Automatic Phase Correction of ^{1}H NMR Spectra and the Baseline Distortion It Creates

Algorithm Families, a Working Python Implementation, and a Case Study on a 60 MHz Benchtop Spectrum of a Hydrocarbon Sample

John C. Edwards, Ph.D.
Process NMR Associates, LLC
84 Patrick Lane, Suite 115, Poughkeepsie, NY 12603
www.process-nmr.com

Document Details

Item	Value
Title	Automatic Phase Correction of ^{1}H NMR Spectra and the Baseline Distortion It Creates
Document type	Technical white paper with accompanying software
Version	1.1
Generation date	11 August 2026
Revision	v1.1 adds Section 9, in which the flank-flatness criterion requested for this spectrum is implemented as an alternating zero-order / first-order refinement, the total flank slope is decomposed into its phase and ramp contributions, the first-order bound is re-expressed in dwell periods and widened, and the adopted solution is revised from φ_1 = +112.80° to +108.30°. Tables and figures after Section 8 are renumbered accordingly.
Filename	Automatic_Phase_Correction_1H_NMR_WhitePaper_v1.1.docx
Author	John C. Edwards, Ph.D., Process NMR Associates, LLC
Subject data file	TalFFTAVG00019.spc (4IR Solutions AI-60, ^{1}H, acquired 9 August 2026)
Accompanying software	nmr_autophase.py v1.1 (Python 3, numpy/scipy)
Accompanying data	TalFFTAVG00019_phased_v2.spc (adopted); TalFFTAVG00019_phased_v2_flatflank.spc (raw flat-flank alternative); TalFFTAVG00019_phase_baseline_full.csv; TalFFTAVG00019_phase_baseline_study.xlsx
Model identity as reported by the session	claude-opus-5
Tools engaged	Literature retrieval and Crossref DOI verification; native GRAMS/Thermo SPC binary parsing; numpy/scipy numerical analysis; matplotlib figure generation; openpyxl workbook generation
Reference verification	Every DOI in the reference list was resolved against the Crossref REST API and the returned title, author list, container title, volume, issue, pagination and year were compared against the citation as written. Items that could not be verified were discarded rather than estimated.
AI assistance statement	Background material, algorithm descriptions, document structure, the Python implementation and the numerical analysis of the supplied spectrum were drafted with AI assistance. The instrument, the sample, the acquisition, the analytical interpretation and the conclusions drawn for client work remain the responsibility of the author. No experimental datum reported here was generated by the AI; all numerical results derive from the supplied SPC file.

Abstract

The frequency-dependent phase error carried by every pulse-Fourier-transform NMR spectrum arises from a physically unavoidable delay between the middle of the excitation pulse and the first sampled point of the free-induction decay, and its removal is among the oldest and least settled problems in NMR data processing. This white paper sets out the theoretical basis of zero-order and first-order phase correction, traces the principal families of automatic phasing algorithm from the numerical Hilbert-transform treatment of Ernst[2] through dispersion-versus-absorption analysis[16,19], lineshape symmetry[22], baseline flatness[20], entropy minimization[27], spectral whitening[29], joint phase-and-baseline optimization[36,37] and the transformer networks now appearing in the metabolomics literature[40,41], and then examines what each family actually delivers when confronted with a low-field spectrum of a complex hydrocarbon mixture. A Python module has been written to accompany the text; it reads and writes GRAMS/Thermo SPC files natively, recovers the quadrature channel from the SPC log binary area, offers five interchangeable phase-objective engines under a physically bounded first-order term, and follows the rotation with polynomial, cubic-spline multipoint, Whittaker or adaptive iteratively reweighted penalized least-squares baseline estimation.

The module has been applied to a 65 536-point ^{1}H spectrum recorded at 60.528437 MHz on a 4IR Solutions AI-60 benchtop instrument, in which a phase anomaly had been reported over the broad aromatic envelope between 4 and 9 ppm while the aliphatic region below 4 ppm appeared correct. Three findings emerged that were not anticipated. First, the aromatic band is not principally mis-phased: when a local baseline anchored on both flanks is refitted at every trial angle, a rotation of ±20° alters the band envelope by no more than 12 % of its own amplitude and leaves a large, angle-independent asymmetry of 0.29 to 0.41, so the shape is that of a genuinely composite aromatic envelope rather than a dispersive admixture. What is being observed instead is a dispersive pedestal, contributed by the intense aliphatic resonances, that is thirty to fifty times larger than the local baseline and that responds violently to small rotations. Second, the residual phase error is small but real and is genuinely frequency dependent: independent measurement in the aliphatic and aromatic bands gives −6.25 ± 1.23° at 1.253 ppm and +3.75 ± 0.42° at 7.111 ppm, from which a first-order term of +112.8 ± 14.7° across the 4000 Hz sweep follows, corresponding to an effective first-point delay of 78.3 µs against a dwell period of 250 µs. Third, three of the five published objective functions tested returned a negative implied first-point delay, which is physically inadmissible because the receiver cannot sample before the pulse; the same three disagreed with one another on the derived aromatic hydrogen content by up to a factor of two.

A second determination was then carried out at the request of the analyst, using the criterion an experienced spectroscopist applies by hand: the first-order term adjusted until the baseline windows flanking the aromatic envelope at 8.5 to 9.5 and 5.5 to 6.2 ppm both lie flat, the zero-order term adjusted until the aliphatic envelope is restored, and the two alternated to convergence. The reading that the upper flank required roughly five degrees more rotation than the lower one is confirmed by measurement, at +5.28 against +2.10 degrees. Implemented against the total flank slope the loop nevertheless converges on a first-order term of +288 degrees, and the reason is general rather than particular: the slope of a phased trace separates exactly into a term that vanishes when the local phase is correct and a term proportional to the ramp rate times the local signal amplitude, and where a dispersive pedestal forty times the local absorption baseline is present the second dominates. The apparent flatness at +288 degrees is a cancellation between two substantial errors; it carries +19.8 degrees of phase onto the aromatic band, costs 15.7 % of its integral and opens a negative lobe on one flank. Nulling the genuine term alone converges in four cycles to a zero-order term of -6.00 and a first-order term of +108.30 degrees, within four degrees of the independent band-based determination, and the residual ramp-induced tilt is then removed at no cost by giving the multipoint baseline enough nodes: thirty nodes bring both flanks flat to the noise while leaving the aromatic integral stable.

The dominant conclusion is quantitative rather than cosmetic. Across the correction strategies examined, the aromatic hydrogen index of this sample — the 6 to 9 ppm integral expressed as a percentage of the summed 0–4, 4.5–6 and 6–9 ppm integrals — ranged from 1.39 % to 3.18 %. Phase correction accounted for approximately 10 % relative change; the uncertainty in the first-order term contributed under 2 % relative; and the selection of baseline anchor windows accounted for the remainder, which was by far the largest single contribution. For low-field quantitative

work on complex mixtures, the anchoring convention is therefore the parameter that must be fixed, documented and applied uniformly across a sample set, and it deserves more attention in a method-validation package than the phasing algorithm itself.

Table of Contents

Document Details 1
Abstract 2
1. Introduction 6
2. Theoretical Basis 6
2.1 Origin of the frequency-dependent phase error 6
2.2 Causality, the Kramers–Kronig relation and the dispersion channel 7
2.3 Why a first-order correction necessarily distorts the baseline 7
3. Families of Automatic Phase Correction Algorithm 8
3.1 Hilbert-transform and peak-based criteria 8
3.2 Dispersion versus absorption 8
3.3 Lineshape symmetry 9
3.4 Baseline flatness 9
3.5 Entropy minimization 9
3.6 Whitening, real-signal maximization and energy-based criteria 9
3.7 Simultaneous phase and baseline correction 10
3.8 Model-based, reference-deconvolution and machine-learning approaches 10
3.9 The search strategy 10
4. Baseline Correction After First-Order Phase Correction 10
4.1 Identification of baseline points 11
4.2 Polynomial models 11
4.3 Cubic-spline multipoint fitting 11
4.4 Penalized least-squares smoothers 12
4.5 Consequences for quantitation 12
5. The Data Set Examined 12
5.1 File structure 12
5.2 Acquisition parameters 13
5.3 Consistency of the stored dispersion channel 14
6. Diagnosis of the Reported Anomaly 14
6.1 The dispersive pedestal 14
6.2 Is the aromatic band actually mis-phased? 15
6.3 Conditioning of the global objective functions 16
7. Determination of the Phase Parameters 16
7.1 The regional method 16
7.2 Result and physical interpretation 16
7.3 Comparison against the published criteria 17
8. Baseline Correction in Practice 17
8.1 Failure of noise-threshold classification 17
8.2 The anchoring convention dominates the result 18
8.3 Functional form: spline against polynomial 18
8.4 The corrected spectrum 19
9. Iterative Refinement Against the Baseline Flanks 20
9.1 The observation 20
9.2 Why the total flank slope is not the phase 20
9.3 The branch ambiguity 21
9.4 The alternating loop and its convergence 21

9.5 What the flanks actually require....22
10. Consequences for Quantitation....23
11. The Python Implementation....25
11.1 File handling....25
11.2 Interface....25
11.3 Design decisions that proved necessary....26
12. Recommendations and Limitations....26
13. References....28
14. Summary....32

1. Introduction

A pulse-Fourier-transform NMR experiment does not deliver an absorption spectrum. What is delivered is a complex function of frequency whose real and imaginary parts are mixtures of absorption and dispersion in a proportion that varies systematically across the spectral width, and the recovery of a pure absorption-mode trace requires that this mixture be undone. The operation is called phase correction, it has been part of the standard processing chain since the earliest days of Fourier-transform spectroscopy[1], and in the routine case it is dispatched in a few seconds by an automatic routine that the spectroscopist rarely inspects. The routine case, however, is a spectrum of a pure compound with well-resolved, narrow lines and generous regions of true baseline. Neither condition holds for the material that occupies most of the working life of an industrial NMR laboratory.

Low-field benchtop instruments have made this discrepancy considerably sharper. At 60 MHz the ^{1}H spectrum of a petroleum fraction, a vegetable oil or a formulated product is not a set of lines but a set of overlapping envelopes, and the assumptions on which most published phasing criteria rest — that isolated resonances exist, that lines are symmetric, that regions of signal-free baseline can be found — are simply not satisfied. It is nevertheless from precisely such spectra that aromaticity, olefinic content, hydrogen distribution and a range of derived properties are calculated, and the numbers that emerge depend on how the spectrum was phased and where its baseline was judged to lie. The purpose of this document is to make that dependence explicit, to describe the algorithms available and the conditions under which each fails, and to provide a working implementation that can be inspected, bounded and audited.

The document is organized as follows. Section 2 develops the physical origin of the phase error and derives the relation between the first-order term and the acquisition timing, which supplies the only genuinely objective constraint on the problem. Section 3 surveys the algorithm families. Section 4 treats baseline correction, which is not an independent operation but the necessary consequence of any nonzero first-order term. Sections 5 through 9 apply the resulting method to a specific spectrum and quantify what is gained and what is merely rearranged. Section 11 describes the accompanying Python module, and Section 12 records the recommendations and the limitations that a reviewer would be entitled to raise.

2. Theoretical Basis

2.1 Origin of the frequency-dependent phase error

Let the transverse magnetization following a hard pulse be represented as a sum of decaying complex exponentials, and let the receiver begin sampling not at the middle of the pulse but at a time t_0 afterwards. The measured signal is then the true signal delayed by t_0, and by the shift theorem the Fourier transform of a delayed signal is the transform of the undelayed signal multiplied by a linear phase ramp. Writing the observed complex spectrum as $S_{obs}(\nu)$ and the desired absorption-mode spectrum as $S_{true}(\nu)$, the relation is

$$S_{obs}(\nu) = S_{true}(\nu) \cdot \exp[\, i\,(\varphi_0 + 2\pi\,\nu\,t_0)\,]$$

in which the constant term φ_0 collects everything that is independent of frequency — the phase of the reference oscillator against the receiver, any fixed offset in the quadrature detector[5], the residual of an imperfect pulse — and the second term carries the entire frequency dependence. Correction consists of the pointwise rotation

$$S_{corr}(\nu) = S_{obs}(\nu) \cdot \exp[\, -i\,(\varphi_0 + \varphi_1\,(\nu - \nu_{piv})\,/\,SW)\,]$$

where SW is the spectral width in hertz, ν_{piv} is the pivot frequency at which the first-order term contributes nothing, and φ_1 is the total phase swept across the full spectral width. Setting the two expressions equal establishes the identity that governs everything that follows:

$$\varphi_1 = 360^{\circ} \cdot SW \cdot t_0$$

This is not a curve-fitting parameter. It is a statement about the acquisition timing, and it means that one complete dwell period of pre-acquisition delay corresponds to exactly 360° of first-order phase. Because t_0 cannot be negative — a receiver cannot sample before the pulse that creates the signal — a fitted φ_1 of the wrong sign is not a marginally worse answer but an inadmissible one, and this simple check turns out to be the most powerful discriminator available between competing automatic routines. The contributions to t_0 are the receiver dead time and probe ringdown, the group delay of any analog and digital filtering in the signal path, one dwell period if the console begins sampling one point late, and a small correction of $2 \cdot pw/\pi$ arising from the finite duration pw of the excitation pulse.

On modern high-field consoles the digital filter chain introduced with oversampling[6] dominates t_0 and produces a group delay of tens of dwell periods, so that the raw first-order phase runs to many thousands of degrees before the vendor software removes it by a circular shift of the FID[7]. Benchtop permanent-magnet instruments generally apply far less digital filtering, and the residual t_0 is correspondingly small — tens of microseconds rather than milliseconds. The consequence for the present work is that a physically defensible bound on φ_1 can be written down before any optimization is attempted, and that bound is tight enough to be useful.

