

Near-axis measures of quasi-isodynamic configurations

Eduardo Rodríguez¹  and Gabriel Plunk¹ 

¹Max Planck Institute for Plasma Physics, Greifswald 17491, Germany

Corresponding author: Eduardo Rodríguez, eduardo.rodriguez@ipp.mpg.de

(Received 5 May 2025; revision received 12 June 2025; accepted 12 June 2025)

We present a number of measures and techniques to characterise and effectively construct quasi-isodynamic stellarators within the near-axis framework, without the need to resort to the computation of global equilibria. These include measures of the reliability of the model (including aspect-ratio limits and the appearance of ripple wells), quantification of omnigenity through ϵ_{eff} , measure and construction of MHD-stabilised fields, and the sensitivity of the field to the pressure gradient. The paper presents, discusses and gives examples of all of these, for which expansions to second order are crucial. This opens the door to the exploration of how key underlying choices of the field design govern the interaction of desired properties (‘trade-offs’) and provides a practical toolkit to perform efficient optimisation directly within the space of near-axis quasi-isodynamic configurations.

Key words: fusion plasma

1. Introduction

The range of stellarator magnetic confinement fields is broad (Spitzer Jr 1958; Boozer 1998; Helander 2014), making it a daunting task both to understand and design them. This is true even when one restricts attention to particular subsets, known to have certain key properties such as omnigenity (Hall & McNamara 1975; Bernardin, Moses & Tataronis 1986; Cary & Shasharina 1997; Helander 2014). Omnigenous fields are capable of confining, by definition, all collisionless charged particle orbits (Northrop 1961; Littlejohn 1983; Wesson 2011; Blank 2004); hence, the particular interest in them. Even though the magnetic field magnitude $|B|$ on their nested flux surfaces requires careful consideration (Boozer 1983; Nührenberg & Zille 1988; Cary & Shasharina 1997; Parra *et al.* 2015), there remains a wide range of choices to be made. This is true whether the problem is attacked from the optimisation perspective (Mynick 2006) or a more controlled analytical approach.

One prominent case of the latter, with which this paper shall be concerned, is the near-axis description of the field (Mercier 1964; Solov’ev & Shafranov 1970;

Lortz & Nührenberg 1976; Garren & Boozer 1991b). Consisting of an asymptotic description of the equilibrium field in the distance (described by r) from its centre (called the magnetic axis), the fields reduce to being represented by only a few functions and parameters. These constitute a model for the field that has proved powerful in advancing the theoretical understanding of optimised omnigeneous stellarators (Mercier 1964; Lortz & Nührenberg 1976; Jorge & Landreman 2020; Landreman & Jorge 2020; Landreman 2021; Rodríguez *et al.* 2022b, 2024) as well as providing a practical tool in stellarator design (Camacho Mata, Plunk & Jorge 2022; Jorge *et al.* 2022; Landreman & Sengupta 2019; Landreman 2022; Rodríguez *et al.* 2023). The latter has however been restricted for the most part to a particularly mature subclass of omnigeneous stellarators: namely, quasisymmetric stellarators (Boozer 1983; Burby, Kallinikos & MacKay 2020; Nührenberg & Zille 1988; Rodríguez *et al.* 2020). Only recently (Camacho Mata *et al.* 2022; Camacho Mata & Plunk 2023; Jorge *et al.* 2022; Plunk, Landreman & Helander 2019; Rodríguez & Plunk 2023; Rodríguez, Plunk & Jorge 2025) has the near-axis framework of quasi-isodynamic (QI) fields (Cary & Shasharina 1997; Helander & Nührenberg 2009; Nührenberg 2010), the other big class of optimised stellarators, been sufficiently developed. The lateness of this development has been a consequence of the complex treatment that these fields demand; while QS fields possess a direction of symmetry in $|\mathbf{B}|$ (either toroidal or helical), the QI fields do not (although they have closed poloidal contours).

The objective of this work is to help enable full use of the recent advances in the near-axis description of QI stellarators. In particular, although Rodríguez *et al.* (2025) verified the construction of approximately QI fields to second order, it left open the issue of how to go about obtaining ‘good’ solutions, i.e. how to best select input parameters, how to evaluate candidate fields and how to navigate the extremely large design space that the near-axis construction opens up. To address this, we develop a set of measures and techniques using second-order near-axis expansion to incorporate notions of neoclassical transport, magnetohydrodynamic (MHD) stability and equilibrium sensitivity into the construction of QI stellarators. With these tools in hand, we are in a position to construct and more thoroughly evaluate stellarator design candidates, and conduct systematic studies exploring things like optimisation trade-offs.

The paper is organised as follows. In §2, essential definitions and notation of near-axis theory are provided. The rest of the paper is divided into sections describing different types of measures. In §3, we introduce the measure of field error δB that quantifies the deviation of a finite aspect ratio consistent equilibrium from its asymptotic near-axis description, following Landreman (2021). This does however not qualify the physical impact of such deviations. Thus, in §4, we present calculations of the neoclassical ϵ_{eff} effective ripple (Nemov *et al.* 1999) within the near-axis expansion, a form to quantify the lack of omnigeneity in which second order is central. In §5, we study the appearance of localised ripple wells in the field, a feature that the near-axis description of ϵ_{eff} cannot capture, focusing on finding at what aspect ratio these first appear. Some questions of MHD stability are addressed in §6, where we explore the sensitivity of the magnetic well (Greene 1997; Landreman & Jorge 2020) of a near-axis field to the choice of second-order shaping. Finally, in §7, we study the role played by the pressure gradient and quantify the sensitivity of a given field to changes in plasma β . We conclude with some closing remarks.

2. Definitions and notation

This paper does not re-derive the inverse-coordinate near-axis framework (Garren & Boozer 1991b; Landreman & Sengupta 2019), but the work is best understood with some familiarity with how the equilibrium equations for an approximately QI field are solved through second order in the expansion. A full detailed and pedagogical account may be found from Rodriguez *et al.* (2025) and references within, but we provide some pointers on notation to help make the presentation as self-contained as possible.

In the vein of the near-axis expansion, all equilibrium quantities are expanded in a Taylor–Fourier series in powers of $r = \sqrt{2\psi/\bar{B}}$ (a pseudo-radial distance from the axis, where ψ is 2π times the toroidal magnetic flux and $\bar{B}$ a reference magnetic field) and cosine (sine) harmonic in the poloidal, θ , (or helical, $\chi = \theta - N\varphi$ for $N \in \mathbb{Z}$) angle. Here, the angles are part of the Boozer coordinate system (Boozer 1981; D’haeseleer *et al.* 2012), which eases the description of the magnetic field. This way, any function f is written as $f = \sum_{n=0}^{\infty} r^n f_n$ and $f_n = \sum_{m=0}^n (f_{nm}^c(\varphi) \cos m\chi + f_{nm}^s(\varphi) \sin m\chi)$, where the latter sum is over even or odd numbers depending on the parity of n .¹ This subscript notation (Rodriguez *et al.* 2025, § 2.3) will be repeatedly used and we refer to n as the order (first, second, etc.) of the expansion.

In the inverse-coordinate framework, a magnetic field is described by only a few functions. First, the φ -functions involved in the expansion of the magnetic field strength, $B = |\mathbf{B}|$. Second, because of the inverse-coordinate nature of the expansion, we describe flux surfaces of the magnetic field through the functions X , Y and Z , which define the relative position of flux surfaces with respect to the magnetic axis in the Frenet–Serret frame (normal, binormal and tangent, respectively) of the latter (Rodriguez *et al.* 2025, (2.2)). Thus, the expansion of these functions also forms part of the description. Finally, the covariant components of the field, $\{G, I, B_\psi\}$, are also explicitly involved (Rodriguez *et al.* 2025, (2.1)), as is the rotational transform of the field, ι .

The governing equilibrium equations impose important relations between the components of all of these quantities, and that is on what solving the near-axis expansion is focused. Those relations are detailed in the context of QI fields by Rodriguez *et al.* (2025). It is however important to emphasise which ingredients are required to uniquely define a field in this framework. An approximately QI near-axis field is uniquely defined to first order by the set $\{\kappa(\ell), \tau(\ell), B_0(\varphi), \bar{d}(\varphi), \alpha_1(\varphi)\}$, where the first two are the curvature and torsion defining the shape of a magnetic axis (which must have vanishing curvature points (Plunk *et al.* 2019; Camacho Mata & Plunk 2023; Rodriguez *et al.* 2025), the third is the magnetic strength along the axis as a function of the toroidal Boozer angle (whose turning points must match the $\kappa = 0$ points), and the latter two define B_1 (Rodriguez *et al.* 2025, (2.4)),

$$B_1 = \kappa(\varphi)B_0(\varphi)\bar{d}(\varphi)\cos[\chi - \alpha_1(\varphi)]. \quad (2.1)$$

In an ideal QI field, κ and $\bar{d}$ must be odd and even functions, respectively, with respect to B_0 about every trapping well minima. The function α_1 may be defined as (if we take $\varphi = 0$ to be the minimum of B_0)

$$\alpha_1 = (\iota_0 - N)\varphi - \pi/2 + \alpha_{\text{buf}}(\varphi), \quad (2.2)$$

¹We may use the following shorthand as well: $F_{11}^c = F_{1c}$, $F_{11}^s = F_{1s}$, $F_{22}^c = F_{2c}$, $F_{22}^s = F_{2s}$ for any function that F may be.

where in the ideal case, $\alpha_{\text{buf}} = 0$. The periodicity constraint on α_1 generally requires finite non-omnigenous deviations close to B_0 maxima (i.e. buffers), which are captured by finite α_{buf} (Plunk *et al.* 2019; Rodríguez & Plunk 2023; Rodríguez *et al.* 2025). These first-order choices also uniquely determine the elliptical shaping of flux surfaces around the magnetic axis.

At second order, the key degrees of freedom (inputs) are $\{p_2, X_{2c}(\varphi), X_{2s}(\varphi)\}$, where the former parameter is the pressure gradient and the latter two are related to triangular shaping. The latter two are directly linked to the shaping of the field strength on the flux surfaces as well. For the detailed role played by each of these and a complete consistent description, we again refer the reader to Rodríguez *et al.* (2025).

3. Field truncation error δB

The construction of QI fields using the near-axis expansion is rather non-trivial already at first order in the expansion. Different choices of magnetic axis and the shaping of elliptical cross-sections require a careful balance. Unless actively sought, such fields will tend to present particularly elongated and twisted shapes (Camacho Mata *et al.* 2022).

Even when suitable parameters are found, we must always remember that the field resulting from the near-axis description is approximate in nature. The study and design of a stellarator field requires, in practice, the full non-expanded equilibrium field. Thus, we generally need to construct a finite aspect ratio version of the near-axis field, which will show some deviations from the purely asymptotic description. Significant deviations can lead to the loss of those properties of the field carefully instilled in the near-axis field. This particularly concerns $|\mathbf{B}|$, the field strength, which controls guiding centre dynamics of particles, and omnigenicity. When such deviations are limited, the near-axis description may be deemed good. This measure of goodness has been previously used (Camacho Mata *et al.* 2022; Camacho Mata & Plunk 2023) to gauge the suitability of first-order near-axis QI fields as interesting candidates for optimised stellarators, and is thus of some interest. This focus on first order allows us to make a first assessment of the near-axis construction without having to explicitly go to higher order. Indeed, depending on the choice of first-order parameters, there is a strong degree of variation of truncation error, which has in the past lead to many unfeasible configurations that do not have the desired field at a reasonable aspect ratio.

To define this goodness more precisely, let us introduce the field strengths B_{global} and $B_{\text{nae}}^{(1)}$. Define B_{global} as the field strength of the global equilibrium at its outermost flux surface $\psi = \psi_a$ parametrised by θ and φ , angular poloidal and toroidal Boozer coordinates (Boozer 1981). This global equilibrium is meant to be a finite aspect ratio representation of the near-axis field and we follow the prescription of Landreman & Sengupta (2019). That is, given a first-order near-axis field, we calculate the shape of the elliptical flux surface at some radius $r = a$ (i.e. $a = \sqrt{2\psi_a/\bar{B}}$) and use it as the outer boundary of a global nested-flux surface, vacuum field. The solution then defines $\mathbf{B}_{\text{global}}$.

The comparison of this global field can then be made to the first-order near-axis one, evaluated at the edge, $B_{\text{nae}}^{(1)} = B_0 + a^2 B_1$. We define the field error as

$$\delta B = \sqrt{\int (B_{\text{global}} - B_{\text{nae}}^{(1)})^2 d\varphi d\theta} / \sqrt{\int B_{\text{global}}^2 d\varphi d\theta}; \quad (3.1)$$

the normalised root-mean-square difference. Large values imply significant differences between the realised field and the near-axis model, and thus the expectation of them behaving quite differently.

Gauging the goodness of a near-axis expansion (NAE) field as defined in (3.1) requires us to solve a global equilibrium. Whether one uses VMEC (Hirshman & Whitson 1983), DESC (Dudt & Kolemen 2020), GVEC (Hindenlang *et al.* 2019) or any other nested flux surface solver, computing δB is numerically costly as compared with the near-axis field, impeding in addition its theoretical investigation. Can we make sense of this goodness of the field from within the near-axis framework?

3.1. Near-axis perspective on δB

By definition, δB measures the discrepancy between the full- B and the truncation of an asymptotic form. That is, δB is, in a sense, a measure of the truncation error of the asymptotic expansion. The truncation error of an asymptotic series (Bender & Orszag 2013, § 3.5) can be gauged by the next order in the expansion,² and thus the calculation to second order could possibly be used to estimate δB . There is however an important caveat: in the near-axis expansion, moving from order N to order $N + 1$ in the expansion brings new degrees of freedom. From first to second, we may shape the triangularity of surfaces in an infinite number of ways.

Amongst this myriad of possibilities, we should pick that which most closely corresponds to our particular scenario. This was dealt with by Landreman & Sengupta (2019) and Landreman (2021), who used the first-order near-axis field of interest at $r = a$ to impose the shape of the outermost boundary of a global equilibrium, which they then asymptotically described. Formally, the description of the global equilibrium field $\tilde{B}$ requires a double expansion in r (a near-axis expansion) and a (a large aspect ratio expansion to deal with the ‘global’ nature of the field). Asymptotically computing this field and evaluating its difference with $B_{\text{nae}}^{(1)}(r = a) = B_0 + aB_1$ can be shown to be (Landreman 2021, § 5.2)³

$$\tilde{B}_2 = a^2 \tilde{B}_0^{(2)} + r^2 (\tilde{B}_{20}^{(0)} + \tilde{B}_{2c}^{(0)} \cos 2\chi + \tilde{B}_{2s}^{(0)} \sin 2\chi). \quad (3.2)$$

The first is a mirror term, (C24) of Landreman (2021), and arises from the finite aspect ratio construction (expansion in a) and is necessary to guarantee a constant toroidal flux $\psi_a = \bar{B}a^2/2$. As such, it is a perturbation on the near-axis input $B_0(\varphi)$, the magnetic field on axis that sets the leading trapping well structure. The r -dependent terms in (3.2) correspond to the second-order near-axis expansion of $\tilde{B}$, whose shaping is chosen to match the outermost surface. They may therefore be found in a form analogous to the regular second-order construction, solving (A41)–(A42) of Landreman & Sengupta (2019) using (C9)–(C12) of Landreman (2021) and then applying (A34)–(A36) of Landreman & Sengupta (2019).

²Note that this is true by definition for sufficiently small r , the expansion variable. In any form, the next order bears information about the truncation error.

³We note that the construction detailed by Landreman (2021), in particular (C30), is not correct. This is apparent from the lack of symmetry of the expression with respect to the mid-point of a field period. The expression should be correctly constructed taking the N -fold symmetry into account, $\bar{f}^{(2)}(\varphi) = \int_0^\varphi \hat{B}(\varphi') d\varphi' + (1/2 - N\varphi/2\pi) \int_0^{2\pi/N} \hat{B}(\varphi') d\varphi' - N/2\pi \int_0^{2\pi/N} d\varphi'' \int_0^{\varphi''} \hat{B}(\varphi') d\varphi'$.

The field error δB may therefore be estimated as

$$\delta B_{\text{ar}} = a^2 \sqrt{\frac{1}{2\pi} \int_0^{2\pi} \left[(\tilde{B}_{20}^{(0)} + \tilde{B}_0^{(2)})^2 + \frac{(\tilde{B}_{2c}^{(0)})^2}{2} + \frac{(\tilde{B}_{2s}^{(0)})^2}{2} \right] d\varphi}, \quad (3.3)$$

where all fields have been evaluated at the boundary $r = a$.

3.2. Numerical implementation

We have implemented the computation of δB_{ar} numerically in `pyQIC`,⁴ following the original work of Landreman (2021), and introducing all the appropriate generalisations that apply to QI fields, especially the non-uniform B_0 and correct handling of half-helicity axes (Camacho Mata & Plunk 2023; Rodríguez *et al.* 2025). In practice, the resources devoted to finding $\tilde{B}$ quantities are analogous to those employed when solving the second-order equations of a near-axis equilibrium. In fact, the very equations involved are quite similar and more specific details may be found in the code itself.

To benchmark δB , we need a set of near-axis configurations and their associated finite aspect ratio equilibria. The latter are obtained running the global flux-surface equilibrium code `VMEC` (Hirshman & Whitson 1983), constructed using the near-axis surface at $r/R = 0.07$ (roughly an aspect ratio of $A = 14$). As for the benchmark set, we use a set of 1680 approximately QI, half-helicity near-axis fields (see some details in Appendix A and a fuller description in an upcoming publication). To compute δB as defined in (3.1), the field strength is computed in Boozer coordinates using `BOOZXFORM` (Sanchez *et al.* 2000). The comparison of δB to the near-axis estimate δB_{ar} is shown in figure 1.

The agreement is excellent (the correlation is in excess of 0.99), especially strong for the lowest δB values. At larger δB , there is a clear systematic deviation, with the near-axis δB_{ar} overestimating the field discrepancy. This overestimation is a result of fields with larger δB having a tendency for increased shaping and sensitivity, and thus smaller radii of convergence. Evaluation at a finite aspect ratio thus tends to overestimate the magnitude of the field. However, δB_{ar} retains a monotonic relation with δB .

The behaviour of the field error is clearly seen to depend on the number of field periods. They appear to form separate clusters, with the larger field period numbers leading to larger deviations (in agreement with the arguments before, with the exception of the special case $N = 1$). Rescaling $\delta B/N^2$ (see inset of figure 1) clusters all different fields together. This quadratic scaling follows the involvement of second derivatives in φ in the construction of the second-order near-axis equilibrium expansion (see Rodríguez *et al.* 2025) and approximate scaling symmetries in axis shapes as discussed by Plunk *et al.* (2025).

Altogether, the error field near-axis measure δB_{ar} is an excellent predictor of δB . The systematic deviation may lead to a failure of a qualitative agreement at larger errors, but even this departure could be refined by using a fit of the curve in figure 1 to map δB_{ar} to δB .

⁴The code may be found in <https://github.com/SebereX/pyQIC.git>, a significantly different version of the original code of Jorge *et al.* (2022), and the particular branch is specified in the Zenodo repository.




FIGURE 1. Benchmark of the error field measure δB_{ar} . The main plot compares the field error δB , (3.1), evaluated using the global equilibrium computed with VMEC, to the near-axis estimate δB_{ar} for the QI near-axis configurations in the benchmark database of Appendix A.1. The black broken line represents $\delta B = \delta B_{nae}$ and the dotted one is the moving average of the scatter. The colours for δB_{ar} denote the number of field periods of the configurations (see legend). The inset plot shows the same data scaled by the field period number as $1/N^2$.

4. Effective ripple: ϵ_{eff}

Although the simplicity of the field error δB makes it appealing as a measure of how ‘good’ a given first-order near-axis field is, it ignores the fact that not all field deviations (even if they may correspond to the same δB) affect the behaviour of the field equally. In particular, some deviations will spoil omnigenity more than others. The field error measure overlooks a second important point, namely that deviations from omnigenity are unavoidable to a certain degree (Cary & Shasharina 1997) even at first order (Plunk *et al.* 2019; Rodríguez & Plunk 2023), so it is important to separately evaluate such deviations at each order. Thus, we are in need of some way of measuring the departure from ideal omnigenity, of which there are many (Mynick 2006; Nemov *et al.* 2008; Rodríguez *et al.* 2022a; Goodman *et al.* 2023; Dudt *et al.* 2024).

