JOURNAL OF MECHANICS OF CONTINUA AND MATHEMATICAL SCIENCES
www.journalimcms.org

ISSN (Online) : 2454 -7190 Vol.-21, No.-08, August (2026) pp 238-270 ISSN (Print) 0973-8975

A CHEBYSHEV SPECTRAL COLLOCATION FRAMEWORK FOR GUIDED WAVES IN FLUID-LOADED CORTICAL BONE: EFFECTS OF ANISOTROPY AND POROSITY

M. S. L. R. Mallika[1], G. Sudheer[2], N. Aparna[3]

[1]Dept. of Mathematics, Anil Neerukonda Institute of Technology & Sciences, Visakhapatnam, AP, India.

[2] Dept. of Mathematics, Gayatri Vidya Parishad College of Engineering for Women, Visakhapatnam, AP, India.

[3] Earth & Climate Sciences Area, National Remote Sensing Centre, Hyderabad, Telangana, India.

Email: [1]radhamallika@yahoo.co.in, [2]sudhwave@gmail.com, [3]ddmsa@nrsc.gov.in

Corresponding Author: G. Sudheer

https://doi.org/10.26782/jmcms.2026.08.00014

(Received: June 10, 2026; Revised: August 02, 2026; Accepted: August 14, 2026)

Abstract

Bone quantitative ultrasound (QUS) relies on the dispersion of guided waves propagating in the cortical shell of long bones. Analytical fluid–solid–fluid trilayer models obtain the dispersion relation from a characteristic determinant whose roots must be located numerically, a procedure that becomes delicate near mode osculations and for anisotropic cortical bone. Here the fluid–solid–fluid waveguide is reformulated as a linear generalized eigenvalue problem using a Chebyshev spectral-collocation method. Each layer is discretized on a Chebyshev–Gauss–Lobatto grid, and the interface and free-surface conditions are imposed directly by row replacement. For a prescribed wavenumber, the eigenvalue is $W = \omega^2$, so every finite discrete eigenmode follows from a single generalized eigensolve without determinant root searching; the physical guided branches are then selected by explicit admissibility and participation criteria. The method is validated against an independently implemented analytical global-matrix solver for a water/aluminum/water trilayer: the fundamental extensional (S_0) and flexural (A_0) phase velocities agree within numerical tolerance (better than 10^{-5} %), and a convergence study with eigen-residual norms confirms spectral accuracy at about fourteen collocation nodes per layer. The same framework is then extended to a transversely isotropic cortical-bone layer and, with one additional volume-fraction field, to a linear elastic material with voids. Transverse isotropy raises the low-frequency extensional plateau by about 8.2 %, whereas increasing void coupling lowers it by up to 15 %. A Hellmann–Feynman sensitivity analysis shows that this plateau constrains only one stiffness combination, so these figures are forward-model discrepancies rather than uniquely recoverable inverse biases.

M. S. L. R. Mallika et al.

Keywords: *guided waves; cortical bone; quantitative ultrasound; spectral collocation; transverse isotropy; elastic materials with voids*

I. Introduction

Quantitative ultrasound is an attractive modality for assessing cortical-bone status because it is non-ionizing, portable, and sensitive to the elastic and geometric determinants of bone strength. In the axial-transmission configuration, a linear array placed along a long bone excites *guided* waves that travel in the cortical shell, and the measured dispersion of these modes is inverted to recover cortical thickness, porosity, and stiffness [IV, V, XIII, XIV]. Because the driving wavelength at typical QUS frequencies is short relative to the radius of curvature of the diaphysis, and because the receiver array spans a region over which the cortex is locally near-uniform, the curved cortical shell is well approximated *locally* by a flat plate; this plate idealization underlies all model-based axial-transmission inversions [IV, XIV]. Reviews of the field document a clear trend from first-arrival-signal velocity, through bare free-plate (Lamb) models, toward layered waveguide models that explicitly include the surrounding soft tissue and the medullary cavity [XII, XIII].

The theoretical backbone of these inversions is the dispersion relation of the waveguide. For a cortical plate coupled with fluid layers, a widely used route is the *partial-wave / global-matrix* method: the displacement and stress fields in each layer are written as sums of upgoing and downgoing partial waves [I, XVIII], the interface and free-surface conditions are assembled into a square matrix $\mathbf{D}(c, \omega)$, and the dispersion relation is the characteristic equation $\det\mathbf{D} = 0$. Nguyen *et al.* [XV] developed exactly such a model for a bone-mimicking plate between two soft-tissue/marrow fluid layers, treating the inviscid fluids as solids with zero shear modulus, deriving a 8×8 characteristic matrix, and additionally computing the amplitudes of the modes excited by a time-harmonic load through elastodynamic reciprocity. Their model is analytically transparent and reduces correctly to known fluid–solid bilayer results; notably, their reported validation uses an *isotropic* aluminum plate in water, not a real anisotropic cortical bone.

Determinant root searching is effective but can be delicate in practice: $\det\mathbf{D} = 0$ is a transcendental equation whose roots must be located by a search at every frequency, and near *osculations*, where two branches approach and repel, the search can skip or double-count modes unless safeguarded. Robust variants exist — contour integration, the argument principle, and SAFE-style mode tracking all mitigate these issues — but each adds bookkeeping, and generalizing the partial-wave construction to an anisotropic cortical layer requires solving a sextic Christoffel problem for the through-thickness slownesses and re-deriving the matrix entries.

Spectral and semi-analytical *discretization* methods avoid root searching by turning the through-thickness equations of motion into a matrix eigenvalue problem whose eigenvalues are the squared frequencies (or wavenumbers). The semi-analytical finite element (SAFE) method is the de facto standard for layered bone waveguides and was used by Nguyen *et al.* [XVI] to compute viscoelastic, transversely isotropic cortical-bone dispersion. Spectral-collocation methods (SCM) pursue the same idea with global Chebyshev polynomials rather than low-order finite elements, achieving exponential convergence for smooth layers: Adamou and Craster [II] introduced the approach for elastic guided waves, Hernando Quintanilla *et al.* [X] extended it to generally anisotropic and viscoelastic multilayers, and recent open-source tools build on these ideas [XI]. The collocation idea has also been used in geophysics to solve the depth-dependent surface-wave eigenproblem [VII].

The purpose of the present work is to bring this machinery to the bone trilayer of Nguyen *et al.* [XV] and to use it as the basis for an anisotropic extension. The principal contributions are:

1. A linear generalized-eigenvalue formulation of the fluid–solid–fluid bone waveguide, replacing determinant root searching. At fixed wavenumber k, the eigenvalue is $W = \omega^2$; because none of the imposed boundary/interface constraints contains W, the matrix pencil is *linear in* W (though k-dependent) and all modes follow from one eigensolve.
2. Phase and group velocities from the eigenpairs using a Hellmann–Feynman relation; the required derivative $\mathrm{d}\mathbf{A}/\mathrm{d}k$ is evaluated by a numerical central difference, so the group velocity is *Hellmann–Feynman-based with a numerical k-derivative* rather than fully closed-form.
3. A robust branch-tracking strategy in which physical modal characteristics (velocity band, parity, solid-displacement fraction) identify the initial branch seed, followed by modal-assurance-criterion (MAC) eigenvector continuation with adaptive wavenumber refinement near rapid modal reorganization, supported by through-thickness mode shapes.
4. A transversely isotropic cortical-bone extension obtained by replacing the two Lamé constants of the solid layer with four sagittal-plane stiffnesses, using a mesoscale stiffness tensor from the recent axial-transmission literature [IX], cited at the point of use.
5. Quantification of the error introduced by the isotropic approximations commonly used in QUS inversion.
6. A constitutive treatment of cortical porosity as a *linear elastic material with voids* [VI, XVII], in which the bone layer carries an independent volume-fraction field coupled to the dilatation, following the bone application of Andreaus *et al.* [III]. The extension adds one field per bone node without disturbing the linear-in-W structure, and quantifies the resulting porosity-induced softening of the fundamental modes.
7. A well-posedness and identifiability analysis of the resulting formulation**.** We establish the thermodynamic admissibility, strong ellipticity, and spectral reality of the augmented bone operator (Section 2.8); prove that the row-replacement enforcement yields a regular descriptor pencil algebraically equivalent to explicit constraint elimination, with quantified conditioning (Appendix B); and derive the parameter sensitivities from the same eigenpairs, using them to establish which constitutive combinations the dispersion actually identifies (Section 4.6).

We validate contributions 1–3 against an independent analytical global-matrix solver for the water/aluminum/water case (Section 3), then present contributions 4–5 for a soft-tissue / TI-cortical-bone / marrow trilayer (Section 4), and develop contribution 6 in Sections 2.7 and 4.5.

II. Spectral-collocation formulation

II.i. Geometry and field representation

We use the geometry of Nguyen *et al.* [XV]: a solid cortical layer Ω of thickness h occupies $-h \leq z \leq 0$, an upper fluid layer (soft tissue) $\hat{\Omega}$ of thickness $\hat{h}$ occupies $0 \leq z \leq \hat{h}$, and a lower fluid layer (marrow) $\tilde{\Omega}$ of thickness $\tilde{h}$ occupies $-(h+\tilde{h}) \leq z \leq -h$. Waves propagate along x; the displacement field of every layer is

$$u_j(x,z,t) = U_j(z)\, e^{ik(x-ct)} = U_j(z)\, e^{i(kx-\omega t)}, \qquad j \in \{x, z\},$$

with wavenumber $k = \omega/c$ and phase velocity c. The unknowns are the through-thickness profiles $U_x(z)$, $U_z(z)$. We restrict attention to **sagittal-plane (in-plane) motion**; for the materials considered (isotropic, or transversely isotropic with principal axes aligned to x and z), the anti-plane (SH) displacement u_y decouples from (u_x, u_z) at the level of both the equations of motion and the boundary conditions, so a sagittal-plane source excites only the Lamb-type S/A families treated here; the SH family would require a separate (scalar) problem.

2.2 Through-thickness operator for an elastic layer

For a layer whose sagittal-plane elasticity is described by the stiffnesses $C_{11}, C_{13}, C_{33}, C_{55}$ (Voigt notation, principal axes aligned with x and z) and density ρ, the in-plane stresses are

$$\sigma_{xx} = C_{11}\, \partial_x u_x + C_{13}\, \partial_z u_z, \qquad \sigma_{zz} = C_{13}\, \partial_x u_x + C_{33}\, \partial_z u_z, \qquad \sigma_{xz} = C_{55}\,(\partial_z u_x + \partial_x u_z).$$

Substituting the modal ansatz ($\partial_x \to ik$, $\partial_z \to \mathrm{d}/\mathrm{d}z \equiv D$) into $\rho\, \partial_t^2 u_j = \partial_x \sigma_{xj} + \partial_z \sigma_{zj}$ and writing $W \equiv \omega^2$ gives $\mathcal{L}\,\mathbf{U} = -\rho\, W\, \mathbf{U}$, i.e.

$$\underbrace{\begin{bmatrix} C_{55}\, D^2 - C_{11} k^2 & ik(C_{13}+C_{55})\, D \\ ik(C_{13}+C_{55})\, D & C_{33}\, D^2 - C_{55} k^2 \end{bmatrix}}_{\mathcal{L}} \begin{bmatrix} U_x \\ U_z \end{bmatrix} = -\rho\, W \begin{bmatrix} U_x \\ U_z \end{bmatrix}. \tag{1}$$

In the assembled problem (Section 2.4), the layer block is written $\mathbf{A}_{\text{layer}} = -\mathcal{L}$ and $\mathbf{B}_{\text{layer}} = \rho\,\mathbf{I}$, so that the global pencil takes the standard form $\mathbf{A}\,\mathbf{x} = W\,\mathbf{B}\,\mathbf{x}$; we state this sign convention explicitly to avoid ambiguity. Before boundary and interface enforcement, each dynamical layer possesses a positive-definite mass block. After the row replacement of Section 2.4, the global mass matrix becomes singular because the rows associated with the algebraic constraints are zeroed, so the assembled system is a *descriptor* generalized eigenvalue problem rather than a definite one (Appendix B). The isotropic special case $C_{11} = C_{33} = \lambda + 2\mu$, $C_{13} = \lambda$, $C_{55} = \mu$ reduces (1) to the standard sagittal operator, and the same form covers transverse isotropy by supplying four independent stiffnesses.

Units: Throughout the implementation, lengths are in micrometres (μm), stiffnesses in GPa, and densities in g/cm³; the wavenumber k is therefore in $\mu\mathrm{m}^{-1}$. With this system $f\,[\mathrm{MHz}] = (10^3/2\pi)\sqrt{W}$, $c\,[\mathrm{m/s}] = 10^3\sqrt{W}/k$, and the group velocity

$c_g\,[\mathrm{m/s}] = 10^3\,(\mathrm{d}W/\mathrm{d}k)/(2\sqrt{W})$. The factor 10^3 in these relations is the μm→mm unit conversion that maps the nondimensional eigenvalue to SI velocity; it likewise appears in the continuation predictor of Section 2.6.

The through-thickness derivative D is the Chebyshev spectral-collocation differentiation matrix [XIX, XX] on the $N+1$ Chebyshev–Gauss–Lobatto nodes $\xi_i = \cos(i\pi/N) \in [-1,1]$,

$$D_{ij}^{\text{cheb}} = \begin{cases} \dfrac{c_i}{c_j}\dfrac{(-1)^{i+j}}{\xi_i - \xi_j}, & i \neq j, \\ -\dfrac{\xi_i}{2(1-\xi_i^2)}, & 1 \leq i = j \leq N-1, \\ \pm\dfrac{2N^2+1}{6}, & i = j = 0\,(+),\ i = j = N\,(-), \end{cases} \qquad c_i = \begin{cases} 2, & i = 0, N \\ 1, & \text{else} \end{cases}.$$

Node ordering: The reference nodes run from $\xi_0 = +1$ to $\xi_N = -1$. Each physical layer is mapped by $z = 1/2\,(z_{\text{hi}} + z_{\text{lo}}) + 1/2\,(z_{\text{hi}} - z_{\text{lo}})\xi$ with $z_{\text{lo}} < z_{\text{hi}}$, so node 0 sits at the *upper* face z_{hi} and node N at the *lower* face z_{lo}, and the operator is scaled by $2/(z_{\text{hi}} - z_{\text{lo}})$. Because this orientation fixes the sign of D and the identity of the face nodes, it must be respected when forming the traction rows below; we follow it consistently across all three layers ($D^2 = D \cdot D$).