2.2 Causality, the Kramers–Kronig relation and the dispersion channel

The absorption and dispersion components of a Fourier-transform NMR spectrum are not independent functions. Because the free-induction decay is causal, vanishing for all times before the pulse, the real and imaginary parts of its transform are related by a Hilbert transform, which is the spectroscopic statement of the Kramers–Kronig relations. Bartholdi and Ernst set this out formally for Fourier spectroscopy[3], and Ernst had already exploited it the other way round, generating the dispersion trace numerically from a stored absorption trace so that phase correction could be applied to a spectrum whose imaginary channel had been discarded[2]. The practical importance is considerable: a spectrum archived as a real trace only is not beyond phase correction, because the missing quadrature partner can be reconstructed.

The reconstruction is nevertheless not free of consequence, and the caveat deserves to be stated plainly because it is routinely overlooked. A Hilbert transform assumes that the trace it is given is a genuine causal absorption spectrum. If the real trace has already been baseline-corrected, offset, truncated at the edges of the sweep, or contains folded intensity, the reconstructed dispersion inherits those defects in a form that has no physical counterpart. Subsequent phase rotation then mixes an artifact into the absorption channel. Where a stored imaginary channel exists it should always be preferred, and where reconstruction is unavoidable the consistency of the result should be examined rather than assumed. The module accompanying this paper reports the correlation between the stored and reconstructed dispersion traces for exactly this purpose[4].

2.3 Why a first-order correction necessarily distorts the baseline

The coupling between the first-order phase term and the spectral baseline is not an implementation defect that better software would remove; it is a property of the discrete Fourier transform. A linear phase ramp applied in the frequency domain is equivalent to a circular shift in the time domain. When the signal being shifted is not periodic within the acquisition window — and an FID never is, since it begins at full amplitude and decays — the wrap-around introduces a discontinuity, and the transform of that discontinuity is a slowly varying ripple across the whole spectrum. Zhu, Torchia and Bax derived the relation between the sampling delay, the required first-order phase and the resulting baseline explicitly, and showed that the baseline is clean only when the delay is an exact integer multiple of the dwell period[8]. Starčuk and co-workers gave the corresponding analysis for simultaneous quadrature sampling, which is the detection mode used by the instrument examined here[9].

A closely related mechanism operates through the first data point. Because the first point of the FID contributes a constant to every frequency in the spectrum, an error in its amplitude appears as a uniform offset, and the familiar prescription of halving it[11] is the correction for the trapezoidal weighting that the discrete transform implies. Marion and Bax showed how initial-point error converts into a distinct baseline roll under real Fourier processing[12], and Tang

separated the resulting distortion into an offset term traceable to first-point amplitude and a rolling term arising from the combination of multi-point corruption with frequency-dependent phase[13]. Schaefer and Stejskal had identified the phenomenon and attributed it to receiver and audio-filter response two decades earlier[10]. Where the corruption extends over several points, backward linear prediction offers a genuine repair rather than a cosmetic one, either in the general parametric form of Barkhuijsen and co-workers[14] or in the phase-constrained variant of Zhu and Bax[15]; both act on the time-domain data and therefore address the cause rather than the symptom. Neither was available for the present work, because the supplied file contains a transformed spectrum and not an FID.

The practical corollary is that phase correction and baseline correction cannot be treated as independent steps to be optimized in sequence. Any objective function that scores the flatness of a phased spectrum, and any baseline model flexible enough to follow a slow roll, will be able to trade one against the other. The degeneracy is real, it is demonstrated numerically in Section 7 below, and it is the reason why the accompanying implementation both restricts the flexibility of the baseline model used inside the phase objective and bounds the first-order term by the acquisition timing.

3. Families of Automatic Phase Correction Algorithm

Every automatic phasing method consists of two separable choices: a criterion that scores how nearly a candidate spectrum resembles pure absorption, and a search strategy that locates the extremum of that criterion in the two-dimensional space spanned by φ_0 and φ_1. The two are frequently conflated in the literature, and Binczyk and co-workers made the useful point that reported differences between published methods often reflect the optimizer rather than the criterion[32]. The families below are distinguished by criterion; the search strategy is treated separately in Section 3.9.

3.1 Hilbert-transform and peak-based criteria

The earliest automatic treatment measures the phase directly at one or more resonances. For an isolated Lorentzian the dispersion component vanishes at the apex, so the argument of the complex spectrum evaluated there gives the local phase immediately; measuring it at two or more well-separated resonances and fitting a straight line yields φ_0 and φ_1 at once. Ernst set out the Hilbert-transform machinery that makes this possible even when only the real trace has been retained[2], and van Vaals and van Gerwen surveyed and extended the peak-based criteria, remaining the best single comparison of the pre-1990 methods[21]. The approach is fast, transparent and physically interpretable, and its weakness is equally clear: contamination of the apex by the dispersive wings of neighboring resonances biases the measurement, and the bias grows as the spectrum becomes more crowded. Džakula addressed this by measuring the phase from integrated peak areas rather than from apex intensities[26]; because the dispersion contribution to an integral taken symmetrically about a band centroid vanishes by oddness, the area-based estimate is markedly more robust, and it forms the basis of the regional engine described in Section 8.

3.2 Dispersion versus absorption

Plotting the dispersion component against the absorption component of a Lorentzian resonance traces a circle passing through the origin; a phase error rotates the circle, and the rotation angle is the phase error. Marshall and Roe introduced this construction for lineshape analysis[16] and it was turned into an autophasing criterion independently and almost simultaneously by Sotak, Dumoulin and Newsham[17] and by Herring and Phillips[18]. The definitive treatment, including the extension to the first-order term, is that of Craig and Marshall[19]. The method is elegant and it is exact for an isolated Lorentzian. It is also, for the same reason, the family least suited to the spectra considered here, because departures from Lorentzian shape and overlap between resonances both deform the circle in ways that the fitting procedure interprets as phase.

3.3 Lineshape symmetry

A pure absorption resonance is even about its own center, whereas a dispersion component is odd. Heuer built an algorithm on this observation, scoring the antisymmetric content of detected lines and minimizing it[22]. The criterion is more forgiving than DISPA because it does not assume a specific lineshape, but it still requires that a line be identifiable and that its center be locatable. Balacco subsequently formulated a criterion that dispenses with the requirement for isolated or symmetric lines altogether[24], which was an important step toward the treatment of crowded spectra. Sterna and Tong, working specifically on the unsaturated carbon analysis of oils, addressed the low-signal, heavily overlapped case that industrial samples present[23]; theirs remains one of the few papers in the phasing literature written from the perspective of a petroleum laboratory.

3.4 Baseline flatness

Brown, Campbell and Moore proposed scoring a candidate phase by the flatness of the resulting baseline, fitting a smooth model to the regions judged signal-free and minimizing the residual[20]. Every subsequent criterion that rewards a flat baseline descends from this work. Its appeal is that it uses the parts of the spectrum where the answer is known — the baseline should be flat and at zero — rather than the parts where it is not. Its hazard is precisely the degeneracy identified in Section 2.3: because a first-order ramp itself produces a baseline roll, and because a flexible baseline model can absorb that roll, the criterion can be satisfied by combinations of phase and baseline that are individually wrong. Koehl, Ling and Lefèvre examined the statistical limits on phase accuracy imposed by noise and peak density and reached a conclusion in the same direction, namely that the achievable precision is bounded by the spectrum rather than by the algorithm[25].

3.5 Entropy minimization

The most widely implemented modern criterion is that of Chen, Weng, Goh and Garland, universally known by its acronym ACME[27]. The objective is the Shannon entropy of the normalized absolute first derivative of the real trace, augmented by a penalty term proportional to the integrated square of any negative intensity. The reasoning is that a correctly phased spectrum is the sparsest representation available: absorption lines are compact and their derivatives concentrate the probability mass into few bins, whereas dispersion contributions spread it. The negative-intensity penalty is essential, because entropy alone is invariant to a 180° rotation and would happily return an inverted spectrum.

ACME is fast, requires no peak detection and no baseline classification, and performs well on well-resolved spectra. Its behavior on crowded spectra is less satisfactory, and Bao and co-workers proposed a refinement aimed specifically at signal-dense data where the derivative entropy loses discrimination[31]. The most useful single reference for a practitioner deciding among criteria remains the benchmark of de Brouwer, who tested the major published objective functions head to head on real data and reported where each fails[28].

3.6 Whitening, real-signal maximization and energy-based criteria

Balacco and Cobas approached the problem through spectral whitening, scoring how nearly the spectrum resembles a signal with a flat power distribution after phase rotation, and showed that the criterion generalizes cleanly to two dimensions[29]. Hardy and co-workers demonstrated that simply maximizing the integral of the real trace is a surprisingly robust criterion[30], which follows from the same oddness argument that underlies area-based phase measurement: the integral of the dispersion component over a symmetric interval vanishes, so the real integral is maximized when the rotation is correct. Zieliński, Duda and Ostrowska developed a fast energy-based MinMax criterion aimed specifically at spectra carrying large linear phase distortion[34], which is the regime created by console group delay.

3.7 Simultaneous phase and baseline correction

The recognition that phase and baseline are coupled rather than sequential is comparatively recent in the phasing literature, although it follows directly from the analysis of Section 2.3. Sawall and co-workers formulated the two as a single multi-objective optimization problem and solved them jointly[36], and Steimers and co-workers validated the resulting method, designated SINC, on quantitative reaction-monitoring data where the accuracy requirement is stringent[37]. Cobas had earlier demonstrated the utility of the Whittaker smoother as the baseline component within such a scheme[35]. The joint formulation is conceptually the correct one, and it is the approach adopted by the engine designated sinc in the accompanying module. It carries a warning that is developed quantitatively in Section 7: unless the baseline component is deliberately made stiff, the joint objective is degenerate, and the optimizer will drive the first-order term to whatever bound it is given.

3.8 Model-based, reference-deconvolution and machine-learning approaches

An alternative to scoring the spectrum is to model it. Reference deconvolution, introduced by Morris, uses the known lineshape of an internal reference resonance to correct instrumental imperfections including phase, and it corrects lineshape and phase together rather than separately[42]. Where a suitable reference signal exists the approach is powerful; for industrial mixtures a clean reference line usually does not.

The machine-learning literature has developed rapidly. Zhao and co-workers described a physics-constrained network in which the phase model itself is embedded as a constraint, aimed at high-throughput metabolomics pipelines[40], and Li and co-workers reported a tandem vision transformer performing fully automated phase correction[41]. These methods report performance superior to the classical criteria on the data on which they were trained. Their applicability to low-field spectra of industrial hydrocarbon mixtures has not been demonstrated, and the training sets described are drawn from high-field biological and metabolomic data. A network that has never seen a 60 MHz spectrum of a lubricant base oil should not be assumed to phase one correctly, and in a regulated or litigated context the inability to state why a particular phase was chosen is a material disadvantage.

It should also be recorded that the two-parameter model is itself an approximation. Jaroszewicz and co-workers developed an automated multi-order routine for ultra-wideline spectra in which φ_0 and φ_1 are demonstrably insufficient[39]. For the narrow sweeps of high-resolution liquids work the higher-order terms are ordinarily negligible, but the assumption is worth stating rather than leaving implicit.

3.9 The search strategy

Every criterion described above possesses multiple local minima in the (φ_0, φ_1) plane, and a local optimizer started from the origin will return whichever basin it happens to occupy. Reproducibility therefore requires either a coarse global scan preceding local refinement, which is the strategy adopted here, or an exhaustive search. Gazziro and co-workers took the latter route, evaluating the objective over the entire (φ_0, φ_1) grid on a graphics processor and thereby eliminating the local-minimum problem outright[38]. With a decimated spectrum the exhaustive approach is entirely practical on ordinary hardware, and the module described in Section 11 performs a coarse grid evaluation before invoking a simplex refinement for exactly this reason.

Reference implementations worth consulting include the phase routines of NMRPipe[43], the ACME and peak-minima implementations in nmrglue[44], the documented commercial algorithm of Zorin, Bernstein and Cobas[33], and the automated preprocessing chains of metabolomics platforms such as rNMR[45] and MetaboLab[46].

4. Baseline Correction After First-Order Phase Correction

Having established that a nonzero φ_1 produces a rolling baseline, the question becomes how that roll should be removed without removing signal with it. Every method proceeds in two stages: the points belonging to baseline are

identified, and a smooth function is fitted through them and subtracted. The two stages fail in different ways and are worth treating separately.

4.1 Identification of baseline points

Dietrich, Rüdel and Neumann introduced the classification of baseline points by the smoothed second derivative, on the reasoning that signal-bearing points possess curvature while baseline points do not[48]. Golotvin and Williams proposed a window-based alternative in which a region is judged baseline when its peak-to-peak span does not exceed a small multiple of the noise[50], which proves markedly more robust on crowded ^{1}H spectra where derivative methods are misled by the flanks of broad bands. Chang, Banack and Shah addressed the genuinely difficult case in which signal-free regions are scarce[54], and Cobas and co-workers combined a continuous-wavelet recognition step with Whittaker smoothing to produce the fully automatic procedure implemented in commercial software[53].

All of these tests share an assumption that deserves scrutiny, namely that baseline is what is indistinguishable from noise. That assumption fails silently once the signal-to-noise ratio becomes large, because the limiting baseline is then not the noise floor but systematic structure — broad components, probe background, residual distortion — lying orders of magnitude above it. Section 7.1 demonstrates the failure quantitatively on the spectrum examined here, where a signal-to-noise ratio near 1.3×10^{5} renders every noise-threshold classifier inoperative.

4.2 Polynomial models

Fitting a polynomial of modest degree through the classified points is the oldest and simplest construction, and for a gentle roll across a narrow sweep it is entirely adequate. Two failure modes recur. The first is that a low-order polynomial cannot follow the structured pedestal that the phase ramp actually produces, leaving a systematic residual under the bands of interest. The second, and more dangerous, is that a high-order polynomial fitted only to sparse classified points is unconstrained between them, and will excurse freely in the gaps. Because the gaps are precisely where the resonances lie, the excursion is subtracted from the signal. Section 7.3 records an instance in which raising the degree from five to nine reduced the residual at the classified points while inflating the aliphatic integral by more than 60 %, which is the diagnostic signature of exactly this failure. Brown demonstrated that Bernstein polynomials, being bounded by their control points, behave better in this respect than the ordinary power basis.

