

Spatial Structure of
Micro-instabilities
in Tokamak Plasmas

Zonal Flows and Global Effects
in Local Gyrokinetic Simulations

Stephen Neil Biggs-Fox

PhD

University of York

Physics

April 2022

Abstract

Micro-instabilities, turbulence, and zonal flows are central to achievable tokamak performance. Knowledge of their poloidal structure would improve theoretical models and experimental interpretation.

The poloidal structure of micro-instabilities may be altered by global effects. Global simulations can be computationally expensive, but global behaviour can be recovered from an array of local simulations via the “local-global method”. This method is, for the first time, demonstrated in the pedestal.

Local gyrokinetic simulations of a JET pedestal reveal that magnetic shear causes narrowing in ballooning angle of kinetic ballooning modes via ideal ballooning physics. Narrowing in ballooning angle is shown to decrease local accuracy, but this may be mitigated by increased shear. The local-global method is applied to a low shear pedestal-like case and compared to global gyrokinetic and MHD simulations. Good agreement suggests the local-global method is valid for toroidal mode numbers $\gtrsim 3$ to 12. Simple models show that global and kinetic effects can affect EPED-like calculations by 0 to 110% in this case depending on the Peeling-Ballooning constraint.

Experimental measurement of zonal flows is difficult due to the limited poloidal extent of relevant diagnostics. Knowledge of the poloidal structure of zonal flow drive would improve interpretation of such data. Nonlinear energy transfer functions calculated from local nonlinear gyrokinetic simulations reveal the poloidal structure of zonal flow drive for the first time. This demonstrates that zonal flows are driven by a broad spectrum of turbulence, and that zonal flows exhibit a limit cycle oscillation (predator-prey) type response with marginal turbulence but enter a quasi-steady state with strong turbulence. In both cases, zonal flow drive peaks on the outboard side, and is correlated with but not proportional to the turbulent energy.

Contents

Title	1
Abstract	2
Contents	3
List of Figures	8
List of Tables	11
Acknowledgements	12
Declaration	13
Part I – Introduction and Background	15
Chapter 1: Introduction	16
1.1 Context	16
1.2 Motivation	19
1.2.1 Global Effects	19
1.2.2 Zonal Flows	21
1.3 Document Structure	22
Chapter 2: Background	24
2.1 Gyrokinetics	24
2.1.1 Introduction	24

2.1.2	Assumptions	25
2.1.3	The δf Approach	27
2.1.4	Coordinate Transform	28
2.1.5	The Gyrokinetic Equation	28
2.2	Ballooning Theory	29
2.2.1	Introduction	29
2.2.2	The Ballooning Transform	33
2.2.3	The Local Approximation	35
2.3	GS2	37
Part II – Global Effects on Micro-Instabilities		39
Chapter 3: The Local-Global Method		40
3.1	Introduction	40
3.2	Obtaining Global Behaviour from Local Simulations	41
Chapter 4: Kinetic Ballooning Mode Narrowing in Ballooning Angle		45
4.1	Introduction	45
4.2	Methods	47
4.3	Results	51
4.3.1	JET Pedestal KBMs Narrow in Ballooning Angle	51
4.4	Discussion of Ballooning Angle Narrowing Mechanism	54
4.4.1	Magnetic Geometry Effects on KBM Ballooning Angle Width	54
4.4.2	The Suydam and Mercier Criteria of Ideal Ballooning Modes	56
4.4.3	Mechanisms of Ideal Ballooning Instability	58
4.4.4	Ballooning Angle Narrowing of Ideal Ballooning Modes	61
4.4.5	Ballooning Angle and Magnetic Shear in Gyrokinetics	65

4.4.6	Additional Kinetic Effects	71
4.5	Discussion of Implications for Local Validity	73
4.5.1	The Relationship Between Ballooning Angle and Wavenumber	73
4.5.2	Conditions for Local Approximation Validity	74
4.6	Conclusions	79
Chapter 5: Local-Global Pedestal Simulations of Kinetic Ballooning Modes		82
5.1	Introduction	82
5.1.1	Chapter Overview	82
5.1.2	Motivation	83
5.1.3	Equilibrium Model	84
5.2	Assessment of the Local-Global Method	85
5.2.1	Local-Global Method – Practical Details	85
5.2.1.1	Running the Simulations	85
5.2.1.2	Obtaining Fit Parameters	87
5.2.2	Local-Global Results	91
5.2.2.1	Local Benchmarking	91
5.2.2.2	Global Mode Structure	93
5.2.2.3	Global Growth Rate	96
5.2.3	Discussion – Local-Global Comparison	98
5.2.3.1	Reduced δB_{\parallel} Model	98
5.2.3.2	Miller Equilibrium Model	99
5.2.3.3	Other Factors	99
5.2.4	Conclusions – Local-Global Comparison	105
5.3	Development of Simple Correction Models	106
5.3.1	Global Correction Model	107
5.3.2	Diamagnetic Correction Model	110
5.3.3	Combined Global Diamagnetic Correction Model	114

5.3.4	Conclusions – Simple Corrective Models	116
5.4	Global and Diamagnetic Correction Models in EPED	117
5.4.1	Method Overview	117
5.4.2	Local MHD Stability	119
5.4.3	β'_{crit} Correction	121
5.4.3.1	Determining the Change in Growth Rate	122
5.4.3.2	Measuring the Stiffness of KBM Onset	123
5.4.3.3	Calculating the Change to β'_{crit}	124
5.4.4	The Effect of the β'_{crit} Correction on the EPED ELM Prediction	126
5.4.4.1	EPED KBM Constraint	126
5.4.4.2	EPED PB Constraint	128
5.4.5	Discussion – Corrective Models in EPED	130
5.4.6	Conclusions – Corrective Models in EPED	133
5.5	Discussion – General	133
5.5.1	Boundary Issues in Global Gyrokinetic Codes	133
5.5.2	Validity of the Local Approximation	134
5.5.3	Future Work	136
5.6	Conclusions – General	137
 Part III – Poloidal Structure of Zonal Flow Drive		139
 Chapter 6: Poloidal Structure of Zonal Flow Drive		140
6.1	Introduction	140
6.2	Methods	146
6.2.1	Simulation Setup	146
6.2.2	Nonlinear Energy Transfer Functions	146
6.3	Results and Discussion	149
6.3.1	Development of Selected Cases	149

6.3.1.1	Nonlinear Benchmarking	149
6.3.1.2	Overview of Selected Cases	152
6.3.2	Transfer Function Spectrum	156
6.3.3	Time History	164
6.3.4	Poloidal structure	168
6.4	Conclusions	173
 Part IV – Conclusions		177
 Chapter 7: Conclusions		178
 Appendices		182
 Appendix A: The Fourier-Ballooning Representation		182
 Appendix B: Linear Convergence		184
 Appendix C: Measuring Frequency and Growth Rate, and Tracking Sub-dominant Modes with GS2		189
 Appendix D: Nonlinear Convergence		191
 Appendix E: Transfer Function Program		197
 References		198

List of Figures

Chapter 2: Background

1	Example flux tube simulation domain	32
---	---	----

Chapter 4: Kinetic Ballooning Mode Narrowing in Ballooning Angle

2	JET equilibrium profiles	48
3	Comparison of CBC and JET flux surface shapes	51
4	Comparison of CBC and JET ballooning space widths	52
5	Evidence that the JET modes are KBMs	53
6	KBM ballooning angle width variation with geometry	55
7	Confirmation of magnetic shear as the key parameter	56
8	Magnetic geometry effects at various ballooning angles	57
9	Illustrations of ideal ballooning physics	62
10	Ideal ballooning mode narrowing with magnetic shear	63
11	Modified s - α diagram as a result of IBM splitting	64
12	The effect of drift scaling at various ballooning angles	67
13	The effect of collisions at various ballooning angles	69
14	Evidence that KBMs are narrowed by ideal ballooning physics	70
15	A comparison of KBM growth rates against IBM stability	72
16	The effect of magnetic shear at various radial wavenumbers	74
17	Variation of local error with ballooning space width	78

Chapter 5: Local-Global Pedestal Simulations of Kinetic Ballooning Modes

18	Pedestal-like Equilibrium Overview	86
19	Local frequency and growth rate contour plots	88
20	Local frequency and growth rate line plots	88
21	Quadratic fit weighting function	91
22	Local benchmarking	92
23	Global mode structure	95
24	Local-global results	97
25	Global growth rate as a function of n_{eff}	107
26	Global correction model fit	108
27	Reverse application of global correction model	109
28	Variation of diamagnetic frequency	112
29	Diamagnetic correction model fit	113
30	Combined global diamagnetic correction model	115
31	Comparison of kinetic and ideal stability	120
32	Uncorrected critical pressure gradient	122
33	Measuring the stiffness of KBM onset	124
34	EPED KBM Constraints	128
35	EPED PB Constraints	131

Chapter 6: Poloidal Structure of Zonal Flow Drive

36	Nonlinear benchmarking via reproduction of the Dimits shift	150
37	Time-history of the cases selected for further analysis	153
38	Spectra of the cases selected for further analysis	155
39	Real-space snapshots of the cases selected for further analysis	155
40	Kinetic energy transfer radial wavenumber spectrum	157
41	Internal energy transfer radial wavenumber spectrum	159
42	Kinetic energy transfer binormal wavenumber spectrum	160

43	The spectrum of kinetic energy transfer into zonal flows . . .	160
44	The spectrum of internal energy transfer into zonal flows . .	161
45	Comparison of turbulence and transfer spectra	163
46	Turbulence and transfer time histories – turbulent case . . .	165
47	Zonal flow and transfer time histories – turbulent case	166
48	Turbulence, zonal flow and transfer time histories – zonal case	166
49	Poloidal structure of zonal flow drive – turbulent case	169
50	Fourier transform of the zonal flow poloidal structure	170
51	Poloidal structure of zonal flow drive – zonal case	172

Appendix B: Linear Convergence

52	Convergence of the JET equilibrium with δt	185
53	Convergence of the JET equilibrium with n_{grid}	185
54	Convergence of the JET equilibrium with n_{gauss}	186
55	Convergence of the JET equilibrium with n_{theta}	186
56	Convergence of the JET equilibrium with n_{period}	187
57	JET equilibrium local mode structures with $n_{period} = 4$. .	187
58	Comparison of CBC Miller and $s-\alpha$ results	188

Appendix D: Nonlinear Convergence

59	Linear benchmarking via wavenumber scan	192
60	Comparison of kinetic and adiabatic electron models	193
61	Nonlinear convergence	193
62	Nonlinear wavenumber grid resolution tests	194
63	The effect of collisions on nonlinear simulations	195
64	The effect of hyper-diffusion on nonlinear simulations	195
65	The effect of wavenumber range on nonlinear simulations . .	196

List of Tables

Chapter 4: Kinetic Ballooning Mode Narrowing in Ballooning Angle	
1 Comparison of parameters for the CBC and JET cases	50
Chapter 5: Local-Global Pedestal Simulations of Kinetic Ballooning Modes	
2 Change in growth rate, $\Delta\gamma$, for each ρ^*	123
3 Corrected critical pressure gradients	125
4 Corrected and uncorrected pedestal heights and widths	132
Appendix B: Linear Convergence	
5 JET equilibrium converged numerical parameters	187
Appendix D: Nonlinear Convergence	
6 Numerical parameter values	196

Acknowledgements

Thank you to all my friends, family, colleagues, mentors, supervisors, and collaborators for supporting me in completing this PhD. Special thanks to: my supervisor, Dr David Dickinson, for teaching me so much; my co-supervisor, Prof Howard Wilson, for helpful advice; my collaborator on Chapter 6, Dr Istvan Cziegler, for invaluable guidance; my Fusion CDT laboratory project collaborator, Prof Steve Tobias, for giving me a different perspective; my new employers at UKAEA, Dr Tom Barrett and Mr Tom Deighan, for being understanding about delays completing my thesis due to the pandemic; my high school physics teacher, Mr Jamie Costello, for bringing the subject to life; my undergraduate tutor, Prof Philip Moriarty, for inspiring me to follow a career in physics; Mum and Dad, for raising me right; my beautiful wife, Sophie, for encouraging me to follow my dreams; and our two wonderful sons, Dougie and Sid, for making it all worthwhile.

This work used: the ARCHER UK National Supercomputing Service (<http://www.archer.ac.uk>) via the Plasma HEC Consortium (EPSRC grant numbers EP/L000237/1 and EP/R029148/1); the EUROfusion High Performance Computer (Marconi-Fusion) via EUROfusion; and the University of York high performance compute facility, the Viking Cluster. This work was funded by the Fusion CDT (EPSRC grant number EP/L01663X/1). Special thanks to the University of York Covid-19 PhD Scholarship Fund for additional funding support.

Declaration

I declare that: this thesis is a presentation of original work; I am the sole author; this work has not previously been presented for an award at this or any other university; all sources are acknowledged as references.

This thesis uses the following software codes. The author's contribution to each code's development is noted below:

- GS2 was used to run gyrokinetic simulations. This was mostly used as-provided, although the author made some minor bug-fixes and general codebase improvement contributions, such as fixing compilation problems and initialisation bugs, removing dead or commented out code, etc. All of these contributions have been merged into the latest version of GS2, which is available online [1].
- `ideal_ball` – an auxiliary program within the GS2 codebase – was used to calculate ideal ballooning stability. The author made some minor bug-fix contributions (e.g. input file format compatibility) that have been merged into the latest version of GS2, which is available online [1]. The author also developed a feature to calculate ideal ballooning stability at arbitrary ballooning angle, η_0 (the as-provided code only calculated this at $\eta_0 = 0$); this feature is available in the branch `feature/ideal-ball-theta0`, which is available online [2].

- Bash scripts were used for submitting GS2 runs. Some of these were provided by the PhD supervisor, Dr David Dickinson, to create a run directory, create a job submission script, submit the job to the batch scheduler, etc. The author made minor modifications to these, and wrote additional scripts to check output files for errors and re-submit failed jobs, prepare a GS2 run for restarting, implement the eigensolver small step method described in Appendix C, etc. These scripts can be provided by request.
- Python code was used for GS2 analysis. This calculates frequencies, growth rates, etc., as described in Appendix C, and produces plots in a standardised format, etc. This was originally developed by the author for a previous degree (MSc Fusion Energy), and was further developed for this thesis. This code is available online [3].
- Python code was used to calculate global results from local simulations. This implements the method described in Chapter 3. Apart from a dependency on a differential operator solver written by the PhD supervisor, Dr David Dickinson, this code was originally developed by the author for a previous degree (MSc Fusion Energy), and was further developed for this thesis. The code is contained within Jupyter Notebooks and can be provided by request.
- Python code was used to calculate nonlinear energy transfer functions. The method implemented is described in Chapter 6 and the functionality of the program is described in Appendix E. All code was written by the author and is available online [4].
- Python code to call the various programs above, import data, perform analysis not covered above, and produce the plots shown in this thesis. All code was written by the author and is contained within Jupyter Notebooks that can be provided by request.

Part I

Introduction and Background

Chapter 1

Introduction

1.1 Context

It is widely accepted that Global Mean Temperature (GMT) has risen compared to pre-industrial levels, primarily due to human activity [5]. Each degree of warming increases the risk of adverse effects [6] including, for example: extinction of plant and animal species; considerable numbers of human lives lost and damage to property due to extreme weather events; displacement of populations due to coastal flooding; decreased availability of food and water; and negative economic impacts.

The rise in GMT is primarily driven by an increase in the concentration of atmospheric carbon dioxide (CO₂) [7, p13], the majority of which comes from burning fossil fuels for energy production [7, pp11–12]. Therefore, significantly reducing or eliminating fossil fuel use in power generation will greatly reduce CO₂ emissions and thus help to mitigate climate change. However, global energy demand is projected to rise significantly over the coming decades [8, p73]. To meet this demand while reducing fossil fuel usage, low-carbon energy sources are required.

Electricity production using nuclear fusion could provide safe, sustainable, widely-available, low-carbon baseload energy without long-lived high-level nuclear waste by the end of this century [9]. While significant progress has been made over previous decades [10], net energy gain has yet to be demonstrated.

The most advanced fusion energy concept, the tokamak [11], generates power by confining a plasma at the temperature (~ 100 million $^{\circ}\text{C}$) and density ($\sim 10^{20} \text{ m}^{-3}$) required for an appreciable fusion reaction rate. Confinement of the plasma is achieved using a magnetic field. A toroidal topology comprised of nested magnetic flux surfaces is used to avoid loss of high-energy particles from the ends of the device as would be experienced with a cylindrical geometry. In addition, a poloidal component of the magnetic field is generated by driving a toroidal current in the plasma to avoid loss of confinement due to particle drifts [12]. Despite these measures, tokamaks still suffer loss of confinement beyond that predicted by collision models in toroidal geometry (neoclassical transport); this so-called “anomalous transport” is believed to be due to turbulent diffusion [11, §4.17]. Therefore, understanding plasma turbulence in tokamaks will enable mitigation of such losses, thus improving device efficiency, which will further progress towards demonstration of net electricity production.

When the heating power applied to a tokamak is increased above a certain threshold, turbulence in the edge of the plasma is suddenly and spontaneously suppressed [13], and the plasma is said to have transitioned from low-confinement mode (L-mode) to high-confinement mode (H-mode) via the low-to-high transition (L-H transition). This results in a region of steep pressure gradient at the edge of the plasma, while the core gradient length-scale remains largely unaffected, as if the core profile has been put up on a pedestal; hence, the region of steep gradients is known as “the pedestal”. Since the required temperature is set by the fusion reaction cross-section and

fusion power scales as density squared [11, Ch.1], a high pedestal increases fusion power. Furthermore, since the core gradient length-scale remains largely unchanged but is at a higher absolute pressure, the core pressure gradient also increases, thus further increasing core pressure and fusion power output. Therefore, maximising pedestal height is crucial for tokamak performance, and predictive models of pedestal growth such as EPED [14] and Europed [15] will be key to achieving that goal. However, such models are based on local (single flux surface) treatment of the underlying plasma instabilities, so do not include global (many flux surface) effects, which may be significant [16, 17], thus reducing model accuracy. This thesis investigates global effects on plasma instabilities involved in the EPED and Europed models (kinetic ballooning modes, KBMs) and how such effects might impact the accuracy of EPED / Europed model predictions.

In addition, turbulence is regulated by radially-sheared plasma flows known as zonal flows [18], which are involved in the L-H transition [19–27]. Therefore, understanding zonal flows is critical for understanding the L-H transition, and thus ensuring entry to H-mode to achieve a high pedestal and maximise tokamak performance. Unfortunately, experimental measurement of zonal flows is difficult due to their symmetry and quasi-stationary nature. While it is easier to experimentally measure the transfer of energy into zonal flows [24], this approach is also difficult due to diagnostic limitations. Therefore, this thesis also presents a simulation-based method to calculate the expected transfer of energy into zonal flows in support of equivalent experimental measurements; this furthers our understanding of zonal flows and their relationship with turbulence, thus making progress towards efficient, high-performance tokamaks.

1.2 Motivation

Plasma micro-instabilities, the turbulence they generate, and their interaction with zonal flows are central to the achievable performance of a tokamak. For computational efficiency, these phenomena are often simulated using local (single flux surface) gyrokinetic models; however, their poloidal structure may be altered in position and intensity by global (many flux surface) effects, with implications for experimental measurements and the accuracy of theoretical models.

This thesis is split into two main parts: Part II studies global effects on micro-instabilities in the pedestal, while Part III investigates the poloidal structure of zonal flow drive. This chapter introduces the key concepts for each of these parts, and describes the document structure.

1.2.1 Global Effects

Turbulent transport is a key loss mechanism that limits tokamak performance. This occurs through anomalous transport (direct turbulent losses) and the interplay between turbulent transport and pedestal growth. Leading ELM prediction models such as EPED [14] and Europed [15], which are crucial to predicting tokamak performance, posit that pedestal growth is constrained by the onset of kinetic ballooning modes (KBMs, one particular class of micro-instability). However, such models do not include global effects, which may be significant [16, 17], thus reducing their accuracy. Furthermore, such models often use ideal MHD ballooning modes (IBMs) as a proxy for KBMs (for performance), thus neglecting kinetic effects such as diamagnetic drift stabilisation, which may further diminish their accuracy. To ensure the accuracy of such models, it is important to assess the impact of global effects on KBMs in the pedestal, and compare this to the IBM behaviour predicted by MHD.

Previous work [17] has used global gyrokinetic simulations to study KBMs subject to global effects in a pedestal-like equilibrium. This demonstrated partial agreement with MHD, but some discrepancies remain unresolved. Addressing these issues is difficult due to the high computational expense of global gyrokinetic simulations, boundary effects, and the magnetic field model in the particular code used (ORB5). Meanwhile, another line of research [28–30] has shown that global effects may be captured by an array of independent local gyrokinetic simulations for which the global coupling is calculated afterwards. This method, referred to here as the local-global method, has significantly reduced computational cost compared to global gyrokinetics, avoids problematic boundary effects, and, used with the local gyrokinetic code GS2 [1], enables the ORB5 magnetic field model to be tested for accuracy. However, the local-global method has so far only been demonstrated for core plasmas, but not in the pedestal. Therefore, it is important, for computational efficiency and investigating new physics, to test the local-global method in the pedestal; this thesis takes a step towards this by testing the local-global model in the pedestal-like equilibrium used in [17], which features a pseudo-pedestal at mid-radius. Demonstration of the local-global method in the pedestal will open up a new analysis method for studying global effects on micro-instabilities in the pedestal, which will enable, among other things, a deeper analysis of the correspondence between KBMs and IBMs, and examination of the consequences for EPED-like models. This may, in turn, improve the accuracy of such models so that we can predict tokamak performance more reliably.

Having not yet tested the local-global method in pedestal-like equilibria, there are a number of uncertainties around the validity of its application. Will the steep gradients of the pedestal invalidate the use of the local approximation? Will there be any other unexpected issues arising from the extreme conditions in the pedestal? And how can one determine the valid-

ity of the local-global method for a given case? Then, how accurately can the local-global method capture global effects in a pedestal-like equilibrium? What are the limits of validity of the local-global method in parameter space? And to what extent does the agreement between the local-global method and MHD meet expectations? And once these questions have been answered, what does this tell us about the behaviour of KBMs in the pedestal relative to IBMs? And what are the implications for EPED-like models?

This thesis answers these questions by first testing local gyrokinetic models of an experimental JET pedestal equilibrium using GS2 (Chapter 4). This highlights a potential issue with local pedestal simulations: that high magnetic shear, common in pedestals, *might* invalidate the use of the local approximation. Chapter 4 discusses the physics driving the magnetic shear problem, the impact on ideal and kinetic ballooning modes, and the implications for the accuracy of local models in regions of high magnetic shear. Chapter 5 completes Part II by applying the local-global method to a pedestal-like equilibrium (with low magnetic shear to avoid any potential issues), comparing to previously published global gyrokinetic and MHD simulations, and analysing the impact on EPED-like models.

1.2.2 Zonal Flows

Turbulence resulting from micro-instabilities is regulated by zonal flows in a “predator-prey” relationship (where zonal flows are the predator and turbulence is the prey) [18]. This governs the L-H transition (e.g. [19–27]) and H-mode transport (e.g. [20, 23, 31]). While these theories are well supported by simulations, experimental evidence is limited due to difficulty measuring zonal flows, as they are poloidally and toroidally symmetric. However, zonal flows are driven by nonlinear interaction of turbulent modes, which *can* be measured in experiments more easily using nonlinear energy transfer functions (NETFs) [22, 24, 32]. Unfortunately, the poloidal extent of

this analysis is limited by the diagnostic measurement area. Therefore, the poloidal structure of zonal flow drive would be invaluable in interpreting zonal flow experiments.

NETFs have been applied to nonlinear micro-turbulence simulations for other purposes (e.g. [33–35]), but the transfer between turbulence and zonal flows has not yet been measured in simulations. Nonlinear transfer into zonal flows could be calculated in simulations *at arbitrary poloidal angles*, thus revealing the poloidal structure of zonal flow drive.

In performing this calculation for the first time, there are a number of questions to address: Which turbulent modes are primarily responsible for driving zonal flows? How does the spectrum of transfer relate to the spectrum of turbulent activity? How do the turbulent and zonal energy levels respond to nonlinear energy transfer? And the key questions: What is the poloidal structure of zonal flow drive? And why does zonal flow drive exhibit the observed structure?

Chapter 6 answers these questions by applying NETFs to local nonlinear gyrokinetic simulations of ion temperature gradient (ITG) turbulence in a reference equilibrium representing a tokamak core plasma (the Cyclone base case, CBC [36]).

1.3 Document Structure

The structure of this thesis is as follows. Part I comprises this chapter and Chapter 2, which describes the local gyrokinetic model used throughout this thesis. Part II begins with Chapter 3, which describes the local-global method used throughout this part. Chapters 4 and 5 investigate global effects on KBMs in the pedestal, as described in Subsection 1.2.1. Part III contains Chapter 6, which studies the poloidal structure of zonal flow drive,

as introduced in Subsection [1.2.2](#). Finally, Part [IV](#) contains Chapter [7](#), which draws together overall conclusions from Parts [II](#) and [III](#). Supplementary information is contained the [Appendices](#).

Chapter 2

Background

2.1 Gyrokinetics

2.1.1 Introduction

Micro-instabilities are ion gyroradius sized plasma waves destabilised by the equilibrium gradients and magnetic geometry, and influenced by kinetic effects such as collisions, wave-particle resonances, trapped particle effects, and finite Larmor radius effects. Therefore, we need a model of plasma dynamics that includes these effects – a kinetic model – to study micro-instabilities.

The simplest kinetic model is to follow the path of each particle subject to the Lorentz equation:

$$m \frac{d\mathbf{v}}{dt} = Ze(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \quad (1)$$

where m is the particle mass, \mathbf{v} its velocity, Z its charge number, e the elementary charge, \mathbf{E} the electric field, and \mathbf{B} the magnetic field; the fields would be calculated from the particle trajectories using Maxwell's equations. Unfortu-

nately, the computational expense of such models is prohibitive for realistic cases due to the large number of particles, six-dimensional phase-space, and the variety of length- and time-scales. However, we can exploit various scale separations [37] to average over the gyromotion; thus, we restrict the problem to the micro-instability length- and time-scales, and eliminate a velocity dimension. This results in a model, known as gyrokinetics, that is solvable on today's supercomputers.

The following subsections summarise the assumptions and key steps required to derive a gyrokinetic model. More detailed derivations are available in [38, 39, and references therein].

2.1.2 Assumptions

First, we evolve a distribution function, f , rather than individual particle trajectories. We assume f varies smoothly in phase-space, which is reasonable for the thermal species of interest (fast particle effects, which are excluded here, require special treatment, e.g. [40]). This leads to the collisional Vlasov equation¹:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{Ze}{m} \{ \mathbf{E} + \mathbf{v} \times \mathbf{B} \} \cdot \nabla_{\mathbf{v}} f = C(f) \quad (2)$$

where t is time and $C(f)$ is the collision operator, which approximates the rate of change of f due to Coulomb collisions.

Next, we assume the plasma is magnetised, i.e. the magnetic field is sufficiently strong relative to the plasma temperature such that: (a) the ion thermal gyroradius, $\rho_i = v_{\text{th}}/\omega_{ci}$, is much smaller than the plasma size, L ; or, equivalently, (b) the transit frequency (the number of times a particle crosses the plasma per unit time), $\omega_t = v_{\text{th}}/L$, is much smaller than the

¹Also known as the Vlasov-Boltzmann or Vlasov-Fokker-Planck equation

ion gyrofrequency (cyclotron frequency), ω_{ci} ; and (c) the ion-ion collision frequency, ν_{ii} , is much less than the gyrofrequency such that gyromotion can occur without being overly impeded by collisions. For completeness, ν_{ii} is given by [1] $\nu_{ii} = \frac{\sqrt{2}\pi n_i Z_i^4 e^4 \ln(\Lambda)}{4\pi\epsilon_0 \sqrt{m_i} T_i^{3/2}}$ where n_i is the ion number density, Λ is the plasma parameter ($\ln(\Lambda)$ is the Coulomb logarithm) given by [41] $\Lambda = \frac{4\pi\epsilon_0^{3/2} T_i^{3/2}}{3e^3 n_i^{1/2}}$, ϵ_0 is the vacuum permittivity, and T_i is the ion temperature.

The assumption of a magnetised plasma defines a fast time-scale $\sim O(\omega_{ci})$, which is removed later by averaging over the gyromotion, and a slow time-scale $\sim O(\omega_t)$, which is retained; i.e. by averaging over a gyro-orbit, dynamics on the fast time-scale are averaged out, such that we no longer have the dynamics of a point-particle as in a fully-kinetic model, but rather the dynamics of charged rings about the particle's guiding centre, which vary on the slow time-scale. Similarly, there is a short length-scale $\sim O(\rho_i)$ that is retained and a long length-scale $\sim O(L)$ that is removed. Thus, we restrict the problem to the length- and time-scales of interest.

In addition, we assume that the drift velocities, v_D , are slow compared to the characteristic velocity of the gyromotion, which is the ion thermal velocity, $v_{th} = \sqrt{2T_i/m_i}$. This restricts the drifts to the slow time-scale, which is retained, so drifts remain in the model to exert their (de)stabilising effects. This assumption is reasonable for studying micro-instabilities that interact with the drifts on the slow time-scale, which are the target here. Situations with drifts on the fast time-scale would require a different model, such as MHD [37] or a fully-kinetic approach (e.g. [42–44]).

Finally, we assume scale separation between the parallel and perpendicular length-scales, L_{\parallel} and L_{\perp} respectively. This captures the fact that particles can move quickly along magnetic field lines, while being constrained in the perpendicular direction.

The above assumptions lead to the gyrokinetic orderings:

$$\frac{\rho_i}{L} \equiv \frac{\omega_t}{\omega_{ci}} \sim \frac{v_{ii}}{\omega_{ci}} \sim \frac{v_D}{v_{th}} \sim \frac{L_{\perp}}{L_{\parallel}} \ll 1 \quad (3)$$

Typical values for tokamak plasmas² are $L \sim \mathcal{O}(1 \text{ m})$ and $\rho_i \sim \mathcal{O}(0.01 \text{ m})$; recall that L is the plasma size, e.g. the tokamak minor radius, and ρ_i is the ion thermal gyroradius, and their ratio, ρ_i/L is similar to that of the perpendicular and parallel length-scales, L_{\perp}/L_{\parallel} (Eq. 3). Therefore, the length- and time-scale ordering conditions are often valid. However, there may be cases where this does not hold. This is discussed in Chapter 5. The collision frequency ordering is also often a valid assumption for tokamak plasmas³, since $\nu_{ii} \sim \mathcal{O}(10^4 \text{ Hz})$ while $\omega_{ci} \sim \mathcal{O}(10^8 \text{ Hz})$.

2.1.3 The δf Approach

Next, we split the fields and distribution function, denoted generically as χ , into equilibrium and fluctuating parts. We assume the equilibrium part, χ_0 , varies on the long length-scale, L , and on a time-scale slower than ω_t so that we may neglect changes to χ_0 , i.e. we assume:

$$\frac{\partial \chi_0}{\partial t} = 0 \quad (4)$$

Furthermore, we assume the equilibrium distribution function, f_0 , is Maxwellian – again, reasonable for thermal species. Meanwhile, we assume the fluctuations, χ_1 , vary on the short length-scale, ρ_i , and the slow time-scale, ω_t , so they are retained. Furthermore, we assume small fluctuations:

$$\frac{\chi_1}{\chi_0} \ll 1 \quad (5)$$

²Assuming $T \sim \mathcal{O}(10 \text{ keV})$, $m \sim \mathcal{O}(1 \text{ amu})$, $Z \sim \mathcal{O}(1)$, $B \sim \mathcal{O}(1 \text{ T})$

³Assuming typical values as above and $n_i \sim \mathcal{O}(10^{20} \text{ m}^{-3})$

This approach is known as δf (or delta- f) gyrokinetics. Full- f methods are available (e.g. [45–47]) for situations where these assumptions do not hold (such as modelling scrape-off layer filaments, or where transport relaxes the distribution function), but are not used in this thesis.

2.1.4 Coordinate Transform

Next, it is convenient to transform to velocity coordinates that isolate the gyrophase angle⁴, α , as this is the dimension that is removed by averaging over the gyromotion. For example, this could be done in terms of the parallel and perpendicular velocities, resulting in $(v_{\parallel}, v_{\perp}, \alpha)$. Alternatively, one could do this in terms of kinetic energy per unit mass, $\varepsilon = v^2/2$, and pitch angle, $\lambda = v_{\perp}^2/(Bv^2)$, resulting in $(\varepsilon, \lambda, \alpha, \sigma)$, where σ is a binary coordinate to indicate the direction of v_{\parallel} since this is lost by using v^2 . The latter is used in the gyrokinetic code GS2 [1] and Equation (6) below.

2.1.5 The Gyrokinetic Equation

Once in suitable coordinates, we further split f_1 (the perturbed distribution function) into a gyrophase independent part, g , and a gyrophase dependent part, \hat{f}_1 , which is removed by gyroaveraging. Carrying out the gyroaveraging (and the associated significant quantity of algebra), one eventually arrives at the gyrokinetic equation; an example (in wavenumber space) is as follows [38, eq. 3.65]:

$$\left[\frac{\partial}{\partial t} + (v_{\parallel} \mathbf{b} + \mathbf{v}_D) \cdot \nabla \right] g = - \left[\frac{\nabla f_0}{B} \cdot \mathbf{b} \times \nabla_{\perp} + \frac{Z_i e}{m_i} \frac{\partial f_0}{\partial \varepsilon} \frac{\partial}{\partial t} + \frac{\nabla g}{B} \cdot \mathbf{b} \times \nabla_{\perp} \right] \left[(\phi_1 - v_{\parallel} A_{1\parallel}) J_0(k\rho_i) + \frac{v_{\perp}}{k} B_{1\parallel} J_1(k\rho_i) \right] \quad (6)$$

⁴This is the angle indicating where around the gyro-orbit the particle is

where \mathbf{b} is the unit vector parallel to the magnetic field, ϕ is the electrostatic potential, A is the magnetic potential ($\mathbf{B} = \nabla \times \mathbf{A}$), J_0 and J_1 are the 0th and 1st order Bessel functions of the first kind (introduced as a consequence of gyroaveraging), k is the wavenumber, and subscripts \parallel and \perp indicate components parallel and perpendicular to the magnetic field respectively.

The left-hand side of Equation (6) indicates that this equation describes the time evolution of g , shown by the $\frac{\partial}{\partial t}$ term and the convective derivative term within the operator (square bracket). On the right-hand side, the first square bracket is the operator and the second square bracket contains the fields that are operated upon. The first two terms of the operator are the linear terms that capture the effect of the background density and temperature gradients respectively in driving changes in the system. The third term in the operator (the ∇g term) is the nonlinear term that captures how the state of g drives further changes in the system. The three terms in the fields bracket represent the $\mathbf{E} + \mathbf{v} \times \mathbf{B}$ part of the Lorentz equation (Equation (1)), with the gradient operators necessary to get from potentials (e.g. ϕ) to fields (e.g. \mathbf{E}) contained within the operator bracket.

2.2 Ballooning Theory

2.2.1 Introduction

The gyrokinetic model presented in Section 2.1 describes the fluctuating fields and distribution function that represent plasma micro-instabilities. These perturbations can be short-circuited by particles moving along magnetic field lines. Therefore, modes aligned with the magnetic field are the most unstable. As a result, the instability of the mode with integer poloidal and toroidal mode numbers m and n peaks on the flux surface with rational safety factor, q , such that:

$$q = m/n \tag{7}$$

For a given n , Equation (7) indicates that different m peak on different flux surfaces, assuming the realistic case of finite magnetic shear, \hat{s} , which is the normalised gradient of q defined as:

$$\hat{s} = \frac{r}{q} \frac{\partial q}{\partial r} \tag{8}$$

where r is the minor radius coordinate.

The various n are linearly independent due to toroidal axisymmetry. Therefore, we may Fourier decompose and ignore the toroidal direction, ζ ; the toroidal behaviour is thus encapsulated by n and described by $\exp[in\zeta]$. However, the poloidal dependence of the equilibrium couples the various m , causing interference and thus generating a global disturbance across many flux surfaces. This also prevents a simple Fourier decomposition of the form $\exp[im\theta]$ in the poloidal direction.

The gyrokinetic equation could be solved numerically across a domain covering many flux surfaces to obtain the global mode directly. This approach is known as global gyrokinetics (e.g. [48–50]). However, such methods are computationally expensive compared to local (single flux-surface) gyrokinetics. This is due to the need for global codes to use [49]: (a) a larger simulation domain (hence a larger computational grid); (b) finite difference methods rather than spectral methods (hence a larger computational grid); and (c) additional interpolation and buffer regions in the radial direction (hence additional computations and a larger computational grid respectively). For example, [49] states that “for a typical two-species, nonlocal and nonlinear trapped electron mode turbulence simulation for a medium-sized tokamak . . . several 100 kCPUh can be expected”. In comparison, the

single-species, local and nonlinear ion temperature gradient mode turbulence simulations reported in Chapter 6 have a run-time $\sim \mathcal{O}(1 \text{ kCPUh})$. While the cases compared above are different from each other, the local computation time is two orders of magnitude less than that of the global simulations, hence the run-time difference is expected to remain significant when comparing two cases that are equivalent to each other. Furthermore, the difference between local and global run-times is even more pronounced in the linear regime, since the use of spectral methods in local simulations means that each wavenumber is independent, so each can be run in parallel or only one need be run if that is all that is of interest; conversely, this decoupling of radial modes is not available in global calculations due to the need to use finite difference methods.

Given the computational expense of global gyrokinetics, we seek a more efficient approach via the so-called “local approximation”. Thankfully, in the limit of high n , there is an additional scale separation we can exploit between the system size, L , and radial width of the global mode, characterised by the rational surface spacing, Δ , given by⁵:

$$\Delta = \frac{1}{nq'} \quad (9)$$

where $q' = \frac{\partial q}{\partial r}$ on the surface with $q = m/n$.

Equation (9) shows that Δ becomes small in the limit of high n . As this happens, the destructive interference of poloidal harmonics occurs within a shorter distance, so the radial width of the global mode becomes small compared to L , and the relevant flux surfaces experience approximately the same equilibrium conditions. Therefore, we seek to exploit the scale separation:

⁵Since $q = m/n$ and $q \pm q'\Delta = (m \pm 1)/n$, assuming a first order Taylor expansion of q in the radial direction near the m/n surface

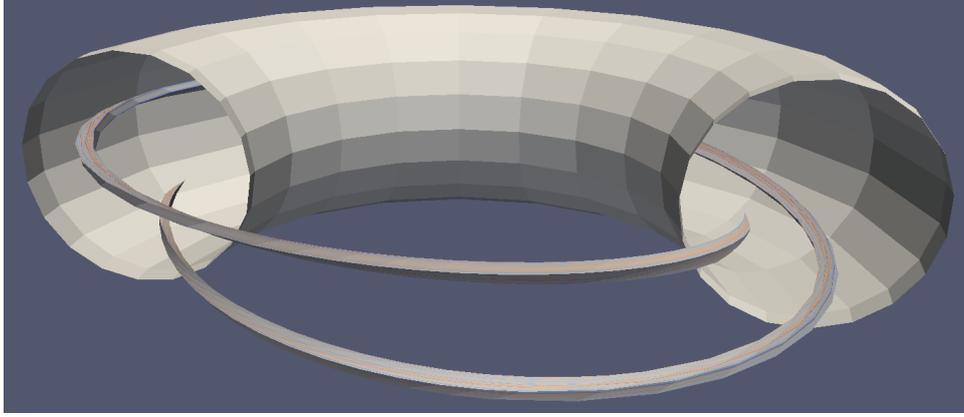


Figure 1: An example flux tube simulation domain. The cut-away shows the last closed flux surface with the flux tube inside.

$$\frac{\Delta}{L} \ll 1 \quad (10)$$

such that we need only consider the equilibrium conditions on a single flux surface. Thus, we obtain a local model, which is the lowest order of an expansion in the small parameter Δ , with the radial equilibrium variation on the length scale L parameterised at this order. This results in a simulation domain, known as a flux tube, that is localised around a single magnetic field line in the perpendicular directions while remaining extended in the parallel direction, as shown in Figure 1. A flux tube is a much smaller domain than that of global gyrokinetics, and thus much less computationally expensive.

The following subsections summarise the theoretical techniques used to derive a local model. More detailed derivations are available in [38, 39, and references therein]. The implications of the local approximation are discussed in Chapters 4 and 5.

2.2.2 The Ballooning Transform

Consider a linear global electrostatic potential fluctuation, $\phi(x, \theta, \zeta, t)$, where $x = r - r_0$ with r_0 a suitable reference minor radius such as that of a rational surface near the centre of the disturbance, θ is the poloidal angle and ζ is the toroidal angle.

First, we assume a separable time dependence:

$$\phi(x, \theta, \zeta, t) = \phi(x, \theta, \zeta) \exp[-i\Omega t] \quad (11)$$

where $\Omega = \omega + i\gamma$ is the global complex frequency with ω the real frequency and γ the growth rate. This assumption is valid since $\phi(x, \theta, \zeta)$ changes on the equilibrium time-scale, which is much slower than Ω .

Next, as noted in Subsection 2.2.1, the ζ dependence is ignorable via toroidal Fourier decomposition by assuming axisymmetry:

$$\phi(x, \theta, \zeta) = \phi(x, \theta) \exp[in\zeta] \quad (12)$$

Also as noted in Subsection 2.2.1, the poloidal modes are coupled. Therefore, we cannot perform a simple Fourier decomposition in the poloidal direction. Instead, using the scale-separation of Equation (10), we seek an eikonal representation of the form:

$$\phi(x, \theta, \zeta) = F(x, \theta) \exp[inS(x, \theta, \zeta)] \quad (13)$$

where the function $F(x, \theta)$ captures the slow variation⁶ on the length-scale L , while the exponential function captures the fast variation on the length-scale Δ , with:

⁶The notation F for slow and S for fast is used for consistency with the literature, e.g. [51]

$$S(x, \theta, \zeta) = \zeta + Y(x, \theta) \quad (14)$$

Unfortunately, this method of achieving an eikonal representation conflicts with the requirement for periodicity in θ for the realistic case of finite magnetic shear [52]. Therefore, we employ the ‘‘ballooning transform’’ [53] to go from the periodic θ domain to the infinite η domain:

$$\phi(x, \theta) = \sum_m \int_{-\infty}^{+\infty} \hat{\phi}(x, \eta) \exp[im(\theta - \eta)] d\eta \quad (15)$$

In doing so, we transform from $\phi(x, \theta)$ to $\hat{\phi}(x, \eta)$. The advantage of this approach is that $\phi(x, \theta)$ is automatically periodic in θ provided $\hat{\phi}(x, \eta) \rightarrow 0$ as $\eta \rightarrow \pm\infty$.

To illustrate the physical significance of the ballooning transform, we use the following identity, which is related to the Poisson summation formula [54]:

$$\sum_N \delta(x + N2\pi) = \frac{1}{2\pi} \sum_M \exp[iMx] \quad (16)$$

where δ is the Dirac delta function. This allows Equation (15) to be expressed as a Dirac comb:

$$\phi(x, \theta) = 2\pi \sum_N \int_{-\infty}^{+\infty} \hat{\phi}(x, \eta) \delta(\theta - \eta + N2\pi) d\eta \quad (17)$$

This form shows that the periodic function $\phi(x, \theta)$ may be constructed by summing samples of the the function $\hat{\phi}(x, \eta)$ on the infinite domain η with a sampling interval of 2π , provided $\hat{\phi}(x, \eta) \rightarrow 0$ as $\eta \rightarrow \pm\infty$, which is required for the integral to converge.

