Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

A Quadruple-Parabola Approximation for the Static Interfacial Properties of Asymmetric Phase-Separated Binary Bose--Einstein Condensates

Hoang Van Quyet

hoangvanquyet@hpu2.edu.vn

Hanoi Pedagogical University 2

Research Article

Keywords: Bose–Einstein condensation, Gross–Pitaevskii theory, Interfacial tension, Quadruple-parabola approximation, Wetting phenomena, Binary quantum fluids

Posted Date: September 29th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11029284/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

A Quadruple-Parabola Approximation for the Static Interfacial Properties of Asymmetric Phase-Separated Binary Bose–Einstein Condensates

Hoang Van Quyet[a,*]

[a]*Department of Physics, Hanoi Pedagogical University 2, Xuan Hoa, Phu Tho, Vietnam*

Abstract

We develop the Quadruple-Parabola Approximation (QPA), an analytical, piecewise-quadratic trial-function model for the static, planar, zero-temperature interface of a phase-separated binary Bose–Einstein condensate (BEC) mixture with unequal healing lengths, $\xi_1 \neq \xi_2$. The construction extends the Double- and Triple-Parabola Approximations (DPA, TPA) by assigning independent curvatures and matching positions to the two components while preserving exact linear solvability. The DPA, TPA, and QPA trial families are nested by construction, so that evaluating every trial profile in the exact Gross–Pitaevskii (GP) functional establishes the variational ordering $\sigma_{\mathrm{DPA}} \geq \sigma_{\mathrm{TPA}} \geq \sigma_{\mathrm{QPA}} \geq \sigma_{\mathrm{GP}}$ for arbitrary ξ; the QPA reproduces the TPA exactly in the symmetric limit $\xi = 1$ and obeys an exact species-exchange covariance, $\sigma(K, \xi) = \xi\, \sigma(K, 1/\xi)$, both verified numerically. Benchmarked against an independent GP boundary-value solver over $K \in \{1.2, 1.5, 2, 3, 5\}$ and $\xi \in \{0.5, 0.7, 1, 1/0.7, 2\}$, the QPA systematically lowers the variational interfacial tension relative to the TPA in every asymmetric case tested – removing up to $\sim 30\%$ of the TPA's residual tension error in the weakest-segregation, most-asymmetric regime, though typically only a few percent elsewhere – while the DPA over-estimates the tension by 9–42% against 6.0–9.6% for the TPA and QPA. This energy improvement does not translate into a smaller pointwise profile error: with a properly L-independent L^2 metric, the QPA profile error exceeds the TPA's in the majority of (species, K, ξ) combinations tested, but the TPA is not universally superior – in 8 of 40 such combinations, concentrated at $K \geq 2$ and related in mirrored pairs by the species-exchange covariance, the QPA is marginally better instead. Within the parameter window, optimisation protocol, and admissible domain explored here, neither the TPA nor the QPA reaches sub-percent tension accuracy. The results clarify both the usefulness and the intrinsic limitations of piecewise-quadratic interfacial models for binary quantum fluids: a larger variational trial family can be a genuinely better energetic approximation without being a better pointwise one.

Keywords: Bose–Einstein condensation, Gross–Pitaevskii theory, Interfacial tension, Quadruple-parabola approximation, Wetting phenomena, Binary quantum fluids

*Corresponding author.
Email address: hoangvanquyet@hpu2.edu.vn (Hoang Van Quyet)

1. Introduction

Multi-component Bose–Einstein condensates (BECs) have been studied intensively since the mean-field predictions of Ho and Shenoy [1] and the experimental realisation in alkali-metal atoms by Myatt *et al.* [2] and Hall *et al.* [3]. When the interspecies repulsion exceeds the geometric mean of the intraspecies couplings, $K \equiv g_{12}/\sqrt{g_{11}g_{22}} > 1$, the mixture phase separates into macroscopic domains bounded by a quantum interface [4, 5, 6, 7, 8, 9]. Exact and asymptotic treatments of this planar boundary in the weak- and strong-segregation limits were obtained early on by Ao and Chui [4] and by Barankov [10], and numerical solutions of the full coupled GP boundary-value problem are now standard [11, 12, 13, 14]. The interfacial tension so defined is not merely a bulk-thermodynamic quantity: it controls the wetting and prewetting behaviour of a condensate mixture at a hard or optical wall [15, 16, 17, 18], the sharp-interface limit of the segregated ground state [19], and the energetics of droplets and capillary excitations at a quantum interface [20, 21, 22, 23, 24, 25, 26]; percent-level changes in σ shift the location of a wetting phase boundary and the equilibrium shape of a droplet exactly as percent-level changes in a classical surface tension shift a contact angle [27]. Analytic control of $\sigma(K, \xi)$ over the full (K, ξ) plane is therefore useful well beyond curve-fitting the GP profile itself. We restrict ourselves throughout to the static, planar, zero-temperature interface; curved interfaces, trapped geometries, and finite temperature [28, 29] are natural extensions left for future work.

Direct numerical integration of the coupled GP equations gives formally exact profiles but must be recomputed for every parameter combination and offers little analytical insight; this motivates piecewise-quadratic trial-function models as a closed-form alternative that remains tied to the exact GP variational principle. The Double-Parabola Approximation (DPA) [17, 30] replaces the quartic GP potential with two shifted harmonic wells matched directly at the interface; it is fully analytic but produces a jump in the local curvature at the matching point and, as confirmed below, over-estimates the interfacial tension by 10–40% in the regime studied here. The Triple-Parabola Approximation (TPA) [31] inserts a single symmetric central region between the two bulk branches, removing the curvature jump and reducing the tension error to below 10% at $\xi_1 = \xi_2$. In the symmetric TPA construction adopted here, this central branch is constrained by a reflection symmetry that limits its ability to assign independent curvatures and matching positions to the two components once $\xi_1 \neq \xi_2$; we make no claim about piecewise-quadratic constructions in general, only about the specific symmetric TPA family used as our point of departure.

Existing analytic approximations of this kind are most valuable when interfacial energies must be evaluated repeatedly across broad parameter ranges, as in studies of wetting, droplet energetics, and phase-boundary problems, where a full GP solve at every grid point is impractical. The symmetric structure of the TPA restricts its flexibility precisely in this regime, since heteronuclear and highly asymmetric mixtures generically have $\xi_1 \neq \xi_2$. The present work asks a specific, falsifiable question: can the additional variational freedom of an asymmetric piecewise-quadratic construction lower the interfacial tension without sacrificing the analytic tractability of the DPA and TPA, and if so, does that energetic gain also improve the pointwise accuracy of the order-parameter profile? We propose the minimal such extension: the single symmetric central parabola of the TPA is replaced by two independent parabolic branches, one biased toward each component, joined at the origin $z_0 = 0$ fixed by translational gauge. Together with the two bulk half-spaces this gives four piecewise-quadratic domains – the Quadruple-Parabola Approximation (QPA) – governed by six physical parameters (two matching positions, two curvatures, two shifts) once twelve linear amplitude coefficients have been eliminated in closed form by continu-

ity matching. The construction is given in Section 3; by design it contains the TPA as the exact sub-case in which the parameters respect the $z \to -z, \phi_1 \leftrightarrow \phi_2$ symmetry available at $\xi = 1$ (Section 3.4). The answer to this two-part question is asymmetric: yes for the tension, not generally for the profile – a distinction that is the central methodological result of this study rather than a secondary observation about a new ansatz.

All quantitative results below are generated by an independently implemented numerical pipeline comprising a boundary-value-problem GP solver (`scipy.integrate.solve_bvp`), the closed-form QPA/TPA/DPA profiles of Section 3 evaluated in the exact GP functional, and a bounded, multi-start optimiser (Section 5.1).

This paper makes six specific contributions: (1) an asymmetric, exactly solvable piecewise-quadratic interfacial construction (the QPA) with six physical parameters, for arbitrary $\xi_1 \neq \xi_2$ (Section 3); (2) a nested variational hierarchy $F_{\mathrm{DPA}} \subset F_{\mathrm{TPA}}(\xi) \subset F_{\mathrm{QPA}}(\xi)$, giving $\sigma_{\mathrm{DPA}} \geq \sigma_{\mathrm{TPA}} \geq \sigma_{\mathrm{QPA}} \geq \sigma_{\mathrm{GP}}$ for every ξ (Section 4.3); (3) an exact and numerically verified reduction of the QPA to the TPA in the symmetric limit $\xi = 1$ (Section 3.4); (4) an exact species-exchange covariance relation, $\sigma(K, \xi) = \xi\, \sigma(K, 1/\xi)$, derived analytically and confirmed to $\sim 10^{-3}\%$ by independently seeded optimisers (Section 3.5); (5) a controlled benchmark against an independent GP boundary-value solver over a 5×5 grid in (K, ξ) (Section 5); and (6) a quantitative dissection of *how much* of the TPA's tension error the QPA removes, and a demonstration that this tension-level gain does not carry over to the pointwise L^2 profile error (Sections 5.3–5.7).

The paper is organised as follows. Section 2 recalls the GP formulation and derives the bulk offset entering the interfacial tension. Section 3 constructs the QPA, defines its admissible parameter domain, and establishes its reduction to the TPA at $\xi = 1$. Section 4 gives the closed-form solutions and the variational tension functional. Section 5 reports the benchmark against the GP equations. Section 6 concludes.