In this paper, we consider the effective ripple ϵ_{eff} as the measure of deviations from omnigenity or ‘omnigenity error’. This is a geometric, dimensionless scalar quantity defined by Nemov *et al.* (1999) whose magnitude (or more precisely, that of $\epsilon_{eff}^{3/2}$) controls the neoclassical electron heat transport across a flux surface in a plasma that balances the radial losses with collisions; i.e. the so-called $1/\nu$ regime. Given its clear physical interpretation, the evaluation of ϵ_{eff} is commonly used to assess modern optimised stellarators, which aim at values below a per cent. Here, we shall be interested in evaluating it within the near-axis framework. Doing so will: (i) avoid having to solve global equilibria and perform neoclassical calculations to assess fields; and (ii) provide insight on how near-axis choices affect it.

Following Nemov *et al.* (1999, (29)), we express ϵ_{eff} in our own notation as

$$\epsilon_{\text{eff}}^{3/2} = \frac{\pi(\bar{R}\bar{B})^2}{8\sqrt{2}} \lim_{L \rightarrow \infty} F(L) \int_{1/B_{\text{max}}}^{1/B_{\text{min}}} d\lambda \lambda \sum_{j=1}^{N_{\text{well}}} \mathcal{E}_j(\lambda), \quad (4.1)$$

where

$$\mathcal{E}_j(\lambda) = \frac{H_j(\lambda)^2}{I_j(\lambda)}, \quad (4.2a)$$

$$H_j(\lambda) = \frac{1}{\bar{B}^2} \int_{\ell_{j,L}}^{\ell_{j,R}} \mathcal{H}(\lambda, B) \frac{\mathbf{B} \times \nabla B \cdot \nabla \psi}{B} d\ell, \quad (4.2b)$$

$$I_j(\lambda) = \int_{\ell_{j,L}}^{\ell_{j,R}} \frac{\sqrt{1 - \lambda B}}{B} d\ell, \quad (4.2c)$$

$$F(L) = \left(\int_0^L \frac{d\ell}{B} \right) \left(\int_0^L |\nabla \psi| \frac{d\ell}{B} \right)^{-2} \quad (4.2d)$$

and

$$\mathcal{H}(\lambda, B) = \frac{\sqrt{1 - \lambda B}}{(B/\bar{B})^2} \left(\frac{4}{\lambda B} - 1 \right). \quad (4.3)$$

The factors $\bar{R}$ and $\bar{B}$ are reference length (major radius) and magnetic field magnitudes that normalise ϵ_{eff} to make it a dimensionless quantity. The effective ripple is then to be understood as a sum over all trapped particles, λ , in all wells, labelled j , along the field line, running from 0 to L (in the limit of $L \rightarrow \infty$). That is, on an irrational surface, over the whole flux surface. Note that this involves a sum over very different classes of trapped particles, including particles that travel over a fraction or many periods of the torus, which makes the calculation particularly hard.

Each of these classes of particles contributes to the total ripple through $\mathcal{E}_j(\lambda) \geq 0$, defined in (4.2a), which for the field to be omnigenous, must vanish for all trapped particles. The non-omnigenous drive comes through H_j , an average between bounce points $\ell_{j,R/L}$ of the off-surface magnetic drift, $\mathbf{v}_d \cdot \nabla \psi$, weighted by the parallel velocity $\propto \sqrt{1 - \lambda B}$. The heat flux, Γ , depends quadratically on H_j . This dependence stems from having $\Gamma \sim \delta f \mathbf{v}_d \cdot \nabla \psi$, where the perturbed particle distribution δf arises from a balance between the average drift and collisions. The competing mixing rate of collisions is related to the function I_j , which is formally similar to the second adiabatic invariant. Finally, $F(L)$ is a geometric normalisation factor that plays the role of normalising the off-surface projection of the drift.

4.1. Identification of the expansion

Many of the complexities involved in evaluating ϵ_{eff} , (4.1), disappear when its asymptotic near-axis form is considered. To see this, let us start by introducing the asymptotic form of the magnetic field strength, which defines trapped particle classes and determines the radial drift. Concerned with the description of a magnetic field

that is approximately QI, we can write $B \approx B_0(\varphi) + rB_1(\chi, \varphi) + r^2B_2(\chi, \varphi)$, where (Rodríguez *et al.* 2025)

$$B_1 = -B_0(\varphi)d(\varphi) \sin[\alpha - \alpha_{\text{buf}}(\varphi)], \quad (4.4a)$$

$$B_2 = B_{20}(\varphi) + B_{2c}(\varphi) \cos 2\chi + B_{2s}(\varphi) \sin 2\chi. \quad (4.4b)$$

To zeroth order in r , $B_0(\varphi)$ defines (by assumption) a single toroidal trapping well per field period. The higher order contributions in $r = \sqrt{2\psi/\bar{B}}$ then deform this underlying well to make distinct ones for different field lines, labelled by $\alpha = \theta - \iota\varphi$. Asymptotically, they do not introduce additional wells. The particular form of B_1 in (4.4a) corresponds to that of an omnigenous field (when taking d to be an odd function respect to B_0 , and $d=0$ at the bottom and top of the trapping well), except where $\alpha_{\text{buf}} \neq 0$, i.e. the buffer regions. These are for the most part unavoidable near the edges of the toroidal domain (Cary & Shasharina 1997; Plunk *et al.* 2019; Rodríguez & Plunk 2023). The second-order field contribution is for now chosen simply to satisfy stellarator symmetry, so B_{20} and B_{2c} are even, and B_{2s} odd.

As a result of the trapping wells becoming labelled by α , the infinite sum over trapping wells in the effective ripple, (4.1), becomes by Weyl's lemma of equidistribution (Weyl 1916) (Stein & Shakarchi 2011, Lemma 2.2) (assuming an irrational rotational transform),

$$\sum_{j=1}^{N_{\text{well}}} f_j \approx \frac{N_{\text{well}}}{2\pi} \int_0^{2\pi} f(\alpha) d\alpha, \quad (4.5)$$

where N_{well} is the number of wells within a field line portion of length L . That is, the sum over wells simply corresponds to an average over α .

Changing variables from ℓ to φ (with the appropriate Boozer Jacobian $\mathcal{J} = (G + \iota I)/B^2$, D'haeseleer *et al.* 2012, (6.6.8b)) and writing flux surface averages in terms of (α, φ) ,

$$\epsilon_{\text{eff}}^{3/2} = \frac{\pi}{8\sqrt{2}} \frac{(\bar{R}\bar{B})^2}{\mathcal{G}^2} \int_{1/B_{\text{max}}}^{1/B_{\text{min}}} \lambda \hat{\mathcal{E}}(\lambda) d\lambda, \quad (4.6)$$

where

$$\hat{\mathcal{E}}(\lambda) = \frac{1}{\pi} \int_0^{2\pi} \frac{\hat{H}(\lambda, \alpha)^2}{\hat{I}(\lambda, \alpha)} d\alpha, \quad (4.7)$$

and all remaining definitions can be found explicitly at the beginning of Appendix B. The functions $\hat{H}$ and $\hat{I}$ (directly linked to the previous H_j and I_j of Nemov *et al.* 1999), (B1a)–(B1b), have been defined to be dimensionless, and $\mathcal{G}^2$ (directly related to F) has dimensions of $[L]^2$, which matches the presence of the normalising $\bar{R}$.

In this form, the asymptotic expansion in powers of r can be carried out explicitly and we do so in Appendix B. The details may be found there, and we emphasise the importance of carefully considering bounce integrals and $1/\sqrt{1-\lambda\bar{B}}$ factors (as in Rodríguez & Mackenbach 2023 and Rodríguez *et al.* 2024). In the main text, we summarise the resulting expressions in $\epsilon_{\text{eff}}^{3/2} \approx \epsilon_{\text{eff}}^{3/2,(0)} + r^2\epsilon_{\text{eff}}^{3/2,(2)} + O(r^4)$, and focus on their significance.

4.2. Contribution from buffer regions: $\epsilon_{\text{eff}}^{3/2,(0)}$

The leading order contribution in the distance from the magnetic axis r to $\epsilon_{\text{eff}}^{3/2}$ is a constant offset that sets a lower bound to the effective ripple value in the limit of $r \rightarrow 0$. The leading order asymptotic form can be written purely in terms of first- (and zeroth-) order quantities (see the details in [Appendix B](#)),

$$\epsilon_{\text{eff}}^{3/2,(0)} = \frac{\pi}{8\sqrt{2}} \frac{(\bar{R}\bar{B})^2}{(\mathcal{G}^{(1)})^2} \int_{1/B_0^{\text{max}}}^{1/B_0^{\text{min}}} \lambda \frac{(h^{(1)})^2}{I^{(0)}} d\lambda, \quad (4.8)$$

where

$$h^{(1)} = \int_{\varphi_-}^{\varphi_+} \frac{B_0}{\bar{B}} \mathcal{H}(\lambda, B_0) d \sin \alpha_{\text{buf}} d\varphi, \quad (4.9a)$$

$$I^{(0)} = \int_{\varphi_-}^{\varphi_+} \frac{\sqrt{1 - \lambda B_0}}{(B_0/\bar{B})^2} d\varphi, \quad (4.9b)$$

$$(\mathcal{G}^{(1)})^2 = 2 \left(\frac{1}{2\pi} \int_0^{2\pi} d\alpha \int_0^{2\pi/N} d\varphi \frac{\Psi^{(1)}}{B_0} \right)^2 \bigg/ \left(\int_0^{2\pi/N} \frac{d\varphi}{B_0^2} \right), \quad (4.9c)$$

where

$$\Psi^{(1)} = [(X_{1c} \sin \chi - X_{1s} \cos \chi)^2 + (Y_{1c} \sin \chi - Y_{1s} \cos \chi)^2]^{1/2}. \quad (4.10)$$

The expression $\epsilon_{\text{eff}}^{3/2,(0)}$, as can be seen by inspecting $h^{(1)}$, measures the impact that buffer regions have on omnigenity. If such buffers were not present (which may only happen if the rotational transform matches the helicity of the axis, $\iota_0 = M$ (Plunk *et al.* 2019; Camacho Mata & Plunk 2023; Rodríguez *et al.* 2025), then $\alpha_{\text{buf}} = 0$ and $h^{(1)} = 0$, implying $\epsilon_{\text{eff}} \rightarrow 0$ as $r \rightarrow 0$. In general though, they will make a finite contribution, and the expression above allows us to gauge its magnitude. It is a measure of the level of non-omnigenity of the first-order construction. Only if this is sufficiently small is it reasonable to proceed to the next order in the expansion (second order) to evaluate the omnigenity at that order (Rodríguez & Plunk 2023; Rodríguez *et al.* 2025).

The expression for $h^{(1)}$, (4.9a), is then a local measure: trapped particles that do not venture into buffer regions near the trapping well tops, such as deeply trapped ones, will not contribute by any amount. Thus, a smaller buffer results in a smaller fraction of contributing particles. Those trapped particles that do, contribute to $h^{(1)}$ only through their radial magnetic drift (proportional to d) in the buffer region. By construction, and to abide by the condition of pseudosymmetry that avoids losses at turning points of the magnetic field, $d = 0$ is chosen to vanish at both the bottom and the top of the trapping well. Hence, we expect $d \approx 0$ near the trapping well tops and thus in the buffer regions, automatically limiting the magnitude of the effective ripple. To make this somewhat more quantitative, consider the buffer to have a toroidal extent δ about the well top. For a magnetic axis with a curvature κ that has a zero of order v (fixing the form of $d = \bar{d}\kappa$, Plunk *et al.* 2019; Rodríguez & Plunk 2023) and a field B_0 with a first derivative of order $u - 1$, following (4.9a), we expect a scaling $h^{(1)} \propto \delta^{u/2} \times \delta^v \times \delta \sim \delta^{u/2+v+1}$. The first factor comes from the slowing down of the particle speed near the tops; the second comes from the low radial drift linked to d itself; and the latter is the size of the buffer domain. The natural impulse to take




FIGURE 2. Benchmark of effective ripple calculation. The plots show, in log scale, a comparison of the effective ripple as calculated using the code NEO on global equilibria (scatter points) and calculated with the near-axis estimate (solid lines), for a number of different benchmark configurations (see Appendix A). The two detail plots on the right show the individual comparison for two of the cases, including the zeroth-order ripple offset $\epsilon_{\text{eff}}^{(0)}$ as reference (dotted line). The ripple is normalised to a reference $\bar{B} = 1$ T and $\bar{R} = 1$ m.

$\delta \rightarrow 0$ has, though, dire consequences on the shaping of the resulting field (Plunk *et al.* 2019; Camacho Mata *et al.* 2022). One could alternatively increase the order of the power of δ , for instance, by making the top of the trapping well flatter. Without the need to resort to any of these, the behaviour of $I^{(0)}$ helps reduce the buffer contribution even further. This second-adiabatic-invariant-like factor is largest for barely trapped particles, reducing the contribution of the buffer to the effective ripple. This physically corresponds to these particles being more strongly diffused by collisions in λ space, to a large extent because they travel for a longer time, and thus their drift contribution is offset to some degree. All in all, the contribution of the buffer tends to be small (see figure 2), setting a negligible lower bound on ϵ_{eff} in the limit of $r \rightarrow 0$.

4.3. Omnigenity to second order: $\epsilon_{\text{eff}}^{3/2, (2)}$

Given the rather benign contribution to the effective ripple from first order, it is important to assess the next order contribution to ϵ_{eff} . The evaluation of the order r^2 correction to the effective ripple is presented in detail in Appendix B. The calculation is significantly more involved than first order, requiring the next two order corrections of many terms and even involving third-order quantities formally. Many of those terms do not, in practice, contribute significantly, and it is a great convenience to retain only second-order contributions.⁵ For this reason, we write a

⁵One could carry out an analysis similar to that in § 3 to find the third-order components corresponding to the finite build of a second-order finite aspect ratio equilibrium construction, but we shall not do that here.

reduced, simplified form of these contributions (see [Appendix B](#) for a more detailed account of this) that focuses on the next order contribution to H ,

$$\epsilon_{\text{eff}}^{3/2,(2)} = \frac{\pi}{8\sqrt{2}} \frac{(\bar{R}\bar{B})^2}{(\mathcal{G}^{(1)})^2} \int_{1/B_0^{\text{max}}}^{1/B_0^{\text{min}}} \lambda \frac{(h^{(2)})^2}{I^{(0)}} d\lambda, \quad (4.11)$$

where

$$h^{(2)} = -2 \int_{\varphi_-}^{\varphi_+} \mathcal{H}(\lambda, B_0) \frac{\Delta B_{2c}^{\text{OI}}}{\bar{B}} d\varphi, \quad (4.12a)$$

$$\Delta B_{2c}^{\text{OI}} = B_{2c}^{\text{OI}} - \frac{1}{4} \partial_\varphi \left(\frac{B_0^2 d^2}{B_0'} \cos 2\alpha_{\text{buf}} \right), \quad (4.12b)$$

and $B_{2c}^{\text{OI}} = -(B_{2c} \cos 2\bar{\iota}\varphi + B_{2s} \sin 2\bar{\iota}\varphi)$. The expression in (4.12b) is equivalent to the omnigenity condition at second order when it vanishes, as derived explicitly by Rodríguez & Plunk (2023 (32c)). Thus, this approximation to the asymptotic $O(r^2)$ behaviour of the effective ripple is driven by the second-order non-omnigenous behaviour and will generally dominate $\epsilon_{\text{eff}}^{3/2,(0)}$ at a finite aspect ratio.

Unlike to leading order, in this case, the non-omnigenous contribution is spread over the entirety of the trapping well, wherever $\Delta B_{2c}^{\text{OI}} \neq 0$, a local measure of radial drift imbalance. An imbalance at some point φ will affect all trapped particles with $\lambda < 1/B_0(\varphi)$ that pass over this point. Although they feel such deviation from omnigenity, this is not to say that they will necessarily behave in a non-omnigenous fashion, as the bounce-integral in (4.12a) can lead to partial cancellations of $\Delta B_{2c}^{\text{OI}}$, which can have either sign. However, it remains true that if $\Delta B_{2c}^{\text{OI}} \neq 0$, then $\epsilon_{\text{eff}}^{3/2,(2)} \neq 0$.

The abovementioned scenario then describes the significance of the omnigenity condition of Rodríguez & Plunk (2023), which is a property of second-order fields. As such, different choices of second-order shaping (namely, triangularity and Shafranov shift) will affect the omnigenous behaviour of the field and their choice becomes key in constructing the stellarator field. This is not to say that the first-order near-axis construction is no longer important, as it will remain to strongly affect the amount and form of second-order shaping needed to make a field more omnigenous. A clear example of this is the radial drift involved in (4.12b), which by increasing axis curvature (and thus d for a controlled elongation of flux surfaces) will generally require stronger shaping.

One may ask how can the shape at second order be chosen to omnigenise a field and if it is an example of an optimisation problem. As shown by Rodríguez *et al.* (2025, (3.6)–(3.7)), there is however a unique, closed form choice of shaping that exactly achieves $\Delta B_{2c}^{\text{OI}} = 0$. This excludes, though, the neighbourhoods (in φ) of flattening points of the magnetic axis, where the shaping necessary to achieve this ideal omnigenous behaviour would diverge unless the first-order construction was specially chosen. Excluding regions around flattening points (see further discussion later) and applying this choice of shaping, we say we have ‘omnigenised’ the construction.




FIGURE 3. Diagram with different contributions to ϵ_{eff} . (a) Deformation of $|B|$ within the principal well violating the equal radial drift condition of omnigenity (depicted by broken line). (b) Appearance of local ripple or secondary wells. (c) Misalignment of field maxima, leading to multiple-well trapped particles.

4.4. Numerical implementation and benchmark

Let us first benchmark the near-axis estimate of the abovementioned effective ripple. To that end, we take an extended set of second-order near-axis equilibria previously used by Rodriguez *et al.* (2025), described in Appendix A, and their associated global equilibria at a number of different aspect ratios. We then compare the near-axis estimate of ϵ_{eff} (as a function of r), implemented in the pyQIC framework,⁶ with the effective ripple computed by the neoclassical code NEO of the finite aspect ratio equilibria. The comparison is presented in figure 2 as a function of aspect ratio $A = R/r$.

The agreement between the predicted near-axis ϵ_{eff} and the finite volume equilibria calculation, although not exact, is excellent over quite a large range of A . The buffer region contribution in all cases sets a rather benign lower bound to the value of ϵ_{eff} at large aspect ratios and seems to be in agreement with the global calculation. This shows the critical role played by the second order in the expansion, which controls the dominant $O(r^{4/3})$ behaviour. At lower aspect ratios, $A \lesssim 10$, the departures grow, as the deformation of the field increases and departs from its asymptotic description (see figure 3). The latter is unable to capture local ripples⁷ and the appearance of new trapped particle classes (as in figure 3b), or misalignment of maxima (as in figure 3c) with the associated abrupt modification of trapped particle behaviour. These effects become increasingly prominent at lower aspect ratios, and thus additional deviations between the near-axis and global calculation are to be expected.

4.5. Measures of omnigenity

With the benchmark in place, we now introduce different measures based on the effective ripple to assess near-axis fields. We propose measures to characterise the effective ripple of second-order constructions, first-order constructions and the shaping involved when the latter are omnigenised at second order.

We start with the fundamental form of the effective ripple.

⁶The implementation makes use of the newer version of the BAD code for bounce averaging.

⁷We do not consider the presence of θ -independent sub-wells (Parra *et al.* 2015), and assume that any ripple will be localised in θ , and thus will be non-omnigenous.

- (a) *Effective ripple from buffer regions*, $\epsilon_{\text{eff}}^{3/2,(0)}$.