Since $\hat{\phi}(x, \eta)$ does not have to be periodic, it can be represented by an eikonal form known as the “ballooning representation” [51, 55]:

$$\hat{\phi}(x, \eta) = \hat{\phi}_0(x, \eta) \exp[-inq'x(\eta - \eta_0)] \quad (18)$$

$\hat{\phi}(x, \eta)$ is the global mode structure in ballooning space, which is related to the local mode structure, $\hat{\phi}_0(x, \eta)$, by the exponential factor⁷. $\hat{\phi}_0(x, \eta)$ varies slowly in x due to its dependence on the equilibrium (which also varies slowly in x). The exponential factor varies slowly in η when x is small and $\eta \sim \eta_0$ but rapidly in η when x and / or $\eta - \eta_0$ are large. Therefore, as x becomes large, the integration over η in Equation (15) will tend towards zero, thus capturing the finite radial extent of the global mode. Combining this with the assumption that nq' is large results in the small radial width of the global mode that was the original basis for seeking an eikonal representation. Furthermore, when x is small, the contribution to the periodic global mode structure comes primarily from the region where $\eta \sim \eta_0$, with η_0 an arbitrary phase that sets the origin of the local mode structure, and hence is known as the ballooning angle, ballooning phase angle or ballooning parameter.

2.2.3 The Local Approximation

As noted in [53], “the calculation of linear oscillations in any axisymmetric system can always be reduced to a two-dimensional eigenvalue problem”:

$$\mathcal{L}(x, \theta)\phi(x, \theta) = \Omega\phi(x, \theta) \quad (19)$$

The ballooning transform allows transformation of Equation (19) from the periodic θ domain to the infinite η domain:

⁷Indeed, the local and global mode structures in ballooning space are equal where $x = 0$ or $\eta = \eta_0$, since the exponential factor is exactly 1.

$$\hat{\mathcal{L}}(x, \eta)\hat{\phi}(x, \eta) = \Omega\hat{\phi}(x, \eta) \quad (20)$$

where Equations (19) and (20) have the same eigenvalue, which is the global complex frequency, Ω .

In transforming Equation (19) to Equation (20), the differential operators in \mathcal{L} are mapped to corresponding operators in $\hat{\mathcal{L}}$ by evaluating their effect upon the ballooning transform (Equation (15)). For example, $\frac{\partial}{\partial x}$ in \mathcal{L} is mapped by evaluating the partial derivative with respect to x of Equation (15) (which is just the partial derivative with respect to x of Equation (18)), to obtain:

$$\frac{\partial}{\partial x} \rightarrow \left(\frac{1}{\hat{\phi}_0} \frac{\partial \hat{\phi}_0}{\partial x} - inq'(\eta - \eta_0) \right) \quad (21)$$

This procedure can then be repeated for the other operators in \mathcal{L} to obtain $\hat{\mathcal{L}}$ and thus the transformed problem.

The first term on the right-hand side of Equation (21) is the radial gradient inverse length scale of the local mode structure, $1/L_{\hat{\phi}_0}$. Multiplying Equation (21) by $\Delta = \frac{1}{nq'}$ shows that this term can be ignored provided $\Delta/L_{\hat{\phi}_0} \ll 1$, which requires that $\hat{\phi}_0$ varies on the equilibrium length scale, which it does. Similar observations can be applied to the other operator mappings to eliminate terms that are small relative to nq' (provided the relevant quantities vary on the long length-scale – this is studied further in Chapter 4). In doing so, the operator to lowest order in Δ , known as the local operator, $\hat{\mathcal{L}}_0$, is derived. We then have the local eigenvalue equation:

$$\hat{\mathcal{L}}_0(\eta)\hat{\phi}_0(\eta) = \Omega_0\hat{\phi}_0(\eta) \quad (22)$$

where the eigenvalue, Ω_0 , is the local complex frequency, which only depends parametrically on x . Furthermore, we now have a one dimensional eigenvalue problem, which is thus easier to solve than the 2D problem of Equation (20). This approach necessarily excludes global effects, which are: (a) poloidal mode coupling, since we only have a single flux surface, hence a single q value, and thus a single m for a given n ; and (b) effects due to radial profile variations, since we have assumed that inverse gradient length scales are small. As such, this model tells us the local behaviour due to the local gradients of the radial profiles.

2.3 GS2

The simulations reported in this thesis use the local δf gyrokinetics code GS2 [1], which solves the gyrokinetic equation in a ballooning space flux tube⁸, coupled with Maxwell's equations to obtain the perturbed gyrophase independent part of the distribution function, g , and the perturbed electromagnetic fields ϕ_1 , $A_{\parallel 1}$ and $B_{\parallel 1}$. GS2 has the following features [56] that are relevant as modelling assumptions for the work presented in this thesis:

- Pseudo-spectral methods in the perpendicular spatial directions for performance and accuracy. Other dimensions use finite difference methods.
- Periodic boundary conditions in the perpendicular directions (valid for box-sizes larger than a few turbulence decorrelation lengths), and a zero incoming particles boundary condition in the parallel direction.
- A range of initial conditions for g .
- Linear or nonlinear mode.

⁸The η and η_0 nomenclature of Section 2.2 is consistent with the literature, e.g. [51]. These quantities are called `theta` and `theta0` in GS2.

- A choice of initial value solver or eigensolver.
- An implicit time-stepping algorithm for linear terms and an explicit algorithm for the nonlinear term.
- Collisions modelled as pitch angle scattering and energy diffusion.
- Adjustable numerical dissipation via spatial and temporal decentering and a hyper-diffusion term.
- Arbitrary species definitions with gyrokinetic or adiabatic treatment for each.
- A variety of analytical magnetic equilibrium models. Various numerical equilibrium formats are also supported.
- User-specified grid resolutions and extents.
- MPI parallel processing with excellent scaling up to thousands of cores.
- Arbitrary equilibrium flow shear (set to zero for all work in this thesis).

Part II

Global Effects on Micro-Instabilities

Chapter 3

The Local-Global Method

3.1 Introduction

The local gyrokinetic model presented in Chapter 2 neglects the global effects of poloidal mode coupling and radial profile variation, which appear at the next order of the expansion in Δ . However, as noted in Chapter 1, global effects may be important [16, 17], but direct global gyrokinetic simulations face issues including computational expense and boundary effects. Conveniently, the local-global method retrieves the global solution from local simulations [28–30]; this chapter describes the method.

To lowest order in Δ (the local model), the operator $\hat{\mathcal{L}}_0(\eta)$ and, hence, the eigenvalue Ω_0 and mode structure $\hat{\phi}_0(\eta)$, have a weak dependence on the equilibrium via x and η_0 , which are free parameters at this order. This is sometimes indicated with semi-colon notation⁹ as follows:

$$\hat{\mathcal{L}}_0(\eta; x, \eta_0)\hat{\phi}_0(\eta; x, \eta_0) = \Omega_0(x, \eta_0)\hat{\phi}_0(\eta; x, \eta_0) \quad (23)$$

⁹Read as, for example, $\hat{\phi}_0$ as a function of η , parameterised by x and η_0

The global eigenvalue, Ω , and mode structure, $\phi(x, \theta)$, are only determined by considering the next order of the expansion in Δ (the global model). Thankfully, as we shall see, the solution at the higher order depends only on the results from the local model, and does not depend on the equilibrium directly. Therefore, we may obtain the global solution from an array of independent local results, with the global coupling calculated separately afterwards.

3.2 Obtaining Global Behaviour from Local Simulations

To show the connection between $\Omega_0(x, \eta_0)$ and Ω , we first combine the ballooning transform (Equation (15)) and the ballooning representation (Equation (18)), and represent Equation (18) by its Fourier transform, which leads to the Fourier-ballooning representation [55]:

$$\phi(x, \theta) = \int_{-\infty}^{+\infty} \hat{\phi}_0(\theta; x, \eta_0) \exp[-in((q_0 + q'x)\theta - q'x\eta_0)] A(\eta_0) d\eta_0 \quad (24)$$

A more complete derivation of Equation (24) is given in Appendix A.

Equation (24) shows that the global mode structure, $\phi(x, \theta)$, is an integration over ballooning phase angles (η_0) of the product of four constituent parts, *viz.* the generic poloidal fluctuation, $\exp[-in(q_0 + q'x)\theta]$, and the generic radial fluctuation, $\exp[inq'x\eta_0]$, weighted by, respectively, the local mode structure along the field line, $\hat{\phi}_0(\theta; x, \eta_0)$, and the so-called “amplitude envelope”, $A(\eta_0)$, that governs the contribution to the global mode from each local mode (i.e. each (x, η_0)).

Calculating the global eigenvalue, Ω , and the global mode structure, $\phi(x, \theta)$, via this method requires determination of the amplitude envelope, $A(\eta_0)$.

To calculate the amplitude envelope, $A(\eta_0)$, we start from the Fourier-ballooning representation, Equation (24), and reintroduce the explicit time dependence:

$$\phi(x, \theta, t) = \int_{-\infty}^{+\infty} \hat{\phi}_0(\theta, t; x, \eta_0) \exp[inq'x\eta_0] \exp[-in(q_0 + q'x)\theta] A(\eta_0) d\eta_0 \quad (25)$$

Then, noting the form of the time dependence in Equation (11), we take the time derivative of Equation (25) to obtain:

$$\int_{-\infty}^{+\infty} \{\Omega - \Omega_0(x, \eta_0)\} \hat{\phi}_0(\theta, t; x, \eta_0) \exp[inq'x\eta_0] \exp[-in(q_0 + q'x)\theta] A(\eta_0) d\eta_0 = 0 \quad (26)$$

Next, we Taylor expand $\Omega_0(x, \eta_0)$ to second order about $x = 0$, which is valid since we have assumed the radial extent of the global mode is small compared to the equilibrium length scale:

$$\Omega_0(x, \eta_0) = \Omega_0(\eta_0)|_{x=0} + \Omega'_0(\eta_0)x + 1/2\Omega''_0(\eta_0)x^2 \quad (27)$$

where the primes indicate radial derivatives at $x = 0$.

Following this, we substitute Equation (27) into Equation (26) to obtain:

$$\int_{-\infty}^{+\infty} \{\Omega - [\Omega_0(\eta_0)|_{x=0} + \Omega'_0(\eta_0)x + 1/2\Omega''_0(\eta_0)x^2]\} \hat{\phi}_0(\theta, t; x, \eta_0) \exp[-in((q_0 + q'x)\theta - q'x\eta_0)] A(\eta_0) d\eta_0 = 0 \quad (28)$$

Since $nq'\eta_0 = k_{x0}$ is the Fourier conjugate variable of x , we can use the standard properties of Fourier transforms to write:

$$x^\ell \rightarrow \left(\frac{i}{nq'}\right)^\ell \frac{\partial^\ell}{\partial \eta_0^\ell} \quad (29)$$

for any integer ℓ , so that Equation (28) becomes:

$$\int_{-\infty}^{+\infty} \left\{ \Omega A(\eta_0) - \left[\Omega_0(\eta_0)|_{x=0} + \frac{i\Omega'_0(\eta_0)}{nq'} \frac{\partial}{\partial \eta_0} - \frac{\Omega''_0(\eta_0)}{2(nq')^2} \frac{\partial^2}{\partial \eta_0^2} \right] A(\eta_0) \right\} \hat{\phi}_0(\theta, t; x, \eta_0) \exp[-in((q_0 + q'x)\theta - q'x\eta_0)] d\eta_0 = 0 \quad (30)$$

For the integral in Equation (30) to be zero in a non-trivial case (i.e. with non-zero $\hat{\phi}_0$), the terms in curly braces must be zero. Therefore, we have the second order differential eigenequation:

$$\left[\Omega_0(\eta_0)|_{x=0} + \frac{i\Omega'_0(\eta_0)}{nq'} \frac{\partial}{\partial \eta_0} - \frac{\Omega''_0(\eta_0)}{2(nq')^2} \frac{\partial^2}{\partial \eta_0^2} \right] A(\eta_0) = \Omega A(\eta_0) \quad (31)$$

that relates the local eigenvalues, $\Omega_0(x, \eta_0)$, to the global eigenvalue, Ω .

Notice the terms in the differential operator of Equation (31) only depend on the the local eigenvalues, $\Omega_0(x, \eta_0)$, and derivatives thereof. Hence, we may collect $\Omega_0(x, \eta_0)$ from an array of independent GS2 simulations, use those results to compute the coefficients for the differential operator, and solve Equation (31) to obtain global behaviour from local simulations.

In addition, there are two practical points. Firstly, to determine the global mode structure, $\phi(x, \theta)$, we evaluate the Fourier-ballooning representation (Equation (24)). This requires $\hat{\phi}_0(\theta; x, \eta_0)$, which may be calculated from $\hat{\phi}_0(\eta; x, \eta_0)$ using Poisson summation, e.g. via a Dirac comb similar to Equation (17), i.e.

$$\hat{\phi}_0(\theta; x, \eta_0) = 2\pi \sum_N \int_{-\infty}^{+\infty} \hat{\phi}_0(\eta; x, \eta_0) \delta(\theta - \eta + N2\pi) d\eta \quad (32)$$

Secondly, for $\phi(x, \theta)$ to be periodic in θ , we must have $A(\eta_0)$ periodic in η_0 , hence we solve Equation (31) numerically with periodic boundary conditions. This means that: (a) Equation (31) is solved using the SciPy sparse

linear algebra function `eigs`, which returns the eigenvalues and eigenvectors of a square matrix; (b) the square matrix passed to the `eigs` function is constructed to represent Equation (31) as a set of simultaneous equations where the differential operators $\frac{\partial}{\partial \eta_0}$ and $\frac{\partial^2}{\partial \eta_0^2}$ acting on $A(\eta_0)$ are computed as discrete central differences, $\left(\frac{\partial A}{\partial \eta_0}\right)_i = \frac{A_{i+1} - A_{i-1}}{2\Delta\eta_0}$ and $\left(\frac{\partial^2 A}{\partial \eta_0^2}\right)_i = \frac{A_{i+1} + A_{i-1} - 2A_i}{\Delta\eta_0^2}$; and (c) the central differences at the limits of the periodic domain $\eta_0 = \pm\pi$ are computed using the elements of $A(\eta_0)$ from the other end of the domain to capture the periodicity of $A(\eta_0)$. As shall be shown in Chapter 5, $A(\eta_0)$ has a Gaussian-like shape that decays to zero at $\eta_0 = \pm\pi$. Therefore, for simplicity, periodic boundary conditions (point (c) above) were achieved in this thesis by setting $A(\eta_0 = \pm\pi) = 0$; this was tested against strictly enforced periodic boundary conditions and was found to give the exact same eigenvalue and eigenfunction.

Chapter 4

Kinetic Ballooning Mode

Narrowing in Ballooning Angle

4.1 Introduction

As discussed in Chapter 1, to ensure the accuracy of EPED-like models, it is important to assess the impact of global effects on KBMs in the pedestal. However, to do so using global gyrokinetic simulations can present issues including computational expense and boundary effects. Thankfully, we can avoid these issues by using the local-global method described in Chapter 3 to retrieve global behaviour from local gyrokinetic simulations.

The local-global method has so far only been demonstrated in core plasmas (for ITGs [28] and KBMs [29]), but not in the pedestal. Therefore, to demonstrate its validity for studying the pedestal, the local-global method needs to be applied to a pedestal (or pedestal-like) equilibrium.

The first step in using the local-global method to study the pedestal is to run local gyrokinetic simulations of such a case. This is done routinely (e.g [16, 57–59]), but often without scanning the ballooning angle, η_0 , since

$\eta_0 = 0$ is *usually* the most unstable. To neglect η_0 is problematic because we cannot distinguish between: (a) stability at all η_0 vs. (b) stability at $\eta_0 = 0$ but instability at $\eta_0 \neq 0$; or (c) instability across a wide range of η_0 about $\eta_0 = 0$ vs. (d) instability only within a very narrow range of η_0 about $\eta_0 = 0$.

Concerning case (a) vs. case (b), it has been known, even since the early days of ballooning theory [60], that an equilibrium is unstable if there is instability at *any* value of η_0 . Therefore, determining stability only at $\eta_0 = 0$ would lead to an incorrect conclusion for equilibria subject to case (b).

With regard to case (c) vs. case (d), since k_{x0} (the radial wavenumber at $\eta = 0$) is linearly proportional to η_0 , instability within a narrow range of η_0 about $\eta_0 = 0$ suggests a global mode that is radially wide. This is potentially in violation of the local approximation, thus calling into question the use of local gyrokinetics for equilibria subject to case (d). At the very least, such equilibria are likely to have a large global correction, such that the resulting global transport is much less than that predicted by the local model.

This chapter begins with local gyrokinetic simulations of a JET pedestal equilibrium, including a scan of η_0 , that exhibit KBMs unstable over a very narrow range of η_0 around $\eta_0 = 0$, i.e. case (d). It appears something similar has been observed in local gyrokinetic simulations before [61], but was dismissed as non-physical. Furthermore, this mode was absent from the corresponding direct global gyrokinetic simulations and was thus excluded from their analysis.

The existence of apparently non-physical modes in local pedestal gyrokinetic simulations presents a problem for applying the local-global method (or, indeed, local gyrokinetics) in the pedestal: Given the local approximation may be violated, is the local approximation (and, hence, the local-global method) a valid tool for such study?

After presenting the JET pedestal results, this chapter addresses this issue by first determining the cause of the observed narrowing of KBMs in η_0 . This is done using local linear gyrokinetic scans of magnetic geometry parameters that differ between the core and pedestal cases. This shows that the narrowing is due to high magnetic shear, \hat{s} , in the JET equilibrium, which is common in pedestal equilibria. The remainder of the chapter then explores how \hat{s} leads to KBM narrowing in η_0 , and discusses the implications of this on the validity of using the local approximation and the local-global method for studying KBMs in the pedestal.

4.2 Methods

To connect to previous work, the JET pedestal equilibrium is based on shot #84795. The equilibrium is described in more detail in [62, and references therein], but the relevant points for this work are:

- Density and temperature profiles were reconstructed by averaging high resolution Thomson scattering (HRTS) measurements over multiple ELM cycles.
- Data was binned by percentage of the time between ELMs. The equilibrium used here is that corresponding to the 80-99% time bin.
- Modified tanh functions were fitted to the profiles.
- The equilibrium was constructed by solving the Grad-Shafranov equation with the tanh profiles as input, the output of which was provided by the authors of [62] in HELENA [63] format. The profiles of safety factor and pressure taken from the equilibrium data file, and the magnetic shear and pressure gradient profiles calculated from them, are shown in Figure 2. Notice the dip in magnetic shear in Figure 2 (c) near the flux-surface chosen for analysis. This is due to the bootstrap

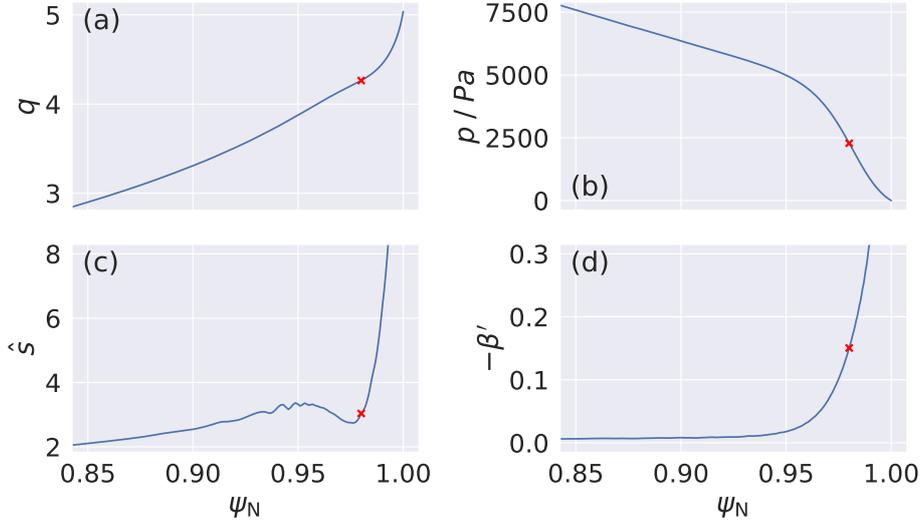


Figure 2: Profiles from the JET equilibrium showing (a) safety factor, (b) pressure, (c) magnetic shear, and (d) normalised pressure gradient inverse length scale (as defined in Equation (33)). The red x markers indicate the $\psi_N = 0.98$ surface used in the analysis.

current, which is proportional to the pressure gradient, hence peaks in the pedestal. This was calculated self-consistently during the HELENA reconstruction [62] using the Koh-Chang pedestal bootstrap current model [64], hence the calculated bootstrap current values contain some uncertainty that comes from the experimental pressure measurements (as noted in [62]).

- The flux surface at normalised poloidal flux $\psi_N = 0.98$ was selected for analysis, consistent with [62], as it was near the peak pressure gradient.

This equilibrium was previously shown [62] to be marginally stable to peeling-ballooning (PB) modes (as expected immediately before an ELM), and the $\psi_N = 0.98$ surface was shown to be in the second stability region for IBMs, and stable to KBMs. However, IBMs and KBMs were found to

be unstable at values of magnetic shear higher than the equilibrium value ($\hat{s} \gtrsim 5.5$ *cf.* the equilibrium value $\hat{s} \sim 3.5$), and instability peaked at normalised pressure gradient inverse length scale, β' , slightly smaller than the equilibrium value, though within experimental uncertainty ($\beta' \sim -0.13$ *cf.* the equilibrium value $\beta' \sim -0.15 \pm 0.03$), where β' is defined as:

$$\beta' = \beta \frac{1}{p} \frac{\partial p}{\partial \psi_N} \quad (33)$$

Therefore, the JET simulation was set to $\hat{s} = 7.92$ and $\beta' = -0.1294$ to promote KBM instability; the equilibrium values were not used as this would put the simulation in the second stability region so there would be no KBMs. While it may seem inconsistent to set \hat{s} so much higher than the equilibrium value, this is justified for three reasons. First, as shown in Figure 2 (c), the gradient of \hat{s} is very steep in the pedestal, thus \hat{s} can be significantly different on nearby flux surfaces. Indeed, the value used ($\hat{s} = 7.92$) occurs on the flux surface at $\psi_N = 0.992$ *cf.* the surface studied at $\psi_N = 0.98$. (The reason to not just use the $\psi_N = 0.992$ surface is that we do not want to use higher \hat{s} *per se*; we want to promote KBM instability to study KBMs, and this was achieved by increasing \hat{s} . The $\psi_N = 0.992$ surface was not tested, and may or may not be unstable to KBMs.) Secondly, notice that the $\psi_N = 0.98$ surface is near the dip in the \hat{s} profile caused by the bootstrap current. This introduces uncertainty in \hat{s} due to the uncertainty of the bootstrap current model, combined with the uncertainty of the Grad-Shafranov solver. Indeed, some numerical error is even apparent in Figure 2 (c) resulting from the calculation of \hat{s} from q . Finally, this is a qualitative study of KBM behaviour, not a quantitative comparison with the experiment. While this does mean an arbitrary equilibrium could have been generated from scratch, this setup based on an experimental equilibrium was used because some unusual KBM behaviour was observed previously in this equilibrium (see Subsection 4.3.1).

Parameter	CBC	JET
κ	1.0	1.56
κ'	0.0	2.21
δ	0.0	0.23
δ'	0.0	0.63
r/a	0.5	0.98
Δ	0.0	-0.28
R_0/a	3.333	3.121
R_a/a	3.333	3.116
q	1.4	4.26
\hat{s}	0.8	7.92

Parameter	CBC	JET
a/L_T	2.54	17.5
a/L_n	0.81	39.6
β	0.02	0.00116
β'	0.0	-0.1294
n	39	24
ρ_*	0.0066	0.0023
ν_{ii}	0.162	0.022
ν_{ee}	0.591	0.996
Z_{eff}	1.0	1.22

Table 1: Comparison of parameters for the CBC and JET cases, where κ is elongation, δ is triangularity, primes indicate radial derivatives, r/a is the normalised minor radius of the flux surface simulated, $\Delta = \frac{r}{a} \frac{dR}{dr}$ is the Shafranov shift, R_0/a and R_a/a are the normalised major radii of the magnetic axis and the last closed flux surface, a/L_T and a/L_n are the normalised temperature and density gradient inverse length scales, ρ_* is the normalised ion gyroradius, and ν_{ii} and ν_{ee} are the normalised ion-ion and electron-electron collision frequencies.

For comparison with the JET pedestal equilibrium, we also run the core equilibrium model reported previously [28, 29] that uses Cyclone Base Case (CBC) parameters [36] with concentric circular flux surfaces using the s - α equilibrium model [65, §III]. The JET and CBC parameters are compared in Table 1, and the flux surface shapes of the two cases are compared in Figure 3.

Table 1 and Figure 3 show that the CBC and JET cases differ in many parameters, most obviously the plasma shaping / magnetic geometry parameters. To facilitate easier investigation of these differences, a Miller equilibrium

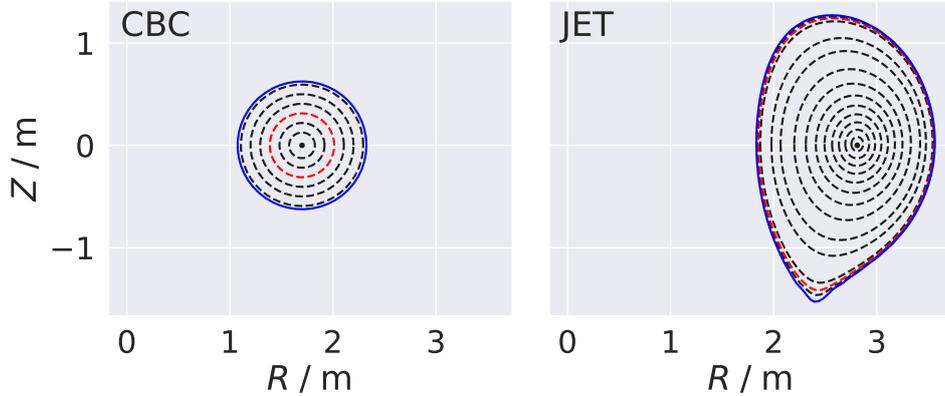


Figure 3: Comparison of CBC and JET flux surface shapes. The blue solid line is the last closed flux surface, and the red dashed line is the flux surface simulated.

model [65, §IV] representation of each case was made. The JET Miller model was tested for numerical convergence and was found to be suitably converged. The CBC setup is based on that previously reported in [29], hence relies on the convergence tests reported there. The JET convergence tests and a check of the CBC Miller model against the previously reported s - α model are shown in Appendix B. Methods for measuring the frequency and growth rate from GS2 data, and for tracking sub-dominant modes using the GS2 eigensolver are described in Appendix C.

4.3 Results

4.3.1 JET Pedestal KBMs Narrow in Ballooning Angle

The ballooning angle, η_0 , was scanned in GS2 for both equilibria. Figure 4 compares the growth rates, frequencies and mode structures from the two cases. This shows that in the CBC setup KBMs were unstable over a wide range $|\eta_0| \lesssim 0.3\pi$, while in the JET case KBMs became sub-dominant within $|\eta_0| \lesssim 0.03\pi$. Thus, Figure 4 shows that KBMs in the JET case are significantly

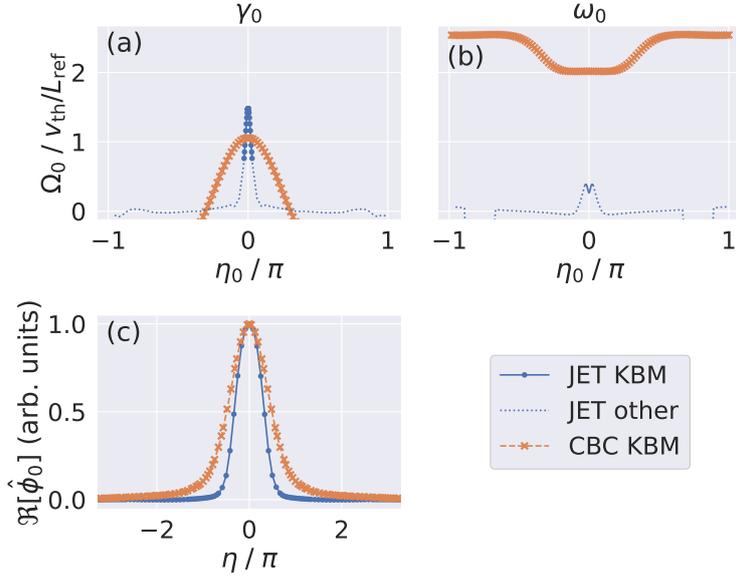


Figure 4: (a) Growth rate, (b) frequency, and (c) mode structure from the JET and CBC simulations. Markers are shown to indicate the resolution used, but are suppressed on the JET line in (b) for clarity.

narrower in η_0 than those in the CBC setup. Given η_0 is related to k_{x0} (the radial wavenumber at $\eta = 0$) by:

$$k_{x0} = -nq'\eta_0 = -k_y\hat{s}\eta_0 \quad (34)$$

this suggests a global mode in the JET case that is radially wide, thus potentially violating the local approximation. However, notice that the magnetic shear, \hat{s} , enters here. As shall be shown, the narrowing is caused by \hat{s} , hence it turns out that narrow in η_0 does not necessarily mean narrow in k_{x0} .

To ensure the observed modes were indeed KBMs, the JET case was rerun with a/L_T and a/L_n swapped; KBMs are pressure gradient driven, so should be largely unaffected by this change. The results are shown in Figure 5, where we see the growth rates before and after the change are

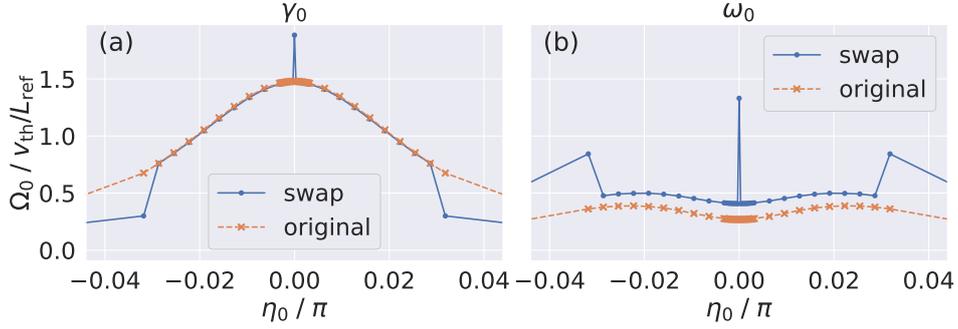


Figure 5: (a) Growth rate, and (b) frequency from the JET simulations before and after swapping a/L_T and a/L_n . Note the narrower range of η_0 shown, hence the wider appearance.

the same for $|\eta_0| \lesssim 0.03\pi$ and the frequencies in that range follow the same trend with a small upshift $\sim 0.14v_{th}/L_{ref}$. These observations are consistent with KBMs in this range, hence the distinction between JET KBMs and other modes in Figure 4. Interestingly, we can see in Figure 5 that with a/L_T and a/L_n swapped, there was another mode at $\eta_0 = 0$ that became sub-dominant within $|\eta_0| \lesssim 0.0003\pi$, which was the smallest non-zero value tested, i.e. the mode has *only* been observed at $\eta_0 = 0$. This mode is unlikely to be a KBM, since it is affected by swapping a/L_T and a/L_n , hence was not investigated further. Nonetheless, this observation shows that there may be effects that only exist at $\eta_0 = 0$, thus highlighting the importance of considering $\eta_0 \neq 0$.

The observation of KBMs narrow in η_0 has not yet been repeated in an alternative experimental pedestal equilibrium (apart from the similar observations reported in [61]). This was attempted in MAST pedestal simulations, but was abandoned due to unidentified numerical issues. Despite the limitation of not directly reproducing this observation in another experimental pedestal equilibrium, the results presented later in this chapter demonstrate that this is a robust effect due to high magnetic shear, so is expected in pedestals in which magnetic shear is high (which is often the

case, e.g. in conventional inductive current drive scenarios [66], though not always the case, e.g. in advanced non-inductive current drive scenarios [66] such as those developed for ITER steady state operations [67] that have low, zero, or negative magnetic shear due to the high bootstrap current fraction) and indeed in *any* equilibria with high magnetic shear (i.e. not necessarily limited to pedestals).

4.4 Discussion of Ballooning Angle Narrowing Mechanism

4.4.1 Magnetic Geometry Effects on KBM Ballooning Angle Width

To investigate the cause of KBM narrowing in η_0 , the parameters that differ between the two cases (Table 1) were varied and η_0 was scanned for each value tested. The width in η_0 was then measured for the resulting $\gamma_0(\eta_0)$ as the full width at half maximum. The scans started from the CBC setup because some parameters had a strong stabilising effect that caused KBMs to be completely stabilised in the JET case. The parameters that affected the width are shown in Figure 6; all other parameters had negligible effect on the width.

Figure 6 shows that the KBM width in η_0 is highly sensitive to \hat{s} , with large changes over a very narrow parameter range compared to the difference between the JET and CBC \hat{s} values. Meanwhile, the other parameters require nearly the full range of the CBC-JET difference to produce a similar change in width. This is a clear indication that \hat{s} is the dominant parameter controlling the KBM width in η_0 .

The identification of \hat{s} as the parameter responsible for KBM narrowing in η_0 was confirmed in two ways. Both confirmation methods varied the magnetic geometry parameters starting from modified setups. The first was

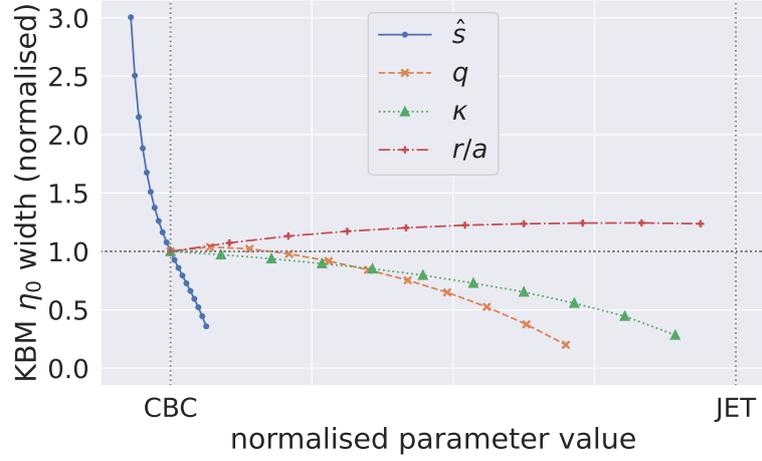


Figure 6: KBM ballooning space width variation with magnetic geometry parameters starting from the CBC setup. The width is normalised to the value obtained for the baseline CBC setup. Each parameter is normalised to the CBC and JET values of that parameter.

a modified CBC setup with higher instability drive (but the same ratio), $a/L_T = 2.54 \rightarrow 7.62$ and $a/L_n = 0.81 \rightarrow 2.43$. This was done so that KBMs were unstable at all $|\eta_0| \leq \pi$ for all parameter values tested to avoid the width measurement being affected by marginal stability. The second modified setup was based on the JET case, but with the CBC values of $\beta = 0.2$ and $\beta' = 0.0$ to avoid KBMs being stabilised completely as the parameter values were changed, which would render the width immeasurable (the JET values of β and β' produced stabilising effects on KBMs). In both cases, the width was measured as before, and the results are shown in Figure 7.

Figure 7 (a) shows that the previously observed narrowing due to q is now a widening, suggesting that the previous width measurements were affected by marginal stability. The \hat{s} measurements, however, show clear narrowing as before. Figure 7 (b) similarly shows that the previously observed narrowing due to κ is negligible compared to that caused by \hat{s} , again suggesting

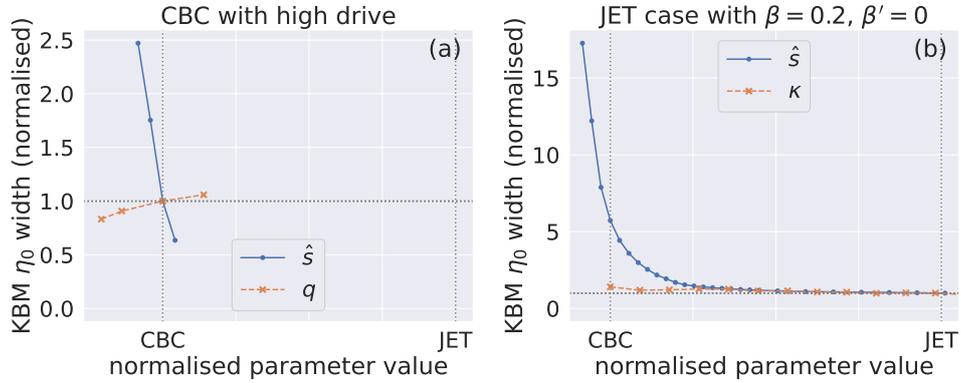


Figure 7: KBM η_0 width variation with magnetic geometry parameters starting from (a) a modified CBC setup, and (b) a modified JET case. The width is normalised to the value obtained for the baseline setup (i.e. (a) CBC and (b) JET). Each parameter is normalised to the CBC and JET values of that parameter.

that the previous κ measurements were affected by marginal stability, while the \hat{s} narrowing is a robust effect that is evident in all three cases (Figure 7 (a) and (b), and Figure 6).

4.4.2 The Suydam and Mercier Criteria of Ideal Ballooning Modes

Having identified \hat{s} as the parameter responsible for the narrowing of KBMs in η_0 observed in the JET case, a mechanism was sought for how \hat{s} causes this narrowing. As an initial source of explanation, we return to two fundamental theories of IBM stability, the Suydam and Mercier criteria. These describe necessary (though not sufficient) criteria for IBM stability in cylindrical and toroidal geometry respectively. While somewhat removed from the KBMs under investigation, these theories provide some insight into the parameters that drive and stabilise IBMs and, thus, are expected to also be involved in KBM instability.

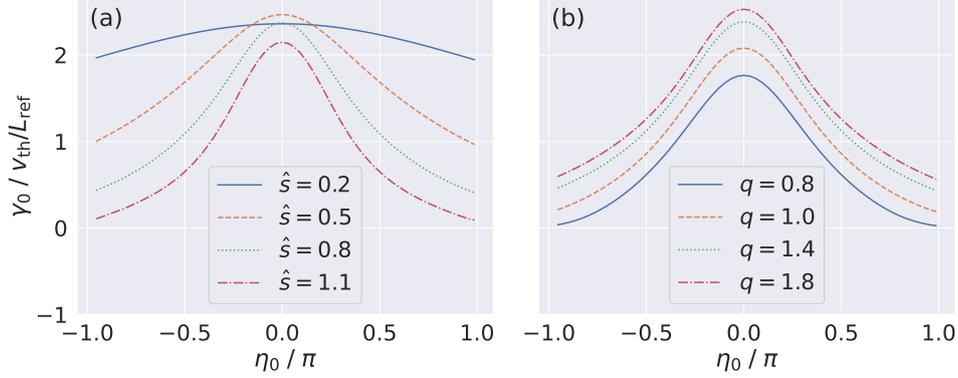


Figure 8: The effect of (a) \hat{s} and (b) q at various ballooning angles.

The Suydam and Mercier criteria [37] are very similar apart from the inclusion of flux surface averages in the Mercier criterion that complicate the expression. It is sufficient, and thus clearer, for our purposes to just discuss the simpler Suydam criterion, which is:

$$q^2 \beta' R^2 \kappa_r < \frac{\hat{s}^2}{4} \quad (\text{for stability}) \quad (35)$$

where κ_r is a coefficient to capture the effect of magnetic field line curvature in cylindrical geometry, and η_0 enters via \hat{s} as per Equation (34).

Equation (35) indicates that q , β' and curvature are destabilising for IBMs, while \hat{s} is stabilising. Checking the effect of these quantities in the CBC setup using GS2 shows that KBMs are subject to the same (de)stabilising effects from these parameters. More importantly, however, examining the effect of these parameters at different η_0 shows that the stabilising effect of \hat{s} affects high η_0 more than low η_0 , while the other parameters affect all η_0 approximately equally. This is shown in Figure 8 for \hat{s} compared to q as an example; similar results were obtained for β' (not shown), while changing the curvature via the R_0/a and R_a/a parameters from the JET to CBC values made negligible difference to both growth rate and width in η_0 .

Figure 8 shows that the destabilising effect of q (and β' , not shown) is insensitive¹⁰ to η_0 . Meanwhile, \hat{s} stabilises high η_0 faster than low η_0 , which leads to the observed narrowing in η_0 .

4.4.3 Mechanisms of Ideal Ballooning Instability

Further to the conditions suggested by the Suydam and Mercier criteria, IBMs are expected [60] in regions with: (1) unfavourable curvature; (2) minimal field line bending; and (3) near zero integrated local magnetic shear. The meaning of these quantities is as follows:

1. Unfavourable curvature refers to the region on the outboard side where toroidal curvature enhances ballooning instabilities since it acts in the same direction as the pressure gradient.
2. Stretched magnetic field lines provide a stabilising effect, so ballooning instabilities are most unstable where bending is minimised.
3. Ballooning instabilities may be considered as an interchange of field lines between neighbouring flux surfaces, hence are most unstable where these are aligned, i.e. where there is near zero integrated local magnetic shear.

These effects are captured in the local IBM eigenvalue equation [60]:

$$\left[-K(\eta) + \frac{\partial}{\partial \eta} \alpha(\eta) \frac{\partial}{\partial \eta} \right] \hat{\phi}_0(\eta) = \omega^2 \hat{\phi}_0(\eta) \quad (36)$$

where the coefficient $\alpha(\eta)$ captures the stabilising effect of magnetic field line bending:

¹⁰The results in Figure 8 also suggest that γ_0 is not periodic in η_0 . This is consistent with γ_0 depending on k_{x0} , and with η_0 being in the infinite η domain.

$$\alpha = \frac{4\pi^2 f(\psi)^2}{R^2 |\nabla\psi|^2} + \frac{4\pi^2 f(\psi)^4 |\nabla\psi|^2}{R^2 B^2} I^2 \quad (37)$$

where $B = |\mathbf{B}|$ with:

$$\mathbf{B} = f(\psi)\nabla\zeta \times \nabla\psi + Rg(\psi)\nabla\zeta \quad (38)$$

where $f(\psi)$ is the poloidal magnetic flux function and $g(\psi)$ the toroidal magnetic flux function with ψ the flux surface label (minor radius coordinate), and I is the integrated local shear:

$$\begin{aligned} I &= \frac{1}{f(\psi)} \int \frac{s}{\mathbf{B} \cdot \nabla\eta} d\eta \\ &= q'(\eta - \eta_0) + q \frac{\nabla\psi \cdot \nabla\eta}{|\nabla\psi|^2} \end{aligned} \quad (39)$$

where s is the local shear¹¹:

$$s = -\frac{\mathbf{B} \times \nabla\psi}{|\nabla\psi|^2} \cdot \nabla \times \frac{\mathbf{B} \times \nabla\psi}{|\nabla\psi|^2} \quad (40)$$

Meanwhile, the coefficient $K(\eta)$ captures the destabilising effect of magnetic field line curvature:

$$K(\eta) = -p'(\kappa_n + \kappa_g) \quad (41)$$

where p is pressure, prime indicates a radial derivative, and κ_n and κ_g are the normal and geodesic components of the curvature, given by:

$$\kappa_n = \frac{v^2 R^2}{R_0^2 B^2} \left[\frac{\partial}{\partial\psi} (2p + B^2) + \frac{\nabla\psi \cdot \nabla\eta}{|\nabla\psi|^2} \frac{\partial}{\partial\eta} (2p + B^2) \right] \quad (42)$$

¹¹The global shear, \hat{s} , is the flux surface average of s

and:

$$\kappa_g = -\frac{v^2 g(\psi)^2}{qB^4} I \frac{\partial}{\partial \eta} (2p + B^2) \quad (43)$$

where v is a geometrical factor defined as:

$$v = \frac{R_0}{2\pi} \int_P \frac{d\tau}{R^2} \quad (44)$$

where the integral is over the plasma volume, and $d\tau$ is the volume element¹².

Equation (39) shows that $I \sim q'(\eta - \eta_0)$, where the ballooning angle, η_0 , is a constant of integration that sets where $I = 0$. Furthermore, Equations (37) and (43) show that $\alpha \sim I^2$ and $\kappa_g \sim I$, hence η_0 also sets where $\kappa_g = 0$ and where α is minimised. The normal component of curvature, κ_n , which is independent of η_0 , primarily sets the region of unfavourable curvature, since the normal components of ∇p and ∇B are aligned on the outboard side.