4.3 Cubic-spline multipoint fitting

Zolnai, Macura and Markley introduced the spline construction for NMR baseline correction[47], and it remains the method that the interactive multipoint tool in every vendor package implements: the analyst nominates a set of points at which the spectrum is judged to be at baseline, and a cubic spline is drawn through them. Its great advantage over the polynomial is locality — an error in one node perturbs the fit only in its immediate neighborhood — and its great advantage over automatic methods is that the judgment is explicit, recorded and reviewable.

Two practical points determine whether the construction succeeds. The first is node siting: a node placed inside a resonance has no honest ordinate, and forcing one there pulls the spline up into the band and removes real integral. Nodes must therefore be sited only where baseline has actually been identified, and the ordinate at each should be a robust statistic — a median rather than a single intensity — taken over a window narrower than the node spacing. If the averaging window is wider than the spacing, adjacent nodes average essentially the same points, return identical ordinates, and the spline degenerates. The second point concerns behavior in wide gaps. A natural cubic spline is free to overshoot between distant knots, and on a spectrum whose only honest anchors lie outside the signal region the gaps are wide by construction. A monotone-preserving interpolant such as PCHIP cannot invent an excursion that the anchors do not support, and is therefore the safer default for quantitative work; the module offers both and the difference between them is quantified in Section 7.3. Sectionally linear interpolation, described for multidimensional spectra by Saffrich and co-workers, is the most conservative choice of all and is worth considering when the anchors are very sparse.

4.4 Penalized least-squares smoothers

The Whittaker smoother, presented in its modern form by Eilers[51], estimates a baseline by minimizing the sum of a weighted fidelity term and a roughness penalty on the second differences. Asymmetric weighting, in which points lying above the current estimate are down-weighted, converts it into a baseline estimator that requires no explicit classification step; the asymmetric formulation is set out by Boelens, Eilers and Hankemeier[52]. Zhang, Chen and Liang generalized the weighting into the adaptive iteratively reweighted scheme known as airPLS, which has become the most widely deployed automatic baseline estimator in open-source spectroscopic pipelines[56]. Fiorucci and co-workers have more recently treated the baseline as a low-frequency intrinsic mode separated by fast iterative filtering, dispensing with baseline-region selection altogether[57].

These methods are attractive because they are automatic, and they perform well on spectra of moderate dynamic range. On a spectrum containing a resonance three orders of magnitude taller than the features of interest they perform poorly, for a structural reason: the smoother attempts to pass beneath every point, the roughness penalty prevents it from diving under the tall resonance, and the estimate is dragged upward across the whole region. The published weight update also overflows in double precision when a single intense resonance makes the residual much larger than its integral, and the implementation accompanying this paper clamps the exponent for that reason. Güntert and Wüthrich took a different route, fitting a truncated sine and cosine series to the signal-free regions in the procedure designated FLATT[49], which behaves better in the presence of high dynamic range because the basis is bounded.

4.5 Consequences for quantitation

That phase and baseline errors propagate into integrals is not in doubt. McLeod and Comisarow derived the bias introduced by discrete summation over a phase- and baseline-imperfect lineshape[58], supplying the quantitative bridge between a processing decision and an analytical result. Malz and Jancke identified phasing and baseline correction among the leading contributors to the uncertainty budget of quantitative NMR[59], a conclusion reinforced by Bharti and Roy[60], by Le Gresley and co-workers for the non-ideal case specifically[61], and by Giraudeau in reviewing the residual barriers to sub-percent trueness[62]. Xi and Rocke quantified how residual baseline propagates into binned feature tables and thence into multivariate models[55], which is directly relevant wherever spectra are reduced to an X matrix for chemometric modeling, and Emwas and co-workers set out consensus processing recommendations for one-dimensional ^{1}H data[63].

A gap in the literature is worth recording. Despite targeted searching, no primary study could be identified that benchmarks multiple baseline-correction algorithms specifically against NMR integration accuracy on a common data set. The nearest available substitutes address the propagation into metabolomics feature tables[55] rather than the integral itself. For a laboratory seeking to justify a particular baseline convention in a method-validation package, this absence means the justification must be constructed in house from its own data, and the sensitivity analysis presented in Section 10 is offered as a template for doing so.

5. The Data Set Examined

5.1 File structure

The spectrum was supplied as TalFFTAVG00019.spc, a 525 196-byte file in the little-endian "new format" GRAMS/Thermo SPC representation identified by the version byte 0x4B. Parsing of the 512-byte main header returns a flag byte of 0x02, indicating evenly spaced abscissa values, a single subfile and 32-bit ordinate storage; an experiment code of 10, denoting NMR; an abscissa type code of 10, denoting parts per million; 65 536 points; and first and last abscissa values of +33.0423202568406 and −33.0423202568406 ppm respectively. The exponent byte holds the value −128, which in this format signifies that ordinates are stored directly as IEEE 32-bit floating-point numbers rather than as scaled integers, so no requantization is involved in reading or writing the file.

Of greater practical interest is the log block, whose offset is recorded at byte 248 of the main header. The log header declares a binary area of 262 144 bytes, which is exactly 65 536 four-byte floats, followed by a 300-byte text area. The binary area proves to contain the dispersion-mode trace. This is a significant convenience: the file is a complete complex data set rather than an archived real trace, and phase rotation can therefore be performed on the measured quadrature pair rather than on a reconstruction. The layout is set out in Table 1 for the benefit of anyone needing to read these files without the vendor software.

Table 1. Layout of the supplied GRAMS/Thermo SPC file. Byte offsets are from the start of the file.

Offset (bytes)	Length (bytes)	Content
0	512	Main header; ftflgs 0x02, fversn 0x4B, fexper 10 (NMR), fexp −128 (float ordinates), fnpts 65 536, ffirst +33.0423202568406 ppm, flast −33.0423202568406 ppm, fxtype 10 (ppm), flogoff 262 688
512	32	Subfile header
544	262 144	Real (absorption) trace, 65 536 × IEEE float32
262 688	64	Log block header (LOGSTC); logbins 262 144, logtxto 262 208
262 752	262 144	Imaginary (dispersion) trace, 65 536 × IEEE float32
524 896	300	Log text block, CR/LF-delimited KEY=VALUE records

5.2 Acquisition parameters

The log text block records the acquisition conditions reproduced in Table 2. The instrument identifier is 4IR, the observed nucleus is ^{1}H, and the observe frequency of 60 528 437 Hz places the measurement on a 60 MHz benchtop permanent-magnet system. The detection mode is recorded as SIM, denoting simultaneous quadrature sampling, which is the case treated analytically by Starčuk and co-workers[9] and which fixes the dwell period at the reciprocal of the spectral width, 250 µs. With 65 536 points across 4000 Hz the digital resolution is 0.061036 Hz per point. The vendor-recorded zero-order phase is zero, from which it may be inferred that the file as supplied has not had an operator phase applied, or that any applied value was not written back to the log.

Table 2. Acquisition parameters as recorded in the SPC log text block, with derived quantities.

Parameter	Value	Note
INSTRUM	4IR	Instrument identifier
NUCLEUS	^{1}H	
NUCFREQ / SPECFREQ	60 528 437 Hz	Observe frequency, 60.528437 MHz
SW_HZ / SPECWIDTH	4000 Hz	Spectral width, equal to 66.0846 ppm
DETMODE	SIM	Simultaneous (complex) quadrature sampling
ACQSIZE / SPECSIZE	65 536	
PULSWD	5×10^{-6} s	Excitation pulse width; $2 \cdot pw/\pi$ = 3.18 µs
DELAY	5 s	Recycle delay
GAIN	30	Receiver gain
ORGN	TAL	Sample origin code
ACQDATE / ACQTIME	Sunday, 9 August 2026, 14:27:28	
ZERO_ORDER_PHASE	0	No vendor zero-order phase recorded
FID	FALSE	File holds a transformed spectrum, not an FID
Dwell period (derived)	250 µs	1/SW; equivalent to 360° of φ_1
Digital resolution (derived)	0.061036 Hz per point	SW / 65 536

One consequence of the recorded parameters constrains the whole analysis and is worth stating at the outset. Because the dwell period is 250 µs, and because a receiver dead time on an instrument of this class is of the order of tens of microseconds, the first-order phase term to be expected lies in the range of a few tens to somewhat over a hundred degrees. A fitted value of several hundred degrees would require a first-point delay approaching or exceeding a full dwell period, and a negative value would require the receiver to sample before the pulse. Both outcomes are excluded on physical grounds before any spectrum is examined.

5.3 Consistency of the stored dispersion channel

The stored imaginary trace was tested against the Hilbert transform of the stored real trace. The correlation coefficient is 0.999999 and the ratio of standard deviations is 1.000000, so the two traces are Hilbert conjugates to within numerical precision. This outcome admits two interpretations, and they are not equivalent. Either the underlying FID is causal and the stored pair is the genuine transform of a one-sided signal, in which case the relation is expected and carries no information; or the vendor generated the dispersion trace by Hilbert transformation of an already-processed real trace, in which case any offset or distortion present in the real channel has been propagated into the imaginary channel non-physically. Nothing in the file distinguishes the two cases. For the purposes of phase rotation the distinction does not alter the arithmetic, but it does bear on how much confidence a large first-order correction deserves, and it is recorded here rather than passed over.

6. Diagnosis of the Reported Anomaly

The spectrum was submitted with the observation that the region between 4 and 9 ppm containing the broad aromatic envelope appeared slightly out of phase while the aliphatic region below 4 ppm appeared correct. That description is consistent with a residual first-order term, and it would be natural to proceed directly to fitting one. Doing so without first testing the premise would have been a mistake, for the reasons developed below.

6.1 The dispersive pedestal

Figure 1(a) shows the real and imaginary traces over the aromatic window as supplied. The real trace carries an aromatic envelope reaching an intensity of 30.4 units at 6.83 ppm on a local baseline near 0.3 units. The imaginary trace over the same interval rises monotonically from 28 units at 9.5 ppm to 56 units at 5.0 ppm, and it does so because the dispersion component of a Lorentzian decays only as the reciprocal of the frequency offset whereas the absorption component decays as its square. The intense aliphatic resonances at 0.86 and 1.24 ppm, reaching 848 and 1440 units respectively, therefore contribute almost nothing to the real channel five ppm away but contribute a substantial and slowly varying pedestal to the imaginary channel.

The numerical consequence is severe and it explains the reported observation. Because the imaginary trace in this window is between thirty and fifty times larger than the real baseline, a rotation of one degree transfers approximately 0.7 units into the real channel, which is comparable to the entire local baseline; a rotation of ten degrees transfers about seven units, which is nearly a quarter of the aromatic peak height. The aromatic region is thus the part of the spectrum in which a small phase error is most visible, not because the error is larger there but because the dispersive reservoir available to be mixed in is larger there. This alone would produce the reported appearance without any frequency dependence of the phase whatsoever.

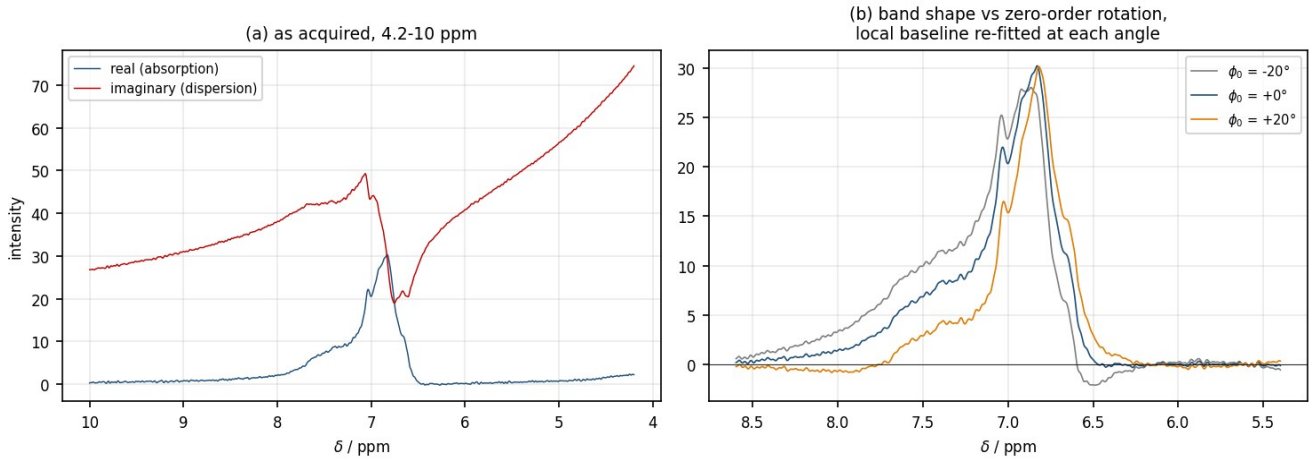

Figure 1. (a) Real and imaginary traces of the supplied spectrum between 4.2 and 10 ppm. The dispersion channel carries a pedestal of 28 to 56 units contributed by the intense aliphatic resonances, against a real baseline near 0.3 units. (b) Envelope of the 5.4 to 8.6 ppm band after zero-order rotation by −20°, 0° and +20°, with a local cubic baseline anchored at 8.6–9.6 ppm and 5.4–6.35 ppm refitted at each angle. The band shape is essentially invariant; what a rotation changes is the pedestal, not the resonance.

6.2 Is the aromatic band actually mis-phased?

A test capable of separating the two explanations is available. A dispersive admixture is antisymmetric about the band center, so it must alter the shape of the band, not merely its pedestal. If the pedestal is removed at each trial angle by refitting a local baseline anchored on both flanks, then any residual change in the band envelope with rotation is attributable to phase and nothing else. The result of that test is given in Table 3 and shown graphically in Figure 1(b).

Table 3. Effect of a pure zero-order rotation on the 6.4–8.4 ppm envelope, with a local cubic baseline anchored at 8.6–9.6 ppm and 5.4–6.35 ppm refitted at each angle. The asymmetry index is the normalized absolute difference between the two halves of the band reflected about its intensity-weighted centroid; it is zero for a perfectly even band.