Leading order contribution to the effective ripple from the buffer regions, see (4.8). Its calculation requires information of first order and the evaluation of bounce integrals between bounce points defined by $B_0(\varphi)$.

- (b) *Second-order effective ripple*, $\epsilon_{\text{eff}}^{3/2,(2)}$.

Approximate, order $O(r^{4/3})$ contribution to the effective ripple due to omnigeneity breaking at second order, see (4.11). It requires information on the second-order construction and once again the evaluation of bounce integrals with bounce points defined by $B_0(\varphi)$.

These combine into a single measure in item (c).

- (c) *Effective ripple at the edge*, $\epsilon_{\text{eff}}^{\text{edge}}$.

Estimation of the effective ripple value at the flux surface corresponding to a reference aspect ratio of $A_{\text{ref}} = 10$ evaluated using the asymptotic expressions

$$\epsilon_{\text{eff}}^{\text{edge}} = \left(\epsilon_{\text{eff}}^{3/2,(0)} + r_{\text{ref}}^2 \epsilon_{\text{eff}}^{3/2,(2)} \right)^{2/3}, \quad (4.13)$$

with $r_{\text{ref}} = R/A_{\text{ref}}$ and R is the major radius calculated using the magnetic axis length, L , as $R = L/2\pi$.

The $\epsilon_{\text{eff}}^{\text{edge}}$ is a dimensionless measure that can be directly interpreted as one would the effective ripple of an optimised stellarator. This is a measure that diagnoses the effective ripple of a fully consistent second-order near-axis construction. However, as discussed previously in this section and by Rodríguez *et al.* (2025), there are certain aspects of the field that are a result of lower-order choices in the near-axis construction. Thus, we would like to have some way to assess this ‘intrinsic’ behaviour. The simplest measure is item (d).

- (d) *Non-omnigeneity mismatch at the well bottom*,⁸ $\Delta B_{\text{min}}^{\text{QI}}$.

Define

$$\Delta B_{\text{min}}^{\text{QI}} = r_{\text{ref}}^2 \left. \frac{\Delta B_{2c}^{\text{QI}}}{\bar{B}} \right|_{\text{min}} \quad (4.14)$$

evaluated at the bottom of the trapping well. Here, $\bar{B}$ is the average value of B_0 and r_{ref} a reference radial value (which we compute for $A_{\text{ref}} = 10$). The measure provides a sense of omnigeneity breaking near the bottom of the trapping well normalised to the average field strength. It only requires first-order quantities, following (3.8) of Rodríguez *et al.* (2025).

To translate this relative breaking of omnigeneity into the language of the effective ripple, we can evaluate $\epsilon_{\text{eff}}^{\text{edge}}$ for a near-axis field constructed using the omnigeneising shaping at second order. That is, the shaping that forces $\Delta B_{2c}^{\text{QI}} = 0$ for all φ , barring some masked region near the flattening points. The shaping needed to achieve this is given by Rodríguez *et al.* (2025, (3.6)–(3.7)). The remaining ripple value is the result of the contribution from the straight sections, we denote as item (e).

⁸We could define an equivalent measure at the top of the well. However, the behaviour at the minimum is of greater importance as all trapped particles experience it.

(e) *Effective ripple of omnigenised field, $\epsilon_{\text{eff}}^{\text{shap}}$.*

Estimate of the effective ripple of an omnigenised near-axis field at r_{ref} . It is defined as $\epsilon_{\text{eff}}^{\text{edge}}$ for a field in which the second-order shaping has been chosen to enforce $\Delta B_{2c}^{\text{OI}} = 0$ everywhere except a 15 % of the toroidal domain around the inflexion points. This evaluation requires solving the second-order near-axis equations with the appropriate omnigenising shaping and computing the appropriate bounce integrals.

To quantify how much shaping is required to construct this omnigenised field, we define item (f).

(f) *Average omnigenising shaping, $\hat{T}^{\text{shap}}$.*

Averaged measure of the amount of shaping introduced at second order. We define

$$\hat{T}^2 = \frac{r_{\text{ref}}^2}{2\pi} \int_0^{2\pi} \left(X_{20}^2 + Y_{20}^2 + Z_{20}^2 + \frac{X_{2c}^2 + X_{2s}^2 + Y_{2c}^2 + Y_{2s}^2 + Z_{2c}^2 + Z_{2s}^2}{2} \right) d\varphi. \quad (4.15)$$

which we have non-dimensionalised with respect to the minor radius r_{ref} , indicating the relative deformation of the elliptical cross-sections needed to omnigenise a field. We define $\hat{T}^{\text{shap}}$ as the shaping value of the omnigenised field at an aspect ratio $A_{\text{ref}} = 10 = R/r_{\text{ref}}$. Note that $\hat{T} = \hat{T}[X_{2c}, X_{2s}]$ may be considered a functional of the two free shaping functions at second order.

The average measure $\hat{T}$ quantifies second-order shaping in absolute terms, but makes no reference to the first-order shape that it modifies. To elucidate how shaped a configuration is, it is important to consider the interaction with first-order shaping at finite radius r . To capture this, we introduce the critical radius r_c (or its reciprocal, the critical aspect ratio $A_c \sim R/r_c$) defined by Landreman (2021). This measure captures the complexity of flux surfaces by reflecting the maximal radial extent at which flux surfaces continue to exist without any unphysical intersection, i.e. before the Jacobian of the coordinate system vanishes. That way, low A_c indicates that the second-order shaping is compatible with lower first-order shaping for relatively compact fields.

(g) *Minimal aspect ratio of omnigenised field, A_c^{shap} .*

Minimum value of the aspect ratio A for which the near-axis construction of the omnigenised field does not present locally self-intersecting flux surfaces. This is based on the definition of Landreman (2021),

$$A_c = R / \min [r \mid \exists \theta, \varphi : \mathcal{J}(r, \theta, \varphi) = 0], \quad (4.16)$$

where $\mathcal{J}$ is the Jacobian of the $\{r, \theta, \varphi\}$ coordinate system.

We gather the abovementioned measures evaluated on the configurations in the benchmark in table 1. The first figure to look at is the effective ripple at the edge, which provides a single measure to contextualise these near-axis constructions amongst the space of optimised stellarators. Configurations with values below a per cent are often regarded as adequate with regards to neoclassical behaviour. It is clear that the second-order ($O(r^{4/3})$) contribution dominates the behaviour of ϵ_{eff} . To understand how much of the behaviour comes from the particular choice

Configs	4.1	4.2	4.3	4.4 (unopt)	4.4 (opt)	# 50	# 76
$\epsilon_{\text{eff}}^{3/2,(0)} \times 10^5$	0.01	1.75	0.63	0.16	0.08	0.33	0.52
$\epsilon_{\text{eff}}^{3/2,(2)} [\text{m}^{-2}]$	0.290	3.079	0.643	0.790	0.313	3.829	0.263
$\epsilon_{\text{eff}}^{\text{edge}} [\%]$	2.46	13.5	3.46	5.47	3.10	11.4	1.91
$\epsilon_{\text{eff}}^{\text{shap}} [\%]$	2.41	3.20	4.12	2.67	0.059	7.74	0.063
$\Delta B_{\text{min}}^{\text{QI}} [\%]$	5.78	0.94	2.53	0.94	0.08	4.17	0.21
A_c	2.09	5.84	2.38	4.46	4.93	3.00	1.59
A_c^{shap}	19.8	22.0	10.9	24.0	19.0	11.2	12.6
$\hat{T}^{\text{shap}}$	2.27	0.78	0.30	0.75	0.79	0.24	0.30

TABLE 1. Shaping configurations for effective ripple. The table shows information regarding the omnigenous nature of the near-axis constructions in the benchmark. The table is separated into two main parts. Top rows show the measures of omnigenicity, in particular, the effective ripple for the original second order field in the benchmark (top), and the omnigenised form of the field (bottom). The lower rows present information regarding the shaping of the original field (top) and the omnigenised field. All the measures are defined in the main text.

at second order and how much is intrinsic to the lower orders, we then look at the properties of the omnigenised field.

The lack of omnigenicity near the inflexion points is manifest in $\epsilon_{\text{eff}}^{\text{shap}}$, which does not approach the baseline value $\epsilon_{\text{eff}}^{(0)}$ for any of the configurations. This is also indicated by the non-zero values of $\Delta B_{\text{min}}^{\text{QI}}$. However, the optimised version of config. 4.4 (by construction) and # 76, both show reduced ripple. What might appear most surprising, though, is that not all omnigenised configurations exhibit smaller ripple than their original forms, that is, $\epsilon_{\text{eff}}^{\text{shap}} > \epsilon_{\text{eff}}^{\text{edge}}$ in some cases. This is a result of a strong non-omnigenous drive near the bottom of the well, which, in the original configuration, partially cancelled the $\Delta B_{2c}^{\text{QI}}$ away from it.

Figure 4 shows in more detail what the act of omnigenising is, presenting the case of configuration # 76. We choose this case because it has an intrinsic near-omnigenous behaviour (small $\Delta B_{\text{min}}^{\text{QI}}$) at the bottom of its trapping well, and thus nicely illustrates the omnigenising effect of shaping. We observe that moderate variations in the cross-sections lead to significant changes in the ripple. The plot in panel (a) shows how gradually increasing the shaping to its perfect omnigenising value (0 meaning unshaped and 1 the completely shaped case) reduces the ripple, while the changes in the surfaces lead to a significant increase in A_c . This reflects a general tendency observed with near-axis QI constructions: second-order shaping to improve omnigenicity is often costly in terms of shaping. Freedom at first order must also be used to effectively construct omnigenous fields.

5. Appearance of secondary trapping wells: A_w

In the previous section, we focused on assessing omnigenicity in terms of the asymptotic behaviour of ϵ_{eff} . However, we noted that this did not account for the appearance of small localised secondary wells, which we now consider (see figure 3b).




FIGURE 4. Evolution of effective ripple with omnigenising second-order shaping. The plots show the evolution of the half-helicity benchmark configuration #76 with changing second-order shaping. (a) Evolution of $\epsilon_{\text{eff}}^{\text{edge}}$ (solid) and A_c (broken) as a function of shaping. A value of 0 for the scaled shape corresponds to the original second-order benchmark configuration, while a value of 1 indicates the omnigenised construction. The shaping is scaled linearly in between. (b) Examples of cross-sections in cylindrical coordinates at three values of shaping, indicated as vertical lines on the left plot.

The formation of these defects is linked to increased neoclassical transport (Ho & Kulsrud 1987) and particle losses (Mynick 1983; Paul *et al.* 2022). It is simple to picture a newly formed class of ripple-trapped particles, which live on those local ripples and thus experience whichever the local non-zero radial drift is. A deeper well results in more trapped particles, and the further from the bottom and top along the trapping well such ripple is, the stronger the local drift is. Having a sense of the presence of such ripples is then important.

5.1. Construction of a ripple well measure, A_w

How do these secondary wells appear in the approximately QI near-axis scenario? By construction choice, in the limit of an infinite aspect ratio equilibrium ($r \rightarrow 0$), there is a single well defined by B_0 and thus no trace of any secondary wells. The ripple well question is then a finite aspect ratio one, which will require evaluating the asymptotic description at a finite A . We define A_w as the largest aspect ratio A for which a ripple well first appears along a magnetic field line within our asymptotic near-axis field model. Trapping wells are identified seeking points at which $\partial_\varphi|_\alpha B = 0$ (i.e. partial toroidal derivative keeping the field line label constant, namely $\partial_\varphi|_\alpha B = (\partial_\varphi + \iota \partial_\theta)B$). To distinguish ripple wells from the main trapping well defined by $B_0(\varphi)$ on axis, we shall picture the appearance of secondary wells as a ‘dynamic’ action in radius. As we look at surfaces of decreasing aspect ratio, ripples may start forming but must do so by first forming inflexion points, $\partial_\varphi|_\alpha^2 B = 0$. We therefore define the ripple-well aspect ratio,

$$A_w = R / \min \{ r \mid \exists \theta, \varphi : \partial_\varphi|_\alpha B(r, \theta, \varphi) = 0, \partial_\varphi|_\alpha^2 B(r, \theta, \varphi) = 0 \}. \quad (5.1)$$

In practice, finding A_w requires a procedure (detailed in Appendix C) similar to the calculation of the critical radius r_c introduced previously and originally by Landreman (2021). Briefly put, the two conditions in the definition of (5.1) written in terms of $B(r, \theta, \varphi) = \sum_{n=0}^2 r^n B_n(\theta, \varphi)$ may be interpreted as simultaneous algebraic




FIGURE 5. Ripple well diagnostic measure $\hat{A}_w$ for configurations 4.1 and 4.3. The plots show (left) the function $1/\hat{A}_w$ as a function of α , the field line label, for configs. (a) 4.1 and (b) 4.3, the spaghetti diagrams. The hatched area represents $\hat{A}_w > A_w$, and the shaded regions the interval of α that have a ripple well. The right plots show the $|\mathbf{B}|$ contours corresponding to the r values indicated by the horizontal silver lines on the left plot. The solid black line in the contour plots shows the direction of a magnetic field line, the broken white line contours of $\partial_\varphi|_\alpha \mathbf{B} = 0$ and the scatter, inflexions captured by the spaghetti diagram.

equations on r and θ , which may be solved for every φ . The resulting multiple roots may then be written as a multi-valued function $\hat{A}_w(\alpha)$ (eliminating φ), of which the maximum is A_w .

5.2. Assessment and application

Let us consider the behaviour of ripple wells in two of the example configurations in the benchmark (see [Appendix A](#)) and plot in [figure 5](#) the multi-valued function $1/\hat{A}_w$. For illustration purposes, we accompany these plots with the corresponding B contours.

The plots of $1/\hat{A}_w$, which we refer to as spaghetti diagrams (borrowing from the condensed matter lingo), are best interpreted from bottom to top. In the interval $0 < 1/\hat{A}_w < 1/A_w$, the field does not present any secondary well, even though it may present topological defects in the contours of $|\mathbf{B}|$, i.e. puddles (Rodríguez & Plunk 2023). At $\hat{A}_w = A_w$, a band is reached; that is, an inflexion point appears

Configs	4.1	4.2	4.3	4.4 (unopt)	4.4 (opt)	# 50	# 76
A_w	27.4	37.0	12.9	32.8	24.4	18.5	16.5
A_c	2.09	5.84	2.38	4.46	4.93	3.00	1.59

TABLE 2. Ripple well distance in the benchmark configurations. The table shows the values of the aspect ratio when the first ripple-well appears, A_w , and that at which the near-axis construction breaks down, A_c , for the configurations used as a benchmark in the paper.

along a field line for the first time. By symmetry, these inflexions occur in pairs. Pairs of bifurcating branches appear and separate with increasing $1/\hat{A}_w$. The shaded areas these curves bound represent every field line (i.e. intervals of α) containing secondary wells. This area grows at lower aspect ratios and additional bifurcations do appear. We must mention the unpaired bifurcations that appear near $\alpha/\pi \approx 1$ of panel (a), which appear to contravene the general behaviour presented previously, as they do not bound field lines with ripples. In fact, field lines in between have extended trapping wells. At least, until these branches cross.

With all this into account, at a finite $A < A_w$, we expect more ripples in config. 4.1 as compared with 4.3. What the spaghetti diagram fails to indicate is the location (in φ) of such ripples and thus how strong is the radial drift experienced by ripple-trapped particles. To fully assess the relevance of ripples, some of that information should be taken into account, but we shall not explore this further here, and leave the relationship between A_w and the level of departure of the near-axis approximation to ϵ_{eff} (see previous section) for future investigation.

The comparison of different configurations can be made more quantitative by looking at A_w , which we present in table 2. The ripple measure A_w exhibits some correlation with the shaping measure A_c (they have a Pearson correlation of $\rho = 0.75$, $p - \text{val} = 0.05$). The more shaping means the more variation in the field, the larger the second-order field terms and thus the larger A_w . However, this is far from being a one-to-one relationship and A_w merits its own use as a field diagnostic. Minimising A_w leads to ‘cleaner’ $|\mathbf{B}|$ contours at finite aspect ratios. From this brief benchmark, it appears that half-helicity configurations (as well as the second-order QI optimised one) are more resilient to the appearance of wells. There are other elements in the near-axis equilibrium that have the potential to significantly influence the appearance of ripple wells, one being the shape of B_0 . We would expect configurations with flatter minima, as those favouring MHD stability and maximum- $\mathcal{J}$ (Plunk *et al.* 2024; Rodríguez *et al.* 2024), to be more susceptible to the appearance of ripples.

6. Magnetic well stability and shaping: W and $\mathcal{G}^W$

There are crucial physical features of the field beyond omnigenicity that also depend on the second-order form of the field. In particular, we are interested in MHD stability, which we capture in this paper through the so-called magnetic well criterion (Greene 1997): a field is deemed unstable (in particular, to interchange instability in the low plasma β limit) if $V'' > 0$, where V is the volume enclosed by flux surfaces and the primes denote derivatives with respect to ψ . In the context of

the near-axis framework, V'' can be written in simple form as (Landreman & Jorge 2020, (3.6))

$$V'' = 2\pi \left| \frac{G_0}{\bar{B}} \right| \int_0^{2\pi} \frac{d\varphi}{B_0^4} \left[3(B_{1s}^2 + B_{1c}^2) - 4B_0 B_{20} - \frac{\mu_0 p_2 B_0^2}{\pi} \int_0^{2\pi} \frac{d\varphi'}{B_0(\varphi')^2} \right] + O(r^2), \quad (6.1)$$

where symbols have their standard near-axis meaning (Landreman & Sengupta 2019; Rodríguez *et al.* 2025). The key ingredient here is B_{20} , the poloidally averaged ‘radial’ gradient of $|\mathbf{B}|$. A positive value is beneficial to stability as it tends to make field-lines curve outwards (from force balance, $\nabla_\perp(B^2/2) = B^2\kappa - \mu_0\nabla p$ for κ the curvature of field lines). The field component B_{20} controls other physics such as precession and the maximum- $\mathcal{J}$ property. We do not delve into this here, but refer the reader to Rodríguez *et al.* (2024) for an in-depth discussion and several useful measures.

Values of V'' are often reported as a relative well measure, given in the form of a percentage (Solov’ev & Shafranov 1970, (14.29)). Following this convention, we define the relative well of a finite aspect ratio equilibrium as $W = \int_0^{\psi_a} V'' d\psi / V'(0)$, where ψ_a is the edge value of ψ . Such a measure might in general be misleading, as the magnetic well criterion is a locally radial one. In the present case of the near-axis expansion, where V'' is constant to leading order, though, W is merely a normalisation choice for V'' . Choosing as a reference $r_{\text{ref}} = R/A_{\text{ref}}$ (unless otherwise stated, with $A_{\text{ref}} \sim 10$), in the near-axis framework, we define

$$W = \frac{\bar{B} r_{\text{ref}}^2}{4\pi G_0} \left[\int_0^{2\pi} \frac{d\varphi}{B_0^2} \right]^{-1} V''. \quad (6.2)$$

With such a measure, we may assess the stability of a prescribed second-order field. We expect the stability to depend both on shaping choices such as triangularity and Shafranov shift (Freidberg 2014, (12.89)), second-order quantities, but also first-order ones. In fact, following the works of Rodríguez *et al.* (2024, 2025), we can look near the straight sections of the field, where second-order shaping has no effect, to study the intrinsic stability contribution of the first order. In particular, the function B_{20} will have some prescribed value there.

(a) *Relative radial gradient at the trapping well bottom, $B_{20,\text{min}}$.*

We define $B_{20,\text{min}}$ as

$$B_{20,\text{min}} = r_{\text{ref}}^2 \left. \frac{B_{20}}{B_0} \right|_{\text{min}} \quad (6.3)$$

at the bottom of the trapping well. This is the poloidal-averaged radial gradient at the bottom of the well, which controls the poloidal precession of deeply trapped particles and contributes to stability if positive (Plunk *et al.* 2024; Rodríguez & Mackenbach 2023; Rodríguez *et al.* 2024). This is an intrinsic feature of the first-order construction. The measure is normalised to serve as a relative field strength change at a reference surface of aspect ratio A_{ref} .