Equations (41) to (43) suggest negative K (unfavourable curvature) on the outboard side, i.e. around $\eta = 0$. Therefore, as $|\eta_0|$ increases away from zero, the location where $I \sim 0$ and α is minimised starts to move out of the region of unfavourable curvature. This is the cause of η_0 stabilisation in general. Meanwhile, the q' dependence of I and α shows that higher q' narrows the region in η where $I \sim 0$ and α is minimised, which is the cause of \hat{s} stabilisation in general. The combination of narrowing in η due to q' and shifting in η due to η_0 leads to an increased sensitivity to η_0 stabilisation at high \hat{s} , which is why \hat{s} stabilises high η_0 faster than low η_0 as observed in Subsection 4.4.2.

¹² v is also related to the Jacobian of the coordinate system, $J = (\nabla\psi \times \nabla\eta \cdot \nabla\zeta)^{-1} = \frac{vR^2}{2\pi R_0}$

These arguments are illustrated in Figure 9 for the CBC setup¹³ by plotting $K(\eta)$, $\alpha(\eta)$ and $I(\eta)$ at $\eta_0 = 0$ and $\eta_0 \neq 0$ for $\hat{s} = 0.8$ and $\hat{s} = 7.92$:

- Figure 9 (a) shows that regions of negative K , minimal α and $I \sim 0$ are aligned since $\eta_0 = 0$. There is a wide region of instability in η since \hat{s} is low, and there is maximally destabilising curvature across the whole region due to the alignment.
- Figure 9 (b) shows how $\eta_0 \neq 0$ shifts the regions of minimal α and $I \sim 0$ out of the region of negative K . There is still a reasonably wide region of instability in η , but curvature is only marginally destabilising across some of it, hence the net instability will be somewhat reduced compared to Figure 9 (a).
- Figure 9 (c) exhibits narrowing in η of the region of minimal α and $I \sim 0$. There is still maximally destabilising curvature across the whole region, but the region is very narrow.
- Finally, Figure 9 (d) shows how the combination of narrowing in η due to q' and shifting in η due to η_0 leads to a very narrow region of minimal α and $I \sim 0$, with only marginally destabilising curvature across the whole region, hence the net result will be marginal stability, compared to only reduced instability in Figure 9 (b).

4.4.4 Ballooning Angle Narrowing of Ideal Ballooning Modes

The ideal ballooning arguments in Subsection 4.4.3 suggest that IBMs should be subject to the same narrowing observed for KBMs in Subsection 4.3.1. This hypothesis was tested using the GS2 auxiliary program

¹³The effects of plasma shaping and geodesic curvature are neglected for simplicity. Similar results are obtained when they are included.

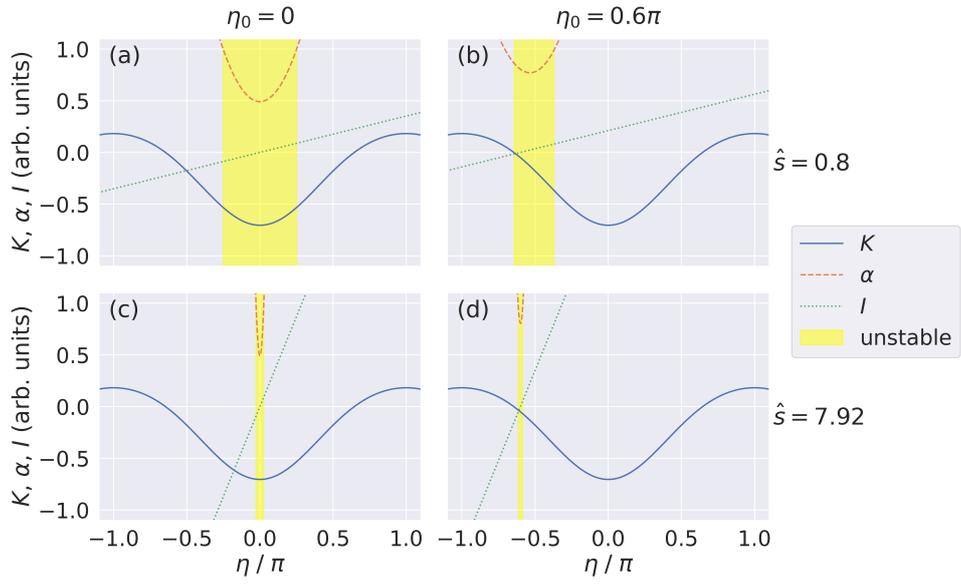


Figure 9: Illustrations of η_0 shifting and \hat{s} narrowing in η of the ideal MHD eigenvalue coefficients K , α and I for the CBC setup at (a) low \hat{s} and $\eta_0 = 0$, (b) low \hat{s} and $\eta_0 \neq 0$, (c) high \hat{s} and $\eta_0 = 0$ and (d) high \hat{s} and $\eta_0 \neq 0$. The yellow region indicates where $\alpha < 1$, $|I| < 0.1$ and $K < 0$ (indicative of unstable region only).

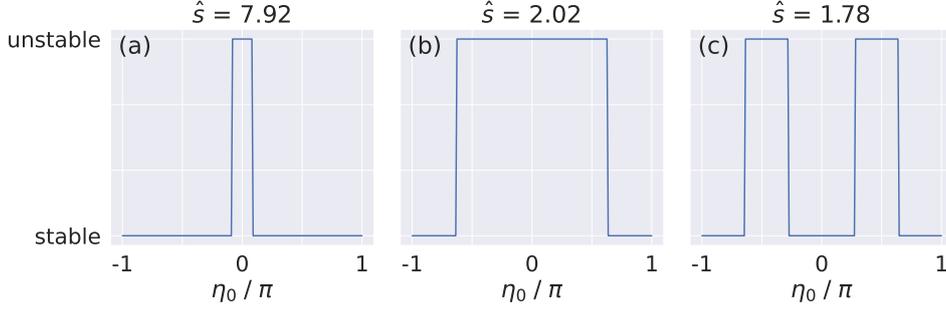


Figure 10: IBM stability as functions of η_0 with $\beta' = -0.13$ showing (a) IBMs narrow in η_0 at high \hat{s} , (b) IBMs wide in η_0 at low \hat{s} , and (c) IBMs stable at $\eta_0 = 0$ but unstable at some other values of η_0 at slightly lower \hat{s}

`ideal_ball`, which solves a local IBM eigenvalue equation equivalent¹⁴ to Equation (36). The `ideal_ball` solver sets the boundary condition $\hat{\phi}_0 = 0$ at one end of the η domain, calculates coefficients equivalent to K , α and $\frac{\partial}{\partial \eta}$, then calculates $\hat{\phi}_0(\eta)$ using first-order finite difference integration. The $\hat{\phi}_0(\eta)$ obtained is unstable if it crosses zero. Thus, `ideal_ball` returns a binary stable / unstable result, but not the eigenvalue.

`ideal_ball` calculates its result at $\eta_0 = 0$. To measure IBM width in η_0 , `ideal_ball` was modified to calculate at arbitrary η_0 . The modified code was run on the JET case with $\beta' \sim -0.13$ (as before) for various \hat{s} over the range $|\eta_0| \leq \pi$. The results, shown in Figure 10, exhibit narrowing of IBMs in η_0 with increased \hat{s} as expected.

Interestingly, at low \hat{s} , the region of η_0 unstable to IBMs splits such that $\eta_0 = 0$ is stable but there is instability at some non-zero values of η_0 (Figure 10 (c)). This is also expected [60], due to the competition of pressure gradient stabilisation through α and destabilisation through K that eventually leads to stabilisation at $\eta_0 = 0$ while instability remains at other values

¹⁴The equation solved by `ideal_ball` differs from Equation (36) due to the normalisations used in GS2

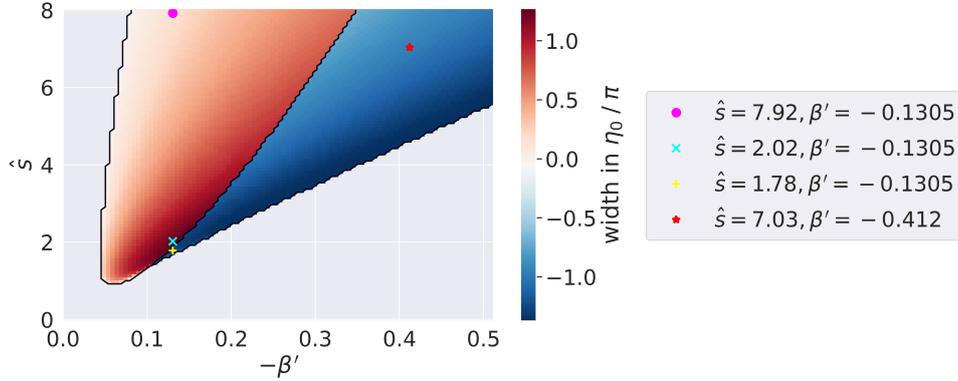


Figure 11: Modified s - α diagram as a result of IBM splitting observed in Figure 10 (c). The red region (positive width) shows the width of the IBM stability region for (\hat{s}, β') unstable at $\eta_0 = 0$. The blue region (negative width) corresponds to (\hat{s}, β') stable at $\eta_0 = 0$ but unstable at some other values of η_0 , where the width covers both unstable regions in η_0 and the stable region in between. The cases from Figures 10 and 15 are indicated.

of η_0 . As the pressure gradient is increased further, the unstable region due to I and α moves further around in η_0 until it moves outside the region of unfavourable curvature, leading to second stability. Therefore, as noted in [60]: “Balloonning instabilities are predicted for those configurations that yield a negative eigenvalue, $\omega^2 < 0$, for some value of η_0 . That is, η_0 must also be varied to obtain a complete test.” Indeed, using the modified `ideal_ball` to also scan β' reveals a region of instability that extends into what would have been the second stability region had η_0 not been scanned, as shown in the s - α diagram in Figure 11.

The results in Figures 10 and 11 show that:

1. IBMs are subject to the narrowing mechanism described in Subsection 4.4.3 as expected, which suggests this mechanism may also be responsible for the narrowing of KBMs observed in Subsection 4.3.1.

2. It is important that $\eta_0 \neq 0$ is tested! In this case, this has resulted in a modified second stability region. While this may not be a concern for cases that are primarily constrained by the first stability boundary, this could explain observations such as those in [62], where the inter-ELM pedestal growth appeared to follow no particular trend and second stability access was found to be important. In such cases, if the second stability boundary is affected by instability at $\eta_0 \neq 0$, then this could be a reason why the IBM proxy appeared to show that pedestal growth was not always KBM limited.
3. The initial onset of IBMs at the first stability boundary is narrow in η_0 for all \hat{s} . This suggests global IBMs that are radially wide, and thus the local approximation may be inappropriate, or at least require a large correction. On the other hand, the width increases rapidly with β' , especially at low \hat{s} , which suggests any invalidity or inaccuracy is quickly mitigated. In any case, this further highlights the importance of considering global effects on IBMs and KBMs.

4.4.5 Ballooning Angle and Magnetic Shear in Gyrokinetics

Having established how the basic mechanisms of IBM stability suggest narrowing in η_0 at high \hat{s} (Subsection 4.4.3) and confirmed that this does indeed cause narrowing of IBMs (Subsection 4.4.4), it is now important to examine how these mechanisms affect KBMs. Since the ideal ballooning coefficients K , α and I are all derived solely from the magnetic geometry, it is expected that these quantities will also affect the gyrokinetic equation. This can be seen in the gyrokinetic equation, where: \hat{s} (or rather q') and η_0 enter via the ballooning representation of ϕ (Equation 18); ϕ is operated upon in the gyrokinetic equation (Equation 6) by $\nabla_{\perp} \sim k_{\perp}$, which is derived from the magnetic geometry, i.e. also depends upon \hat{s} and η_0 (this relates to the I ideal ballooning coefficient); and the time rate of change of the distribution

function, $\frac{\partial g}{\partial t}$, also depends on the convective derivative, which includes a contribution from the magnetic drifts, $\mathbf{v}_D \cdot \nabla$, which also depends on the magnetic geometry (∇p , ∇B^2 , etc., i.e. relates to the K ideal ballooning coefficient). This was confirmed by inspecting the GS2 source code to see how the parameters \hat{s} and η_0 enter the gyrokinetic equation, details of which are given in the next two paragraphs. This revealed that these parameters both enter via the curvature and ∇B drifts, and via the perpendicular wavenumber, k_\perp , both of which are calculated from the same magnetic geometry variables that are used in `ideal_ball`. As shall be shown, the drifts effectively perform the role of K , while k_\perp produces an effect similar to that of I .

In GS2, the curvature and ∇B drifts are represented by the operators `cvdrift`, `cvdrift0`, `gbdrift` and `gbdrift0`, which relate to the terms of Equation (41) as follows:

- `cvdrift` represents the ∇p terms of κ_n and the η_0 independent part of κ_g
- `cvdrift0` represents the ∇p term of the η_0 dependent part of κ_g
- `gbdrift` represents the ∇B^2 terms of κ_n and the η_0 independent part of κ_g
- `gbdrift0` represents the ∇B^2 term of the η_0 dependent part of κ_g

These are then combined into `wdrift_func` (drift frequency function), which effectively performs the role of K , and then enters the “source” variable, which represents the convective derivative term ($\mathbf{v} \cdot \nabla$) and right-hand side of the gyrokinetic equation (Equation (6)) and is used to update the distribution function at each timestep. This is how curvature enters GS2 and, coupled with the kinetic gradient drives, leads to the condition of unfavourable curvature required to destabilise ballooning modes.

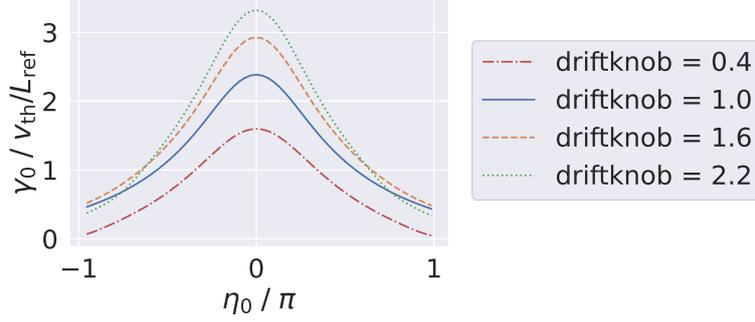


Figure 12: The effect of drift scaling at various ballooning angles

The effect of the drifts in GS2 was tested by varying the `driftknob` parameter, which applies a multiplication factor to the drift operators (`cvdrift` etc.). The results, presented in Figure 12, show that increasing the effect of the drifts increases instability near $\eta_0 = 0$, while leaving the growth rates largely unchanged at higher η_0 . This is consistent with a deepening of the negative well of $K(\eta)$, such that peak instability near $\eta_0 = 0$ increases, while the marginally destabilising curvature ($K \sim 0$) at higher η_0 does not change significantly. This demonstrates that \hat{s} does not cause narrowing of KBMs via the drifts.

Meanwhile, the perpendicular wavenumber is defined as:

$$k_{\perp}^2 = k_y^2 + k_x^2 \quad (45)$$

where k_y is set¹⁵ by the n chosen in Table 1 and k_x is given by:

$$k_x = nq'(\eta - \eta_0) = k_y \hat{s}(\eta - \eta_0) \quad (46)$$

¹⁵The calculation of k_y is non-trivial as it depends on the magnetic geometry. This is dealt with automatically by GS2. As an example, for concentric circular flux surfaces without Shafranov shift we have $k_y = \frac{nq}{r} \cos \left[\arctan \left(\frac{r}{a} \frac{\epsilon}{q} \right) \right]$ where $\epsilon = a/R_0$ is the inverse aspect ratio.

The $q'(\eta - \eta_0)$ factor that also appears in Equation 39 further illustrates the connection between k_{\perp} and I .

The perpendicular wavenumber then enters the gyrokinetic equation in two ways: (1) via the collision operator, both directly and via the Bessel functions¹⁶; and (2) via the Bessel functions in the “source” variable.

Collisions in GS2 are a kinetic effect, since they spread the distribution function in velocity space. However, Subsections 4.4.2 to 4.4.4 suggest that narrowing is caused by ideal ballooning effects. Therefore, it is expected that \hat{s} causes narrowing of KBMs via the Bessel functions in the “source” variable, but not by collisions.

The hypothesis that collisions do not cause narrowing of KBMs in η_0 was tested by repeating the $\hat{s} = 0.8$ simulations of Figure 8 (a) with collisions turned off. A comparison of the results with and without collisions is shown in Figure 13; similar results were obtained for other values of \hat{s} . This shows that collisions produce an effect similar to that of the drifts observed in Figure 12, whereby collisions cause increased peak instability while leaving higher η_0 largely unaffected. In comparison, the narrowing due to \hat{s} observed in Figure 8 (a) leaves the peak instability at $\eta_0 \sim 0$ largely unaffected, while stabilising higher η_0 more quickly. This demonstrates that the use of \hat{s} via the collision operator is not responsible for the \hat{s} dependence observed in Figure 8 (a).

Finally, the hypothesis that \hat{s} causes narrowing of KBMs via the Bessel functions in the “source” variable was tested. This is difficult in GS2 since \hat{s} and η_0 also affect κ_g but the `driftknob` parameters affect both κ_n and κ_g . Similarly, the `bess_fac` (Bessel factor) parameter applies a multiplication factor to the argument of the Bessel functions that includes k_{\perp} , but this

¹⁶Recall that Bessel functions feature in the gyrokinetic equation (Equation (6)) as a result of gyroaveraging

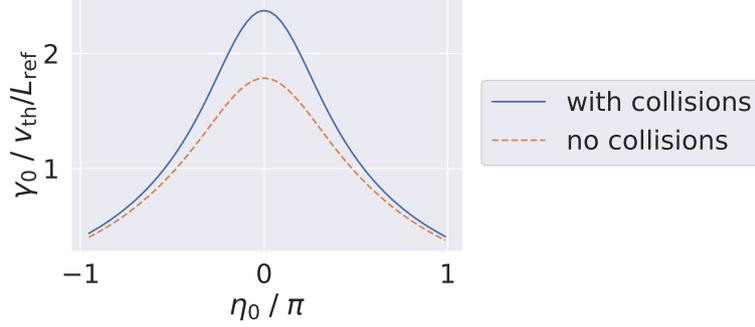


Figure 13: The effect of collisions at various ballooning angles

affects both the \hat{s} dependent (k_x) and \hat{s} independent (k_y) part. Therefore, to isolate the effect of \hat{s} via the Bessel functions in the “source” variable, it was necessary to switch back to the s - α equilibrium model, where GS2 has the flexibility to turn off the κ_g terms of the drift operators, thus eliminating the effect of \hat{s} via that channel.

In the default version of the s - α equilibrium, the drift operators are defined as follows:

$$\text{cvdrift} = 2\epsilon \{ \cos(\eta) + [\hat{s}\eta - \Delta \sin(\eta)] \sin(\eta) \} \quad (47)$$

$$\text{cvdrift0} = -2\epsilon\hat{s} \sin(\eta) \quad (48)$$

$$\text{gbdrift} = \text{cvdrift} \quad (49)$$

$$\text{gbdrift0} = \text{cvdrift0} \quad (50)$$

where $\epsilon = \frac{a}{R_0}$ is the inverse aspect ratio.

The \hat{s} and η_0 scans of Figure 8 (a) were re-run using this version of the s - α equilibrium model with collisions turned off as a baseline. The results, shown in Figure 14 (a), exhibit narrowing of KBMs in η_0 as expected.

The `normal_only` version of the s - α equilibrium redefines the drift operators as follows:

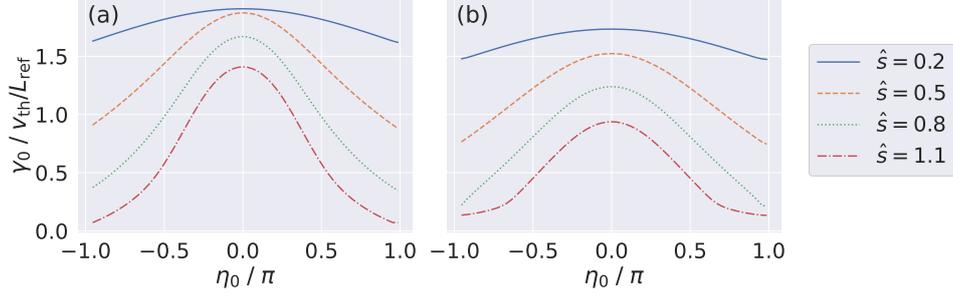


Figure 14: The effect of \hat{s} at various η_0 (a) including and (b) excluding the effect of \hat{s} via the drifts. Both data sets are collisionless.

$$\text{cvdrift} = 2\epsilon \cos(\eta) \quad (51)$$

$$\text{cvdrift0} = 0 \quad (52)$$

$$\text{gbdrift} = \text{cvdrift} \quad (53)$$

$$\text{gbdrift0} = \text{cvdrift0} \quad (54)$$

Thus, the κ_g terms have been removed from the drift operators and the effect of \hat{s} removed from the drifts. Now, the *only* way that \hat{s} enters the gyrokinetic equation is via the Bessel functions in the “source” variable (provided collisions are turned off, which they were).

Re-running the \hat{s} and η_0 scans using this version of the s - α equilibrium model yields the results shown in Figure 14 (b). These results exhibit narrowing of KBMs in η_0 , thus demonstrating that \hat{s} does indeed, as expected, exert its narrowing effect via k_\perp in the argument to the Bessel functions, which affects the distribution function via the “source” variable. While the Bessel functions themselves do not appear in the IBM calculation, the perpendicular wavenumber, k_\perp , that forms part of the argument to the Bessel functions is derived from the same magnetic geometry parameters that do go into the IBM code. Therefore, KBM narrowing in η_0 is primarily an ideal

ballooning effect, rather than a kinetic effect, as one might expect given the IBM physics discussed in Subsections 4.4.3 and 4.4.4.

Notice, however, that the narrowing in Figure 14 (b) is less pronounced than in Figure 14 (a), which suggests that the narrowing effect is enhanced by the destabilising effect of \hat{s} at $\eta_0 \sim 0$ via the drifts and collisions. Since collisions are a kinetic effect, this suggests that KBM narrowing may be modified as compared to IBM narrowing. Nonetheless, these results show that the narrowing of KBMs in η_0 is primarily driven by an increased rate of k_x stabilisation in η_0 . As shall be discussed in Subsection 4.5.1, this is consistent with the α and I stabilisation of IBMs discussed in Subsection 4.4.3, and with the assertion that KBM narrowing in η_0 is primarily an ideal ballooning effect, rather than a kinetic effect.

4.4.6 Additional Kinetic Effects

Subsections 4.4.1 to 4.4.5 show that \hat{s} causes narrowing of KBMs and IBMs, and that KBMs are primarily narrowed by the same ideal ballooning physics that narrows IBMs. However, upon checking the split IBM cases (those similar to Figure 10 (c), i.e. in the blue region of Figure 11) to see if KBMs also exhibit splitting, a divergence of IBM and KBM behaviour was observed, whereby KBMs were unstable at $\eta_0 = 0$ even though IBMs were stable there.

To avoid the possibility of this being caused by proximity to the original or modified second stability boundaries, a case was selected near the middle of the modified second stability region, with $\hat{s} = 7.0$ and $\beta' = -0.41$ (as indicated in Figure 11). GS2 simulations were run to measure KBM growth rates. The GS2 simulations were checked to confirm that the instabilities observed are indeed KBMs (by swapping a/L_T and a/L_n as before), which confirmed that KBMs are active over the full range of η_0 .

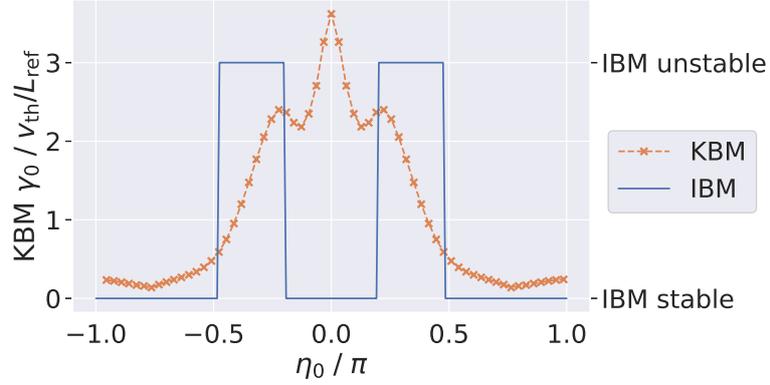


Figure 15: A comparison of KBM growth rates against IBM stability for the JET case with $\hat{s} = 7.0$ and $\beta' = -0.41$.

A comparison of the KBM growth rates from GS2 against the IBM stability from `ideal_ball` is presented in Figure 15. This shows that KBMs exhibit a large peak $\gamma_0 \sim 3.6 v_{th}/L_{ref}$ at $\eta_0 = 0$ where IBMs are stable, with smaller secondary peaks ($\gamma_0 \sim 2.4 v_{th}/L_{ref}$) aligned with the region where IBMs are unstable ($\pi/4 \lesssim |\eta_0| \lesssim \pi/2$). This suggests that the secondary peaks are driven by ideal ballooning physics, while the peak at $\eta_0 = 0$ is driven by additional kinetic effects.

The origin of the additional kinetic effects was not investigated due to time constraints. Nonetheless, this shows that additional kinetic effects can be important for KBMs in the modified second stability region. This further highlights that using IBMs as a proxy for KBMs does not always give the complete picture, especially if only $\eta_0 = 0$ is considered.

4.5 Discussion of Implications for Local Validity

4.5.1 The Relationship Between Ballooning Angle and Wavenumber

As discussed in Subsection 4.4.3, ballooning modes are destabilised where there is alignment of magnetic field lines on neighbouring flux surfaces. Furthermore, as noted in Subsection 2.2.1, modes aligned with the magnetic field are the most unstable. The combination of these ideas suggests that the stabilisation of ballooning modes with η_0 due to the shifting of α and I in η is equivalent to stabilisation with increasing k_{x0} (Equation (34)), which shifts in η the region where $k_x = 0$ (Equation (46)). Furthermore, in the ideal ballooning description, increased \hat{s} causes a narrowing of the region where α is minimised and $I \sim 0$, which leads to an increased sensitivity to η_0 stabilisation. In terms of wavenumber, this translates to increased \hat{s} causing a narrowing of the region where $k_x \sim 0$ – indeed, Subsection 4.4.5 showed that \hat{s} exerts its narrowing effect on KBMs via k_\perp . This raises the question of whether \hat{s} also causes narrowing in k_{x0} , which may be important for determining the real-space radial width of the corresponding global mode and, thus, whether the local approximation is valid for a given case.

To address this question, the results in Figure 8 (a) (shown again in Figure 16 (a) for ease of reference) were re-plotted against k_{x0} instead of η_0 , where k_{x0} was calculated using Equation (34) and the value $k_y = 0.1678 / \rho_i$ was obtained from the GS2 output.

The plot against k_{x0} is presented in Figure 16 (b). This shows that:

1. There is no narrowing in k_{x0} .
2. We observe an apparent narrowing in η_0 because higher \hat{s} maps the fixed range of $|\eta_0| \leq \pi$ to a larger range of k_{x0} , such that the fixed rate of stabilisation in k_{x0} happens faster in η_0 .

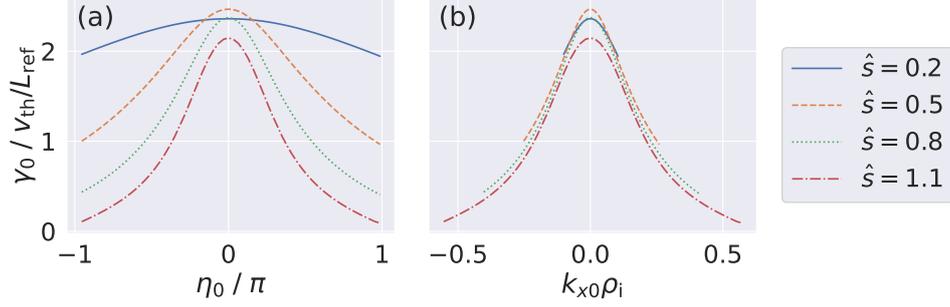


Figure 16: The effect of \hat{s} (a) at various ballooning angles (this is the same data as Figure 8 (a) and is shown again here for ease of reference); and (b) re-plotted against k_{x0} .

Therefore, while modes that are narrow in η_0 may at first *appear* to be in violation of the local approximation, the fact that they remain the same width in k_{x0} as \hat{s} increases suggests that they may in fact be the same width in real-space. This raises the question of whether η_0 or k_{x0} is the important parameter to determine whether γ_0 varies sufficiently slowly such that the local approximation is valid. This question is addressed in Subsection 4.5.2.

4.5.2 Conditions for Local Approximation Validity

To answer the question of whether η_0 or k_{x0} is the important parameter for local approximation validity, we return to the method of deriving a local model, as described in Subsection 2.2.3. The transform of the radial derivative operator, $\frac{\partial}{\partial x}$, has already been evaluated in Subsection 2.2.3 to reveal the mapping:

$$\frac{\partial}{\partial x} \rightarrow -in q'(\eta - \eta_0) \quad (55)$$

provided:

$$\frac{1}{\hat{\phi}_0} \frac{\partial \hat{\phi}_0}{\partial x} \ll nq' \quad (56)$$

Applying a similar procedure for the radial scaling operator (i.e. multiplication by x), results in the mapping:

$$x \rightarrow -\frac{i}{nq'} \frac{\partial}{\partial \eta_0} \quad (57)$$

provided:

$$\frac{1}{\hat{\phi}_0} \frac{\partial \hat{\phi}_0}{\partial \eta_0} \ll nq'x \quad (58)$$

Equation (34) shows that $\eta_0 = -\Delta k_{x0}$. Therefore, re-formulating Equation (58) in terms of k_{x0} , we find the mapping:

$$x \rightarrow i \frac{\partial}{\partial k_{x0}} \quad (59)$$

provided:

$$\frac{1}{\hat{\phi}_0} \frac{\partial \hat{\phi}_0}{\partial k_{x0}} \ll x \quad (60)$$

Therefore, we have different but equivalent conditions depending on which parameter is used, η_0 or k_{x0} . For simplicity, the remainder of this discussion will focus on η_0 .

To assess the implications of Equations (56) and (58) on any requirements for γ_0 to be slowly varying in x and η_0 , we consider the amplitude of $\hat{\phi}_0$. Assuming all local mode structures start at the same amplitude at an arbitrary reference time (which is reasonable for an initial seed perturbation of random noise), then the amplitude varies as:

$$\hat{\phi}_0 \sim \exp[\gamma_0(x, \eta_0)t] \quad (61)$$

Since time is arbitrary, we set $t = 1$ to simplify the analysis. Substituting Equation (61) into Equations (56) and (58) reveals the conditions:

$$\frac{\partial \gamma_0}{\partial x} \ll nq' \quad (62)$$

and:

$$\frac{\partial \gamma_0}{\partial \eta_0} \ll nq'x \quad (63)$$

This chapter has shown that $\frac{\partial \gamma_0}{\partial \eta_0}$ increases with $\hat{s} = \frac{r}{q}q'$. However, notice that the right-hand side of Equation (63) also contains q' . Therefore, as $\hat{s} \sim q'$ acts to increase $\frac{\partial \gamma_0}{\partial \eta_0}$, the right-hand side of Equation (63) also increases. Furthermore, this chapter has shown that $\frac{\partial \gamma_0}{\partial k_{x0}}$ does not change with q' , and that the change in $\frac{\partial \gamma_0}{\partial \eta_0}$ is a consequence of the fixed range of $|\eta_0| \leq \pi$ being mapped to a larger range of $k_{x0} = -nq'\eta_0$ (Equation (34)). Therefore, neglecting other effects such as the enhancement of narrowing by the drifts, the change in $\frac{\partial \gamma_0}{\partial \eta_0}$ is linearly proportional to q' . Hence, if Equation (63) is satisfied by KBMs wide in η_0 at low q' , then it is also satisfied by the same KBMs narrow in η_0 at high q' . In other words, KBMs narrow in η_0 are not *necessarily* wide in real-space nor do they necessarily contradict the assumptions of the local approximation – it all depends on the corresponding value of q' .

The observation above suggests the impact of \hat{s} narrowing of KBMs on local approximation validity may be mitigated by the corresponding increase in nq' . To demonstrate this effect the local-global method was applied to artificial but representative data, which was constructed as follows.

The local growth rate, γ_0 , was set as an inverted quadratic in x multiplied by a Gaussian in η_0 peaking with a value of 1 at $(x, \eta_0) = (0, 0)$:

$$\gamma_0(x, \eta_0) = \left(1 - \left(\frac{x}{w_x}\right)^2\right) \exp\left[-\frac{\eta_0^2}{2c^2}\right] \quad (64)$$

where $w_x \sim 1/\frac{\partial\gamma_0}{\partial x}$ is a measure¹⁷ of the width of γ_0 in the x direction such that $\gamma_0 = 0$ at $|x| = w_x$, and:

$$c = \frac{w_{\eta_0}}{2\sqrt{2\ln 2}} \quad (65)$$

where $w_{\eta_0} \sim 1/\frac{\partial\gamma_0}{\partial\eta_0}$ is a measure¹⁸ of the width of γ_0 in the η_0 direction such that $\gamma_0 \sim 0$ for $|\eta_0| \gtrsim w_{\eta_0}$.

The frequency was set as $\omega_0 = \gamma_0$, w_x was set as 0.2, the safety factor at $x = 0$ was chosen as $q_0 = \frac{4}{3}$, the initial magnetic shear was set as $\hat{s} = 1$ and the minor radius as 0.5, which results in $q' = \frac{8}{3}$, and the toroidal mode number was chosen as $n = 54$, which results in $nq' = 144$ and $m_0 = nq_0 = 72$. Local mode structures, $\hat{\phi}_0(\eta; x, \eta_0)$, were set as Gaussians centred on $\eta = \eta_0$ with a full width at half maximum of $\frac{\pi}{2}$.

The global eigenvalue, Ω , and the global mode structure, ϕ , were calculated from these γ_0 , ω_0 and $\hat{\phi}_0$ data using the local-global method as described in Chapter 3. Three tests were run: varying w_{η_0} , varying q' , and varying q' and w_{η_0} consistently (constant $\frac{\partial\gamma_0}{\partial k_{x0}}$). Similar results were obtained when varying w_x (not shown).

For each value tested within each of the three tests, the ‘‘local error’’ was calculated as the difference between the local eigenvalue, Ω_0 , at $(x, \eta_0) = (0, 0)$ and the global eigenvalue, Ω . In the results that follow, this is given

¹⁷Considering only $x = 0$ and $x = w_x$

¹⁸Similarly, considering only $\eta_0 = 0$ and $\eta_0 = w_{\eta_0}$

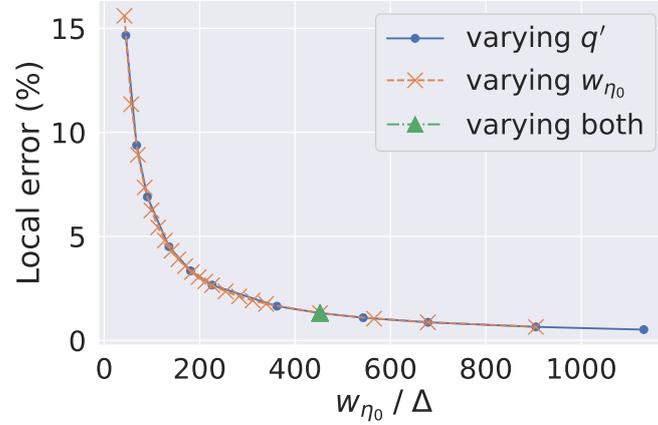


Figure 17: Local error as a function of $\frac{w_{\eta_0}}{\Delta}$. The “varying both” line comprises 9 data points at the same location.

as a percentage of Ω , with a positive value indicating that Ω_0 was an over-estimate compared to Ω . The smaller the local error, the more accurate and, hence, more valid the local approximation.

Given $w_{\eta_0} \sim 1/\frac{\partial\gamma_0}{\partial\eta_0}$ and $\Delta = \frac{1}{nq'}$, and since Equation (63) requires $\frac{\partial\gamma_0}{\partial\eta_0} \ll nq'x$, we therefore expect the local approximation to be more accurate when $\frac{w_{\eta_0}}{\Delta} \gg 1$. The local error from each of the three tests is plotted against this metric in Figure 17.

Figure 17 shows that the local error increases as $\frac{w_{\eta_0}}{\Delta}$ becomes small, as expected by Equation (63). This is consistent with the original assertion that $\gamma_0(\eta_0)$ narrowing *could* affect local validity. However, the “varying both” results show that all 9 data points have the same $\frac{w_{\eta_0}}{\Delta}$ and the same local error. That is, the local error does not change when both q' and w_{η_0} are varied consistently. This is because, as expected by Equation (63), local validity depends on $\frac{w_{\eta_0}}{\Delta}$, which does not change when both q' and w_{η_0} are varied consistently since $w_{\eta_0} \propto \Delta$. This demonstrates that the effect of \hat{s} narrowing on local validity *is* mitigated by the corresponding increase in q' in the absence of other effects such as enhancement of the narrowing by the

drifts. Therefore, KBMs that are narrow in η_0 are not *necessarily* in violation of the local approximation, especially if they are in a region of high \hat{s} (as was the case for those observed in Chapter 4); indeed, KBMs that are *wide* in η_0 may be in violation of the local approximation if \hat{s} is low. Hence, it is essential for confidence in the accuracy of local models to check the validity of the local approximation by testing the sensitivity to η_0 (and x) relative to nq' .

4.6 Conclusions

This chapter has observed KBMs narrow in ballooning angle, η_0 , in local gyrokinetic simulations of a JET pedestal. This suggests global modes that are radially wide, which may significantly reduce the accuracy of the local approximation, or even result in these local modes being non-physical.

The observed narrowing was shown to be caused by magnetic shear, \hat{s} , which stabilises high η_0 faster than low η_0 , and this was shown to be a robust effect in both the CBC and JET equilibria. Therefore, \hat{s} narrowing is expected in *any* equilibria with high \hat{s} , which includes most pedestal equilibria, where \hat{s} is often high. This may explain previous observations of apparently non-physical KBMs in local gyrokinetic JET pedestal simulations [61].

An explanation for how \hat{s} causes narrowing was provided by ideal ballooning physics [60], whereby \hat{s} narrows the region of ballooning space where field line bending is minimised and field lines on neighbouring flux surfaces are aligned. This effect was then shown to cause narrowing of both IBMs and KBMs. Furthermore, this physics causes splitting of the IBM stability region in η_0 , leading to a modification of the second stability region, which highlights the importance of considering $\eta_0 \neq 0$.

The narrowing effect of \hat{s} was shown to enter the gyrokinetic equation via the perpendicular wavenumber, k_{\perp} . This is consistent with the ideal ballooning description of increased \hat{s} causing a narrowing of the region where field lines on neighbouring flux surfaces are aligned; hence, KBM narrowing is primarily caused by ideal ballooning physics. However, in the split IBM case, additional kinetic effects appear at $\eta_0 \sim 0$ to result in unstable KBMs where IBMs are stable. This may explain observations where pedestal evolution appeared to not be constrained by IBM stability at all times [62], and recent work that identified weakly unstable KBMs at $\eta_0 = 0$ where IBMs were stable [68].

While \hat{s} causes narrowing of KBMs in η_0 , re-plotting this data against k_{x0} (the radial wavenumber at $\eta = 0$) shows that there is no narrowing in k_{x0} . Rather, stabilisation with k_{x0} occurs at the same rate, but higher \hat{s} maps the fixed range of $|\eta_0| \leq \pi$ to a larger range of k_{x0} . As a result, k_{x0} stabilisation happens faster in η_0 , hence we observe an apparent narrowing in η_0 . Examining the assumptions of the local approximation reveals that these parameters lead to different but equivalent requirements for γ_0 to be slowly varying in η_0 or k_{x0} (and also in x). These conditions suggest that, in the absence of other effects such as the enhancement of narrowing by the drifts, the impact of \hat{s} narrowing may be mitigated by the corresponding increase in nq' (neglecting other effects). This effect was demonstrated by applying the local-global method to artificial but representative data, which confirmed that the increase in local error due to \hat{s} narrowing *is* mitigated by the corresponding increase in nq' . Therefore, KBMs that are narrow in η_0 are not necessarily in violation of the local approximation, nor do KBMs wide in η_0 necessarily satisfy the local approximation – it all depends on the corresponding value of q' .

To summarise, this work raises some important warnings:

1. While $\eta_0 = 0$ is *often* the most unstable, this is not always the case. Therefore, it is important to consider $\eta_0 \neq 0$, especially in the second stability region.
2. It is important to consider kinetic effects that may impact the validity of using IBMs as a proxy for KBMs.
3. Local calculations, whether gyrokinetic or ideal ballooning, need to consider the impact of global effects on local validity. This is especially true in regions of high \hat{s} , whether pedestal or otherwise, where \hat{s} narrowing may make local modes non-physical (or just *appear* non-physical) and this may or may not be mitigated by the corresponding increase in nq' depending on the strength of other effects such as enhancement of narrowing by the drifts.

Finally, the conditions for local validity derived in this chapter require knowledge of $\frac{\partial \gamma_0}{\partial x}$ and $\frac{\partial \gamma_0}{\partial \eta_0}$. Therefore, it is not necessarily possible to determine *a priori* whether the local approximation is valid for a given case; rather, one must assume local validity, be sure to test the variation of γ_0 with x and η_0 , and, thus, verify local validity afterwards. This is the approach taken in Chapter 5, where the local-global method is applied to KBM simulations in a pedestal-like equilibrium.

Chapter 5

Local-Global Pedestal Simulations of Kinetic Ballooning Modes

5.1 Introduction

5.1.1 Chapter Overview

Chapter 4 showed that there are potential issues with local simulations of the pedestal, and that it is not necessarily possible to determine *a priori* whether the local approximation is valid for a given case. This chapter presents application of the local-global method to a pedestal-like case. An artificial case with low magnetic shear was used for consistency with previous work [17] and to avoid the potential issues presented in Chapter 4. The validity of the local approximation in this case is examined as part of the analysis.

The results are compared (favourably) to previously published [17] global gyrokinetic and global MHD simulations. The results are then used to de-

velop simple correction models for local MHD and assess their utility in an EPED-like calculation.

The structure of this chapter is as follows. First, Subsection 5.1.2 motivates the comparison of global (or local-global) gyrokinetics with MHD, before Subsection 5.1.3 describes the equilibrium model used to do so. Sections 5.2 to 5.4 then cover, respectively, the comparison of the new local-global results against previous work, development of simple corrective models, and application to an EPED-like calculation. These sections include section-specific discussion and conclusions subsections. Finally, Section 5.5 discusses points that are general to the chapter as a whole, before Section 5.6 draws conclusions for the chapter as a whole.

5.1.2 Motivation

Due to the gyrokinetic orderings of Equation (3), gyrokinetic turbulence is inherently Larmor radius scale. However, MHD does not include finite Larmor radius effects; indeed, this length scale does not feature in MHD. Therefore, MHD effects are inherently system sized. As a result, MHD is typically used to study large-scale (low n) instabilities, while gyrokinetics is used to study small-scale (high n) modes. The exception is local MHD, where the small length scale $\frac{1}{nq}$ is introduced with an assumption of large n and finite q' to provide the scale separation necessary to derive a local model (see Subsection 2.2.3).

Given the differing wavelength regimes of gyrokinetics and MHD, it is useful to compare these models to confirm agreement in the intermediate wavelength regime as expected. However, it is yet more useful to assess how well MHD-scale modes are captured by gyrokinetics and *visa versa*. To what extent does MHD remain valid as n gets larger? To what extent does

gyrokinetics remain valid as n gets smaller? And how far can we go in n and still get accurate results from these models?