Rotation φ₀ / degrees	Peak height	RMS change vs 0°	as % of peak height	Asymmetry index
−20	28.03	3.148	10.41	0.3642
−10	29.18	1.625	5.38	0.3830
−5	29.79	0.826	2.73	0.3957
0	30.23	0.000	0.00	0.4063
+5	30.48	0.853	2.82	0.3957
+10	30.55	1.732	5.73	0.3659
+20	30.16	3.567	11.80	0.2874

The asymmetry index remains between 0.29 and 0.41 across a forty-degree range of rotation. Were the band a single resonance carrying a dispersive admixture, the index would pass through a sharp minimum near zero at the correct angle; it does not approach zero anywhere. The asymmetry is therefore structural. The band rises gradually from 8.5 ppm to a maximum near 6.9 ppm and falls sharply below 6.6 ppm, which is the signature of a composite envelope in which several distinct aromatic environments contribute at different shifts, not of a mis-phased single line. Over the same range the RMS change in the envelope is under 12 % of the peak height even at ±20°, confirming that rotation principally moves the pedestal rather than reshaping the band.

This is the first substantive correction to the working hypothesis. The appearance that prompted the enquiry is dominated by a pedestal-and-baseline effect, and treating it as a phase problem alone would leave the analyst adjusting a parameter that is not the principal cause. That said, a genuine residual phase error is also present, as the next section establishes; the two effects coexist and must be separated rather than conflated.

6.3 Conditioning of the global objective functions

Before adopting any automatic criterion it is worth examining how sharply it determines the answer. The ACME objective, the integrated negative area, and a simultaneous phase-and-baseline objective were each evaluated over a grid spanning ±40° in φ_0 and ±360° in φ_1, with the spectrum decimated threefold and the pivot placed on the most intense resonance. The ACME surface possesses a defined minimum in φ_0 near −8° but is nearly flat along φ_1 over a range of at least ±100°, the objective changing by roughly one part in a thousand across that interval. The negative-area surface is worse: an extended region of the plane produces no negative intensity at all and therefore scores identically zero, so the criterion cannot discriminate within it. The simultaneous objective, when configured with a flexible Whittaker baseline, decreases monotonically toward the edge of the search domain and returns whatever bound it is given, which is the degeneracy predicted in Section 2.3 appearing in practice.

The conclusion is that on this spectrum the global objective functions determine φ_0 usefully and φ_1 hardly at all. This is not a criticism of the algorithms, which were designed for spectra containing resolved lines distributed across the sweep; it is a statement about the information content of a spectrum in which 97 % of the intensity lies within a 3 ppm interval. A method that measures phase where the signal actually is, and derives the first-order term from the separation between those measurements, is better conditioned, and that is the approach adopted.

7. Determination of the Phase Parameters

7.1 The regional method

The procedure adopted measures the optimal rotation independently within each of a small number of nominated bands and then fits a straight line through the resulting phase-versus-frequency points. Within each band the real trace is recomputed at every trial angle, a polynomial baseline is fitted to anchor windows on both flanks and subtracted, and the residual negative area within the band is scored. Refitting the local baseline at every trial is what makes the estimate robust: the slowly varying dispersive pedestal contributed by distant resonances is removed rather than being mistaken for a phase error, which is precisely the failure that defeats apex-phase estimators here. The construction is a bounded-window relative of the area-based measurement of Džakula[26] and of the multi-frequency measurement described by van Vaals and van Gerwen[21].

The polynomial order and the anchor windows were both varied — three orders crossed with three choices of each flanking window, giving twenty-seven determinations per band — and the spread of the resulting optima was taken as the uncertainty. This is a more honest error estimate than the convergence tolerance of the optimizer, which is smaller by an order of magnitude and measures nothing of interest.

7.2 Result and physical interpretation

The measurements are collected in Table 4. The aliphatic envelope between 0.4 and 3.0 ppm, whose intensity-weighted centroid falls at 1.253 ppm, requires a rotation of −6.25 ± 1.23°. The aromatic envelope between 6.4 and 8.4 ppm, centroid 7.111 ppm, requires +3.75 ± 0.42°. The two centroids are separated by 354.6 Hz, and the phase difference of 10.00° across that interval extrapolates to a first-order term of +112.8° across the full 4000 Hz spectral width. Weighted least squares through the two points, with the pivot placed at 1.2398 ppm, gives φ_0 = −6.27° and φ_1 = +112.80° with a standard uncertainty on φ_1 of 14.7°.

Table 4. Regional phase measurements and the derived first-order term. Uncertainties are the standard deviation over twenty-seven combinations of baseline polynomial order and anchor-window selection.

Band	Interval / ppm	Centroid / ppm	Offset / Hz	Optimal rotation / degrees
Aliphatic envelope	3.00 – 0.40	1.253	75.8	−6.25 ± 1.23
Aromatic envelope	8.40 – 6.40	7.111	430.4	+3.75 ± 0.42

Band	Interval / ppm	Centroid / ppm	Offset / Hz	Optimal rotation / degrees
Difference	—	—	354.6	+10.00

Applying the identity of Section 2.1 in reverse, a first-order term of 112.80° across a 4000 Hz sweep corresponds to an effective first-point delay of 78.3 µs. That figure is entirely physical. It is 31 % of one dwell period, it exceeds the pulse-width correction of 3.18 µs by a wide margin, and the residual of approximately 75 µs is a reasonable receiver dead time and filter group delay for a benchtop permanent-magnet spectrometer. The standard uncertainty of 14.7° on φ_1 corresponds to ±10 µs on t_0. The internal consistency of this result — a phase determination made purely from spectral shape returning a delay that matches what the hardware would be expected to exhibit — is the strongest evidence available that the fitted parameters are meaningful rather than merely optimal.

7.3 Comparison against the published criteria

The same spectrum was phased using each of the objective functions implemented in the accompanying module, with the identical bound on φ_1 corresponding to one full dwell period, and each result was followed by the identical baseline treatment so that the comparison isolates the phase criterion. The outcomes are given in Table 5, and the implied first-point delays are plotted in Figure 6(a).

Table 5. Phase parameters returned by each objective function on the same spectrum, with the identical first-order bound and identical subsequent baseline treatment. The aromatic hydrogen index used throughout this document is defined as the integral from 6 to 9 ppm expressed as a percentage of the sum of the integrals over 0–4, 4.5–6 and 6–9 ppm; the 4.0–4.5 ppm interval is excluded from the normalization because it carries the glyceride and heteroatom-adjacent methylene region, and its inclusion would change the index by 0.011 percentage points.

Engine	φ_0 / deg	φ_1 / deg	Implied t_0 / µs	Physically admissible	Aromatic H index / %
regional	−6.27	+112.80	78.3	yes	2.264
acme	−8.52	−74.99	−52.1	no	2.543
sinc	+2.20	+342.58	237.9	marginal	1.389
dispmin	+170.84	+362.63	251.8	no (spectrum inverted)	1.463
negarea	+0.00	−141.63	−98.4	no	2.800

Three of the five criteria return a negative implied first-point delay. No amount of algorithmic sophistication makes that admissible: it asserts that the receiver sampled the signal before the pulse produced it. The dispersion-minimization engine additionally settles on a zero-order value near 171°, which is a near-inversion of the spectrum, and its negative-area fraction of 0.986 confirms that it has done so; the DISPA family requires isolated Lorentzian resonances, and on an envelope spectrum it has nothing to work with. The simultaneous engine converges close to a full dwell period, which is admissible in principle but sits at the extreme end of the plausible range and reflects the degeneracy discussed above rather than a genuine determination.

The analytical consequence is set out in the final column. Depending only on which published phasing criterion is applied, with everything else held constant, the derived aromatic hydrogen index of this sample ranges from 1.389 % to 2.800 %, a spread of a factor of two. It would be difficult to defend any of these numbers in a technical dispute without also being able to state why that particular criterion had been chosen, and the physical-admissibility test on t_0 is the only defensible basis for making that choice that emerged from this work.

8. Baseline Correction in Practice

8.1 Failure of noise-threshold classification

The point noise of the phased trace, estimated robustly from the median absolute deviation of the point-to-point difference, is 0.0108 intensity units against a maximum intensity of 1449, giving a signal-to-noise ratio of approximately 1.3×10^5. At that level no noise-threshold classifier functions. The Golotvin–Williams window span

test[50], applied with a threshold of three standard deviations over sixty-four-point windows, accepts 6.0 % of the analytical window and, when tightened, accepts under 1 %; the Dietrich-type derivative cross-check[48] rejects nearly everything, because the reference deviation it computes is dominated by the vast empty wings of the 66 ppm sweep. Every region within the analytical window carries structure many hundreds of times the noise, and the classifier correctly reports that none of it is indistinguishable from noise.

The correct response is not to loosen the threshold until something is accepted, which merely converts a well-posed test into an arbitrary one. It is to recognize that the question being asked is wrong. What matters for baseline correction is not whether a region is at the noise level but whether the analyst is prepared to treat it as containing no signal of interest, and that is a judgment about chemistry rather than about statistics. The module therefore provides an explicit anchor-window mode, in which the spectroscopist nominates the intervals to be treated as baseline, and a scale-free envelope-quantile alternative for use where no judgment is available. Both are used below, and they agree.

8.2 The anchoring convention dominates the result

Six baseline treatments were applied to the identically phased spectrum, differing only in which intervals were treated as baseline and in the functional form fitted through them. Table 6 records the outcome, and Figure 6(c) displays it.

Table 6. Effect of the baseline anchoring convention on derived integrals, all applied to the identically phased spectrum (φ_0 = −6.27°, φ_1 = +112.80°). Sets A to D use a monotone multipoint interpolant; sets E and F use a polynomial through the anchors of set B.

Baseline strategy	Anchor intervals / ppm	∫ 6–9 ppm	∫ 4.5–6 ppm	∫ 0–4 ppm	Aromatic H index / %
A outer only	11.5–9.3; −0.70 to −1.90	23 067	12 411	690 969	3.175
B + mid anchor	A plus 6.30–5.75	15 735	2 584	676 578	2.264
C + olefinic anchor	B plus 5.40–4.60	15 578	149	665 523	2.287
D envelope quantile	automatic, quietest 30 %	14 143	−57	650 499	2.128
E polynomial order 3	anchors of set B	14 111	5 374	670 723	2.044
F polynomial order 5	anchors of set B	14 002	3 129	676 599	2.018

The single decision that matters most is whether the interval between 5.75 and 6.30 ppm is treated as baseline. Including it moves the derived aromatic hydrogen index from 3.175 % to 2.264 %, a relative change of 29 %, and it moves the nominal olefinic integral between 4.5 and 6 ppm by a factor of five. Without that anchor the interpolant runs beneath the entire 4.5 to 9 ppm region and the aromatic band is integrated on top of an artificially depressed baseline. Once the anchor is included, three quite different strategies — the three-window set, the four-window set, and the fully automatic envelope quantile, which uses no analyst judgment at all — converge on 2.13 % to 2.29 %. That convergence is reassuring, and it identifies set B as a defensible default.

It should be recorded that treating 5.75 to 6.30 ppm as baseline removes a small amount of genuine intensity, since the Lorentzian wing of the aliphatic envelope does extend into that interval. The choice is therefore a convention rather than a truth, and its only real defense is that it is stated, reproducible, and applied identically across every sample in a set. For comparative work — a calibration series, a batch of production samples, an interlaboratory study — a stated convention applied uniformly is worth considerably more than an unstated one applied optimally to each spectrum in turn.

8.3 Functional form: spline against polynomial

Sets E and F of Table 6 fit polynomials of degree three and five through the same anchors used by set B. Both depress the aromatic integral relative to the multipoint interpolant, and the reason is visible in the residuals: a polynomial constrained only at the anchors is free between them, and the analytical window contains a single gap of nearly seven ppm across which no constraint exists. Raising the degree makes this worse rather than better. In an earlier stage of

this work, with baseline points classified automatically rather than nominated, raising the polynomial degree from five to nine reduced the residual at the classified points from 0.066 to 0.052 while inflating the aliphatic integral from 590 000 to 1 032 000 units, an increase of 75 %, because the fitted curve had excursed to −233 units beneath the principal resonances. Degree eleven produced an aliphatic integral of 5.79×10^{6}. The residual at the fitted points is therefore worthless as a figure of merit for a polynomial baseline, and any implementation reporting only that statistic is capable of reporting excellent agreement while destroying the result.

Within the spline family the choice of interpolant matters less but is not negligible. A natural cubic spline through the anchors of set B returns an aromatic content of 2.281 % against 2.264 % for the monotone PCHIP interpolant, a relative difference of 0.8 %, the cubic overshooting slightly in the wide gap. The monotone interpolant is preferred on the principle that a baseline model should not introduce an excursion the anchors do not support. Node placement proved more consequential than interpolant selection: when nodes were spaced more closely than the window over which their ordinates were averaged, adjacent nodes returned identical values, and the interpolant collapsed toward a straight line over that stretch. This is an easy defect to introduce and a difficult one to notice, and the implementation now ties the averaging half-window to the node spacing for that reason.

8.4 The corrected spectrum

Figure 2 presents the result of the complete procedure. Panel (c) shows the aromatic window at each stage and makes the sequence explicit: the phase rotation, taken alone, depresses the region below 8 ppm by roughly 3.7 units and lifts it above 5 ppm, because the frequency-dependent rotation mixes an increasing proportion of the dispersive pedestal into the real channel as the offset from the pivot grows. Panel (d) shows the baseline model that results. The correction required is −15.95 units at −1 ppm and −3.66 units at +10 ppm, against an aromatic band whose peak height is 30 units. The magnitude of that correction is the price of the first-order term, and it is entirely as anticipated: a φ_1 of 113° cannot be applied to this spectrum without incurring it.