II.iii. Fluid layers as $\mu = 0$ solids and the treatment of spurious modes

Following [XV], the inviscid soft tissue and marrow are treated as degenerate solids with $\mu = 0$, so $C_{55} = 0$ and $C_{11} = C_{33} = C_{13} = \lambda = \rho c_L^2$; the layer then supports only the dilatational (pressure) wave, and the same discretized operator (1) applies. With $C_{55} = 0$, however, the fluid's tangential displacement U_x no longer appears in any stress and is dynamically unconstrained except through inertia, so the assembled pencil admits, in addition to the physical guided modes, (i) the dispersionless bulk-fluid branch at $c = c_g = c_L$ and (ii) a set of near-static / shear-free fluid modes with negligible solid participation. These are not numerical artifacts of an ill-posed pencil — the eigen-residual norms reported in Section 3 remain at the 10^{-10} level for all retained modes — but they are non-guided and must be excluded. We do so with three explicit rejection criteria:

1. **Eigenvalue admissibility / residual-norm check:** Retain only eigenpairs with a real, positive eigenvalue ($\Re W > 0$ and $|\Im W| < 10^{-4}\,|\Re W|$) that solve the discrete pencil to a small normalized residual $\| \mathbf{Ax} - W\mathbf{Bx} \| / \| \mathbf{Ax} \|$ at the round-off level (verified for every retained mode in Table 3). This discards the ill-conditioned, complex, or non-positive eigenvalues that a dense QZ factorization can return for a near-singular pencil.
2. **Fluid/solid displacement-norm fraction:** Reject the non-guided bulk and near-static fluid modes, which have negligible solid participation, by requiring a non-trivial solid displacement-norm fraction $\eta_{\text{solid}} = \| \mathbf{x}_{\text{solid}} \|^2 / \| \mathbf{x} \|^2$ (equivalently, a low fluid fraction), and by excluding the dispersionless bulk-fluid line at $c \approx c_g \approx c_L$ (Section 2.6). A physical velocity band $c_{\min} \leq c \leq c_{\max}$ and a frequency band $f \leq f_{\max}$ are applied alongside.

3. **Continuity (MAC) test:** In branch tracking, require a high modal-assurance-criterion overlap between the eigenvector at successive wavenumber steps (Section 2.6); this rejects spurious candidates and prevents branch jumps near osculations.

All filter thresholds are stated explicitly so that no physical mode is silently discarded. Unless noted otherwise, the computations below use $|\Im W| < 10^{-4}|\Re W|$, a solid displacement-norm fraction floor $\eta_{\text{solid}} \geq 0.2$, and velocity/frequency windows chosen per figure and quoted in the corresponding caption or text; the plateau extractions of Sections 4.3 and 4.5 additionally restrict to $2.4 \leq c \leq 6.5$ km/s and $f \leq 0.06$ MHz. The parity indicator P is used only as a diagnostic, and the MAC is used as a continuity diagnostic and adaptive-refinement trigger (Section 2.6), not as a direct mode-rejection filter; both are heuristic indicators of modal identity rather than proofs of it.

II.iv. Global eigenvalue problem and boundary conditions

Stacking the two nodal vectors of every layer into a global unknown $\mathbf{x}$ (ordered $[U_x(z_0), \dots, U_x(z_N), U_z(z_0), \dots, U_z(z_N)]$ per layer), the layer operators (1) assemble into block matrices $\mathbf{A} = \text{blkdiag}\left(-\mathcal{L}^{(\ell)}\right)$ and $\mathbf{B} = \text{blkdiag}\left(\rho^{(\ell)}\mathbf{I}\right)$, giving

$$\mathbf{A}(k)\,\mathbf{x} = W\,\mathbf{B}\,\mathbf{x}. \tag{2}$$

Selected rows of (2) — those associated with the outer faces and the interface nodes — are then *replaced* by the boundary and interface conditions of the analytical model [XV, Eq. (13)]:

- Free fluid surfaces (top of soft tissue $z = \hat{h}$; bottom of marrow $z = -\left(h + \tilde{h}\right)$): vanishing normal stress, $\sigma_{zz} = 0$. This is one scalar condition at each of the two outer faces.
- Each fluid–solid interface ($z = 0$ and $z = -h$): continuity of normal displacement u_z, continuity of normal stress σ_{zz}, and a *free tangential traction* $\sigma_{xz} = 0$ on the solid face. These are three scalar conditions at each of the two interfaces. The inviscid fluid cannot sustain shear, so the tangential displacement u_x is not required to be continuous (slip); the fluid's u_x degree of freedom is retained.

The total number of replaced scalar rows is therefore $2 + 3 + 3 = 8$ (two outer free-surface conditions plus three conditions at each of two interfaces), matching the eight scalar equations of the analytical 8×8 system. When the void field of Section 2.7 is active, two further rows are replaced by the free-pore conditions at the bone faces, giving 10 replaced rows in the augmented problem. In collocation form, the tractions on a face node i of layer ℓ are

$$\sigma_{zz}|_i = ik\,C_{13}^{(\ell)}\,U_x(z_i) + C_{33}^{(\ell)}\left(D^{(\ell)}\,U_z\right)_i, \qquad \sigma_{xz}|_i = C_{55}^{(\ell)}\left[\left(D^{(\ell)}\,U_x\right)_i + ik\,U_z(z_i)\right],$$

each a linear row in the global unknowns; the corresponding rows of $\mathbf{B}$ are zeroed so the constraint is algebraic. The resulting pencil therefore has a singular mass matrix, and its spectrum comprises finite eigenvalues together with infinite eigenvalues associated with the algebraic constraints; the QZ algorithm used below is appropriate precisely because it handles singular $\mathbf{B}$. Appendix B shows that, when the imposed constraints are independent, the pencil is a regular descriptor problem whose finite spectrum coincides exactly with that of explicit constraint elimination. Crucially, every constraint row is independent of $W=\omega^2$ (the traction rows depend on k through the explicit ik factors, but not on W), so (2) remains a *linear* generalized eigenproblem, solved directly with the QZ algorithm. This is the property that distinguishes the present approach from determinant-based formulations requiring a root search at each frequency.

2.5 Phase and group velocity

The phase velocity follows from each eigenvalue, $c=\sqrt{W}/k$. The group velocity is obtained by differentiating (2) with respect to k and applying the Hellmann–Feynman theorem. With $\mathbf{y},\mathbf{x}$ the left and right eigenvectors of a simple eigenvalue W,

$$\frac{\mathrm{d}W}{\mathrm{d}k}=\frac{\mathbf{y}^H\,(\mathrm{d}\mathbf{A}/\mathrm{d}k)\,\mathbf{x}}{\mathbf{y}^H\,\mathbf{B}\,\mathbf{x}},\qquad c_g=\frac{\mathrm{d}\omega}{\mathrm{d}k}=\frac{1}{2\sqrt{W}}\frac{\mathrm{d}W}{\mathrm{d}k}. \tag{3}$$

Here $\mathbf{B}$ is k-independent, so no $\mathrm{d}\mathbf{B}/\mathrm{d}k$ term appears; the normalization $\mathbf{y}^H\mathbf{Bx}$ is computed for each eigenpair returned by the generalized eigensolver. The only quantity that is differentiated is $\mathbf{A}(k)$, which we evaluate by a central difference $\mathrm{d}\mathbf{A}/\mathrm{d}k\approx[\mathbf{A}(k+\delta k)-\mathbf{A}(k-\delta k)]/2\delta k$ with $\delta k=10^{-4}k$; because $\mathbf{A}(k)$ is a low-degree polynomial in k, this difference is accurate to round-off. We therefore describe (3) as a Hellmann–Feynman group velocity with a numerical k-derivative of $\mathbf{A}$, rather than a fully closed-form expression. Group velocities near modal cutoffs and turning points should nevertheless be interpreted with caution: there the eigenvalue spacing becomes small, and the mode shape reorganizes rapidly, so the eigenvector normalization $\mathbf{y}^H\mathbf{Bx}$ in (3) is more sensitive to eigenvalue spacing and the resulting c_g is correspondingly less robust, even though $\mathbf{A}(k)$ itself remains smooth. For a repeated or near-degenerate eigenvalue, the simple-eigenvalue formula (3) is not used directly; such pairs are flagged by a small $|\mathbf{y}^H\mathbf{Bx}|$ and resolved by the MAC continuation of Section 2.6.

II.vi. Mode classification and branch tracking

Two scalar diagnostics are computed from each eigenvector. The solid displacement-norm fraction $\eta_{\mathrm{solid}}=\|\mathbf{x}_{\mathrm{solid}}\|^2/\|\mathbf{x}\|^2$ (a Euclidean displacement norm, not a physical energy) distinguishes guided plate modes from the bulk/near-static fluid modes of Section 2.3. The mid-plane parity of U_x, $P=\Re\left(\sum_i U_x(z_i)\,\overline{U_x(z_{N-i})}\right)/\sum_i |U_x(z_i)|^2$, serves as an *approximate* branch label for the symmetric (S_0, $P\approx+1$) and antisymmetric (A_0, $P\approx-1$) fundamentals; because the two fluid layers have unequal thickness and properties, the trilayer is not exactly mirror-symmetric about the solid mid-plane, so P is used only as a label, not as a strict symmetry classifier.

Each fundamental branch is followed by *continuation in* k: $\omega(k)$ is smooth even where $c=\omega/k$ falls steeply, so the next eigenvalue is predicted as $\sqrt{W_{\mathrm{pred}}}=\sqrt{W}+\left(c_g/10^3\right)\Delta k$ (the 10^3 being the same μm→mm conversion noted in Section 2.2, since c_g[m/s]$=10^3\,\mathrm{d}\sqrt{W}/\mathrm{d}k$) and matched to the nearest eigenvalue, with the bulk-fluid line excluded and c_g used as a tie-breaker.

MAC verification: To prevent branch jumps near osculations, velocity, displacement-participation and parity criteria are used **only to identify the initial branch seed**. Thereafter, the branch is continued by maximum eigenvector overlap (MAC), because the physical character of a mode may change substantially near a fluid-loading cusp; applying participation or bulk-wave filters at every continuation step can otherwise remove the genuine continuation and induce a silent branch switch. At each step, we therefore select the candidate maximizing the modal assurance criterion

$$\mathrm{MAC}(\mathbf{x}_m, \mathbf{x}_n) = \frac{|\mathbf{x}_m^H \mathbf{x}_n|^2}{(\mathbf{x}_m^H \mathbf{x}_m)(\mathbf{x}_n^H \mathbf{x}_n)},$$

and report the minimum and mean MAC along the branch. Where a consecutive-step MAC falls below 0.8, the wavenumber interval is bisected and the step repeated, so the criterion is enforced by adaptive refinement rather than by rejection. In practice, refinement is rarely required and always sufficient: over the void sweep of Section 4.5, a single bisection is triggered, at the strongest coupling, after which the minimum consecutive MAC is 0.899 ($\hat{g} = 0.9$), 0.956 (0.4), 0.970 (0.1) and 0.971 (0), so the 0.8 criterion is met throughout without being relaxed. Should refinement ever fail to reach the threshold, the value is reported and the mode shapes inspected directly rather than the step being rejected automatically. We emphasize that MAC branch continuity does not imply preservation of parity character: a branch may remain continuous in the eigenvector sense while its physical character changes substantially (Section 4.4). A MAC near unity between consecutive steps provides strong evidence of continuous modal evolution; the only departure occurs, as expected, at the sharp low-frequency S_0 cusp.

II.vii. Cortical porosity as a linear elastic material with voids

The operator (1) treats the cortex as a fully dense elastic solid. Cortical bone is, however, a porous medium — vascular porosity 5–10% in healthy cortex, substantially higher with age and disease — and porosity is among the principal targets of the QUS inverse problem. To represent it at the *constitutive* level, rather than only through an effective stiffness, we model the bone layer as a **linear elastic material with voids** in the sense of Goodman–Cowin [VIII] and Cowin–Nunziato [VI, XVII], following the bone-remodelling application of Andreaus, Giorgio and Madeo [III]. The bulk density is written as the product of a matrix density and a volume-fraction field, and the change of volume fraction $\xi(x,z,t)$ from the reference configuration is carried as an independent kinematic field alongside the displacement.

For an isotropic matrix, the stored-energy density adds three terms to the classical elastic energy [III, XVII],

$$\mathcal{E} = 1/2\,\lambda(\operatorname{tr} E)^2 + \mu\, E_{ij}E_{ij} + 1/2\, K_1\, \xi^2 + 1/2\, K_2\, \xi_{,i}\xi_{,i} + K_3\, \xi \operatorname{tr} E,$$

where K_1 is a pore (volume-fraction) stiffness, K_2 penalizes gradients of ξ over a microstructural length, and K_3 couples the void field to the dilatation; positive-definiteness of the augmented energy requires, in addition to $\mu > 0$ and $K_1, K_2 > 0$, the Cowin–Nunziato coupling restriction $K_3^2 < (\lambda + 2/3\,\mu)K_1$, equivalently $(3\lambda + 2\mu)K_1 > 3K_3^2$ [III, XVII]; this is derived, and its consequences for the parameter range used here are established in Section 2.8. The Cauchy stress gains a void contribution $\sigma_{ij} = \lambda(\mathrm{tr}\, E)\delta_{ij} + 2\mu E_{ij} + K_3\,\xi\,\delta_{ij}$, and the dynamic balance for the void field carries an equilibrated inertia $\rho\kappa$ [XVII]. With the modal ansatz ($\partial_x \to ik$, $\partial_z \to D$) the bone layer becomes a three-field operator on (U_x, U_z, ξ):

$$\begin{aligned} [\mu D^2 - (\lambda + 2\mu)k^2]U_x + ik(\lambda + \mu)D\,U_z + ik\,K_3\,\xi \quad &= -\rho\,W\,U_x, \\ ik(\lambda + \mu)D\,U_x + [(\lambda + 2\mu)D^2 - \mu k^2]U_z + K_3\,D\,\xi \quad &= -\rho\,W\,U_z, \\ K_2(D^2 - k^2)\xi - K_1\,\xi - K_3(ik\,U_x + D\,U_z) \quad &= -\rho\kappa\,W\,\xi. \end{aligned} \tag{4}$$

The void field adds one nodal unknown per bone node, so the bone block grows from $2(N_b + 1)$ to $3(N_b + 1)$ unknowns; the differentiation matrices act on ξ exactly as on the displacements, and — crucially — the assembled pencil (2) remains linear in W, since the new inertia $\rho\kappa$ enters only the mass matrix $\mathbf{B}$ and every constraint row stays W-free. At each traction-free bone face, the void field obeys a free-pore condition, vanishing equilibrated stress $h_z = K_2\,\partial_z\xi = 0$ [III]; this replaces the bone ξ-row at the two interface nodes, while the displacement interface conditions of Section 2.4 are unchanged except that σ_{zz} on the bone faces now includes the $K_3\,\xi$ term. The fluids carry no void field, and the bone ξ field is not continued into them.