2. Gross–Pitaevskii theory and interfacial energy

We consider a phase-segregated binary BEC mixture at zero temperature. In the mean-field framework, the bulk energy density is

$$\mathcal{E}_b = \sum_{j=1}^{2} \left[\frac{\hbar^2}{2m_j} |\nabla \psi_j|^2 - \mu_j |\psi_j|^2 + \frac{g_{jj}}{2} |\psi_j|^4 \right] + g_{12} |\psi_1|^2 |\psi_2|^2, \tag{1}$$

where m_j, μ_j are the atomic mass and chemical potential of species j, and g_{jj}, g_{12} are the intra- and interspecies coupling constants [4, 30]. For a flat interface perpendicular to z, with real $\psi_j = \psi_j(z)$, lengths scaled by $\xi_1 = \hbar/\sqrt{2m_1\mu_1}$ and dimensionless order parameters $\phi_j = \psi_j/\sqrt{n_{j0}}$, $n_{j0} = \mu_j/g_{jj}$, the reduced potential is

$$V(\phi_1, \phi_2) = -\frac{\phi_1^2}{2} + \frac{\phi_1^4}{4} - \frac{\phi_2^2}{2} + \frac{\phi_2^4}{4} + \frac{K}{2} \phi_1^2 \phi_2^2, \tag{2}$$

with $K = g_{12}/\sqrt{g_{11} g_{22}} > 1$. The coupled GP equations are

$$-\frac{\mathrm{d}^2 \phi_1}{\mathrm{d}z^2} - \phi_1 + \phi_1^3 + K \phi_2^2 \phi_1 = 0, \tag{3}$$

$$-\xi^2 \frac{\mathrm{d}^2 \phi_2}{\mathrm{d}z^2} - \phi_2 + \phi_2^3 + K \phi_1^2 \phi_2 = 0, \tag{4}$$

with $\xi = \xi_2/\xi_1 = \sqrt{m_1\mu_1/(m_2\mu_2)}$ and boundary conditions $\phi_1(-\infty) = 0$, $\phi_1(+\infty) = 1$, $\phi_2(-\infty) = 1$, $\phi_2(+\infty) = 0$. These follow from the Lagrangian density

$$\mathcal{L} = \frac{1}{2}\phi_1'^2 + \frac{\xi^2}{2}\phi_2'^2 - V(\phi_1, \phi_2). \tag{5}$$

Bulk offset. The reduced interfacial tension is $\sigma = \int[\mathcal{L}_{\rm kin} + V - P_0]\,\mathrm{d}z$, with P_0 the bulk value of V chosen so the integrand vanishes as $z \to \pm\infty$. Evaluating Eq. (2) at either bulk minimum, $(\phi_1, \phi_2) = (1, 0)$ or $(0, 1)$, gives $V_{\rm bulk} = -1/4$ for any K, so

$$P_0 = -\frac{1}{4}, \qquad \sigma = \int_{-\infty}^{+\infty}\left[\frac{1}{2}\phi_1'^2 + \frac{\xi^2}{2}\phi_2'^2 + V(\phi_1, \phi_2) + \frac{1}{4}\right]\mathrm{d}z. \tag{6}$$

This constant is used consistently throughout.

3. The Quadruple-Parabola Approximation

3.1. Construction

The DPA and TPA replace V by piecewise-quadratic sheets, rendering the Euler–Lagrange equations linear in each domain. In the TPA family adopted here, the two share a single, symmetric central region whose curvature and matching position are tied together by construction and so cannot be assigned independently to the two components when $\xi_1 \neq \xi_2$. The QPA replaces this single central parabola by two independent parabolic sheets, one biased toward each species, giving four domains (Table 1):

$$V_{\rm QPA} = \begin{cases} V_{\rm I}, & z \geq z_1, \\ V_{\rm II}, & 0 \leq z < z_1, \\ V_{\rm III}, & z_2 \leq z < 0, \\ V_{\rm IV}, & z < z_2, \end{cases} \qquad z_2 < 0 < z_1, \tag{7}$$

$$V_{\rm I} = (\phi_1 - 1)^2 + \lambda_1\,\phi_2^2, \qquad V_{\rm II} = \gamma_1(\phi_1 - \beta_1)^2 + \lambda_2\,\phi_2^2, \tag{8}$$

$$V_{\rm III} = \lambda_3\,\phi_1^2 + \gamma_2(\phi_2 - \beta_2)^2, \qquad V_{\rm IV} = \lambda_4\,\phi_1^2 + (\phi_2 - 1)^2. \tag{9}$$

Table 1: The four QPA domains ($z_0 = 0$ fixed by translational gauge). In each domain one field carries a free "majority" curvature and the other the exact bulk-quadratic "minority" curvature of Eqs. (12).

Dom.	Interval	Majority field	Free params
I	$z \geq z_1$	$\phi_1 \to 1$ (fixed)	–
II	$0 \leq z < z_1$	ϕ_1, curv. γ_1, shift β_1	γ_1, β_1, z_1
III	$z_2 \leq z < 0$	ϕ_2, curv. γ_2, shift β_2	γ_2, β_2, z_2
IV	$z < z_2$	$\phi_2 \to 1$ (fixed)	–

3.2. Calibration of the bulk curvatures

Matching the Taylor expansion of the exact potential $V(\phi_1, \phi_2)$, Eq. (2), about the two bulk minima fixes the curvatures $\lambda_1, \ldots, \lambda_4$. Since $V(\phi_1, \phi_2)$ contains no explicit ξ-dependence (only the kinetic term $\frac{\xi^2}{2}\phi_2'^2$ in Eq. (5) does), each λ_i is fixed purely by the field-space expansion of V and is independent of ξ. At the right bulk $(\phi_1, \phi_2) = (1, 0)$,

$$V + \frac{1}{4} = (\phi_1 - 1)^2 + \frac{K-1}{2}\phi_2^2 + O(\delta\phi^3), \tag{10}$$

fixing $\lambda_1 = (K-1)/2$; continuity of the ϕ_2 curvature across z_0 requires the same value for λ_2. At the left bulk $(\phi_1, \phi_2) = (0, 1)$, the $\phi_1 \leftrightarrow \phi_2$ symmetry of V gives

$$V + \frac{1}{4} = \frac{K-1}{2}\phi_1^2 + (\phi_2 - 1)^2 + O(\delta\phi^3), \tag{11}$$

fixing $\lambda_4 = (K-1)/2$ and, by the same argument, $\lambda_3 = (K-1)/2$. Hence

$$\lambda_1 = \lambda_2 = \lambda_3 = \lambda_4 = \frac{K-1}{2}, \tag{12}$$

with no factor of ξ^2. The healing-length ratio ξ enters the problem only through the kinetic term $\frac{\xi^2}{2}\phi_2'^2$, and hence only through the decay rate with which ϕ_2 approaches its fixed point in each domain, never through the curvature of V itself. In Domains III and IV, where ϕ_1 is the minority field, its equation is $\phi_1'' = 2\lambda_3\phi_1 = (K-1)\phi_1$, with decay rate $\kappa_0 = \sqrt{K-1}$ (no ξ); in Domains I and II, where ϕ_2 is the minority field, its equation is $\xi^2\phi_2'' = 2\lambda_2\phi_2 = (K-1)\phi_2$, i.e. $\phi_2'' = (K-1)\phi_2/\xi^2$, with decay rate $\kappa_\xi = \kappa_0/\xi$. Section 4.1 makes this bookkeeping explicit for all four domains. Because $\xi = 1$ trivially satisfies $\kappa_0 = \kappa_\xi$, a symmetric-only benchmark cannot validate the ξ-dependence of the construction; the species-exchange covariance of Section 3.5, verified to $\sim 10^{-3}\%$ in Table 5, provides the required $\xi \neq 1$ check.

3.3. Admissible domain

The physical parameter vector is $\boldsymbol{\lambda} = (z_1, z_2, \gamma_1, \gamma_2, \beta_1, \beta_2)$ with

$$z_2 < 0 < z_1, \quad \gamma_1, \gamma_2 > 0, \quad 0 \le \beta_1, \beta_2 \le 1, \tag{13}$$

the last condition reflecting that $\beta_{1,2}$ interpolate between full depletion ($\beta = 0$) and the bulk value ($\beta = 1$) of the majority field. Equation (13) is enforced directly as box constraints, $z_1 \in [0.02, 6]$, $z_2 \in [-6, -0.02]$, $\gamma_{1,2} \in [0.05, 40]$, $\beta_{1,2} \in [0.005, 0.995]$, in the bounded optimiser of Section 5.1.

3.4. Symmetric limit: the QPA ansatz contains the TPA exactly

At $\xi = 1$, Eqs. (2)–(4) are invariant under $z \to -z$, $\phi_1 \leftrightarrow \phi_2$. The QPA ansatz respects this symmetry, and Domains II and III become mirror images of one another, precisely when

$$z_1 = -z_2 \equiv z_c, \qquad \gamma_1 = \gamma_2 \equiv \gamma, \qquad \beta_1 = \beta_2 \equiv \beta. \tag{14}$$

This three-parameter subspace, (z_c, γ, β), is identified with the TPA. Two distinct statements follow: (i) the six-parameter QPA family contains the TPA exactly as this reflection-symmetric subfamily, by construction; (ii) the unconstrained global optimum of the QPA functional lies on this subfamily at $\xi = 1$. Statement (ii) does not follow automatically from the symmetry of the functional – a symmetric functional can in general admit symmetry-broken or degenerate optima – and is established here as a numerical result. We therefore optimise the TPA and QPA independently (Section 5.1); the unconstrained six-parameter QPA optimiser recovers $z_1 = -z_2$, $\gamma_1 = \gamma_2$, and $\beta_1 = \beta_2$ to four significant figures at every $\xi = 1$ point studied, with σ_{QPA} and σ_{TPA} agreeing to better than 10^{-6} (relative).