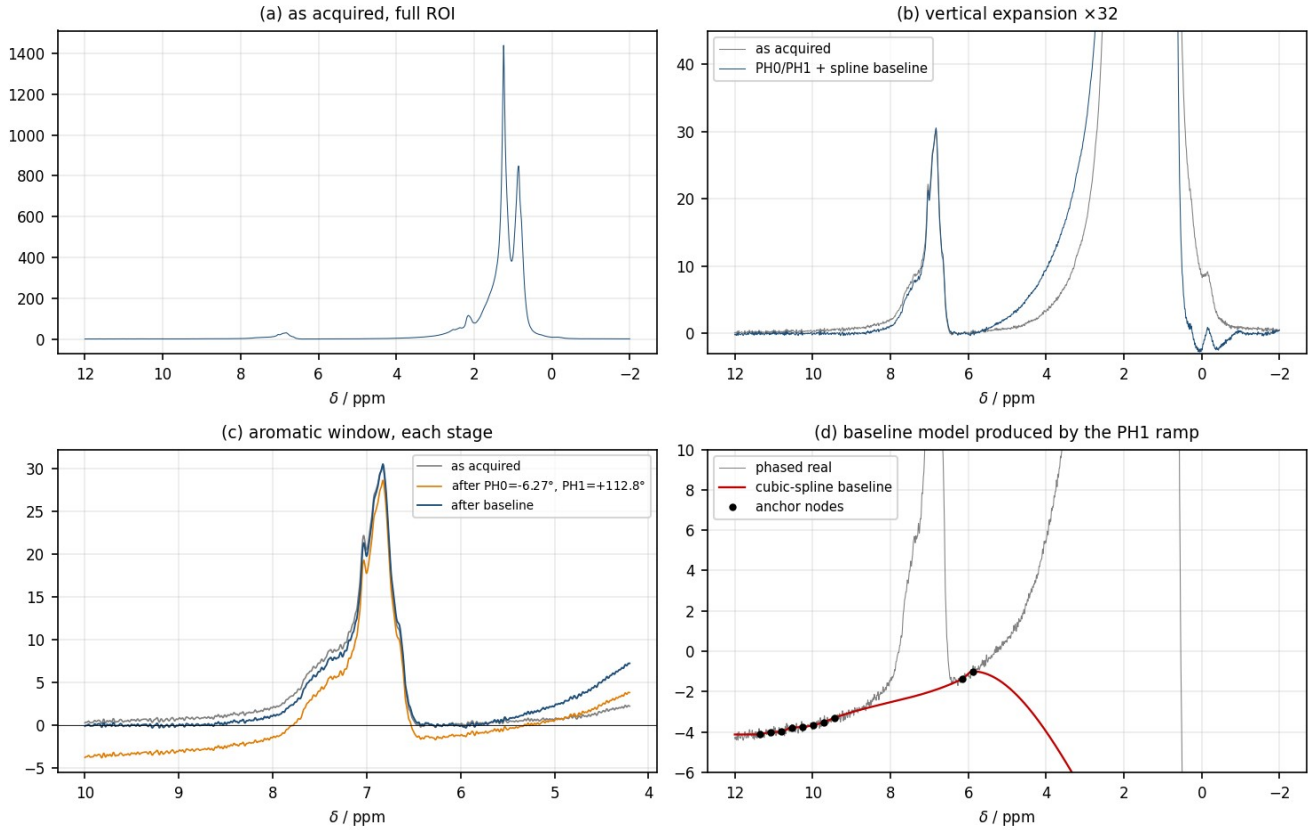

Figure 2. (a) The supplied spectrum over the analytical window. (b) The same, expanded thirty-two-fold, before and after the complete correction. (c) The aromatic window at each stage of processing, showing that the phase rotation alone introduces a pronounced tilt which the baseline step then removes. (d) The phased real trace with the monotone multipoint baseline model and its anchor nodes; the model must remove almost sixteen units at −1 ppm, which is half the height of the aromatic band.

9. Iterative Refinement Against the Baseline Flanks

The determination described in Section 7 measures phase inside the bands. An experienced spectroscopist phasing the same spectrum by hand does something different: the two baseline windows flanking the band of interest are watched, φ_1 is turned until both lie flat, φ_0 is turned until the intense reference band is restored, and the two adjustments are alternated until neither moves. That procedure was requested for this spectrum, with the flanks taken as 8.5 to 9.5 ppm above the aromatic envelope and 5.5 to 6.2 ppm below it, and it was implemented and run. It proved more instructive than expected, and it is reported here in full because the outcome bears directly on how much first-order phase such a criterion should be permitted to demand.

9.1 The observation

The initial reading of the spectrum was that the upper flank required roughly five degrees more rotation than the lower one before both appeared flat. That reading is confirmed by measurement. For a constant rotation the slope of the real trace within a window vanishes when tan φ equals the ratio of the gradients of the real and imaginary traces there, and evaluating that condition gives +5.28 degrees for the 8.5 to 9.5 ppm window and +2.10 degrees for the 5.5 to 6.2 ppm window. The visual estimate was therefore accurate to within a fraction of a degree.

The difficulty lies in what follows from it. The two flank centers are 9.000 and 5.850 ppm, separated by 190.7 Hz, which is 4.8 % of the 4000 Hz sweep. A phase difference of 3.18 degrees across that interval extrapolates to a first-order term of only +66.6 degrees across the full width, which is smaller than the +112.8 degrees obtained from the bands rather than larger. A criterion built on two closely spaced windows has a very short lever arm, and a small error in either measurement is magnified by the reciprocal of the separation.

9.2 Why the total flank slope is not the phase

Running the alternating loop against the observed slope nevertheless drives φ_1 to +288 degrees, not +67. The reason is worth setting out algebraically, because it is a general property of first-order phase correction and not a peculiarity of this data set. After a ramp has been applied the real trace is R cos φ − I sin φ, and differentiating with respect to frequency gives exactly two terms:

$$dr/d\nu \;=\; [\,R'\cos\varphi - I'\sin\varphi\,] \;-\; [\,R\sin\varphi + I\cos\varphi\,]\cdot\varphi'$$

The first bracket vanishes when the local phase is correct; it is the quantity a phase correction is entitled to null. The second is proportional to the ramp rate φ′ multiplied by the local signal amplitude, and it is nonzero whenever φ_1 is nonzero no matter how perfectly the spectrum is phased. In the flank windows examined here the mean modulus of the dispersion trace is 31 units against 0.75 units for the absorption trace, a ratio of 41 to 1, so the second term is large wherever the dispersive pedestal is large. A criterion that nulls the total slope is therefore not selecting the φ_1 that makes the phase correct; it is selecting the φ_1 at which the two terms happen to cancel.

Table 7 gives the decomposition at the two solutions. At +112.8 degrees the genuine phase terms are +0.16 and −0.14, essentially zero, while the ramp terms of −0.92 and −1.28 leave the flanks visibly tilted. At +288 degrees the genuine terms have grown to +2.04 and +3.26 and the ramp terms to −2.11 and −3.18, so the totals are −0.07 and +0.08 and the flanks look flat. The flatness at +288 degrees is a cancellation of two substantial errors, and the price is a local phase of +28 degrees at the upper flank and +19.8 degrees at the band centroid.

Table 7. Decomposition of the flank slope into the genuine phase term and the ramp term at the two candidate first-order values, with φ_0 at its converged value in each case. A flank is genuinely correctly phased only when the first column is zero.

Flank window	φ₁ / deg	Local phase / deg	Genuine term	Ramp term	Total slope
8.5–9.5 ppm	+112.8	+6.72	+0.156	−0.921	−0.765
5.5–6.2 ppm	+112.8	+1.55	−0.136	−1.283	−1.419
8.5–9.5 ppm	+288.1	+28.08	+2.042	−2.111	−0.069
5.5–6.2 ppm	+288.1	+14.34	+3.257	−3.180	+0.078

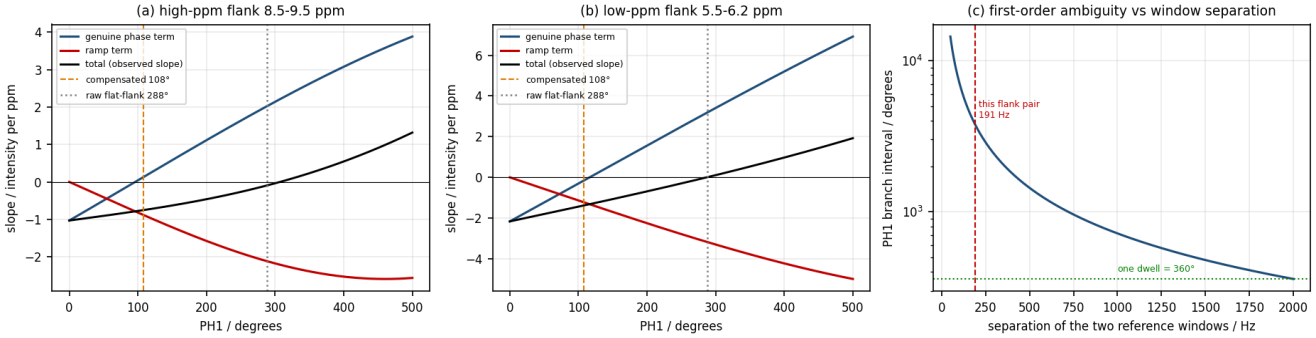

Figure 3. (a, b) The two contributions to the slope within each flank as a function of the first-order term, at the converged zero-order value. The total crosses zero near +288 degrees only because the genuine and ramp contributions cancel there; the genuine term alone crosses near +108 degrees. (c) Degrees of first-order phase corresponding to one 180-degree branch ambiguity, as a function of the separation of the two reference windows. For the flank pair used here the interval is 3776 degrees.

9.3 The branch ambiguity

A slope-nulling phase is defined only modulo 180 degrees, since rotating a further half turn inverts the trace without changing where its derivative vanishes. A pair of reference windows separated by a frequency interval therefore fixes the first-order term only modulo 180 degrees times the ratio of the sweep to that interval. For the flanks used here the interval is 3776 degrees, and the neighboring solutions correspond to first-point delays of −2576 and +2669 microseconds against a dwell period of 250 microseconds. Neither is remotely admissible, so in this instance the ambiguity is harmless; but the same calculation carried out with more closely spaced windows, or on a wider sweep, can leave several branches inside the physically plausible range, and it should be evaluated rather than assumed away. The routine reports the interval alongside every result for that reason.

The bound on the first-order term was widened at the same time, and is now expressed in dwell periods rather than in degrees so that the physical meaning stays visible. Permitting up to four dwell periods, or 1440 degrees, changes nothing: the alternating loop still converges to +288 degrees, because the next branch of the flatness criterion lies far outside even that range. The extra freedom is available where an instrument genuinely warrants it — a console applying substantial digital filtering will need it — but on this spectrum it is not the constraint that was limiting the answer.

9.4 The alternating loop and its convergence

Two modes were implemented. In the raw mode the first-order step nulls the total flank slope, which is the criterion as ordinarily practiced. In the compensated mode it nulls the genuine phase term alone, leaving the ramp-induced tilt to be removed by the baseline step where it belongs. Both alternate with a zero-order step that restores the aliphatic envelope, measured by the negative-area criterion of Section 7 with its local baseline refitted, and the pivot is placed on that envelope so that the coupling between the two steps is weak.

Convergence is rapid in both modes, as Figure 4(a) shows. The compensated mode settles in four cycles at $\varphi_0 = -6.00$ and $\varphi_1 = +108.30$ degrees, corresponding to a first-point delay of 75.2 microseconds. The raw mode settles in five cycles at $\varphi_0 = -5.75$ and $\varphi_1 = +288.06$ degrees, a delay of 200.0 microseconds or four fifths of a dwell period. The zero-order term moves by only half a degree between the two, which confirms that the pivot placement has done its job; the

anticipated need to re-adjust the zero order after a large first-order change is real but small here, and the loop absorbs it automatically.

The compensated result deserves emphasis. It was obtained from the flanks alone, by a criterion that never looks inside a band, and it lands within four degrees of the +112.8 degrees obtained in Section 7 from the bands alone, by a criterion that never looks at a baseline. Two independent measurements of the same quantity agreeing to that tolerance is considerably stronger evidence than either taken singly.

9.5 What the flanks actually require

The remaining question is whether a first-order term of nearly 300 degrees is needed to deliver flat flanks in the finished spectrum. It is not, and the demonstration is straightforward. When the flank windows are themselves adopted as baseline anchors and the multipoint interpolant is given enough nodes to follow them, the residual flank slope falls to the noise level for every candidate phase. Figure 4(b) records the node-count scan: with fourteen nodes the compensated solution leaves the flanks tilted by 0.041 and 0.311 units per ppm, but by thirty nodes the figures are 0.016 and 0.025 against a noise-equivalent threshold of 0.032, and the aromatic integral is stable to a few tenths of a percent anywhere between twenty-four and forty-eight nodes. Beyond about sixty nodes the interpolant begins to follow the noise and the integral drifts upward again.

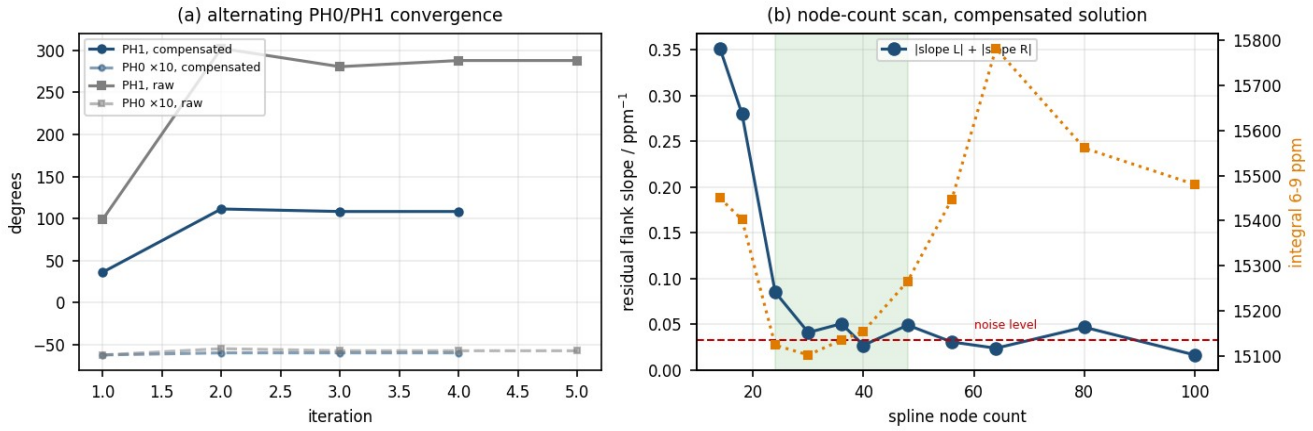

Figure 4. (a) Convergence of the alternating zero-order and first-order refinement in both modes; the zero-order trace is magnified tenfold to make it visible on the same axis. (b) Node-count scan for the compensated solution: the sum of the absolute residual flank slopes against the number of spline nodes, with the aromatic integral overlaid. The shaded band marks the range within which the flanks are flat to the noise and the integral is stable.