Away from these special points, it is possible to choose the second-order shaping to make the field MHD stable, but how should this shape be chosen and how much does that shaping depend on the lower orders?

6.1. Sensitivity of vacuum well

To determine the sensitivity of the magnetic well to the shaping at second order, we compute the variation of the magnetic well defined in (6.1) with respect to the input functions X_{2c} and X_{2s} on which the integrand depends. Following the general prescription detailed in Appendix E for computing the shape gradient $\mathcal{G}^S = (\delta S/\delta X_{2c}, \delta S/\delta X_{2s})$ of a general functional of the form

$$S = \int_0^{2\pi} f(B_{20}, B_{2c}, B_{2s}, \varphi) d\varphi, \quad (6.4)$$

we may compute the shape gradient associated with the magnetic well, $\mathcal{G}^W$. For the magnetic well, we need

$$\frac{\partial f}{\partial B_{20}} = -\frac{2r_{\text{ref}}^2}{B_0^3} \left[\int_0^{2\pi} \frac{d\varphi}{B_0^2} \right]^{-1}, \quad \frac{\partial f}{\partial B_{2c}} = 0, \quad \frac{\partial f}{\partial B_{2s}} = 0, \quad (6.5)$$

which may be used to compute explicitly the gradient, as in Appendix E.

Although explicit, the resulting gradient involves the inversion of a differential linear operator (see Appendix D). This is a result of B_{20} being part of a self-consistent equilibrium solution. The necessary calculation to find the gradient is then computationally equivalent to solving a modified, second-order near-axis equilibrium. The strength of this gradient calculation is that, because the magnetic well is linear on B_{20} , and thus also on the shaping, the gradient is independent of second order. Once we have computed the gradient, we know precisely how the magnetic well will change under changes in the shaping.

Before exploiting this knowledge of the gradient to concoct some diagnosing measures of the fields, we can gain some perspective on $\mathcal{G}^W$ by considering the simplest limit we can: that of an up-down symmetric tokamak field.⁹ In this limit, any toroidal derivative drops out and (see details in Appendix E.1)

$$\mathcal{G}_{2c}^W = \frac{3r_{\text{ref}}^2 \bar{d}^4 - 1}{\pi R \bar{d}^4 - 3}, \quad \mathcal{G}_{2s}^W = 0, \quad (6.6)$$

where $\bar{d} = 1$ corresponds to a tokamak with circular cross-sections.¹⁰ For the case of circular cross-sections, the magnetic well is insensitive to changes in shaping (triangularity) (Rodríguez 2023). For a vertically elongated cross-section ($\bar{d} < 1$), an increase in X_{2c} (i.e. of positive triangularity) favours the vacuum well, matching the well-known tokamak intuition (Freidberg 2014; Rodríguez 2023). The divergence as $\bar{d}^4 \rightarrow 3$ is worrisome, as it leads to ill-behaviour for a very particular horizontally elongated elliptically shaped tokamak. Such behaviour was previously described and analysed by Rodríguez *et al.* (2023) in the context of quasisymmetric stellarators, where the generalisation of the above observations is true. The divergence is indicative of only one particular form of shaping X_{2c} being physical and consistent with

⁹The considerations presented here hold beyond the context of fields with poloidally closed $|\mathbf{B}|$ contours. They extend to any general magnetohydrostatic equilibrium near-axis field with the appropriate adjustments.

¹⁰ It is common to use $\bar{\eta}$ instead of $\bar{d}$ in the tokamak and quasisymmetric literature (Garren & Boozer 1991a; Landreman & Sengupta 2019), but we here try to remain consistent with the rest of the notation in the paper.

how the construction is being carried out. This unphysicality can be seen on a similar divergence of triangularity, (D11). Fortunately, configurations of interest tend to live away from this singularity.

Having a knowledge of the gradient $\mathcal{G}^W$ provides complete insight into the sensitivity of the magnetic well of a configuration. In particular, we are interested in knowing what is the minimal amount of shaping that would make a given first-order near-axis configuration achieve a magnetic well (mirroring the practical approach of Plunk *et al.* (2024)). Because we have computed the shape gradient, we know precisely which shaping combinations will make a configuration stable: we simply need to choose X_{2c} and X_{2s} such that

$$\mathcal{C} = \int_0^{2\pi} (\mathcal{G}_{2c}^W X_{2c} + \mathcal{G}_{2s}^W X_{2s}) \, d\varphi + W_{\text{ref}} = 0, \quad (6.7)$$

where W_{ref} is the magnetic well value of the second-order near-axis construction for $X_{2c} = 0 = X_{2s}$. This imposes marginal stability, and in the rare case of $W_{\text{ref}} > 0$, it would require relaxing the stability of the field. The ‘problem’, though, is that there is no unique way of choosing the shaping to satisfy this constraint. We must then impose some regularising choice, which we take to be one that minimises shaping in the form of an integral over the sum of squares of all the second-order shaping functions, as defined in (4.15).¹¹ The problem of finding the minimal stabilising shaping can then be formulated as one of finding X_{2c} and X_{2s} that minimise the following constrained variational problem:

$$T^2[X_{2c}, X_{2s}] = \hat{T}^2[X_{2c}, X_{2s}] - \lambda \mathcal{C}[X_{2c}, X_{2s}], \quad (6.8)$$

where λ is a Lagrange multiplier. We denote the resulting shaping as X_{2c}^{mhd} and X_{2s}^{mhd} , for which explicit expressions may be found in Appendix F. These expressions involve the shape gradients computed previously, as well as some additional matrix multiplications, and could in principle be studied analytically. The main power of the approach is however providing a prescription to construct marginally stable approximately QI near-axis fields given a choice of axis, magnetic field strength on axis and first-order information.

To assess such stabilised constructions, we introduce some simple measures. One of the simplest is the following.

(b) *Average second-order shaping of stabilised field, $\hat{T}^{\text{mhd}}$.*

The root-mean-square value of the consistent stabilised second-order shaping, evaluated using (4.15) using the minimal, stabilising choice of shaping, $\hat{T}^{\text{mhd}} = \hat{T}[X_{2c}^{\text{mhd}}, X_{2s}^{\text{mhd}}]$. It is therefore a simple measure of the amount of shaping necessary to stabilise a field, normalised to some reference r_{ref} .

A more accurate measure of the shaping that takes toroidal variation into consideration may also be introduced as follows.

¹¹This is a natural regularising choice, L^2 norm, but others would also be possible (especially those that include penalties penalising large toroidal variations). Simpler forms could also be constructed only minimising the choice of input functions, which does however capture the notion of second-order shaping less accurately.

Configs	4.1	4.2	4.3	4.4 (unopt)	4.4 (opt)	# 50	# 76
$B_{20,\min}$ [%]	−2.15	0.80	−2.10	−1.37	−0.99	−5.27	−0.52
W [%]	5.78	5.56	4.52	9.12	8.52	6.64	4.12
$ W_{\text{mhd}} \times 10^{-16}$	1.53	1.39	2.10	1.78	2.97	5.00	0.48
A_c	2.09	5.84	2.38	4.46	4.93	3.00	1.59
A_c^{mhd}	5.17	6.10	4.70	6.43	6.23	5.91	4.82
$\hat{T}^{\text{mhd}}$	0.20	0.17	0.16	0.22	0.19	0.22	0.13
$ \Delta W/\Delta A_c $ [%]	1.30	8.48	1.25	2.43	3.29	1.37	1.20

TABLE 3. Magnetic well sensitivity of benchmark configurations. The table shows the values of the magnetic well and re-shaped fields for the benchmark configurations, the shape $\hat{T}$ and r_c^{mhd} second-order shaping measures.

(c) *Critical aspect ratio of stabilised field, A_c^{mhd} .*

The minimum aspect ratio for which the near-axis construction of the stabilised configuration is physical (non-intersecting flux-surfaces). That is, A_c for the near-axis field evaluated with X_{2c}^{mhd} and X_{2s}^{mhd} . It is a dimensionless quantity.

Finally, to get a sense of the sensitivity of the shaping and the magnetic well, we also define $\Delta W/\Delta A_c$ as follows.

(d) *Magnetic well sensitivity, $\Delta W/\Delta A_c$.*

Change in the magnetic well W with the change of shaping as measured by A_c around the marginally stable point. Here, we define it by finite differencing as

$$\frac{\Delta W}{\Delta A_c} = \frac{W(V'' = 1 \text{ T}^{-2} \text{ m}^{-1})}{A_c(V'' = 1 \text{ T}^{-2} \text{ m}^{-1}) - A_c(V'' = 0)}, \quad (6.9)$$

a dimensionless quantity. A larger value indicates increased sensitivity to shaping. We leave a more precise (perhaps analytic) form of this quantity to future work.

6.2. Numerical implementation and benchmark

We apply the abovementioned measures to the configurations that constitute the benchmark for this paper. We compute the shape gradient for all such configurations, compute their stabilised versions and evaluate the scalar measures presented previously. A summary of the relevant vacuum-well-related properties is gathered in table 3. We start by noting that all configurations in the benchmark are MHD unstable, which appears to be the overwhelming tendency with constructed QI fields, as is the case with the other classes of omnigeneous fields (Landreman & Jorge 2020; Landreman 2022; Rodríguez *et al.* 2023). All examples encountered require some concerted shaping to push them towards MHD stability, which we illustrate in figure 6. The correctness of the approach is evidenced in the vanishing of the magnetic well machine precision for the reshaped configurations, W_{mhd} , in table 3.

Beyond the qualitative assessment of the shaping needed from the cross-sections, both the measures $\hat{T}$ and A_c^{mhd} quantitatively describe the amount of shaping, with the simpler form correlating with the latter. Whenever the re-shaped configuration




FIGURE 6. Stabilised configurations and sensitivity of vacuum magnetic well to shaping. (a,b) Plots show (left) the shaping input for stabilising the near-axis fields and (right) the resulting re-shaped cross-sections for the first three configurations of the benchmark. The black, solid lines represent the re-shaped configuration, while the dotted one denotes the starting point. (c) Change in the magnetic well as a function of the shaping measure A_c , illustrating the sensitivity of the fields to shaping.

has a small value for the critical radius, it means that significant shaping is present and if there is a large difference from the original benchmark field construction, A_c in table 3, we can conclude that the magnetic well is quite insensitive to shaping. Such sensitivity is more precisely measured by $\Delta W/\Delta A_c$. Note its large value in the case of config. 4.2, figure 6(b), where minor changes in the shape of the cross-sections are enough to stabilise the field (see figure 6c). This configuration is also special in that $B_{20,\min} > 0$, as we may expect from the help of having a finite plasma beta in that case. The high sensitivity is a potentially problematic feature of the field if small changes in its shaping can change the response significantly. Note however that this sensitivity does change with W (see figure 6c) so that a given configuration could perhaps be made less sensitive if sufficiently stabilised. This merits further study.

If one was interested in finding the most MHD stable configurations, these measures could be used as a target of optimisation, which could be helped by the exploitation of ideas from Plunk *et al.* (2024) and Rodríguez *et al.* (2024). Before moving on, we note that this notion of MHD stability is only a minimal one, but not necessarily sufficient to achieve true MHD stability. Behaviour such as localised ballooning modes will react differently to the field geometry. In addition, this stabilising reshaping necessarily causes other second-order behaviours to change, and the ensuing trade-offs are an important venue of future work.

7. β -sensitivity of Shafranov shift

So far in this paper, we have made little explicit reference to plasma pressure. As we change the plasma pressure, though, we expect the magnetic field to change,

as the Lorentz forces should continuously balance pressure gradients. One of the most notable consequences is the change in the Shafranov shift, the relative rigid displacement of nested flux surfaces with respect to each other. To describe this shift, we must define a reference point, a ‘centre’, for each cross-section at different radii, and define the Shafranov shift to be the ψ -derivative of the position of that mid-point. In the context of the near-axis description, the Shafranov shift can be defined as (Rodríguez 2023, (3.5); Landreman 2021)

$$\Delta = \begin{pmatrix} \Delta_x \\ \Delta_y \end{pmatrix} = \begin{pmatrix} X_{20} + X_{2c} \\ Y_{20} + Y_{2c} \end{pmatrix}, \quad (7.1)$$

where the displacements are in the normal (x) or binormal (y) direction as defined by the signed Frenet–Serret frame of the magnetic axis. This expression matches the known definition in the tokamak limit, but one should bear in mind that there is no unique way in which to define the ‘centre’ of a non-elliptical cross-section.

As the pressure gradient increases, so does the Shafranov shift, which results from the plasma pushing against the magnetic field (the poloidal field and hence ι , in a tokamak). Formally, as plasma pressure increases, so does B_{20} linearly, and with it, the rigid displacement of surfaces X_{20} , eventually leading to two consecutive flux surfaces to touch (see figure 9). This breakdown limits the existence of the equilibrium and thus it is important to gauge how sensitive a field is to changes in pressure. In the context of the near-axis expansion, the sensitivity of a field to changes in the pressure is most naturally formulated keeping the shape of the magnetic axis fixed. This is unlike the scenario of how an actual stellarator might be operated experimentally, namely by fixing some given coil currents (and thus a vacuum field), and letting plasma pressure grow. Modelling the latter, which generally will require the magnetic axis shape to vary with β , would involve mathematically separating the vacuum from the plasma field and performing a double expansion, and we do not do that here.

7.1. Shafranov shift shape gradient, $\mathcal{S}$

The variations problem is now one of $\delta\Delta$, a change in the shift, with δp_2 , a change in the pressure gradient. As a result of the change in the plasma pressure, from the definition in (7.1),

$$\delta\Delta = \underbrace{\begin{pmatrix} \delta X_{20} \\ \delta Y_{20} \end{pmatrix}}_{\text{Rigid shift}} + \underbrace{\begin{pmatrix} 0 \\ \delta Y_{2c} \end{pmatrix}}_{\text{Triangular shaping}}, \quad (7.2)$$

where we distinguish between the shift of the ‘cross-section’ due to a rigid displacement and due to triangular shaping. Although both contribute to the Shafranov shift, we shall here focus on the θ -average, rigid shift, most useful when thinking of the problem in terms of shifted ellipses (see figure 9).

The variation required then simply involves the second-order equations in Appendix D, explicitly in (D3). In this case, unlike for the magnetic well sensitivity, we need the variation respect to δp_2 , a scalar,

$$\hat{\mathbb{L}} \begin{pmatrix} \delta X_{20} \\ \delta Y_{20} \end{pmatrix} = \frac{\delta f}{\delta p_2} \delta p_2, \quad (7.3)$$

the variation of $\hat{\mathbb{L}}$ vanishing exactly. Solving this linear system, we may define

$$\mathcal{S} = \hat{\mathbb{L}}^{-1} \frac{\delta \mathbf{f}}{\delta p_2}, \quad (7.4)$$

which is a vector function of φ . This gradient holds for all second-order choices, as the equation is linear on p_2 . That is, $\mathcal{S}$ is a property of the first-order field.

It is natural to introduce some normalisation for the gradient now. With $[\Delta] = [L]^{-1}$ and $[p_2] = [p][L]^{-2}$, introducing the dimensionless plasma beta parameter $\beta_{p,2} = \mu_0 p_2 / (\bar{B}^2/2)$ for some reference magnetic field and the reference $r_{\text{ref}} = R/A_{\text{ref}}$, we define $\hat{\mathcal{S}} = (\bar{B}^2/2\mu_0 r_{\text{ref}})\mathcal{S}$. We should therefore interpret this gradient as a fractional displacement of the surface at a reference aspect ratio A_{ref} due to a change in plasma beta. We define the following.

(a) *Maximum sensitivity of Shafranov shift, $\hat{\mathcal{S}}_{\text{max}}$.*

Maximum relative value of the Shafranov shift gradient with respect to changes of plasma β at a reference aspect ratio A_{ref} ,

$$\hat{\mathcal{S}}_{\text{max}} = \max_{\varphi} \left[\frac{\bar{B}^2}{2\mu_0 r_{\text{ref}}} |\mathcal{S}| \right]. \quad (7.5)$$

It is illuminating to consider the familiar tokamak limit. Simplifying the system (see details in [Appendix D.3](#)), and considering for simplicity an up-down symmetric configuration,

$$\hat{\mathcal{S}}_{\text{max}} = \frac{A_{\text{ref}}}{\iota_0^2} \frac{\bar{d}^4}{\bar{d}^4 - 3}, \quad (7.6)$$

where all quantities have their usual meaning. The sensitivity grows with aspect ratio, as the same displacement becomes relatively larger compared with the minor radius. The scaling with $\propto \iota_0^{-2}$ responds to the physical picture of the poloidal field balancing the pressure gradient. A stronger poloidal component means the more resilient the field is to the push of the pressure. The role of elongation is more involved, although it guarantees for vertically elongated configurations, $\bar{d} < 1$, $\hat{\mathcal{S}}_x < 0$. That is, there will be a bunching of surfaces on the outboard side. It also shows the artificial divergence discussed in the previous section.

So far, the gradient measure defined in (7.5), $\hat{\mathcal{S}}_{\text{max}}$, provides a scalar measure of sensitivity without distinguishing direction. The significance of any given shift does however depend on the direction in which it occurs. It is not the same to rigidly shift nested ellipses in the direction of the minor axis or the major axis. For directions in which surfaces are closer to each other, the bunching of surfaces is easier and so, potentially, is their intersection. We define an associated critical β as that value for which, due to the rigid Shafranov shift, the underlying first-order ellipses just touch at a radius of r_{ref} . This definition directly uses $\hat{\mathcal{S}}$, as described in detail in [Appendix G](#) and illustrated in [figure 9](#).

(b) *Estimate of critical plasma β , β_{Δ} .*

Estimate of the critical plasma β_p at which, due to the rigid part of the Shafranov shift, the first-order near-axis elliptical flux surfaces of an equilibrium just touch at $r = r_{\text{ref}}$. The value of the plasma beta is defined as a dimensionless scalar β_{Δ} , for which a full definition is provided in [\(G12\)](#).




FIGURE 7. Shafranov shift sensitivity to plasma β . The plots show the shape gradient $\hat{S}_x$ and $\hat{S}_y$ for configurations (a) 4.1 and (b) 4.3 in the benchmark set.

Large values of β_Δ indicate the resilience of the Shafranov shift to changes in plasma β .

The metric β_Δ condenses the information of the gradient S into a single physically meaningful measure. However, it does so under certain simplifying assumptions. In particular, when it comes to describing when flux surfaces touch each other, it ignores the triangular shaping of flux surfaces, approximating them as ellipses. Thus, β_Δ is only an estimate of the true critical plasma β , which we may compute numerically.

(c) *Numerical critical plasma β , β_c .*

Critical value of $\beta_p = r_{\text{ref}}^2 \beta_{p,2}$, for an assumed aspect ratio A_{ref} , at which the flux surfaces of the near-axis construction just touch. That is, for $\beta_p = \beta_c$, we have $A_c = A_{\text{ref}}$.

In other words, fixing $\beta_p = \beta_c$, the near-axis construction will not succeed for aspect ratios below A_{ref} . The value of β_c may be found by a root search (and thus is more numerically costly) and includes all elements of shaping up to (and including) second order. We emphasise that this quantity, although defined in exact terms, may still not be the most realistic β limit, because it is calculated taking the axis fixed, allowing flux surfaces to move around it as the plasma beta is changed.

7.2. Implementation and examples

We assess the Shafranov shift sensitivity in the configurations of the benchmark. Figure 7 shows the shape gradient of the Shafranov shift for two different examples and table 4 summarises the scalar features of the configurations.¹²

The sensitivity S_{max} shows that there exists a wide range of sensitivities to changes in plasma β . The gradient as a function of φ delves deeper into these differences showing the stark difference in the binormal shift of the surfaces (which the half-helicity field 4.3 appears to avoid to a large extent). The gradients also vanish at certain points in φ , corresponding to the directions that would break the up-down

¹²The gradient calculation can be shown by finite differencing to be correct to machine precision.