The answers to these questions will be useful for: (a) improving MHD-based models used as proxies for gyrokinetic models where computational efficiency is crucial, e.g. EPED [14] and Europol [15]; and (b) extending gyrokinetics to capture MHD-scale physics to make gyrokinetics a more complete model, rather than having to run both gyrokinetics and MHD and risk missing interactions between Larmor radius scale and MHD-scale modes.

Global gyrokinetics and global MHD have been compared previously [17], where reasonable agreement was shown. However, the results also showed some significant differences and a number of unanswered questions remain about the origin of these differences. The work presented in this chapter applies the local-global method (see Section 3.2) to the same equilibrium used in [17]. The local-global method is an attractive technique for such cases, as it should be able to overcome some of the limitations and computational expense of global gyrokinetics. Furthermore, the separation of local and global effects, and the use of an alternative gyrokinetic code will allow deeper investigation of the previously observed differences. In doing so, this work will provide further evidence to explain the length scales at which these models agree / disagree and why. This knowledge can then be used to improve both models for use in various applications.

5.1.3 Equilibrium Model

The simulations presented in this chapter use the same equilibrium as [17], which was provided by the authors. This equilibrium is consistent with the Grad-Shafranov equation, and has a pseudo-pedestal at mid-radius, low magnetic shear (to promote growth of ballooning modes), approximately

circular flux surfaces, and non-zero Shafranov shift. The equilibrium also features a flat density profile (at $7.37 \times 10^{19} \text{ m}^{-3}$) so that the pressure pedestal is governed solely by the temperature profile as this is expected to yield better agreement with MHD [69]. This equilibrium was used for ease-of-comparison with previous work [17] and to avoid issues associated with high magnetic shear in the pedestal. Previous work [17] developed this equilibrium with its pseudo-pedestal at mid-radius to avoid boundary effects in their global gyrokinetic simulations. Such issues could potentially be avoided by using the local-global method (this is discussed in Section 5.5). An overview of the equilibrium is given in Figure 18.

The global domain was selected as the region where KBMs are locally unstable. Sensitivity tests showed that a larger domain did not significantly affect the results.

The original equilibrium was produced by the CHEASE equilibrium code [70]. The CHEASE equilibrium is used throughout this chapter. Sensitivity to the equilibrium model is tested in Subsection 5.2.3.2 using a Miller equilibrium model [65] fit to the CHEASE equilibrium.

5.2 Assessment of the Local-Global Method

5.2.1 Local-Global Method – Practical Details

5.2.1.1 Running the Simulations

A reference surface was chosen with $q = m/n = 4/3$ to define the radial origin $x = 0$. This value was chosen as it is near to the point of maximum temperature gradient magnitude, which is expected to be near the centre of the global disturbance. Sensitivity studies showed that the results are insensitive to the exact location of the radial origin, provided it is indeed near to the centre of the global disturbance.

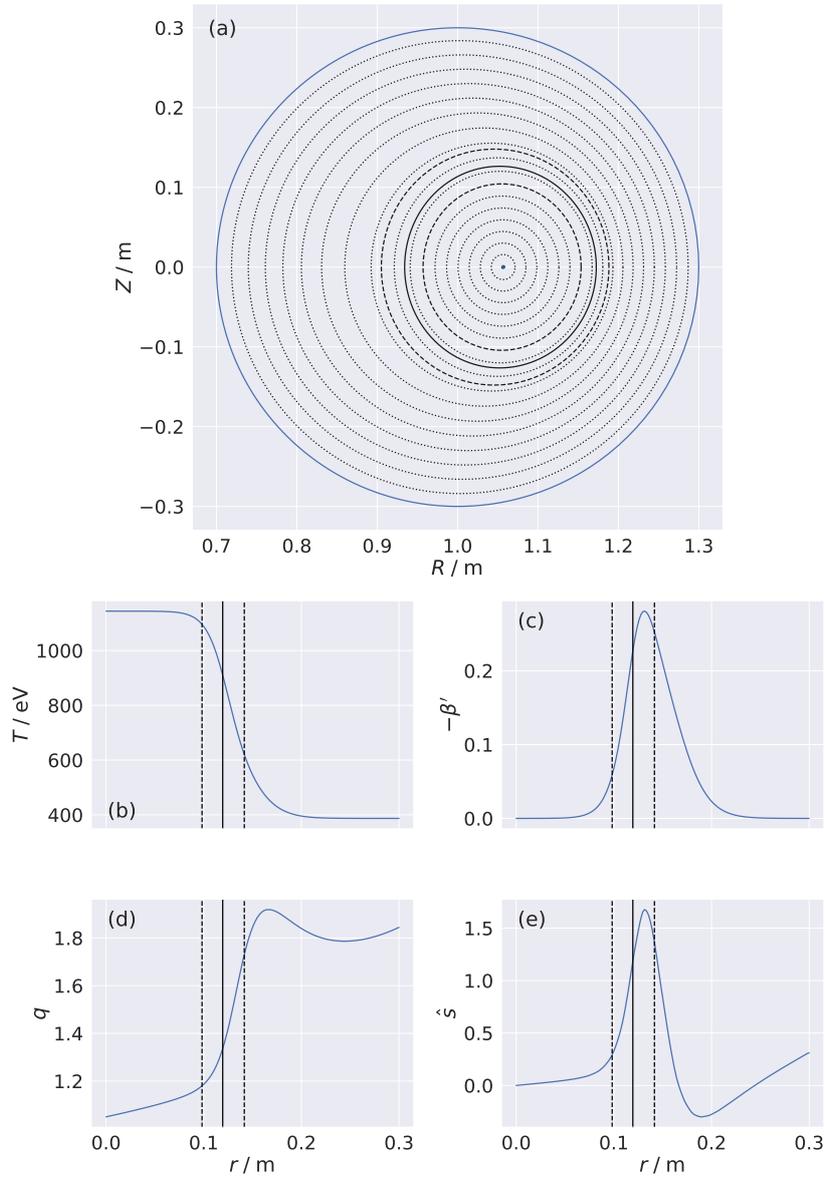


Figure 18: An overview of the pedestal-like equilibrium showing (a) flux surfaces, (b) temperature profile with a pseudo-pedestal at mid-radius, (c) pressure gradient inverse length scale, (d) safety factor profile and (e) magnetic shear profile. In all cases, the local-global simulation domain is indicated by the dashed lines and the origin ($x = 0$) by the solid black line.

The equilibrium normalised Larmor radius is $\rho^* \sim 1/70$. The previous study [17] tested $\rho^* = 1/200, 1/400$ and $1/800$ to ensure the gyrokinetic assumption of small ρ^* was satisfied. For consistency with [17], this thesis also uses $\rho^* = 1/200, 1/400$ and $1/800$.

Local GS2 simulations were run for a range of toroidal mode numbers with normalised Larmor radius, $\rho^* = 1/200$. This ρ^* was used in GS2 to allow flexibility to maintain integer n analysis as ρ^* is changed in the calculation of the global mode for the various ρ^* listed above.

The simulations were run for all of the surfaces within the global domain indicated in Figure 18 that were available within the equilibrium file. For each surface, the ballooning parameter, η_0 , was scanned from 0.0 in steps of 0.2 (determined by preliminary tests to be a suitable value to balance accuracy and performance) up to π or until run-times became prohibitive due to KBMs being stable.

An example of the local frequency and growth rate results obtained are shown in Figure 19. This example is from the $n = 7, \rho^* = 1/200$ case, as this results in $n = 28$ when analysed using $\rho^* = 1/800$, and this is the closest n available to the baseline case from [17] with $n = 30$ and $\rho^* = 1/800$.

5.2.1.2 Obtaining Fit Parameters

Having obtained $\Omega_0(x, \eta_0)$, the coefficients for the differential operator of Equation (31) were calculated by fitting quadratics to the real and imaginary parts of $\Omega_0(x)$ for each η_0 . An example of $\Omega_0(x)$ at $\eta_0 = 0$ for the data from Figure 19 is shown in Figure 20.

Figure 20 shows that we have approximately quadratic behaviour near $x = 0$ as required. Furthermore, the growth rate peaks near $x = 0$ at $x \sim 0.007$ and is nearly symmetric, so we expect a quadratic fit of the growth rate to have a relatively large x^2 coefficient and a relatively small x coefficient.

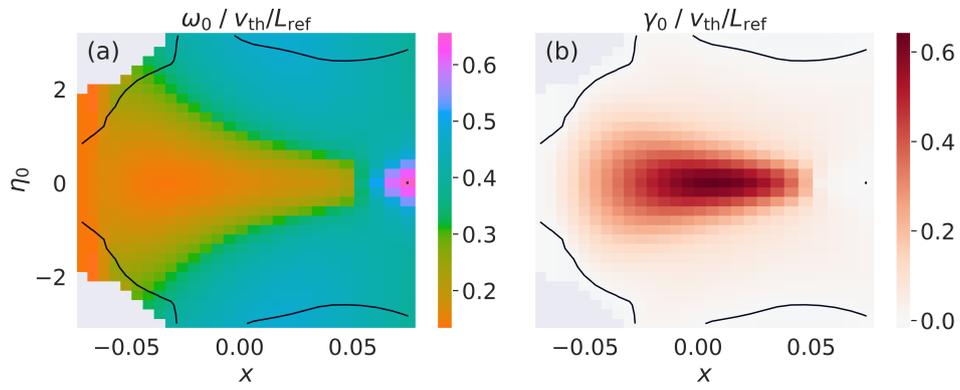


Figure 19: An example of the local (a) frequency and (b) growth rate results as calculated by GS2 over the global domain indicated in Figure 18. The solid black line indicates marginal stability.

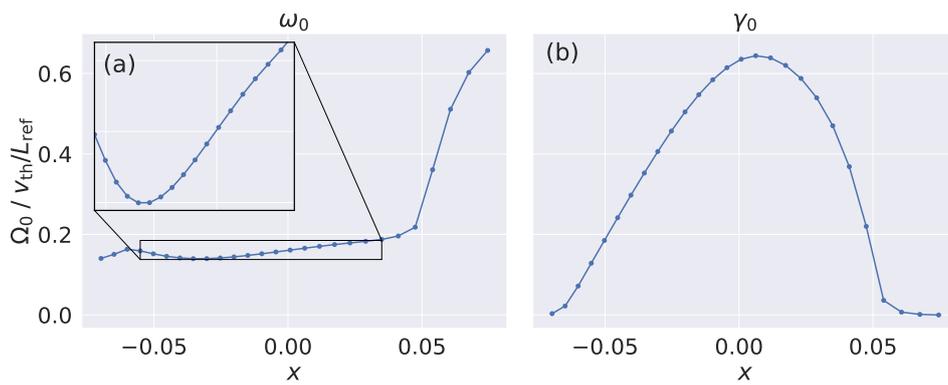


Figure 20: An example of the local (a) frequency and (b) growth rate results as functions of x at $\eta_0 = 0$.

Meanwhile, the turning point in the frequency is somewhat offset from $x = 0$ being instead at $x \sim -0.035$. Therefore, we expect a quadratic fit of the frequency *in the region where x is small* to have a relatively small x^2 coefficient and a relatively large x coefficient. As a result of this coefficient difference between the frequency and growth rate fits, which is itself a consequence of the differing turning point locations, we expect the global mode to take the form of a “mixed mode” with the global peak shifted away from the outboard midplane slightly, as opposed to an “isolated mode” that peaks at the outboard midplane (and occurs with matching turning point locations) or a “general mode” that peaks at the top or bottom (and occurs when both the frequency and growth rate have small x^2 coefficients and large x coefficients) [30]. The reason for the differing turning point locations has not been investigated due to time constraints. However, the Cyclone base case equilibrium used in [29] produced isolated modes. Therefore, it is reasonable to assume that the difference in expected mode type is due to one or more of the differences between the two cases. For example, the case in [29] varies a/L_T while holding T constant, whereas a/L_T and T are varied consistently in the equilibrium used here.

Figure 20 also shows that the quadratic frequency behaviour breaks down far from $x = 0$ as the growth rate approaches marginal stability. Such frequency behaviour near marginal stability is common for KBMs (e.g. [71]). Furthermore, in this case:

- β is relatively high, so ITGs are expected to be largely suppressed in line with [72];
- the frequency does not change sign, which excludes a switch to an electron driven mode such as the trapped electron mode (TEM) or electron temperature gradient mode (ETG);

- the local mode structures do not change from twisting (even) parity to tearing (odd) parity, which excludes a switch to micro-tearing modes (MTMs); and
- ω and γ are continuous in x , which is consistent with the same mode operating at all values of x tested.

Therefore, this appears to be genuine KBM behaviour. In any case, this data is likely to significantly affect the quadratic fit coefficients that are based on an assumption of approximately quadratic behaviour near $x = 0$. Therefore, these data have to be excluded from the fitting. This is justified because the corresponding growth rates are approximately zero so these local modes will not contribute significantly to the global mode.

To avoid having the quadratic fits skewed by data at large values of $|x|$, the fitting was weighted by a Gaussian weighting function to favour data points near $x = 0$, consistent with the assumption of a Taylor expansion about $x = 0$. The weighting function used and an example of the fits obtained with and without the weighting function are shown in Figure 21.

Figure 21 shows that the quadratic fit is greatly improved in the region near $x = 0$ using the Gaussian weighting function as compared to the fit obtained without any weighting. The width of the weighting function was chosen such that the weighting function decayed to approximately zero by the edge of the region where $\Omega_0(x)$ exhibits approximately quadratic behaviour. A sensitivity study with an alternative weighting function (a clipped inverse quadratic) showed the fitting to be insensitive to the fine details of the weighting function.

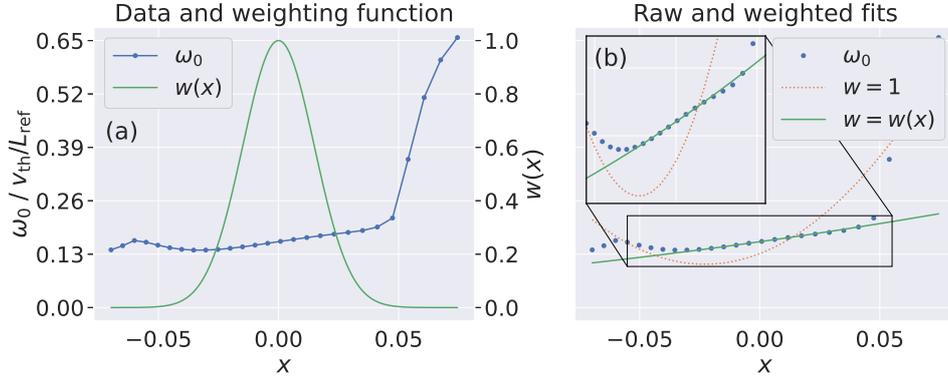


Figure 21: (a) The local frequency data, ω_0 , from Figure 20 alongside the Gaussian weighting function, $w(x)$, used to favour data points near $x = 0$ when fitting a quadratic function to the data; and (b) the same ω_0 data alongside quadratic fits obtained with and without use of the Gaussian weighting function ($w = w(x)$ and $w = 1$ respectively).

5.2.2 Local-Global Results

5.2.2.1 Local Benchmarking

The local results were benchmarked against previously published data obtained by digitising [17, Fig. 9 (a)], which shows the maximum local growth rate found within the global domain. A comparison of that data with the equivalent data from the new simulations is shown in Figure 22. The effective toroidal mode number, n_{eff} , was calculated as:

$$n_{\text{eff}} = n(\rho^*/\rho_0^*) \quad (66)$$

where ρ_0^* is the equilibrium normalised Larmor radius with a value of $\rho_0^* \sim 1/70$. The GS2 normalised units of growth rate were converted to physical units for ease of comparison with the data from [17].

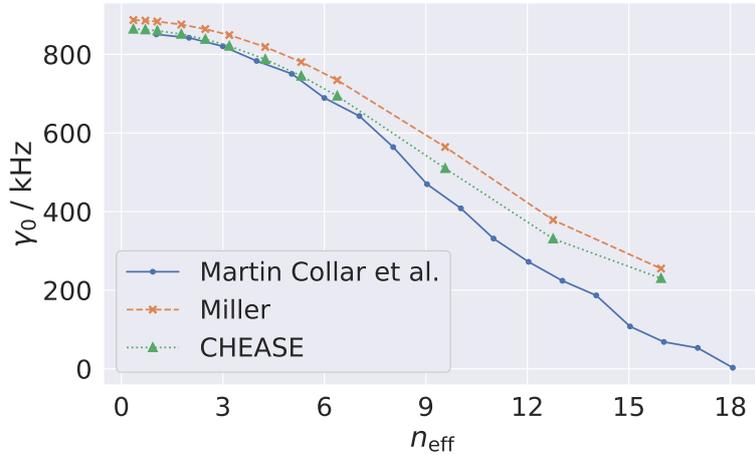


Figure 22: Comparison of the maximum local growth rate found within the global domain against equivalent data from [17]

This comparison shows that using the CHEASE equilibrium results in local growth rates that match those of [17] very well up to $n_{\text{eff}} \sim 7$ and match reasonably well at higher n_{eff} with slightly higher growth rates in the new data. The reason for the mismatch at high n_{eff} was investigated through convergence and sensitivity studies but a explanation was not found. However, the difference is not expected to be significant for the local-global analysis because the local growth rates are relatively small at these values of n_{eff} , so these mismatching local modes have a relatively small contribution to the global mode.

Figure 22 also shows that using the Miller equilibrium results in local growth rates that match the trend of the new CHEASE results, though with a slightly higher growth rate at all n_{eff} . A small difference is expected when using a different equilibrium model, so the data in Figure 22 can be taken as evidence that the Miller representation is a reasonable approximation of the numerical CHEASE equilibrium in this case. The effect of this in the local-global analysis is tested in Subsection 5.2.3.2.

5.2.2.2 Global Mode Structure

Having obtained the coefficients for the differential operator of Equation (31) (Subsection 5.2.1.2) and performed local benchmarking (Subsection 5.2.2.1), the amplitude envelope, $A(\eta_0)$, and global eigenvalue, Ω , were calculated by solving Equation (31) with periodic boundary conditions. The q' value used in the analysis was that on the reference surface at $x = 0$ with a value of $q' \sim 3.9$. Sensitivity to this was not tested as this is the equilibrium value, but the results in Chapter 4 suggest that the local correction should decrease with increased q' .

The global mode structure, $\phi(x, \theta)$, was then re-constructed using $A(\eta_0)$ and $\phi_0(\eta; x, \eta_0)$. The local mode structures calculated by GS2 are of arbitrary absolute magnitude and phase and are of arbitrary magnitude and phase relative to each other; only the relative magnitude and phase within a given local mode structure is non-arbitrary. The correct relative amplitude comes in part from the growth rates. The correct relative amplitude also comes from $A(\eta_0)$ as does the correct relative phase; the absolute magnitude and phase remain arbitrary. Therefore, the local mode structures calculated by GS2 were normalised such that the data point with the peak magnitude had a value of exactly 1.0 and was exactly real, i.e. such that all local mode structures are of the same amplitude and phase. These were then scaled by the corresponding local growth rates such that the peak amplitude was equal to e^{γ_0} , thus effectively assuming that all local modes have grown from the same starting size at their individual growth rates for one normalised time unit. This choice of time is arbitrary and was chosen for convenience. Sensitivity tests show that the global mode structure is insensitive to the time chosen for small times.

The resulting normalised and scaled $\phi_0(\eta; x, \eta_0)$ was converted to $\phi_0(\theta; x, \eta_0)$ using Poisson summation (Equation (32)). Further correction of the ampli-

tude and phase is then provided by $A(\eta_0)$ when the global mode structure is calculated.

The global mode structure, $\phi(x, \theta)$, was then calculated by substitution of $\phi_0(\theta; x, \eta_0)$ and $A(\eta_0)$ into the Fourier-ballooning representation (Equation (24)), direct evaluation of the generic poloidal and radial fluctuations, $\exp[-in(q_0 + q'x)\theta]$ and $\exp[inq'x\eta_0]$ respectively, and evaluation of Equation (24). An example of the global mode structure obtained is shown in Figure 23 alongside the global mode structure from [17] for comparison. Both cases use $\rho^* = 1/800$, and the case from [17] has $n = 30$, while the case from the current work has $n = 28$ as this was the closest n available in the current data due to running GS2 with $\rho^* = 1/200$ and integer n . When preparing Figure 23 (b), the poloidal angle θ was plotted as the negative of the “poloidal angle-like coordinate”, χ , as defined in [17] and the CHEASE equilibrium, because comparison of Subsection 2.2.2 and [17] show that $\theta \equiv -\chi$, with the minus due to a differing sign convention.

Figure 23 (a) shows that the amplitude envelope, $A(\eta_0)$, is as expected, i.e. an approximately Gaussian shape, in this case peaking slightly away from $\eta_0 = 0$ consistent with the amplitude envelope expected for a mixed mode. Figure 23 (b) shows that the global mode structure looks as expected, i.e. has a relatively narrow radial width, has the expected poloidal mode number, and the peak is slightly shifted away (downwards in this case) from the outboard midplane (notice the slight asymmetry about $\theta = 0$) consistent with a mixed mode as predicted. These results (Figure 23 (a) and (b)) are qualitatively consistent with previous local-global work studying a different case [28]. Finally, Figure 23 (b) and (c) show that the global mode structure obtained from this analysis has some similarities and differences compared to that obtained from ORB5 [17]. Both have a similar poloidal mode number and radial width, as expected. However, the structures shear towards the low-field side in the local-global case, but towards the top / bottom in

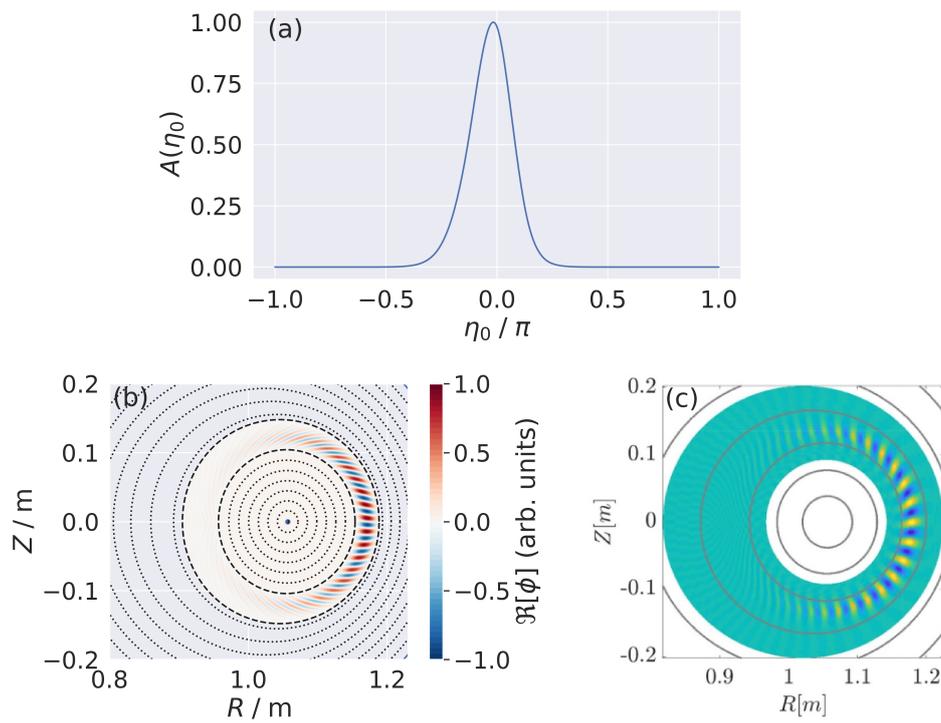


Figure 23: An example of the global mode structure obtained from this analysis showing (a) the amplitude envelope, $A(\eta_0)$; (b) the global mode structure; and (c) a similar global mode structure from [17] for comparison.

the ORB5 case. This difference was investigated, including examining the possibility of a difference in sign convention in the poloidal angle coordinate or magnetic shear definition, but the discrepancy could not be resolved.

5.2.2.3 Global Growth Rate

As noted in Subsection 5.2.1.1, the simulations were run for a range of toroidal mode numbers, n . Therefore, $A(\eta_0)$ and Ω could be calculated for each n to obtain $\Omega(n)$. Furthermore, n and $\rho^* = 1/200$ were specified in the GS2 input files, but GS2 uses k_y internally, which is calculated from n and ρ^* as follows:

$$k_{y,GS2} = n\rho^* \frac{d\rho}{d\psi} \quad (67)$$

where $\frac{d\rho}{d\psi}$ is calculated within GS2 from the magnetic geometry. Therefore, the analysis may be repeated for different values of ρ^* using the same data (same k_y) by changing the values of n used in the calculation of the coefficients of Equation (31) accordingly, e.g. double n to half ρ^* . This was done and the results are shown in Figure 24 labelled ‘‘GS2 (CHEASE, full)’’ alongside the global MHD (MISHKA) and global gyrokinetics (ORB5) data from [17] for comparison. Additional lines (CHEASE, reduced and Miller, full) are shown for discussion in Subsection 5.2.3.

Figure 24 shows that the global growth rate obtained with the local-global method matches that from direct global gyrokinetic simulations reasonably well across a wide range of n and ρ^* . However, the local-global results have a higher peak growth rate (at intermediate n), a lower growth rate at high n and a higher growth rate again at very high n compared to the direct global gyrokinetic results. Furthermore, the local-global results match the global MHD results more closely than the match between the global MHD and global gyrokinetic results. The reason for this is investigated in

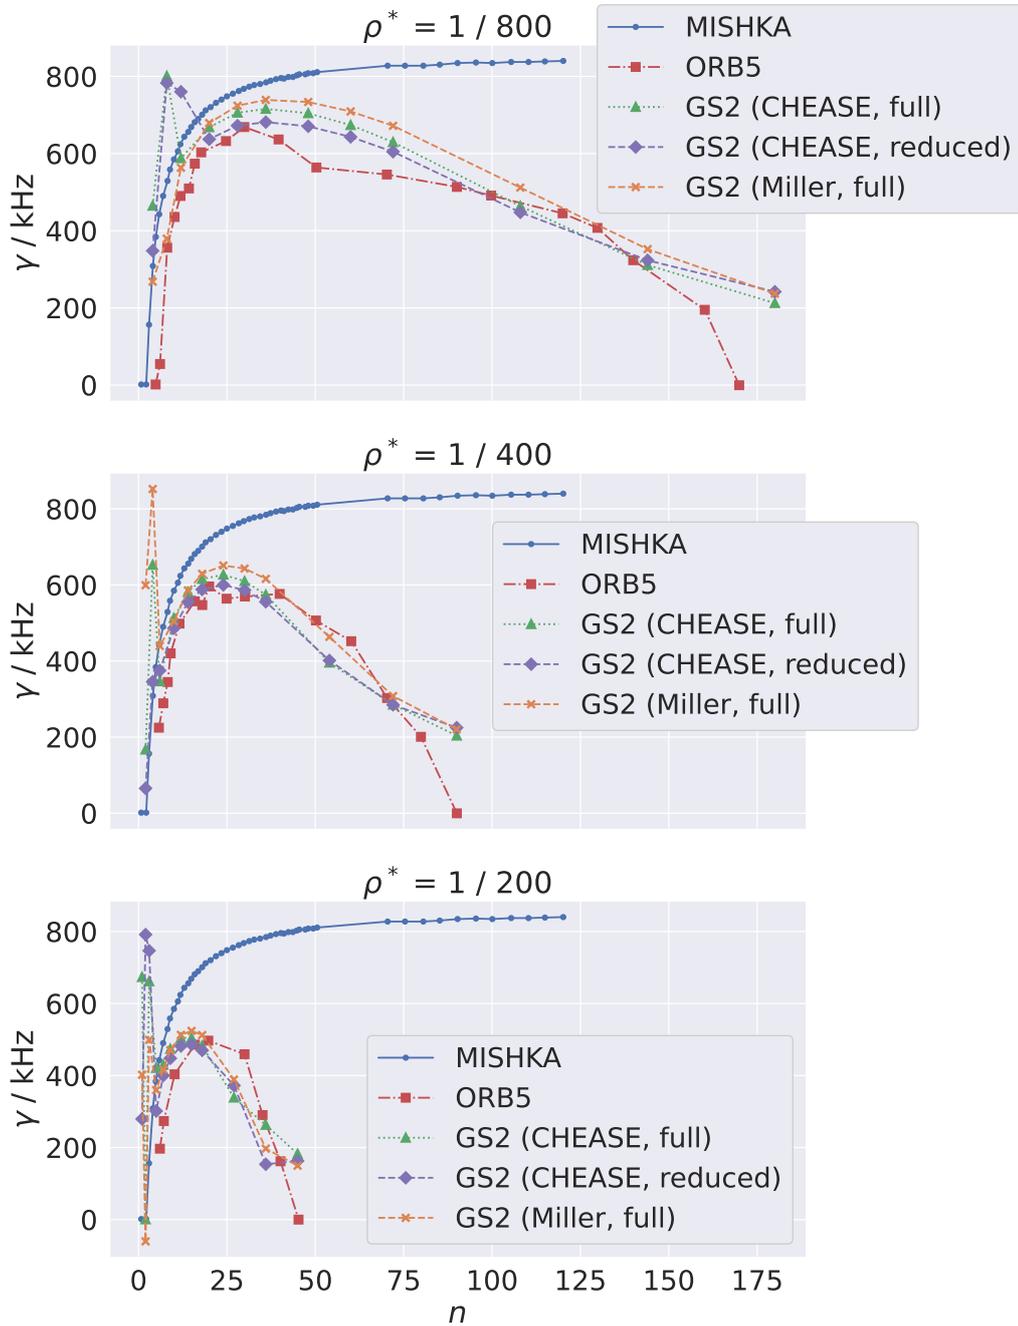


Figure 24: Local-global results for various ρ^* from GS2 using the CHEASE equilibrium with the full δB_{\parallel} model, reduced δB_{\parallel} model, and the Miller equilibrium model, compared against global MHD (MISHKA) and global gyrokinetics (ORB5) data digitised from [17, Fig. 7].

Subsection 5.2.3. Finally, the local-global method breaks down at low n as expected. However, in this case, this occurs at *surprisingly low* values of $3 \lesssim n \lesssim 12$ depending on the value of ρ^* .

5.2.3 Discussion – Local-Global Comparison

5.2.3.1 Reduced δB_{\parallel} Model

Figure 24 shows that there is a difference between the global growth rates predicted by the local-global method compared to the direct global gyrokinetic simulations. The reason for this was investigated and found to be due in part to the handling of parallel magnetic field fluctuations, δB_{\parallel} , which are neglected in ORB5. As a result, the global gyrokinetic simulations reported in [17] use an *ad-hoc* model to capture the effect of δB_{\parallel} , whereby the ∇B drift was set equal to the curvature drift and the pressure gradient doubled. GS2 models δB_{\parallel} explicitly, so this is a potentially important difference between the two sets of simulations. The flexibility of GS2 allows the handling of δB_{\parallel} to be turned off (by setting `fbpar = 0.0`) and the *ad hoc* model of [17] to be tested directly (via the `gb_to_cv` option). This was done and the results are also shown in Figure 24, labelled “GS2 (CHEASE, reduced)”.

The “GS2 (CHEASE, reduced)” lines in Figure 24 shows that the use of the reduced δB_{\parallel} model improves the match to the ORB5 data in the region of peak growth rate (intermediate n) and towards lower n until the local-global method breaks down. This suggests that the discrepancy between the GS2 local-global and ORB5 results in this region is due in part to the different handling of δB_{\parallel} in the two codes.

5.2.3.2 Miller Equilibrium Model

In addition to investigating the δB_{\parallel} model as above, the Miller equilibrium model was also tested. The results are also shown in Figure 24, labelled “GS2 (Miller, full)”.

The “GS2 (Miller, full)” lines in Figure 24 shows that the use of the Miller equilibrium causes a slight increase in the local-global growth rate at intermediate and high n for all ρ^* . Meanwhile, there is little effect at low n . The results continue to break down at very low n , except in the $\rho^* = 1/800$ case, where the local-global method using the Miller equilibrium appears to remain valid down to $n = 4$, which was the lowest value tested. The intermediate and high n increase to the growth rate when using the Miller equilibrium appears to be due to the difference in the local results shown in Figure 22. This suggests that the global results are sensitive to the equilibrium treatment. Therefore, while the GS2 local-global CHEASE results use the same equilibrium as the ORB5 results, it is likely that the two codes handle the equilibrium magnetic geometry differently and it is possible that this is an additional source of the observed differences between the GS2 local-global results and the ORB5 results.

5.2.3.3 Other Factors

The discrepancy between the GS2 local-global and ORB5 results at high n was investigated, but a definitive reason was not found. Notice that there is also disagreement at high n_{eff} shown in Figure 22 between the new GS2 and previously published GS2 data. As noted in Subsection 5.2.2.1, this was also investigated but a reason was not found. However, the fact that there is disagreement at high n even between two different uses of GS2 suggests that the local results at least are sensitive to one or more input parameters in this region and / or the code version used. Therefore, it is conceivable

that the ORB5 simulations are also sensitive to one or more input parameters in this region. In any case, the results in this region are less significant for global KBM stability than those near the peak global growth rate, so this discrepancy was not investigated further.

A number of other possible reasons for the observed difference between the global MHD and global gyrokinetic results at low n were proposed in [17, §9]. These were shown in the current work to be insignificant. The proposed reasons are each addressed in turn in the following subsections with an explanation of how each was eliminated.

Kink Drive

[17, §9] proposed the absence of kink drive from ORB5 but not from MISHKA as a possible reasons for the observed difference between the results at low n . However, kink drive is absent from both ORB5 and GS2, and an improved match to the MISHKA data at low n was achieved using GS2. This suggests that the absence of kink drive is not significant in this case.

Gyrokinetics Invalid at Low n

[17, §9] suggests that gyrokinetics may be fundamentally invalid for long wavelengths, i.e. low n . This is certainly true of the local approximation, which assumes large n . The consequence of this is evident in the break-down of the local-global method at low n in Figure 24. However, for gyrokinetics, the picture is more complicated. The gyrokinetic orderings of Equation (3) assume $\rho_i \ll L$, thus limiting the validity to fluctuations with a wavelength:

$$\lambda \ll L \tag{68}$$

The wavelength can be approximated as:

$$\lambda \approx \frac{2\pi r_0}{nq} \quad (69)$$

where we typically have $L/2 \lesssim r_0 \lesssim L$. Taking $r_0 \sim L/2$ as is the case in the current work and combining that with Equations (68) and (69), we find the requirement on n to be:

$$n \gg \pi/q \quad (70)$$

or, since this is all rather approximate, simply:

$$n \gg 1 \quad (71)$$

Equation (71) shows that it is technically true that gyrokinetics is strictly invalid as $n \rightarrow 1$. But how much greater than 1 does n need to be for gyrokinetics to remain valid? Figure 24 shows that the local-global method and hence both the local approximation and gyrokinetics remains valid down to $3 \lesssim n \lesssim 12$ depending on the value of ρ^* and the equilibrium treatment. In all cases, it appears that the local-global method remains valid as n is decreased until the local-global results are approximately equal to the global MHD results, and the local-global method breaks-down for n lower than that. This suggests a cross-over regime around these values of n where gyrokinetics and MHD agree, above which gyrokinetics is more accurate than MHD, and below which MHD is more accurate than gyrokinetics.

Parallel Magnetic Field Perturbations

[17, §9] suggests that the work of [73] indicates that the δB_{\parallel} treatment is unlikely to explain the observed difference between the global MHD and global gyrokinetic results. However, Figure 24 shows that use of the reduced δB_{\parallel} model accounts for most of the observed difference. Furthermore, this

also causes the local-global method to break down at a higher n in the $\rho^* = 1/800$ case compared to use of the full δB_{\parallel} model ($12 \lesssim n \lesssim 16$ rather than $8 \lesssim n \lesssim 12$). In addition, closer inspection of [73] reveals that they find reduced δB_{\parallel} models can result in significant artificial stabilisation due to perturbed magnetic curvature, consistent with the findings in the current work.

Trapped Particles and Collisions

Effects related to trapped particles are cited in [17, §9] as a potential explanation for the observed difference between global MHD and global gyrokinetics at low n , as they are “absent in the collisional limit of MHD” (and, indeed, in MHD models in general). Collisional effects more broadly are also cited as a possible reason. However, simulations during the current work varying trapped particle treatment and the handling of collisions in numerous ways (not shown) indicate that the results are insensitive to these parameters. This eliminates these suggestions from being the cause of the observed discrepancies.

Expected Wavelength for Agreement

Another possible reason given in [17, §9] for the mismatch between global MHD and global gyrokinetics is that, “The derivation in the short-wavelength, used to show that MHD and gyrokinetics should agree in certain parameter regimes, does not apply for long-wavelength modes.” It is not clear what derivation this comment refers to. However, some observations of the actual and expected wavelength for agreement between global MHD and the new local-global results are discussed below.

The gyrokinetic orderings (Equation (3)) order the drifts slow compared to the thermal velocity ($\frac{v_D}{v_{th}} \ll 1$) and allows fast fluctuations ($\frac{\chi_1}{\chi_0} \ll 1$). Conversely, MHD assumes no fast fluctuations ($\chi_1 = 0$) and orders the drifts

as fast as the thermal velocity (i.e the MHD ordering, $\frac{v_D}{v_{th}} \sim 1$). Therefore, these models are fundamentally aimed at different regimes, so it would appear that they are not necessarily expected to agree except where the orderings become equal.

The fluctuation ordering becomes equivalent as $\rho_i \rightarrow 0$. Indeed, Figure 24 shows that the local-global results approach the global MHD results as ρ^* gets smaller.

To examine the velocity ordering, it is useful to consider the example of the drift velocity as the $\mathbf{E} \times \mathbf{B}$ drift velocity and to recall the form of the $\mathbf{E} \times \mathbf{B}$ drift and thermal velocities, which lead to:

$$\frac{v_{\mathbf{E} \times \mathbf{B}}}{v_{th}} = \frac{\mathbf{E} \times \mathbf{B} \sqrt{m_i}}{B^2 \sqrt{2T_i}} \quad (72)$$

For a given equilibrium, the only variable in this is the electric field, which is related to the electrostatic potential by:

$$\mathbf{E} = -\nabla\phi \quad (73)$$

In the short-wavelength limit, ϕ varies rapidly in space. Therefore, \mathbf{E} averages to a small value over a relatively short distance. This puts us in the gyrokinetic regime. Conversely, in the long-wavelength limit, turning points in ϕ are separated by large distances in space. Therefore, the electric field acts in the same direction over a large distance so does not average to a small value, even over relatively large distances. This puts us in the MHD regime. The contrast of these two limits implies a cross-over regime at intermediate n . This is consistent with the data in Figure 24, where the local-global gyrokinetic and global MHD results converge towards the same value at some intermediate n of $3 \lesssim n \lesssim 12$ depending on the value of ρ^* .

At lower n , the global MHD results can be taken as more accurate, and *visa versa* at higher n .

MISHKA Plasma Inertia Approximation

It is suggested in [17, §9] that the source of the discrepancy between global MHD and global gyrokinetics at low n could be the approximation for the plasma inertia used in MISHKA. However, the fact that the GS2 local-global results agree well with the MISHKA results at low-to-intermediate n suggests that the cause of the discrepancy is more likely to be something that is absent or approximated in ORB5 rather than MISHKA, such as the δB_{\parallel} treatment.

Perpendicular vs. Poloidal Wavenumber

It is noted in [17, §9] that ORB5 uses the poloidal wavenumber, k_p , in place of the perpendicular wavenumber, k_y , as an approximation, and suggests this could be the source of the discrepancy between global MHD and global gyrokinetics at low n . However, the angle, w , between the corresponding wavevectors is:

$$\begin{aligned} w &= \arctan \left[\frac{2\pi r/m}{2\pi R/n} \right] \\ &= \arctan \left[\frac{\varepsilon}{q} \right] \end{aligned} \tag{74}$$

Therefore, the ratio $\frac{k_p}{k_y}$ is:

$$\begin{aligned} \frac{k_p}{k_y} &= \cos[w] \\ &= \cos \left[\arctan \left[\frac{\varepsilon}{q} \right] \right] \end{aligned} \tag{75}$$

In this case, we have $\varepsilon \sim 0.1$ and $q \sim 4/3$, hence we have $\frac{k_p}{k_y} \sim 0.997$. Therefore, this approximation makes a difference of less than 1% to the

wavenumber in this case, which will have a similarly small impact on the local-global results. Thus, we can conclude that the effect of this approximation is insignificant in this case. Indeed, for this approximation to be significant would require large ε (e.g. spherical tokamaks) and low q , so is unlikely to affect realistic pedestal simulations where q is generally higher than that used here.

Equilibrium Treatment

Finally, [17, §9] suggests the equilibrium treatment as a potential source of discrepancy between the global MHD and global gyrokinetic results at low n . As shown by Figure 24, the GS2 local-global results are quite sensitive to the equilibrium treatment (i.e. the way in which the code calculates geometric factors for a given equilibrium model), giving different results for different models of the same equilibrium (i.e. the local-global growth rate increases by $\sim 3\%$ when switching from the CHEASE to Miller equilibrium model, which may be attributed to the equilibrium treatment, as discussed in Subsection 5.2.3.2). Therefore, it is conceivable that the MISHKA and ORB5 results may also be sensitive to any equilibrium treatment differences between those codes.

5.2.4 Conclusions – Local-Global Comparison

Based on the evidence presented in Subsection 5.2.3, it appears that the previously observed difference between global MHD and global gyrokinetics is mainly accounted for by the reduced δB_{\parallel} model used in ORB5 and possibly also the equilibrium treatment. Other factors suggested in [17, §9] have been shown to have negligible effect.

The new local-global gyrokinetic results show an improved match to the global MHD results due to the use of a fully self-consistent calculation of δB_{\parallel} effects. The local-global method appears to remain valid in this case

down to relatively low n of $3 \lesssim n \lesssim 12$ depending on the value of ρ^* . This suggests that, in principle, it may be possible to include additional MHD-scale physics (e.g. kink drive) in gyrokinetics to make a more complete model that captures the interactions between Larmor radius scale and MHD-scale modes. Furthermore, the apparent validity of the local-global method down to low n suggests that the separation of local effects from global effects is reasonable in this case down to these values of n . This is of interest for pedestal simulations where global codes have concerns relating to boundary conditions (see Subsection 5.5.1 for further discussion). The results also demonstrate explicitly that low n stabilisation is due to global effects (i.e. destructive interference of poloidal harmonics) and high n stabilisation is due to gyrokinetic effects, namely the diamagnetic drift stabilisation cited in [17], which is investigated further in Section 5.3. This is evident in Figures 22 and 24, where the low n stabilisation is only present in the global results (global MHD, global gyrokinetics and local-global gyrokinetics), while the high n stabilisation is only present in the gyrokinetic results (local gyrokinetics, global gyrokinetics and local-global gyrokinetics), which include diamagnetic drift stabilisation, but not in the global MHD results, which do not include this effect.

5.3 Development of Simple Correction Models

Section 5.2 showed that local and global effects are separable, that low n stabilisation is due to global effects, and that high n stabilisation is due to gyrokinetic effects. The latter is expected to be the diamagnetic drift stabilisation cited in [17]. To study this further, simple models of the global and diamagnetic corrections are developed and applied to the gyrokinetic and MHD data to see if the results of Section 5.2 can be reproduced using only the simple corrective models. This will improve our understanding of

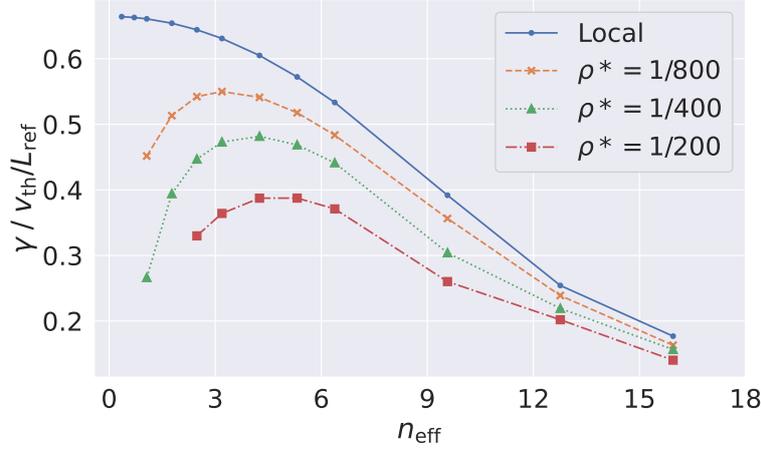


Figure 25: Global growth rate as a function of n_{eff} for various ρ^* alongside the corresponding local growth rates.

the underlying physics, and assist us in improving EPED-like models that rely on rapid calculations covering many different inputs.