Table 8. The four candidate treatments carried through the identical baseline procedure — flank-anchored multipoint interpolant with forty nodes. Flank slopes are measured after baseline correction; the noise-equivalent threshold is 0.032 units per ppm.

Treatment	φ_0 / deg	φ_1 / deg	t_0 / µs	slope 8.5–9.5	slope 5.5–6.2	φ at 7.1 ppm	∫ 6–9 ppm	Aromatic H index / %
No phase correction	0.00	0.00	0.0	−0.026	−0.043	0.00	15 686	2.393
Regional, band criterion	−6.27	+112.80	78.3	−0.009	−0.018	+3.75	15 132	2.184
Iterative, compensated	−6.00	+108.30	75.2	−0.010	−0.017	+3.62	15 154	2.192
Iterative, raw flat-flank	−5.75	+288.06	200.0	+0.004	−0.031	+19.84	12 781	1.851

Every treatment in Table 8 delivers flat flanks. What separates them is what happens to the band between the flanks. The raw flat-flank solution costs 15.7 % of the aromatic integral relative to the compensated solution, drives the minimum of the corrected band from −0.28 to −0.55 units, and reduces the asymmetry index from 0.41 to 0.29 by skewing the envelope rather than by symmetrizing it. That last point is the decisive one: as shown in Section 6.2 the asymmetry index of this band falls monotonically with rotation and passes through no minimum, so a rotation that

lowers it is over-rotating an already asymmetric composite, not correcting a dispersive admixture. The appearance of a negative lobe on one side of the band is the classical signature of over-phasing, and it appears here.

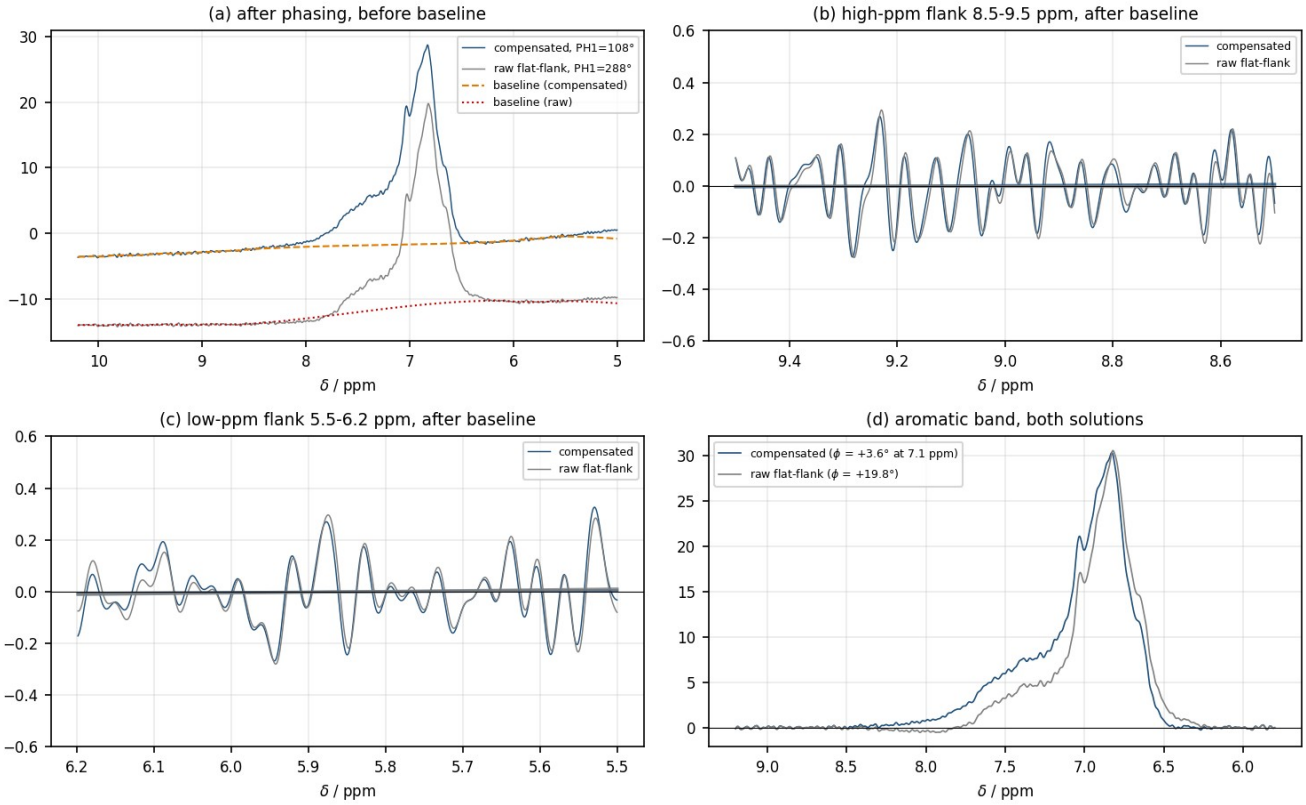

Figure 5. (a) The phased traces and their baseline models before baseline subtraction; the raw flat-flank solution sits some fourteen units lower and carries a much flatter pedestal. (b, c) The two flank windows after baseline correction with forty nodes, with the fitted straight line overlaid; both solutions are flat to the noise. (d) The corrected aromatic band. The raw flat-flank solution has lost the shoulder between 7.2 and 8.0 ppm, which is where the 15.7 % of integral has gone.

The conclusion is that the flank-flatness criterion, applied to the total slope, is a reliable guide only where the dispersive pedestal is small compared with the signal. On a spectrum such as this one, where an intense aliphatic envelope projects a dispersion tail forty times the local absorption baseline into the aromatic window, the criterion is systematically biased toward large first-order terms and should be used in its compensated form. The tilt it is trying to remove is real, but it is a baseline artifact created by the ramp, and the baseline step removes it without charge; removing it with the first-order term instead charges 15.7 % of the aromatic integral for the same cosmetic result.

10. Consequences for Quantitation

The purpose of correcting a spectrum is to obtain numbers from it, and the appropriate test of a correction is what it does to those numbers rather than how the display looks. Table 9 accumulates the effect of each stage on the integrals of interest, and Figure 6 summarizes the sensitivities.

Table 9. Cumulative effect of each correction stage. All rows use the anchor set B baseline of Section 8.2 with a monotone multipoint interpolant and fourteen nodes, so that differences between rows are attributable to the phase treatment alone except where noted. The anchoring differs from that adopted in Section 9, which uses the flanks themselves as anchors with forty nodes; the two conventions are compared in Table 8 and the integrals are not interchangeable between the two tables.

Treatment	φ_0 / deg	φ_1 / deg	Baseline at −1 ppm	Baseline at +10 ppm	∫ 6–9 ppm	∫ 0–4 ppm	Aromatic H index / %
As acquired, baseline only	0.00	0.00	+0.88	+0.37	16 638	639 828	2.533

Treatment	φ₀ / deg	φ₁ / deg	Baseline at −1 ppm	Baseline at +10 ppm	∫ 6–9 ppm	∫ 0–4 ppm	Aromatic H index / %
Zero order only	−6.27	0.00	−9.68	+3.28	17 161	674 589	2.470
Zero and first order, band criterion	−6.27	+112.80	−15.95	−3.66	15 735	676 578	2.264
Band criterion, natural cubic spline	−6.27	+112.80	−15.78	−3.66	16 149	689 333	2.281
φ₁ at −1 standard uncertainty	−6.27	+98.10	−15.13	−2.77	15 987	676 333	2.300
φ₁ at +1 standard uncertainty	−6.27	+127.50	−16.76	−4.56	15 455	676 820	2.224

Reading down the table, the phase correction as a whole moves the aromatic hydrogen index from 2.533 % to 2.264 %, a relative change of 10.6 %. Of that, the zero-order term accounts for a decrease of 2.5 % relative and the first-order term for the remaining 8.3 %. The uncertainty in the first-order term, propagated by recomputing the entire chain at φ₁ ± 14.7°, moves the result between 2.300 % and 2.224 %, which is ±1.7 % relative. The choice between a monotone and a natural cubic interpolant contributes 0.8 % relative. Set against these, the anchoring convention of Table 6 spans 2.018 % to 3.175 %, or 45 % of the mean.

Substituting the adopted first-order term of Section 9, +108.30° rather than +112.80°, moves the index in this table by 0.008 percentage points, which is smaller than every other entry in the comparison and is the reason the revision does not disturb the conclusions drawn here. The hierarchy is therefore unambiguous, and it inverts the attention these steps ordinarily receive. The baseline convention dominates; the phase correction matters at the ten-percent level; the uncertainty in the first-order term, which is what a phasing algorithm actually determines, matters at under two percent. For a laboratory establishing a method for aromatic hydrogen determination at low field, the effort is far better spent fixing and documenting the integration and baseline conventions than in selecting among phasing algorithms. This conclusion is consistent with the qNMR uncertainty-budget literature[59,61] and with the analysis of discrete-integration bias[58], though it is stated here rather more strongly because the low-field envelope case is more severe than the high-field cases those studies examined.

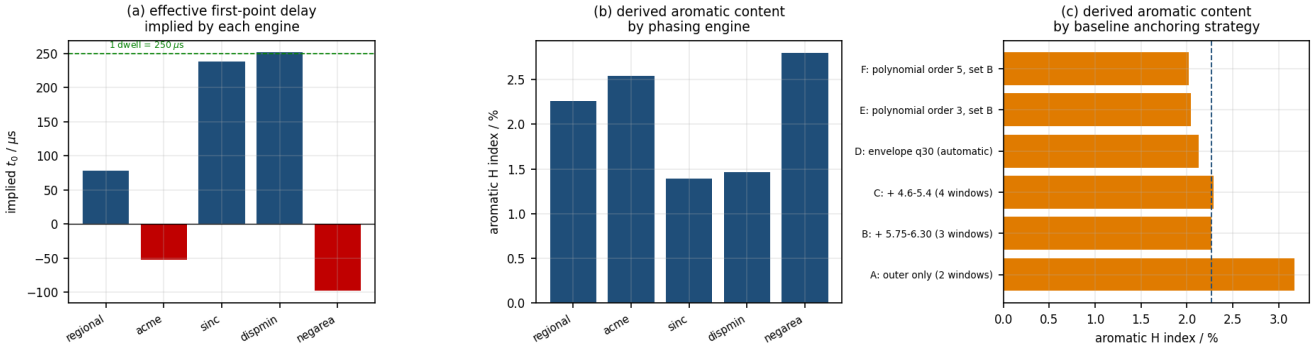

Figure 6. (a) Effective first-point delay implied by the first-order term returned by each phasing engine; bars below zero are physically inadmissible, and the broken line marks one dwell period. (b) Derived aromatic hydrogen index as a function of the phasing engine, all other processing held constant. (c) Derived aromatic hydrogen index as a function of the baseline anchoring strategy, phase held constant; the broken line marks the adopted value.

One further caution belongs here. Where spectra are reduced to an X matrix for multivariate modeling, a systematic baseline convention applied uniformly across the set is largely absorbed by mean-centering and by the latent structure of the model, whereas an inconsistent convention introduces variance that correlates with nothing chemical and will be modeled as though it were real[55]. In a calibration transfer or an interlaboratory comparison the same inconsistency appears as a bias between instruments. The processing chain, once fixed, should be applied by script rather than interactively for exactly this reason, and the module described below was written to make that practical.

11. The Python Implementation

The accompanying module, nmr_autophase.py, is a single self-contained file requiring only numpy and scipy, with openpyxl and matplotlib needed solely for the optional export and plotting utilities. It deliberately avoids a dependency on nmrglue[44] so that the SPC handling can be inspected and so that files written by it round-trip into GRAMS/AI and into the SPC X-matrix chemometrics workflow without an intermediate conversion.

11.1 File handling

The reader parses the little-endian new-format SPC header directly, handles sixteen-bit, thirty-two-bit and IEEE-float ordinate storage, and recovers the dispersion trace from the log binary area when one is present. Where no stored dispersion channel exists the trace is regenerated by discrete Hilbert transform, and the consistency between a stored and a reconstructed channel is reported so that the caveat of Section 5.3 can be evaluated rather than ignored. The writer preserves the source header and the complete log text, writes ordinates as IEEE floats so that no requantization occurs, restores the corrected dispersion trace to the log binary area so that the output remains a complete complex data set, and appends the applied parameters to the log text as PROC_PH0, PROC_PH1, PROC_PIVOT_PPM, PROC_PHASE_ENGINE and PROC_BASELINE records. The corrected file supplied with this document was verified to round-trip with a maximum ordinate difference of exactly zero in both channels, and the appended records are readable by the same parser.

11.2 Interface

Table 10. Principal entry points of nmr_autophase.py.