Configs	4.1	4.2	4.3	4.4 (unopt)	4.4 (opt)	# 50	# 76
$\hat{S}_{\max}$	17.82	2.52	4.30	7.02	6.02	2.40	3.31
$\beta_{\Delta}[\%]$	3.98	24.30	20.20	11.57	16.15	47.61	22.51
$\beta_c [\%]$	8.30	15.29	23.66	14.13	19.31	48.65	22.47
$ \iota_0 $	0.110	1.230	0.768	0.385	0.359	0.943	0.918

TABLE 4. Shafranov shift sensitivity to plasma pressure. The table shows the values of the maximum sensitivity of the Shafranov shift, the estimate and numerical critical β as well as the reference rotational transform on axis. The half-helicity fields appear to be the most resilient in the benchmark set. The comparison of β_c to β_{Δ} shows that although β_c includes some key elements of the full phenomena, it can deviate significantly.

symmetry of the cross-sections at stellarator symmetric points. From this benchmark, the lowest number of field period configuration (config. 4.1) is the most sensitive, which also has the lowest value of rotational transform. This is true of β_{Δ} as well, showing that half-helicity configurations (namely configs. 4.3, # 50 and # 76) consistently exhibit larger beta limits.

The tools introduced here make it possible, in principle, to explore the generality of such observations, but an in-depth analysis is left for future work.

8. Conclusions

Recent works have demonstrated that it is possible to find approximately quasi-isodynamic stellarator equilibria to second order using near-axis theory, and faithfully reproduce their properties in global equilibria at finite aspect ratio. Such ‘directly constructed’ stellarator equilibria can in principle be used as the basis for further integrated optimisation to develop new stellarator designs.

However, to make the near-axis construction really practically useful, both for developing fundamental understanding and as a tool for stellarator optimisation, it is critical to develop techniques to navigate the space of second-order solutions. The present work tackles this task by defining a set of measures designed to assess the key properties for which a quasi-isodynamic stellarator strives, including low neoclassical transport, quality of omnigenicity and robust stability. Other measures already implemented (although not mentioned in the text) include the coil complexity proxy $L_{\nabla B}$ (Landreman 2021; Kappel, Landreman & Malhotra 2024), the effective Rosenbluth–Hinton residual (Rosenbluth & Hinton 1998; Rodríguez & Plunk 2025; Zhu, Lin & Bhattacharjee 2025) and proxies to evaluate the maximum- J property (Rodríguez *et al.* 2024). Virtually any other measure that depends on magnetic geometry, for instance, MHD ballooning stability or micro-turbulence, can be readily implemented as well in the future.

The procedure for obtaining fields to second order involves fixing a set of free functions at zeroth, first and second order. Unsurprisingly, there is a certain interaction between these choices and the metrics have been formulated in a manner to help assess these choices. We also propose constructive approaches to make these choices, e.g. a way to obtain a minimally shaped, marginally stable (in the magnetic well sense) field or to construct an optimally ‘omnigenised’ configuration given a first-order construction.

The toolset presented here, consisting of both absolute measures (e.g. ϵ_{eff} on edge) and sensitivity measures (shape gradient of the magnetic well), now opens the way to exploratory studies of the broad parameter space of QI stellarators, and systematic investigation of ‘trade-offs’ – i.e. the compatibility of desired stellarators properties, underlying reasons and strategies for effectively striking compromises.

Acknowledgements

We thank Carolin Nührenberg for her assistance with magnetic well definitions.

Editor Peter Catto thanks the referees for their advice in evaluating this article.

Data availability

The data and scripts that support the findings of this study are openly available at the Zenodo repository with DOI/URL 10.5281/zenodo.15340967.

Declaration of interests

The authors report no conflict of interest.

Funding

E.R. was partially supported by a grant by Alexander-von-Humboldt-Stiftung, Bonn, Germany, through a postdoctoral research fellowship.

Appendix A. Benchmark configuration summary

Throughout the paper, we use a number of near-axis fields as a way of benchmarking or illustrating the various measures introduced. In this section, we summarise of what those configurations consist, indicating their main features and referring to the right piece of literature where appropriate.

A.1. QI database

In §3, we use a large set of QI configurations as a way to benchmark the δB measure. This benchmark set had previously been used in work like Rodriguez & Plunk (2025), and we now present some of its details. A full description of its construction and the theoretical elements that go into it will be the focus of future publication (Plunk *et al.* 2025). For a specific example, see Plunk *et al.* (2024).

The benchmark set consists of first-order QI configurations with a number of field periods ranging from $N = 1$ to 6, constructed as a three-parameter family as follows. The first parameter is related to the shape of the magnetic axis, which is described by its curvature, κ , and torsion, τ . For technical reasons, these must be defined differently for the case $N = 1$, as compared with $N \geq 2$. For the former, we define

$$\kappa = \cos^2\left(\frac{N\ell}{2}\right) \sin\left(\frac{N\ell}{2}\right) (\kappa_1 \sin(N\ell) + \kappa_2 \sin(2N\ell)), \quad (\text{A1})$$

$$\tau = \tau_0 + \tau_1 \cos(N\ell) + \tau_2 \cos(2N\ell), \quad (\text{A2})$$

while in the latter case ($N \geq 2$), we set $\kappa_2 = 0$ and $\tau_2 = 0$. This difference is because four degrees of freedom are here used to close the curve smoothly for the case $N = 1$, while only two degrees of freedom are required for $N \geq 2$. Thus, a single parameter remains to define a one-parameter family of curves. The variable ℓ is the length along the curve. This describes a stellarator symmetric axis with zeroes of

second order at the tops and third order at the bottom. This is consistent with a curve of half-helicity (Camacho Mata & Plunk 2023; Rodríguez *et al.* 2025). The total length of the axis is taken to be $L = 2\pi$.

The magnetic field on axis B_0 is chosen to match the zeroes of κ ,

$$B_0 = 1 + 0.25 \cos(N\ell) + 0.0625 \cos(2N\ell), \quad (\text{A3})$$

and to have a mirror ratio of $\sim 24\%$.¹³ The breaking of omnigenicity necessary at first order is done through a smooth buffer region as described by Rodríguez *et al.* (2025, § B.2.3) of order $k = 3$.

The first-order construction is finalised by the choice of the elongation profile,

$$\rho = \rho_0 + \rho_1 \cos(N\varphi), \quad (\text{A4})$$

which is directly related to the more common near-axis quantities of Landreman & Sengupta (2019) and Rodríguez *et al.* (2025)] by $\rho = \bar{e} + (1 + \sigma^2)/\bar{e}$, where $\bar{e} = \bar{d}^2 B_0 / \bar{B}$ (which is intimately related to the true elongation of the flux surface) (Plunk *et al.* 2024). These two parameters $\{\rho_0, \rho_1\}$, together with τ_1 , are our three parameters that span the QI database set considered.

Considering different values of these parameters, the QI database is constructed with a total of 1680 configurations distributed in $\{63, 502, 358, 244, 250, 263\}$ between number of field periods $\{1, 2, 3, 4, 5, 6\}$. Further details on the procedure followed to construct such configurations will follow in a future paper.

A.2. Benchmark configurations from Rodríguez *et al.* (2025)

In the rest of the paper, we use a reduced set of benchmark configurations, mainly based on those used and explored by Rodríguez *et al.* (2025, § 4) to test the correct near-axis construction to second order. Here, we briefly summarise the key properties of each of those configurations,

- (a) **Configuration 4.1:** minimally shaped second-order stellarator symmetric, vacuum configuration with $N = 2$, a smooth buffer region with $k = 5$, mirror ratio 15 % and first-order zeroes of curvature (zero helicity axis).
- (b) **Configuration 4.2:** shaped second-order stellarator symmetric, finite beta and current configuration with $N = 3$, a smooth buffer region with $k = 5$, mirror ratio 25 % and first-order zeroes of curvature (zero helicity axis).
- (c) **Configuration 4.3:** minimally shaped second-order stellarator symmetric, vacuum configuration with $N = 3$, a smooth buffer region with $k = 5$, mirror ratio 25 %, and third-order zero of curvature at the bottom and two at the top (half-helicity axis).
- (d) **Configuration 4.4:** minimally shaped second-order stellarator symmetric, vacuum configuration with $N = 3$, a smooth buffer region with $k = 5$, mirror ratio 25 %, and third-order zero of curvature at the bottom and first-order zeroes of curvature (zero helicity axis). The first-order choices were then optimised to minimise the second-order QI error in the central 20 % of the toroidal domain. In this paper, we include both the unoptimised and optimised versions.

Figure 8 shows three-dimensional renditions of these configurations for reference.

¹³Here, mirror ratio is defined as $\Delta_{\text{mirr}} = (B_{0,\text{max}} - B_{0,\text{min}})/(B_{0,\text{max}} + B_{0,\text{min}})$, which is the most used form.




FIGURE 8. Three-dimensional rendition of configurations in the benchmark set constructed for $r = 0.1$ and with the colourmap denoting the strength of the magnetic field $|\mathbf{B}|$.

A.3. Additional half-helicity configurations

In addition to the abovementioned findings, we also include two additional half-helicity fields constructed much in the same way as Config. 4.4, but corresponding to other constructions in the same parameter family.

- (a) Config. # 50: $\kappa_1 = -7.451287$, $\tau_0 = 1.379097$, $\tau_1 = -0.444444$, $\rho_0 = 4.5$, $\rho_1 = -0.9$.
- (b) Config. # 76: $\kappa_1 = -8.938611$, $\tau_0 = 1.284143$, $\tau_1 = -0.177778$, $\rho_0 = 4.5$, $\rho_1 = -0.9$.

The case of # 76 is particularly interesting, as it has a rather omnigenous intrinsic second-order behaviour. See figure 8 for the configurations.

Appendix B. Details on the near-axis form of ϵ_{eff}

In this appendix, we present the details of the asymptotic treatment of the effective ripple of a magnetic field with poloidal contours of $|\mathbf{B}|$ within the near-axis framework. The expression we must evaluate asymptotically is the following:

$$\epsilon_{\text{eff}}^{3/2} = \frac{\pi}{8\sqrt{2}} \frac{(\bar{R}\bar{B})^2}{\mathcal{G}^2} \int_{1/B_{\text{max}}}^{1/B_{\text{min}}} \lambda \hat{\mathcal{E}}(\lambda) d\lambda, \quad (4.6)$$

where

$$\hat{\mathcal{E}}(\lambda) = \frac{1}{\pi} \int_0^{2\pi} \frac{\hat{H}(\lambda, \alpha)^2}{\hat{I}(\lambda, \alpha)} d\alpha, \quad (4.7)$$

$$\hat{H}(\lambda, \alpha) = \frac{1}{\bar{B}} \int_{\varphi_-}^{\varphi_+} \frac{\mathcal{H}(\lambda, B)}{B^2} \mathbf{B} \times \nabla B \cdot \nabla \psi d\varphi, \quad (\text{B1a})$$

$$\hat{I}(\lambda, \alpha) = \int_{\varphi_-}^{\varphi_+} \frac{\sqrt{1 - \lambda \bar{B}}}{(B/\bar{B})^2} d\varphi, \quad (\text{B1b})$$

$$\mathcal{G}^2 = 2 \left(\frac{1}{2\pi} \int_0^{2\pi} d\alpha \int_0^{2\pi/N} d\varphi \frac{|\nabla\psi|^2}{B^2} \right) / \left(\frac{1}{2\pi} \int_0^{2\pi} d\alpha \int_0^{2\pi/N} \frac{d\varphi}{B^2} \right), \quad (\text{B1c})$$

$$\mathcal{H}(\lambda, B) = \frac{\sqrt{1 - \lambda \bar{B}}}{(B/\bar{B})^2} \left(\frac{4}{\lambda B} - 1 \right), \quad (4.3)$$

where we have explicitly applied the considerations in § 4.1 to write the expressions in the most near-axis friendly form.

B.1. Expansion of $\hat{H}$

For the expansion of $\hat{H}$, let us write using Boozer coordinates

$$\mathbf{B} \times \nabla B \cdot \nabla \psi = \frac{B^2}{G + \iota I} (I \partial_\varphi - G \partial_\theta) B, \quad (\text{B2})$$

which may be directly expanded,

$$\frac{\mathbf{B} \times \nabla B \cdot \nabla \psi}{B^2} \approx r B_0 d \cos(\alpha - \alpha_{\text{buf}}) + \quad (\text{B3})$$

$$+ r^2 \left(B_{2s}^{\text{QI}} \cos 2\alpha - B_{2c}^{\text{QI}} \sin 2\alpha + \frac{I_2 B'_0}{G_0} \right) + O(r^3), \quad (\text{B4})$$

where the order $O(r^3)$ term will include third-order elements of the magnetic field magnitude and thus we shall not include them.

With this expansion of the radial drift and noting that $\mathcal{H}$ vanishes at bounce points, the perturbation of $\hat{H}$ is a combination of the asymptotic form of the drift and that of the integrand $\mathcal{H}$. At first order, we easily obtain $\hat{H} \approx r h^{(1)} \sin \alpha$,

$$h^{(1)} = \int_{\varphi_-}^{\varphi_+} \frac{B_0}{\bar{B}} \mathcal{H}(\lambda, B_0) d \sin \alpha_{\text{buf}} d\varphi, \quad (4.9b)$$

where the odd part of the radial drift does not contribute, as its contributions from each side of the well precisely cancel out.

At second order, we must consider both the contributions from the modification in the drift, but also the measure $\mathcal{H}$. To avoid the explicit appearance of the singular integrand $1/\sqrt{1 - \lambda \bar{B}_0}$, we integrate the resulting expression by parts (the boundary term being guaranteed to vanish for a field without puddles, Rodríguez & Plunk 2023), and thus, after enforcing parity in φ and collecting terms $\hat{H}$ from both the perturbed $\hat{H} \approx r h^{(1)} \sin \alpha + r^2 h^{(2)} \sin 2\alpha$,

$$h^{(2)} = -2 \int_{\varphi_-}^{\varphi_+} \mathcal{H}(\lambda, B_0) \frac{\Delta B_{2c}^{\text{QI}}}{\bar{B}} d\varphi, \quad (4.12a)$$

$$\Delta B_{2c}^{\text{QI}} = B_{2c}^{\text{QI}} - \frac{1}{4} \partial_\varphi \left(\frac{B_0^2 d^2}{B'_0} \cos 2\alpha_{\text{buf}} \right). \quad (4.12b)$$

B.2. Expansion of $\hat{I}$

We may now proceed in a similar fashion with the asymptotic evaluation of $\hat{I}$. It is useful to note that $\hat{I}$ is almost the expression for the second adiabatic invariant $\mathcal{J}_{\parallel}$ for trapped particles labelled by λ . Fortunately, the asymptotic treatment of such expressions (especially at higher order), with the careful handling of the boundary contributions, has already been done elsewhere and thus we simply need to get the appropriate expressions from Rodríguez *et al.* (2024).

The leading order asymptotic form is simple, $\hat{I} \approx I^{(0)}$,

$$I^{(0)} = \int_{\varphi_-}^{\varphi_+} \frac{\sqrt{1 - \lambda B_0}}{(B_0/\bar{B})^2} d\varphi. \quad (4.9b)$$

For the next orders, we can directly draw from (A14) of Rodríguez *et al.* (2024), and thus write separating $\hat{I} \approx I^{(0)} + r I^{(1)} \cos \alpha + r^2 (\bar{I}^{(2)} + \tilde{I}^{(2)} \cos 2\alpha)$,

$$I^{(1)} = -2 \int_{\varphi_-}^{\varphi_+} F^*(\lambda, B_0) d \sin \alpha_{\text{buf}} d\varphi, \quad (B6)$$

and

$$\bar{I}^{(2)} = -2 \int_{\varphi_-}^{\varphi_+} \frac{F^*(\lambda, B_0)}{B_0} \left[B_{20} - \frac{1}{4} \partial_{\varphi} \left(\frac{B_0^2 d^2}{B'_0} \right) \right] d\varphi, \quad (B7a)$$

$$\tilde{I}^{(2)} = 2 \int_{\varphi_-}^{\varphi_+} \frac{F^*(\lambda, B_0)}{B_0} \left[B_{2c}^{\text{QI}} - \frac{1}{4} \partial_{\varphi} \left(\frac{B_0^2 d^2}{B'_0} \cos 2\alpha_{\text{buf}} \right) \right] d\varphi, \quad (B7b)$$

where

$$F^*(\lambda, B_0) = \frac{1}{\sqrt{1 - \lambda B_0}} \frac{1 - 3\lambda B_0/4}{(B_0/\bar{B})^2}. \quad (B8)$$

Note that all the α dependence of $\hat{I}$ vanishes in the QI limit, as it must do given the similarity of $\hat{I}$ with $\mathcal{J}_{\parallel}$; the second adiabatic invariant is a flux function in an omnigenous field (Bernardin *et al.* 1986). The first-order correction is purely driven by the buffer region, while the second comes from the second-order QI mismatch. In that omnigenous limit, $\bar{I}^{(2)}$ is the only higher order correction left to second order, and it represents the change in the field line length (B_{20}), but also in the velocity of the trapped particle along the bounce trajectory.

B.3. Constructing $\mathcal{E}$

Let us then consider now using the asymptotic definitions $\hat{H} \approx r h^{(1)} \sin \alpha + r^2 h^{(2)} \sin 2\alpha$ and $\hat{I} \approx I^{(0)} + r I^{(1)} \cos \alpha + r^2 (\bar{I}^{(2)} + \tilde{I}^{(2)} \cos 2\alpha)$,

$$\frac{1}{2\pi} \int_0^{2\pi} \frac{\hat{H}^2}{\hat{I}} d\alpha \approx \frac{r^2}{2} \left[\frac{(h^{(1)})^2}{I^{(0)}} + r^2 \left(\frac{(h^{(2)})^2}{I^{(0)}} - \underbrace{\frac{h^{(2)} h^{(1)} I^{(1)}}{I^{(0)} I^{(0)}}}_{\textcircled{1}} \right) \right] \quad (B9)$$

$$+ \underbrace{\frac{(h^{(1)})^2}{I^{(0)}} \left[\left(\frac{I^{(1)}}{2I^{(0)}} \right)^2 - \frac{\bar{I}^{(2)} - \tilde{I}^{(2)}/2}{I^{(0)}} \right]}_{\textcircled{2}} \right]. \quad (B10)$$

A careful consideration of this expansion will show that we have *a priori* unfairly omitted a contribution that is also of order r^4 , and that comes from beating of $h^{(1)}$ with what would be $h^{(3)}$, and in particular its $\sin \alpha$ harmonic. That term involves higher order magnetic field components, notably B_{31}^c , and thus goes beyond the second-order considerations here. Setting this third component aside, the term will also have components inherited from lower orders, which can be evaluated explicitly to assess their magnitude. In practice, these contributions appear to be small. All in all, we shall ignore such a contribution.

In fact, in practice, at order $O(r^4)$, the main contribution comes from $h^{(2)}$; namely, the breaking of omnigenity at second order. Thus, in practice, and as explicitly presented in the text, we will focus on this term.