3.5. Species-exchange covariance

A stronger consistency test, applicable for all ξ, follows from the observation that relabelling the two components maps a solution at (K, ξ) onto a solution at $(K, 1/\xi)$, since the labelling "component 1" fixes only the length unit, not the physics. Because $V(\phi_1, \phi_2)$ in Eq. (2) is exactly symmetric under $\phi_1 \leftrightarrow \phi_2$, if $(\phi_1(z), \phi_2(z))$ solves Eqs. (3)–(4) at (K, ξ), then

$$\phi_1'(z') = \phi_2(\xi z'), \qquad \phi_2'(z') = \phi_1(\xi z'), \qquad z' = z/\xi, \tag{15}$$

solves the same equations at $(K, 1/\xi)$: substituting Eq. (15) into Eqs. (3)–(4) and using $\mathrm{d}/\mathrm{d}z = \xi^{-1}\mathrm{d}/\mathrm{d}z'$ maps the ϕ_1 equation at (K, ξ) onto the ϕ_2' equation at $(K, 1/\xi)$ and vice versa, term by term. Substituting into the tension functional Eq. (6) and changing variables from z to z' gives, after using the $\phi_1 \leftrightarrow \phi_2$ symmetry of V,

$$\boxed{\sigma(K, \xi) = \xi\, \sigma(K, 1/\xi)}. \tag{16}$$

This relation is exact for the true GP minimiser, using only the symmetry of V and a linear change of variables. If a trial-function family is constructed with the corresponding reflection covariance – as the DPA, TPA, and QPA are here, via the explicit $1 \leftrightarrow 2$, I $\leftrightarrow$ IV, II $\leftrightarrow$ III mapping of Table 1 – it must hold for the variational optimum of each family provided the optimiser correctly finds that optimum independently at ξ and at $1/\xi$. Our optimisers for ξ and $1/\xi$ are run independently (different random seeds, no shared state), so Eq. (16) constitutes a genuine numerical test, verified in Section 5.6. An immediate corollary is that the relative tension error is invariant under $\xi \to 1/\xi$,

$$\varepsilon_\sigma(K, \xi) = \frac{\sigma_{\mathrm{model}}(K, \xi) - \sigma_{\mathrm{GP}}(K, \xi)}{\sigma_{\mathrm{GP}}(K, \xi)} = \varepsilon_\sigma(K, 1/\xi), \tag{17}$$

which we exploit in Section 5.7 to halve the cost of the full (K, ξ) scan.

4. Analytical solutions and interfacial tension

4.1. Linear field equations and general solutions

With $\alpha_0 = \sqrt{2}$, $\kappa_0 = \sqrt{K-1}$, $\nu_1 = \sqrt{2\gamma_1}$, $\kappa_\xi = \kappa_0/\xi$, $\alpha_\xi = \alpha_0/\xi$, and $\nu_2 = \sqrt{2\gamma_2}/\xi$, the Euler–Lagrange equations of Eqs. (8)–(9) are linear and second-order in each domain. Every rate constant associated with ϕ_2 carries a compensating $1/\xi$ $(\kappa_\xi, \alpha_\xi, \nu_2)$, while rate constants associated with ϕ_1 never do $(\alpha_0, \kappa_0, \nu_1)$, regardless of domain or of whether ϕ_1 or ϕ_2 plays the majority role. The general solutions are

$$\text{I } (z \geq z_1): \quad \phi_1 = 1 - A_1 e^{-\alpha_0 (z - z_1)}, \;\; \phi_2 = B_1 e^{-\kappa_\xi (z - z_1)}, \tag{18}$$

$$\text{II } (0 \leq z < z_1): \quad \phi_1 = \beta_1 + A_2 \cosh(\nu_1 z) + B_2 \sinh(\nu_1 z),$$
$$\phi_2 = C_2 \cosh(\kappa_\xi z) + D_2 \sinh(\kappa_\xi z), \tag{19}$$

$$\text{III } (z_2 \leq z < 0): \quad \phi_1 = A_3 \cosh(\kappa_0 z) + B_3 \sinh(\kappa_0 z),$$
$$\phi_2 = \beta_2 + C_3 \cosh(\nu_2 z) + D_3 \sinh(\nu_2 z), \tag{20}$$

$$\text{IV } (z < z_2): \quad \phi_1 = A_4 e^{\kappa_0 (z - z_2)}, \;\; \phi_2 = 1 - B_4 e^{\alpha_\xi (z - z_2)}. \tag{21}$$

Note the asymmetry between Eqs. (18)–(19) and Eqs. (20)–(21): ϕ_1 decays with the bare rate κ_0 in Domains III/IV, while ϕ_2's frequency in Domain III is $\nu_2 = \sqrt{2\gamma_2}/\xi$, for the reason given in Section 3.2.

4.2. Species-decoupled matching system

Continuity of ϕ_1, ϕ_2 and their first derivatives at $z_1, 0, z_2$ gives 12 equations for 12 amplitudes. Because Eqs. (8)–(9) contain no $\phi_1\phi_2$ cross-term, the field equations never mix ϕ_1 and ϕ_2, and the system splits exactly into two independent 6×6 (species-block-diagonal) blocks, $\mathbf{M}^{(1)}\mathbf{x}^{(1)} = \mathbf{b}^{(1)}$ for $\mathbf{x}^{(1)} = (A_1, A_2, B_2, A_3, B_3, A_4)$ and $\mathbf{M}^{(2)}\mathbf{x}^{(2)} = \mathbf{b}^{(2)}$ for $\mathbf{x}^{(2)} = (B_1, C_2, D_2, C_3, D_3, B_4)$, each block sparse but coupling the amplitudes of neighbouring domains. The explicit matrices are given in Appendix A; both blocks are invertible for every admissible parameter set tested, and the amplitudes follow from two direct linear solves, $\mathbf{x}^{(i)} = (\mathbf{M}^{(i)})^{-1}\mathbf{b}^{(i)}$.

What is, and is not, continuous at the matching points. The twelve conditions above enforce continuity of ϕ_1, ϕ_2 and of ϕ_1', ϕ_2' at $z_1, 0, z_2$, but not of the second derivatives ϕ_1'', ϕ_2''. This is a deliberate, structural feature rather than an oversight: because the trial potential V_{QPA} is piecewise-quadratic with a different curvature in each domain, $\phi_j'' \propto \partial V_{\mathrm{QPA}}/\partial\phi_j$ is generically discontinuous wherever the curvature changes, i.e. at every one of the three matching points. Demanding C^2 matching in addition to C^1 matching would over-determine the linear system of Section 4.2, which already uses all available continuity conditions to fix the twelve amplitudes; the residual curvature jump is the mechanism by which a piecewise-quadratic potential can approximate the smooth quartic potential V of Eq. (2) using a finite number of parameters, exactly as in the DPA and TPA. Because σ is evaluated with the exact GP potential rather than the piecewise surrogate (Section 4.3), this curvature discontinuity does not by itself bias the reported tension; it does, however, leave a nonzero pointwise residual of the true GP equation,

$$R_1(z) = -\phi_1'' + \frac{\partial V}{\partial \phi_1}, \qquad R_2(z) = -\xi^2\phi_2'' + \frac{\partial V}{\partial \phi_2}, \tag{22}$$

evaluated on the QPA (or TPA, DPA) trial profile, which is exactly zero only for the true GP minimiser. In this sense the QPA is not literally "smoother" than the TPA or DPA at the matching points – all three retain a curvature jump by construction – but it has more matching points at which to distribute that jump, and hence more freedom to keep $R_j(z)$ small over the bulk of the interface where $|\phi_1'\phi_2|$-weighted contributions to σ are largest. $R_j(z)$ is not computed quantitatively here; a residual-based diagnostic of this kind is a natural extension of the profile-error analysis of Section 5.3.

4.3. Interfacial tension: a genuine variational bound

Functional evaluated. Every interfacial tension reported in this paper – σ_{DPA}, σ_{TPA}, σ_{QPA} – is obtained by substituting the closed-form trial profiles (18)–(21) into

$$\sigma[\phi_1, \phi_2] = \int_{-\infty}^{+\infty} \left[\frac{1}{2}\phi_1'^2 + \frac{\xi^2}{2}\phi_2'^2 + V(\phi_1, \phi_2) + \frac{1}{4}\right] \mathrm{d}z, \tag{23}$$

i.e. Eq. (6) with the true GP potential V of Eq. (2), never the piecewise surrogate V_{QPA} of Eq. (7). This is what makes the variational principle applicable: for any admissible trial pair (ϕ_1, ϕ_2) obeying the boundary conditions of Eqs. (3)–(4), $\sigma[\phi_1, \phi_2] \geq \sigma_{\mathrm{GP}}$, with equality only for the true GP minimiser.

Ordering for every ξ. Define the three trial-function families as subsets of the six-dimensional QPA parameter space $\boldsymbol{\lambda} = (z_1, z_2, \gamma_1, \gamma_2, \beta_1, \beta_2)$, for fixed but arbitrary ξ:

$$F_{\mathrm{QPA}}(\xi) = \{\boldsymbol{\lambda} : \boldsymbol{\lambda} \text{ admissible, Eq. (13)}\}, \tag{24}$$

$$F_{\mathrm{TPA}}(\xi) = \{\boldsymbol{\lambda} \in F_{\mathrm{QPA}}(\xi) : z_1 = -z_2,\ \gamma_1 = \gamma_2,\ \beta_1 = \beta_2\}, \tag{25}$$

$$F_{\mathrm{DPA}} = \{\boldsymbol{\lambda} \in \overline{F_{\mathrm{TPA}}(\xi)} : z_1 = -z_2 \to 0\}. \tag{26}$$

The constraints defining $F_{\mathrm{TPA}}(\xi)$ do not reference ξ and are equally well posed whether or not $\xi = 1$, so $F_{\mathrm{TPA}}(\xi) \subset F_{\mathrm{QPA}}(\xi)$ and $F_{\mathrm{DPA}} \subset \overline{F_{\mathrm{TPA}}(\xi)}$ for every ξ (the closure accounting for the zero-width limit at which $\gamma_{1,2}, \beta_{1,2}$ become irrelevant, reproducing the parameter-free DPA profile).