5.3.1 Global Correction Model

Figure 25 shows the local growth rates from Figure 22 alongside the local-global growth rates from Figure 24 plotted against n_{eff} to facilitate comparison. The data is now shown in GS2 units for convenience. The data at low n where the local-global method breaks down has been removed for clarity. Figure 25 highlights the nature of the global correction, i.e. that the global growth rate is reduced compared to the peak local growth rate, and that this effect is stronger for smaller n and larger ρ^* .

The local-global method's second order differential eigenequation (Equation (31)) shows we can expect the reduction to the local growth rate due to global effects to be an $\mathcal{O}\left(\frac{1}{nq'}\right)$ correction. Therefore, we start by fitting the following model to the local-global GS2 results for each ρ^* to obtain the fit parameter a :

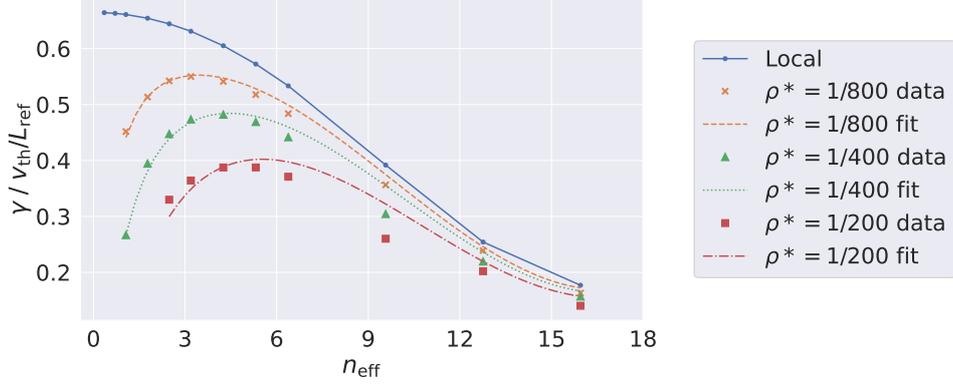


Figure 26: Fits of the global correction model (Equation (76)) to the data from Figure 25 using a of 18.9, 21.0 and 20.7 for the ρ^* of 1/800, 1/400 and 1/200 respectively

$$\gamma = \gamma_0 \left(1 - \frac{1}{nq'}\right)^a \quad (76)$$

This results in values for a of 18.9, 21.0 and 20.7 for ρ^* of 1/800, 1/400 and 1/200 respectively. Plotting the model of Equation (76) with these values of a results in the curves shown in Figure 26.

Figure 26 shows that the global correction model of Equation (76) produces a good approximation of the global growth rates for all ρ^* tested, although less so as ρ^* increases. This includes a prediction of the peak global growth rate that agrees with the local-global method to within 1%, 1% and 4% for ρ^* of 1/800, 1/400 and 1/200 respectively.

Figure 26 also shows that the magnitude of the global correction, $\Delta\gamma_{\text{global}}$, decreases with ρ^* , as expected. The reason for this is evident from examination of Equation (76). Applying the binomial expansion, we find that $\Delta\gamma_{\text{global}} \sim \frac{a}{nq'}$. Indeed, this form of Equation (76) also produces a reasonable fit (not shown) to the data for $nq' \gg 1$ but not at low n . This is because the power law form of Equation (76) also includes the gyrokinetic $O((nq')^2)$

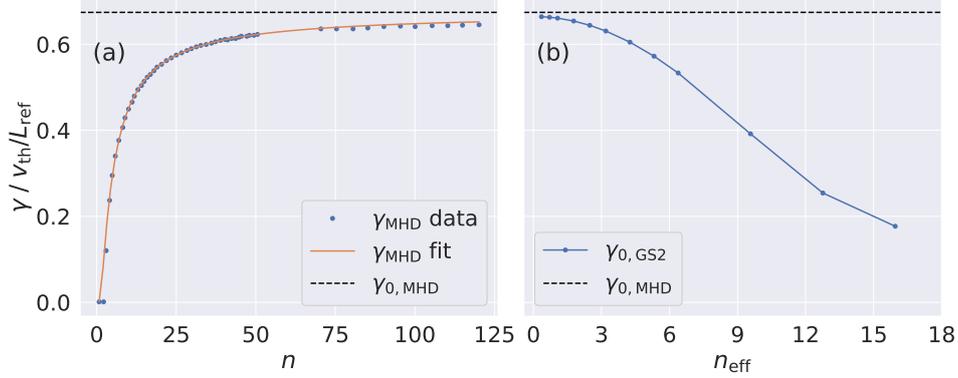


Figure 27: (a) Reverse application of the global correction model (Equation (76), solid orange line) to the global MHD data (blue dots). This also shows the effective local MHD growth rate (black dashed line). (b) The local GS2 growth rates (blue dots and solid line) for comparison.

corrections¹⁹, so can remain valid to lower n . Given the relationship between n and n_{eff} (Equation (66)), this suggests that $\Delta\gamma_{\text{global}} \propto \rho^*$, as observed.

The global correction model of Equation (76) may be tested further using the global MHD data, γ_{MHD} . The local MHD growth rate, $\gamma_{0, \text{MHD}}$, is independent of n . Therefore, we substitute the names γ_{MHD} and $\gamma_{0, \text{MHD}}$ into Equation (76) in place of γ and γ_0 respectively (the values are still unknown), and fit the model with a and $\gamma_{0, \text{MHD}}$ as the fit parameters. The fit was weighted to favour the region of the peak global gyrokinetic growth rate of $10 \lesssim n \lesssim 40$ since this is the primary region of interest. This was done using maximum weighting for all values within the desired range with the weighting decreasing linearly with n outside that range. The result of the fit is that we obtain: the effective $\gamma_{0, \text{MHD}}$ corresponding to the global MHD data, with a value of $\gamma_{0, \text{MHD}} = 0.674 v_{\text{th}} / L_{\text{ref}}$; and the fit parameter a , with a value of $a = 15.6$. The result of this fit is shown in Figure 27.

¹⁹The power law form does not include other $O(n^2)$ corrections that are ordered out of gyrokinetics, such as current gradient drive

Figure 27 shows that the global correction model of Equation (76) also fits the global MHD data very well. Furthermore, this produces an effective $\gamma_{0,\text{MHD}}$ that is slightly higher than the maximum γ_{MHD} and $\gamma_{0,\text{GS2}}$ as expected. Therefore, we can conclude that the global model of Equation (76) is a suitably accurate approximation in this case.

5.3.2 Diamagnetic Correction Model

To develop a simple corrective model to capture the effect of diamagnetic drift stabilisation, we use the relationship between the gyrokinetic complex frequency, $\Omega = \omega + i\gamma$, and the MHD eigenvalue, Ω_{MHD}^2 , noted in [17, §3]:

$$\Omega_{\text{MHD}}^2 = \Omega(\Omega + \omega_*) \quad (77)$$

where ω_* is the diamagnetic frequency, which is the product of the diamagnetic drift velocity and the wavenumber in that direction (which happens to be the binormal direction). Therefore, ω_* is given by:

$$\omega_{*j} = \frac{T_j[\text{eV}]}{Z_j} \left(\frac{1}{T_j} \frac{\partial T_j}{\partial \rho} + \frac{1}{n_j} \frac{\partial n_j}{\partial \rho} \right) \frac{1}{B} k_y \quad (78)$$

where j is the species label. Alternatively, ω_* may be given in terms of n thus:

$$\omega_{*j} = \frac{T_j[\text{eV}]}{Z_j} \left(\frac{1}{T_j} \frac{\partial T_j}{\partial \rho} + \frac{1}{n_j} \frac{\partial n_j}{\partial \rho} \right) \frac{n}{Ba} \frac{\partial \rho}{\partial \psi} \quad (79)$$

Due to the Hermitian nature of MHD, and hence the requirement for real MHD eigenvalues, MHD eigenvalues with $\Omega_{\text{MHD}}^2 > 0$ are stable while those with $\Omega_{\text{MHD}}^2 < 0$ are unstable and we have the conditions:

$$\Omega_{\text{MHD}}^2 > 0 \Rightarrow \gamma_{\text{MHD}} = 0 \Rightarrow \Omega_{\text{MHD}} = \omega_{\text{MHD}} \quad (80)$$

$$\Omega_{\text{MHD}}^2 < 0 \Rightarrow \omega_{\text{MHD}} = 0 \Rightarrow \Omega_{\text{MHD}} = i\gamma_{\text{MHD}} \quad (81)$$

with ω_{MHD} and γ_{MHD} the real frequency and growth rate respectively.

Considering only unstable modes (i.e. Equation (81)) and combining this with Equation (77) results in the following diamagnetic correction model:

$$\gamma_0 = \sqrt{\gamma_{0,\text{MHD}}^2 - \frac{\omega_*^2}{4}} \quad (82)$$

with the subscript 0 indicating that this will be applied to the local MHD growth rate obtained in Subsection 5.3.1. This model is consistent with the low-frequency correction noted in [74, §4].

Equation (82) shows that γ_0 will be reduced compared to $\gamma_{0,\text{MHD}}$, as required. Furthermore, since the reduction depends on ω_* , which is proportional to n , the correction will increase with n , as required. Finally, the model is strictly only valid for $\gamma_{0,\text{MHD}}^2 \geq \frac{\omega_*^2}{4}$. Therefore, in this study, γ_0 is set to zero when this condition is violated, i.e. the mode is completely stabilised when ω_* is sufficiently high, e.g. at high n . This is consistent with the plot of the [74] model shown in [17, Fig. 9(a)].

Equation (79) shows that there is a problem with the model of Equation (82), which is that ω_* varies across flux surfaces (since T_j , n_j , B and $\frac{\partial \rho}{\partial \psi}$ vary across flux surfaces) and ω_* varies along each flux surface (since B varies along each flux surface). This variation is illustrated in Figure 28.

The variation of ω_* across and along flux surfaces shown in Figure 28 raises the question of which value of ω_* to use. This was investigated, and it was found that the best fit of Equation (82) to the data was obtained using the

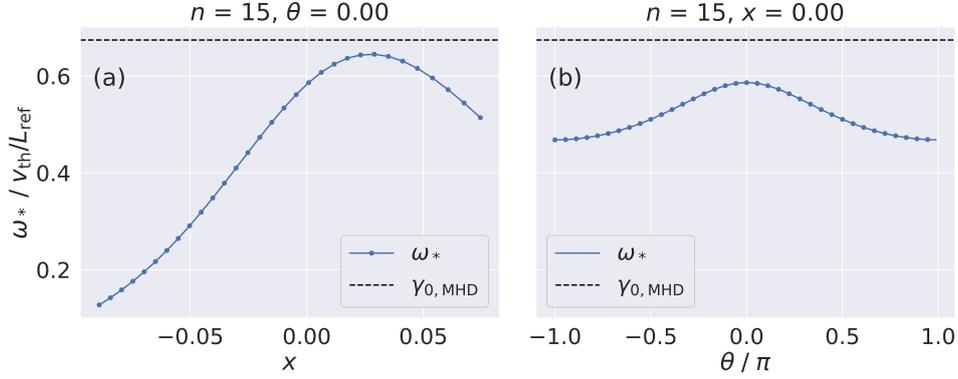


Figure 28: The variation of ω_* at $n = 15$ (a) across flux surfaces (at $\theta = 0$) and (b) along a flux surface (at $x = 0$). The effective local MHD growth rate (black dashed line) is shown for comparison. Other n show similar results.

maximum ω_* found within the global domain along with the introduction of a fit parameter, b , to fine tune the amount of diamagnetic correction:

$$\gamma_0 = \sqrt{\gamma_{0,\text{MHD}}^2 - b \frac{\omega_{*max}^2}{4}} \quad (83)$$

Equation (83) was fitted to the local GS2 data as a function of n_{eff} . The results of this fit is shown in Figure 29.

Figure 29 shows a near perfect fit to the data excluding the final three data points. This cutoff corresponds to $k_y \rho_i \sim 0.5$. The reason for this was investigated using an effective diamagnetic frequency, ω_{*eff} , i.e. the ω_* that would be required to produce the observed correction at each n . It was found that ω_{*eff} follows an approximately bi-linear variation with n , i.e. varies linearly with n for $k_y \rho_i \lesssim 0.5$ and then linearly with n with a different gradient for $k_y \rho_i \gtrsim 0.5$. The fit could be improved to remain near perfect at low n and to become reasonable at high n by using a more complicated bi-linear model for ω_* . However, since we are primarily interested in the peak global growth rate, which occurs within the region of near perfect fit with

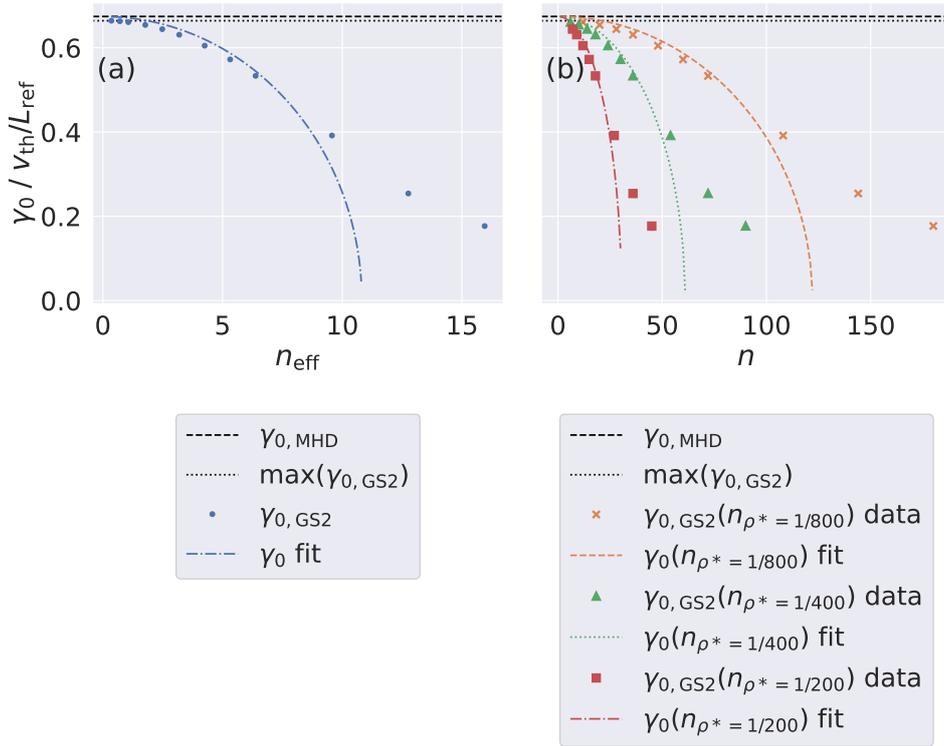


Figure 29: Fits of the diamagnetic correction model (Equation (83)) to the local data from Figure 25 using $b = 8.52$ plotted against (a) n_{eff} and (b) the n corresponding to the ρ^* values tested

the model of Equation (83), the existing model was retained for simplicity and the final three data points were excluded from the fitting process.

Notice that Figure 29 shows the $\gamma_{0,\text{MHD}}$ line at $0.674 v_{\text{th}}/L_{\text{ref}}$ as determined in Subsection 5.3.1 (black dashed line) and the maximum GS2 local growth rate line, $\max(\gamma_{0,\text{GS2}})$, at $0.664 v_{\text{th}}/L_{\text{ref}}$ (black dotted line). The fits obtained using $\gamma_{0,\text{MHD}}$ were excellent, but the fits using $\max(\gamma_{0,\text{GS2}})$ were even better. The reduction from $\gamma_{0,\text{MHD}}$ to $\max(\gamma_{0,\text{GS2}})$ was less than 2%, which is similar to both the level of uncertainty in the global correction model and the difference between the CHEASE and Miller equilibrium models. Therefore, it is reasonable to use $\max(\gamma_{0,\text{GS2}})$ to generate the fits, which was done for those shown in Figure 29.

Figure 29 (b) shows that the magnitude of the diamagnetic correction, $\Delta\gamma_{\text{diamag}}$, decreases with ρ^* , as expected. As with the global correction, the reason for this is evident from examination of the model equation (in this case, Equation (83)). Again, applying the binomial expansion, we find this time that $\Delta\gamma_{\text{diamag}} \sim \omega_{*,\text{max}}^2$. Given the relationship between ω_* and k_y (Equation (78)), this suggests that $\Delta\gamma_{\text{diamag}} \propto k_y^2$, and, hence, $\Delta\gamma_{\text{diamag}} \propto (n\rho^*)^2$. Thus, for a given n , we find that $\Delta\gamma_{\text{diamag}} \propto \rho^{*2}$, as observed. Therefore, in the MHD limit of $\rho^* \rightarrow 0$, we find that $\Delta\gamma_{\text{diamag}} \rightarrow 0$ even as $n \rightarrow \infty$, consistent with the local MHD and global MHD results.

5.3.3 Combined Global Diamagnetic Correction Model

Having developed and tested separate models for the global and diamagnetic corrections, they may now be tested in combination. This was done by substitution of the diamagnetic correction model (Equation (83)) into the global correction model (Equation (76)) to obtain:

$$\gamma = \sqrt{\gamma_{0,\text{MHD}}^2 - b \frac{\omega_{*,\text{max}}^2}{4} \left(1 - \frac{1}{nq'}\right)^a} \quad (84)$$

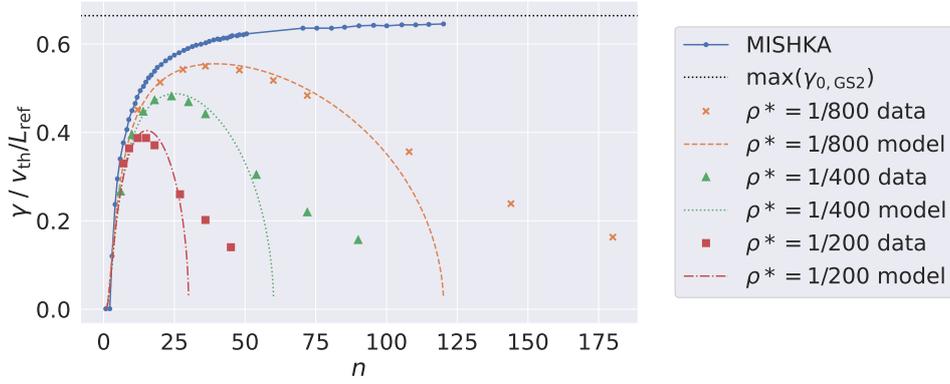


Figure 30: Global growth rates obtained from the combined global diamagnetic correction model (Equation (84)) alongside the GS2 local-global growth rates for comparison.

Global growth rates were calculated using the model of Equation (84) with the fit parameters for a of 18.9, 21.0 and 20.7 for the ρ^* of 1/800, 1/400 and 1/200 respectively, $b = 8.52$, and $\gamma_{0,\text{MHD}} = \max(\gamma_{0,\text{GS2}}) = 0.664 v_{\text{th}}/L_{\text{ref}}$, as obtained in Subsections 5.3.1 and 5.3.2. These values are not necessarily universal; this point is discussed at the end of this subsection. The results of these calculations are shown in Figure 30.

Figure 30 shows that the model of Equation (84) is an accurate approximation to the local-global data, especially in the region of the peak global growth rate. The model does not predict the last three data points well. As discussed in Subsection 5.3.2, the fit to these data points can be improved through the use of a bi-linear diamagnetic frequency model, but the model of Equation (84) was retained for simplicity since we are mainly interested in the peak global growth rate. Indeed, the model calculates the peak global growth rate to within 1%, 2% and 5% of the local-global value for the ρ^* of 1/800, 1/400 and 1/200 respectively.

The values for a and b have been obtained using only this case, so it is unknown whether these values will apply to other cases. Therefore, it is

recommended that this model is tested for a range of equilibria in future work. It is expected that expanding the analysis to include other cases will lead to at least a refinement of these values to maintain a good fit to all cases. However, it may not be possible to find a single set of values that provide a good fit to all cases, so the parameters may need to be varied from case to case. If that turns out to be true, then a model for determining the values of a and b based on the physical parameters of a given case would be needed to avoid carrying out a lengthy and computationally expensive analysis as above each time. The $\gamma_{0,\text{MHD}}$ parameter is expected to vary across cases and would need to be determined each time, but this can be done for relatively low analyst effort and computational cost using MHD stability codes.

5.3.4 Conclusions – Simple Corrective Models

The evidence presented in Subsections 5.3.1 to 5.3.3 again highlight that low n stabilisation is due to global effects, while high n stabilisation is due to gyrokinetic effects. Furthermore, Subsection 5.3.2 shows explicitly that the high n stabilisation is accounted for by diamagnetic drift stabilisation, as expected. Finally, Subsection 5.3.3 demonstrates that the combined global diamagnetic model of Equation (84) produces accurate estimates of the global growth rates across a wide range of n , including the n at which the global growth rate peaks, and this can be calculated using only the local MHD ($n = \infty$) growth rate, $\gamma_{0,\text{MHD}}$, equilibrium parameters ($\omega_{*,\text{max}}$, q' , etc.), and the model parameters a and b . In this work, $\gamma_{0,\text{MHD}}$ has been obtained from reverse application of the global model to global MHD data (which varies with n) and was later taken as just the maximum local GS2 growth rate, $\max(\gamma_{0,\text{GS2}})$ (i.e. the maximum value from a function of n). However, $\gamma_{0,\text{MHD}}$ could conceivably be calculated directly by a local MHD stability code (i.e. at $n = \infty$ only), which would be even more computationally efficient than global MHD stability codes such as MISHKA (which calculate

stability at various finite n). In addition, the model parameters a and b have been obtained here by fitting the data, and these are likely to depend on the equilibrium. Future work may wish to investigate whether these fit parameters can be derived from first principles, or otherwise generalised from a wider range of equilibria.

The results in Subsection 5.3.3 show that the inclusion of global effects and diamagnetic drift stabilisation via a simple (i.e. computationally inexpensive) corrective model can predict the peak global growth rate to within 5% or less of the local-global growth rate depending on the value of ρ^* . Meanwhile, these peak global growth rates are significantly reduced compared to $\gamma_{0,\text{MHD}}$ by 18%, 28% and 40% for the ρ^* of 1/800, 1/400 and 1/200 respectively. It is not yet clear how this reduction in growth rate compared to $\gamma_{0,\text{MHD}}$ will affect the corresponding critical pressure gradient and, thus, the EPED prediction of pedestal height and width at which an ELM will occur. However, the global diamagnetic correction model of Equation (84) allows this to be investigated, which is done in Section 5.4.

5.4 Global and Diamagnetic Correction Models in EPED

5.4.1 Method Overview

Having developed the simple corrective model of Equation (84), the impact of this on EPED-like calculations may now be assessed. This is done by determining EPED-like KBM constraints with and without the corrections of Equation (84), combining these with EPED-like PB constraints, and thus determining the consequent impact on the final EPED-like prediction of pedestal height and width at which an ELM would be triggered.

The EPED [14] KBM constraint represents the critical pedestal height for the onset of KBMs as a function of pedestal width, which is determined

using the “ballooning critical pedestal” (BCP) technique. This is done by defining the pedestal profiles using modified tanh functions that include the pedestal height and width as parameters, thus allowing modification of these parameters to generate new equilibria. The modified tanh functions also include the pedestal location as a parameter, which, combined with the width parameter, defines the “EPED pedestal” region. Local IBM stability (as a proxy for KBMs) is determined for each flux surface within the EPED pedestal for each pedestal height and width tested. The critical height at a given width is then taken as the value at which the central 50% of the EPED pedestal is unstable to IBMs.

The method used here is built around a modified BCP technique as follows:

1. A local MHD stability code, `ideal_ball` (an auxiliary program within the GS2 code-base, as used in Chapter 4), is used to obtain the critical pressure gradient length scale, β'_{crit} , for the onset of local IBMs. This method is described in Subsection 5.4.2, including details of how to deal with second stability. As with the BCP technique, the critical value is that at which the central 50% of the EPED pedestal is unstable to IBMs.
2. As an extra step to account for global and diamagnetic effects, Equation (84) is used to calculate a change in growth rate, $\Delta\gamma$, and local gyrokinetic simulation data is used to determine the stiffness of KBM onset, which is represented by the gradient $\frac{d\gamma_0}{d\beta'}$. These are combined to obtain the corresponding change to β'_{crit} that accounts for global and diamagnetic effects.
3. The corrected and uncorrected values of β'_{crit} are used to calculate the corresponding EPED KBM constraints. As a modification to the BCP technique (for convenience), only one equilibrium is tested to determine β'_{crit} and this is used to calculate the critical pedestal height for

various pedestal widths, rather than measuring the critical pedestal height directly from various equilibria with different widths. This is a valid approach since this is a qualitative study of the impact of global and diamagnetic corrections, which are expected to exceed any corrections from varying the equilibrium consistently.

4. Finally, we go beyond the BCP technique to the other aspect of EPED-like calculations, the PB constraint. An approximate form is used to calculate a range of representative PB constraints. These are compared with the corrected and uncorrected KBM constraints, thus revealing the potential impact of global and diamagnetic effects of the KBM on EPED-like ELM predictions.

More detail about the method of each of these steps is given in the corresponding subsections that follow, along with results of the application of the method to the equilibrium used in this chapter.

5.4.2 Local MHD Stability

As noted in Chapter 4, the GS2 auxiliary program `ideal_ball` calculates local MHD stability for a given equilibrium and returns a binary stable / unstable result. Therefore, `ideal_ball` can be used to find the critical pressure gradient length scale for the onset of local MHD ballooning modes, $\beta'_{\text{crit,MHD}}$. This was done by running `ideal_ball` for each surface in the equilibrium file within the global domain for a range of β' values that were known from the GS2 runs to cover the regions of first and second stability, and peak instability. This runs in approximately 1 hour 40 minutes on 30 processors using a somewhat inefficient parallelisation algorithm for 662 β' values across 38 surfaces, resulting in 25,156 (x, β') points in total, i.e. very high resolution. The high resolution ensured a well-defined first and second stability boundary, even on surfaces where these were close together

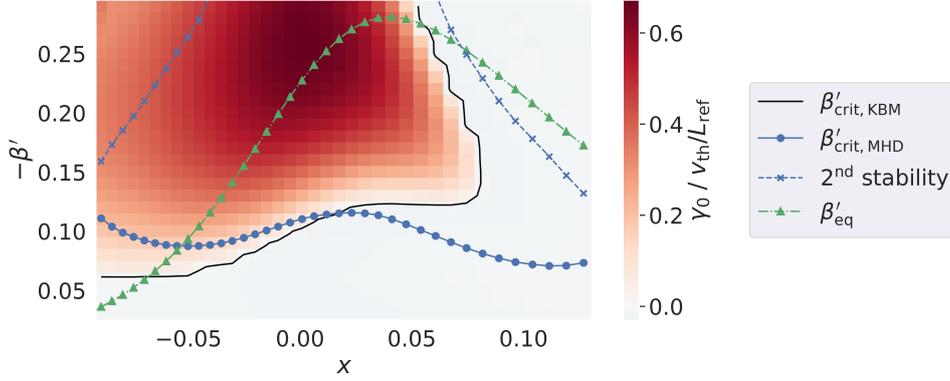


Figure 31: Comparison of β'_{crit} for KBMs and IBMs (from GS2 and `ideal_ball` respectively) overlaid on KBM growth rates. Also shown are the IBM second stability boundary, and the equilibrium β' .

in β' . The run-time was deemed fast enough for the current investigation but would be too slow for a real EPED-like application. However, tests showed that this may be reduced to less than 5 minutes on 1 processor by using a more efficient search algorithm (e.g. bisection) or even further with improved parallel utilisation if required in the future. The bisection method tested was not used to generate the results reported in this thesis due to issues on some surfaces in finding the stability boundaries. This occurred because the starting points for the bisection were arbitrary so it was possible to miss the unstable region if the first and second stability boundaries were close together in β' . Further work would be needed to improve the bisection method if it were to be adopted in the future.

The results of $\beta'_{\text{crit,MHD}}$ as a function of x as obtained from `ideal_ball` are shown in Figure 31. Alongside this are the KBM growth rates as measured with GS2 by scanning β' , and the contour of $\beta'_{\text{crit,KBM}}$ where the GS2 KBM growth rate is zero (found by linear interpolation). The GS2 results were generated using $(n, \rho^*) = (1, 1/70)$ to minimise diamagnetic drift stabilisation and thus be as close as possible to the MHD case.

Figure 31 shows that $\beta'_{\text{crit,MHD}}$ varies with x , but only by a relatively small amount compared to β'_{eq} . Furthermore, β'_{eq} is greater than $\beta'_{\text{crit,MHD}}$ for the majority of the pedestal region (which covers most of the x range shown, depending on how the pedestal region is defined – see Figure 32). This is consistent with the finding that KBMs are unstable over most of this region. Finally, the $\beta'_{\text{crit,MHD}}$ values obtained from `ideal_ball` are similar to the $\beta'_{\text{crit,KBM}}$ values from GS2. Interestingly, there is a region at high positive x where KBMs are stable for all β' tested while IBMs are not. This was not investigated due to time constraints, but does offer an example of where using KBMs as a proxy for IBMs may not be sufficiently accurate.

Figure 32 shows the same $\beta'_{\text{crit,MHD}}$ data alongside the local-global method and EPED pedestal domains. This highlights the differing definitions of these regions. The short vertical dotted lines indicate the central 50% of these domains that could be used to calculate a single value for β'_{crit} using a BCP-like technique. The EPED definitions of the pedestal region and the requirement for the central 50% to be unstable are both somewhat arbitrary. Therefore, sensitivity to this was tested by requiring the middle 50% of the local-global domain to be unstable or *any* 50% of the local-global domain to be unstable. However, these alternative (also arbitrary) definitions made, respectively, no difference or negligible (< 1%) difference. Therefore, for simplicity and consistency with EPED, the original BCP requirement for the central 50% of the EPED pedestal to be unstable was used. This results in a value of $\beta'_{\text{crit,BCP}} = 0.116$, as indicated in Figure 32.

5.4.3 β'_{crit} Correction

The correction to β'_{crit} is calculated in three steps:

1. Determine the change in growth rate, $\Delta\gamma$, due to Equation (84) at the n at which Equation (84) peaks for each ρ^*

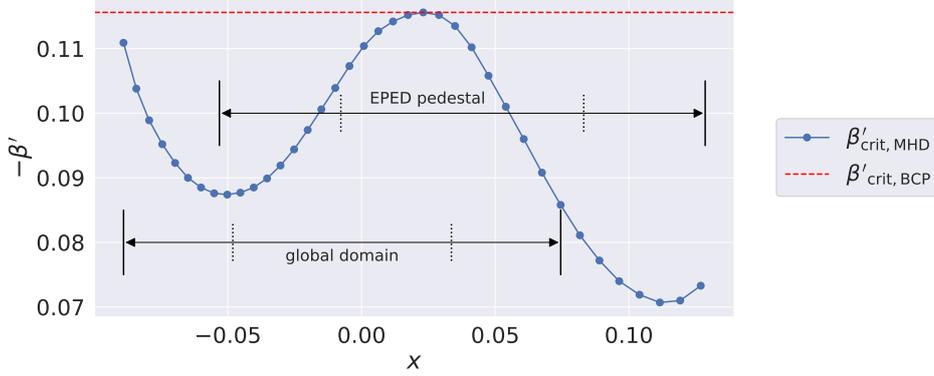


Figure 32: Uncorrected β'_{crit} (dashed red line) resulting from the BCP technique, defined as the minimum β' required for the central 50% of the EPED pedestal region (indicated) to be unstable, i.e. the maximum $\beta'_{\text{crit,MHD}}$ in that region (solid blue line, circle markers, same data as corresponding line in Figure 31). Also shown is the local-global method simulation domain for comparison.

2. Measure the stiffness of KBM onset to get the gradient $\frac{d\gamma_0}{d\beta'}$
3. Calculate the change to β'_{crit} as $\Delta\gamma / \frac{d\gamma_0}{d\beta'}$

These steps are described in more detail alongside the corresponding results in the following subsections.

5.4.3.1 Determining the Change in Growth Rate

The n at which Equation (84) peaks is the most unstable toroidal mode for the corresponding ρ^* . Since this mode will go unstable first as β' increases, it is this mode that will lead to increased turbulent transport, thus limiting the pedestal pressure gradient. Therefore, the change in growth rate, $\Delta\gamma$, due to the global and diamagnetic effects of the KBM is taken as the difference between the local MHD growth rate and the peak corrected growth rate produced by Equation (84).

ρ^*	$\Delta\gamma / v_{\text{th}}/L_{\text{ref}}$
1 / 800	0.11
1 / 400	0.18
1 / 200	0.26

Table 2: Change in growth rate, $\Delta\gamma$, for each ρ^*

It is possible to find $\Delta\gamma$ analytically by substituting Equation (79) into Equation (84), and finding the roots of the derivative with respect to n . However, the resulting analytical form is rather long and complicated and this accuracy is not needed anyway since n must be an integer. Therefore, it is clearer to simply evaluate Equation (84) at each integer n from 1 upward until the square root factor becomes imaginary, and then select the maximum value. This was done, resulting in the $\Delta\gamma$ values shown in Table 2, which are consistent with Figure 30.

5.4.3.2 Measuring the Stiffness of KBM Onset

Now that we have $\Delta\gamma$, the corresponding change to β'_{crit} , denoted $\Delta\beta'_{\text{crit}}$, may be calculated using $\frac{d\gamma_0}{d\beta'}$.

Figure 32 shows that, in this case, we have $\beta'_{\text{crit}} \sim 0.1$. Since we are calculating the correction to β'_{crit} , we therefore measure $\frac{d\gamma_0}{d\beta'}$ in this region. First, β' was scanned using GS2 for each surface in the equilibrium file within the global domain to obtain $\gamma_0(x, \beta')$ (the heat map in Figure 31). Next, the gradient between each pair of neighbouring data points in the β' direction was calculated using a first order finite difference. Finally, the selected value was the first gradient in the stiff region. An example of the results (at $x = 0$) are shown in Figure 33 (a) and (b) alongside $\frac{d\gamma_0}{d\beta'}$ as a function of x in Figure 33 (c). As in Subsection 5.4.2, the GS2 runs use $(n, \rho^*) = (1, 1/70)$ to minimise diamagnetic drift stabilisation and thus be as close as possible to

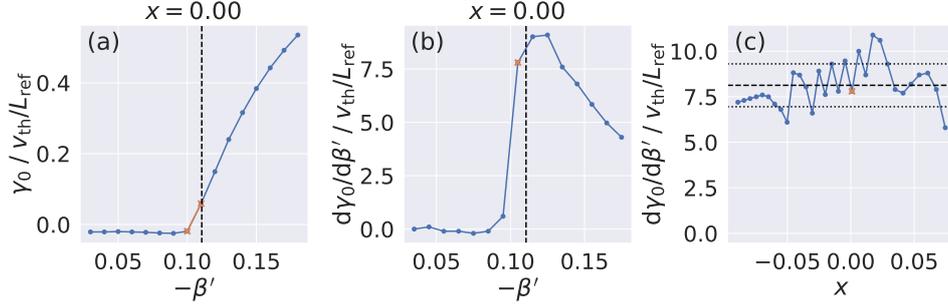


Figure 33: (a) γ_0 and (b) $\frac{d\gamma_0}{d\beta'}$, as functions of β' at $x = 0$ for $(n, \rho^*) = (1, 1/70)$ with β'_{crit} (black dashed line) and the first gradient in the stiff region (orange line / x marker) indicated. (c) $\frac{d\gamma_0}{d\beta'}$ as a function of x with the $x = 0$ value (orange x marker) from (b), mean (black dashed line), and standard deviation (black dotted lines) indicated.

the MHD case; other n show similar results, though less clearly due to the less well resolved KBMs as a result of the lower growth rates.

Figure 33 (a) shows that KBMs exhibit a stiff onset in this case, as expected. Furthermore, Figure 33 (a) and (b) show that the gradient $\frac{d\gamma_0}{d\beta'}$ is relatively constant in the region of β'_{crit} , suggesting that a measurement of $\frac{d\gamma_0}{d\beta'}$ is appropriate to calculate $\Delta\beta'_{\text{crit}}$ from $\Delta\gamma$. Finally, Figure 33 (c) shows that $\frac{d\gamma_0}{d\beta'}$ is relatively constant across different surfaces, with a mean value of $8.1 v_{\text{th}}/L_{\text{ref}}$ and a standard deviation of $1.2 v_{\text{th}}/L_{\text{ref}}$ (14.5%). This indicates that it may be appropriate to use the radially averaged value for simplicity.

5.4.3.3 Calculating the Change to β'_{crit}

So far, we have obtained a single value for the uncorrected β'_{crit} (0.116, Subsection 5.4.2), a value of $\Delta\gamma$ for each ρ^* (Table 2, Subsection 5.4.3.1) and a surface averaged $\frac{d\gamma_0}{d\beta'}$ ($8.1 v_{\text{th}}/L_{\text{ref}}$, Subsection 5.4.3.2). Therefore, a β'_{crit} value that is corrected for global and diamagnetic effects, denoted $\beta'_{\text{crit,corr}}$ may now be calculated as simply:

ρ^*	β'_{crit}	% change
Uncorrected	0.116	n/a
1/800	0.129	12%
1/400	0.137	19%
1/200	0.147	27%

Table 3: Critical pressure gradients for the onset of MHD ballooning modes corrected for global and diamagnetic effects for various values of ρ^*

$$\beta'_{\text{crit,corr}} = \beta'_{\text{crit}} + \Delta\beta'_{\text{crit}} \quad (85)$$

where:

$$\Delta\beta'_{\text{crit}} = \frac{\Delta\gamma}{d\gamma_0/d\beta'} \quad (86)$$

This results in the corrected β'_{crit} values shown in Table 3.

More advanced methods were tested for accuracy such as: (a) adjusting the KBM $\gamma_0(x)$ by $\Delta\gamma$ and applying the BCP technique to the corrected $\gamma_0(x)$ to account for the variation of $\frac{d\gamma_0}{d\beta'}$ with β' ; and (b) applying the per surface $\frac{d\gamma_0}{d\beta'}$ (Figure 33 (c)) to the per surface β'_{crit} (Figure 31, blue solid line, circle markers) before applying the BCP technique to account for the variation of $\frac{d\gamma_0}{d\beta'}$ with x . However, these methods made negligible difference ($\lesssim 1\%$ and $\lesssim 4\%$ respectively). Therefore, we may conclude that the global and diamagnetic effects are dominant over the effect of the variability of $\frac{d\gamma_0}{d\beta'}$ with x and β' . Hence, the original version of the method (i.e. correcting the single value obtained from the BCP technique using a surface averaged $\frac{d\gamma_0}{d\beta'}$ as described above) was used for simplicity.

The use of $\beta'_{\text{crit,MHD}}$ as a proxy for $\beta'_{\text{crit,KBM}}$ is quoted [14] as being “accurate (within 10%) . . . , due to the partial offset of destabilising . . . and stabilis-

ing [kinetic] effects". However, the percentage changes presented in Table 3 appear to contradict that finding. This suggests that global effects, which are not included in the justification cited above, may produce differences between the values of $\beta'_{\text{crit,MHD}}$ and $\beta'_{\text{crit,KBM}}$ beyond the 10% expected in EPED-like models. However, the important result of an EPED-like calculation is the pedestal height and width at which an ELM is triggered. It is unclear from the results presented so far how any inaccuracy in the value of β'_{crit} used in the model as a result of the above proxy will carry through to the pedestal height and width prediction. Therefore, this is investigated in Subsection 5.4.4.

5.4.4 The Effect of the β'_{crit} Correction on the EPED ELM Prediction

The corrections described so far are to β'_{crit} . However, the primary output of EPED-like models is a prediction of pedestal height and width at which an ELM will be triggered. The EPED prediction depends on both the KBM and PB constraints, the former of which depends on the β'_{crit} values corrected in Subsection 5.4.3. Therefore, it is unclear from the results presented so far how much these changes to β'_{crit} will affect the EPED ELM prediction, or even whether they will lead to an increase or a decrease in the pedestal height and width predicted by EPED-like models. Hence, the corresponding KBM constraints and some representative PB constraints are calculated in this subsection to assess the impact of the β'_{crit} correction upon the EPED ELM prediction.

5.4.4.1 EPED KBM Constraint

Snyder *et al.* [75] define tanh based functions to capture the general shape of pedestal temperature and density profiles in terms of a number of parameters such as pedestal height, width and location. For the current work, the

following simplified form is sufficient to describe the profile shapes used in this work with an approximately flat profile outside of the pedestal region:

$$T(x) = T_{\text{sep}} + \left(\frac{T_{\text{ped}}}{2}\right) \left\{ \tanh \left[2 \frac{(1 - x_{\text{mid}})}{\Delta} \right] - \tanh \left[2 \frac{(x - x_{\text{mid}})}{\Delta} \right] \right\} \quad (87)$$

where T_{sep} , T_{ped} , Δ and x_{mid} are the fit parameters representing, respectively, separatrix temperature, pedestal height (temperature on axis minus T_{sep}), pedestal width and location of the pedestal centre. Fitting Equation (87) to the equilibrium temperature profile results in values of $T_{\text{sep}} = 391$ eV and $T_{\text{ped}} = 758$ eV in approximate agreement with the expected values of $T_{\text{sep}} = 387$ eV and $T_{\text{ped}} = 757$ eV [17]. The corresponding location and width parameters are $x_{\text{mid}} = 0.038$ and $\Delta = 0.182$, resulting in the EPED pedestal region indicated in Figure 32.

Holding x_{mid} and T_{sep} constant, we may now use Equation (87) to generate profiles of arbitrary height and width. The EPED KBM constraint is the critical pedestal height for the onset of KBMs (using IBMs as a proxy) as a function of pedestal width. We determine the critical height as the lowest for which *any* 50% of the EPED pedestal has β' higher than β'_{crit} . The use of any 50%, rather than the middle 50% has already been shown to have negligible impact (Subsection 5.4.2) and was chosen here for convenience.

The critical height may be calculated using the uncorrected β'_{crit} for a range of pedestal widths to map out the uncorrected EPED KBM constraint. This may then be repeated for the various corrected values of $\beta'_{\text{crit,corr}}$ to get the EPED KBM constraints corrected for global and diamagnetic effects. This was done for a typical range of widths used in EPED calculations and the results are shown in Figure 34.

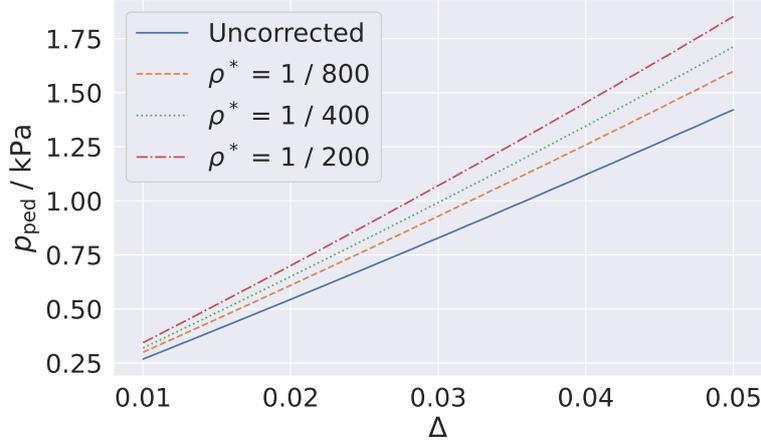


Figure 34: EPED KBM constraints showing critical pressure pedestal height, p_{ped} , as a function of pedestal width, Δ , calculated using the uncorrected β'_{crit} and the corrected $\beta'_{\text{crit,corr}}$ for each ρ^*

5.4.4.2 EPED PB Constraint

Fully self-consistent calculation of the PB constraint for the equilibrium studied in this chapter is beyond the scope of this thesis²⁰. This is because that would involve using a Grad-Shafranov code to generate many alternative versions of the equilibrium with pedestals of different heights and widths and running these through an MHD stability code to determine the PB stability. However, Snyder *et al.* [14] note that the functional dependence of the PB constraint is $p_{\text{ped}} \sim \Delta^{3/4}$. Therefore, to illustrate the potential impact of the KBM constraint correction, we may generate a range of representative PB constraints using the following model:

$$p_{\text{ped}} = c\Delta^{3/4} + d \quad (88)$$

²⁰While this is not necessarily any more difficult than than performing the MISHKA calculations used earlier in this chapter, the MISHKA simulations were already done as part of previous work [17] and the calculations using that data reported here were necessary to develop the simple corrective models.