B.4. Expansion of geometric factor $\mathcal{G}$

Let us write the geometric factor $\mathcal{G}^2 = 2D^2/L$ and consider the asymptotic expansion of the integrals. Following Jorge & Landreman (2020, (33)) and carrying it out to higher order, we write

$$\begin{aligned} D &= \frac{1}{2\pi} \int d\alpha \int \frac{d\varphi}{B^2} |\nabla\psi| \\ &\approx r \underbrace{\frac{1}{2\pi} \int d\alpha \int \frac{d\varphi}{B_0} T_1}_{D_1} + r^2 \underbrace{\frac{1}{2\pi} \int d\alpha \int \frac{d\varphi}{B_0} T_1 \left(\frac{T_2}{T_1} - 2 \frac{B_1}{B_0} \right)}_{D_2} + \\ &\quad + r^3 \underbrace{\frac{1}{2\pi} \int d\alpha \int \frac{d\varphi}{B_0^2} T_1 \left(\frac{T_3}{T_1} - 2 \frac{B_1 T_2}{B_0 T_1} + 3 \left(\frac{B_1}{B_0} \right)^2 - 2 \frac{B_2}{B_0} \right)}_{D_3}, \end{aligned} \quad (\text{B11})$$

where $|\nabla\psi|/B_0 \approx rT_1 + r^2T_2 + \dots$, and we may explicitly write

$$T_1 = \sqrt{(\partial_\chi X_1)^2 + (\partial_\chi Y_1)^2}, \quad (\text{B12a})$$

$$T_2 = \frac{1}{T_1} \left[T_1^2 \frac{B_1}{B_0} + \partial_\chi X_1 \partial_\chi X_2 + \partial_\chi Y_1 \partial_\chi Y_2 \right], \quad (\text{B12b})$$

and we omit the explicit form of T_3 for brevity. Note that $D_2 \approx 0$ as we have odd α -harmonics that average to zero (even if there is a complicated α -dependence in T_1). To evaluate the α -average of the other quantities, we must explicitly write their poloidal angle dependence. For instance, T_1 in the text (4.10), (Jorge & Landreman 2020, (33)),

$$|\nabla\psi|^2 \approx r^2 B_0^2 [(X_{1c} \sin \chi - X_{1s} \cos \chi)^2 + (Y_{1c} \sin \chi - Y_{1s} \cos \chi)^2]. \quad (\text{B13})$$

Because we have $|\nabla\psi|$ instead of $|\nabla\psi|^2$, the integral over χ must be performed numerically (or related to elliptic integrals, Gradshteyn & Ryzhik 2014, (3.670.1)).

The expansion of L is simpler and only the expansion in B is necessary. Considering the average over α to eliminate the bare harmonics of χ , we write

$$L = \underbrace{\frac{1}{2\pi} \int d\alpha \int \frac{d\varphi}{B^2}}_{L_0} \approx \underbrace{\int \frac{d\varphi}{B_0^2} + r^2 \frac{1}{2\pi} \int d\alpha \int \frac{d\varphi}{B_0^2} \left[3 \left(\frac{B_1}{B_0} \right)^2 - 2 \frac{B_2}{B_0} \right]}_{L_2}. \quad (\text{B14})$$

Configs	4.1	4.2	4.3	4.4 (unopt)	4.4 (opt)	# 50	# 76
①	9.99×10^{-06}	2.37×10^{-04}	-2.99×10^{-04}	4.80×10^{-05}	2.35×10^{-05}	-8.30×10^{-05}	2.31×10^{-04}
②	-1.72×10^{-06}	-2.69×10^{-05}	-5.76×10^{-05}	-1.66×10^{-05}	-1.83×10^{-05}	-7.11×10^{-06}	-9.91×10^{-05}
③	6.26×10^{-08}	-1.85×10^{-06}	3.25×10^{-06}	-3.31×10^{-07}	-1.04×10^{-06}	1.98×10^{-07}	6.68×10^{-06}

TABLE 5. Relative magnitude of additional asymptotic terms in $\epsilon_{\text{eff}}^{3/2, (2)}$. The table shows the relative contribution to $\epsilon_{\text{eff}}^{3/2, (2)}$ of the asymptotic terms ①, ② and ③. The smallness of these contributions throughout the benchmark configurations supports, along with the role of QI breaking, the expression for the effective ripple measure used.

With both of these,

$$\frac{1}{\mathcal{G}^2} \approx \frac{L_0}{G_0 D_1^2} \left[1 + r^2 \underbrace{\left(\frac{L_2}{L_0} - \frac{G_2}{G_0} - 2 \frac{D_3}{D_1} \right)}_{\textcircled{3}} \right]. \quad (\text{B15})$$

In practice, it is the leading order piece that matters and which we defined as $\mathcal{G}^{(0)}$ in the main text.

B.5. Assessment of approximation

With the approximations and expansions considered previously, we may then construct the leading order forms of $\epsilon_{\text{eff}}^{3/2}$ as presented in §4. We had argued that some of the terms involved in the asymptotic expression of the effective ripple are, in practice, subsidiary, even though not in the $O(r)$ sense. We now provide some numerical evidence in such regard using the benchmark configurations in the paper as a test-bed. In table 5, we summarise the relative magnitude of different contributions in an attempt to argue the correctness of the approximation to ϵ_{eff} . We show by ①, ② and ③ the relative contribution by the terms denoted by these symbols in the above-mentioned asymptotic expansions to $\epsilon_{\text{eff}}^{3/2, (2)}$. The smallness of these terms gives us an argument to focus on the contribution of the second-order QI breaking and use it as a part of our effective ripple measure. The benchmark in figure 2 also seconds this.

Appendix C. Details of the ripple well measure A_w calculation

In this appendix, we develop the formalism necessary to efficiently evaluate the ripple well measure A_w presented in the main text, §5; namely,

$$A_w = R / \min \{ r \mid \exists \theta, \varphi : \partial_\varphi|_\alpha B(r, \theta, \varphi) = 0, \partial_\varphi^2|_\alpha B(r, \theta, \varphi) = 0 \}. \quad (5.1)$$

The multi-valued form of A_w for each fixed φ was defined as $\hat{A}_w$, with $A_w = \max_\varphi \hat{A}_w$.

C.1. First-order construction: too simple

Let us consider first, as a way of introduction, A_w for the simplest case of a first-order field: $B(\theta, \varphi) = B_0(\varphi) + r B_1(\theta, \varphi)$, where $B_1(\theta, \varphi)$ is given in (4.4a). The two

conditions in (5.1) to solve simultaneously are

$$B'_0 + r(S' \cos \alpha - C' \sin \alpha) = 0, \quad (\text{C1a})$$

$$B''_0 + r(S'' \cos \alpha - C'' \sin \alpha) = 0, \quad (\text{C1b})$$

where $S = B_0 d \sin \alpha_{\text{buf}}$, $C = B_0 d \cos \alpha_{\text{buf}}$ and primes denote derivatives in φ . Combining these equations and assuming $r \neq 0$, one may solve for $\tan \alpha$ and substitute it back to retain a real, positive root for r ,

$$R/\hat{A}_w^{(1)} = \left| \frac{B'_0}{S' - C' \tan \alpha} \sqrt{1 + \tan^2 \alpha} \right|, \quad (\text{C2})$$

where α is given by

$$\tan \alpha = \frac{(S'/B'_0)'}{(C'/B'_0)'}, \quad (\text{C3})$$

defined in the appropriate quadrant so that $1/\hat{A}_w^{(1)}$ solves the original set of equations. Note that in obtaining an expression for $\tan \alpha$, we have assumed that the equation is not trivially satisfied by $r = 0$. This would have worked for $B'_0 = 0 = B''_0$, which now instead sees $\hat{r}_w \rightarrow \infty$.¹⁴ This should make it clear that the appearance of secondary wells is not the same as the breakdown of $|\mathbf{B}|$ contour topology.

In practice, the appearance of secondary wells due to the first-order variation of the field is too simplistic. It is far from capturing the evolution of a second-order near-axis field construction. We must therefore turn to the more complex higher order consideration.

C.2. Higher order form

The procedure at higher orders resembles the procedure followed in the construction of the critical radius r_c by Landreman (2021). The latter corresponds to the smallest distance from the axis at which flux surfaces become ill-behaved; namely, the first instance of a vanishing coordinate Jacobian $\mathcal{J} = \partial_\psi \mathbf{x} \times \partial_\theta \mathbf{x} \cdot \partial_\varphi \mathbf{x} = 0$. That problem requires the solution of a set of equations formally analogous to those for A_w , see Landreman (2021, § 4). We exploit that analogy and adapt that work to our problem.

Using the helical angle $\chi = \theta - N\varphi$, where N is the helicity of the magnetic axis (Rodríguez *et al.* 2022b; Camacho Mata & Plunk 2023; Rodríguez *et al.* 2025), and writing $B_1(\chi, \varphi) = B_{1c}(\varphi) \cos \chi + B_{1s}(\varphi) \sin \chi$ and $B_2(\chi, \varphi) = B_{20}(\varphi) + B_{2c}(\varphi) \cos 2\chi + B_{2s}(\varphi) \sin 2\chi$, the first condition that $\hat{A}_w$ must satisfy is $\partial_\varphi B|_\alpha = 0$, which explicitly reads

$$g_0 + r(g_{1c} \cos \chi + g_{1s} \sin \chi) + r^2(g_{20} + g_{2c} \cos 2\chi + g_{2s} \sin 2\chi) = 0, \quad (\text{C4})$$

where

$$g_0 = B'_0, \quad (\text{C5a})$$

$$g_{1c} = B'_{1c} + \bar{l}B_{1s}, \quad g_{1s} = B'_{1s} - \bar{l}B_{1c}, \quad (\text{C5b})$$

$$g_{20} = B'_{20}, \quad g_{2c} = B'_{2c} + 2\bar{l}B_{2s}, \quad g_{2s} = B'_{2s} - 2\bar{l}B_{2c}. \quad (\text{C5c})$$

¹⁴To show this, it is sufficient to consider the no-buffer limit of (C2), which at the bottom of the well requires $B'_0 d' = 0$ and $r \sin \alpha = B'_0/C'$, so that $r_w = B'_0/C' = 1/d + B'_0/B_0 d'$. Regardless of how one chooses the orders of B_0 and d , this expression diverges as the bottom of the well is approached.

The second condition, namely $\partial_\varphi^2 B|_\alpha = 0$, can also be found explicitly and written in a form analogous to (C4). We shall not give it explicitly for brevity. Given these two equations, we approach their solution by isolating r for every φ and eliminating it from the equations, reducing the system to a single equation on harmonics of χ . We shall call this the K -equation. Following this procedure, we may write

$$r = \frac{f_{1c} \cos \chi + f_{1s} \sin \chi}{f_{20} + f_{2c} \cos 2\chi + f_{2s} \sin 2\chi}, \quad (\text{C6})$$

where

$$f_{1c} = B'_0(B''_{1c} + 2\bar{t}_0 B'_{1s} - \bar{t}_0^2 B_{1c}) - B''_0(B'_{1c} + \bar{t}_0 B_{1s}), \quad (\text{C7a})$$

$$f_{1s} = B'_0(B''_{1s} - 2\bar{t}_0 B'_{1c} - \bar{t}_0^2 B_{1s}) + B''_0(-B'_{1s} + \bar{t}_0 B_{1c}), \quad (\text{C7b})$$

$$f_{20} = B''_0 B'_{20} - B'_0 B''_{20}, \quad (\text{C7c})$$

$$f_{2c} = -B'_0(B''_{2c} + 4\bar{t}_0 B'_{2s} - 4\bar{t}_0^2 B_{2c}) + B''_0(B'_{2c} + 2\bar{t}_0 B_{2s}), \quad (\text{C7d})$$

$$f_{2s} = B'_0(-B''_{2s} + 4\bar{t}_0 B'_{2c} + 4\bar{t}_0^2 B_{2s}) + B''_0(B'_{2s} - 2\bar{t}_0 B_{2c}). \quad (\text{C7e})$$

Eliminating r from (C4), we are left with the following K -equation in the coordinates χ and φ ,

$$K_0 + K_{2c} \cos 2\chi + K_{2s} \sin 2\chi + K_{4s} \sin 4\chi + K_{4s} \sin 4\chi = 0, \quad (\text{C8})$$

where

$$\begin{aligned} K_0 = & \frac{1}{4} [2g_0(f_{2c}^2 + 2f_{20}^2 + f_{2s}^2) + f_{1c}(f_{2c}g_{1c} + f_{2s}g_{1s}) + f_{1s}(f_{2s}g_{1c} - f_{2c}g_{1s}) \\ & + 2f_{20}(f_{1c}g_{1c} + f_{1s}g_{1s}) + 2(f_{1c}^2 + f_{1s}^2)g_{20} \\ & + (f_{1c}^2 - f_{1s}^2)g_{2c} + 2f_{1c}f_{1s}g_{2s}], \end{aligned} \quad (\text{C9a})$$

$$\begin{aligned} K_{2c} = & \frac{1}{2} [f_{1c}f_{2c}g_{1c} + f_{20}(4f_{2c}g_0 + f_{1c}g_{1c} - f_{1s}g_{1s}) + f_{1c}^2(g_{20} + g_{2c}) \\ & + f_{1s}(f_{2c}g_{1s} - f_{1s}g_{20} + f_{1s}g_{2c})], \end{aligned} \quad (\text{C9b})$$

$$\begin{aligned} K_{2s} = & \frac{1}{2} [f_{1c}f_{2s}g_{1c} + f_{1s}f_{2s}g_{1s} + f_{20}(4f_{2s}g_0 + f_{1s}g_{1c} + f_{1c}g_{1s}) \\ & + 2f_{1c}f_{1s}g_{20} + (f_{1c}^2 + f_{1s}^2)g_{2s}], \end{aligned} \quad (\text{C9c})$$

$$\begin{aligned} K_{4c} = & \frac{1}{4} [2(f_{2c}^2 - f_{2s}^2)g_0 - f_{2s}(f_{1s}g_{1c} + f_{1c}g_{1s}) + f_{2c}(f_{1c}g_{1c} - f_{1s}g_{1s}) \\ & + (f_{1c}^2 - f_{1s}^2)g_{2c} - 2f_{1c}f_{1s}g_{2s}], \end{aligned} \quad (\text{C9d})$$

$$\begin{aligned} K_{4s} = & \frac{1}{4} [f_{2c}(4f_{2s}g_0 + f_{1s}g_{1c} + f_{1c}g_{1s}) + f_{1c}(f_{2s}g_{1c} + 2f_{1s}g_{2c}) \\ & + (f_{1c}^2 - f_{1s}^2)g_{2s} - f_{1s}f_{2s}g_{1s}]. \end{aligned} \quad (\text{C9e})$$

With the equation cast this way, we may then proceed the way that is detailed by Landreman (2021, §4.2). That is, we solve the K -equation for $\sin 2\chi$ which is a quartic, being careful about the sign of the trigonometric functions. As a result, we have up to four real roots for each value of φ considered in the toroidal domain.

Each one of these roots corresponds to a value of r using (C6), and thus represents a triplet (r, χ, φ) . Finding $\alpha = \chi - \bar{\iota}\varphi$, we may then represent the values of r for this multitude of roots as a function of α , namely $\hat{A}_w$. This is what we represent in figure 5 in § 5. The largest of all these roots is A_w .

Appendix D. Linear system of equations at second order

The construction of the near-axis field at second order has been detailed elsewhere, most notably by Landreman & Sengupta (2019), and in the QI scenario by Rodríguez *et al.* (2025), where a thorough and pedagogical description is provided. Here, we shall not reproduce the derivation of the construction nor the way in which this is solved. However, we shall for completeness write some of the key equations involved in the near-axis construction to second order. In particular, the equations that the different pieces of the second-order shaping satisfy, focusing on the operator notation that we use in this paper.

D.1. Expressions for Y_2

The harmonics of Y_2 are defined through (A32)–(A33) of Landreman & Sengupta (2019) as

$$\begin{pmatrix} Y_{2c} \\ Y_{2s} \end{pmatrix} = \underbrace{\begin{pmatrix} \mathcal{Y}_{2c}^0 \\ \mathcal{Y}_{2s}^0 \end{pmatrix}}_{\mathcal{Y}_0} + \underbrace{\begin{pmatrix} \mathcal{Y}_{2c}^{X_{20}} \mathcal{Y}_{2c}^{Y_{20}} \\ \mathcal{Y}_{2s}^{X_{20}} \mathcal{Y}_{2s}^{Y_{20}} \end{pmatrix}}_{\tilde{\mathbf{Y}}} \begin{pmatrix} X_{20} \\ Y_{20} \end{pmatrix} + \underbrace{\begin{pmatrix} \mathcal{Y}_{2c}^{X_{2c}} \mathcal{Y}_{2c}^{X_{2s}} \\ \mathcal{Y}_{2s}^{X_{2c}} \mathcal{Y}_{2s}^{X_{2s}} \end{pmatrix}}_{\tilde{\mathbf{Y}}} \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix}, \quad (\text{D1})$$

where the components of each matrix are

$$\mathcal{Y}_{2c}^0 = \frac{\kappa \bar{B}}{B_0} \frac{X_{1c} X_{1s}}{X_{1c}^2 + X_{1s}^2}, \quad (\text{D2a})$$

$$\mathcal{Y}_{2c}^{X_{20}} = \frac{X_{1s} Y_{1s} - X_{1c} Y_{1c}}{X_{1c}^2 + X_{1s}^2}, \quad (\text{D2b})$$

$$\mathcal{Y}_{2c}^{Y_{20}} = -1 + 2 \frac{X_{1c}^2}{X_{1c}^2 + X_{1s}^2}, \quad (\text{D2c})$$

$$\mathcal{Y}_{2c}^{X_{2c}} = \frac{X_{1c} Y_{1c} + X_{1s} Y_{1s}}{X_{1c}^2 + X_{1s}^2} = \mathcal{Y}_{2s}^{X_{2s}}, \quad (\text{D2d})$$

$$\mathcal{Y}_{2s}^0 = \frac{\kappa \bar{B}}{2B_0} \frac{X_{1s}^2 - X_{1c}^2}{X_{1c}^2 + X_{1s}^2}, \quad (\text{D2e})$$

$$\mathcal{Y}_{2s}^{X_{20}} = -\frac{X_{1s} Y_{1c} + X_{1c} Y_{1s}}{X_{1c}^2 + X_{1s}^2}, \quad (\text{D2f})$$

$$\mathcal{Y}_{2s}^{Y_{20}} = 2 \frac{X_{1c} X_{1s}}{X_{1c}^2 + X_{1s}^2}, \quad (\text{D2g})$$

$$\mathcal{Y}_{2c}^{X_{2s}} = \frac{X_{1c} Y_{1s} - X_{1s} Y_{1c}}{X_{1c}^2 + X_{1s}^2} = -\mathcal{Y}_{2s}^{X_{2c}}. \quad (\text{D2h})$$

This is a purely algebraic equation, which involves no derivatives.

D.2. Equation for X_{20} and Y_{20}

To complete the construction at second order, using X_{2c} and X_{2s} as inputs to the problem, we need to solve for X_{20} and Y_{20} . The equations to solve are (A41)–(A42) of Landreman & Sengupta (2019), which constitute two coupled first-order ordinary differential equations in φ to be solved simultaneously. Explicitly solving such a system requires rewriting all second-order quantities explicitly in terms of X_{20} , Y_{20} , X_{2c} , X_{2s} and their derivatives, which involves (D1).

Instead of writing all the components of these operators out explicitly, we show how one should systematically proceed to find these through one particular example. The fully fleshed components may be found in the numerical implementation of the near-axis construction presented in the repository published alongside this work. In its operator form, the full expression is

$$\hat{\mathbb{L}} \begin{pmatrix} X_{20} \\ Y_{20} \end{pmatrix} = -\hat{\mathbb{F}} \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix} + \mathbf{f}, \quad (\text{D3})$$

and now we focus on constructing explicitly, say, $\mathbb{L}_{00}$. That is, the piece of the operator acting on X_{20} in (A41) of Landreman & Sengupta (2019), not just explicitly, but also through other second-order functions. Here, both $\hat{\mathbb{L}}$ and $\hat{\mathbb{F}}$ represent linear differential operators of first order that act on a vector of dimension 2, as shown. The vector $\mathbf{f}$ represents a combination of first-order terms.