Proposition 1 (Variational ordering)**.** *Let $F_{\mathrm{DPA}} \subset F_{\mathrm{TPA}}(\xi) \subset F_{\mathrm{QPA}}(\xi)$ be the admissible trial-profile families defined above, and let every trial profile be evaluated in the exact GP functional $\sigma[\phi_1, \phi_2]$ of Eq.* (23)*. If each optimisation attains the global minimum of its family, then for every ξ,*

$$\sigma_{\mathrm{DPA}} \ \geq\ \sigma_{\mathrm{TPA}}(\xi) \ \geq\ \sigma_{\mathrm{QPA}}(\xi) \ \geq\ \sigma_{\mathrm{GP}}(\xi). \tag{27}$$

Proof. Minimising a fixed functional over a larger set cannot give a larger minimum than minimising over a subset of that set: for $A \subset B$, $\min_B f \leq \min_A f$. Applying this twice to $F_{\mathrm{DPA}} \subset F_{\mathrm{TPA}}(\xi) \subset F_{\mathrm{QPA}}(\xi)$ gives $\min_{F_{\mathrm{QPA}}(\xi)} \sigma \leq \min_{F_{\mathrm{TPA}}(\xi)} \sigma \leq \min_{F_{\mathrm{DPA}}} \sigma$. Because $\sigma[\phi_1, \phi_2]$ is the exact GP functional evaluated on trial pairs satisfying the correct boundary conditions, $\sigma[\phi_1, \phi_2] \geq \sigma_{\mathrm{GP}}(\xi)$ for every admissible trial pair, with equality only at the true GP minimiser; taking the infimum over $F_{\mathrm{QPA}}(\xi)$ gives $\min_{F_{\mathrm{QPA}}(\xi)} \sigma \geq \sigma_{\mathrm{GP}}(\xi)$. Chaining the two inequalities gives Eq. (27). □

The argument uses only the set inclusion $F_{\mathrm{DPA}} \subset F_{\mathrm{TPA}}(\xi) \subset F_{\mathrm{QPA}}(\xi)$ and the fact that σ is evaluated with the exact GP potential rather than a piecewise surrogate; it is independent of the reflection symmetry available only at $\xi = 1$. The global-minimum proviso in Proposition 1 is addressed empirically via the multi-start protocol of Section 5.1, and the concrete parameter values at which F_{QPA} reduces to F_{TPA} and F_{TPA} reduces to F_{DPA} are given explicitly by Eq. (14) and by $z_1 = -z_2 \to 0$, respectively; both reductions preserve the continuity matching conditions of Section 4.2 in the corresponding limit.

Reflection symmetry instead yields a logically separate, strictly stronger statement at $\xi = 1$: not merely $\sigma_{\mathrm{QPA}}(1) \leq \sigma_{\mathrm{TPA}}(1)$, already implied by Proposition 1, but the equality $\sigma_{\mathrm{QPA}}(1) = \sigma_{\mathrm{TPA}}(1)$, because the unconstrained global optimum of the six-parameter family lies exactly on the three-parameter subfamily $F_{\mathrm{TPA}}(1)$. This is verified numerically in Section 5.2.

5. Results

Throughout this section $\xi \equiv \xi_2/\xi_1$; lengths z are measured in units of ξ_1; σ denotes the reduced interfacial tension of Eq. (6), with $P_0 = -1/4$, so that $\sigma = \gamma/(|P_0|\xi_1)$ in terms of the physical tension γ. GP and model profiles are compared on the same z/ξ_1 grid, with $z_0 = 0$ fixed by the gauge choice of Section 3.3.

5.1. Numerical pipeline

GP benchmark. Equations (3)–(4) are solved as a boundary-value problem with `scipy.integrate.solve_bvp` [32], using fourth-order collocation with residual-based adaptive mesh refinement (relative/absolute tolerance 10^{-9}, up to 3×10^5 nodes, with automatic retry on a

larger box and coarser tolerance – up to 2×10^6 nodes – for the most asymmetric cases) on $[-L, L]$ with $L = 10$–27 chosen adaptively so profiles reach their bulk values to within 10^{-6}; σ_{GP} is obtained by Simpson integration of Eq. (6) on a 6000-point grid. The initial mesh uses 2000 equally spaced points on $[-L, L]$, seeded with the analytic DPA profile of Section 3 (a tanh-like guess) as the initial guess for ϕ_1, ϕ_2, so that the collocation solver starts already close to the correct monotonic interface shape; boundary conditions are imposed as Dirichlet values $\phi_1(-L) = \phi_2(L) = 0$, $\phi_1(L) = \phi_2(-L) = 1$ rather than as asymptotic matching conditions. Convergence is declared when `solve_bvp` reports `status=0` (residual below tolerance on the final mesh); on failure or non-convergence, the domain is enlarged by 50% and, if that also fails, the tolerance is coarsened as described above. No case in the reported (K, ξ) grid required more than one retry.

DPA. Fully analytic (Section 3 with Domains II, III absent, i.e. Domains I and IV matched directly at $z = 0$); no optimisation required.

TPA. The three-parameter symmetric sub-case of Eq. (14), optimised independently.

QPA. The full six-parameter model, optimised independently.

Both the TPA and QPA optimisations minimise σ_{trial}, Eq. (6), using bounded L-BFGS-B (box constraints of Section 3.3) with 11 multi-start initial points per case (three fixed seeds plus eight randomised seeds), keeping the lowest-tension optimum found; this protocol is necessary because an unconstrained, single-seed search can stagnate or diverge for strongly asymmetric ξ. The (K, ξ) grid scan of Section 5.7 uses a lighter, four-seed variant to keep the computation tractable; two spot-checked grid points agreed with the full protocol to better than 10^{-9} (relative). Table 2 summarises the pipeline settings.

Across the 17 independently optimised (K, ξ) cases underlying Table 3 and the grid of Section 5.7, only one case ($K = 1.5, \xi = 0.5$) has a parameter ($\beta_2 = 0.995$) within 0.1% of its bound; every other case converges to an interior point ($\gamma_{1,2} \leq 30.8$, $z_1 \leq 2.7$, $z_2 \geq -4.4$, against bounds of 40, 6, −6). The boundary-touching case shows no anomalous tension error relative to its neighbours ($\varepsilon_\sigma^{\mathrm{QPA}} = 6.11\%$).

Table 2: Reproducibility summary for the numerical pipeline of Section 5.1.

Quantity	Value
GP domain $[-L, L]$	$L = 10$–27, adaptive
GP solver tolerance	rtol/atol = 10^{-9} (coarsened to 10^{-6} on retry)
GP boundary residual	profiles reach bulk values to 10^{-6}
GP mesh budget	up to 3×10^5 nodes (2×10^6 on retry)
TPA/QPA optimiser	bounded L-BFGS-B (`scipy.optimize.minimize`)
Optimiser tolerance	`ftol`= 10^{-13}, `gtol`= 10^{-10}
Starts per case (Table 3)	11 (3 fixed + 8 random)
Starts per case (grid scan, Section 5.7)	4 (3 fixed + 1 random)
Parameter bounds	Eq. (13)
Boundary-touching optima	1 of 17
Light- vs. full-protocol agreement	$< 10^{-9}$ relative (2-point spot check)

5.2. Symmetric regime ($\xi = 1$): QPA ≡ TPA

Table 3 (upper block) and Figs. 1 and 2 summarise the symmetric case. The DPA overestimates σ_{GP} by 9–40%, worst near the miscibility threshold $K \to 1^+$, where the interface is

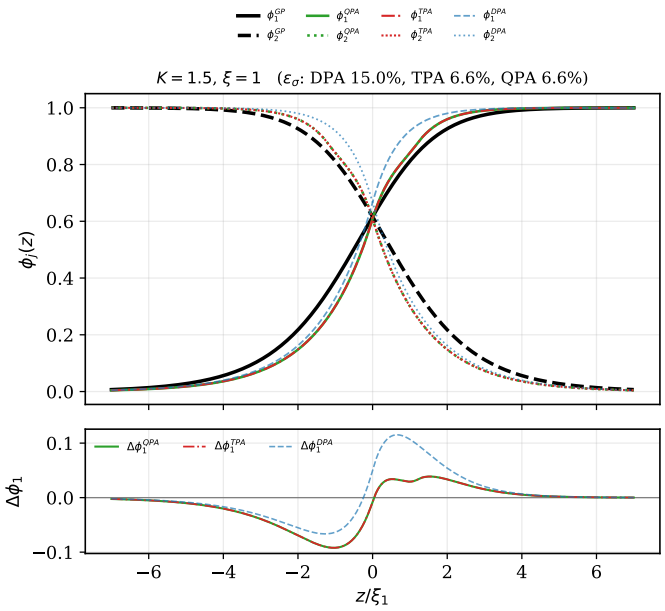


Figure 1: Order-parameter profiles at $K = 1.5$, $\xi = 1$: GP benchmark (black), QPA (green), TPA (red, dash-dot), DPA (blue, dashed). Lower panel: residuals $\Delta\phi_1 = \phi_1^{\text{model}} - \phi_1^{\text{GP}}$. The QPA and TPA curves are, as proved in Section 3.4, numerically indistinguishable ($\sigma_{\text{QPA}} = \sigma_{\text{TPA}}$ to six digits); relative tension errors are DPA 15.0%, TPA/QPA 6.6%.

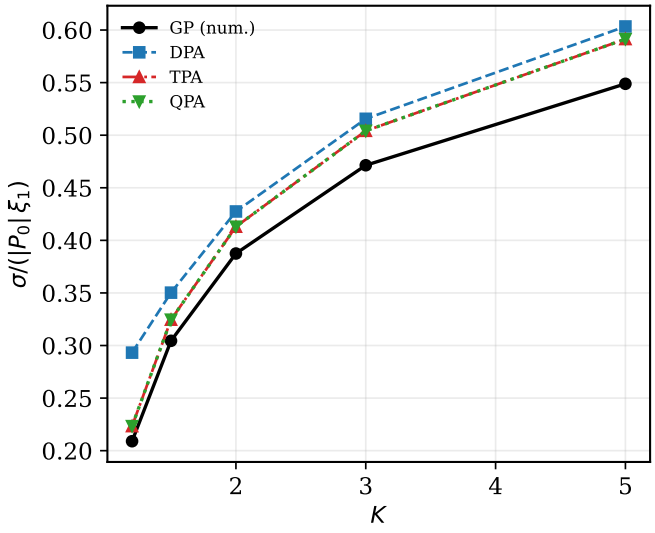


Figure 2: Reduced interfacial tension $\sigma/(|P_0|\xi_1)$ vs. K at $\xi = 1$ (Table 3, upper block). GP: black circles; DPA: blue squares; TPA: red triangles; QPA: green inverted triangles. The TPA and QPA curves overlap exactly, as required by Section 3.4.

wide and a single bulk-matched parabola per side is least appropriate. The TPA reduces this to 7–8% across $K \in [1.2, 5]$, and the QPA gives identical results at $\xi = 1$, confirming both the analytic reduction of Section 3.4 and the consistency of the numerical pipeline.