Values for the coefficients c and d were chosen to produce three limiting cases:

1. A PB constraint with a positive gradient, resulting in a reduction as ρ^* is increased of both the pedestal height and width at which an ELM is predicted to occur.
2. A PB constraint with zero gradient, resulting in a reduction as ρ^* is increased of the pedestal width at which an ELM is predicted to occur but no effect on the pedestal height.
3. A PB constraint with a negative gradient, resulting in a reduction of the pedestal width prediction but an increase in the pedestal height as ρ^* is increased.

These three cases are shown alongside the KBM constraints in Figure 35 (a) to (c). This shows that (in this case and with the representative PB constraints developed), the correction of the KBM constraint to account for global and diamagnetic effects can potentially have a significant impact on the EPED ELM prediction. In all cases, the uncorrected KBM constraint over-predicts the pedestal width at which an ELM will occur relative to the corrected constraints by 0 to 110%. Furthermore, whether the correction results in a significant increase, decrease or a negligible change to the ELM-triggering pedestal height (ranging from -19% to +69%) depends crucially on the details of the PB constraint. The relative changes to the height and width predictions as a consequence of accounting for global and diamagnetic effects are summarised in Table 4.

Comparing the representative PB constraints developed here to previous self-consistently calculated PB constraints (e.g. [14, 76]) suggests that a PB constraint with a positive gradient (Figure 35 (a)) is the most likely. In this case, we have the counter-intuitive result that an ELM is triggered at a *lower* pedestal height even though the equilibrium is more stable to KBMs

than calculations without global and diamagnetic effects suggest. This is because the raising of the KBM stability boundary allows the pedestal to grow steeper during the KBM constrained phase such that the PB constraint is reached at a lower pedestal height. This may cause smaller ELMs to be triggered more frequently, but may also limit the achievable pedestal height. Previous EPED predictions (e.g. [14, 76]) were examined to see if there was a systematic over-prediction of the pedestal height and width to corroborate this finding, but this investigation was inconclusive due to the large uncertainty on the experimental measurements.

5.4.5 Discussion – Corrective Models in EPED

The method of including global and kinetic effects in EPED-like models presented in this chapter introduces three fit parameters: the coefficient in the $1/n$ global correction model; the coefficient in the diamagnetic correction model; and the gradient of growth rate with respect to β' to measure the stiffness of KBM onset and thus calculate $\Delta\beta'_{\text{crit}}$. This is undesirable since the EPED philosophy is to calculate the prediction from first principles without fit parameters. Furthermore, the method developed here introduces a requirement to run more detailed simulations to calculate the above fit parameters. This is also undesirable since computational efficiency is required for EPED-like calculations to rapidly assess a range of equilibria.

To address these issues, future work should investigate the variability of the fit parameters for different equilibria. The aim of this would be to develop a first principles model of the global and kinetic corrections without fit parameters, thus removing the requirement to run more detailed simulations.

Nonetheless, the results presented in this section show that global and kinetic effects can potentially have a significant impact on EPED-like ELM

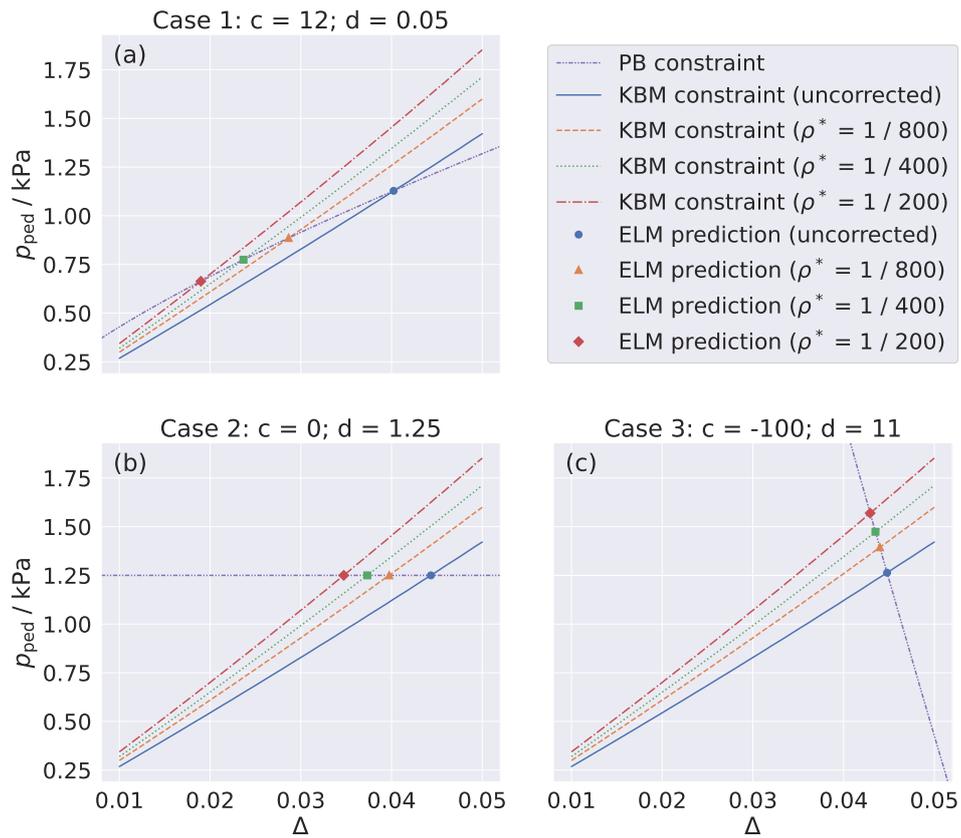


Figure 35: (a) - (c) Three representative PB constraints generated using Equation (88) and the coefficient values shown, alongside the EPED KBM constraints from Figure 34, resulting in the EPED ELM predictions indicated by the markers.

Case 1		
ρ^*	$p_{\text{ped}} / \text{kPa}$	Δ
Uncorrected	1.13 (0%)	0.0403 (0%)
1 / 800	0.89 (+27%)	0.0287 (+40%)
1 / 400	0.78 (+46%)	0.0237 (+70%)
1 / 200	0.67 (+69%)	0.0192 (+110%)
Case 2		
ρ^*	$p_{\text{ped}} / \text{kPa}$	Δ
Uncorrected	1.25 (0%)	0.0443 (0%)
1 / 800	1.25 (0%)	0.0398 (+11%)
1 / 400	1.25 (0%)	0.0374 (+19%)
1 / 200	1.25 (0%)	0.0348 (+27%)
Case 3		
ρ^*	$p_{\text{ped}} / \text{kPa}$	Δ
Uncorrected	1.26 (0%)	0.0448 (0%)
1 / 800	1.39 (-9%)	0.0440 (+2%)
1 / 400	1.47 (-14%)	0.0435 (+3%)
1 / 200	1.56 (-19%)	0.0430 (+4%)

Table 4: Pedestal height and width of the corrected and uncorrected EPED ELM prediction for each ρ^* . The value in brackets is the percentage error of using the uncorrected value rather than the corrected value.

predictions, depending on the PB constraint, so should be considered in future EPED-like studies.

5.4.6 Conclusions – Corrective Models in EPED

This section has developed a method of including global and kinetic effects on the KBM in an EPED-like calculation. This is based on EPED’s “ballooning critical pedestal” (BCP) technique, with additional steps to correct the local IBM β'_{crit} for global and kinetic effects. The consequent impact on the EPED KBM constraint has been calculated and compared against a range of representative PB constraints. This shows that including global and kinetic corrections in EPED-like calculations can potentially have a significant impact on EPED-like ELM predictions, depending on the PB constraint. Therefore, the impact of such effects should be considered in future EPED-like studies.

5.5 Discussion – General

5.5.1 Boundary Issues in Global Gyrokinetic Codes

Global gyrokinetic codes simulate many flux surfaces and thus require boundary conditions on the surfaces at the edge of the simulation domain. It is expected that these boundary conditions will have some impact on the behaviour for regions within the simulation domain near the boundary [49]. Therefore, previous work [17] developed the equilibrium studied in this chapter with a pseudo-pedestal at mid-radius, rather than at the edge as in real tokamaks, to avoid any influence from the boundary conditions on the region of interest.

Local gyrokinetic codes, meanwhile, simulate within a flux tube simulation domain centred around a single magnetic field line. The radial size of the simulation domain is small compared to the equilibrium variation scale

length but large compared to the decorrelation scale length of the turbulence simulated. This allows the use of periodic radial boundary conditions, which are not expected to adversely affect the simulation results. Therefore, local gyrokinetics and, thus, local-global gyrokinetics are expected to be relatively unaffected by the boundary issues faced by global gyrokinetic codes. Furthermore, local-global gyrokinetics can accurately capture the global behaviour of gyrokinetic instabilities (at least in this case), as shown in this thesis. Therefore, local-global gyrokinetics may provide a method for studying global gyrokinetic modes at the edge of magnetically confined plasmas, including those in tokamak pedestals, without the complications of simulation boundary effects (such as those present in, for example, [77, 78], where their effect was, respectively, not mentioned, or discussed briefly but otherwise dismissed).

5.5.2 Validity of the Local Approximation

Chapter 4 concluded that it is not possible to determine *a priori* whether the local approximation is valid for a given case; rather, one should assume validity, run the simulations, and confirm validity afterwards. There are two conditions to check: the gyrokinetic assumption, $\rho^* \ll L$; and the local assumption, $\Delta \ll L$. The question is, “how much less is *much less*?”

In this case, the local-global results show excellent agreement with global gyrokinetics and MHD within the expected range. Therefore, we can conclude that the gyrokinetic and local approximations are valid, and we can compare the above parameters (ρ^* , Δ and L) to begin answering the question of how much less is “much less” for the benefit of future studies that do not have global gyrokinetics and MHD to compare against.

This work shows that the local-global method breaks down at low toroidal mode numbers, n , as expected. In this case, this occurs at surprisingly low

n , with excellent agreement down to $3 \lesssim n \lesssim 12$ depending on ρ^* . We also have $q' \sim 12 \pm 1$ (mean \pm standard error of the mean across the local-global domain). Therefore, we can calculate the limiting rational surface spacing as $\Delta = 1/(nq') = 1/(3 \times 11) \sim 0.03$. This needs to be small compared to the system length scale, L . Taking L as the plasma minor radius (1 in normalised units), results in the condition $\Delta/L \sim 0.03 \ll 1$ being sufficient. If we instead use L as the pedestal width (~ 0.167 in this case), which is arguably a more appropriate equilibrium variation length scale, then Δ/L could be a factor of ~ 6 larger, i.e. $\Delta/L \sim 0.18 \ll 1$ may be sufficient. This starts to provide a quantitative answer to the question of how much less is “much less” for the local approximation to be valid.

But what about the gyrokinetic assumption that requires $\rho^* \ll L$? The equilibrium in this case has $\rho^* \sim 1/70 \sim 0.014$. However, the largest ρ^* tested in [17] was $1/200$ to ensure that the condition $\rho^* \ll L$ was satisfied. Indeed, the calculations presented in this chapter show excellent agreement with global gyrokinetics and MHD for all of the ρ^* values used in [17], *viz.* $1/200$, $1/400$ and $1/800$. However, the local-global calculation produces non-physical results (not shown) when the equilibrium value $\rho^* \sim 1/70$ is used, suggesting that this ρ^* value is too large. Therefore, taking $L = 1$ suggests that $\rho^*/L = 0.005 \ll 1$ is sufficient, while $\rho^*/L \sim 0.014$ is not. If we instead use L as the pedestal width, then $\rho^*/L \sim 0.03 \ll 1$ may be sufficient but $\rho^*/L \sim 0.08$ would not be sufficient. This also starts to provide a quantitative answer to the question of how much less is “much less” for gyrokinetics to be valid.

Interestingly, the analysis above suggests the gyrokinetic condition is more stringent than the local condition by roughly an order of magnitude, hence $\Delta \gg \rho^*$ is allowable. This makes sense as we are ultimately calculating the global gyrokinetic result, so even if the locality condition is not *that* well satisfied, then it doesn't matter too much in the end – we just get a larger

global correction. Indeed, as n increases and so Δ decreases, the locality condition is more strongly satisfied and the global result tends towards the local result, as expected.

5.5.3 Future Work

Important future directions for this research include: (1) developing the local-global method such that it can be used to study experimental pedestal equilibria; and (2) incorporating the simple corrective model of Equation (84) into EPED and / or Europed to improve their predictions of pedestal height and width.

Point (1) would give many benefits including an understanding of the expected degree of poloidal shifting of the global mode structure due to the transition from isolated to general modes, and the consequent improved interpretation of experimental measurements of turbulent fluctuations. This would also support generalisation of the simple corrective model of Equation (84) to a wider range of equilibria in preparation for point (2). Once point (2) is complete, it would then be beneficial to compare the improved pedestal predictions to experimental measurements as validation of the corrective model.

An important next step towards the goal of point (1) is to apply the local-global method to a more realistic case where (a) the pedestal is at the edge and (b) \hat{s} is high. Case (a) should be possible because local gyrokinetics (and, hence, the local-global method) do not suffer from boundary effects in the same way that global gyrokinetics can, as discussed in Subsection 5.5.1. The benefit of this is that it allows us to check effects that exist at the edge, such as a significant trapped fraction or smaller surface aspect ratio (although the tests reported in Subsection 5.2.3.3 and Chapter 4 suggest that these particular examples are not expected to have a significant effect). Case (b)

is useful to confirm that high \hat{s} does not adversely affect the validity of the local approximation due to the mitigation of narrowing in η_0 by increased q' as expected from Chapter 4.

5.6 Conclusions – General

This chapter has presented local-global simulations of KBMs in a pedestal-like equilibrium. The results were compared favourably to global gyrokinetic and MHD simulations, which suggests that the local-global method is valid in this case. This comparison also showed that the global growth rate is sensitive to the parallel magnetic field fluctuation model and the equilibrium treatment, but largely insensitive to the other factors tested such as trapped particles and collisions.

The comparison of the local-global results with global gyrokinetics and MHD suggests that the local-global method is valid in this case down to surprisingly low toroidal mode numbers of $3 \leq n \leq 12$ depending on ρ^* . This in turn suggests that it may be possible, in principle, to incorporate additional physics such as kink drive into gyrokinetics.

Analysis of the local validity of this case suggests that $\Delta/L \sim 0.03 \ll 1$ may be sufficient for local validity taking L as the plasma minor radius, or $\Delta/L \sim 0.18 \ll 1$ taking L as the pedestal width (arguably a more appropriate equilibrium variation length scale). Similarly, $\rho^*/L \sim 0.005 \ll 1$ may be sufficient for gyrokinetic validity with L as the plasma minor radius, while $\rho^*/L \sim 0.014$ may not be sufficient; or $\rho^*/L \sim 0.03 \ll 1$ may be sufficient with L as the pedestal width, while $\rho^*/L \sim 0.08$ would not be sufficient. These values start to provide a quantitative answer to the question of how much less is “much less” in terms of local and gyrokinetic validity.

This chapter also developed a simple corrective model to capture global and diamagnetic effects, and this was shown to give a good approximation to the local-global results in the region of the peak growth rate. This shows that the observed low n stabilisation of global KBMs is caused by global effects while high n stabilisation is caused by kinetic effects (namely diamagnetic drift stabilisation). Further work is required to generalise the model's fit parameters. Despite this limitation, this shows that it should be possible to approximate global and diamagnetic effects using such a model for applications where computational efficiency is paramount.

Finally, the utility of the simple corrective model was demonstrated in an EPED-like calculation. Further work is required to generalise the stiffness of KBM onset used in the method. Nonetheless, this showed that global and diamagnetic corrections can potentially have a significant impact on the ELM prediction of EPED-like models, depending on the PB constraint. Therefore, future EPED-like calculations should consider the impact of global and diamagnetic effects, which can be done using the method presented here.

Part III

Poloidal Structure of Zonal Flow Drive

Chapter 6

Poloidal Structure of Zonal Flow Drive

6.1 Introduction

Zonal flows in tokamaks are radially-sheared flows constant in the poloidal and toroidal directions, and are driven by nonlinear energy transfer from turbulent micro-instabilities. They are central to achievable performance since they regulate turbulence and, thus, govern the L-H transition (e.g. [19–27]) and H-mode transport (e.g. [20, 23, 31]).

To measure zonal flows experimentally is difficult due to the symmetry of the flow and the small poloidal measurement area of the relevant diagnostics (discussed below). However, the poloidal structure of the drive mechanisms of zonal flows is not necessarily symmetric. Therefore, interpretation of zonal flow experiments would be greatly enhanced by knowledge of the poloidal structure of zonal flow drive. For example, a factor ~ 50 discrepancy between the nonlinear energy transfer and turbulent energy loss was observed on Alcator C-Mod and investigated using gas-puff

imaging (GPI) diagnostic data [24]. The GPI study improved the agreement between the measured nonlinear transfer and turbulence loss, but a factor ~ 3.5 discrepancy remained due to the limited poloidal extent of the GPI diagnostic. However, while the zonal flow is constant in the poloidal and toroidal directions, the turbulence is not, hence there is no reason to assume that the nonlinear transfer and turbulence loss are poloidally and toroidally symmetric. A method to quantitatively determine the poloidal structure of zonal flow drive would allow the GPI analysis reported in [24] to be compared against corresponding measurements from simulations, thus facilitating a more in depth study of the observed discrepancy and validation of the experimental nonlinear transfer measurements.

Zonal flow properties relevant for this thesis, including driving and damping mechanisms, are summarised below. A detailed review is given in [18].

Since zonal flows in tokamaks are toroidally and poloidally constant, radially-sheared $\mathbf{E} \times \mathbf{B}$ flows on a magnetic flux surface (i.e. excluding equilibrium shear flow), they constitute the zero-frequency, $n = m = 0$ (or $k_{\parallel} = k_y = 0$) component of the electrostatic potential perturbation with $k_x \neq 0$. As a result of this symmetry, zonal flows cannot directly access the free energy in the equilibrium density and temperature gradients and are, thus, linearly stable. Instead, they are driven by the nonlinear interaction of plasma micro-instabilities (which do extract free energy from the gradients).

The poloidal and toroidal symmetry of zonal flows also makes direct experimental measurements of the flow difficult, since there is no poloidal variation. However, the transfer of energy into zonal flows from turbulence (i.e. zonal flow drive) is not necessarily subject to the same symmetry, and thus can be measured in experiments more easily. This is done using nonlinear energy transfer functions (NETFs), which are discussed in more detail

below. This chapter aims to determine the poloidal structure of zonal flow drive by applying NETFs to nonlinear gyrokinetic zonal flow simulations.

The primary drive mechanisms of zonal flows are parametric decay and modulational instability²¹. Both processes constitute three-wave coupling, whereby the convective nonlinearity, $\mathbf{u} \cdot \nabla f$, of a fluctuating quantity, f , (such as the density or velocity field) transfers energy between modes of differing wavevectors. The convention used here is that the transfer is to a target mode with wavevector \mathbf{k} , from a source mode with wavevector \mathbf{k}' , and this occurs via their interaction with a mediator with $\mathbf{k} - \mathbf{k}'$ as its wavevector²². Any of the three modes may be turbulent fluctuations or zonal flows, hence these processes include the nonlinear transfer of energy among turbulent modes, among zonal flows, and between turbulence and zonal flows. The two scenarios examined in this chapter are where $\mathbf{k} - \mathbf{k}'$ is a zonal flow (i.e. the action of zonal flows on turbulence) or where \mathbf{k} is a zonal flow (i.e. nonlinear transfer into zonal flows).

The primary damping mechanism of zonal flows is collisional friction. If the rate of collisional damping is slower than the zonal flow growth rate due to the drive mechanisms described above (as is often the case for hot and, thus, low collisionality fusion plasmas), then zonal flows can develop. Furthermore, once established, zonal flows can persist in the absence of turbulence to provide the drive, and will decay on a time-scale set by the collisionality. The simulations reported in this chapter are collisionless, but include an equivalent damping mechanism via a hyper-diffusion term. In addition, zonal flows may be damped by various nonlinear processes such as tertiary instability (e.g. Kelvin-Helmholtz (KH) instability or KH-like instabilities),

²¹There may be additional drive from poloidal asymmetry in the particle flux, known as Stringer spin-up, but this is not examined directly in this thesis

²²In the case of parametric decay, this is via the decay of \mathbf{k}' into \mathbf{k} and $\mathbf{k} - \mathbf{k}'$, while in the case of modulational instability this is by the scattering of \mathbf{k}' off of $\mathbf{k} - \mathbf{k}'$ into \mathbf{k}

nonlinear wave-packet scattering and nonlinear wave-packet trapping [18, §3.4], all of which are included via the triad interactions described above. These would manifest as reduced (or negative) transfer in the case where \mathbf{k} is a zonal flow, i.e. nonlinear transfer *out of* zonal flows.

Zonal flows, being $\mathbf{E} \times \mathbf{B}$ flows, occur readily in electrostatic ($\beta = 0$) simulations and are, as such, primarily an electrostatic phenomenon (i.e. due to turbulence generated by a fluctuating electric field without magnetic field fluctuations). Electromagnetic effects (i.e. with magnetic perturbations) can provide additional damping at low β or drive at high β , but this is neglected in this thesis for simplicity. Furthermore, zonal flows are stronger for ion-scale turbulence than electron-scale, due to the screening of electron-scale zonal flow drive by the ion response. For these reasons (and to simplify the numerics), the simulations reported in this chapter are of electrostatic ion temperature gradient (ITG) turbulence with adiabatic electrons.

Zonal flows regulate turbulence in two ways. The first, is by the shearing of turbulent eddies into progressively smaller structures. This constitutes an increase in the radial wavenumber, known as a “forward cascade”. This continues until the turbulence at small scales is damped by viscosity; this is also represented by the hyper-diffusion term in the simulations reported here. The second way in which zonal flows regulate turbulence is by nonlinear transfer. Given zonal flows are driven by a transfer of energy from turbulence, this constitutes a corresponding decrease in the turbulent intensity. This leads to the “predator-prey” model of interaction between turbulence and zonal flows, where the turbulence is the prey and the zonal flows the predator: as the prey population (turbulent intensity) increases, this feeds the predator population (zonal flow intensity), which thus increases soon after. This, in turn, reduces the prey population and the predator population then also decreases, thus allowing the prey population to grow again.

Depending on available turbulence drive, the predator-prey relationship may result in a quasi-steady state solution or a limit cycle oscillation (LCO) type response. The latter is thought to be responsible for the Dimits shift [36], whereby the nonlinear onset of turbulent transport occurs at a higher pressure gradient than the linear onset of instability. Two scenarios are studied in this chapter: one from within the Dimits region (zonal flow dominated), and one from beyond it (turbulence dominated). Correspondingly, both solutions (LCO-type response and quasi-steady state) are observed.

There is a growing body of experimental evidence to support the theory of zonal flows, including experimental confirmation of their existence, and their interaction with and regulation of turbulence and transport [79, and references therein]. However, quantitative comparisons are few in number. Furthermore, most experiments have focused on the related and more easily measured Geodesic Acoustic Mode (GAM), which includes coupling to the $m = 1$ and $m = 2$ sidebands and thus features an oscillation; consequently, measurements of the pure zonal flow, which is quasi-stationary (hence why GAMs are easier to measure), are relatively limited.

Another reason that zonal flow measurements are relatively few in number is the diagnostic requirements. The strongest evidence of zonal flows is measurement of the $\mathbf{E} \times \mathbf{B}$ velocity, radial electric field and/or electrostatic potential. These data provide direct evidence of zonal flows and enable the calculation of nonlinear energy transfer for corroboration against simulations. Such measurements are available from various diagnostics including Doppler back scattering (DBS), heavy ion beam probe (HIBP), charge exchange recombination spectroscopy (CHERS), indirectly from motional Stark effect (MSE) and beam emission spectroscopy (BES) [18], and gas-puff imaging (GPI) using time-delay-estimation (TDE) velocimetry [24, 32]. However, for various reasons (such as the balance of resolution *vs* cost, and competition for space with other diagnostics), these diagnostics typically

have a poloidally localised measurement volume. Therefore, knowledge of the poloidal structure of zonal flow drive would facilitate corroboration of these poloidally limited measurements among the diagnostics listed above and validation of nonlinear transfer models against experiments. Furthermore, knowledge of the relationship between the poloidal structure of zonal flow drive and the corresponding turbulent intensity would allow corroboration against turbulent fluctuation measurements from other diagnostics.

The rate of transfer due to the nonlinear driving and damping processes described above can be measured (in experiments and simulations) using nonlinear energy transfer functions (NETFs). Such methods have been in development since at least the 1980s [80, 81], and have matured significantly since then [82–85]. More recently, these methods have been applied in various studies [20, 34, 35, 86, 87], including measurement of nonlinear energy transfer *due to* GAMs in experiments and simulations [33], *into* GAMs in experiments [32], and *into* zonal flows *in experiments* [24]. However, none of the literature to date measures nonlinear energy transfer into zonal flows in simulations, nor does any measure the poloidal variation of nonlinear energy transfer. Since simulations are not limited in their poloidal measurement area, such calculations would reveal the poloidal structure of zonal flow drive in support of experimental measurements of zonal flows.

Given the importance of zonal flows and the issues with experimental measurements discussed above, this chapter determines the poloidal structure of zonal flow drive and its relationship to turbulent intensity. This is done using NETFs applied to nonlinear electrostatic local gyrokinetic simulations of core ITG turbulence. Section 6.2 describes the gyrokinetic simulation setup and the method of calculating NETFs. Then, the poloidal structure of the zonal flow drive of two selected cases (zonal flow dominated and turbulence dominated) is presented and discussed in Section 6.3 before drawing conclusions in Section 6.4.

6.2 Methods

6.2.1 Simulation Setup

The simulations in this chapter are based on a Miller equilibrium model [65] representation of the Cyclone Base Case (CBC) [36], as described in Chapter 4. All parameters are the same as those used in Chapter 4 (Table 1) with the exception of:

- $R_0/a = R_a/a = 2.72$ *cf.* 3.333 (minor difference, already shown in Chapter 4 to be insignificant)
- $\hat{s} = 0.78$ *cf.* 0.8 (minor difference, insignificant)
- $\beta = 0$ *cf.* 0.02 (to make the simulations electrostatic – the A_{\parallel} and B_{\parallel} fields were also explicitly turned off)
- $\nu_{ii} = \nu_{ee} = 0$ (to make the simulations collisionless – the collision operator was also explicitly turned off)
- multiple n are included (since the simulations are nonlinear)
- various a/L_T were tested (to obtain the zonal flow dominated and turbulence dominated cases).

The simulations here also treat electrons adiabatically rather than kinetically to simplify the numerical setup, and use a hyper-diffusion term to provide numerical damping in the absence of collisions. The justification for these changes, plus evidence of nonlinear convergence is given in Appendix D. Equilibrium shear flows are also neglected for simplicity.

6.2.2 Nonlinear Energy Transfer Functions

Following the methods of previous NETF studies [24, 32, 33, 82, 85], we evaluate the NETF, denoted as T , of a generic fluctuating field, $f(\mathbf{x})$, as the

partial time derivative at wavevector \mathbf{k} of the spectral power, $|\hat{f}(\mathbf{k})|^2$, due to that mode's nonlinear interaction with modes of wavevector \mathbf{k}' and $\mathbf{k} - \mathbf{k}'$:

$$T_f(\mathbf{k}, \mathbf{k}') \equiv \left[\left(\frac{\partial |\hat{f}(\mathbf{k})|^2}{\partial t} \right)_{\text{NL}} \right]_{\mathbf{k}'} \quad (89)$$

where $\hat{f}(\mathbf{k}) \equiv \mathcal{F}\{f(\mathbf{x})\}$, the round brackets and subscript NL on the right-hand side indicate that we only consider nonlinear effects, and the square brackets and subscript \mathbf{k}' indicate that we only consider a single \mathbf{k}' rather than an integration over all \mathbf{k}' .

To derive a calculable form for T , we begin by considering the fluctuating field, $f(\mathbf{x})$, in real space. Retaining the rate of change due only to the convective nonlinearity $\mathbf{u} \cdot \nabla f$, the total derivative is zero, hence we have:

$$\left(\frac{\partial f(\mathbf{x}, t)}{\partial t} \right)_{\text{NL}} = -\mathbf{u}(\mathbf{x}, t) \cdot \nabla f(\mathbf{x}, t) \quad (90)$$

where \mathbf{u} is the $\mathbf{E} \times \mathbf{B}$ velocity²³.

Fourier transforming to wavenumber space, and applying the convolution theorem, we obtain:

$$\frac{\partial \hat{f}(\mathbf{k})}{\partial t} = - \sum_{\alpha} \int_{-\infty}^{+\infty} \hat{u}_{\alpha}(\mathbf{k} - \mathbf{k}') \partial_{\alpha} \hat{f}(\mathbf{k}') d\mathbf{k}' \quad (91)$$

where α are the directions of the coordinate system, \hat{u}_{α} is the component of $\hat{\mathbf{u}}$ in the α direction, $\partial_{\alpha} \hat{f} \equiv \mathcal{F}\{\nabla_{\alpha} f(\mathbf{x})\}$ and the $(\dots)_{\text{NL}}$ notation has been dropped for brevity.

Next, using the identity:

²³Hence \mathbf{u} is in the plane perpendicular to \mathbf{B} , so the parallel component of the gradient operator is not required

$$\frac{\partial |\hat{f}(\mathbf{k})|^2}{\partial t} \equiv 2\Re \left[\hat{f}^*(\mathbf{k}) \frac{\partial \hat{f}(\mathbf{k})}{\partial t} \right] \quad (92)$$

we find:

$$\frac{\partial |\hat{f}(\mathbf{k})|^2}{\partial t} = -2\Re \left[\sum_{\alpha} \hat{f}^*(\mathbf{k}) \int_{-\infty}^{+\infty} \hat{u}_{\alpha}(\mathbf{k} - \mathbf{k}') \partial_{\alpha} \hat{f}(\mathbf{k}') d\mathbf{k}' \right] \quad (93)$$

where $\hat{f}^*(\mathbf{k})$ is the complex conjugate of $\hat{f}(\mathbf{k})$ and \Re indicates the real component.

Since we desire the change due to a specific source mode, $\hat{f}(\mathbf{k}')$, we may drop the integration, which yields the general form of a NETF:

$$T_f(\mathbf{k}, \mathbf{k}') = -2\Re \left[\sum_{\alpha} \hat{f}^*(\mathbf{k}) \hat{u}_{\alpha}(\mathbf{k} - \mathbf{k}') \partial_{\alpha} \hat{f}(\mathbf{k}') \right] \quad (94)$$

As noted above, the relevant flow velocity is the $\mathbf{E} \times \mathbf{B}$ velocity. Orienting a magnetic field line following rectangular coordinate system such that $\mathbf{B} = B\hat{\mathbf{z}}$ (as per GS2), we find:

$$\hat{\mathbf{u}} = \frac{1}{B}(\hat{E}_y \hat{\mathbf{x}} - \hat{E}_x \hat{\mathbf{y}}) = \frac{1}{B}(\partial_x \hat{\phi}_0 \hat{\mathbf{y}} - \partial_y \hat{\phi}_0 \hat{\mathbf{x}}) \quad (95)$$

Next, we use:

$$\partial_{\alpha} \hat{f}(\mathbf{k}) = ik_{\alpha} \hat{f}(\mathbf{k}) \quad (96)$$

where k_{α} is the component of \mathbf{k} in the α direction²⁴.

²⁴Since we do not require the parallel component of the gradient operator, we hence do not require k_{\parallel}

Finally, substituting the local density fluctuation obtained from GS2, $\hat{f} = \hat{n}_0$, we obtain the density transfer function²⁵:

$$T_n(\mathbf{k}, \mathbf{k}') = \frac{2(\mathbf{k} \times \mathbf{k}' \cdot \hat{\mathbf{z}})}{B} \Re \left[\hat{n}_0^*(\mathbf{k}) \hat{\phi}_0(\mathbf{k} - \mathbf{k}') \hat{n}_0(\mathbf{k}') \right] \quad (97)$$

Substituting instead the flow velocity, $\hat{f} = \hat{\mathbf{u}}$, which we obtain from GS2's local electrostatic potential, $\hat{\phi}_0$, via Equation (95), yields the velocity (or kinetic energy) transfer function:

$$T_u(\mathbf{k}, \mathbf{k}') = \frac{2(\mathbf{k} \times \mathbf{k}' \cdot \hat{\mathbf{z}})(\mathbf{k} \times \hat{\mathbf{z}}) \cdot (\hat{\mathbf{z}} \times \mathbf{k}')}{B^3} \Re \left[\hat{\phi}_0^*(\mathbf{k}) \hat{\phi}_0(\mathbf{k} - \mathbf{k}') \hat{\phi}_0(\mathbf{k}') \right] \quad (98)$$

Evaluating Equations (97) and (98) presents a computationally demanding challenge, since \mathbf{k} and \mathbf{k}' are each 2D, B , \hat{n}_0 and $\hat{\phi}_0$ are functions of poloidal angle, and \hat{n}_0 and $\hat{\phi}_0$ are also functions of time. Therefore, a computer program [4] was written to efficiently calculate transfer function spectra at a single arbitrary poloidal angle, or net transfer into zonal flows as a function of poloidal angle. Both functions can be calculated with or without time averaging. The program is described in more detail in Appendix E.

6.3 Results and Discussion

6.3.1 Development of Selected Cases

6.3.1.1 Nonlinear Benchmarking

Having developed a numerically robust and well converged simulation setup (see Appendix D), the physics of the simulations could be benchmarked against previous work. To do this, we measure the critical tempera-

²⁵The density transfer function has also been referred to as the “internal fluctuation energy” [33] and “internal energy” [86] transfer function

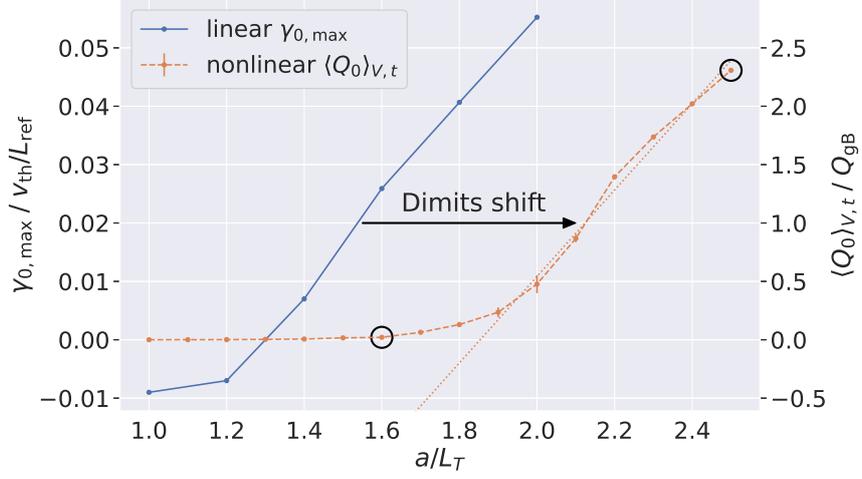


Figure 36: Maximum linear growth rate, $\gamma_{0,\max}$, and the volume- and time-averaged nonlinear heat flux, $\langle Q_0 \rangle_{V,t}$, as functions of normalised temperature gradient length-scale, a/L_T , exhibiting Dimits shift in agreement with [36]. The circled points indicate the two nonlinear simulations selected for further analysis.

ture gradient for the linear and nonlinear onset of ITG turbulence and, thus, reproduce the Dimits shift [36].

To measure the critical temperature gradient for the linear onset of ITGs, the simulation setup was run linearly with various values for the normalised temperature gradient length-scale, a/L_T . For each value of a/L_T , the binormal wavenumber, k_y , was also varied. Then, the maximum growth rate, $\gamma_{0,\max}$, was determined by spline interpolation of $\gamma_0(k_y)$ for each a/L_T value. The results are shown in Figure 36. The linear critical temperature gradient was then determined by linear interpolation as the value of a/L_T at which $\gamma_{0,\max}(a/L_T)$ was zero.

In a similar way, the critical temperature gradient for the nonlinear onset of ITGs was measured by re-running the simulation setup nonlinearly with

various values for a/L_T . Since this was done with nonlinear simulations, they already include many k_y and their nonlinear interactions.

The nonlinear heat flux was used to measure the nonlinear critical temperature gradient. This was done using the GS2 output quantity `hflux_tot`, which is related to the heat flux density tensor, \mathbf{Q}_s , which is defined as [41]:

$$\mathbf{Q}_s(\mathbf{x}, t) = \frac{1}{2} m_s \int_{-\infty}^{+\infty} w_s^2 \mathbf{w}_s f_s(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \quad (99)$$

where subscript s is the species label (ions in this case), m_s is the species mass, $\mathbf{w}_s = \mathbf{v} - \mathbf{u}_s$ is the velocity of each particle relative to the bulk flow velocity, and f_s is the species distribution function. `hflux_tot` is the magnitude of the tensor contraction of the component of \mathbf{Q}_s due to the non-adiabatic part of the perturbed distribution function averaged over the flux tube volume, denoted as $\langle Q_0 \rangle_V$ in this thesis. `hflux_tot` is normalised in the GS2 output and $\langle Q_0 \rangle_V$ is normalised in this thesis; both are normalised to the gyro-Bohm heat flux [56]:

$$Q_{\text{gB}} = \frac{n_{\text{ref}} T_{\text{ref}} v_{\text{th}} \rho_{\text{ref}}^2}{2\sqrt{2} L_{\text{ref}}^2} \quad (100)$$

Heat flux was used because this provides a measure of energy transport due to turbulence; the higher the heat flux, the faster energy is being lost from the magnetic confinement. However, if the modes are linearly unstable but the heat flux is low, this suggests that some other mechanism (namely, zonal flow) is preventing the unstable modes from generating sufficient turbulence to cause appreciable losses.

$\langle Q_0 \rangle_V$ was averaged over time for each value of a/L_T . For simplicity, the initial transient due to the overshoot of the linear growth phase before nonlinear saturation was not excluded from the time average as this is expected to have negligible impact for sufficiently long simulation times.

The resulting volume- and time-averaged heat flux, $\langle Q_0 \rangle_{V,t}$, as a function of a/L_T is shown in Figure 36 alongside the linear $\gamma_{0,\max}(a/L_T)$. Figure 36 exhibits the Dimits shift, as expected.

The nonlinear critical temperature gradient was then determined by linear extrapolation of the steep gradient section of $\langle Q_0 \rangle_{V,t}$ as a function of a/L_T , as indicated in Figure 36. The critical value was taken as the value of a/L_T at which the extrapolated $\langle Q_0 \rangle_{V,t}$ was zero.

The results of this analysis indicate that the linear onset of ITGs is $a/L_T = 1.30$, while the nonlinear onset is $a/L_T = 1.85$, i.e. the nonlinear onset is shifted up by $\sim 40\%$ compared to the linear onset. In comparison, Dimits *et al.* [36] calculate the linear onset as $a/L_T = 1.6$ for the most similar of the available equilibrium models and the nonlinear onset as $a/L_T = 2.4$, corresponding to a shift of $\sim 50\%$. These results are in reasonable agreement, albeit with some quantitative differences that are attributed to the use of different gyrokinetic codes, equilibrium models and critical gradient calculation methods. Nonetheless, this comparison shows that the nonlinear simulations presented in this thesis are displaying the expected physical behaviour to a degree sufficient to support conclusions regarding the poloidal structure of zonal flow drive, which are anyway qualitative for this first study.

6.3.1.2 Overview of Selected Cases

Having developed a numerically and physically justified set of nonlinear simulations, two cases were selected for further analysis as indicated in Figure 36. These cases were selected because the one at $a/L_T = 1.6$ is within the Dimits region, so is expected to be dominated by zonal flows, whereas the one at $a/L_T = 2.5$ is beyond the Dimits region, so is expected to be dominated by turbulence.

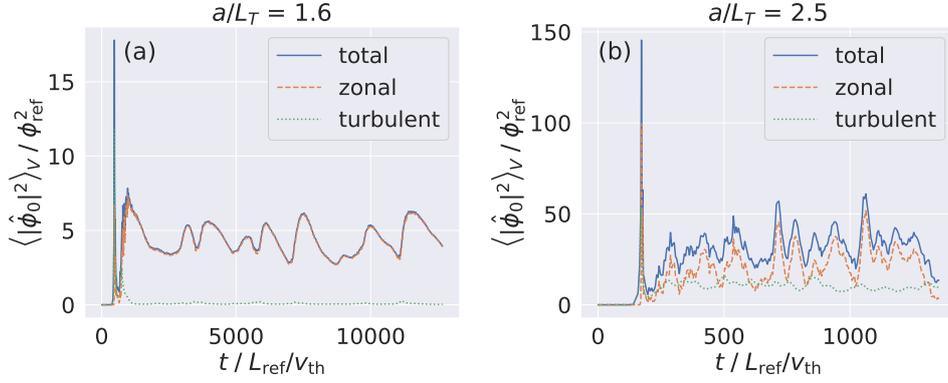


Figure 37: Nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t , alongside the zonal and turbulent components thereof for the zonal ($a/L_T = 1.6$) and turbulent ($a/L_T = 2.5$) cases. Notice the differing t and $\langle |\hat{\phi}_0|^2 \rangle_V$ scales.

The selected cases were re-run with additional, more expensive simulation diagnostics enabled to facilitate the transfer function calculations detailed in Subsection 6.2.2. This resulted in dynamically different but statistically similar data compared to those reported in Subsection 6.3.1.1. This is due to the use of different initial conditions (random numbers from different processors), different floating point and MPI operation order, *etc.*

Figure 37 shows the nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t . Also shown are the zonal and turbulent components thereof²⁶. This shows that the $a/L_T = 1.6$ case is indeed dominated by zonal flows, while the $a/L_T = 2.5$ case has a much more significant contribution from the turbulent modes. Therefore, these cases shall be referred to as the zonal case and the turbulent case, respectively.

²⁶Calculated by summing over the relevant elements of the GS2 output quantity `phi2_by_mode` taking account of GS2's weighting of $k_y = 0$ and $k_y \neq 0$ components

To further demonstrate the zonal and turbulent nature of these cases, one may define the zonality, Z , as the fraction of the total energy in the electrostatic potential fluctuation field that is in the zonal flow component:

$$Z = \frac{\langle |\hat{\phi}_0|^2 \rangle_{V,\text{zonal}}}{\langle |\hat{\phi}_0|^2 \rangle_{V,\text{total}}} \quad (101)$$

Time-averaging Z excluding the initial transient due to the overshoot of the linear growth phase before nonlinear saturation (i.e. for $t \gtrsim 2000L_{\text{ref}}/v_{\text{th}}$ for the zonal case and for $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$ for the turbulent case) shows that the zonal case has $Z = 0.9808 \pm 0.0005$, while the turbulent case has $Z = 0.65 \pm 0.01$.

The spectra of the time- and field-line-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_{t,\theta}$, for the zonal and turbulent cases are shown in Figure 38. These show that the turbulence has largely the same spectral structure in both cases, in agreement with the linear result (not shown) that the structure of the k_y spectrum does not change significantly with a/L_T . However, the balance between turbulent and zonal modes is different, with noticeably more turbulent activity in the turbulent case, in agreement with the results shown in Figure 37.