Function	Purpose
`read_spc(path)`	Parse a new-format SPC file; returns real and imaginary traces, the abscissa, the decoded log parameters and the verbatim headers
`write_spc(spec, path, real, imag, append_log)`	Write a corrected SPC preserving header and log text, with the dispersion trace restored to the log binary area
`hilbert_partner(real)`	Regenerate the dispersion trace from the absorption trace
`hilbert_consistency(real, imag)`	Report correlation and scale ratio between a stored and a reconstructed dispersion trace
`ph1_bound_from_timing(sw_hz, max_delay_s)`	Return the physically defensible bound on $\lvert\varphi_1\rvert$ in degrees
`flank_slope_terms(x, z, window, ph0, ph1, ...)`	Decompose the slope inside a baseline window into its genuine-phase and ramp contributions, and report the slope-nulling phase
`ph1_branch_interval(hi_ppm, lo_ppm, sf, sw)`	Degrees of φ_1 corresponding to one 180° ambiguity between two reference windows
`autophase_iterative(x, z, flanks, anchor_band, ...)`	Alternating zero-order / first-order refinement against the baseline flanks; mode 'compensated' or 'raw', bound set in dwell periods
`region_optimal_phase(x, z, band, anchors)`	Measure the optimal rotation within one band, with an uncertainty from anchor and order variation
`autophase_regional(x, z, bands, ...)`	Fit φ_0 and φ_1 from two or more regional measurements; reports the implied first-point delay
`autophase(z, engine, ...)`	Global objective-function phasing; engines acme, negarea, baseline, dispmin, sinc
`apply_phase_ppm(x, z, ph0, ph1, pivot_ppm, sf, sw)`	Apply a frequency-referenced phase ramp
`anchors_to_mask(x, anchors)`	Build a baseline mask from analyst-nominated ppm windows
`envelope_baseline_mask(r, quantile)`	Scale-free automatic baseline classification by slope envelope
`classify_baseline_points(r, ...)`	Golotvin–Williams window span test with a Dietrich derivative cross-check
`baseline_correct(x, r, method, ...)`	Baseline estimation; methods spline (natural or PCHIP), poly, whittaker, airpls
`spectrum_metrics(x, r, mask, regions, roi_ppm)`	Figures of merit and region integrals restricted to the analytical window

Function	Purpose
`process_file(path, ...)`	End-to-end driver; also available as a command-line entry point

A representative invocation for the spectrum discussed here nominates the two bands, bounds the first-order term at one dwell period, and applies the set B anchors with a monotone interpolant:

```
python nmr_autophase.py TalFFTAVG00019.spc --engine regional \
       --baseline spline --nodes 14 --max-delay-us 250 \
       --pivot-ppm 1.2398 --region 12 -2 -o TalFFTAVG00019_phased
```

11.3 Design decisions that proved necessary

Several features of the implementation exist because their absence produced wrong answers during development, and they are recorded because the same traps will be met by anyone writing equivalent code. The first-order term is bounded by the acquisition timing rather than by a convenient constant, because otherwise the simultaneous engine runs to whichever bound it is given. The baseline model used inside the phase objective is a stiff low-order Chebyshev fit rather than a flexible smoother, and a ridge penalty on φ_1 is added, because a flexible baseline and a linear phase ramp span nearly the same function space. The averaging half-window for spline node ordinates follows the node spacing, because otherwise adjacent nodes return identical values. The airPLS weight update is clamped, because the published exponential overflows in double precision when a single resonance is three orders of magnitude taller than the baseline. Baselines are held flat outside the classified span rather than extrapolated, because both cubic splines and high-order polynomials diverge rapidly beyond their outermost constraint. Figures of merit are computed within the analytical window, because on a 66 ppm sweep the empty wings otherwise dominate every summary statistic and make quite different corrections appear identical.

12. Recommendations and Limitations

For spectra of this class the following procedure is recommended. The dispersion channel should be taken from the file where one exists and reconstructed only when it does not, with the consistency of any reconstruction reported. The first-order term should be bounded before optimization by the identity $\varphi_1 = 360° \cdot SW \cdot t_0$, using a maximum delay consistent with the instrument rather than a convenient number, and any fitted value implying a negative delay should be rejected outright. Phase should be measured where signal exists, in two or more bands separated as widely as the spectrum permits, with a local baseline refitted at every trial angle; the uncertainty should be taken from the spread over anchor and order variation rather than from the optimizer. Baseline anchors should be nominated explicitly, recorded, and applied identically to every spectrum in a set. A monotone multipoint interpolant should be preferred to a natural cubic spline and both to a polynomial. Where the flanks either side of a band are used to judge the phase, the observed slope should be decomposed before it is nulled, because on a spectrum carrying a substantial dispersive pedestal the ramp contribution will otherwise be mistaken for a phase error and a first-order term two to three times too large will be adopted; the tilt that remains after the genuine term has been nulled belongs to the baseline step, and a multipoint interpolant with sufficient node density removes it without touching the integral. Finally, the derived analytical quantity should be recomputed at the bounds of the phase uncertainty and under at least one alternative anchoring convention, and the resulting spread reported alongside the result.

The limitations of the present work should be equally explicit. The analysis rests on a single spectrum, and while the mechanisms identified are general the numerical values are not transferable to another sample or instrument. The regional method requires the analyst to nominate bands, and a poor nomination will produce a confident and wrong answer; it is better conditioned than the global criteria on this spectrum but it is not automatic in the sense that ACME is. Only two bands were available here, and the two baseline flanks used in Section 9 are separated by less than five per cent of the sweep, so in both cases the straight line through the phase-versus-frequency points is exactly determined with no redundancy and no test of linearity; a combined weighted fit over all six reference windows returns a reduced chi-square of 8.2, which states plainly that a two-parameter linear phase model does not account for the scatter among

the individual measurements and that the first-order term is correspondingly less well determined than any single fit suggests; a spectrum offering three or more usable bands would permit the linearity of the phase model to be checked rather than assumed, and that check should be made wherever the data allow. Because the supplied file contains a transformed spectrum rather than an FID, the superior remedies of backward linear prediction[14,15] and of first-point scaling[11] could not be applied; if the raw FID can be retained by the acquisition software, repairing the initial points before transformation would address the cause and would reduce both the first-order term and the baseline roll it creates. Higher-order phase terms were assumed negligible and not tested[39]. The identity of the sample was not disclosed and no assignment beyond the generic aliphatic and aromatic regions has been attempted; the aromatic envelope has been treated as a composite on the evidence of its angle-independent asymmetry, but confirming that interpretation would require either a higher-field measurement or a deconvolution against known components.

A reviewer or opposing expert would most likely press on two points. The first is the treatment of the 5.75 to 6.30 ppm interval as baseline, since that single convention moves the reported aromatic content by 29 % relative and it is a judgment rather than a measurement; the defense available is that three independent strategies including a fully automatic one converge once it is adopted, and that the convention is stated rather than concealed. The second is the adoption of a first-order term that the global objective functions do not support, since ACME and the negative-area criterion both prefer the opposite sign. The defense there is stronger: those criteria return physically inadmissible first-point delays on this spectrum and are demonstrably ill-conditioned in φ_1, whereas the adopted value derives from measurements made where signal exists and returns a delay consistent with the hardware. Both defenses depend on the sensitivity analysis having been performed and reported, which is the practical argument for performing it.

13. References

Every digital object identifier below was resolved against the Crossref REST API and the returned title, author list, container title, volume, issue, pagination and year compared against the citation as written. Items that could not be verified in that way were discarded rather than estimated. Two records required correction against the publisher metadata: the title of reference 11 is deposited in corrupted form by the publisher and is given here in its correct printed form, and the author names of reference 9 carry diacritics that the deposited record omits.

1. Ernst, R. R.; Anderson, W. A. Application of Fourier Transform Spectroscopy to Magnetic Resonance. Rev. Sci. Instrum. 1966, 37 (1), 93–102. DOI: 10.1063/1.1719961.
2. Ernst, R. R. Numerical Hilbert Transform and Automatic Phase Correction in Magnetic Resonance Spectroscopy. J. Magn. Reson. 1969, 1 (1), 7–26. DOI: 10.1016/0022-2364(69)90003-1.
3. Bartholdi, E.; Ernst, R. R. Fourier Spectroscopy and the Causality Principle. J. Magn. Reson. 1973, 11 (1), 9–19. DOI: 10.1016/0022-2364(73)90076-0.
4. Hoch, J. C.; Stern, A. S. NMR Data Processing; Wiley-Liss: New York, 1996; 196 pp. ISBN 978-0-471-03900-6.
5. Hoult, D. I.; Chen, C.-N.; Sank, V. J. Quadrature Detection in the Laboratory Frame. Magn. Reson. Med. 1984, 1 (3), 339–353. DOI: 10.1002/mrm.1910010305.
6. Delsuc, M. A.; Lallemand, J. Y. Improvement of Dynamic Range in NMR by Oversampling. J. Magn. Reson. 1986, 69 (3), 504–507. DOI: 10.1016/0022-2364(86)90164-2.
7. Moskau, D. Application of Real Time Digital Filters in NMR Spectroscopy. Concepts Magn. Reson. 2002, 15 (2), 164–176. DOI: 10.1002/cmr.10031.
8. Zhu, G.; Torchia, D. A.; Bax, A. Discrete Fourier Transformation of NMR Signals. The Relationship between Sampling Delay Time and Spectral Baseline. J. Magn. Reson., Ser. A 1993, 105 (2), 219–222. DOI: 10.1006/jmra.1993.1274.
9. Starčuk, Z.; Bartušek, K.; Starčuk, Z., Jr. First-Data-Point Problem and the Baseline Distortion in Fourier-Transform NMR Spectroscopy with Simultaneous Sampling. J. Magn. Reson., Ser. A 1994, 108 (2), 177–188. DOI: 10.1006/jmra.1994.1108.
10. Schaefer, J.; Stejskal, E. O. Baseline Artifacts in High-Resolution Fourier Transform NMR Spectra. J. Magn. Reson. 1974, 15 (1), 173–176. DOI: 10.1016/0022-2364(74)90189-9.
11. Otting, G.; Widmer, H.; Wagner, G.; Wüthrich, K. Origin of t_1 and t_2 Ridges in 2D NMR Spectra and Procedures for Suppression. J. Magn. Reson. 1986, 66 (1), 187–193. DOI: 10.1016/0022-2364(86)90122-8.
12. Marion, D.; Bax, A. Baseline Distortion in Real-Fourier-Transform NMR Spectra. J. Magn. Reson. 1988, 79 (2), 352–356. DOI: 10.1016/0022-2364(88)90230-2.
13. Tang, C. G. An Analysis of Baseline Distortion and Offset in NMR Spectra. J. Magn. Reson., Ser. A 1994, 109 (2), 232–240. DOI: 10.1006/jmra.1994.1160.
14. Barkhuijsen, H.; de Beer, R.; Bovée, W. M. M. J.; van Ormondt, D. Retrieval of Frequencies, Amplitudes, Damping Factors, and Phases from Time-Domain Signals Using a Linear Least-Squares Procedure. J. Magn. Reson. 1985, 61 (3), 465–481. DOI: 10.1016/0022-2364(85)90187-8.
15. Zhu, G.; Bax, A. Improved Linear Prediction for Truncated Signals of Known Phase. J. Magn. Reson. 1990, 90 (2), 405–410. DOI: 10.1016/0022-2364(90)90150-8.
16. Marshall, A. G.; Roe, D. C. Dispersion versus Absorption: Spectral Line Shape Analysis for Radiofrequency and Microwave Spectrometry. Anal. Chem. 1978, 50 (6), 756–763. DOI: 10.1021/ac50027a023.
17. Sotak, C. H.; Dumoulin, C. L.; Newsham, M. D. Automatic Phase Correction of Fourier Transform NMR Spectra Based on the Dispersion versus Absorption (DISPA) Lineshape Analysis. J. Magn. Reson. 1984, 57 (3), 453–462. DOI: 10.1016/0022-2364(84)90260-9.
18. Herring, F. G.; Phillips, P. S. Automatic Phase Correction in Magnetic Resonance Using DISPA. J. Magn. Reson. 1984, 59 (3), 489–496. DOI: 10.1016/0022-2364(84)90083-0.
19. Craig, E. C.; Marshall, A. G. Automated Phase Correction of FT NMR Spectra by Means of Phase Measurement Based on Dispersion versus Absorption Relation (DISPA). J. Magn. Reson. 1988, 76 (3), 458–475. DOI: 10.1016/0022-2364(88)90350-2.

20. Brown, D. E.; Campbell, T. W.; Moore, R. N. Automated Phase Correction of FT NMR Spectra by Baseline Optimization. J. Magn. Reson. 1989, 85 (1), 15–23. DOI: 10.1016/0022-2364(89)90315-6.

21. van Vaals, J. J.; van Gerwen, P. H. J. Novel Methods for Automatic Phase Correction of NMR Spectra. J. Magn. Reson. 1990, 86 (1), 127–147. DOI: 10.1016/0022-2364(90)90216-V.

22. Heuer, A. A New Algorithm for Automatic Phase Correction by Symmetrizing Lines. J. Magn. Reson. 1991, 91 (2), 241–253. DOI: 10.1016/0022-2364(91)90189-Z.

23. Sterna, L. L.; Tong, V. P. Automated Method of Phasing Difficult Nuclear Magnetic Resonance Spectra with Application to the Unsaturated Carbon Analysis of Oils. Fuel 1991, 70 (8), 941–945. DOI: 10.1016/0016-2361(91)90048-F.

24. Balacco, G. A New Criterion for Automatic Phase Correction of High-Resolution NMR-Spectra Which Does Not Require Isolated or Symmetrical Lines. J. Magn. Reson., Ser. A 1994, 110 (1), 19–25. DOI: 10.1006/jmra.1994.1175.

25. Koehl, P.; Ling, C.; Lefèvre, J. F. Automatic Phase Correction of NMR Spectra: Statistics and Limits. J. Chim. Phys. 1995, 92, 1929–1938. DOI: 10.1051/jcp/1995921929.

26. Džakula, Ž. Phase Angle Measurement from Peak Areas (PAMPAS). J. Magn. Reson. 2000, 146 (1), 20–32. DOI: 10.1006/jmre.2000.2123.

27. Chen, L.; Weng, Z.; Goh, L.; Garland, M. An Efficient Algorithm for Automatic Phase Correction of NMR Spectra Based on Entropy Minimization. J. Magn. Reson. 2002, 158 (1–2), 164–168. DOI: 10.1016/S1090-7807(02)00069-1.