5.3. Asymmetric regime ($K = 2$, $\xi \in [0.5, 2]$)

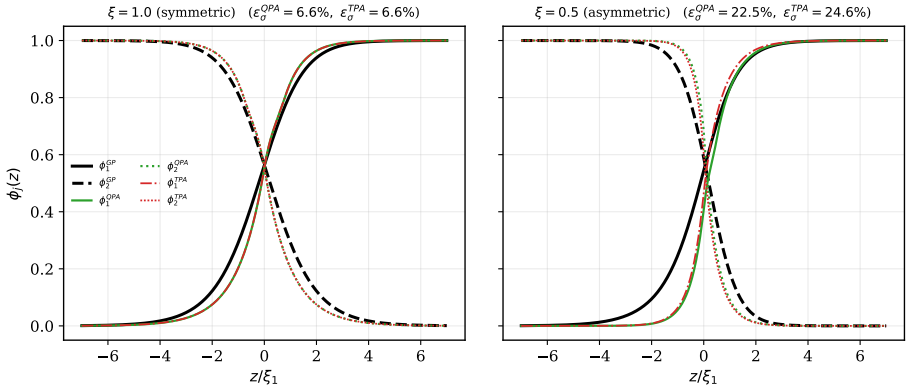


Figure 3: Order-parameter profiles at $K = 2$ for $\xi = 1$ (left, symmetric) and $\xi = 0.5$ (right, asymmetric). The QPA (green) assigns independent curvatures to the two components and attains a lower tension error than the TPA (red) in both panels (Table 3), but the pointwise L^2 profile error is slightly larger for the QPA than for the TPA in the asymmetric panel.

Table 3: Interfacial tension and relative tension error, $\varepsilon_\sigma = 100 \times (\sigma_{\text{model}} - \sigma_{\text{GP}})/\sigma_{\text{GP}}$, for the nine representative points from the two one-dimensional cuts (K at fixed $\xi = 1$; ξ at fixed $K = 2$) of Sections 5.2–5.3 ($P_0 = -1/4$ throughout). These nine points are a subset of the full $5\times5 = 25$-point (K, ξ) grid scan of Table 5 and Fig. 7 (Section 5.7). $\varepsilon_\sigma^{\text{TPA}} = \varepsilon_\sigma^{\text{QPA}}$ at $\xi = 1$ by construction (Section 3.4). The $\xi = 1.4$ row is the literal, directly optimised point $\xi = 1.4$ (not the exact reciprocal $1/0.7 \approx 1.4286$ used for the covariance grid of Table 5 and Fig. 7); σ_{GP} therefore differs slightly between the two (0.4684 at $\xi = 1.4$ vs. 0.4744 at $\xi = 1/0.7$, from $\sigma(2, 0.7)/0.7$), while the relative errors ε_σ agree with the covariance-derived $1/0.7$ values to within 0.02 percentage points, consistent with the spot check reported in the caption of Fig. 7.

K	ξ	σ_{GP}	$\varepsilon_\sigma^{\text{DPA}}$	$\varepsilon_\sigma^{\text{TPA}}$	$\varepsilon_\sigma^{\text{QPA}}$
1.2	1.0	0.2090	40.3%	6.9%	6.9%
1.5	1.0	0.3045	15.0%	6.6%	6.6%
2.0	1.0	0.3875	10.3%	6.6%	6.6%
3.0	1.0	0.4714	9.4%	7.0%	7.0%
5.0	1.0	0.5489	9.9%	7.8%	7.8%
2.0	0.5	0.2993	11.0%	6.5%	6.0%
2.0	0.7	0.3321	10.5%	6.6%	6.4%
2.0	1.4	0.4684	10.5%	6.6%	6.4%
2.0	2.0	0.5985	11.0%	6.5%	6.0%

Table 3 (lower block), Fig. 3, and Fig. 4 give the asymmetric comparison. The ordering $\sigma_{\text{DPA}} \geq \sigma_{\text{TPA}} \geq \sigma_{\text{QPA}} \geq \sigma_{\text{GP}}$ holds in every case, consistent with Section 4.3. The DPA, TPA, and QPA errors are comparatively uniform across $\xi \in [0.5, 2]$ at fixed $K = 2$: DPA remains near 11%, and TPA/QPA remain in the narrow band 6.0–6.6%, essentially unchanged from their $\xi = 1$

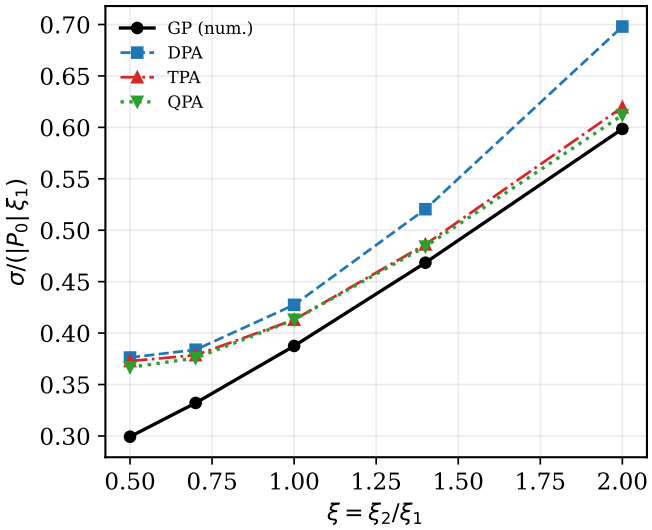


Figure 4: Reduced interfacial tension vs. $\xi = \xi_2/\xi_1$ at $K = 2$ (Table 3, lower block). All three approximations track the GP growth of σ with ξ reasonably uniformly, with no strong asymmetry between $\xi < 1$ and $\xi > 1$.

values. The QPA's advantage over the TPA is real but small and essentially ξ-independent in this range: 0.1–0.5 percentage points in absolute terms (equivalently 3–8% relative to the TPA's own error). This is consistent with the ξ-bookkeeping of Section 3.2, under which no domain of the model has a built-in reason to behave qualitatively differently between $\xi < 1$ and $\xi > 1$.

Two limitations follow from this benchmark. First, the TPA/QPA errors do not fall below the $\sim$ 6–8% level across the entire (K, ξ) range studied; whether sub-percent accuracy is achievable with a different optimiser, a finer parameterisation, or lies outside the reach of a four-domain piecewise-quadratic ansatz remains open. Second, the QPA's advantage over the TPA is concentrated in the integrated tension rather than the pointwise profile.

How much of the TPA's error does the QPA remove?. A percentage-point difference in ε_σ is not, by itself, informative about how large the QPA's gain is relative to what remained to be corrected. We therefore define the improvement fraction

$$G_\sigma(K, \xi) = 100 \times \frac{\sigma_{\mathrm{TPA}} - \sigma_{\mathrm{QPA}}}{\sigma_{\mathrm{TPA}} - \sigma_{\mathrm{GP}}} = 100 \times \frac{\varepsilon_\sigma^{\mathrm{TPA}} - \varepsilon_\sigma^{\mathrm{QPA}}}{\varepsilon_\sigma^{\mathrm{TPA}}}, \tag{28}$$

the second equality following directly from the definition of ε_σ in Table 3. At fixed $K = 2$, G_σ ranges from 2.7% ($\xi = 0.7$) to 7.7% ($\xi = 0.5$ or 2.0): the QPA removes less than a tenth of the TPA's residual tension error in this cut. Evaluated over the full 5×5 grid of Section 5.7, G_σ is far from uniform: it is largest, 31%, at the weakest-segregation, most-asymmetric corner ($K = 1.2, \xi = 0.5$ or 2.0), and falls to a few percent or less elsewhere, including essentially 0% at $\xi = 1$ by construction. G_σ therefore isolates the regime in which the QPA's additional parameters are doing genuine work – weak segregation combined with strong healing-length asymmetry – from the much larger part of the (K, ξ) plane in which the QPA is a small, essentially cosmetic refinement of the TPA.

Table 4 reports the per-species L^2 profile errors, $\varepsilon_{\phi,j} = \|\phi_j^{\text{model}} - \phi_j^{\text{GP}}\|_2$, separately for $j = 1, 2$, at $K = 2$. The TPA gives a better pointwise fit than the QPA for the majority of individual (species, ξ) combinations, but not all: at $\xi = 0.7$ the QPA is marginally *better* for ϕ_1, and at $\xi = 1/0.7$ – its covariance partner – the QPA is marginally better for ϕ_2. This is not an isolated artefact of $K = 2$. Repeating the computation over the full 5×5 grid of Section 5.7 (Fig. 8) shows the QPA matching or beating the TPA in 8 of the 40 (species, K, ξ) combinations at $\xi \neq 1$, and every one of these reversals occurs at $K \geq 2$ (never at $K = 1.2$ or 1.5) and respects the species-exchange covariance exactly: a QPA advantage for ϕ_1 at some $\xi_0 < 1$ is always mirrored by a QPA advantage for ϕ_2 at $\xi = 1/\xi_0$, and never the other way around. This rules out a clean, universal statement that the QPA is pointwise worse in every case, but it does not weaken the central point: the reversals are a modest, structured, strong-segregation phenomenon confined to a minority of cases, while the QPA's tension advantage G_σ is largest precisely in the opposite corner of parameter space – weak segregation, $K \to 1^+$ (Section 5.3) – so the regime in which the QPA earns its largest energy gain is not the regime in which it also earns a pointwise gain. The QPA's tension gain and its pointwise cost are, in general, uncorrelated in both sign and magnitude across the (K, ξ) plane, even though the QPA's tension is lower and, by Proposition 1, more accurate in the variational sense everywhere.