Finally, the Fourier coefficients of the electrostatic potential fluctuation field at the mid-point of the flux tube²⁷, $\hat{\phi}_0(t, k_y, k_x, \theta = 0)$, can be Fourier transformed back to real-space to obtain $\phi_0(t, x, y, \theta = 0)$ and thus visualise the structure of the turbulence in real-space by taking a snapshot in t . The results of this operation are presented in Figure 39, which shows clear zonal banding in the zonal case and a mixture of smaller scale eddies in the turbulent case.

²⁷From the GS2 output quantity `phi0`

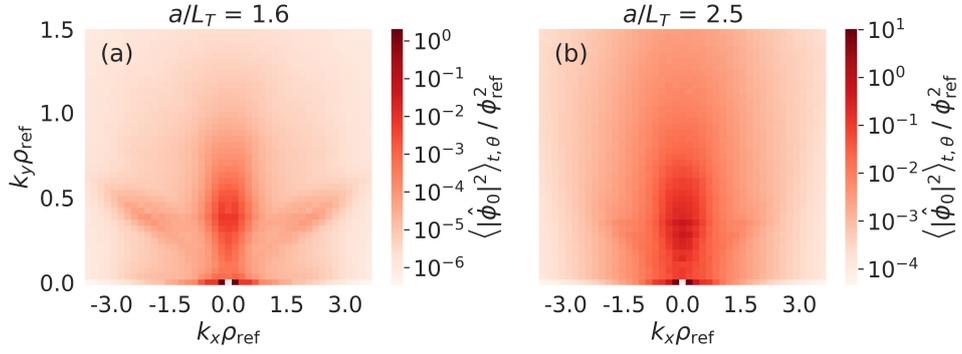


Figure 38: Spectra of the time- and field-line-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_{t,\theta}$, for the zonal ($a/L_T = 1.6$) and turbulent ($a/L_T = 2.5$) cases selected for further analysis. Notice the differing $\langle |\hat{\phi}_0|^2 \rangle_{t,\theta}$ scales.

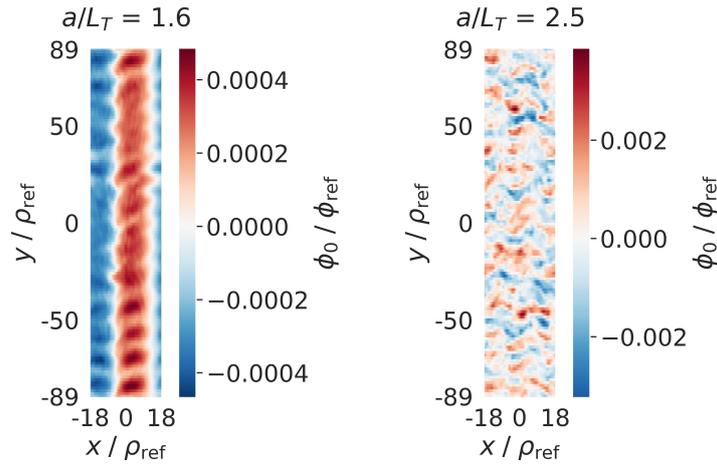


Figure 39: Real-space snapshots of the electrostatic potential fluctuation field, $\phi_0(x, y)$, at $\theta = 0$ for the zonal ($a/L_T = 1.6$) and turbulent ($a/L_T = 2.5$) cases selected for further analysis. In both cases, the snapshot is at the last available time point. Notice the differing ϕ_0 scales.

6.3.2 Transfer Function Spectrum

To verify the correctness of the transfer functions and the subsequent computer implementation described in Subsection 6.2.2, the full 4-D transfer function spectra, $T(k_x, k_y, k'_x, k'_y)$, were calculated (at $\theta = 0$ as a representative location) for the cases selected in Subsection 6.3.1.2. As an illustrative example, this subsection discusses the transfer function spectra for the turbulent case time-averaged over the steady state period ($t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$); the transfer functions for the zonal case show similar results.

By the reality condition, $\hat{f}^*(\mathbf{k}) = \hat{f}(-\mathbf{k})$, we expect the transfer functions to be anti-symmetric about $\mathbf{k} = \mathbf{k}'$. To demonstrate the required symmetry, the kinetic energy transfer function²⁸, $T_{\mathbf{u}}(k_x, k'_x)$, for a given fluctuation scale in the binormal direction, $k_y = k'_y = \pm 0.35$, is shown in Figure 40. The chosen binormal scale is the peak k_y for the dominant (first) non-zero k_x in the $\langle |\hat{\phi}_0|^2 \rangle_{t,\theta}$ spectrum as shown in Figure 38 (b). Inspection of the transfer spectrum for different values of $k_y = k'_y$ shows similar results. Figure 40 shows transfer from k'_x (vertical axis) to k_x (horizontal axis) where the colour represents the energy transfer rate with positive values (red) indicating transfer from k'_x to k_x and negative values (blue) indicating transfer in the opposite direction.

Each coloured box in Figure 40 represents a $(\mathbf{k}, \mathbf{k}')$ pair. Since the transfer function spectrum in Figure 40 is shown for a given fluctuation scale in the binormal direction, the box that is opposite across the line $k_x = k'_x$ therefore represents the same transfer channel but with source and target waves swapped (i.e. $\mathbf{k} \rightarrow \mathbf{k}'$ and $\mathbf{k}' \rightarrow \mathbf{k}$). Therefore, the transfer function must be anti-symmetric about the line $k_x = k'_x$, which is indeed the case in Figure 40. Furthermore, the transfer must be zero along the line $k_x = k'_x$ (modes cannot transfer energy to themselves), which is also the case in Figure 40. Finally, if

²⁸Calculated using the GS2 output quantity `phi_t` as $\hat{\phi}_0$ in Equation (98)

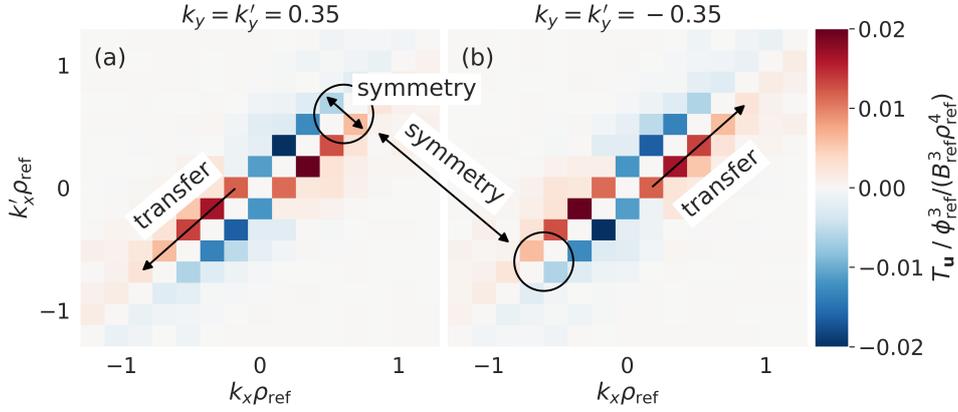


Figure 40: Kinetic energy transfer spectrum, $T_{\mathbf{u}}(k_x, k'_x)$ for the turbulent case at $k_y = k'_y = \pm 0.35$, $\theta = 0$ and time-averaged over $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$. The short arrow highlights elements that indicate the same transfer channel but with source and target waves swapped and the long arrow highlights elements that indicate equivalent interactions due to waves travelling in opposite directions.

we consider the interaction between three arbitrary wavevectors, \mathbf{k} , \mathbf{k}' and $\mathbf{k} - \mathbf{k}'$, and then reverse the sign on those wavevectors, then we have an interaction between three waves of the same scales and relative directions as the original interaction but the waves are now all moving in the opposite direction. The combination of this fact with the reality condition means that we expect the transfer function to be invariant under the transformation $\mathbf{k} \rightarrow -\mathbf{k}$ and $\mathbf{k}' \rightarrow -\mathbf{k}'$. This property can also be seen in the transfer function spectrum by comparison of Figures 40 (a) and 40 (b).

The transfer spectrum in Figure 40 shows a cascade of energy from large scales (small wavenumbers) towards small scales (large wavenumbers), as indicated by positive $T_{\mathbf{u}}$ values for $|k_x| > |k'_x|$ and negative values for $|k_x| < |k'_x|$. Since the data in Figure 40 is for $k_y = k'_y$, the interaction is mediated by the zonal flow ($k_y - k'_y = 0$). Therefore, this shows that the action of the zonal flow is to transfer kinetic energy from low k_x velocity fluctuations towards

high k_x velocity fluctuations. This can be understood physically as zonal flow shearing of turbulent eddies into progressively smaller structures. In addition, being a cascade, this process is dominated by “local” transfer (i.e. from one scale to the next), as opposed to “non-local” transfer (i.e. directly from the largest scales to the smallest scales) as indicated by the strongest transfer channels being those closest to the line $k_x = k'_x$. The observed symmetry and the fact that we have a local cascade is in agreement with previous work [33, 82].

The density transfer function²⁹, $T_n(k_x, k'_x)$, for the same fluctuation scale in the binormal direction as used for $T_u(k_x, k'_x)$ above, $k_y = k'_y = \pm 0.35$, is shown in Figure 41. The symmetry and local cascade structure of the transfer function is largely the same for $T_n(k_x, k'_x)$ in Figure 41 as compared to $T_u(k_x, k'_x)$ in Figure 40. However, there is one key difference: $T_n(k_x, k'_x)$ shows an inverse cascade towards large scales, whereas $T_u(k_x, k'_x)$ shows a forward cascade towards small scales, as indicated by negative T_n values for $|k_x| > |k'_x|$ (where T_u is positive) and positive T_n values for $|k_x| < |k'_x|$ (where T_u is negative). This observation of a dual cascade in opposing directions is expected and in agreement with previous work [82]. The fact that T_n shows an inverse cascade (and, indeed, the structure of the T_n spectrum more generally) is in qualitative agreement with previous work [33]³⁰.

While the T_n values can be compared across different cases (and likewise for T_u values), a direct comparison of the T_n and T_u values with each other is not strictly valid as they are in different normalised units. The conversion to comparable physical units is non-trivial and is not of great significance to the arguments presented here so has been omitted. Despite this limitation on direct comparison, the normalised units have been chosen such that quantities are typically $\sim O(1)$. Therefore, the order of magnitude of the

²⁹Calculated using the GS2 output quantity `density_t` as \hat{n} in Equation (97)

³⁰A quantitative comparison is not possible as [33] gives results in arbitrary units.

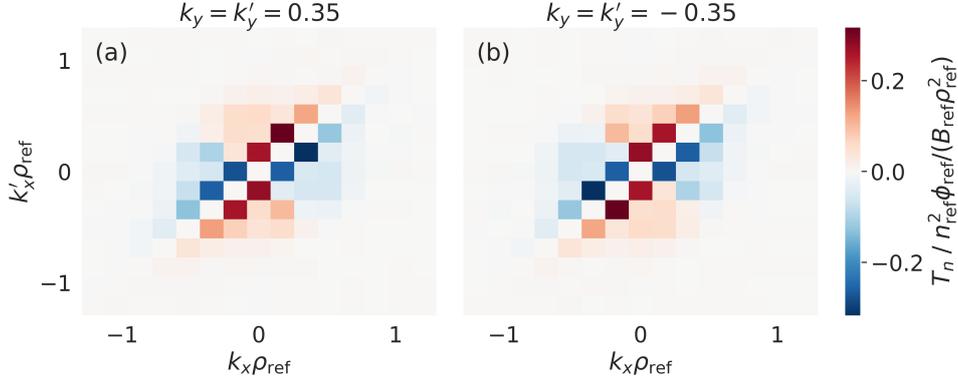


Figure 41: Internal energy transfer spectrum, $T_n(k_x, k'_x)$, for the turbulent case at $k_y = k'_y = \pm 0.35$, $\theta = 0$ and time-averaged over $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$.

transfer in normalised units relative to $\sim O(1)$ may give an approximate indication of the significance of the transfer. In this context, the transfer spectra in Figures 40 and 41 suggest that the velocity transfer rate between individual modes may be comparable to that of the density transfer.

To further verify the transfer function calculation and to investigate the spectrum of which modes are transferring energy into the zonal flow, a plot of $T_{\mathbf{u}}(k_y, k'_y)$ for a given fluctuation scale in the radial direction, $k_x = k'_x = \pm 0.17$, is shown in Figure 42. The chosen radial scale is the dominant (first) non-zero k_x . Inspection of the transfer spectrum for different values of $k_x = k'_x$ shows similar results. This again exhibits the required symmetry (anti-symmetric about $k_y = k'_y$, zero along $k_y = k'_y$ and invariant under the transform $\mathbf{k} \rightarrow -\mathbf{k}$ and $\mathbf{k}' \rightarrow -\mathbf{k}'$). Figure 42 also shows a forward cascade, albeit much weaker than that in Figure 40. However, this is overshadowed by non-local transfer from a spectrum of turbulent modes directly into the zonal flows. The density transfer spectrum shows similar results but with the colours reversed.

To further investigate the spectrum of transfer into the zonal flows, the $k_y = 0$ component of $T_{\mathbf{u}}(|k_x|, |k_y|, |k'_x|, |k'_y|)$ is shown in Figure 43. The plot

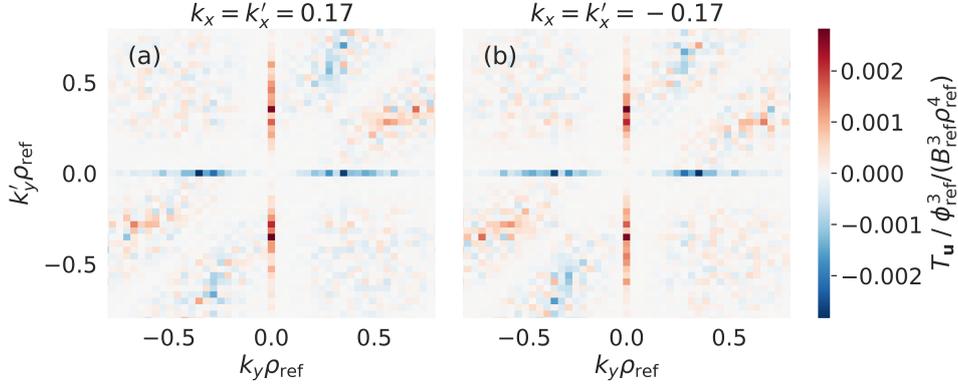


Figure 42: Kinetic energy transfer spectrum, $T_{\mathbf{u}}(k_y, k'_y)$ for the turbulent case at $k_x = k'_x = \pm 0.17$, $\theta = 0$ and time-averaged over $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$.

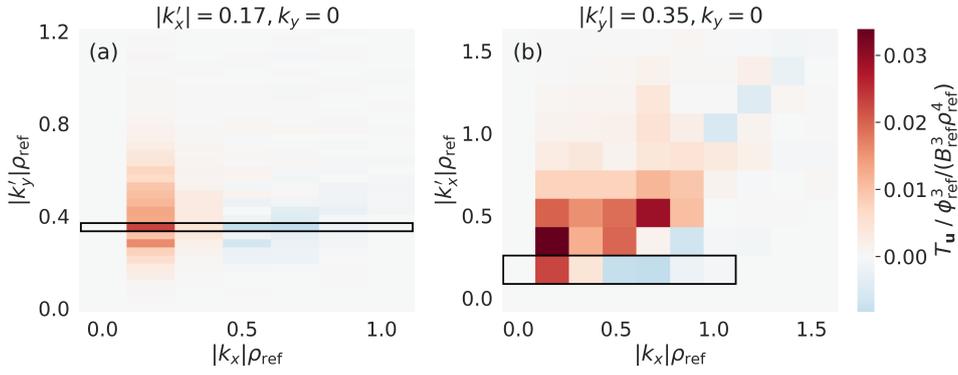


Figure 43: The spectrum of kinetic energy transfer into zonal flows, $T_{\mathbf{u}}(|k_x|, k_y = 0, |k'_x|, |k'_y|)$ for the turbulent case at $\theta = 0$ and time-averaged over $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$ for (a) $|k'_x| = 0.17$ and (b) $|k'_y| = 0.35$. Different values of $|k'_x|$ and $|k'_y|$ show similar results. The highlighted elements represent the same $(\mathbf{k}, \mathbf{k}')$ pairs.

is no longer expected to exhibit symmetry in the same way as Figures 40 to 42, as we are no longer plotting the same modes on both axes; rather, we have zonal flows on the horizontal axis and turbulent modes on the vertical axis. To simplify the plot in Figure 43, $T_{\mathbf{u}}$ values have been summed over all interactions involving fluctuations of the same scale.

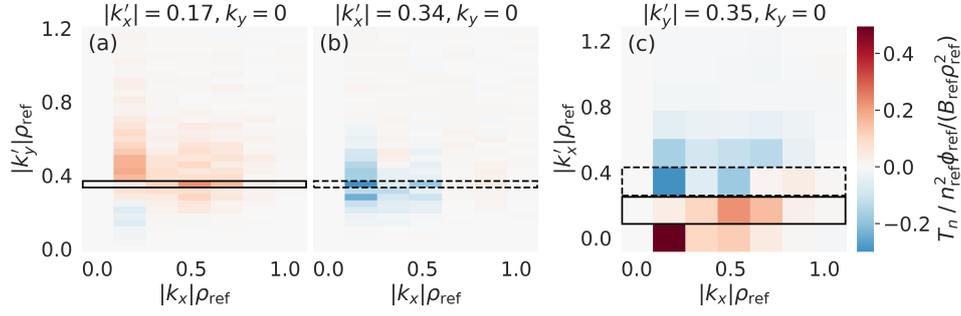


Figure 44: Internal energy transfer spectrum from turbulent density fluctuations into the zonal component of the density fluctuations, $T_n(|k_x|, k_y = 0, |k'_x|, |k'_y|)$ for the turbulent case at $\theta = 0$ and time-averaged over $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$ for (a) $|k'_x| = 0.17$, (b) $|k'_x| = 0.34$ and (c) $|k'_y| = 0.35$. Different values of $|k'_x|$ show results similar to (b) and different values of $|k'_y|$ show results similar to (c). The highlighted elements represent the same $(\mathbf{k}, \mathbf{k}')$ pairs. Those within the solid lines are the same pairs highlighted in Figure 43.

Figure 43 (a) shows that energy is transferred into the zonal flows with the largest scales from a broad spectrum of turbulence in $|k'_y|$ and that smaller scale zonal flows are actually losing energy to those same $|k'_y|$ modes. Meanwhile, Figure 43 (b) shows that energy is transferred into a broad range of zonal flow scales from a broad range of modes in $|k'_x|$.

Plots equivalent to those in Figure 43 but for T_n are shown in Figure 44 (a) and (c); an extra $|k'_x|$ value is shown in Figure 44 (b) as that shown in (a) is not typical of the other $|k'_x|$ values while that shown in (b) is.

Figure 44 (a) shows that internal fluctuation energy is transferred into a broad $|k_x|$ range of zonal components of the density fluctuations from a broad spectrum of turbulence in $|k'_y|$ at the largest $|k'_x|$ scale apart from some relatively small transfer out of the largest scale zonal components of the density fluctuations into the largest scales of turbulence in $|k'_y|$. However,

Figure 44 (b) shows that a broad $|k_x|$ range of zonal components of the density fluctuations are *losing* internal fluctuation energy to a broad range of modes in $|k'_y|$. Meanwhile, Figure 44 (c) shows that the dominant transfer of internal energy into the zonal components of the density fluctuations is that from low $|k'_x|$ modes (radial streamers) and that these zonal modes are simultaneously losing internal fluctuation energy to high $|k'_x|$ modes. Therefore, there is likely to be some cancellation between transfer in at low k'_x and transfer out at moderate k'_x .

Figures 40 to 44 suggest a rather complicated physical picture of transfer to / from zonal flows and the zonal component of the density fluctuations. However, this picture is simplified by comparison of the transfer spectrum against the turbulent activity level, which is shown in Figure 45. This was done by plotting the transfer into the largest scale zonal flow / zonal component of the density fluctuations with $|k_x| = 0.17$ alongside the corresponding elements extracted from the $\langle |\hat{\phi}_0|^2 \rangle_{t,\theta}$ spectrum in Figure 38 (b).

Figure 45 shows that the structure of the $T_u(|k'_y|)$, $T_u(k'_x)$ and $T_n(|k'_y|)$ transfer spectra closely match that of the turbulent activity level apart from regions of reduced transfer from low wavenumber modes, while the $T_n(k'_x)$ transfer spectrum is dominated by transfer from the low k'_x modes. This can be understood in relation to the structure of the general form of a NETF (Equation (94)), where we have: (i) the product of three fluctuating fields, hence the strong dependence on the turbulent activity level; and (ii) a gradient operator (or, when using potential as a proxy for velocity, *two* gradient operators in the case of T_n or *four* in the case of T_u), hence the reduced transfer at low wavenumbers. In the case of $T_n(k'_x)$, the leading k -factor is $k_x k'_y - k_y k'_x$. Therefore, with $k_y = 0$, we have just $k_x k'_y$, which is constant for the data shown in Figure 45 (b), hence we only have the turbulent activity level dependence on \hat{n}_0 (not shown) and $\hat{\phi}_0$. Indeed, normalising the trans-

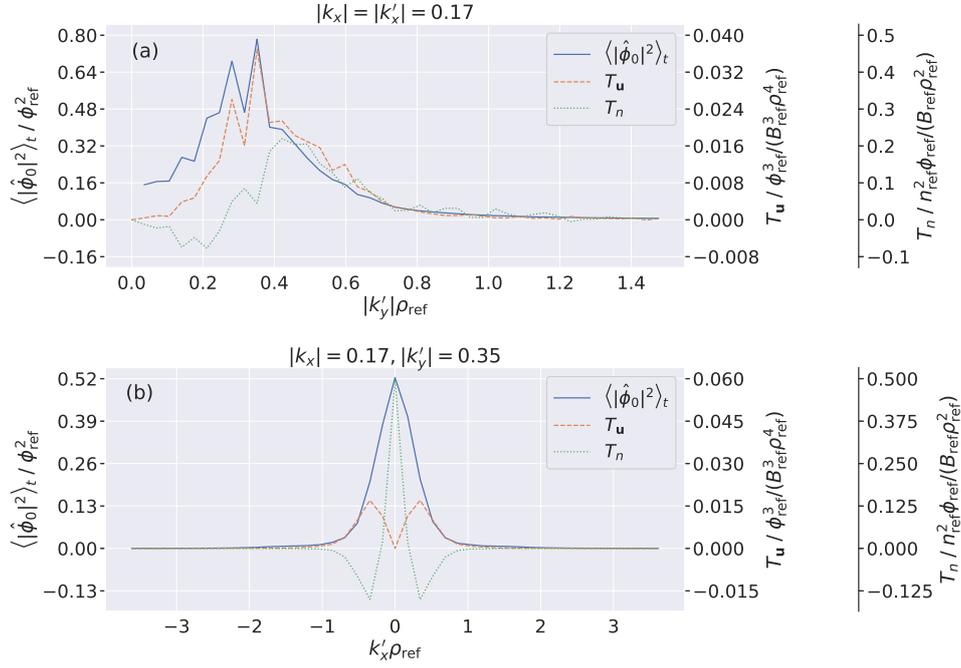


Figure 45: (a) $|k'_y|$ and (b) k'_x spectrum of turbulent activity level (solid blue line) compared against transfer into the largest scale zonal flow (dashed orange line) and transfer into the largest scale zonal component of the density fluctuations (dotted green line). All lines are taken from the turbulent case at $\theta = 0$ time-averaged over $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$.

fer to the turbulent activity level (also not shown) indicates that the high wavenumber interactions are more efficient as transfer channels (effectively highlighting the wavenumber dependence). Therefore, we have a balance between more efficient transfer channels at high wavenumber against more energy available for transfer at low wavenumber, resulting in a peak at moderate wavenumbers.

In the case of $T_n(k'_x)$, transfer at $k'_x = 0$ is non-zero because the T_n wavenumber weighting is independent of k'_x . It is unclear why the sign reverses at moderate wavenumber. This shows that large scale radially elongated structures put energy into zonal flows, structures of moderate radial scale take

energy out of zonal flows, and small scale structures exchange little energy. One possible explanation is that large scale radially elongated structures extend across different zonal bands so they can be sheared apart to feed the zonal flow, while structures of moderate radial scale are the appropriate size to be accelerated by the zonal flow in such a way that they take energy out of the zonal flow. Meanwhile, the small scale structures are too small compared to the zonal flow to have any significant interaction. In any case, this change in sign indicates that there is likely to be significant cancellation when summing over all source radial scales (k'_x).

6.3.3 Time History

The transfer function spectra presented in Subsection 6.3.2 were, for computational performance reasons, calculated only at $\theta = 0$. However, since we are specifically interested in the transfer into / out of zonal flows, we may calculate the reduced 3-D transfer function, $T(k_x, k'_x, k'_y)$ at $k_y = 0$, thus eliminating a significant fraction of the computation. This can then be summed over k_x , k'_x and k'_y to get the net transfer into / out of zonal flows at the current θ location and time point. This procedure can then be repeated along the field line and through the time-history of the simulation to get $T(t, \theta)$. Finally, this can be summed over θ to get the net transfer as a function of time, $T(t)$, or time-averaged to obtain $T(\theta)$.

The procedure described above was applied to both of the cases selected in Subsection 6.3.1.2 to obtain $T(t)$. Continuing from the discussion in Subsection 6.3.2, let us first consider the turbulent case. Figure 46 shows the time history of T_u and T_n alongside the time history of the turbulent component of $\langle |\hat{\phi}_0|^2 \rangle_V$ (the dotted green line from Figure 37 (b)). As with the spectrum comparison in Figure 45, this shows that the transfer rate largely follows the turbulent activity level.

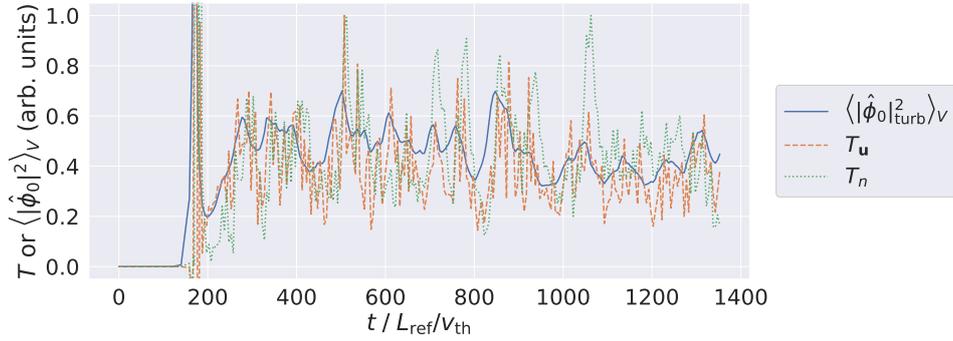


Figure 46: Time history of the turbulent case showing turbulent activity level, $\langle |\hat{\phi}_0|^2_{\text{turb}} \rangle_V$ (solid blue line), compared against the net transfer rate into zonal flows (dashed orange line) and into the zonal components of the density fluctuations (dotted green line). All lines have been arbitrarily scaled to fit on the same axes and to highlight the similarity in their structures.

Figure 47 shows the same time histories of T_u and T_n but this time alongside the time history of the *zonal* component of $\langle |\hat{\phi}_0|^2 \rangle_V$ (the dashed orange line from Figure 37 (b)). This shows that the energy in the zonal flows rises and falls with the rate of transfer into the zonal flows.

Figures 46 and 47 show that the transfer rate follows the turbulent activity level and that the zonal flow energy follows the transfer rate, albeit not that decisively. However, this observation is much clearer when we examine the time history of the zonal case, which is shown in Figure 48.

Figure 48 shows that the transfer rate follows the turbulent activity level very closely. Furthermore, the energy in the zonal flows rises sharply when there is increased transfer and then decays on a slower time-scale during periods of reduced transfer. These observations constitute the LCO-type response predicted by the predator-prey model of the zonal flow–turbulence relationship:

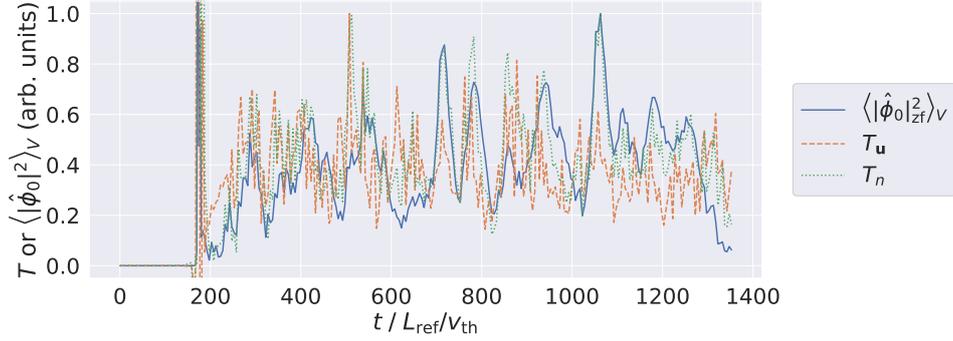


Figure 47: Time history of the turbulent case showing energy in the zonal flows, $\langle |\hat{\phi}_0|_{\text{zf}}^2 \rangle_V$ (solid blue line), compared against the net transfer rate into zonal flows (dashed orange line) and into the zonal components of the density fluctuations (dotted green line). All lines have been arbitrarily scaled to fit on the same axes and to highlight the similarity in their structures.

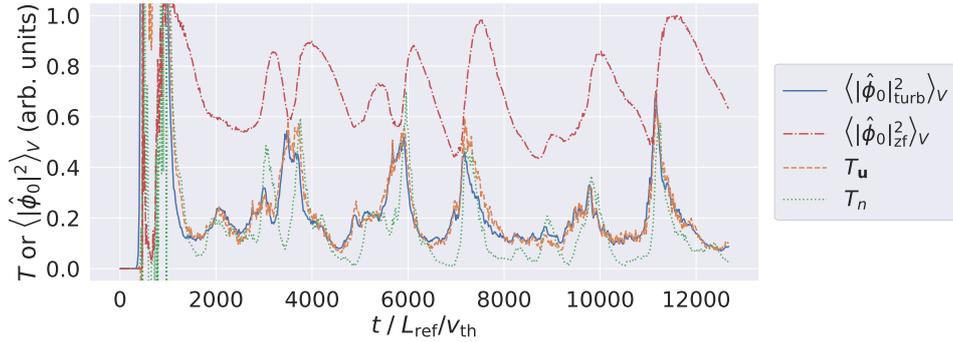


Figure 48: Time history of the zonal case showing turbulent activity level, $\langle |\hat{\phi}_0|_{\text{turb}}^2 \rangle_V$ (solid blue line), and energy in the zonal flows, $\langle |\hat{\phi}_0|_{\text{zf}}^2 \rangle_V$ (dash-dotted red line), compared against the net transfer rate into zonal flows (dashed orange line) and into the zonal components of the density fluctuations (dotted green line). All lines have been arbitrarily scaled to fit on the same axes and to highlight the similarity in their structures.

1. When there is little energy in the zonal flows, the turbulent energy increases due to extraction of free-energy from the equilibrium gradi-

ents via linear instabilities in the k_y range indicated in Figure 38. This is relatively unimpeded by zonal flow shearing (since the zonal flows are weak) and by nonlinear transfer (since the transfer rate is low).

2. Next, when there is sufficient energy in the turbulent modes, the rate of energy transfer into the zonal flows increases at the scales where there is more turbulent energy available for transfer, consistent with the observations in the turbulence and transfer spectra in Figure 45. This causes the energy in the zonal flow to increase, as shown by Figure 48.
3. This transfer reduces the energy in the turbulent modes. Plus, when there is sufficient energy in the zonal flows, this leads to radial shearing of turbulent eddies that transfers turbulent energy towards small scales (where energy is dissipated, e.g. by viscosity), consistent with the forward cascade observed in Figures 40 and 42.
4. Following on from that, when the energy in the turbulent modes has dropped such that there is insufficient energy available for transfer, the transfer rate also decreases, as seen in Figure 48.
5. Finally, when there is an established zonal flow and a low rate of energy transfer into the zonal flows, the zonal energy decays (e.g. due to collisional damping, represented here by hyper-diffusion) and the cycle repeats, as can be seen in Figure 48.

In the turbulent case, the linear modes are driven so strongly by the steeper equilibrium temperature gradient that the turbulent activity level remains high even with the transfer of energy out of turbulence and into zonal flows and with the radial shearing effect of the zonal flows. As a result, there is a continuous supply of turbulent energy available to feed the zonal flows such that there are no periods of low turbulent activity, low transfer and zonal flow decay. Instead, an equilibrium is reached between the rate of

energy gain by the linear modes, the rate of transfer into zonal flows and towards small scales, and the rates of dissipation at small scales and zonal flow damping. Hence, a quasi-steady state is achieved.

6.3.4 Poloidal structure

Instead of summing $T(t, \theta)$ over θ to get $T(t)$ as was done in Subsection 6.3.3, one can instead average over t to get the time averaged $T(\theta)$. This was done for both of the cases selected in Subsection 6.3.1.2.

Figure 49 shows, from the turbulent case, the poloidal structure of $T_{\mathbf{u}}$ and T_n alongside the poloidal structure of the turbulent and zonal components of $\langle |\hat{\phi}_0|^2 \rangle_{x,y,t}$, where the subscripts indicate averaging over the steady state period, $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$, and over the perpendicular directions (i.e. summing over all perpendicular wavenumbers). Similar results (not shown) are observed for the zonal case, for alternative time averaging windows and for instantaneous poloidal structures (i.e. considering only a single time-step with no time averaging).

Figure 49 shows that the rate of transfer into the zonal flows ($T_{\mathbf{u}}$, dashed orange line) is strongly correlated with the poloidal structure of the turbulent activity level (solid blue line), consistent with the turbulence and transfer spectra presented in Figure 45 and the time histories in Subsection 6.3.3. Importantly, however, the transfer rate is not proportional to the turbulent activity level, even in this relatively simple and symmetrical case. The reason for this correlation without proportionality is unclear. However, the data shows that, as the turbulent activity level decreases with θ , the turbulent power shifts to a different part of wavenumber space due to the effect of magnetic shear. Therefore, one possible reason is that this shift in the spectrum moves the energy to transfer channels that are more or less efficient than those where the energy is at $\theta = 0$. In any case, this correlation

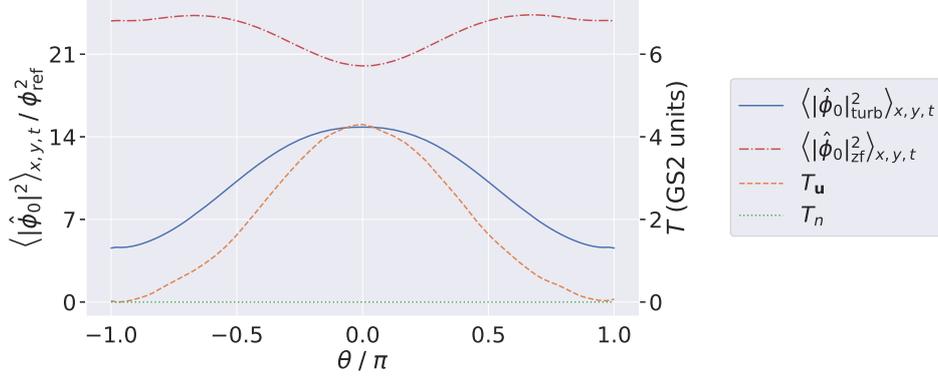


Figure 49: Poloidal structures from the turbulent case showing the turbulent activity level, $\langle |\hat{\phi}_0|_{\text{turb}}^2 \rangle_{x,y,t}$ (solid blue line), and energy in the zonal flows, $\langle |\hat{\phi}_0|_{\text{zf}}^2 \rangle_{x,y,t}$ (dash-dotted red line), compared against the transfer rate into zonal flows (dashed orange line) and into the zonal components of the density fluctuations (dotted green line). All lines have been averaged over the perpendicular directions and the steady state period, $t \gtrsim 400L_{\text{ref}}/v_{\text{th}}$.

without proportionality suggests that analysis such as this will be beneficial for zonal flow experiments to relate the small measured region to the full poloidal structure of turbulence and nonlinear transfer. For example, this could explain the remaining factor ~ 3.5 discrepancy between experimental nonlinear transfer and turbulence loss measurements reported in [24].

The poloidal distribution of zonal flow energy (dash-dotted red line) has a structure different from the turbulence. This is expected as the zonal flow is, by definition, constant in the poloidal direction. Therefore, we expect a flat line ($m = 0$) for the pure zonal flow with some coupling to the $m = 1$ and $m = 2$ sidebands (the GAM). A Fourier transform of the zonal flow structure in the θ direction, as shown in Figure 50, confirms that this is indeed the case, with the majority of the energy in the pure zonal flow, a little in the GAM and a negligible amount in the higher m sidebands.

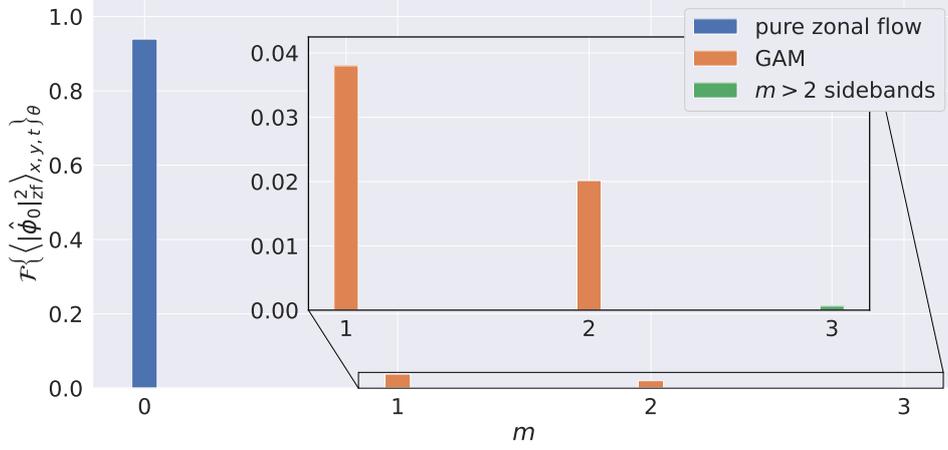


Figure 50: Fourier transform of the zonal flow poloidal structure

The net density transfer is effectively zero at all θ at the scale shown. Zooming in on the poloidal structure of the density transfer (not shown) reveals that $T_n \sim \mathcal{O}(10^{-5})$ at all θ with an apparent increase towards large $|\theta|$. Comparing this to the T_n values in Figures 41 (inverse cascade), 44 (spectra of density transfer into zonal flows) and 45 (turbulence and transfer spectra) that show the density transfer from individual k'_x modes is $\sim \mathcal{O}(0.1)$, suggests that there may be significant cancellation of transfer into and out of the zonal component of the density fluctuations. Furthermore, comparing Figure 43 (spectra of velocity transfer into zonal flows) with Figure 44 and Figure 45 (a) with Figure 45 (b), notice that there is more negative transfer (i.e. transfer out of zonal flows) in the density transfer than there is in the velocity transfer. Indeed, further examination (not shown) of $T_n(k'_x)$ for selected individual values or sums over the other dimensions (k'_y , k_x and θ), confirms that the net density transfer $\sim \mathcal{O}(10^{-5})$ is a consequence of the transfer into the zonal component of the density fluctuations at low $|k'_x|$ being cancelled out by the transfer out of the zonal component of the density fluctuations at moderate $|k'_x|$. Furthermore, the apparent increase in T_n at large $|\theta|$ may be explained as follows. The turbulent activity level decreases

with increasing $|k'_x|$ and with increasing $|\theta|$. Meanwhile, there is some noise in $T(k'_x)$ at high $|k'_x|$ due to the noise in the ϕ and n data at these scales. As the turbulent activity at low and moderate $|k'_x|$ decreases with increasing $|\theta|$, this allows the noise at large $|k'_x|$ to increase. This increases the transfer from this part of the spectrum, which results in an apparent increase in the net transfer at large $|\theta|$. However, this apparent poloidal structure should be considered as numerical error and the net density transfer should be taken as effectively zero at all θ , as shown in Figure 49. Finally, while T_u and T_n have different normalised units ($\phi_{\text{ref}}^3/(B_{\text{ref}}^3\rho_{\text{ref}}^4)$ cf. $n_{\text{ref}}^2\phi_{\text{ref}}/(B_{\text{ref}}\rho_{\text{ref}}^2)$ respectively) so a direct comparison is not strictly correct, these are normalised units so one would expect values $\sim \mathcal{O}(1)$ to represent significant transfer. Therefore, a difference of 5 orders of magnitude is sufficient to consider the density transfer as negligible compared to the velocity transfer.

The reason why there is cancellation of the density transfer to a negligible level is unclear. However, one possible reason is the use of adiabatic electrons in the simulations, which is known to prevent particle flux. It may be that adiabatic electrons are, in a similar way, preventing net density transfer. This was not tested because including kinetic electrons in the simulations requires a much larger computational grid in k_x and k_y , hence significantly increased computational cost.

Figure 51 shows data from the zonal case equivalent to that of Figure 49 (which was for the turbulent case). This data was time averaged over the entire quasi-periodic phase of the simulation, $t \gtrsim 2000L_{\text{ref}}/v_{\text{th}}$. Alternative time averaging windows and instantaneous poloidal structures (not shown) exhibit similar results.

Figure 51 exhibits many of the same features as Figure 49 including: the poloidal structure of the rate of transfer into the zonal flows is strongly correlated with the poloidal structure of the turbulent activity level but

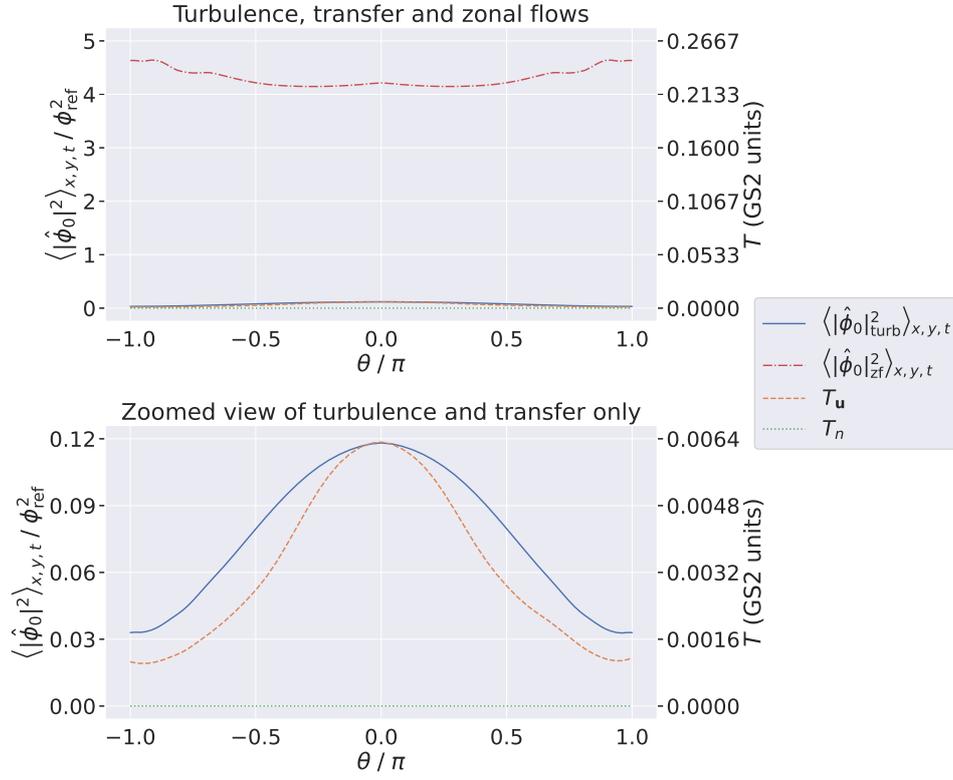


Figure 51: As Figure 49 but for the zonal case. All lines have been averaged over the entire quasi-periodic phase of the simulation, $t \gtrsim 2000L_{\text{ref}}/v_{\text{th}}$.

they are not proportional; the poloidal distribution of zonal flow energy is dominated by the $m = 0$ pure zonal flow with a little activity in the $m = 1$ and $m = 2$ sidebands (the GAM) and negligible energy in the $m > 2$ sidebands (confirmed by a Fourier transform in the θ direction - not shown); and the density transfer is effectively zero at all θ . This shows that these observations hold across the two cases examined, which suggest they may hold for other cases too, thus supporting experimental searches for zonal flows.

