf4e

Zaritsky, D., Harris, J., Thompson, I. B., & Grebel, E. K. 2004, AJ, 128, 1606, doi: 10.1086/423910