Two limits bound the extension and serve as built-in checks. As $K_3 \to 0$ the void field decouples, and the displacement spectrum reduces to that of the elastic trilayer of Sections 2–3 (verified to four significant figures in Section 4.5). We claim reduction of the displacement spectrum only, not full three-field spectral identity: in the implementation, the void degrees of freedom are removed altogether at exactly $K_3 = 0$, so that case is an elastic specialization rather than a decoupled three-field solution. And the void field supports its own dilatational porosity branch with cutoff $\omega_c = \sqrt{K_1/(\rho\kappa)}$; for cortical parameters this cutoff is approximately 7.0 MHz (Section 2.8(d)), well above the 0–1.0 MHz range of guided-wave results reported here, so in-band the void field relaxes quasi-statically and its only effect is to replace λ by the equivalent Lamé parameter $\lambda_{\mathrm{eq}} = \lambda - K_3^2/K_1$ [III, Appendix 1].

The equilibrated inertia $\rho\kappa$ has no counterpart in the static remodelling model [III] and must be supplied for a dynamic problem; we estimate $\kappa \sim \ell^2$ from a vascular-pore length scale ℓ. The remaining coefficients follow [III]: $K_1 = K_1^{\max}(\varsigma^*)^\alpha$ with $K_1^{\max} \approx 11.8$ GPa scaled by the solid volume fraction, and $K_3 = \sqrt{\hat{g}\,\lambda K_1}$ with $\hat{g} \in [0,1]$ a dimensionless coupling that plays the role of the parameter A_{k_3} of [III]. We report results below as a sweep over $\hat{g}$ at fixed K_1, K_2, κ, isolating the mechanical effect of the void coupling.

II.viii. Admissibility, ellipticity and spectral reality of the augmented bone layer

Because the void extension (4) modifies the stiffness operator of the cortical layer, we verify explicitly that the augmented bone layer remains thermodynamically admissible, strongly elliptic, and spectrally real over the whole parameter range explored, rather than relying on the numerical solution to reveal instability *a posteriori*. Throughout this subsection, the statements concern the augmented bone operator; the fluid layers, being $\mu = 0$ solids (Section 2.3), are deliberately degenerate and are treated separately in remark (e).

(a) Positive definiteness of the augmented strain energy: Decompose the strain into deviatoric and spherical parts, $E_{ij} = E'_{ij} + 1/3\,\theta\delta_{ij}$ with $\theta = \operatorname{tr} E$, so that $E_{ij}E_{ij} = E'_{ij}E'_{ij} + \theta^2/3$. The stored energy of Section 2.7 then separates into three uncoupled quadratic forms,

$$\mathcal{E} = \mu\, E'_{ij}E'_{ij} \;+\; 1/2\, K_2\, \xi_{,i}\xi_{,i} \;+\; 1/2 \begin{bmatrix}\theta\\ \xi\end{bmatrix}^{\top} \begin{bmatrix}K & K_3\\ K_3 & K_1\end{bmatrix} \begin{bmatrix}\theta\\ \xi\end{bmatrix}, \qquad K \equiv \lambda + 2/3\,\mu,$$

with K the matrix bulk modulus. Hence $\mathcal{E}$ is positive definite if and only if

$$\mu > 0, \qquad K > 0, \qquad K_1 > 0, \qquad K_2 > 0, \qquad KK_1 - K_3^2 > 0, \tag{5}$$

the last of which is the Cowin–Nunziato coupling restriction

$$K_3^2 < (\lambda + 2/3\,\mu)K_1 \qquad \Leftrightarrow \qquad (3\lambda + 2\mu)K_1 > 3K_3^2. \tag{6}$$

Condition (6) is considerably stronger than the frequently quoted requirement $K_1, K_2 > 0$ alone, and it is the binding constraint on the void coupling. Expressed through the dimensionless coupling $\hat{g} = K_3^2/(\lambda K_1)$ of Section 2.7, (6) becomes a bound independent of K_1,

$$\hat{g} < \hat{g}^* = \frac{K}{\lambda} = 1 + \frac{2\mu}{3\lambda}.$$

For the cortical constants of Section 4.5 ($\lambda = 17.46$, $\mu = 8.22$ GPa) this gives $K = 22.94$ GPa and $\hat{g}^* = 22.94/17.46 \approx 1.314$, so the entire sweep $\hat{g} \in [0,0.9]$ used in Section 4.5 is thermodynamically admissible with margin. Equivalently, since $KK_1 - K_3^2 = K_1\left(\lambda_{\text{eq}} + 2/3\,\mu\right)$ with $\lambda_{\text{eq}} = \lambda(1 - \hat{g})$admissibility is precisely the statement that the in-band equivalent bulk modulus remains positive: the void coupling may soften the cortex but not to the point of negative compressibility. Table 1 lists the margin at each sampled coupling; at the strongest coupling explored, the equivalent bulk modulus is still 7.23 GPa, 31% of its dense value.

Table 1. Thermodynamic admissibility margin over the void-coupling sweep of Section 4.5 ($K_1 = 9$ GPa, $\lambda = 17.46$, $\mu = 8.22$ GPa, $K = 22.94$ GPa). Admissibility by (5)–(6) requires the third and fourth columns to be positive.

$\hat{g}$	λ_{eq} **(GPa)**	$\lambda_{\text{eq}} + 2/3\,\mu$ **(GPa)**	$KK_1 - K_3^2$ **(GPa²)**	$K_3^2/(KK_1)$
0.0	17.46	22.94	206.5	0.000
0.1	15.71	21.19	190.8	0.076
0.4	10.48	15.96	143.6	0.305
0.9	1.75	7.23	65.0	0.685

(b) Strong ellipticity of the augmented bone operator: The highest-order part of the three-field system (4) comprises the elastic term $C_{ijkl}u_{k,lj}$ and the void-gradient term $K_2\nabla^2\xi$; the K_3 coupling contributes only *first* spatial derivatives and therefore does not enter the principal symbol, which is block diagonal,

$$\mathbf{P}(\mathbf{n}) = \begin{bmatrix}\mu|\mathbf{n}|^2\mathbf{I} + (\lambda + \mu)\,\mathbf{n}\otimes\mathbf{n} & \mathbf{0}\\ \mathbf{0}^{\top} & K_2|\mathbf{n}|^2\end{bmatrix}. \tag{7}$$

Strong ellipticity of the augmented bone layer therefore reduces to the ordinary elastic conditions together with a positive gradient stiffness,

$$\mu > 0, \qquad \lambda + 2\mu > 0, \qquad K_2 > 0, \tag{8}$$

and is, notably, *independent of* K_3: the void coupling cannot destroy strong ellipticity, though by (6) it can destroy positive definiteness of the energy. Conditions (8) are therefore the complete strong-ellipticity requirement for the augmented bone layer. As a separate positivity and stability check — not a further ellipticity condition — we note that the quasi-static reduced elastic response in band is governed by λ_{eq} (Section 2.7), and verify that $\lambda_{\text{eq}} + 2\mu > 0$, satisfied throughout the sweep (18.19 GPa at $\hat{g} = 0.9$). We stress that (8) is a statement about the bone layer only; the complete three-layer displacement operator is deliberately degenerate because the fluids have $C_{55} = 0$ (remark (e)).

(c) Reality of the continuous spectrum: The bone operator in (4) is complex as written because of the explicit ik factors, but it is *unitarily equivalent to a real symmetric one*. Under the diagonal transformation $\mathbf{T} = \text{diag}(1, i, i)$ acting on (U_x, U_z, ξ) — that is, writing $U_z = iV$ and $\xi = i\eta$ — system (4) becomes, after dividing the second and third rows by i,

$$\begin{aligned} [\mu D^2 - (\lambda + 2\mu)k^2]U_x - k(\lambda + \mu)D\,V - kK_3\,\eta &= -\rho W U_x, \\ k(\lambda + \mu)D\,U_x + [(\lambda + 2\mu)D^2 - \mu k^2]V + K_3 D\,\eta &= -\rho W V, \\ K_2(D^2 - k^2)\eta - K_1\eta - K_3(kU_x + DV) &= -\rho\kappa W\eta, \end{aligned} \tag{9}$$

with every coefficient real. Under the natural inner product, $D = \mathrm{d}/\mathrm{d}z$ is skew-adjoint and D^2 self-adjoint on the free-pore and traction conditions of Sections 2.4 and 2.7, so the $(1,2)$ and $(2,1)$ blocks of (9) are mutually adjoint, as are the $(2,3)$ and $(3,2)$ blocks, while the $(1,3)/(3,1)$ coupling $-kK_3$ is symmetric. The continuous conservative bone problem is therefore self-adjoint with a positive-definite inertia $\text{diag}(\rho, \rho, \rho\kappa)$ for $\rho > 0$, $\kappa > 0$, so its spectrum is real, and under (5) the stiffness form is positive definite for $k \neq 0$, giving $W > 0$ and real propagating frequencies $\omega = \sqrt{W}$. No growing modes exist in the admissible parameter range.

It is important to state the limits of this conclusion. The collocation discretization *approximates* this spectrum, and the row-replaced discrete pencil of Section 2.4 is not symmetric (Appendix B); it therefore does not automatically inherit exact matrix self-adjointness, and the discrete eigenvalues are not guaranteed to be exactly real by the continuum argument alone. Accordingly, we do not claim exact reality of the computed spectrum: small imaginary parts serve as *numerical diagnostics*, are required to satisfy $|\Im W| < 10^{-4}|\Re W|$ by criterion 1 of Section 2.3, and vanish under refinement.

(d) The porosity branch: The branch introduced by the void field has cutoff $\omega_c = \sqrt{K_1/(\rho\kappa)}$, real for all admissible couplings by (5). For the parameters of Section 4.5, this is $f_c \approx 7$ MHz, about 13× the highest frequency shown in Figure 7 (0.55 MHz), so direct resonance with the void branch is not expected within the investigated QUS band; this is consistent with the absence of an observed avoided crossing in that figure.

(e) The fluid subspaces are degenerate by construction: No positive-definiteness claim is made for the complete fluid–solid stiffness operator. Because the fluids are represented as zero-shear elastic solids ($C_{55} = 0$, Section 2.3), the displacement formulation admits divergence-free fluid displacement fields that carry inertia but no elastic energy; the operator is consequently positive *semi*definite and degenerate on those subspaces. This is a known and intended consequence of representing an inviscid acoustic fluid in displacement form; it is the origin of the zero and near-zero shear-free fluid modes discussed in Section 2.3, and those modes are removed from the guided spectrum by the explicit admissibility and participation filters of Section 2.3 rather than by any property of the operator.

(f) An independent closed-form check: Because the void field relaxes quasi-statically in band, the low-frequency S_0 plateau of the voids model must coincide with the classical extensional plate velocity evaluated with λ_{eq},

$$c_{S_0} = \sqrt{\frac{C_{11}^{\mathrm{eq}} - (C_{13}^{\mathrm{eq}})^2 / C_{33}^{\mathrm{eq}}}{\rho}}, \qquad C_{11}^{\mathrm{eq}} = C_{33}^{\mathrm{eq}} = \lambda_{\mathrm{eq}} + 2\mu, \quad C_{13}^{\mathrm{eq}} = \lambda_{\mathrm{eq}}. \tag{10}$$

Equation (10) predicts 3649.6, 3617.7, 3494.7 and 3104.1 m/s at $\hat{g} = 0,\ 0.1,\ 0.4,\ 0.9$, against the computed 3645.6, 3614.2, 3492.4 and 3103.9 m/s of Section 4.5 — agreement to better than 0.11% across the sweep. This is an independent analytical check on the augmented assembly, complementary to the $K_3 \to 0$ reduction test.

III. Validation against an independent analytical solver

To make the validation quantitative rather than visual, we implemented the analytical *partial-wave global-matrix* solver from scratch and independently of the SCM code, in SI units. It assembles the 8×8 matrix $\mathbf{D}(c, f)$ from the partial-wave fields and the eight boundary/interface conditions of [XV] and locates modes as minima of $\log|\det\mathbf{D}|$ refined by golden-section search. As a correctness check on this reference itself, in the light-fluid limit (fluid density reduced by 10^3) its S_0/A_0 velocities reproduce the exact Rayleigh–Lamb free-plate solutions of an aluminum plate to better than 0.1 m/s (e.g. at 0.10 MHz, $S_0 = 5375.8$ and $A_0 = 1875.8$ m/s by both routes).

We then compare the SCM (collocation orders $N = (22{,}18{,}44)$) against this analytical solver for the water/aluminum/water trilayer with the properties of Nguyen *et al.* [XV] — water $\rho = 1000\ \mathrm{kg/m^3}$, $\lambda = 2.25$ GPa, $\mu = 0$; aluminum $\rho = 2700\ \mathrm{kg/m^3}$, $\lambda = 55.25$ GPa, $\mu = 25.94$ GPa ($c_L = 6299$, $c_T = 3100$ m/s); thicknesses $3/5/10$ mm. Table 2 lists the S_0 and A_0 phase velocities at six frequencies.

Table 2. Quantitative validation: S_0 and A_0 phase velocities (m/s) of the water/aluminum/water trilayer from the independent analytical global-matrix solver and from the SCM, with percentage error.

f (MHz)	S_0 analytical	S_0 SCM	error (%)	A_0 analytical	A_0 SCM	error (%)
0.05	3309.77	3309.77	<0.001	1250.59	1250.59	<0.001
0.10	1851.46	1851.46	<0.001	1515.17	1515.17	<0.001
0.15	1872.22	1872.22	<0.001	1521.76	1521.76	<0.001
0.20	1760.34	1760.34	<0.001	1511.04	1511.04	<0.001
0.30	1605.31	1605.31	<0.001	1501.61	1501.61	<0.001
0.40	1552.07	1552.07	<0.001	1498.57	1498.57	<0.001

In each solver, and at each frequency independently, the two fundamentals are identified without any shape heuristic. For real wavenumber, the underlying continuous conservative problem is self-adjoint (Section 2.8(c)), so its physical eigenvalues are real; the row-replaced discrete pencil is not symmetric (Appendix B.4), so reality of the computed spectrum is verified rather than assumed. After removal of the nonphysical and fluid-degenerate eigenvalues listed below, ordering the admissible eigenvalues at each k provides a robust numerical labelling of the two lowest branches over the parameter range considered; continuity of the resulting A_0 and S_0 branches was additionally verified by MAC continuation (Section 2.6) and from the computed dispersion curves and modal characteristics (Section 4.4). We do not rely on a general no-crossing theorem. The lowest two admissible branches are A_0 and S_0: as the wavenumber tends to zero, the first tends to zero phase velocity and the second to the plate (extensional) velocity, 5396.7 m/s for this aluminum. Before ordering, two families of numerical artifacts are removed: eigenvalues that are infinite, non-positive, or appreciably complex, and the degenerate cluster carried entirely by the zero-shear fluid layers, which is discarded by requiring a solid-participation fraction above 10^{-6}. For the analytical solver, the corresponding roots are the first and second roots of the 8×8 characteristic determinant above the numerical noise floor, ordered in phase velocity at fixed frequency. Neither solver's result is used to select the other's mode. The SCM entry is obtained by bisecting on the wavenumber so that the identified mode sits at the tabulated frequency exactly, rather than by interpolating a sampled branch.