Table 4: Per-species L^2 profile error $\varepsilon_{\phi,j} = \|\phi_j^{\text{model}} - \phi_j^{\text{GP}}\|_2 = (\int_{-L}^{L} (\phi_j^{\text{model}} - \phi_j^{\text{GP}})^2 \, dz)^{1/2}$ at $K = 2$, in units of $\sqrt{\xi_1} = 1$, comparing TPA and QPA component by component. Unlike a GP-norm-normalised ratio $\|\Delta\phi_j\|_2 / \|\phi_j^{\text{GP}}\|_2$, which necessarily $\to 0$ as $L \to \infty$ because $\|\phi_j^{\text{GP}}\|_2 \propto \sqrt{L}$ diverges while $\|\Delta\phi_j\|_2$ saturates, the unnormalised norm reported here is a genuine, L-independent property of the pair of profiles: it is converged to five significant figures for any $L \gtrsim 10$, since $\phi_j^{\text{model}} - \phi_j^{\text{GP}}$ decays exponentially away from the interface (verified numerically: $\|\Delta\phi_1\|_2 = 0.204255$ at $L = 7$ vs. 0.204258 at $L = 20$, for TPA at $\xi = 2$). All entries below use $L = 20$. Boldface marks the two cases in this cut where the QPA is marginally better than the TPA (Section 5.3).

ξ	$\varepsilon_{\phi,1}^{\text{TPA}}$	$\varepsilon_{\phi,1}^{\text{QPA}}$	$\varepsilon_{\phi,2}^{\text{TPA}}$	$\varepsilon_{\phi,2}^{\text{QPA}}$
0.5	0.0737	0.0849	0.1445	0.1736
0.7	0.0919	**0.0910**	0.1310	0.1438
1.4	0.1539	0.1694	0.1101	**0.1085**
2.0	0.2043	0.2454	0.1042	0.1201

Minimising σ_{trial} is a variational statement about the energy: a lower trial energy need not coincide with a smaller pointwise deviation of the profile, particularly for a six-parameter, tension-only objective with no explicit incentive to fit $\phi_j(z)$ pointwise. Consequently, if the scientific goal is an accurate order-parameter profile – for example for surface-mode spectra that depend on local curvature – the QPA as optimised here is not automatically preferable to the simpler TPA: it is more accurate in the variational tension and less accurate in the L^2 profile metric. A joint tension-and-profile objective is a natural extension of the present approach.

5.4. QPA inner structure

Figure 5 shows the optimised QPA structure at $K = 1.5, \xi = 1$: $z_1 \approx 1.193$, $z_2 \approx -1.193$, $\gamma_1 \approx \gamma_2 \approx 1.693$, $\beta_1 \approx \beta_2 \approx 0.827$, reproducing the independently optimised TPA parameters $(z_c, \gamma, \beta) \approx (1.193, 1.693, 0.827)$, confirming Eq. (14).

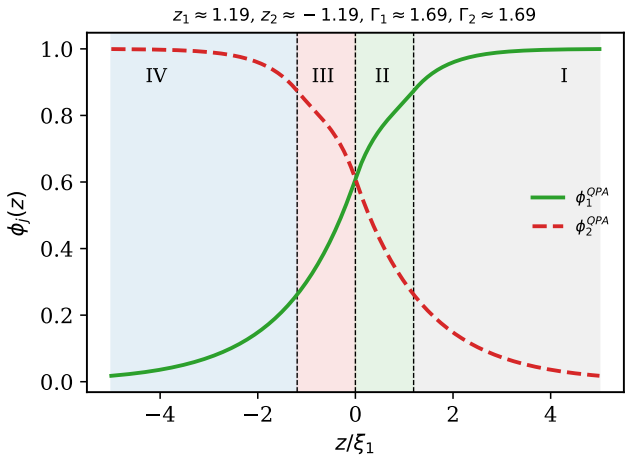


Figure 5: QPA domain structure at $K = 1.5$, $\xi = 1$, with the optimised matching points and majority curvatures reported in the title. Consistent with Section 3.4, $z_1 \approx -z_2$ and $\gamma_1 \approx \gamma_2$ to four significant figures.

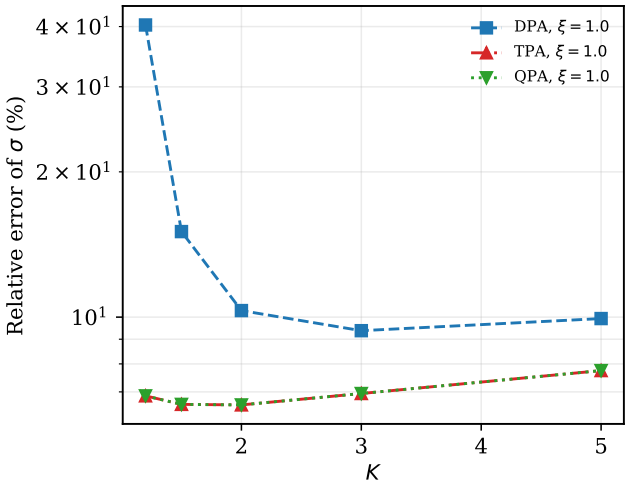


Figure 6: Relative tension error vs. K at $\xi = 1$, log scale. DPA (blue) is consistently the least accurate, worst near $K \to 1^+$; TPA and QPA coincide exactly at this symmetric point and both plateau near 7–8% for $K \lesssim 5$.

Table 5: Numerical test of the exact covariance $\sigma(K,\xi) = \xi\,\sigma(K,1/\xi)$, Eq. (16), at $K = 2$. The GP column uses two independent boundary-value solves; the DPA column is a fully analytic, optimiser-free evaluation; the TPA and QPA columns use two independently seeded optimisations. The relative discrepancy is reported as $100\times|\sigma(\xi)-\xi\sigma(1/\xi)|/\sigma(\xi)$.

Model	$\xi = 0.5$ vs. $1/\xi = 2.0$	$\xi = 0.7$ vs. $1/\xi \approx 1.429$
GP	$2\times10^{-10}\%$	–
DPA	$1.0\times10^{-3}\%$	$5\times10^{-4}\%$
TPA	$1.1\times10^{-3}\%$	$6\times10^{-4}\%$
QPA	$1.3\times10^{-3}\%$	$6\times10^{-4}\%$

5.5. *Relative accuracy across K*

Figure 6 shows that the TPA/QPA tension error is not monotonic in K: it is smallest near $K \approx 1.5$–2 (6.6%) and increases slightly toward both $K \to 1^+$ (6.9% at $K = 1.2$) and $K = 5$ (7.8%), a departure from the naive expectation that the strong-segregation limit is most favourable to parabolic approximations. This likely reflects the fixed six-parameter capacity of the model becoming a tighter constraint as the true GP profile develops the exponential $\phi_1 \sim e^{-\sqrt{2}z}$, $\phi_2 \sim e^{-\sqrt{K-1}z/\xi}$ tails of Eq. (12) more sharply at large K.

5.6. *Verification of the species-exchange covariance*

Table 5 verifies Eq. (16) directly. The GP benchmark satisfies the covariance relation to $2\times10^{-10}\%$, effectively at the solver's numerical tolerance, as expected since Eq. (16) is an exact property of the underlying field theory. The DPA (closed-form, optimiser-free), and the independently optimised TPA and QPA, satisfy the same relation to $\sim 10^{-3}\%$ – far tighter than their 6–7% accuracy relative to the GP benchmark – confirming that (i) the bulk curvatures of Section 3.2 are genuinely covariant under species exchange, and (ii) the multi-start optimisers converge to consistent, reproducible optima rather than to run-dependent local minima.

5.7. *Full (K, ξ) scan*

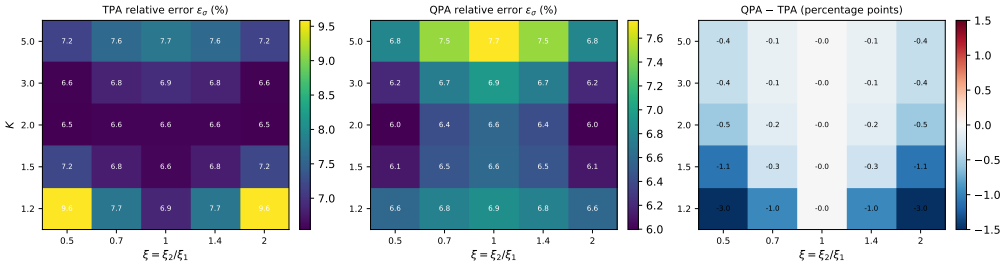


Figure 7: Relative tension error ε_σ (%, colour and in-cell value) of the TPA (left) and QPA (centre) across a 5×5 grid, rows $K \in \{1.2, 1.5, 2, 3, 5\}$ and columns $\xi \in \{0.5, 0.7, 1, 1/0.7, 2\}$, on the same colour scale for direct comparison; and their difference (right, $\varepsilon_\sigma^{\mathrm{QPA}} - \varepsilon_\sigma^{\mathrm{TPA}}$ in percentage points, blue = QPA better, white = no difference). The QPA is better than or equal to the TPA in every cell; the largest improvement, −3.0 percentage points, is at the bottom corners ($K = 1.2$, $\xi = 0.5$ or 2) – the weakest-segregation, most-asymmetric region of the grid. Columns for $\xi > 1$ use direct optimisation at $\xi \le 1$ combined with the exact covariance of Eq. (17), so only 15 of the 25 grid points required independent optimisation; the remaining 10 follow from Eq. (17), with two spot-checked points ($K = 2, \xi = 1.4$ and $K = 2, \xi = 2$) agreeing with direct optimisation to < 0.02 percentage points.

Figure 7 extends the two one-dimensional cuts of Sections 5.2–5.3 to a two-dimensional scan. The TPA and QPA error surfaces are both symmetric under $\xi \to 1/\xi$ by construction (Eq. (17)), visible as the left-right mirror structure of the first two panels about the $\xi = 1$ column. The QPA-minus-TPA difference (right panel) is negative (QPA better) at every grid point, with the largest absolute improvement, ~ 3.0 percentage points, at the weakest-segregation, most-asymmetric corner ($K = 1.2, \xi = 0.5$ or equivalently $\xi = 2$) – where the interface is widest and the two components' healing lengths differ most in absolute terms. The improvement shrinks toward the strong-segregation corner ($K = 5$), where the TPA error is already close to its $\xi = 1$ value ($\sim$ 7–8%) and the extra curvature freedom of the QPA has comparatively little further mismatch to correct. Averaged over all 25 grid points, the QPA reduces the TPA's tension error by 0.56 percentage points (mean), corresponding to a mean improvement fraction $\overline{G_\sigma} \approx 7\%$ (Eq. (28)); the distribution is strongly skewed, with the single weakest-segregation/most-asymmetric corner cell ($G_\sigma = 31\%$) contributing disproportionately to both the mean absolute and mean fractional gain.