Let us be more explicit and consider the contributions one by one. First, the explicit X'_{20} components,

$$A_{X'_{20}} = -X_{1s} \frac{d}{d\varphi}, \quad (\text{D4a})$$

and explicit in X_{20} ,

$$A_{X_{20}} = -Y_{1s} \left(\tau \ell' - 2 \frac{I_2}{B} \ell' \right) - 4Y_{1c} \frac{G_0}{B} Z_{2c} - 4Y_{1s} \frac{G_0}{B} Z_{2s}. \quad (\text{D4b})$$

We must then also consider the pieces that depend on Y_{2c} and Y_{2s} , which we know are also related to X_{20} by (D1). So we write again,

$$A_{Y_{2c}} = 4X_{1s} \frac{G_0}{B} Z_{2s} - X_{1c} \left(4 \frac{G_0}{B} Z_{20} - 2\beta_0 \ell' \right) + X_{1s} \left(\tau \ell' - 2\ell' \frac{I_2}{B} \right) - 2Y_{1c} \bar{t}_0, \quad (\text{D4c})$$

$$A_{Y_{2s}} = -4X_{1s} \frac{G_0}{B} Z_{2c} - X_{1c} \left(\tau \ell' - 2\ell' \frac{I_2}{B} \right) - X_{1s} \left(4 \frac{G_0}{B} Z_{20} - 2\beta_0 \ell' \right) - 2Y_{1s} \bar{t}_0, \quad (\text{D4d})$$

and finally, the terms that depend on the derivatives of these,

$$A_{Y'_{2c}} = -Y_{1s} \frac{d}{d\varphi}, \quad (\text{D4e})$$

$$A_{Y'_{2s}} = Y_{1c} \frac{d}{d\varphi}. \quad (\text{D4f})$$

With these expressions, we may then write $\hat{\mathbb{L}}_{00}$ explicitly,

$$\begin{aligned}\hat{\mathbb{L}}_{00} = & A_{X_{20}} + \mathcal{Y}_{2c}^{X_{20}} A_{Y_{2c}} + \mathcal{Y}_{2s}^{X_{20}} A_{Y_{2s}} - Y_{1s} (\mathcal{Y}_{2c}^{X_{20}})' + Y_{1c} (\mathcal{Y}_{2s}^{X_{20}})' \\ & - (X_{1s} + Y_{1s} \mathcal{Y}_{2c}^{X_{20}} - Y_{1c} \mathcal{Y}_{2s}^{X_{20}}) \frac{d}{d\varphi}.\end{aligned}\quad (\text{D5})$$

Following a similar procedure with the other components (which we spare the reader from), the remainder terms of $\hat{\mathbb{L}}$ may be found, as well as those of $\hat{\mathbb{F}}$ and $\mathbf{f}$.

D.3. Tokamak limit simplification

Even though the topology of $|\mathbf{B}|$ contours over flux surfaces is different in tokamaks compared with QI configurations, the equilibrium construction remains largely the same. For the tokamak limit expressions evaluated in the main text, it is then appropriate to consider the abovementioned tokamak limit of the operators. The toroidal coordinate φ being one of symmetry simplifies the evaluation of the expressions significantly, especially because the toroidal derivatives may be taken to vanish exactly.

In the near-axis tokamak notation, we also have $X_{1c} = \bar{d}$ and $X_{1s} = 0$, with $Y_{1s} = 1/\bar{d}$, and define the major radius $R_0 = G_0/B_0$. Taking this into consideration and going through the equations, the following is true in the tokamak limit. The operator $\hat{\mathbb{L}}$ becomes

$$\hat{\mathbb{L}}_{\text{tok}} = -\bar{l}_0 \frac{B_0}{\bar{d}} \begin{pmatrix} \bar{d}^2 - \frac{3}{\bar{d}^2}(1 + \sigma^2)4\sigma & \\ 0 & 2 \end{pmatrix}, \quad (\text{D6})$$

which becomes diagonal in the case of an up-down symmetric tokamak (i.e. $\sigma = 0$). Note here the possibility of a singular matrix if, in the up-down symmetric scenario, $\bar{d}^4 = 3$. This singularity was discussed in the main text and is directly related to how the equilibrium problem is being solved. This is however seldom a problem given the typical values of elongation considered.

For $\hat{\mathbb{F}}$,

$$\hat{\mathbb{F}}_{\text{tok}} = -\bar{l}_0 \frac{3B_0}{\bar{d}^3} \begin{pmatrix} \bar{d}^4 - 1 + \sigma^2 & 2\sigma \\ 2\sigma & -(\bar{d}^4 - 1 + \sigma^2) \end{pmatrix}, \quad (\text{D7})$$

which is a symmetric matrix. Finally,

$$\mathbf{f}_{\text{tok}} = \frac{B_0 \bar{l}_0}{\bar{d} R_0} \begin{pmatrix} \frac{1}{2}[1 + 2(\bar{d}^4 + \sigma^2)] - \left(\frac{\bar{d}}{\bar{l}_0}\right)^2 \frac{\mu_0 p_2}{B_0^2/2} & \\ 5 & \frac{5}{2}\sigma \end{pmatrix}. \quad (\text{D8})$$

The matrix operators can be easily inverted and thus used for the shape gradient calculations in the paper.

D.4. Triangularity

It is often illuminating to use more geometrical definitions of the shaping other than those directly involved in the near-axis construction. At second order, we may define an approximation to triangularity following Rodríguez (2023),

$$\delta = 2 \left(\frac{Y_{2s}}{Y_{1s}} - \frac{X_{2c}}{X_{1c}} \right), \quad (\text{D9})$$

which through (D1) depends on all X_{20} , X_{2s} , X_{2c} and Y_{20} . For the purpose of this paper, we are interested in the variations of this expression with respect to the second-order shape choices. That is,

$$\delta\delta = -\frac{2}{Y_{1s}} \left[\begin{pmatrix} \mathcal{Y}_{2s}^{X_{20}} \\ \mathcal{Y}_{2s}^{Y_{20}} \end{pmatrix}^T \hat{\mathbb{L}}^{-1} \hat{\mathbb{F}} - \begin{pmatrix} \mathcal{Y}_{2s}^{X_{2c}} \\ \mathcal{Y}_{2s}^{X_{2s}} \end{pmatrix}^T \right] \begin{pmatrix} \delta X_{2c} \\ \delta X_{2s} \end{pmatrix} - \frac{2}{X_{1c}} \delta X_{2c}, \quad (\text{D10})$$

which is rather complicated. In the tokamak limit, using the expressions in (D6) and (D7), the expression simplifies significantly to give

$$\delta\delta_{\text{tok}} = \frac{2}{\bar{d}} \frac{\bar{d}^4 + 3}{\bar{d}^4 - 3} \delta X_{2c}, \quad (\text{D11})$$

as used in the main text. Note the appearance of the diverging denominator here, which signals the huge variation in the geometric triangularity due to a change in the input function X_{2c} when close to the resonance, emphasising the perilous balance there.

Appendix E. General sensitivity theory

In this appendix, we consider the evaluation of the shape gradient of some functional S with respect to the shaping degrees of freedom at second order. Let such a functional be of the form,

$$S[B_{20}, B_{2c}, B_{2s}] = \int_0^{2\pi} f(B_{20}, B_{2c}, B_{2s}, \varphi) d\varphi, \quad (\text{E1})$$

a functional that depends directly on the second-order magnetic field functions. As we change the second-order shaping, we are interested in how S changes, call it δS . Assuming f to be differentiable with respect to the three components of the second-order magnetic field magnitude,

$$\delta S = \int_0^{2\pi} \left(\frac{\partial f}{\partial B_{20}} \delta B_{20} + \frac{\partial f}{\partial B_{2c}} \delta B_{2c} + \frac{\partial f}{\partial B_{2s}} \delta B_{2s} \right) d\varphi, \quad (\text{E2})$$

and the partial derivatives are meant to be taken keeping the other second-order magnetic field components fixed.¹⁵ To proceed further, we must relate the variation of these magnetic field components to the shaping. To that end, we draw from the presentation of the second-order construction in § 2.4 of Rodriguez *et al.* (2025) and the explicit equations in Appendix A of Landreman & Sengupta (2019).

The shaping freedom is on the functions X_{2s} and X_{2c} , which represent in some form the triangularity that we may exploit to modify S . More precisely, they describe the $m = 2$ variation of the distance of the flux surfaces from the magnetic axis along the normal of the latter. As the flux surface is deformed, the magnetic field magnitude must change accordingly, and thus B_{2c} and B_{2s} are quite directly modified. From (2.7) of Rodriguez *et al.* (2025) or (A35)–(A36) in Landreman & Sengupta (2019),

$$\delta B_{2c} = \kappa B_0 \delta X_{2c}, \quad (\text{E3})$$

$$\delta B_{2s} = \kappa B_0 \delta X_{2s}. \quad (\text{E4})$$

¹⁵It truly is a variation in the functional theory sense.

A larger curvature of the magnetic axis results in a stronger supported perpendicular gradient of the magnetic field, and thus the easier the magnitude of the magnetic field can be shaped.

The variation of B_{20} ((A34) of Landreman & Sengupta (2019)) depends directly on X_{20} , which must be self-consistently solved alongside Y_{20} in (D3). The dependence on X_{2c} and X_{2s} (and their derivatives) do therefore involve the inversion of differential operators, so in the notation of Appendix D.2,

$$\delta B_{20} = -\kappa B_0 (10) \hat{\mathbb{L}}^{-1} \hat{\mathbb{F}} \begin{pmatrix} \delta X_{2c} \\ \delta X_{2s} \end{pmatrix}. \quad (\text{E5})$$

Once again, we encounter the proportionality to κ , which denotes the partial inability of the shaping to modify the field, especially near inflexion points. However, this does not mean that the shaping at these points has no influence on S . Effects lose locality through the inversion of $\hat{\mathbb{L}}$.

With this, we write

$$\delta S = \int_0^{2\pi} \kappa B_0 \underbrace{\left[- \begin{pmatrix} \frac{\partial f}{\partial B_{20}} & 0 \end{pmatrix} \hat{\mathbb{L}}^{-1} \hat{\mathbb{F}} + \begin{pmatrix} \frac{\partial f}{\partial B_{2c}} & \frac{\partial f}{\partial B_{2s}} \end{pmatrix} \right]}_{=\mathcal{G}^T} \delta X_2 \, d\varphi, \quad (\text{E6})$$

where $\mathcal{G}$ can be thought of as the shape gradient $\mathcal{G} = (\mathcal{G}_{2c}, \mathcal{G}_{2s})$ and $\delta X_2 = (\delta X_{2c}, \delta X_{2s})$. The function $\mathcal{G}$ can then be interpreted as the so-called shape gradient of S with respect to the shaping X_2 . That is, it measures the amount by which an infinitesimal local variation of second-order shaping would change the non-local value of S .

A priori, and with all the quantities in $\mathcal{G}$ being known (they belong to the first-order NAE), we could directly calculate it. However, it is in practice convenient to avoid having to invert $\hat{\mathbb{L}}$ explicitly.¹⁶ Instead, we may formulate the calculation of $\mathcal{G}$ through an adjoint problem. For simplicity, let us discretise the equation, as we will do anyway in practice. Defining a grid $\{\varphi_0, \varphi_1, \dots, \varphi_{N-1}\}$ and discretising integrals with the appropriate numerical scheme weight $\int f \, d\varphi \approx \sum w_i f(\varphi_i)$, we write

$$\delta S = \sum_{i=0}^{2(N-1)} w_i \mathcal{G}_i \delta X_{2,i}, \quad (\text{E7})$$

where

$$\mathcal{G}_i = -\frac{1}{w_i} (\mathbb{F}^T)_{ij} \mathbf{y}_j + \kappa B_0 \begin{pmatrix} \frac{\partial f}{\partial B_{2c}} \\ \frac{\partial f}{\partial B_{2s}} \end{pmatrix}, \quad (\text{E8})$$

and $\mathbf{y}$ is the solution to the following adjoint problem:

$$\mathbb{L}^T \mathbf{y} = \begin{pmatrix} \kappa B_0 w \frac{\partial f}{\partial B_{20}} \\ 0 \end{pmatrix}. \quad (\text{E9})$$

¹⁶Avoiding explicit inversion of the matrix in solving a linear system tends to be numerically faster and more stable (Cormen *et al.* 2022, Chap. 28).

Thus, to compute the shape gradient $\mathcal{G}$, the key is to perform an adjoint solve to find $\mathbf{y}$. The system to solve is nevertheless almost the same as for the direct second-order solve problem, (D3), as it involves $\mathbb{L}$, and thus it involves a complexity analogous to a second-order solve itself. This construction presented in generality may be used for a variety of different measures.

The answer will depend on how we discretise the integral and thus on the weights w . This is the price to pay for not having to integrate by parts. In its most basic form, using the standard trapezoidal scheme, the weight (Süli & Mayers 2003, p. 202)

$$w_i = \begin{cases} \frac{1}{2} (P + \varphi_1 - \varphi_{N-1}), & (i = 0) \\ \frac{1}{2} (\varphi_{i+1} - \varphi_{i-1}), & (0 < i < N) \end{cases} \quad (\text{E10})$$

where P is the period of the domain (2π if the whole toroidal angle extent is considered).

E.1. Tokamak limit

Let us consider the simplifying limit of a tokamak, where the toroidal derivatives vanish, and thus the procedure is rather simple (in fact, without the need to use a particular collocation grid). In this limit, using the expressions in Appendix D.3, we may write for the magnetic well,

$$\mathcal{G}^W = -\kappa B_0 \frac{\partial f}{\partial B_{20}} (\mathbb{L}^{-1} \mathbb{F})^T \begin{pmatrix} 1 \\ 0 \end{pmatrix} = -3\kappa B_0 \frac{\partial f}{\partial B_{20}} \begin{pmatrix} \frac{\bar{d}^4 - 1 - 3\sigma^2}{\bar{d}^4 - 3(1 + \sigma^2)} \\ -\frac{2\sigma(\sigma^2 + \bar{d}^4)}{3(1 + \sigma^2) - \bar{d}^4} \end{pmatrix}. \quad (\text{E11})$$

Using the derivative of f for W explicitly, (6.5), and taking the up-down symmetric limit for simplicity ($\sigma = 0$), we are left with

$$\mathcal{G}^W = \frac{3r_{\text{ref}}^2}{\pi R} \frac{\bar{d}^4 - 1}{\bar{d}^4 - 3} \begin{pmatrix} 1 \\ 0 \end{pmatrix}. \quad (\text{E12})$$

This is the expression presented in the main text, (6.6).

Appendix F. Variational re-shaping problem

In this appendix, we show how the variational problem of § 6 leads to a very particular form for the shaping inputs at second order in the near-axis construction that allows us to reliably construct stable approximately QI fields. Let us rewrite here the equation in the main text,

$$\begin{aligned} T[X_{2c}, X_{2s}] &= \frac{1}{2\pi} \int_0^{2\pi} \left(X_{20}^2 + Y_{20}^2 + \frac{X_{2c}^2 + X_{2s}^2 + Y_{2c}^2 + Y_{2s}^2}{2} \right) d\varphi - \\ &\quad - \frac{\lambda}{2\pi} \left[\int_0^{2\pi} (\mathcal{G}_{2c}^W X_{2c} + \mathcal{G}_{2s}^W X_{2s}) d\varphi + W_{\text{ref}} \right], \end{aligned} \quad (\text{6.8})$$

of which we need to take variations. We shall construct the variational considerations of the discretised form of the problem. We could do it in the continuous case, but

that would involve integrating by parts, and ultimately, in practice, a discrete version is really needed.

Let us then focus on writing T in the discrete form, thereon, a vector $\mathbf{v} = (a(\varphi), b(\varphi))$ is defined as an element of $\mathbb{R}^{2N}$, with $(a_0, a_1, \dots, a_{N-1}, b_0, \dots, b_{N-1})$, where the subscripts correspond to the function evaluated at φ_i , the discretised φ -grid. Writing the weight of the discretised integral as $\mathbf{W} = \text{diag}(\mathbf{w}, \mathbf{w})$, where for the trapezoidal rule, $\mathbf{w}$ is defined in (E10), and defining a convenient scalar product $\langle \mathbf{a}, \mathbf{b} \rangle_{\mathbf{W}} = \mathbf{a}^T \mathbf{W} \mathbf{b}$, we may then write (6.8) as

$$2\pi T[X_{2c}, X_{2s}] = \frac{1}{2} \left\langle \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix}, \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix} \right\rangle_{\mathbf{W}} + \left\langle \begin{pmatrix} X_{20} \\ Y_{20} \end{pmatrix}, \begin{pmatrix} X_{20} \\ Y_{20} \end{pmatrix} \right\rangle_{\mathbf{W}} + \frac{1}{2} \left\langle \begin{pmatrix} Y_{2c} \\ Y_{2s} \end{pmatrix}, \begin{pmatrix} Y_{2c} \\ Y_{2s} \end{pmatrix} \right\rangle_{\mathbf{W}} - \lambda \left[\left\langle \begin{pmatrix} \mathcal{G}_{2c}^W \\ \mathcal{G}_{2s}^W \end{pmatrix}, \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix} \right\rangle_{\mathbf{W}} + W_{\text{ref}} \right]. \quad (\text{F1})$$

We must now take variations of these expressions with respect to the second-order input functions X_{2c} and X_{2s} . For the first term, this is trivial, but for the next two, we must take the variation through the equations that define first X_{20} and Y_{20} , and then Y_{2c} and Y_{2s} . For the former, we need (D3) and for the latter, both that and (D1). We may summarise these variations as

$$\begin{pmatrix} \delta X_{20} \\ \delta Y_{20} \end{pmatrix} = - \underbrace{\hat{\mathbb{L}}^{-1} \mathbb{F}}_{\mathbf{X}} \begin{pmatrix} \delta X_{2c} \\ \delta X_{2s} \end{pmatrix}, \quad \begin{pmatrix} \delta Y_{2c} \\ \delta Y_{2s} \end{pmatrix} = \underbrace{(\hat{\mathbf{Y}} - \bar{\mathbf{Y}} \hat{\mathbb{L}}^{-1} \mathbb{F})}_{\mathbf{Y}} \begin{pmatrix} \delta X_{2c} \\ \delta X_{2s} \end{pmatrix}, \quad (\text{F2})$$

with all matrices defined in Appendix D.

Expressing all the remaining terms explicitly in terms of the second-order input functions, the variation δT can be written as

$$2\pi \delta T = \left\langle \mathbf{M} \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix} - \mathbf{\Lambda} - \lambda \mathbf{G}, \begin{pmatrix} \delta X_{2c} \\ \delta X_{2s} \end{pmatrix} \right\rangle_{\mathbf{W}}, \quad (\text{F3})$$

where

$$\mathbf{G} = \begin{pmatrix} \mathcal{G}_{2c}^W \\ \mathcal{G}_{2s}^W \end{pmatrix}, \quad (\text{F4a})$$

$$\mathbf{M} = \mathbf{1} + 2\mathbf{W}^{-1} \mathbf{X}^T \mathbf{W} \mathbf{X} + \mathbf{W}^{-1} \mathbf{Y}^T \mathbf{W} \mathbf{Y}, \quad (\text{F4b})$$

$$\mathbf{\Lambda} = [2\mathbf{W}^{-1} \mathbf{X}^T \mathbf{W} - \mathbf{W}^{-1} \mathbf{Y}^T \mathbf{W} \bar{\mathbf{Y}}] \mathbb{L}^{-1} \mathbf{f} - \mathbf{W}^{-1} \mathbf{Y}^T \mathbf{W} \mathcal{Y}_0, \quad (\text{F4c})$$

and it should be clear that $\mathbf{M}$ is a matrix of dimensions $2N \times 2N$, while $\mathbf{\Lambda}$ is a $2N$ vector.

We are then in a position to require the variation (F3) to vanish for all variations, implying

$$\mathbf{M} \begin{pmatrix} X_{2c} \\ X_{2s} \end{pmatrix} = \mathbf{\Lambda} + \lambda \mathbf{G}. \quad (\text{F5})$$

Using this in the constraint equation leads to an expression for the Lagrange multiplier λ ,

$$\lambda = - \frac{W_{\text{ref}} + \langle \mathbf{G}, \mathbf{M}^{-1} \mathbf{\Lambda} \rangle_{\mathbf{W}}}{\langle \mathbf{G}, \mathbf{M}^{-1} \mathbf{G} \rangle_{\mathbf{W}}}. \quad (\text{F6})$$

With that, the required precise form of the second-order shaping is

$$\begin{pmatrix} X_{2c}^{\text{mhd}} \\ X_{2s}^{\text{mhd}} \end{pmatrix} = \mathbf{M}^{-1} (\mathbf{\Lambda} + \lambda \mathbf{G}). \quad (\text{F7})$$

This is the answer we sought.