At all six frequencies, the two independent methods agree to better than 10^{-5} %; the largest absolute difference is 5×10^{-5} m/s, which is itself at the level of the golden-section tolerance used to refine the analytical roots, so the tabulated agreement should be read as a bound rather than as a resolved discrepancy. This is a sharper level of agreement than a branch-interpolated comparison can show, because the SCM value is evaluated at the wavenumber that places the mode exactly at the tabulated frequency. The two independent methods therefore agree both visually and numerically. Figure 1 shows the full phase- and group-velocity spectra; the SCM curves overlie the analytical dispersion of [XV, Fig. 2]. Note that S_0 falls steeply through a cusp near 0.045 MHz, from the 5396.7 m/s plateau to below 1900 m/s, and thereafter decreases slowly toward the bulk-water speed; the tabulated values at and above 0.10 MHz lie on this post-cusp portion of the branch.

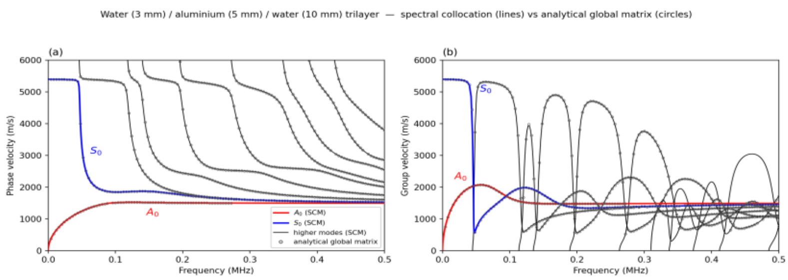

Fig. 1. Dispersion of the water/aluminum/water trilayer (3/5/10 mm) by the spectral-collocation method (lines) compared with the independent analytical global-matrix solution (open circles): (a) phase velocity, (b) group velocity. S_0 (blue), A_0 (red), higher modes (black).

The dense horizontal accumulation at 1500 m/s is the bulk-water branch (Section 2.3), excluded from the S_0/A_0 identification. In both panels, A_0 and S_0 are the first- and second-rank-ordered admissible eigenbranches at each wavenumber, so the plotted curves are mathematically continuous and match the values reported in Table 2 exactly. S_0 leaves the 5396.7 m/s extensional plateau through a sharp cusp near 0.045 MHz, where the mode reorganizes from a plate-extensional to a fluid-loaded character; the group-velocity panel shows the corresponding minimum. The open circles are the independent analytical global-matrix solution. Over the plotted range, the two solvers agree to better than 10^{-5} %, the same agreement quantified numerically in Table 2.

III.i. Convergence and eigen-residuals

Table 3 reports the first five mode frequencies of the water/aluminum/water trilayer at two fixed wavenumbers as the collocation orders are refined, together with the maximum normalized eigen-residual $\max_m \| \mathbf{A}\mathbf{x}_m - W_m\mathbf{B}\mathbf{x}_m \| / \| \mathbf{A}\mathbf{x}_m \|$ over those five modes and the error of the lowest frequency relative to the finest grid $N = (28,26,56)$.

Percentage errors in Table 2 are computed from unrounded velocities; the tabulated values are rounded for display, so recomputing the percentages from the printed figures may differ in the last digit.

Table 3. Convergence with eigen-residual norms. First five mode frequencies (MHz), maximum normalized residual, and lowest-mode error vs. the finest grid, at two wavenumbers.

$(N_{\text{soft}}, N_{\text{bone}}, N_{\text{marrow}})$	first five f (MHz)	max $\| r \|$	$\lvert f_1 - f_1^{\text{ref}} \rvert$
$k = 8 \times 10^{-4}\ \mu\text{m}^{-1}$			
(10,8,20)	0.19099, 0.19099, 0.19255, 0.21849, 0.22125	2.7×10^{-10}	5.5×10^{-12}
(14,12,28)	0.19099, 0.19099, 0.19255, 0.21849, 0.22125	3.3×10^{-10}	5.5×10^{-12}
(18,16,36)	0.19099, 0.19099, 0.19255, 0.21849, 0.22125	8.6×10^{-10}	5.3×10^{-12}
(22,18,44)	0.19099, 0.19099, 0.19255, 0.21849, 0.22125	7.0×10^{-10}	5.5×10^{-12}
(28,26,56)	0.19099, 0.19099, 0.19255, 0.21849, 0.22125	1.2×10^{-9}	reference
$k = 1.6 \times 10^{-3}\ \mu\text{m}^{-1}$			
(10,8,20)	0.38170, 0.38197, 0.38197, 0.39562, 0.39643	2.1×10^{-11}	4.4×10^{-6}
(14,12,28)	0.38170, 0.38197, 0.38197, 0.39561, 0.39642	7.5×10^{-11}	9.1×10^{-9}
(18,16,36)	0.38170, 0.38197, 0.38197, 0.39561, 0.39642	1.5×10^{-10}	3.7×10^{-12}
(22,18,44)	0.38170, 0.38197, 0.38197, 0.39561, 0.39642	1.2×10^{-10}	3.5×10^{-14}
(28,26,56)	0.38170, 0.38197, 0.38197, 0.39561, 0.39642	4.2×10^{-10}	reference

Two points follow. First, the normalized residuals stay at the 10^{-10}–10^{-9} level for *every* retained mode at *every* order, confirming that the retained eigenpairs accurately satisfy the discretized pencil. A small residual by itself does not establish that a mode is physical — a spurious discretization mode can also possess a tiny residual — so physicality is supported additionally by grid convergence, solid participation, MAC continuity, mode shapes, and the analytical comparison of Table 2 (cf. Section 2.3). Second, the lowest-mode error decays exponentially: at the higher wavenumber it falls from 4.4×10^{-6} at (10,8,20) to 3.5×10^{-14} at (22,18,44), the hallmark of spectral convergence. (At $k = 8 \times 10^{-4}$ the lowest mode is already converged to round-off even at (10,8,20).) The figures use $N = (22,18,44)$, a comfortably converged grid. The transversely isotropic trilayer of Section 4 behaves identically: at

$k = 8 \times 10^{-4}$ its lowest mode converges from $f_1 = 0.1535627$ MHz (error 2.2×10^{-7} at the coarsest grid) to 0.15356292 MHz, error $\sim 10^{-13}$.

To be precise about the convergence claim: the tabulated quantity $|f_1 - f_1^{\text{ref}}|$ is the absolute error of the lowest eigenfrequency (in MHz) against the fine-grid reference solution $N = (28,26,56)$, taken as converged since refining beyond it changes f_1 only at the 10^{-14} level. Dividing by f_1 gives the **relative** eigenfrequency error (e.g. $3.5 \times 10^{-14}/0.396 \approx 9 \times 10^{-14}$ at the finest tabulated grid). On this relative measure, the lowest eigenfrequency reaches the 10^{-10} level by $N \approx (14,12,28)$ at the lower wavenumber ($5.5 \times 10^{-12}/0.191 \approx 3 \times 10^{-11}$) and by $N \approx (18,16,36)$ at the higher wavenumber; the abstract claim refers to this lowest-mode relative error, not to every mode at every wavenumber.

IV. Transversely isotropic cortical bone

IV.i. Material model and axis mapping

Real cortical bone is well approximated at the tissue scale as transversely isotropic (hexagonal symmetry) with its stiff axis along the diaphysis, and several QUS studies argue that ignoring this anisotropy biases the recovered properties [IV, XIV, XVI]. We use a mesoscale TI stiffness tensor of the kind employed in modern axial-transmission inversions, derived from micromechanical homogenization of the vascular porosity [IX] and tabulated in the guided-wave bone literature. These constants are not arbitrary: they correspond to mesoscale cortical-bone properties commonly used in axial-transmission simulations and inversions [IV, IX, XIV], so that the dispersion predicted here is directly comparable to that literature. Taking representative periosteal values, the material constants are, **in the material frame whose axis 3 is the bone long axis**,

$$C_{33} = 33.90, \quad C_{11} = C_{22} = 26.10, \quad C_{12} = 10.64, \quad C_{13} = C_{23} = 11.16, \quad C_{44} = C_{55} = 8.22, \quad C_{66} = 7.73 \text{ (GPa)},$$

with density $\rho = 1.87\ \text{g/cm}^3$. The tensor is internally consistent — $C_{66} = (C_{11} - C_{12})/2 = (26.10 - 10.64)/2 = 7.73$ — and shows the expected cortical ordering $C_{33} > C_{11} > C_{44} > C_{66}$.

Axis mapping: In axial transmission, the wave propagates along the bone long axis, so the *material* axis 3 is aligned with the *spatial* propagation direction x, the spatial thickness direction z being one of the transverse material axes. The sagittal operator (1) requires the four spatial constants $\left(C_{11}^{\text{sp}}, C_{13}^{\text{sp}}, C_{33}^{\text{sp}}, C_{55}^{\text{sp}}\right)$, obtained from the material tensor by the index map (axis 3) $\to x$:

$$C_{11}^{\text{sp}} = C_{33}^{\text{mat}} = 33.90, \quad C_{33}^{\text{sp}} = C_{11}^{\text{mat}} = 26.10, \quad C_{13}^{\text{sp}} = C_{13}^{\text{mat}} = 11.16, \quad C_{55}^{\text{sp}} = C_{44}^{\text{mat}} = 8.22 \text{ (GPa)},$$

i.e. the plate is stiffest along the propagation direction ($C_{11}^{\text{sp}} > C_{33}^{\text{sp}}$). The associated bulk speeds are an axial longitudinal $c_L^{\text{ax}} = \sqrt{C_{11}^{\text{sp}}/\rho} = 4258$ m/s, an axial shear $c_T^{\text{ax}} = \sqrt{C_{55}^{\text{sp}}/\rho} = 2097$ m/s, and a thickness-direction longitudinal $\sqrt{C_{33}^{\text{sp}}/\rho} = 3736$ m/s. The surrounding layers are soft tissue ($\rho = 1000\ \text{kg/m}^3$, $c_L = 1540$ m/s) and marrow ($\rho = 980\ \mathrm{kg/m^3}$, $c_L = 1450$ m/s); thicknesses are $4/3/10$ mm (soft tissue / cortical bone/marrow), the 3-mm cortex being in the range reported for the radius and tibia.

For comparison, we define an **isotropic-from-axial** bone — the model an isotropic inversion would adopt if calibrated to the bone's *axial* speeds: $\mu = C_{55}^{\mathrm{sp}} = 8.22$ GPa and $\lambda + 2\mu = C_{11}^{\mathrm{sp}} = 33.90$ GPa, hence $\lambda = 17.46$ GPa, $\mu = 8.22$ GPa, same density $\rho = 1.87\ \mathrm{g/cm^3}$. By construction, it shares the bone's fast-axis speeds but raises the thickness stiffness C_{33} to the axial value.

IV.ii. Dispersion of the soft-tissue / TI-bone / marrow trilayer

Figure 2 shows the computed modal spectrum over 0–1.0 MHz, with each mode coloured by the fraction of its kinetic energy carried in the bone layer. The fundamental S_0 branch leaves its 3943 m/s extensional plateau through a cusp near 0.05 MHz and then decreases slowly toward the marrow speed, while the A_0 branch rises from zero and asymptotes toward the slower fluid speed; both are obtained by rank ordering the admissible eigenvalues, as in Section 3. The colouring makes the underlying physics explicit: bone-dominated energy is not confined to a single branch above the cusp but is handed successively to higher modes at each avoided crossing with the soft-tissue/marrow resonance spectrum, tracing a staircase that remains near 3900–4100 m/s. It is this staircase, and not a single eigenbranch, that a shape-based or velocity-window selection rule picks out; we therefore do not use such a rule. The flat, essentially bone-free lines at the soft-tissue and marrow bulk speeds (1540, 1450 m/s) are the zero-shear fluid cluster and are excluded from S_0/A_0 identification.

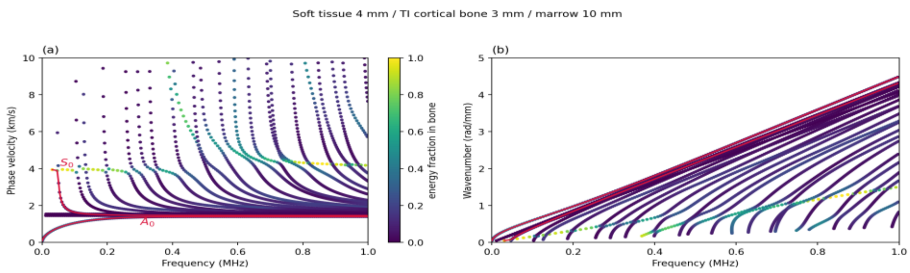

Fig. 2. Modal spectrum of the soft-tissue / TI-cortical-bone / marrow trilayer (4/3/10 mm), computed with the Section 4.1 constants: (a) phase velocity and (b) wavenumber against frequency.

Each point is one admissible eigenmode, coloured by the fraction of its energy carried in the bone layer. The red curves are the rank-ordered fundamentals S_0 and A_0. The bone-dominated (yellow-green) points above the S_0 cusp lie on a succession of different higher modes, not on a single branch.

IV.iii. Transverse isotropy versus isotropy

Figure 3 overlays the fundamental S_0 and A_0 branches of the TI bone against the isotropic-from-axial bone. The most consequential difference is the low-frequency S_0 extensional plateau — the quantity most directly tied to plate stiffness and a primary target of QUS inversion.