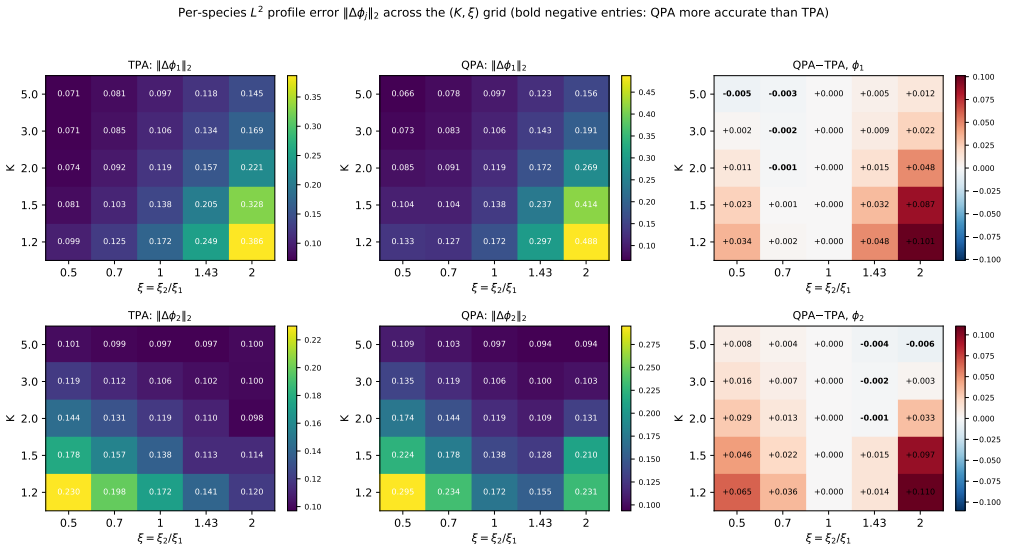


Figure 8: Per-species L^2 profile error $\|\Delta\phi_j\|_2$ (Table 4) across the same 5×5 grid as Fig. 7, for ϕ_1 (top row) and ϕ_2 (bottom row): TPA (left), QPA (centre), and their difference QPA−TPA (right; blue/bold = QPA more accurate, red = TPA more accurate). Unlike the tension heatmap, the difference panels are not single-signed: the QPA is pointwise *more* accurate (bold blue) in a structured minority of cells, confined to $K \geq 2$ and occurring in mirrored pairs related by $\xi \to 1/\xi$ together with $1 \leftrightarrow 2$ – an advantage for ϕ_1 at ξ_0 is matched by an advantage for ϕ_2 at $1/\xi_0$, visible as the point-reflection symmetry between the top-right and bottom-right difference panels. Elsewhere, and increasingly so toward weak segregation ($K \to 1.2$), the QPA is pointwise worse, most severely at the most asymmetric points ($\xi = 0.5, 2$) – the opposite corner of the grid from where its tension advantage G_σ is largest.

Figure 8 makes this tension–profile mismatch spatially explicit. Where the QPA's tension gain over the TPA is largest ($K \to 1.2$, $\xi \in \{0.5, 2\}$; the bottom corners of Fig. 7), its pointwise profile is markedly *worse* than the TPA's (the same corners in Fig. 8 are the most saturated red). Conversely, the cells where the QPA is pointwise better ($K \geq 2$, one species at $\xi = 0.7$ or $1/0.7$) are cells where G_σ is small (Fig. 7, centre-right region, error differences within 0.1–0.3 percentage points). The two effects are anticorrelated in location across the (K, ξ) plane: the

parameter regime that most rewards the QPA energetically is, empirically, the regime that most penalises it pointwise.

Two spot checks beyond the main 5×5 window ($1.2 \leq K \leq 5$, $0.5 \leq \xi \leq 2$) – $K = 10$ at $\xi = 1$, and $\xi = 0.25$ at $K = 2$ (with $\xi = 4$ following from Eq. (17)) – extend the trend. At $K = 10$, $\varepsilon_\sigma^{\mathrm{TPA}} = \varepsilon_\sigma^{\mathrm{QPA}} = 9.1\%$, consistent with the non-monotonic, saturating behaviour of Section 5.5. At $\xi = 0.25$ (and $\xi = 4$), $\varepsilon_\sigma^{\mathrm{TPA}} = 6.3\%$ and $\varepsilon_\sigma^{\mathrm{QPA}} = 4.7\%$, a 1.6 percentage point ($\sim 25\%$ relative) reduction – the largest relative improvement of the QPA over the TPA found in this study, and further evidence that the QPA's advantage grows with asymmetry.

6. Conclusions

The Quadruple-Parabola Approximation (QPA) is a six-parameter, exactly solvable, piecewise-quadratic trial-function model for the static planar interface of a phase-separated binary BEC, benchmarked here against an independently implemented numerical GP solver. Four results follow. (1) $F_{\mathrm{DPA}} \subset F_{\mathrm{TPA}}(\xi) \subset F_{\mathrm{QPA}}(\xi)$ by explicit parameter-space embedding, for every ξ; since every reported tension is evaluated with the exact GP functional, this containment makes $\sigma_{\mathrm{DPA}} \geq \sigma_{\mathrm{TPA}} \geq \sigma_{\mathrm{QPA}} \geq \sigma_{\mathrm{GP}}$ a variational consequence for the global optimum of each family, at every ξ (Proposition 1). (2) At $\xi_1 = \xi_2$, the additional reflection symmetry of the field theory forces the unconstrained six-parameter optimum onto the TPA subfamily exactly, giving $\sigma_{\mathrm{QPA}}(1) = \sigma_{\mathrm{TPA}}(1)$; this is confirmed numerically to four significant figures in the parameters and six digits in the tension at every symmetric case tested. (3) The QPA's tension-accuracy advantage over the TPA is real and reproducible but modest and strongly concentrated in parameter space: the improvement fraction G_σ, Eq. (28), is 7% on average over the (K, ξ) grid studied, but reaches 31% at the weakest-segregation, most-asymmetric corner tested and falls to a few percent or less over most of the remaining plane (against a baseline TPA/QPA error of 6.0–9.6%). (4) This advantage does not extend uniformly to the pointwise order-parameter profile: with a properly L-independent L^2 metric (Table 4), the QPA profile error exceeds the TPA's in most (species, K, ξ) combinations tested, but in 8 of 40 such combinations – confined to $K \geq 2$ and occurring in mirrored pairs related by the species-exchange covariance – the QPA is instead marginally more accurate pointwise, even though its tension is lower and hence more accurate in the variational sense everywhere.

Point (4), together with the concentration of G_σ revealed in point (3), is the central methodological finding of this study: a trial-function model can be a genuinely better energetic (variational) approximation – and, in a bounded region of parameter space, a meaningfully better one – without being a uniformly better profile approximation, and the two should be reported and optimised separately rather than conflated under the single label "accuracy." The QPA is best understood not as a general-purpose replacement for the TPA, but as a controlled case study of this energy–profile distinction, together with a practical recommendation: the additional parameters of the QPA are worth the extra bookkeeping specifically in the weak-segregation, strongly asymmetric regime where G_σ is largest, and are close to redundant elsewhere. Within the tested window ($K \in \{1.2, 1.5, 2, 3, 5\}$, $\xi \in \{0.5, 0.7, 1, 1/0.7, 2\}$), neither the TPA nor the QPA reaches sub-percent tension accuracy with the present optimisation protocol and admissible parameter domain; whether a different protocol, a wider domain, or a fundamentally different piecewise-polynomial construction is required for sub-percent accuracy is left open by the present study.

Two further results place the ξ-dependence of the construction on firm footing. The species-exchange covariance relation, $\sigma(K, \xi) = \xi\, \sigma(K, 1/\xi)$ (Section 3.5), is satisfied to $\sim 10^{-3}\%$ by

the DPA, TPA, and QPA when optimised completely independently at ξ and at $1/\xi$, confirming that the construction treats the two components consistently under $\xi \to 1/\xi$. The full 5×5 grid in (K, ξ) (Section 5.7), obtained at half the computational cost by exploiting this covariance, confirms that the QPA's modest advantage over the TPA is a stable feature across the segregation-strength/healing-length-ratio plane, largest at weak segregation and strong asymmetry and shrinking toward the strong-segregation limit.

Open directions include a joint tension-and-profile objective for the variational optimisation, to test whether the QPA's structural flexibility can be converted into a genuine pointwise-profile advantage; a finer and wider (K, ξ) scan to determine whether the 6–10% error plateau found here is intrinsic to the four-domain ansatz or an artefact of the optimiser; and extension of the present static, planar construction to trapped geometries and finite temperature.