28. de Brouwer, H. Evaluation of Algorithms for Automated Phase Correction of NMR Spectra. J. Magn. Reson. 2009, 201 (2), 230–238. DOI: 10.1016/j.jmr.2009.09.017.

29. Balacco, G.; Cobas, C. Automatic Phase Correction of 2D NMR Spectra by a Whitening Method. Magn. Reson. Chem. 2009, 47 (4), 322–327. DOI: 10.1002/mrc.2394.

30. Hardy, E. H.; Hoferer, J.; Mertens, D.; Kasper, G. Automated Phase Correction via Maximization of the Real Signal. Magn. Reson. Imaging 2009, 27 (3), 393–400. DOI: 10.1016/j.mri.2008.07.009.

31. Bao, Q.; Feng, J.; Chen, L.; Chen, F.; Liu, Z.; Jiang, B.; Liu, C. A Robust Automatic Phase Correction Method for Signal Dense Spectra. J. Magn. Reson. 2013, 234, 82–89. DOI: 10.1016/j.jmr.2013.06.012.

32. Binczyk, F.; Tarnawski, R.; Polanska, J. Strategies for Optimizing the Phase Correction Algorithms in Nuclear Magnetic Resonance Spectroscopy. BioMed. Eng. OnLine 2015, 14 (Suppl. 2), S5. DOI: 10.1186/1475-925X-14-S2-S5.

33. Zorin, V.; Bernstein, M. A.; Cobas, C. A Robust, General Automatic Phase Correction Algorithm for High-Resolution NMR Data. Magn. Reson. Chem. 2017, 55 (8), 738–746. DOI: 10.1002/mrc.4586.

34. Zieliński, T. P.; Duda, K.; Ostrowska, K. Fast MinMax Energy-Based Phase Correction Method for NMR Spectra with Linear Phase Distortion. J. Magn. Reson. 2017, 281, 104–117. DOI: 10.1016/j.jmr.2017.05.014.

35. Cobas, C. Applications of the Whittaker Smoother in NMR Spectroscopy. Magn. Reson. Chem. 2018, 56 (12), 1140–1148. DOI: 10.1002/mrc.4747.

36. Sawall, M.; von Harbou, E.; Moog, A.; Behrens, R.; Schröder, H.; Simoneau, J.; Steimers, E.; Neymeyr, K. Multi-Objective Optimization for an Automated and Simultaneous Phase and Baseline Correction of NMR Spectral Data. J. Magn. Reson. 2018, 289, 132–141. DOI: 10.1016/j.jmr.2018.02.012.

37. Steimers, E.; Sawall, M.; Behrens, R.; Meinhardt, D.; Simoneau, J.; Münnemann, K.; Neymeyr, K.; von Harbou, E. Application of a New Method for Simultaneous Phase and Baseline Correction of NMR Signals (SINC). Magn. Reson. Chem. 2020, 58 (3), 260–270. DOI: 10.1002/mrc.4964.

38. Gazziro, M.; Amorim, M. L. M.; Cavallari, M. R.; Carmo, J. P.; Tannus, A.; Ando Junior, O. H.; Schwiebert, L. Automatic Phase Correction of NMR Spectra Using Brute-Force GPU Method. Inventions 2025, 10 (2), 21. DOI: 10.3390/inventions10020021.

39. Jaroszewicz, M. J.; Altenhof, A. R.; Schurko, R. W.; Frydman, L. An Automated Multi-Order Phase Correction Routine for Processing Ultra-Wideline NMR Spectra. J. Magn. Reson. 2023, 354, 107528. DOI: 10.1016/j.jmr.2023.107528.

40. Zhao, C.; Chen, G.; Liu, C.; Li, Z.; Zhao, J.; Sun, P.; Chu, H.; Zhou, X.; Liu, M.; Song, P.; Bao, Q.; Liu, C. Phase Model-Driven Deep Learning for Robust Phase Correction in High-Throughput NMR-Based Metabolomics. J. Phys. Chem. Lett. 2026, 17 (4), 1245–1253. DOI: 10.1021/acs.jpclett.5c03529.

41. Li, D.-W.; Bruschweiler-Li, L.; Lee, K.; Brüschweiler, R. DEEP Phaser: A Deep Learning Tandem Vision Transformer for Fully Automated NMR Phase Correction. J. Phys. Chem. Lett. 2026, 17 (19), 5520–5526. DOI: 10.1021/acs.jpclett.6c00770.

42. Morris, G. A. Compensation of Instrumental Imperfections by Deconvolution Using an Internal Reference Signal. J. Magn. Reson. 1988, 80 (3), 547–552. DOI: 10.1016/0022-2364(88)90253-3.

43. Delaglio, F.; Grzesiek, S.; Vuister, G. W.; Zhu, G.; Pfeifer, J.; Bax, A. NMRPipe: A Multidimensional Spectral Processing System Based on UNIX Pipes. J. Biomol. NMR 1995, 6 (3), 277–293. DOI: 10.1007/BF00197809.

44. Helmus, J. J.; Jaroniec, C. P. Nmrglue: An Open Source Python Package for the Analysis of Multidimensional NMR Data. J. Biomol. NMR 2013, 55 (4), 355–367. DOI: 10.1007/s10858-013-9718-x.

45. Lewis, I. A.; Schommer, S. C.; Markley, J. L. rNMR: Open Source Software for Identifying and Quantifying Metabolites in NMR Spectra. Magn. Reson. Chem. 2009, 47 (S1), S123–S126. DOI: 10.1002/mrc.2526.

46. Ludwig, C.; Günther, U. L. MetaboLab — Advanced NMR Data Processing and Analysis for Metabolomics. BMC Bioinf. 2011, 12 (1), 366. DOI: 10.1186/1471-2105-12-366.

47. Zolnai, Z.; Macura, S.; Markley, J. L. Spline Method for Correcting Baseplane Distortions in Two-Dimensional NMR Spectra. J. Magn. Reson. 1989, 82 (3), 496–504. DOI: 10.1016/0022-2364(89)90213-8.

48. Dietrich, W.; Rüdel, C. H.; Neumann, M. Fast and Precise Automatic Baseline Correction of One- and Two-Dimensional NMR Spectra. J. Magn. Reson. 1991, 91 (1), 1–11. DOI: 10.1016/0022-2364(91)90402-F.

49. Güntert, P.; Wüthrich, K. FLATT — A New Procedure for High-Quality Baseline Correction of Multidimensional NMR Spectra. J. Magn. Reson. 1992, 96 (2), 403–407. DOI: 10.1016/0022-2364(92)90095-O.

50. Golotvin, S.; Williams, A. Improved Baseline Recognition and Modeling of FT NMR Spectra. J. Magn. Reson. 2000, 146 (1), 122–125. DOI: 10.1006/jmre.2000.2121.

51. Eilers, P. H. C. A Perfect Smoother. Anal. Chem. 2003, 75 (14), 3631–3636. DOI: 10.1021/ac034173t.

52. Boelens, H. F. M.; Eilers, P. H. C.; Hankemeier, T. Sign Constraints Improve the Detection of Differences between Complex Spectral Data Sets: LC–IR As an Example. Anal. Chem. 2005, 77 (24), 7998–8007. DOI: 10.1021/ac051370e.

53. Cobas, J. C.; Bernstein, M. A.; Martín-Pastor, M.; Tahoces, P. G. A New General-Purpose Fully Automatic Baseline-Correction Procedure for 1D and 2D NMR Data. J. Magn. Reson. 2006, 183 (1), 145–151. DOI: 10.1016/j.jmr.2006.07.013.

54. Chang, D.; Banack, C. D.; Shah, S. L. Robust Baseline Correction Algorithm for Signal Dense NMR Spectra. J. Magn. Reson. 2007, 187 (2), 288–292. DOI: 10.1016/j.jmr.2007.05.008.

55. Xi, Y.; Rocke, D. M. Baseline Correction for NMR Spectroscopic Metabolomics Data Analysis. BMC Bioinf. 2008, 9 (1), 324. DOI: 10.1186/1471-2105-9-324.

56. Zhang, Z.-M.; Chen, S.; Liang, Y.-Z. Baseline Correction Using Adaptive Iteratively Reweighted Penalized Least Squares. Analyst 2010, 135 (5), 1138–1146. DOI: 10.1039/b922045c.

57. Fiorucci, L.; Bruno, F.; Ricci, M.; Cicone, A.; Ravera, E. Applications of Fast Iterative Filtering in NMR Spectroscopy: Baseline Correction. Magn. Reson. Chem. 2025, 63 (9), 737–747. DOI: 10.1002/mrc.70004.

58. McLeod, K.; Comisarow, M. B. Systematic Errors in the Discrete Integration of FT NMR Spectra. J. Magn. Reson. 1989, 84 (3), 490–500. DOI: 10.1016/0022-2364(89)90115-7.

59. Malz, F.; Jancke, H. Validation of Quantitative NMR. J. Pharm. Biomed. Anal. 2005, 38 (5), 813–823. DOI: 10.1016/j.jpba.2005.01.043.

60. Bharti, S. K.; Roy, R. Quantitative ^{1}H NMR Spectroscopy. TrAC, Trends Anal. Chem. 2012, 35, 5–26. DOI: 10.1016/j.trac.2012.02.007.

61. Le Gresley, A.; Fardus, F.; Warren, J. Bias and Uncertainty in Non-Ideal qNMR Analysis. Crit. Rev. Anal. Chem. 2015, 45 (4), 300–310. DOI: 10.1080/10408347.2014.944971.

62. Giraudeau, P. Challenges and Perspectives in Quantitative NMR. Magn. Reson. Chem. 2017, 55 (1), 61–69. DOI: 10.1002/mrc.4475.

63. Emwas, A.-H.; Saccenti, E.; Gao, X.; McKay, R. T.; dos Santos, V. A. P. M.; Roy, R.; Wishart, D. S. Recommended Strategies for Spectral Processing and Post-Processing of 1D ^{1}H-NMR Data of Biofluids with a Particular Focus on Urine. Metabolomics 2018, 14 (3), 31. DOI: 10.1007/s11306-018-1321-4.

14. Summary

The phase error carried by a Fourier-transform NMR spectrum has a physical origin in the delay between excitation and the first sampled point, and that origin supplies the one constraint on the problem that is not a matter of taste: the first-order phase and the first-point delay are related by $\varphi_1 = 360° \cdot SW \cdot t_0$, and t_0 cannot be negative. Applied to the spectrum examined here, that constraint disqualifies three of the five published objective functions tested, each of which returned a negative implied delay, and it disqualifies them on grounds that do not depend on how the resulting spectrum looks.

The reported anomaly over the aromatic region proved on examination to have two separable causes, only one of which is phase. The dominant cause is a dispersive pedestal, thirty to fifty times the local baseline, contributed by the intense aliphatic resonances and responding sharply to small rotations; the aromatic band itself retains a large asymmetry that does not vary with rotation and is therefore structural rather than dispersive. A genuine frequency-dependent phase error is nonetheless present, and measuring it independently in the aliphatic and aromatic bands with a locally refitted baseline gives $\varphi_0 = -6.27°$ and $\varphi_1 = +112.80 \pm 14.7°$, corresponding to an effective first-point delay of 78.3 ± 10 µs against a dwell period of 250 µs. The internal consistency of that figure with the hardware is the principal evidence that the determination is sound, since the global objective functions are demonstrably almost blind to φ_1 on a spectrum in which nearly all the intensity lies within a 3 ppm interval.

Applying the first-order term necessarily produced the baseline curvature that was anticipated at the outset, and its magnitude is substantial: the baseline model must remove almost sixteen intensity units at −1 ppm against an aromatic band thirty units tall. Removing it with a monotone multipoint interpolant through explicitly nominated anchors proved reliable, whereas polynomial models were unreliable in a way that their own residual statistics conceal, and penalized smoothers behaved poorly against the three-order-of-magnitude dynamic range of this spectrum. Automatic noise-threshold classification of baseline points was found to be inoperative at a signal-to-noise ratio of 1.3×10^5, which is not a defect of the classifiers but a consequence of asking whether a region is at the noise level when the limiting baseline is systematic structure far above it.

The result that matters analytically is the hierarchy of sensitivities. The complete phase correction changed the derived aromatic hydrogen index by 10.6 % relative; the uncertainty in the first-order term contributed under 2 % relative; the choice of interpolant contributed under 1 %; and the baseline anchoring convention spanned 45 % of the mean, from 2.018 % to 3.175 % on the aromatic hydrogen index. Attention devoted to selecting among phasing algorithms is therefore misdirected relative to attention devoted to fixing, documenting and uniformly applying the baseline and integration conventions. For quantitative work at low field on complex mixtures, the convention is the method, and it should be treated with the formality that implies.

Repeating the determination against the criterion an analyst uses by hand — flatten the flanks either side of the aromatic band, restore the aliphatic envelope, alternate until neither moves — sharpened rather than overturned that picture. The observation that prompted it was accurate: the upper flank does require about five degrees more rotation than the lower one, +5.28 against +2.10. But the two flanks lie only 191 Hz apart, so that difference extrapolates to +67 degrees of first-order phase rather than to a larger value, and driving the loop against the total flank slope instead converges on +288 degrees only because the genuine phase term and the ramp term cancel there. Nulling the genuine term alone gives +108.30 degrees, agreeing to within four degrees with the wholly independent band-based figure of +112.80. The flat flanks that motivated the exercise are obtained at either value once the multipoint baseline carries thirty or more nodes; obtaining them from the first-order term instead costs 15.7 % of the aromatic integral and opens a negative lobe on the band. The bound on the first-order term has been widened and re-expressed in dwell periods, but on this spectrum the extra range changes nothing, the next branch of the flatness criterion lying some 3776 degrees away.

The accompanying software implements the procedure described, reads and writes GRAMS SPC files natively so that corrected spectra return unchanged into an existing chemometric workflow, records the applied parameters in the

file log, and reports the diagnostics — Hilbert consistency, implied first-point delay, baseline-point count, region integrals — by which a correction should be judged rather than merely admired. The study data, comprising the full corrected spectrum, the stage-by-stage comparison, the engine comparison and the anchoring comparison, are supplied as a workbook alongside this document.