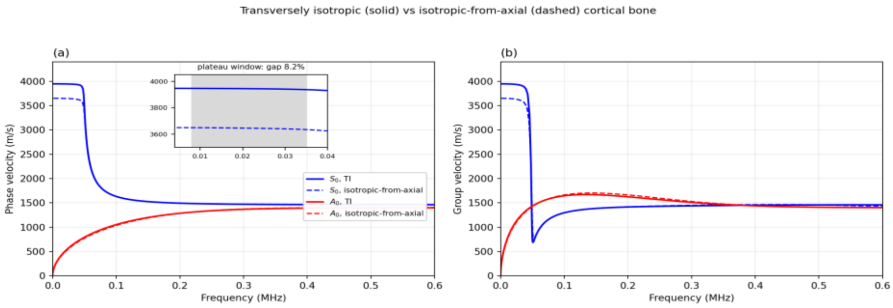

Fig. 3. Transversely isotropic (solid) versus isotropic-from-axial (dashed) cortical bone: (a) phase velocity and (b) group velocity of the S_0 (blue) and A_0 (red) fundamentals, both rank-ordered eigenbranches as in Figures 1–2.

The inset magnifies the sub-cusp plateau window (shaded, 0.008–0.035 MHz) over which the 8.2 % gap is quoted. Above the cusp near 0.05 MHz, the two S_0 branches collapse onto each other and onto the fluid speed, so the anisotropy signature is confined to the plateau region.

The averaging window 0.008–0.035 MHz is chosen because it lies *below* the S_0 group-velocity cusp (near 0.05 MHz): over this interval the S_0 phase velocity is essentially flat, and the mode is a genuine extensional plate wave (its through-thickness shape is the axial-dominated profile of Figure 4a), so the plateau value is a well-defined material descriptor rather than a point on a steep transition. Averaged over this window,

$$c_{S_0}^{\mathrm{TI}} = 3943\ \mathrm{m/s}, \qquad c_{S_0}^{\mathrm{iso}} = 3645\ \mathrm{m/s}, \qquad \frac{c_{S_0}^{\mathrm{TI}} - c_{S_0}^{\mathrm{iso}}}{c_{S_0}^{\mathrm{iso}}} = 8.2\%.$$

The TI value (3943 m/s) sits between the thickness-direction speed (3736 m/s) and the axial speed (4258 m/s), as expected for an in-plane extensional mode governed by the axial modulus but relaxed by the lower out-of-plane stiffness through C_{13}. The isotropic-from-axial model raises the thickness stiffness to the axial value and so *underestimates* the plate velocity. Because the S_0 velocity maps onto an effective plate stiffness through $c^2 \propto C$, the forward-model-equivalent stiffness discrepancy between the TI model and an isotropic surrogate calibrated to the fast-axis speeds is

$$\frac{\Delta C}{C} \approx \left(\frac{c_{S_0}^{\mathrm{TI}}}{c_{S_0}^{\mathrm{iso}}}\right)^2 - 1 = \left(\frac{3943}{3645}\right)^2 - 1 \approx 17\%.$$

A representative sampling of the two fundamentals (verified to be correctly labelled by the MAC continuation of Section 4.4) is given in Table 4.

Table 4: S_0 and A_0 phase velocities of the TI and isotropic-from-axial cortical bone.

f (MHz)	S_0 TI (m/s)	S_0 iso (m/s)	A_0 TI (m/s)	A_0 iso (m/s)
0.05	3460	3333	792	773
0.10	1631	1630	1051	1035
0.20	1492	1492	1286	1283
0.40	1464	1469	1391	1398

S_0 and A_0 in this table are the second- and first-rank-ordered admissible eigenbranches at each frequency (the same criterion as Table 2 and Figures 1–3), not shape-identified and not continuity-tracked from a seed; the values therefore match Figures 2 and 3 exactly. The TI/isotropic gap is strongly frequency dependent: it is 8.2 % on the sub-cusp plateau and 3.8 % at 0.05 MHz, but has fallen to 0.1 % by 0.10 MHz and is below 0.5 % thereafter. This is the expected behaviour and not a defect of the model: above the cusp, both S_0 branches are fluid-loaded, and their phase velocity is controlled by the surrounding fluids rather than by the cortical stiffness, so sensitivity to anisotropy is lost. The practical consequence is that the anisotropy information usable for inversion resides in the low-frequency plateau, consistent with the QUS literature recommending TI waveguide models for reliable cortical-property recovery [IV, XIV, XVI]. The approximately 17 % figure is a discrepancy between two forward models at the stated parameter sets, evaluated on that plateau; because the forward map is locally many-to-one, it should not be read as a uniquely recoverable inversion error. Section 4.6 quantifies which parameter combinations the plateau actually constrains.

IV.iv. Mode-shape and MAC evidence for the tracked branches

Two objects must not be conflated here. The first is the single mathematically continuous eigenbranch labelled S_0 in Table 2, Table 4 and Figures 1–3, obtained by rank ordering at each wavenumber; for the cases studied here, seeding at the extensional mode and continuing outward by eigenvector overlap (MAC) gives the same branch, as verified below — the two procedures are not guaranteed to coincide in general near avoided crossings, so this is an empirical check rather than a claim of equivalence. This is what Figure 4 displays. The second is the locus of bone-dominated, extensional-looking modes visible as the yellow-green staircase in Figure 2, which a shape-based or velocity-window selection rule would return. These coincide below the cusp and diverge above it. The S_0 eigenbranch's phase velocity collapses close to the fluid speed already near 0.10 MHz; this is expected and is not an error, since the branch is continuing correctly but its physical character is no longer extensional beyond the cusp. What must be verified is only that the continuation itself does not silently jump to an unrelated mode. We rule that out in two independent ways.

First, we re-tracked each fundamental purely by eigenvector overlap (MAC), independently of the velocity predictor. Over 219 continuation steps in k, the minimum and mean consecutive-step MAC are

$$S_0\text{: minMAC} = 0.927,\ \ \overline{\text{MAC}} = 0.999; \qquad A_0\text{: minMAC} = 0.9999,\ \ \overline{\text{MAC}} = 1.000.$$

The A_0 shape is essentially invariant; the single S_0 dip to 0.927 occurs exactly at the sharp low-frequency cusp, where the mode shape physically reorganizes from extensional to fluid-loaded — yet it remains the highest-overlap continuation, so no branch jump occurs.

Second, Figure 4 shows the through-thickness displacement profiles of the tracked S_0 branch below the cusp (0.02 MHz, $c = 3944$ m/s) and just above it (0.12 MHz, $c = 1571$ m/s). Below the cusp, the tracked branch is a bone-dominated extensional mode, axial-dominated and nearly uniform across the cortex with $P \simeq +1$ — the signature of an extensional (S_0) plate wave; across the cusp the same MAC-continuous branch undergoes a pronounced redistribution of displacement into the surrounding fluids, so that at ≈ 0.12 MHz the solid displacement-norm fraction falls to about 4% and the bone-midplane parity indicator is no longer a meaningful symmetry classifier. The profiles vary smoothly and support continuous MAC tracking of the same fundamental branch through this reorganization.

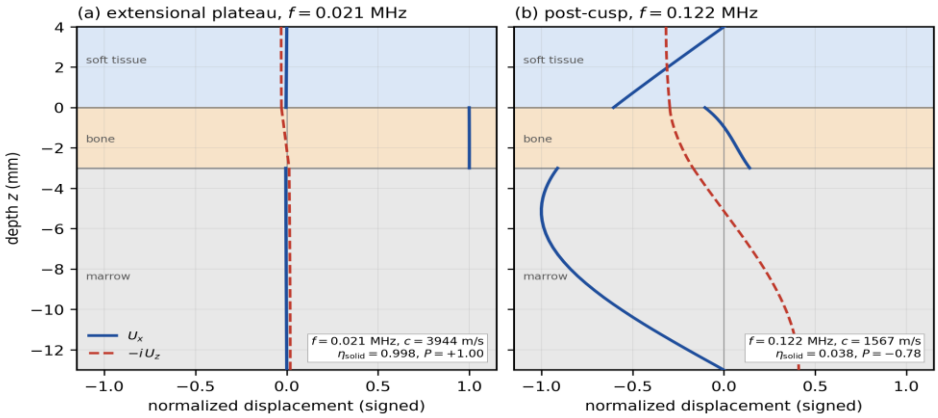

Fig. 4. Through-thickness phase-fixed signed mode shapes (U_x, solid; $-iU_z$, dashed) of the tracked S_0 branch (TI bone): (a) extensional plateau, $f = 0.021$ MHz; (b) post-cusp, $f = 0.122$ MHz.

The global phase is fixed on U_x; because U_x and U_z are in quadrature under the $e^{i(kx-\omega t)}$ convention, the pair $(U_x, -iU_z)$ is simultaneously real (Section 2.8(c)). The profiles support continuous MAC tracking of the same fundamental branch through its extensional-to-fluid-loaded reorganization.

IV.v. Porosity-induced softening of the S_0 plate wave

To quantify how cortical porosity affects the guided spectrum, we apply the elastic-material-with-voids model of Section 2.7 to the isotropic-from-axial bone of Section 4.1 ($\lambda = 17.46$, $\mu = 8.22$ GPa, $\rho = 1.87$ g/cm³), holding the elastic constants fixed and varying only the void coupling. We take $K_1 = 9$ GPa (the Andreaus pore stiffness $K_1^{\max} \approx 11.8$ GPa scaled to a cortical solid fraction), a pore length $\ell = 50$ μm giving $K_2 = K_1\ell^2$ and $\kappa = \ell^2$, and sweep $\hat{g} = K_3^2/(\lambda K_1) \in [0{,}0.9]$. For these parameters, the porosity-branch cutoff is $f_c = \omega_c/2\pi \approx 7.0$ MHz, roughly 140× the top of the 0–0.05 MHz window used below, over which the void field relaxes quasi-statically, and its sole in-band effect is to replace λ by $\lambda_{\text{eq}} = \lambda - K_3^2/K_1$.

The consequence is a downward shift of the S_0 extensional plateau (Figure 5). The branch is isolated at low wavenumber ($k = 8 \times 10^{-6}$–8×10^{-5} μm^{-1}) as the positive-parity, bone-dominated mode (parity $P \to +1$ and solid displacement-norm fraction $\eta_{\text{solid}} \to 1$ in the low-frequency plateau region used to seed and characterize the branch). At $\hat{g} = 0$ the plateau is 3645.6 m/s, identical to the elastic value of Section 4.3 — the $K_3 \to 0$ reduction confirmed to four significant figures, which validates the augmented assembly. As the coupling increases, the plateau falls monotonically: 3614.2 m/s at $\hat{g} = 0.1$, 3492.4 m/s at $\hat{g} = 0.4$, and 3103.9 m/s at $\hat{g} = 0.9$ (the maximum coupling explored in [III]), a 14.8% reduction. Because the plate velocity maps onto an effective stiffness through $c^2 \propto C$, neglecting the void coupling incurs a forward-model-equivalent stiffness discrepancy $(c_0/c)^2 - 1$ — the ratio of the plate modulus implied by the dense model to that implied by the porous one — reaching $\approx 9\%$ at $\hat{g} = 0.4$ and $\approx 38\%$ at $\hat{g} = 0.9$ (Figure 6). Section 2.8 confirms that the whole sweep satisfies the thermodynamic admissibility condition (6).

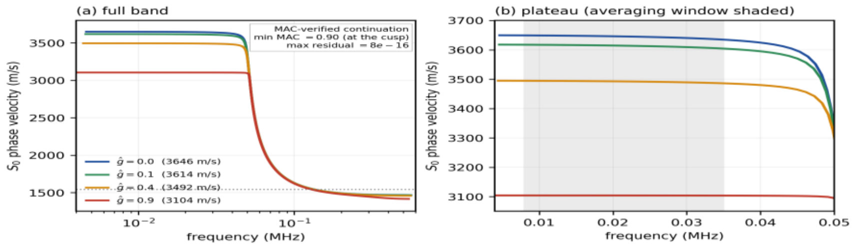

Fig. 5. Verified S_0 extensional branch of the soft-tissue / voids-cortical-bone / marrow trilayer for void couplings $\hat{g} = 0$ (elastic), 0.1, 0.4, 0.9. The plateau is the low-frequency material descriptor of Section 4.3; increasing $\hat{g}$ lowers it from 3645.6 to 3103.9 m/s. The $\hat{g} = 0$ curve reproduces the elastic result of Section 4.3 (reduction check).

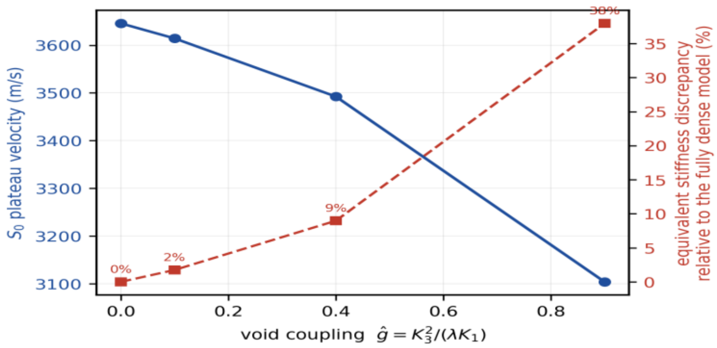

Fig. 6. Porosity-coupling sweep. Left axis (circles): S_0 plateau velocity versus void coupling $\hat{g}$. Right axis (squares): the forward-model-equivalent stiffness discrepancy $(c_0/c)^2 - 1$ incurred by a fully dense model that neglects the void coupling.

The void coupling reshapes the full modal spectrum, not only the S_0 plateau. Figure 7 shows the parity-resolved modal cloud at $\hat{g} = 0.4$ over 0–0.55 MHz: the S-like S_0 plateau near 3500 m/s, the A-like A_0 branch rising from zero toward the fluid speed, and the higher symmetric/antisymmetric branches with their cutoffs — the familiar trilayer spectrum, now computed with the augmented three-field bone operator.