Appendix A. Explicit matching matrices

With $c_1 = \cosh(\nu_1 z_1)$, $s_1 = \sinh(\nu_1 z_1)$, $c_3 = \cosh(\kappa_0 z_2)$, $s_3 = \sinh(\kappa_0 z_2)$ (note: κ_0, not κ_ξ, per Section 4.1), the ϕ_1 block is $\mathbf{M}^{(1)}\mathbf{x}^{(1)} = \mathbf{b}^{(1)}$ with $\mathbf{x}^{(1)} = (A_1, A_2, B_2, A_3, B_3, A_4)^\mathsf{T}$,

$$\mathbf{M}^{(1)} = \begin{pmatrix} -1 & -c_1 & -s_1 & 0 & 0 & 0 \\ \alpha_0 & -\nu_1 s_1 & -\nu_1 c_1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 \\ 0 & 0 & -\nu_1 & 0 & \kappa_0 & 0 \\ 0 & 0 & 0 & -c_3 & -s_3 & 1 \\ 0 & 0 & 0 & -\kappa_0 s_3 & -\kappa_0 c_3 & \kappa_0 \end{pmatrix}, \quad \mathbf{b}^{(1)} = \begin{pmatrix} -1+\beta_1 \\ 0 \\ \beta_1 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \tag{A.1}$$

The six rows enforce, in order: continuity of ϕ_1 at z_1; continuity of ϕ_1' at z_1; continuity of ϕ_1 at $z_0 = 0$; continuity of ϕ_1' at $z_0 = 0$; continuity of ϕ_1 at z_2; continuity of ϕ_1' at z_2. With $c_2^{(1)} = \cosh(\kappa_\xi z_1)$, $s_2^{(1)} = \sinh(\kappa_\xi z_1)$, $c_2^{(2)} = \cosh(\nu_2 z_2)$, $s_2^{(2)} = \sinh(\nu_2 z_2)$ (note: $\nu_2 = \sqrt{2\gamma_2}/\xi$, per Section 4.1), the ϕ_2 block is $\mathbf{M}^{(2)}\mathbf{x}^{(2)} = \mathbf{b}^{(2)}$ with $\mathbf{x}^{(2)} = (B_1, C_2, D_2, C_3, D_3, B_4)^\mathsf{T}$,

$$\mathbf{M}^{(2)} = \begin{pmatrix} -1 & c_2^{(1)} & s_2^{(1)} & 0 & 0 & 0 \\ -\kappa_\xi & -\kappa_\xi s_2^{(1)} & -\kappa_\xi c_2^{(1)} & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 \\ 0 & 0 & -\kappa_\xi & 0 & \nu_2 & 0 \\ 0 & 0 & 0 & -c_2^{(2)} & -s_2^{(2)} & -1 \\ 0 & 0 & 0 & -\nu_2 s_2^{(2)} & -\nu_2 c_2^{(2)} & -\alpha_\xi \end{pmatrix}, \quad \mathbf{b}^{(2)} = \begin{pmatrix} 0 \\ 0 \\ -\beta_2 \\ 0 \\ -1+\beta_2 \\ 0 \end{pmatrix}. \tag{A.2}$$

Both matrices follow by direct substitution of Eqs. (18)–(21) into the continuity conditions at $z_1, 0, z_2$, and were verified by direct numerical continuity testing: for representative asymmetric parameters (e.g. $K = 2, \xi = 0.5$), the piecewise profile and its first derivative agree from both sides to within finite-difference precision ($\sim 10^{-6}$–10^{-7}) at all three matching points. $|\det \mathbf{M}^{(1)}|$ and $|\det \mathbf{M}^{(2)}|$ exceeded 38 at every one of the 17 directly optimised parameter sets used in this paper.

References

[1] T.-L. Ho, V. B. Shenoy, *Binary Bose-Einstein condensates*, Phys. Rev. Lett. **77** (1996) 3276. DOI: 10.1103/PhysRevLett.77.3276.

[2] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell, C. E. Wieman, *Production of two overlapping Bose-Einstein condensates by sympathetic cooling*, Phys. Rev. Lett. **78** (1997) 586. DOI: 10.1103/PhysRevLett.78.586.

[3] D. S. Hall, M. R. Matthews, J. R. Ensher, C. E. Wieman, E. A. Cornell, *Dynamics of component separation in a binary mixture of Bose-Einstein condensates*, Phys. Rev. Lett. **81** (1998) 1539. DOI: 10.1103/PhysRevLett.81.1539.

[4] P. Ao, S. T. Chui, *Binary Bose-Einstein condensate mixtures in weakly and strongly segregated phases*, Phys. Rev. A **58** (1998) 4836. DOI: 10.1103/PhysRevA.58.4836.

[5] B. D. Esry, C. H. Greene, J. P. Burke, J. L. Bohn, *Hartree-Fock theory for double condensates*, Phys. Rev. Lett. **78** (1997) 3594. DOI: 10.1103/PhysRevLett.78.3594.

[6] E. V. Goldstein, P. Meystre, *Quasiparticle excitations in binary mixtures of Bose-Einstein condensates*, Phys. Rev. A **55** (1997) 2935. DOI: 10.1103/PhysRevA.55.2935.

[7] K. Mølmer, *Bose-Einstein condensates in traps: A solution of the multi-component Gross-Pitaevskii equation*, Phys. Rev. Lett. **80** (1998) 1804. DOI: 10.1103/PhysRevLett.80.1804.

[8] E. Timmermans, *Phase separation of Bose-Einstein condensates*, Phys. Rev. Lett. **81** (1998) 5718. DOI: 10.1103/PhysRevLett.81.5718.

[9] H. Pu, N. P. Bigelow, *Properties of two-species Bose condensates*, Phys. Rev. Lett. **80** (1998) 1130. DOI: 10.1103/PhysRevLett.80.1130.

[10] R. A. Barankov, *Boundary of two mixed Bose-Einstein condensates*, Phys. Rev. A **66** (2002) 013612. DOI: 10.1103/PhysRevA.66.013612.

[11] H. Shi, W.-M. Zheng, S.-T. Chui, *Interface profiles of a binary Bose-Einstein condensate mixture*, Phys. Rev. A **61** (2000) 063613. DOI: 10.1103/PhysRevA.61.063613.

[12] A. Roy, S. Gautam, D. Angom, *Interfacial tension and structure of phase-separated two-component Bose-Einstein condensates*, Phys. Rev. A **89** (2014) 013617. DOI: 10.1103/PhysRevA.89.013617.

[13] C. Ticknor, *Interfacial properties of binary Bose mixtures*, Phys. Rev. A **89** (2014) 053601. DOI: 10.1103/PhysRevA.89.053601.

[14] S. Gautam, P. Muruganandam, D. Angom, *Excitation spectra of two-component Bose-Einstein condensates in displaced harmonic traps*, J. Phys. B **46** (2013) 055301. DOI: 10.1088/0953-4075/46/5/055301.

[15] J. O. Indekeu, B. Van Schaeybroeck, *Extraordinary wetting phase diagram for mixtures of Bose-Einstein condensates*, Phys. Rev. Lett. **93** (2004) 210402. DOI: 10.1103/PhysRevLett.93.210402.

[16] B. Van Schaeybroeck, J. O. Indekeu, *Critical wetting, first-order wetting, and prewetting phase transitions in binary mixtures of Bose-Einstein condensates*, Phys. Rev. A **91** (2015) 013626. DOI: 10.1103/PhysRevA.91.013626.

[17] N. V. Thu, T. H. Phat, P. T. Song, *Wetting phase transition of two segregated Bose-Einstein condensates restricted by a hard wall*, Phys. Lett. A **380** (2016) 1164. DOI: 10.1016/j.physleta.2016.02.021.

[18] B. Van Schaeybroeck, P. Navez, J. O. Indekeu, *Interface potential and line tension for Bose-Einstein condensate mixtures near a hard wall*, Phys. Rev. A **105** (2022) 053309. DOI: 10.1103/PhysRevA.105.053309.

[19] M. Goldman, J. Royo-Letelier, *Sharp interface limit for two components Bose-Einstein condensates*, ESAIM: Control, Optimisation and Calculus of Variations **21** (2015) 603. DOI: 10.1051/cocv/2014040.

[20] B. Van Schaeybroeck, *Interface properties of a phase-separated Bose-Einstein condensate mixture*, Phys. Rev. A **78** (2008) 023624. DOI: 10.1103/PhysRevA.78.023624.

[21] H. Takeuchi, K. Kasamatsu, *Nambu-Goldstone modes in segregated Bose-Einstein condensates*, Phys. Rev. A **88** (2013) 043612. DOI: 10.1103/PhysRevA.88.043612.

[22] D. A. Takahashi, M. Kobayashi, M. Nitta, *Topological defects and surface tension in multi-component BECs*, Phys. Rev. B **91** (2015) 184501. DOI: 10.1103/PhysRevB.91.184501.

[23] T. H. Phat, L. V. Hoa, N. T. Anh, N. V. Long, *Phase separation and surface state of binary Bose-Einstein condensates*, Ann. Phys. (N.Y.) **324** (2009) 2074. DOI: 10.1016/j.aop.2009.06.007.

[24] I. E. Mazets, *Phase segregation in a trapped binary Bose-Einstein condensate*, Phys. Rev. A **65** (2002) 033618. DOI: 10.1103/PhysRevA.65.033618.

[25] K. Kasamatsu, M. Tsubota, M. Ueda, *Vortices and multi-component Bose-Einstein condensates*, Int. J. Mod. Phys. B **19** (2005) 1835. DOI: 10.1142/S021797920503031X.

[26] D. S. Petrov, *Quantum mechanical stabilization of a collapsing Bose-Bose mixture*, Phys. Rev. Lett. **115** (2015) 155302. DOI: 10.1103/PhysRevLett.115.155302.

[27] D. Bonn, J. Eggers, J. O. Indekeu, J. Meunier, E. Rolley, *Wetting and spreading*, Rev. Mod. Phys. **81** (2009) 739. DOI: 10.1103/RevModPhys.81.739.

[28] M. J. Edmonds, K. L. Lee, N. P. Proukakis, *Thermal effects on the interface of phase-separated binary condensates*, Phys. Rev. A **91** (2015) 011602(R). DOI: 10.1103/PhysRevA.91.011602.

[29] H. Takeuchi, Y. Mizuno, K. Dehara, *Spontaneous interface formation and wetting in multi-component ultracold gases*, Phys. Rev. A **92** (2015) 043608. DOI: 10.1103/PhysRevA.92.043608.

[30] J. O. Indekeu, C.-Y. Lin, N. V. Thu, B. V. Schaeybroeck, T. H. Phat, *Static interfacial properties of Bose-Einstein-condensate mixtures*, Phys. Rev. A **91** (2015) 033615. DOI: 10.1103/PhysRevA.91.033615.

[31] Z. Deng, B. Van Schaeybroeck, C.-Y. Lin, N. V. Thu, J. O. Indekeu, *Interfacial tension and wall energy of a Bose-Einstein condensate binary mixture: Triple-parabola approximation*, Physica A **444** (2016) 1027. DOI: 10.1016/j.physa.2015.10.072.

[32] P. Virtanen *et al.* (SciPy 1.0 Contributors), *SciPy 1.0: fundamental algorithms for scientific computing in Python*, Nat. Methods **17** (2020) 261. DOI: 10.1038/s41592-019-0686-2.