To gauge the computational complexity involved to reach the necessary shaping in (F7), we may consider roughly the operations that will be needed. Starting from the last equation, (F7), we here will need two linear system solves. These very operations may be recycled in the calculation of the Lagrange multiplier in (F6). However, we are hiding equally costly operations in the form of matrix multiplication and inversion (which are numerically of equal complexity (Golub & Van Loan 2013, Chap. 3; Cormen *et al.* 2022, Chap. 28), which are needed to evaluate both $\mathbf{M}$ and $\mathbf{\Lambda}$. To avoid taking explicit inverses especially to avoid numerical instability, the necessary operation could be broadcasted from right to left, so that operations on vectors are used at all times. However, in practice, we see that a naïve implementation is well behaved.

Appendix G. Estimation of critical β

In § 7, we introduced a measure of a critical plasma β that leads to flux surfaces in the near-axis field construction touching. In this appendix, we present the details of said calculation and the exact meaning of the elements that go into it.

G.1. Problem set-up

The problem we want to describe is as follows. Consider a plane normal to the axis at some φ value, as we shall describe the problem slice by slice. To first order in r , the shape of the flux surfaces in such a plane corresponds to ellipses. At second order, these ellipses are rigidly shifted as well as triangularly shaped. For the estimation in this appendix, we will simply ignore that triangular shaping and model the cross-sections as ellipses that are rigidly shifted.

Let us be a bit more precise. The ellipses are defined in the signed Frenet–Serret frame (where X and Y correspond to Cartesian coordinates aligned with the normal and binormal directions, respectively) by

$$X/r = X_{1c} \cos \chi + X_{1s} \sin \chi, \quad (\text{G1a})$$

$$Y/r = Y_{1c} \cos \chi + Y_{1s} \sin \chi. \quad (\text{G1b})$$

This defines an implicit ellipse, following (Rodríguez 2023, B1)

$$(1 + \sigma^2)X^2 - 2XY\sigma\mathring{\mathcal{E}} + \mathring{\mathcal{E}}^2Y^2 = \frac{\bar{B}}{B_0}\mathring{\mathcal{E}}, \quad (\text{G2})$$

where $\mathring{\mathcal{E}} = B_0 \bar{d}^2 / \bar{B}$. Note that this form is analogous to the quasisymmetric case of Rodríguez (2023), where $\mathring{\mathcal{E}} = (\eta/\kappa)^2$. This is true with the exception of the right-hand side, due to the variation of the elliptical cross-section with the variation of the field strength B_0 in the current QI case. However, this means that the description of the ellipses is identical to that there. Namely, we may characterise an ellipse by its elongation $\mathcal{E} = \tan e$, defined as the ratio of major to minor radius, and θ , the angle that the major axis makes with X . Explicitly,

$$\sin 2e = \frac{2\mathring{\mathcal{E}}}{\mathring{\mathcal{E}}^2 + 1 + \sigma^2}, \quad \tan 2\theta = \frac{2\sigma\mathring{\mathcal{E}}}{\mathring{\mathcal{E}}^2 - (1 + \sigma^2)}. \quad (\text{G3})$$




FIGURE 9. Illustrating diagrams for the estimation of the critical β . (a) Elliptical cross-section indicating the rotation angle θ and major and minor axes, a and b , respectively. (b) Shift of ellipses along the direction D and the M -transformed scenario involving circles. The diagram shows the geometric meaning of d as defined in (G7).

We note that the inversion of the tangent must be carried out with some care to correctly capture the described physical angle.

Now, we consider these ellipses to be shifted rigidly at second order, by an amount that we may refer to as

$$D = r^2 \begin{pmatrix} X_{20} \\ Y_{20} \end{pmatrix}, \quad (\text{G4})$$

in the $(\hat{X}, \hat{Y})$ basis. Note that each value of r defines a different ellipse. The goal of this appendix is to precisely describe when this displacement is such that two consecutive ellipses touch.

G.2. Making circles

The problem of touching ellipses can be made significantly simpler by appropriately rotating and scaling the plane. First, we shall consider a clockwise rotation by θ , the angle of the ellipse, (G3), so that all ellipses become ‘aligned’ with the axes. By align, we mean that the major and minor radii of all ellipses are parallel (or orthogonal) to the signed Frenet–Serret coordinate system. Then we scale the major axis direction by its magnitude and similarly for the minor radius so that the result is a set of displaced circles of different radii. That is, we apply the linear map defined on this basis by

$$M = \begin{pmatrix} r/a & 0 \\ 0 & r/b \end{pmatrix} \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}, \quad (\text{G5})$$

to all (X, Y) , so that ellipses are mapped to circles of radius r . Using $\mathcal{E} = a/b$ and the definition of the toroidal flux $\psi = \text{Area} \times B_0/2\pi$, so that $ab = \bar{B}r^2/B_0$, then

$$a/r = \sqrt{\frac{\bar{B}}{B_0}} \mathcal{E}, \quad b/r = \sqrt{\frac{\bar{B}}{B_0 \mathcal{E}}}. \quad (\text{G6})$$

This transformation may be applied to the shift D , so that we have an effective distance $\hat{D} = MD$ (see figure 9).

Through this transformation, we have converted the problem of ellipses into a set of circles with varying radii displaced along the same direction by different amounts. The discussion of touching surfaces may thus be reduced to a discussion along this line.

G.3. *Touching circles*

Take the line on the rescaled plane defined by $\hat{\mathbf{D}}$ and define a function $f(r)$ that gives the distance from the origin to the point of intersection of the circle of radius r with this line. Thanks to all being nicely stacked circles,

$$f(r) = r \left(1 - r|\hat{\mathbf{D}}| \right), \quad (\text{G7})$$

where we have defined the point of intersection in the half-line in which circles are bunching.

The point at which circles will touch corresponds to $\partial_{rf} = 0$, which gives

$$A_{\Delta} = 2R|\hat{\mathbf{D}}|, \quad (\text{G8})$$

where A corresponds to the aspect ratio with R the effective major radius defined in terms of the length of the magnetic axis. Below this aspect ratio, flux surfaces will intersect and the near-axis construction is invalid.

G.4. *Critical β*

The question is then how much does this critical aspect ratio change with the pressure in the problem? We know from the main text,

$$\frac{\delta \mathbf{D}}{\delta \beta} = \frac{\hat{\mathbf{S}}}{r_{\text{ref}}}, \quad (\text{G9})$$

where $\beta = \beta_2 r_{\text{ref}}^2$, and because it is linear in β , we may write exactly

$$\mathbf{D}_{\beta} = \mathbf{D}_0 + \frac{\delta \mathbf{D}}{\delta \beta} \beta, \quad (\text{G10})$$

where the subscript zero denotes the Shafranov shift at zero plasma beta.¹⁷

Let us now define the critical β_{Δ} as the value of β at which the touching of surfaces occurs at a reference aspect ratio $A_{\text{ref}} \approx 10$. Putting (G8) and (G10) together,

$$A_{\text{ref}}^2 \stackrel{!}{=} \left(\hat{\mathbf{D}}_0 + \frac{\delta \hat{\mathbf{D}}}{\delta \beta} \beta_{\Delta} \right)^2, \quad (\text{G11})$$

where the hats indicate transformed quantities by the matrix $\mathbf{M}$, (G5). The critical β_{Δ} then corresponds to the largest negative root of this equation. That is, for every φ , we may define

$$\beta_{\Delta}(\varphi) = \min \left[0, \beta_{\pm} : \left(\frac{\delta \hat{\mathbf{D}}}{\delta \beta} \right)^2 \beta_{\pm}^2 + 2 \frac{\delta \hat{\mathbf{D}}}{\delta \beta} \cdot \hat{\mathbf{D}}_0 \beta_{\pm} + \left(\hat{\mathbf{D}}_0^2 - \frac{1}{4r_{\text{ref}}^2} \right) = 0 \right], \quad (\text{G12})$$

and $\beta_{\pm}$ may be computed explicitly with the quadratic formula. The single scalar β_{Δ} is then the minimum value of this function in the whole φ domain.

¹⁷If a known construction has a finite beta β_0 , then one may simply modify the above with $\beta \mapsto \beta - \beta_0$. The problem is linear.

REFERENCES

- BENDER, C.M. & ORSZAG, S.A. 2013 *Advanced Mathematical Methods for Scientists and Engineers I: Asymptotic Methods and Perturbation Theory*. Springer Science & Business Media.
- BERNARDIN, M.P., MOSES, R.W. & TATARONIS, J.A. 1986 Isodynamical (omnigenous) equilibrium in symmetrically confined plasma configurations. *Phys. Fluids* **29**, 2605–2611.
- BLANK, H.DE 2004 Guiding center motion. *Fusion Sci. Technol.* **45**, 47–54.
- BOOZER, A.H. 1981 Plasma equilibrium with rational magnetic surfaces. *Phys. Fluids* **24**, 1999–2003.
- BOOZER, A.H. 1983 Transport and isomorphic equilibria. *Phys. Fluids* **26**, 496–499.
- BOOZER, A.H. 1998 What is a stellarator? *Phys. Plasmas* **5**, 1647–1655.
- BURBY, J.W., KALLINIKOS, N. & MACKAY, R.S. 2020 Some mathematics for quasi-symmetry. *J. Math. Phys.* **61**, 093503.
- CAMACHO MATA, K. & PLUNK, G.G. 2023 Helicity of the magnetic axes of quasi-isodynamic stellarators. *J. Plasma Phys.* **89**, 905890609.
- CAMACHO MATA, K., PLUNK, G.G. & JORGE, R. 2022 Direct construction of stellarator-symmetric quasi-isodynamic magnetic configurations. *J. Plasma Phys.* **88**, 905880503.
- CARY, J.R. & SHASHARINA, S.G. 1997 Omnigenity and quasihelicity in helical plasma confinement systems. *Phys. Plasmas* **4**, 3323–3333.
- CORMEN, T.H., LEISERSON, C.E., RIVEST, R.L. & STEIN, C. 2022 *Introduction to Algorithms*. MIT Press.
- D'HAESLEER, W.D., HITCHON, W.N.G., CALLEN, J.D. & SHOHET, J.L. 2012 *Flux Coordinates and Magnetic Field Structure: A Guide to a Fundamental Tool of Plasma Theory*. Springer Science & Business Media.
- DUDT, D.W. & KOLEMEN, E. 2020 DESC: A stellarator equilibrium solver. *Phys. Plasmas* **27**, 102513.
- DUDT, D.W., GOODMAN, A.G., CONLIN, R., PANICI, D. & KOLEMEN, E. 2024 Magnetic fields with general omnigenity. *J. Plasma Phys.* **90**, 905900120.
- FREIDBERG, J.P. 2014 *Ideal MHD*. Cambridge University Press.
- GARREN, D.A. & BOOZER, A.H. 1991a Existence of quasihelically symmetric stellarators. *Phys. Fluids B: Plasma Phys.* **3**, 2822–2834.
- GARREN, D.A. & BOOZER, A.H. 1991b Magnetic field strength of toroidal plasma equilibria. *Phys. Fluids B: Plasma Phys.* **3**, 2805–2821.
- GOLUB, G.H., VAN, L. & CHARLES, F. 2013 *Matrix Computations*. JHU Press.
- GOODMAN, A.G. *et al.* 2023 Constructing precisely quasi-isodynamic magnetic fields. *J. Plasma Phys.* **89**, 905890504.
- GRADSHTEYN, I.S. & RYZHIK, I.M. 2014 *Table of Integrals, Series, and Products*. Academic Press.
- GREENE, J.M. 1997 A brief review of magnetic wells. *Comment. Plasma Phys. Controll. Fusion* **17**, 389–402.
- HALL, L.S. & MCNAMARA, B. 1975 Three-dimensional equilibrium of the anisotropic, finite-pressure guiding-center plasma: theory of the magnetic plasma. *Phys. Fluids* **18**, 552–565.
- HELANDER, P. 2014 Theory of plasma confinement in non-axisymmetric magnetic fields. *Rep. Prog. Phys.* **77**, 087001.
- HELANDER, P. & NÜHRENBURG, J. 2009 Bootstrap current and neoclassical transport in quasi-isodynamic stellarators. *Plasma Phys. Controll. Fusion* **51**, 055004.
- HINDENLANG, F., MAJ, O., STRUMBERGER, E., RAMPP, M. & SONNENDRÜCKER, E. 2019 *Gvec: A Newly Developed 3d Ideal Mhd Galerkin Variational Equilibrium Code*. Simons Collaboration on Hidden Symmetries and Fusion Energy.
- HIRSHMAN, S.P. & WHITSON, J.C. 1983 Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria. *Phys. Fluids* **26**, 3553–3568.
- HO, D.D.-M. & KULSRUD, R.M. 1987 Neoclassical transport in stellarators. *Phys. Fluids* **30**, 442–461.
- JORGE, R. & LANDREMAN, M. 2020 The use of near-axis magnetic fields for stellarator turbulence simulations. *Plasma Phys. Controll. Fusion* **63**, 014001.
- JORGE, R., PLUNK, G.G., DREVLAK, M., LANDREMAN, M., LOBSIEN, J.-F., CAMACHO MATA, K. & HELANDER, P. 2022 A single-field-period quasi-isodynamic stellarator. *J. Plasma Phys.* **88**, 175880504.

- KAPPEL, J., LANDREMAN, M. & MALHOTRA, D. 2024 The magnetic gradient scale length explains why certain plasmas require close external magnetic coils. *Plasma Phys. Controll. Fusion* **66**, 025018.
- LANDREMAN, M. 2021 Figures of merit for stellarators near the magnetic axis. *J. Plasma Phys.* **87**, 905870112.
- LANDREMAN, M. 2022 Mapping the space of quasisymmetric stellarators using optimized near-axis expansion. *J. Plasma Phys.* **88**, 905880616.
- LANDREMAN, M. & JORGE, R. 2020 Magnetic well and Mercier stability of stellarators near the magnetic axis. *J. Plasma Phys.* **86**, 905860510.
- LANDREMAN, M. & SENGUPTA, W. 2019 Constructing stellarators with quasisymmetry to high order. *J. Plasma Phys.* **85**, 815850601.
- LITTLEJOHN, R.G. 1983 Variational principles of guiding centre motion. *J. Plasma Phys.* **29**, 111–125.
- LORTZ, D. & NÜHRENBURG, J. 1976 Equilibrium and stability of a three-dimensional toroidal mhd configuration near its magnetic axis. *Zeitschrift für Naturforschung A* **31**, 1277–1288.
- MERCIER, C. 1964 Equilibrium and stability of a toroidal magnetohydrodynamic system in the neighbourhood of a magnetic axis. *Nucl. Fusion* **4**, 213–226.
- MYNICK, H.E. 1983 Improved theory of collisionless particle motion in stellarators. *Phys. Fluids* **26**, 1008–1017.
- MYNICK, H.E. 2006 Transport optimization in stellarators. *Phys. Plasmas* **13**, 058102.
- NEMOV, V.V., KASILOV, S.V., KERNBICHLER, W. & HEYN, M.F. 1999 Evaluation of $1/\nu$ neoclassical transport in stellarators. *Phys. Plasmas* **6**, 4622–4632.
- NEMOV, V.V., KASILOV, S.V., KERNBICHLER, W. & LEITOLD, G.O. 2008 Poloidal motion of trapped particle orbits in real-space coordinates. *Phys. Plasmas* **15**, 052501.
- NORTHROP, T.G. 1961 The guiding center approximation to charged particle motion. *Ann. Phys.* **15**, 79–101.
- NÜHRENBURG, J. & ZILLE, R. 1988 Quasi-helically symmetric toroidal stellarators. *Phys. Lett. A* **129**, 113–117.
- NÜHRENBURG, J. 2010 Development of quasi-isodynamic stellarators. *Plasma Phys. Controll. Fusion* **52**, 124003.
- PARRA, F.I., CALVO, I., HELANDER, P. & LANDREMAN, M. 2015 Less constrained omnigenous stellarators. *Nucl. Fusion* **55**, 033005.
- PAUL, E.J., BHATTACHARJEE, A., LANDREMAN, M., ALEX, D., VELASCO, J.L. & NIES, R. 2022 Energetic particle loss mechanisms in reactor-scale equilibria close to quasisymmetry. *Nucl. Fusion* **62**, 126054.
- PLUNK, G.G., DREVLAK, M., RODRIGUEZ, E., BABIN, R., GOODMAN, A. & HINDENLANG, F. 2024 Back to the figure-8 stellarator. arXiv:2411.16411.
- PLUNK, G.G., LANDREMAN, M. & HELANDER, P. 2019 Direct construction of optimized stellarator shapes. Part 3. Omnigenity near the magnetic axis. *J. Plasma Phys.* **85**, 905850602.
- PLUNK, G.G. *et al.* 2025 A geometric approach to constructing quasi-isodynamic fields. In preparation.
- RODRÍGUEZ, E. 2023 Magnetohydrodynamic stability and the effects of shaping: a near-axis view for tokamaks and quasisymmetric stellarators. *J. Plasma Phys.* **89**, 905890211.
- RODRÍGUEZ, E. & PLUNK, G.G. 2023 Higher order theory of quasi-isodynamicity near the magnetic axis of stellarators. *Phys. Plasmas* **30**, 062507.
- RODRÍGUEZ, E., HELANDER, P. & BHATTACHARJEE, A. 2020 Necessary and sufficient conditions for quasisymmetry. *Phys. Plasmas* **27**, 062501.
- RODRÍGUEZ, E., HELANDER, P. & GOODMAN, A.G. 2024 The maximum- j property in quasi-isodynamic stellarators. *J. Plasma Phys.* **90**, 905900212.
- RODRÍGUEZ, E. & MACKENBACH, R.J.J. 2023 Trapped-particle precession and modes in quasisymmetric stellarators and tokamaks: a near-axis perspective. *J. Plasma Phys.* **89**, 905890521.
- RODRÍGUEZ, E., PAUL, E.J. & BHATTACHARJEE, A. 2022a Measures of quasisymmetry for stellarators. *J. Plasma Phys.* **88**, 905880109.
- RODRÍGUEZ, E. & PLUNK, G.G. 2025 The zonal-flow residual does not tend to zero in the limit of small mirror ratio. *J. Plasma Phys.* **91**, E102.

- RODRIGUEZ, E., PLUNK, G.G. & JORGE, R. 2025 Near-axis description of stellarator-symmetric quasi-isodynamic stellarators to second order. *J. Plasma Phys.* **91**, E59.
- RODRÍGUEZ, E., SENGUPTA, W. & BHATTACHARJEE, A. 2022*b* Phases and phase-transitions in quasisymmetric configuration space. *Plasma Phys. Controll. Fusion* **64**, 105006.
- RODRÍGUEZ, E., SENGUPTA, W. & BHATTACHARJEE, A. 2023 Constructing the space of quasisymmetric stellarators through near-axis expansion. *Plasma Phys. Controll. Fusion* **65**, 095004.
- ROSENBLUTH, M.N. & HINTON, F.L. 1998 Poloidal flow driven by ion-temperature-gradient turbulence in tokamaks. *Phys. Rev. Lett.* **80**, 724–727.
- SANCHEZ, R., HIRSHMAN, S.P., WARE, A.S., BERRY, L.A. & SPONG, D.A. 2000 Ballooning stability optimization of low-aspect-ratio stellarators. *Plasma Phys. Controll. Fusion* **42**, 641–653.
- SOLOV'EV, L.S. & SHAFRANOV, V.D. 1970 In *Reviews of Plasma Physics*, vol. 5. Consultants Bureau.
- Jr SPITZER, L. 1958 The stellarator concept. *Phys. Fluids* **1**, 253–264.
- STEIN, E.M. & SHAKARCHI, R. 2011 *Fourier Analysis: an Introduction*. vol. 1. Princeton University Press.
- SÜLI, E. & MAYERS, D.F. 2003 *An Introduction to Numerical Analysis*. Cambridge University Press.
- WESSON, J. 2011 Tokamaks. In *International Series of Monographs On Physics*. 4th edn Oxford Univ. Press.
- WEYL, H. 1916 Über die gleichverteilung von zahlen mod. eins. *Mathematische Annalen* **77**, 313–352.
- ZHU, H., LIN, Z. & BHATTACHARJEE, A. 2025 Collisionless zonal-flow dynamics in quasisymmetric stellarators. *J. Plasma Phys.* **91**, E28.