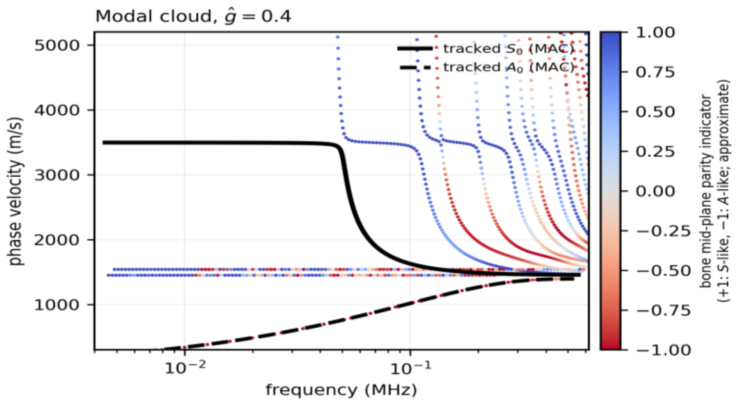

Fig. 7. Parity-resolved modal cloud of the soft-tissue / voids-cortical-bone / marrow trilayer at $\hat{g} = 0.4$. Colour encodes the approximate bone-midplane parity indicator ($+1$: S-like; -1: A-like).

The S_0 plateau, the A_0 branch, and the higher modes are all recovered from the single linear eigensolve of the three-field operator (4).

This is the *dynamic* counterpart of the static stiffness loss reported by Andreaus *et al.* [III]: the same volume-fraction coupling that lowers the quasi-static modulus of a remodelling cortex lowers the velocity of the guided wave used to probe it. The effect acts on the same descriptor as the anisotropy discrepancy of Section 4.3 but in the *opposite* sense: transverse isotropy raises the S_0 plateau relative to an isotropic surrogate (3943 against 3645 m/s), whereas void coupling lowers it (3103.9 against 3645.6 m/s at the strongest coupling). The two mis-specifications therefore displace the implied plate modulus in opposite directions and partially cancel rather than accumulate, so we do not claim that they compound; what they share is the S_0 plateau as a descriptor jointly sensitive to anisotropy and porosity — and, as Section 4.6(d) shows, one that cannot separate them. Matching a measured S_0 plateau to the two stiffness limits (λ at high frequency, λ_{eq} at low frequency) is, in principle, a route to estimating porosity-related parameters from guided-wave data, though we do not pursue inversion here.

IV.vi. Local sensitivity, practical identifiability and parameter correlation

The discrepancies quantified in Sections 4.3 and 4.5 are properties of the *forward* map: they state how the dispersion changes when anisotropy or porosity is imposed at known parameter values. Whether the same parameters can be *recovered* from dispersion data is a separate question, and guided-wave inversion is well known to be ill-conditioned: distinct parameter combinations can generate nearly indistinguishable dispersion. Because the numbers reported above are frequently read as inversion errors, we quantify the local identifiability of the constitutive parameters directly within the present framework. The analysis quantifies local (practical) identifiability and parameter confounding about the adopted reference parameter set; it is not intended as a global structural-identifiability proof.

(a) Sensitivities at no additional cost: The Hellmann–Feynman relation (3) used for the group velocity generalizes to any model parameter θ. Differentiating (2) with respect to θ and projecting onto the left eigenvector of a simple eigenvalue gives

$$\frac{\partial W}{\partial \theta} = \frac{\mathbf{y}^H(\partial \mathbf{A}/\partial \theta - W\ \partial \mathbf{B}/\partial \theta)\mathbf{x}}{\mathbf{y}^H \mathbf{B}\mathbf{x}}, \qquad \frac{\partial c}{\partial \theta} = \frac{1}{2k\sqrt{W}}\frac{\partial W}{\partial \theta}. \qquad (11)$$

Because $\mathbf{A}$ is *affine* in each stiffness $C_{11}, C_{13}, C_{33}, C_{55}$ and in K_1, K_2, K_3 — the layer blocks and the traction rows all depend linearly on them — the derivatives $\partial \mathbf{A}/\partial \theta$ are exact and require no finite differencing. Likewise $\partial \mathbf{A}/\partial \rho = \mathbf{0}$ while $\partial \mathbf{B}/\partial \rho$ is the indicator of the layer's degrees of freedom. Only the thickness derivative requires a difference quotient, through the affine map of Section 2.2. The complete sensitivity matrix is thus obtained from eigenpairs already computed, at the cost of one matrix–vector product per parameter per mode. (Sensitivity to the coupling $\hat{g}$ is best evaluated by differentiating with respect to K_3, in which $\mathbf{A}$ is affine, and converting by the chain rule through K_3^2; since $K_3 = \sqrt{\hat{g}\lambda K_1}$, differentiation with respect to $\hat{g}$ is singular at $\hat{g} = 0$ whereas W is even in K_3 and therefore smooth in K_3^2.)

(b) A closed-form degeneracy at low frequency: The structure of the ill-conditioning can be exhibited analytically. In the long-wavelength limit, the fundamental symmetric mode of a transversely isotropic plate is the extensional plate wave with

$$c_{\text{ext}}^2 = \frac{M}{\rho}, \qquad M \equiv C_{11} - \frac{C_{13}^2}{C_{33}}. \qquad (12)$$

Equation (12) reproduces the computed plateaux of Sections 4.3 and 4.5 to better than 0.11% (TI: 3946.7 against 3943 m/s; isotropic-from-axial: 3649.6 against 3645.6 m/s). Its logarithmic sensitivities are

$$\frac{\partial \ln c}{\partial \ln C_{11}} = \frac{C_{11}}{2M}, \quad \frac{\partial \ln c}{\partial \ln C_{13}} = -\frac{C_{13}^2}{MC_{33}}, \quad \frac{\partial \ln c}{\partial \ln C_{33}} = \frac{C_{13}^2}{2MC_{33}}, \quad \frac{\partial \ln c}{\partial \ln C_{55}} = 0,$$
$$\frac{\partial \ln c}{\partial \ln \rho} = -1/2, \quad \frac{\partial \ln c}{\partial \ln h} = 0. \qquad (13)$$

Three consequences follow immediately. First, the three stiffness sensitivities sum to $+1/2$, exactly cancelling the density sensitivity: the simultaneous rescaling $C_{ij} \to \alpha C_{ij}$, $\rho \to \alpha\rho$ leaves the plateau unchanged, so absolute stiffness and density are not separately identifiable from the plateau. For a free plate, this degeneracy is exact at all frequencies; fluid loading breaks it only weakly, through the fixed impedance contrast with the surrounding media, so it survives as a small singular value rather than a strict null direction. Second, C_{55} has *identically zero* plateau sensitivity and can be constrained only by A_0 or by higher branches. Third, cortical thickness does not appear in (12) at all: h is invisible in the plateau value and is constrained only through the frequency scale at which dispersion sets in.

Evaluating (11) numerically on the tracked S_0 branch over the averaging window of Section 4.3 reproduces (13) to three decimal places, which is a stringent joint check on the sensitivity implementation and on the long-wavelength interpretation (Table 5).

Table 5. Analytic and computed logarithmic sensitivities of the S_0 plateau velocity.

	C_{11}	C_{13}	C_{33}	C_{55}	ρ	h	$\sum_{\text{stiff}}$
analytic, Eq. (13)	+0.582	-0.164	+0.082	0	-0.500	0	+0.500
computed, Eq. (11)	+0.582	-0.165	+0.083	+0.000	-0.499	-0.001	+0.4998

(c) Jacobian rank and Fisher information: To quantify what multi-mode data recovers, we assemble the log-sensitivity Jacobian

$$J_{ij} = \frac{\partial \ln c_{m_i}(f_i)}{\partial \ln \theta_j}, \qquad \boldsymbol{\theta} = (C_{11}, C_{13}, C_{33}, C_{55}, \rho, h),$$

over 150 wavenumbers per branch, spanning 0.003–0.63 MHz on S_0 and 0.001–0.53 MHz on A_0, with branches followed by the MAC continuation of Section 2.6 seeded at $k = 3 \times 10^{-4}\ \mu\text{m}^{-1}$ and tracked outward in both directions. Assuming independent, zero-mean relative phase-velocity errors of standard deviation σ — equivalently, to first order, errors of the same standard deviation in ln c, consistent with the logarithmic Jacobian above — the Fisher information is $\mathbf{F} = \sigma^{-2} J^{\top} J$, the Cramér–Rao covariance bound is $\mathbf{F}^{-1}$, and the parameter correlation matrix is $R_{ij} = (\mathbf{F}^{-1})_{ij} / \sqrt{(\mathbf{F}^{-1})_{ii} (\mathbf{F}^{-1})_{jj}}$; we take $\sigma = 1\%$, that is a 1% relative error in phase velocity, representative of axial-transmission phase-velocity estimation. The correlations are independent of σ; the bounds scale linearly with it.

Table 6: Normalized singular values of the log-sensitivity Jacobian, condition number, and numerical rank at a relative tolerance of 10^{-3}.

observation set	σ_2/σ_1	σ_3/σ_1	σ_4/σ_1	σ_5/σ_1	σ_6/σ_1	$\kappa(J)$	rank
S_0, plateau window only	2.5×10^{-3}	2.6×10^{-5}	5.6×10^{-7}	1.2×10^{-7}	6.6×10^{-12}	1.52×10^{11}	**2**
S_0, full band	7.0×10^{-2}	3.4×10^{-2}	2.9×10^{-2}	5.3×10^{-3}	1.5×10^{-3}	6.71×10^{2}	6
$S_0 + A_0$, full band	4.9×10^{-1}	1.2×10^{-1}	2.5×10^{-2}	1.2×10^{-2}	4.3×10^{-3}	2.35×10^{2}	6

The plateau window alone is numerically of rank 2 out of 6, with condition number 1.5×10^{11}: it constrains the single combination M/ρ of (12), exactly as the closed form predicts, and its least-identifiable direction is C_{55} to five decimal places. Broadening to the full S_0 band restores nominal full rank but leaves the problem poorly conditioned; adding A_0 improves the conditioning by a further factor of about three. The near-degenerate direction predicted in (b) behaves as anticipated: measuring $\| J\mathbf{v} \| / \| J \|_2$ for $\mathbf{v} \propto (1,1,1,1,1,0)$ gives 4.1×10^{-4} on the plateau (effectively null), 1.6×10^{-2} over the full S_0 band, and 1.0×10^{-1} for $S_0 + A_0$, so fluid loading and multi-mode data progressively — but only partially — break the stiffness–density degeneracy.

Table 7. Cramér–Rao relative standard deviations at $\sigma = 1\%$, and the parameter correlation matrix for the joint $S_0 + A_0$ observation set.

	C_{11}	C_{13}	C_{33}	C_{55}	ρ	h
CRLB, S_0 alone	44.7%	25.3%	37.1%	57.0%	42.7%	42.4%
CRLB, $S_0 + A_0$	5.4%	16.9%	9.0%	4.1%	4.3%	1.1%

R	C_{11}	C_{13}	C_{33}	C_{55}	ρ	h
C_{11}	1.000	0.702	0.559	0.388	0.722	-0.734
C_{13}	0.702	1.000	0.911	-0.357	0.028	-0.043
C_{33}	0.559	0.911	1.000	-0.368	-0.029	0.027
C_{55}	0.388	-0.357	-0.368	1.000	0.900	-0.894
ρ	0.722	0.028	-0.029	0.900	1.000	-0.992
h	-0.734	-0.043	0.027	-0.894	-0.992	1.000

Three features matter for inversion practice. Joint $S_0 + A_0$ inversion reduces the attainable uncertainties by roughly a factor of five relative to S_0 alone, the largest gains being in thickness (42.4% → 1.1%) and in C_{55} (57.0% → 4.1%) — precisely the two quantities to which the plateau is blind, which is a quantitative argument for multi-mode rather than plateau-only inversion. The residual worst-determined parameter is C_{13} (16.9%), and the least-identifiable direction is dominated by C_{13} and C_{33}, the pair entering (12) only through the ratio C_{13}^2/C_{33}, with a mutual correlation of 0.911. Finally, ρ and h remain almost perfectly anticorrelated (-0.992), so an independent constraint on either propagates directly into the recovered stiffnesses.

(d) Confounding of the void coupling: The porosity extension is subject to a related limitation. In the quasi-static in-band limit, the void field enters only through $\lambda_{\text{eq}} = \lambda - K_3^2/K_1$ (Section 2.7), so K_3 is locally confounded with the volumetric Lamé stiffness λ: evaluating (11) on the S_0 plateau of the voids model, the sensitivity columns for $\hat{g}$ and $\ln\lambda$ are antiparallel to within 4×10^{-5} relative at every sampled wavenumber and at $\hat{g} = 0.1$, 0.4, 0.9, and the corresponding direction is null to 1.9×10^{-6} in $\| J\mathbf{v} \| / \| J \|_2$. Consequently, porosity-related coupling cannot be uniquely separated from elastic-constitutive uncertainty in volumetric stiffness using the S_0 plateau alone. We state this deliberately in terms of volumetric stiffness rather than of anisotropy in general: general transverse isotropy alters C_{11}, C_{13}, C_{33} and C_{55} in several distinct ways and may retain distinguishable signatures across multiple modes and frequencies, whereas the isotropic void coupling adopted here acts on a single volumetric combination. Separating the two would require data approaching the porosity cutoff $f_c \approx 7$ MHz, an independent estimate of ρ or h, or a prior linking $\hat{g}$ to density through the solid volume fraction.

(e) Consequences for the reported discrepancies: The analysis above does not alter any computed result of Sections 4.3 and 4.5, but it fixes their interpretation. The ≈ 17% and ≈ 38% figures are forward-model-equivalent stiffness discrepancies — the difference in the plate modulus implied by two forward models at the stated parameter sets — and not claims that these quantities are uniquely recoverable by inversion. The map $c = c(C_{11}, C_{13}, C_{33}, C_{55}, \rho, h, \ldots)$ is locally many-to-one, and (13) shows that anisotropy and porosity both act on the same identifiable combination in the plateau. The defensible inferential statement is therefore conditional: *if* thickness and density are constrained independently — as they are in practice by site-matched radiography, DXA, or a companion pulse-echo measurement — and *if* multiple modes are inverted jointly rather than the plateau alone, then neglecting transverse isotropy or porosity introduces stiffness discrepancies of the magnitudes reported here. We regard this conditional framing, together with the sensitivity machinery of (11), as the more useful contribution: the same eigenpairs that give the dispersion also give the Jacobian needed to assess whether a proposed inversion is well posed before it is attempted.

V. Discussion

Relation to the analytical model: The spectral-collocation pencil (2) is the discrete counterpart of the analytical characteristic equation of [XV]. Where the analytical model encodes the through-thickness physics in partial-wave amplitudes and locates the zeros of a 8×8 determinant, the collocation model encodes the same physics in differentiation matrices and reads off the modes as eigenvalues. The two are equivalent in the continuum limit — the agreement of Table 2 to better than 10^{-5} % confirms this — but the eigenvalue route returns all modes at once and, with the MAC continuation of Section 2.6, does not skip or double-count branches near osculations.

Why the linear formulation matters: Expressed as $\omega(k)$, the eigenvalue is $W = \omega^2$ and every boundary row is W-free, so (2) is a genuine linear generalized eigenproblem solved by one QZ factorization per wavenumber. Casting the problem as $k(\omega)$ would make the operator quadratic in the eigenparameter and require a companion-matrix linearization; the $\omega(k)$ form avoids this and makes the Hellmann–Feynman group velocity (3) immediate.

Anisotropy: The TI extension required no change to the solver — four stiffnesses in place of two Lamé constants. Orthotropy, an off-axis (monoclinic) cut, or depth-graded porosity could be incorporated by supplying the appropriate (possibly z-dependent) coefficients, since the differentiation matrices accommodate continuous property variation [VII, X]. The 8.2% S_0 plate-velocity gap provides a *representative estimate* of the modelling error of an isotropic inversion under a deliberately favourable isotropic surrogate — one tuned to the bone's own fast-axis speeds. We do not claim it as a bound: a different isotropic surrogate, or compensating changes in density and thickness, could make the discrepancy either larger or smaller, as the correlations of Section 4.6 make explicit.

V.i. Limitations

The present model is deliberately reduced, and several limitations bound its quantitative use:

- **Inviscid fluids and elastic bone:** Both fluids are inviscid, and the bone is purely elastic; absorption — significant in cortical bone and modelled elsewhere with complex (Kelvin–Voigt) moduli [XVI] — is omitted, so the curves are real-valued dispersion only and predict no attenuation. Complex moduli would make (2) complex but leave the linear-in-W structure intact.
- **Homogeneous cortex:** The cortex is treated as homogeneous; real cortical bone is porosity-graded across its thickness, which the collocation framework could represent through z-dependent stiffnesses (and z-dependent void coefficients) but which is not done here.
- **Elastic void field:** Cortical porosity is now represented constitutively through the elastic-material-with-voids extension of Sections 2.7 and 4.5, but the void field is purely elastic: viscous pore dissipation and the interstitial saturating fluid are not modelled [III], so the model still predicts no

amplitude decay. The void coefficients are taken from the remodelling literature [III] rather than fitted to QUS data, and the reported softening is a parametric sweep rather than a calibrated figure.

- **Planar geometry:** The trilayer is planar with finite fluid layers and free outer faces; in vivo the cortex is curved, and the surrounding media are effectively half-spaces. The flat-plate approximation is justified locally (Section 1) but introduces error for thick cortices or strongly curved sites.
- **Dispersion only; no forced response:** This work reproduces the *dispersion* analysis of [XV] but not its reciprocity-based computation of the *modal amplitudes* excited by a time-harmonic load. The collocation eigenvectors could serve as the free-wave fields in those reciprocity integrals — a natural extension.
- **Local, not global, identifiability:** Section 4.6 quantifies first-order (Fisher) identifiability about the reference parameter set, under an assumed 1% relative velocity uncertainty and with the fluid layers treated as known. Global non-uniqueness over the full physiological parameter range is not excluded and would require a sampling-based study; the reported Cramér–Rao values are therefore best-case bounds rather than expected inversion errors.
- **Isotropic void coupling:** The void field is coupled to the dilatation through the single scalar K_3, so in band it acts on one volumetric combination only (Section 4.6(d)). A void coupling anisotropic in the sense of the matrix symmetry is not considered, and the confounding reported in Section 4.6(d) is specific to the isotropic coupling adopted here.
- **No experimental validation:** The validation here is against an independent analytical solver, not measured data; comparison with axial-transmission measurements remains future work.

VI. Conclusions

We have reformulated the fluid–solid–fluid bone-waveguide model of Nguyen *et al.* [XV] as a linear generalized eigenvalue problem using a Chebyshev spectral-collocation discretization, with inviscid fluids represented as zero-shear solids coupled to the cortical layer through slip interface conditions imposed by replacing eight scalar rows of the assembled pencil (ten when the void field is active). Because those rows are algebraic, the resulting system is a descriptor pencil with a singular mass matrix; Appendix B shows it to be regular whenever the imposed constraints are independent, and its finite spectrum to coincide exactly with that of explicit constraint elimination. The method returns the finite discrete spectrum in a single generalized eigensolve at each wavenumber, from which the physical guided modes are selected by the admissibility and participation filters of Section 2.3; group velocities follow from a Hellmann–Feynman relation with a numerical k-derivative of the system matrix. Validated against an independently implemented analytical global-matrix solver for the water/aluminum/water trilayer, the fundamental extensional (S_0) and flexural (A_0) phase velocities agree to better than 10^{-5} %, both branches being identified in each solver independently by rank ordering the admissible eigenvalues at fixed wavenumber rather than by any mode-shape heuristic, and the eigenfrequencies converge spectrally with round-off-level residuals.

Extending the solid layer to a transversely isotropic human cortical bone — without altering the numerical solution procedure, apart from supplying the four sagittal-plane stiffness coefficients — shows that anisotropy raises the low-frequency S_0 extensional plate velocity to 3943 m/s, 8.2% above the best isotropic surrogate built from the same axial speeds, a forward-model-equivalent plate-stiffness discrepancy of $\approx 17\%$; MAC overlap, and through-thickness mode shapes confirm the fundamental branches are tracked continuously. Representing cortical porosity constitutively — by modelling the bone as a linear elastic material with voids [III, VI, XVII], which adds one volume-fraction field per bone node without disturbing the linear-in-W structure — further lowers the S_0 plateau by up to 14.8% as the void coupling spans the range used in bone-remodelling studies, a forward-model-equivalent plate-stiffness discrepancy of $\approx 38\%$ at the strongest coupling explored, which Section 2.8 confirms to be thermodynamically admissible. This dynamic softening is the guided-wave counterpart of the static stiffness loss of Andreaus *et al.* [III]. It acts in the opposite sense to the anisotropy discrepancy, since transverse isotropy raises the same plateau that porosity lowers, so the two partially cancel rather than compound. A local sensitivity and identifiability analysis (Section 4.6) shows further that the plateau constrains the single combination essentially $(C_{11} - C_{13}^2/C_{33})/\rho$, with a much weaker second direction induced by fluid loading (numerical rank two in six parameters), that thickness and C_{55} are invisible in it, and that the void coupling is locally confounded with volumetric stiffness; the two quantities above are therefore forward-model discrepancies, and their recovery by inversion requires multi-mode data together with independent constraints on thickness and density. The framework is a compact, robust, and extensible basis for forward modelling in cortical-bone quantitative ultrasound. We emphasize that its clinical-inversion relevance is contingent on extensions not addressed here — viscoelastic absorption, porosity grading, realistic cylindrical geometry, measurement uncertainty, and experimental validation — and the quantitative anisotropy-induced forward-model discrepancy reported above should be read as motivation for TI inversion models rather than as a calibrated clinical figure.

Appendix A: Assembly details and row-replacement map

The map below is for the elastic trilayer, in which eight rows are replaced. When the void field of Section 2.7 is active, the bone block carries a third nodal field, and two further rows are replaced by the free-pore conditions $\partial_z \xi = 0$ at the two bone faces, giving ten replaced rows; the eight displacement constraints are unchanged except that σ_{zz} on the bone faces acquires the $K_3 \xi$ term.

For a layer with $M = N + 1$ nodes, the unknowns are ordered $[\,U_x(z_0), \dots, U_x(z_N),\, U_z(z_0), \dots, U_z(z_N)\,]$, with global offset off_ℓ for layer ℓ. With I the $M \times M$ identity, the $2M \times 2M$ layer block of (2) is

$$\mathbf{A}_{\text{layer}} = -\begin{bmatrix} C_{55}D^2 - C_{11}k^2 I & ik(C_{13} + C_{55})D \\ ik(C_{13} + C_{55})D & C_{33}D^2 - C_{55}k^2 I \end{bmatrix}, \qquad \mathbf{B}_{\text{layer}} = \rho \begin{bmatrix} I & 0 \\ 0 & I \end{bmatrix}.$$

The traction/displacement rows that replace selected equation-of-motion rows are, for node i of layer ℓ (with $\mathrm{ux}_\ell(i) = \mathrm{off}_\ell + i$, $\mathrm{uz}_\ell(i) = \mathrm{off}_\ell + M + i$):

$$[\sigma_{zz}]_i\colon\ ik\,C_{13}^{(\ell)}\,U_x(z_i) + C_{33}^{(\ell)}\sum_j D_{ij}^{(\ell)}\,U_z(z_j), [\sigma_{xz}]_i\colon\ C_{55}^{(\ell)}\left(\sum_j D_{ij}^{(\ell)}\,U_x(z_j)\,ik\,U_z(z_i) + ik\,U_z(z_i)\right),$$

$$[u_z]_i\colon U_z(z_i).$$

Duplicate interface nodes: Each fluid–solid interface is a coincident pair of grid nodes — the fluid layer's near-face node and the solid layer's near-face node both lie at the same physical z. They are *kept as separate unknowns* (so tangential slip is admitted); the three interface conditions couple them. The exact eight replaced global rows are:

#	replaced global row	condition imposed
1	$\mathrm{uz}_{\mathrm{soft}}(0)$	$\sigma_{zz}^{\mathrm{soft}} = 0$ (top free surface)
2	$\mathrm{uz}_{\mathrm{marrow}}(N)$	$\sigma_{zz}^{\mathrm{marrow}} = 0$ (bottom free surface)
3	$\mathrm{uz}_{\mathrm{soft}}(N)$	$u_z^{\mathrm{soft}} - u_z^{\mathrm{bone}} = 0$ (interface 1)
4	$\mathrm{uz}_{\mathrm{bone}}(0)$	$\sigma_{zz}^{\mathrm{bone}} - \sigma_{zz}^{\mathrm{soft}} = 0$ (interface 1)
5	$\mathrm{ux}_{\mathrm{bone}}(0)$	$\sigma_{xz}^{\mathrm{bone}} = 0$ (interface 1)
6	$\mathrm{uz}_{\mathrm{marrow}}(0)$	$u_z^{\mathrm{marrow}} - u_z^{\mathrm{bone}} = 0$ (interface 2)
7	$\mathrm{uz}_{\mathrm{bone}}(N)$	$\sigma_{zz}^{\mathrm{bone}} - \sigma_{zz}^{\mathrm{marrow}} = 0$ (interface 2)
8	$\mathrm{ux}_{\mathrm{bone}}(N)$	$\sigma_{xz}^{\mathrm{bone}} = 0$ (interface 2)

The matched rows of $\mathbf{B}$ are zeroed so the constraints are algebraic, preserving the linearity of (2) in W. In pseudocode: assemble block-diagonal $\mathbf{A}, \mathbf{B}$ from (1); for each of the eight rows above, overwrite $\mathbf{A}[\mathrm{row},:]$ with the constraint vector and set $\mathbf{B}[\mathrm{row},:] = 0$; solve $\mathbf{Ax} = W\mathbf{Bx}$.

Appendix B: Algebraic properties of the row-replaced pencil

Section 2.4 imposes the boundary and interface conditions by overwriting selected rows of $\mathbf{A}$ and zeroing the matched rows of $\mathbf{B}$. Because that operation destroys both the symmetry and the definiteness of the mass matrix, we establish here that it nevertheless preserves the algebraic structure required for the generalized eigenproblem to be solvable, and that it introduces no spurious finite modes. Throughout, q denotes the number of replaced rows: $q = 8$ for the elastic trilayer of Sections 2.1–2.6, and $q = 10$ when the void field is active, the two additional rows being the free-pore conditions $\partial_z\xi = 0$ at the bone faces (Section 2.7).

B.1 The constrained problem is a descriptor pencil

Before boundary and interface enforcement, each dynamical layer possesses a positive-definite mass block $\rho^{(\ell)}\mathbf{I}$. After row replacement, the global mass matrix is **singular**, because the q rows associated with the algebraic constraints are zeroed. The assembled problem $\mathbf{A}(k)\mathbf{x} = W\mathbf{Bx}$ is therefore not a definite generalized eigenproblem but a **descriptor** (singular-$\mathbf{B}$) generalized eigenproblem, whose spectrum comprises finite eigenvalues together with infinite eigenvalues associated with the algebraic constraints. This is the correct setting for the analysis below, and the QZ algorithm used in Section 2.4 is appropriate precisely because it handles singular $\mathbf{B}$.

Let $\mathcal{J} \subset \{1, \dots, n\}$ be the set of q replaced row indices and $\mathcal{K}$ its complement, and permute the unknowns as $\mathbf{x} = (\mathbf{x}_{\mathcal{J}}, \mathbf{x}_{\mathcal{K}})$. The pencil then reads

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{\mathcal{K}\mathcal{J}} & \mathbf{A}_{\mathcal{K}\mathcal{K}} \\ \mathbf{C}_{\mathcal{J}} & \mathbf{C}_{\mathcal{K}} \end{bmatrix}, \qquad \mathbf{B} = \begin{bmatrix} \mathbf{0} & \mathbf{R} \\ \mathbf{0} & \mathbf{0} \end{bmatrix},$$

where $\mathbf{C} = [\mathbf{C}_{\mathcal{J}} \;\; \mathbf{C}_{\mathcal{K}}]$ collects the q constraint rows of Appendix A and $\mathbf{R} = \mathrm{diag}(\rho_p)_{p\in\mathcal{K}}$ is the surviving, strictly positive part of the mass matrix. The zero block in the first column of $\mathbf{B}$ is exact: $\mathbf{B}$ is diagonal, so a retained row $p \in \mathcal{K}$ has its only nonzero entry in column $p \in \mathcal{K}$.

B.2 A sufficient regularity criterion

Since only the $n-q$ retained rows carry W, $\det(\mathbf{A} - W\mathbf{B})$ is a polynomial in W of degree at most $n-q$, and expansion gives its leading coefficient as

$$[W^{\,n-q}]\,\det(\mathbf{A} - W\mathbf{B}) = \pm\det(\mathbf{R})\,\det(\mathbf{C}_{\mathcal{J}}), \qquad \text{(B.1)}$$

with $\det\mathbf{R} = \prod_{p\in\mathcal{K}} \rho_p > 0$. This yields a criterion that is checkable at assembly time at negligible cost:

Sufficient regularity criterion: If the constraint submatrix $\mathbf{C}_{\mathcal{J}}$ associated with the replaced degrees of freedom is nonsingular, then the row-replaced pencil is regular, possesses $n-q$ finite generalized eigenvalues counting algebraic multiplicity, together with q infinite eigenvalues associated with the imposed algebraic constraints, and its finite spectrum coincides with that of the reduced constraint-eliminated problem (B.2), whose mass matrix is positive definite (Appendix B.3); the finite eigenvalue count is therefore exact.

We state this as a *sufficient* condition only. The converse does not follow: $\det\mathbf{C}_{\mathcal{J}} = 0$ makes the leading coefficient in (B.1) vanish and hence lowers the degree of $\det(\mathbf{A} - W\mathbf{B})$, which implies additional infinite eigenvalues, but it does not by itself render the descriptor pencil irregular. Establishing necessity would require a finer descriptor-system argument that we do not need, since the criterion is verified directly in every computation reported here.

With the row/degree-of-freedom pairing of Appendix A, the diagonal entries of $\mathbf{C}_{\mathcal{J}}$ are dominated by $C_{33}D_{00}$, $C_{55}D_{00}$, K_2D_{00} or unity. For the TI cortical trilayer we find $\sigma_{\min}(\mathbf{C}_{\mathcal{J}}) \geq 7.7\times10^{-2}$ and $\mathrm{cond}_2(\mathbf{C}_{\mathcal{J}}) \leq 20$ over the entire wavenumber sweep and every material configuration reported, the criterion becoming *better* conditioned under refinement ($\sigma_{\min} = 0.078,\ 0.192,\ 0.310$ at $N = (14{,}12{,}28),\ (22{,}18{,}44),\ (28{,}26{,}56)$). No marginal case arises at the orders used.

B.3 Equivalence to explicit constraint elimination

When $\mathbf{C}_{\mathcal{J}}$ is nonsingular, the constraint rows $\mathbf{C}_{\mathcal{J}}\mathbf{x}_{\mathcal{J}} + \mathbf{C}_{\mathcal{K}}\mathbf{x}_{\mathcal{K}} = \mathbf{0}$ can be solved for the replaced unknowns, $\mathbf{x}_{\mathcal{J}} = -\mathbf{C}_{\mathcal{J}}^{-1}\mathbf{C}_{\mathcal{K}}\mathbf{x}_{\mathcal{K}}$, and substitution into the retained equation-of-motion rows gives the reduced problem

$$\widehat{\mathbf{A}}\,\mathbf{x}_{\mathcal{K}} = W\,\mathbf{R}\,\mathbf{x}_{\mathcal{K}}, \qquad \widehat{\mathbf{A}} = \mathbf{A}_{\mathcal{K}\mathcal{K}} - \mathbf{A}_{\mathcal{K}\mathcal{J}}\,\mathbf{C}_{\mathcal{J}}^{-1}\mathbf{C}_{\mathcal{K}}, \qquad \text{(B.2)}$$

of dimension $n-q$ with a positive-definite mass matrix $\mathbf{R}$. Equation (B.2) is an algebraic identity, not an approximation: row replacement is not arbitrarily altering the finite physical problem but is algebraically equivalent to explicit constraint elimination. This is the central answer to the concern that row replacement might not preserve the intended spectrum.

Two consequences follow. First, $W = 0$ can be an eigenvalue only if $\widehat{\mathbf{A}}$ is singular, that is, only if the constrained stiffness operator admits a non-trivial static mode; row replacement introduces no artificial null mode of its own. We do not claim that $\widehat{\mathbf{A}}$ is positive definite — it cannot be, since the zero-shear fluid subspaces are degenerate by construction (Section 2.8(e)) — and we therefore verify the absence of artificial null modes numerically rather than by appeal to definiteness. Across every grid and wavenumber tested, the number of finite eigenvalues returned by QZ is exactly $n - q$, and the smallest finite eigenvalue is in each case identifiable as a physical mode: for the TI trilayer at $k = 8 \times 10^{-4}\ \mu\mathrm{m}^{-1}$ it is $W = 9.31 \times 10^{-7}$, that is $f_1 = 0.15357$ MHz, the lowest mode already tabulated in Section 3.1; in the voids configuration at small wavenumber it is the fundamental flexural branch, identified by the scalings $c \propto k$ and $W \propto k^4$ characteristic of A_0.

Second, we confirmed the identity (B.2) directly by forming $\widehat{\mathbf{A}}$ and solving the reduced problem independently. The two spectra contain the same number of admissible modes and agree to within 4.5×10^{-12} and 3.2×10^{-10} in relative eigenvalue difference over the twenty lowest modes, at every collocation order of Table 3.

B.4 Symmetry and conditioning

Row replacement destroys the symmetry that the unconstrained operator possesses in the transformed variables of Section 2.8(c): the Schur complement in (B.2) mixes constraint rows into the interior equations, so $\widehat{\mathbf{A}}$ is generally nonsymmetric. As noted in Section 2.8(c), the discrete pencil therefore does not inherit exact matrix self-adjointness, and reality of the computed eigenvalues is a property to be verified rather than assumed; this is why the imaginary-part test of Section 2.3 is imposed and why residual norms are reported for every retained mode.

The assembled $\mathbf{A}$ is severely ill-conditioned in the matrix norm, $\mathrm{cond}_2(\mathbf{A}) \sim 10^{22}$–$10^{23}$, owing to the $O(N^4)$ growth of Chebyshev second-derivative matrices together with a unit system spanning many orders of magnitude. This is intrinsic to the discretization and is *not* caused by the interface treatment: equilibrating the constraint rows to unit ∞-norm before the eigensolve, which leaves the spectrum unchanged, does not reduce it, confirming that the constraint rows are not the source.

The matrix condition number is, however, a pessimistic bound that does not transfer to the quantities actually reported. The operative measure is the condition number of the individual eigenvalues,

$$\kappa(W_m) = \frac{\| \mathbf{y}_m \|_2 \| \mathbf{x}_m \|_2}{|\mathbf{y}_m^H \mathbf{B} \mathbf{x}_m|},$$

and for the twenty lowest retained modes we find $\kappa(W_m) \leq 1.47$ at every collocation order and both wavenumbers of Table 3, indicating that the retained eigenvalues are individually well conditioned despite the large matrix condition number. This resolves the apparent tension between the large $\mathrm{cond}_2(\mathbf{A})$ and the round-off-level residuals of Table 3: the retained modes are well separated, and their left and right eigenvectors are far from orthogonal, so the $O(N^4)$ growth in the matrix norm does not propagate into the reported eigenfrequencies. Consistently, $\kappa(W_m)$ rises only near the S_0 cusp, precisely where the MAC dip of Section 4.4 occurs and where Section 2.5 already cautions against over-interpreting c_g.

Table B.1. Algebraic diagnostics of the row-replaced pencil for the TI cortical-bone trilayer ($q = 8$).

$(N_{\text{soft}}, N_{\text{bone}}, N_{\text{marrow}})$	k (μm^{-1})	n	finite eigs (expected)	$\sigma_{\min}(\mathbf{C}_g)$	$\text{cond}_2\mathbf{C}_g$	$\max_m \kappa(W_m)$	(B.2) agreement
(14,12,28)	8×10^{-4}	114	106 (106)	7.77×10^{-2}	20	1.43	4.5×10^{-11}
(14,12,28)	1.6×10^{-3}	114	106 (106)	7.77×10^{-2}	20	1.47	4.5×10^{-12}
(22,18,44)	8×10^{-4}	174	166 (166)	1.92×10^{-1}	11	1.43	3.8×10^{-11}
(22,18,44)	1.6×10^{-3}	174	166 (166)	1.92×10^{-1}	11	1.45	8.0×10^{-11}
(28,26,56)	8×10^{-4}	226	218 (218)	3.10×10^{-1}	13	1.44	3.2×10^{-10}
(28,26,56)	1.6×10^{-3}	226	218 (218)	3.10×10^{-1}	13	1.44	1.1×10^{-10}

Finite-eigenvalue counts are exact at every entry; $\boldsymbol{\kappa}$ is the maximum eigenvalue condition number over the twenty lowest retained modes; the last column is the maximum relative eigenvalue difference against independent null-space elimination via (B.2).

Appendix C. Reproducibility

Environment. Python 3; NumPy 2.4.4; SciPy 1.17.1; Matplotlib. Dense generalized eigensolves use scipy.linalg.eig.

Verification of the void extension: The Chebyshev first-derivative operator used in the void block was checked against the through-thickness reflection $z \to -z$ (under which $D \to -D$ while D^2 is invariant), so that the eigenfrequencies of the symmetric extensional branch are independent of the operator's orientation; the S_0 branch is seeded in the low-frequency region using its positive-parity, bone-dominated character and subsequently continued by MAC; neither parity nor solid displacement participation is required to remain fixed after the cusp (Figure 5 and Section 4.4). The $\hat{g} \to 0$ limit reproduces the elastic plateau of Section 4.3 to four significant figures. For the void-augmented branches, the normalized eigen-residuals $\| \mathbf{Ax} - W\mathbf{Bx} \|_2 / (\| \mathbf{A} \|_2 \| \mathbf{x} \|_2)$ remain at the 10^{-16} level over the entire tracked sweep, i.e., at round-off. We note that normalizing instead by $\| \mathbf{Ax} \|_2$ is misleading here: the spectrum contains very slow, nearly static modes for which $\mathbf{Ax}$ is itself almost null (at $\hat{g} = 0.1$, $k = 1.12 \times 10^{-4}\ \mu\text{m}^{-1}$ we measure $\| \mathbf{Ax} \|_2 = 6.2 \times 10^{-10}$ against $\| \mathbf{A} \|_2 \| \mathbf{x} \|_2 = 1.1 \times 10^2$), so that ratio inflates a backward error of 2.5×10^{-16} to 4.4×10^{-5}. All residuals quoted in this paper use the $\| \mathbf{A} \|_2 \| \mathbf{x} \|_2$ normalization.

Conflict of Interest

The authors declare that they have no conflict of interest.

References

I. J. D. Achenbach, *Wave Propagation in Elastic Solids*, North-Holland, Amsterdam (1973).

II. A. T. I. Adamou, R. V. Craster, *Spectral methods for modelling guided waves in elastic media*, Journal of the Acoustical Society of America **116**(3) (2004) 1524–1535. 10.1121/1.1777871

III. U. Andreaus, I. Giorgio, A. Madeo, *Modeling of the interaction between bone tissue and resorbable biomaterial as linear elastic materials with voids*, Zeitschrift für angewandte Mathematik und Physik (ZAMP) **66**(1) (2015) 209–237. 10.1007/s00033-014-0403-z

IV. N. Bochud, Q. Vallet, Y. Bala, H. Follet, J.-G. Minonzio, P. Laugier, *Genetic algorithms-based inversion of multimode guided waves for cortical bone characterization*, Physics in Medicine and Biology **61**(19) (2016) 6953–6974. 10.1088/0031-9155/61/19/6953

V. A. Chaboty, V.-H. Nguyen, G. Haïat, P. Bélanger, *Cortical bone plate properties assessment using inversion of axially transmitted low-frequency ultrasonic guided waves*, Journal of the Acoustical Society of America **156**(2) (2024) 954–967. 10.1121/10.0028173

VI. S. C. Cowin, J. W. Nunziato, *Linear elastic materials with voids*, Journal of Elasticity **13**(2) (1983) 125–147. 10.1007/BF00041230

VII. M. A. Denolle, E. M. Dunham, G. C. Beroza, *Solving the surface-wave eigenproblem with Chebyshev spectral collocation*, Bulletin of the Seismological Society of America **102**(3) (2012) 1214–1223. 10.1785/0120110183

VIII. M. A. Goodman, S. C. Cowin, *A continuum theory for granular materials*, Archive for Rational Mechanics and Analysis **44**(4) (1972) 249–266. 10.1007/BF00284326

IX. M. Granke, Q. Grimal, A. Saïed, P. Nauleau, F. Peyrin, P. Laugier, *Change in porosity is the major determinant of the variation of cortical bone elasticity at the millimeter scale in aged women*, Bone **49**(5) (2011) 1020–1026. 10.1016/j.bone.2011.08.002

X. F. Hernando Quintanilla, Z. Fan, M. J. S. Lowe, R. V. Craster, *Guided waves' dispersion curves in anisotropic viscoelastic single- and multi-layered media*, Proceedings of the Royal Society A **471**(2183) (2015) 20150268. 10.1098/rspa.2015.0268

XI. D. A. Kiefer, G. Watzl, K. Burgholzer, M. Ryzy, C. Grünsteidl, *Electroelastic guided wave dispersion in piezoelectric plates: spectral methods and laser-ultrasound experiments*, Journal of Applied Physics **137**(11) (2025) 114502. 10.1063/5.0250494

XII. P. Laugier, G. Haïat (eds.), *Bone Quantitative Ultrasound*, Springer (2010). 10.1007/978-94-007-0017-8

XIII. L. H. Le, V.-H. Nguyen, T. N. H. T. Tran, *Ultrasonic guided waves in bone: a decade of advancement in review*, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control **69**(9) (2022) 2540–2557.

XIV. J.-G. Minonzio, N. Bochud, Q. Vallet, Y. Bala, D. Ramiandrisoa, H. Follet, D. Mitton, P. Laugier, *Bone cortical thickness and porosity assessment using ultrasound guided waves: an ex vivo validation study*, Bone **116** (2018) 111–119. 10.1016/j.bone.2018.07.018

XV. H. Nguyen, D. Le, E. Plan, S. T. Dang, H. Phan, *Theoretical model of guided waves in a bone-mimicking plate coupled with soft-tissue layers*, Vietnam Journal of Mechanics **43**(1) (2021) 91–104. 10.15625/0866-7136/15774

XVI. V.-H. Nguyen, T. N. H. T. Tran, M. D. Sacchi, S. Naili, L. H. Le, *Computing dispersion curves of elastic/viscoelastic transversely-isotropic bone plates coupled with soft tissue and marrow using a semi-analytical finite element (SAFE) method*, Computers in Biology and Medicine **87** (2017) 371–381. 10.1016/j.compbiomed.2017.06.001

XVII. J. W. Nunziato, S. C. Cowin, *A nonlinear theory of elastic materials with voids*, Archive for Rational Mechanics and Analysis **72**(2) (1979) 175–201. 10.1007/BF00249363

XVIII. J. L. Rose, *Ultrasonic Guided Waves in Solid Media*, Cambridge University Press (2014).

XIX. L. N. Trefethen, *Spectral Methods in MATLAB*, SIAM, Philadelphia (2000). 10.1137/1.9780898719598

XX. J. A. C. Weideman, S. C. Reddy, *A MATLAB differentiation matrix suite*, ACM Transactions on Mathematical Software **26**(4) (2000) 465–519. 10.1145/365723.365727