6.4 Conclusions

This chapter has presented development of a method to measure the non-linear transfer of energy between turbulence and zonal flows. This has been successfully applied to nonlinear gyrokinetic simulations of ITG turbulence to reveal, for the first time, the poloidal structure of zonal flow drive.

The results confirm the existence of a number of previously known phenomena, including the Dimits shift, and a dual cascade (forward in velocity, inverse in density). The predator-prey relationship between zonal flows and turbulence respectively has been observed previously in simplified analytical models [18] and Limit Cycle Oscillations in experiments (e.g. [88, 89]), and has been fitted to gyrokinetic simulations of turbulence and zonal flows [89]. However, this work (Figure 48) goes beyond showing that the zonal energy and turbulent intensity (predator and prey populations) rise and fall in a synchronised fashion, and demonstrates directly for the first time that the mechanism driving these predatory-prey dynamics is the non-linear transfer of energy from turbulence to zonal flows (in effect, observing the predators catching their prey for the first time). This may be related to previously reported gyrokinetic observations of turbulence suppression following a momentary peak in the rate of Reynolds work [90]. In addition, while the dual cascade seen in Figures 40 and 41 is in qualitative agreement with previous work [33], the observation of transfer from a broad range of turbulence directly into zonal flows (Figure 42) has not been shown previously.

In the cases studied, the density transfer is effectively zero due to cancellation of transfer into zonal flows at low $|k'_x|$ with transfer out of zonal flows at moderate $|k'_x|$. The exact reason why there is such cancellation is unclear but one possible reason is that the simulations modelled electrons adiabatically, which is known to result in zero particle flux. Future work may wish

to address the more computationally expensive case of including kinetic electrons to see if that allows for significant density transfer. Furthermore, this observation of cancellation may be an instance of the so-called “fictitious” interactions cited by Maeyama *et al.* [91] as a potential issue with the non-symmetrized style of transfer functions used here, whereby energy is circulated amongst the three waves involved in a triad coupling. Further analysis of the transfer functions calculated here may reveal that this is indeed the case. Future work may wish to investigate whether this analysis could be improved by using the recently developed symmetrized and/or directional transfer functions presented in [91].

In addition to investigating the effect of kinetic electrons, future work may wish to test other parameters to progress towards more realistic simulations, including plasma shaping, up-down asymmetry, electromagnetic effects, etc. At the same time, it will also be important to apply this method to experimental data in the search for zonal flows. This will require experimental measurements of the velocity field using techniques such as Cross-Correlation Time Delay Estimation [92] and Dynamic Time Warping [93]. The ultimate aim of these two aspects of future studies is to work towards comparison of experiment and simulation through the transfer function method to identify experimental zonal flows.

To facilitate the above future work, two specific improvements to the method developed here are recommended. The first is to address performance issues in the transfer function calculation code. Due to compatibility issues between the selective loading and parallelisation modules employed in the software implementation [4] of the transfer function calculation (`xarray` and `multiprocessing` respectively), full arrays of ϕ and n still had to be loaded on all processors before dispatching parallel tasks. As a result, the program had to be run on a computer with a large amount of memory available. This could be improved by only loading the data required for each pro-

cessor into the corresponding memory. This would require modification of xarray and/or multiprocessing, or a different method of selective loading and/or parallelisation, e.g. porting the code to a high performance language such as C++. The second specific recommendation is to derive the conversion from GS2 units to physical units such that the velocity and density transfer rates can be compared directly and so that simulation results can be compared with experimental data.

One further consideration for future application of this work is as follows. While this method has been developed to aid the experimental search for zonal flows, it may have another application. If zonal flows can be identified by other means, this technique could be used as a diagnostic for turbulence without having to run expensive non-linear gyrokinetic simulations. For such an application, the zonal flows could be measured experimentally, then the method presented in this chapter could be used to work out the turbulence that must have existed to generate those zonal flows and thus infer the turbulence without direct simulation, although it is unclear at this stage whether there would be a unique solution.

Finally, an additional avenue for future research could be the following. The transfer functions employed in this chapter constitute a fluid model of nonlinear energy transfer since they are based on the first two fluid moments of the distribution function, n and \mathbf{u} . While these transfer functions are used for the current investigation due to connection to existing diagnostics and previous work, it may be interesting for future studies to examine the nonlinear transfer of the spectral power of the distribution function itself – essentially a kinetic model of nonlinear energy transfer. This has been studied in the literature [94, 95], though not extensively. Such methods have the added complication that the distribution function is not necessarily advected by the fluid velocity but rather by the kinetic velocity, i.e. different parts of the distribution function are advected at different velocities. As a

result, the calculation of the kinetic model transfer function is likely to be quite similar to the calculation of the nonlinear term in gyrokinetic codes such as GS2, which may provide a useful starting point for such work.

Most importantly, this analysis has, for the first time, revealed the poloidal structure of zonal flow drive. For the velocity transfer, this is strongly correlated with the turbulent activity level. However, the rate of transfer into zonal flows is not proportional to the turbulent activity level. This is an important result for zonal flow experiments to relate the small measured region to the plasma behaviour outside that region. The correlation is due to the dependence of the transfer function on the magnitude of the fluctuating fields, while the reason for the lack of proportionality is unclear. A hypothesis related to the shifting of the turbulence spectrum due to magnetic shear with increasing $|\theta|$ has been offered but further work beyond the scope of this thesis is required to investigate this. Nonetheless, this shows that zonal flow drive is not constant in the poloidal direction, and while it is correlated with the turbulent activity level it is not proportional to it. These findings may provide an explanation for the remaining discrepancy between experimental nonlinear transfer and turbulence loss measurements reported in [24]. Furthermore, the method presented here may provide a route to more detailed analysis of the discrepancy reported in [24] and similar experimental observations by facilitating quantitative measurement of the poloidal structure of nonlinear transfer in such cases and, thus, quantitative comparison with the GPI or other diagnostic measurements and validation of nonlinear energy transfer models.

Part IV

Conclusions

Chapter 7

Conclusions

This thesis has investigated global effects on micro-instabilities and the poloidal structure of zonal flow drive. This was done using the local gyrokinetic simulation code GS2.

Regarding global effects, the aim was to study global KBMs in the pedestal. As a first step, a JET pedestal equilibrium was simulated using local gyrokinetics. This revealed KBMs that were narrow in ballooning angle, η_0 , which suggests a global mode that is radially wide and thus potentially in violation of the local approximation. The narrowing was shown to be caused by high magnetic shear, \hat{s} , which was in turn shown to exert its narrowing effect on both IBMs and KBMs via ideal ballooning physics. This leads to a splitting of the IBM instability region in η_0 , and hence a modification to the IBM s - α stability diagram, which in this case affects the second stability region. This highlights the importance of considering $\eta_0 \neq 0$ in local simulations, whether gyrokinetic, MHD or otherwise. While IBMs exhibit splitting, it was shown that the KBM in such cases may be subject to additional kinetic effects, thus highlighting the importance of considering kinetic effects

when using IBMs as a proxy for KBMs, which may affect the accuracy of EPED-like models.

Analysis of the local approximation revealed validity conditions that require the local frequency and growth rate to be slowly varying with respect to minor radius, x , and ballooning angle, η_0 , compared to the rational surface spacing, $\Delta = \frac{1}{nq}$. Re-formulating the η_0 condition in terms of k_{x0} (the radial wavenumber at $\eta = 0$, where η is the ballooning-space coordinate) reveals a different but equivalent condition. Combined with the fact that the KBM simulations discussed above showed no variation of the width in k_{x0} when \hat{s} was increased, this suggests that the impact on local validity of \hat{s} narrowing in η_0 may be mitigated by the corresponding increase in q' . This effect was demonstrated through application of the local-global method to artificial but representative data, which confirmed that local accuracy is reduced by narrowing in η_0 , but that this is mitigated by a corresponding increase in q' in the absence of other effects such as the enhancement of narrowing by the drifts. This highlights the importance of considering global effects in local simulations, especially in regions of high \hat{s} , which can be low but is often high in the pedestal.

Next, the local-global method was applied to a pedestal-like equilibrium. This was compared favourably to global gyrokinetic and MHD simulations from previous work [17], which suggests that the local-global method is valid in this case down to surprisingly low toroidal mode numbers of $3 \leq n \leq 12$ depending on the normalised Larmor radius. Simple corrective models were developed to capture global and diamagnetic effects, and these were shown to give a good approximation to the local-global results. The utility of these models was demonstrated in an EPED-like calculation, which suggests that global and diamagnetic corrections can potentially have a significant impact on the ELM prediction of EPED-like models, depending on the peeling-ballooning constraint. Therefore, future EPED-like calculations

should consider the impact of global and diamagnetic effects, which can be done using the method presented in this thesis. Future work may wish to develop a more general model for these corrections, where the fit parameters that were found empirically here are derived from first principles if possible, or are otherwise measured from a broader range of equilibria and reduced to a response surface for fast correction.

This thesis has also developed a method to measure the nonlinear transfer of energy from turbulence to zonal flows, which has been successfully applied to nonlinear gyrokinetic zonal flow simulations. This has, for the first time, demonstrated the mechanism underlying the predator-prey model (namely, the nonlinear transfer of energy from turbulence to zonal flows) through observations of a limit cycle oscillation type response in the zonal flow dominated (marginally turbulent) case. Also for the first time, this has shown clear and direct evidence of zonal flows being driven by the direct transfer of energy from a broad range of turbulence. Finally, this has also revealed, for the first time, the poloidal structure of zonal flow drive, which is peaked on the outboard side and is correlated with but not proportional to the turbulent activity level. This is an important result for zonal flow experiments to relate the small measured region to the plasma behaviour outside that region. Future work may wish to apply the transfer function method from this thesis to simulations using experimental equilibria and correlate with experimental measurements provide observational evidence of zonal flows.

Appendices

Appendix A

The Fourier-Ballooning Representation

To derive the Fourier-ballooning representation, we first substitute the ballooning representation (Equation (18)) into the ballooning transform (Equation (15)) to obtain:

$$\phi(x, \theta) = \sum_m \int_{-\infty}^{+\infty} \hat{\phi}_0(\eta; x) \exp[im(\theta - \eta)] \exp[-inq'x(\eta - \eta_0)] d\eta \quad (102)$$

Next, for convenience, we assume that the x and η dependencies of $\hat{\phi}_0$ are separable³¹, i.e. $\hat{\phi}_0(\eta; x) \rightarrow \alpha(x)\hat{\phi}_0(\eta)$. Then, taking the factors independent of η and m outside the integral and summation gives:

$$\phi(x, \theta) = \alpha(x) \exp[inq'x\eta_0] \sum_m \int_{-\infty}^{+\infty} \hat{\phi}_0(\eta) \exp[im(\theta - \eta)] \exp[-inq'x\eta] d\eta \quad (103)$$

³¹This assumption is valid since these dependencies vary on different length scales

The summation and integral is of the same form as Equation (15), so Equation (103) can be written as:

$$\phi(x, \theta) = \alpha(x) \hat{\phi}_0(\theta) \exp[-in q' x (\theta - \eta_0)] \quad (104)$$

Next, we represent $\alpha(x)$ by its Fourier transform:

$$\alpha(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{A}(k_{x0}) \exp[ik_{x0}x] dk_{x0} \quad (105)$$

where k_{x0} is the radial wavenumber at $\eta = 0$ given by $k_{x0} = -nq'\eta_0$ (Equation (34)), which allows Equation (105) to be written in terms of η_0 as follows:

$$\alpha(x) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \hat{A}(-nq'\eta_0) \exp[inq'x\eta_0] nq' d\eta_0 \quad (106)$$

For convenience, we re-define $\hat{A}(-nq'\eta_0)$ as just a function of η_0 (since $-n$ and q' on the reference surface are constants), absorb the $\frac{1}{2\pi}$ and nq' factors into that definition, and split the result into two separate functions so that one of them may be prescribed, i.e. $\frac{1}{2\pi} nq' \hat{A}(-nq'\eta_0) \rightarrow \check{A}(\eta_0) \rightarrow A(\eta_0)B(\eta_0)$, to write:

$$\alpha(x) = \int_{-\infty}^{+\infty} A(\eta_0)B(\eta_0) \exp[inq'x\eta_0] d\eta_0 \quad (107)$$

For convenience and consistency with previous work [39, 55], we prescribe $B(\eta_0) = \exp[-in(q_0\theta + q'x\eta_0)]$, where $q_0 = q(x = 0)$, so that we have the Fourier-ballooning representation (Equation (24)):

$$\phi(x, \theta) = \int_{-\infty}^{+\infty} \hat{\phi}_0(\theta; x, \eta_0) \exp[-in((q_0 + q'x)\theta - q'x\eta_0)] A(\eta_0) d\eta_0$$

Appendix B

Linear Convergence

The JET equilibrium used in Chapter 4 was tested for numerical convergence and was found to be suitably converged. Convergence test results for timestep (`de1t`), energy grid resolution (`negrid`), pitch angle grid resolution (`ngauss`), and parallel grid resolution (`ntheta`) are shown in Figures 52 to 55 respectively. These all show clear convergence. The test of parallel grid extent (`nperiod`) was less clearly converged, as shown in Figure 56, so the local mode structures, $\hat{\phi}_0(\eta)$, were also checked. These comfortably meet the requirement that $\hat{\phi}_0(\eta) \rightarrow 0$ as $|\eta| \rightarrow \infty$ with `nperiod` = 4, as shown in Figure 57, hence this value was used. The converged numerical parameters listed in Table 5 were used for the simulations reported in Chapter 4.

The CBC setup is based on that previously reported in [29], hence relies on the convergence tests reported there. The new Miller representation is compared to the previous s - α results as a benchmark in Figure 58. This shows that the results are suitably similar such that we can rely on the convergence tests reported previously. The differences are due to the different geometry models used and do not affect the conclusions of Chapter 4.

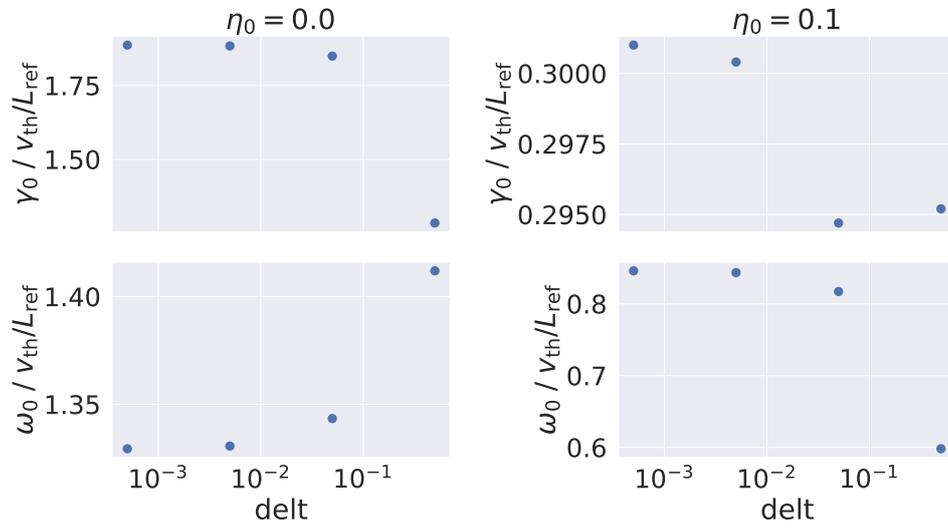


Figure 52: Convergence of the JET equilibrium with delt

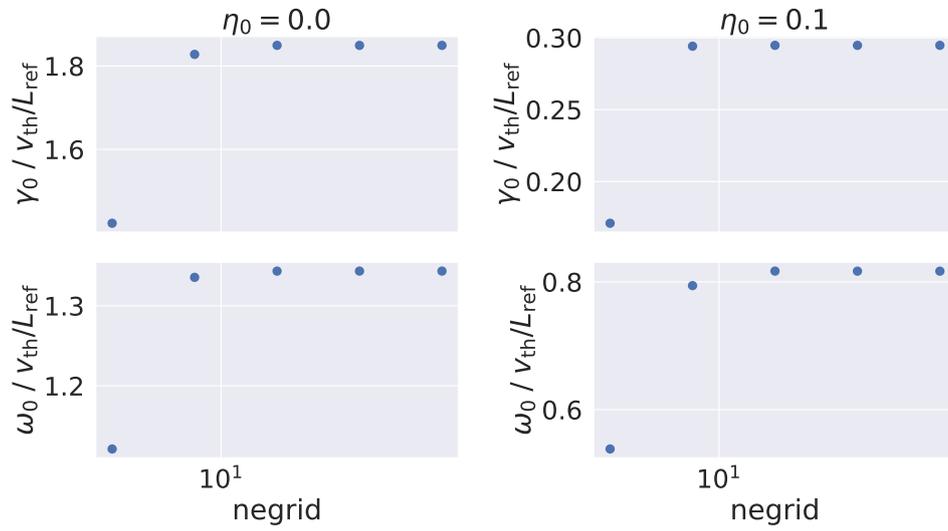


Figure 53: Convergence of the JET equilibrium with negrid

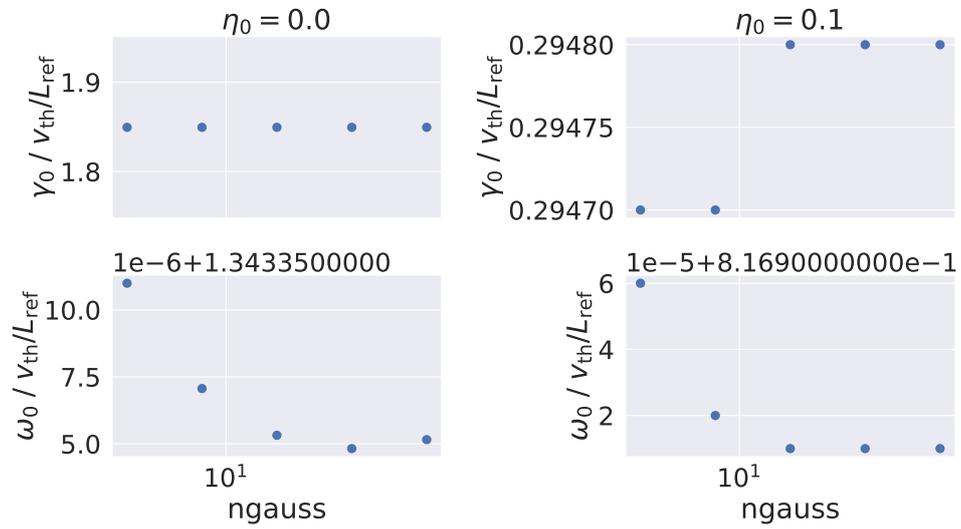


Figure 54: Convergence of the JET equilibrium with ngauss

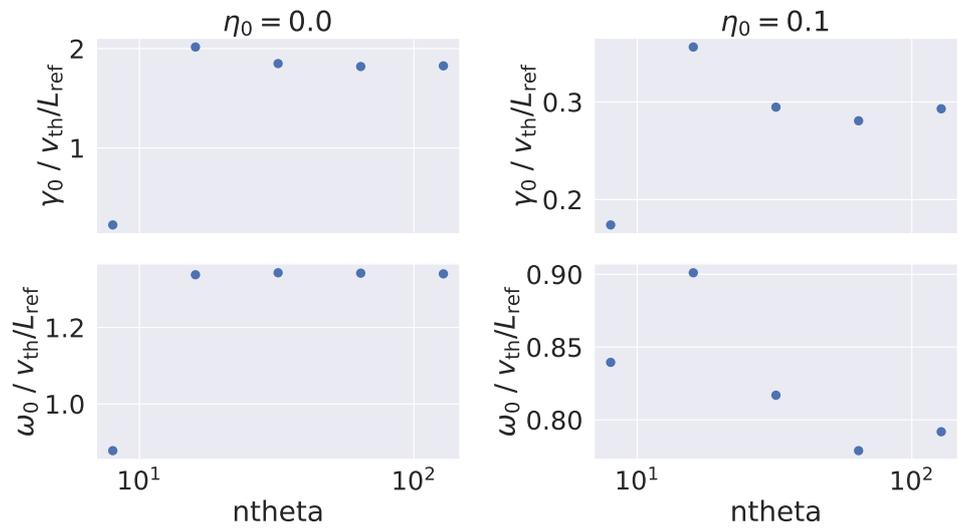


Figure 55: Convergence of the JET equilibrium with ntheta

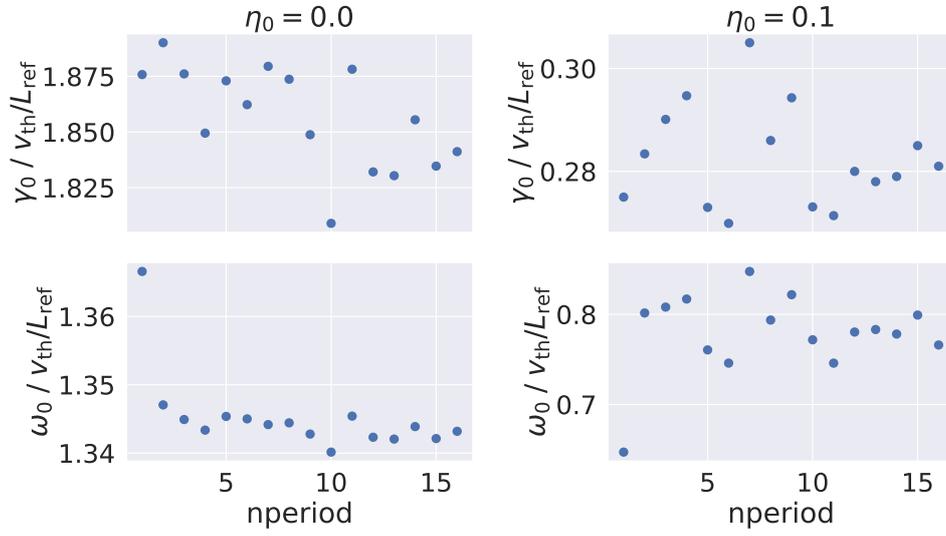


Figure 56: Convergence of the JET equilibrium with $nperiod$

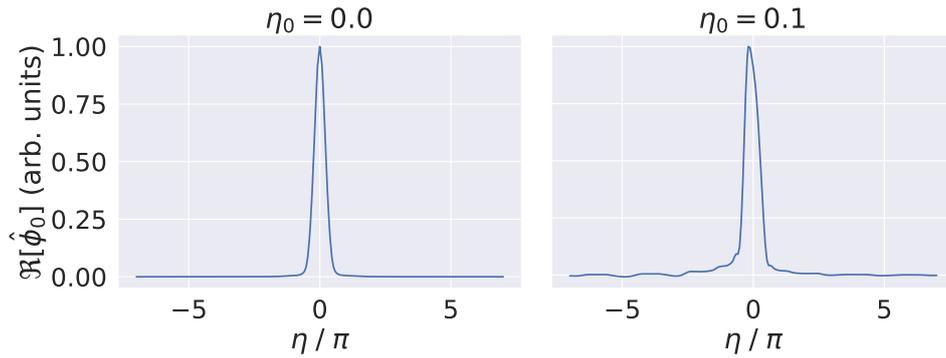


Figure 57: JET equilibrium local mode structures with $nperiod = 4$

Parameter	Value
delt	0.005
negridd	16
ngauss	8
ntheta	32
nperiod	4

Table 5: Converged numerical parameters used for the JET simulations in Chapter 4

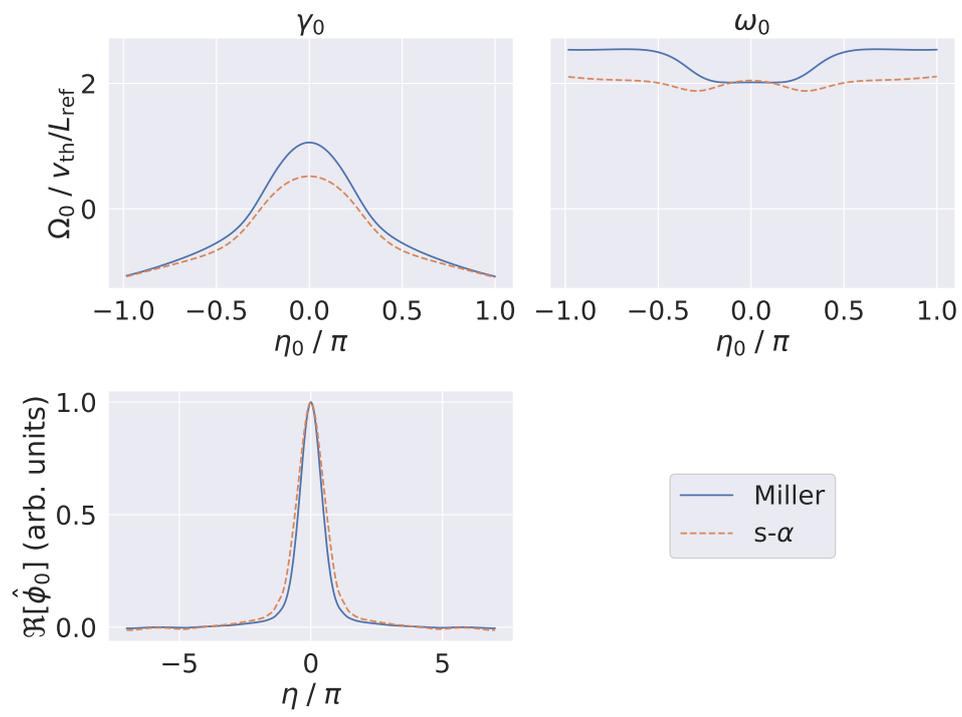


Figure 58: Comparison of CBC Miller and $s-\alpha$ results

Appendix C

Measuring Frequency and Growth Rate, and Tracking Sub-dominant Modes with GS2

The GS2 initial value (IV) solver reports (amongst other data) the frequency (ω_0), growth rate (γ_0) and mode structure ($\hat{\phi}_0(\eta)$) of the most unstable mode. However, more reliable ω_0 and γ_0 results are obtained using the method described in [29], which is, in brief:

- γ_0 is obtained by fitting a straight line to the natural logarithm of the flux tube volume averaged squared magnitude of the electrostatic potential fluctuation, $\ln \left[\left\langle |\hat{\phi}_0|^2 \right\rangle_V \right]$, as a function of time, t , and dividing the gradient by two since $\hat{\phi}_0 \sim \exp[i\Omega_0 t]$ where $\Omega_0 = \omega_0 + i\gamma_0$.
- ω_0 is obtained by fitting a constant to the final 20% of the GS2 $\omega_0(t)$ data.

The GS2 eigensolver is used to obtain ω_0 , γ_0 and $\hat{\phi}_0(\eta)$ of sub-dominant modes. Eigensolver results use the reported ω_0 and γ_0 data directly.

The CBC setup reported in Chapter 4 was run using the GS2 eigensolver small step method described in [29], which is, in brief:

- Run the IV solver at $\eta_0 = 0$ to get a preliminary measure of the frequency and growth rate.
- Re-run $\eta_0 = 0$ using the eigensolver to get the eigensolver ω_0, γ_0 and $\hat{\phi}_0(\eta)$.
- Run the next simulation with a small step in $\eta_0 \sim \pi/60$ using the previously obtained ω_0 and γ_0 as a starting point in the search for a solution, since the actual solution will be similar due to the small step in η_0 . This offers improved performance compared to taking large steps in η_0 thanks to the eigensolver starting near to the solution and thus finding the actual solution very quickly.

In addition, the following improvements to the above method were developed as part of the work presented in this thesis:

- Rather than using the previous ω_0 and γ_0 , the previous two values were used (from the third point onward) to perform a linear extrapolation to obtain more accurate starting values, as this was found to improve the tracking of sub-dominant modes. Extrapolation past marginal stability with a larger step was also tested to improve tracking of sub-dominant modes near marginal stability, although this was only sometimes beneficial so was only used as required.
- As well as starting from extrapolated ω_0 and γ_0 values, the eigensolver runs were also started from the previous mode structure to further improve run-times.
- The previous version of the method was also coupled in the radial direction. However, this was found to be unnecessary, so was eliminated to improve parallelisability.

Appendix D

Nonlinear Convergence

Figure 59 shows the local linear frequency and growth rate as functions of binormal wavenumber, and the mode structure of the most unstable mode. These results are consistent with ITGs and are in qualitative agreement with previously published work [28].

Figure 60 shows that the effect of the adiabatic electron model is significant, resulting in a change in growth rate by a factor ~ 2 . However, running nonlinear simulations with kinetic ions *and* kinetic electrons is computationally expensive. Therefore, as a simple first case, the simulations presented in Chapter 6 use the adiabatic electron model.

Figure 61 shows that the nonlinear setup is converged with respect to the baseline values of `delt`, `ngrid` and `ngauss`, while `ntheta = 64` was required for numerical stability.

Figure 62 shows that the number of points in the k_x direction, `nx = 64`, was converged, but in the k_y direction, `ny = 128` was required for numerical stability.

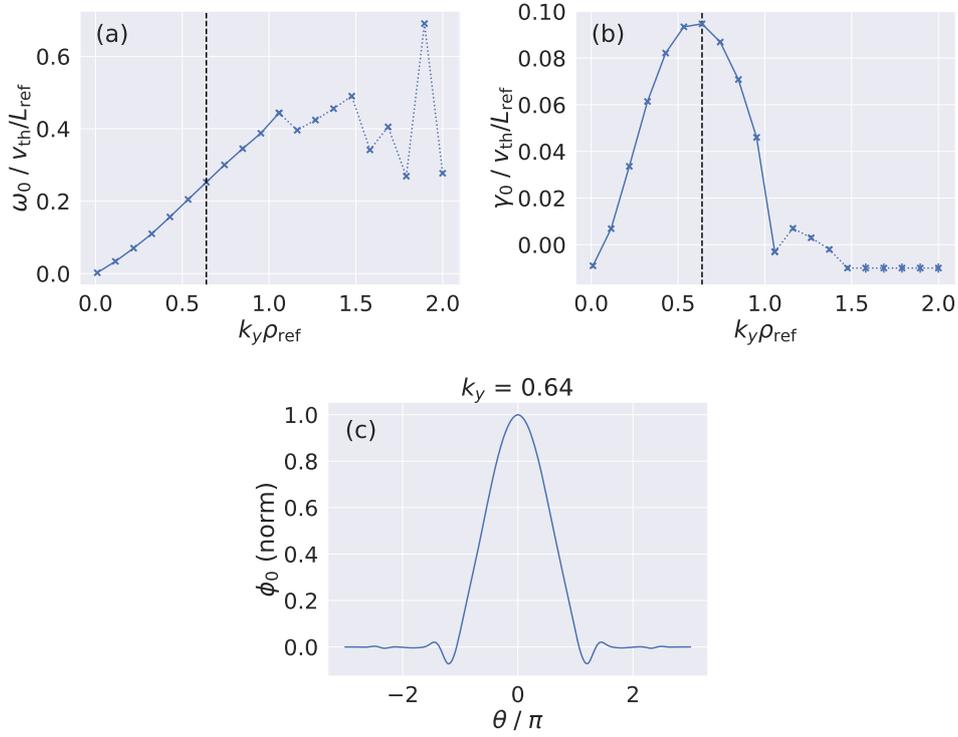


Figure 59: Local linear (a) frequency and (b) growth rate as functions of binormal wavenumber, k_y , and (c) mode structure of the most unstable mode (corresponding to the vertical line in (a) and (b)). The solid line indicates where ITGs have been captured reliably, while the dotted line indicates where the simulations have switched to a different mode.

The `cf1` parameter controls when GS2 will change the timestep to avoid violating the Courant-Friedrichs-Lewy (CFL) condition, with a higher `cf1` value resulting in a larger safety margin. `cf1 = 1.5` exhibited improved numerical stability and computational performance over the baseline value of 0.1, so 1.5 was used for the nonlinear simulations presented in Chapter 6.

Figure 63 shows that the nonlinear results are statistically insensitive to collisions. Furthermore, this test showed that the collisionless case was faster (computationally) than the collisional case by a factor of ~ 2.3 . Therefore, the nonlinear simulations reported in Chapter 6 are collisionless.

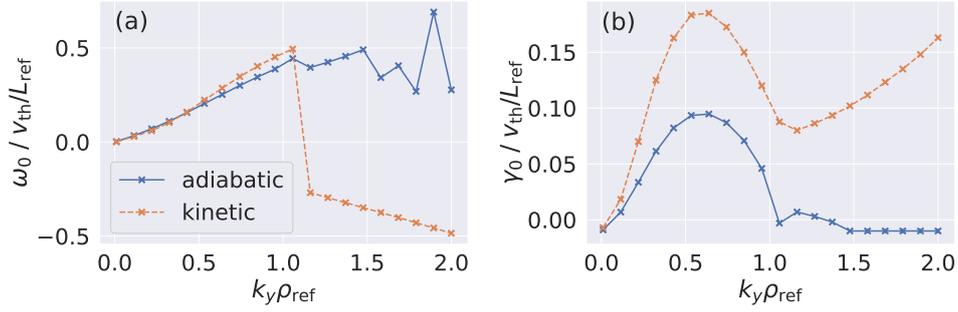


Figure 60: Local linear (a) frequency and (b) growth rate as functions of binormal wavenumber, k_y , for adiabatic and kinetic electron models

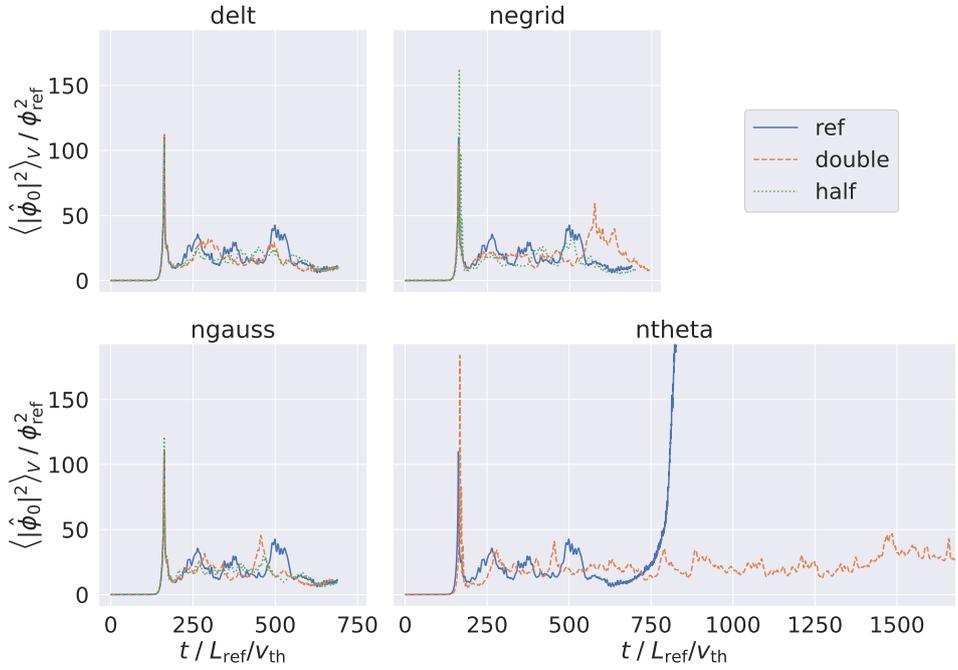


Figure 61: Nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t , for various values of the GS2 numerical parameters

Figure 64 shows that the nonlinear results are statistically insensitive to the hyper-diffusion coefficient, D , for the values tested. Furthermore, this test showed that the $D = 10^{-2}$ case was marginally faster (speedup of \sim

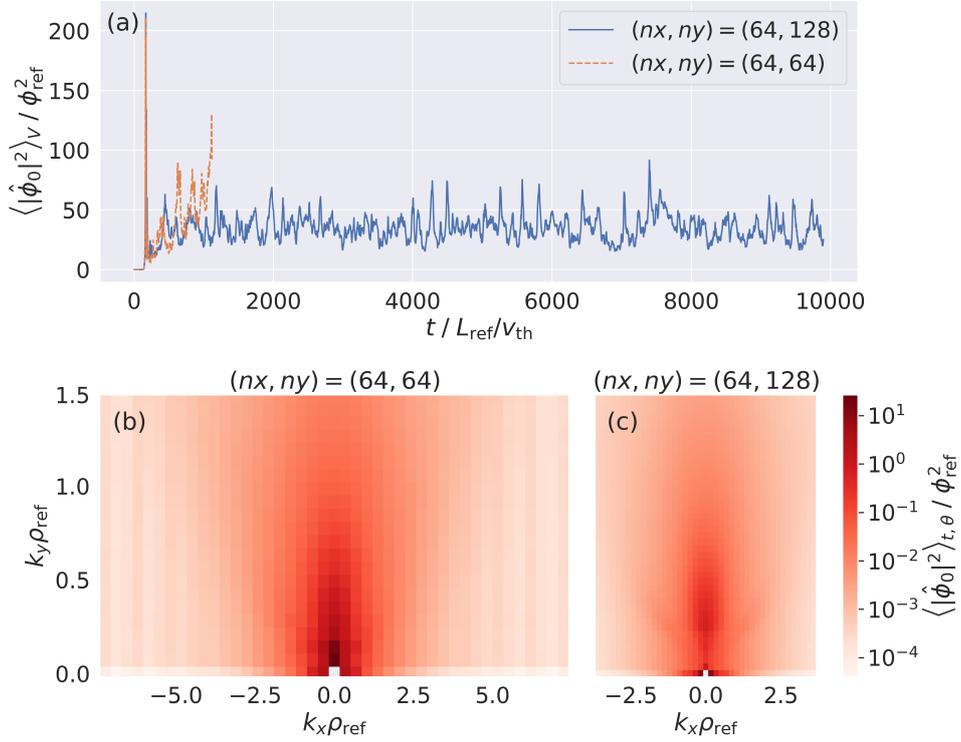


Figure 62: (a) Nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t , for two different wavenumber grid resolutions; and (b) and (c) the corresponding spectral energy density spectra averaged over $t \gtrsim 200L_{\text{ref}}/v_{\text{th}}$ and all θ (1.03 compared to the other cases). Therefore, $D = 10^{-2}$ was used for the nonlinear simulations reported in Chapter 6.

Figure 65 shows that the nonlinear results are converged with respect to wavenumber range for $k_{y, \text{max}} \rho_{\text{ref}} = 1.5$ and 2.0 but not for $k_{y, \text{max}} \rho_{\text{ref}} = 1.0$. Since $k_{y, \text{max}} \rho_{\text{ref}} = 1.5$ also maximises k_y resolution for a given number of k_y grid points compared to $k_{y, \text{max}} \rho_{\text{ref}} = 2.0$, the range used for the nonlinear simulations presented in Chapter 6 is $0.0 \leq k_y \rho_{\text{ref}} \leq 1.5$.

The final set of converged parameters are listed in Table 6.

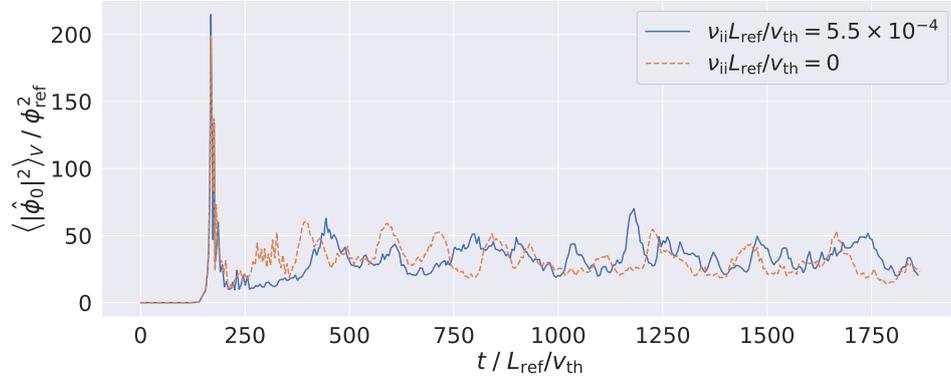


Figure 63: Nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t , for collisionless and collisional nonlinear cases

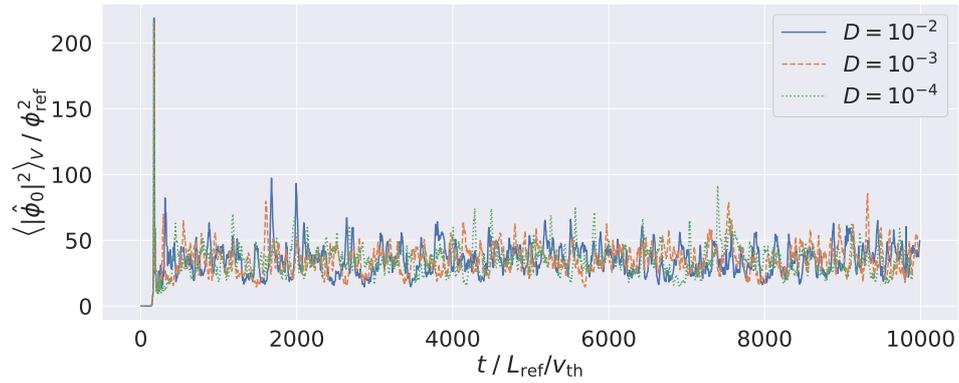


Figure 64: Nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t , for simulations with various values of hyper-diffusion coefficient, D

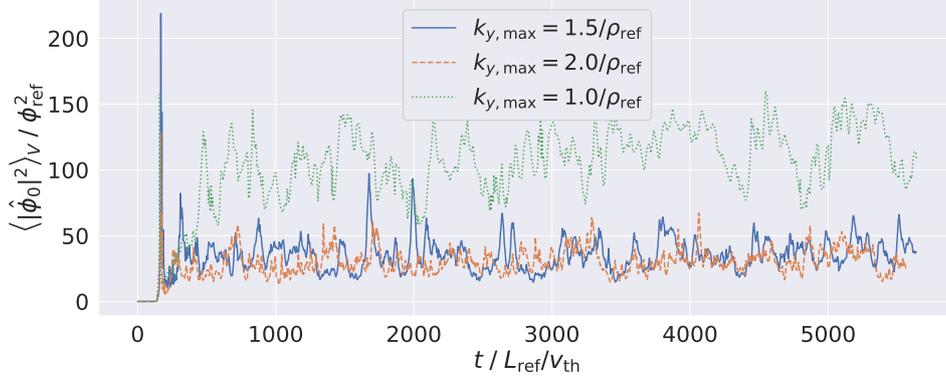


Figure 65: Nonlinear volume-averaged spectral energy density of the electrostatic potential fluctuation field, $\langle |\hat{\phi}_0|^2 \rangle_V$, as a function of time, t , for simulations with various k_y ranges (for the same number of k_y points, hence various k_y resolutions)

Parameter	Reference value for convergence tests	Value used in nonlinear simulations
delt	0.2	0.2
nperiod	1	1
ngrid	16	16
ngauss	2	2
ntheta	32	64
nx	64	64
ny	64	128
$k_{y, \text{max}} \rho_{\text{ref}}$	1.0	1.5
cfl	0.1	1.5
d_hyper	10^{-4}	10^{-2}

Table 6: Numerical parameter values

Appendix E

Transfer Function Program

To address the computational challenges of evaluating NETFs, a program [4] was written in Python that includes the following features:

- Selective loading and use of function scope for memory performance
- Pre-calculation of leading k -factors to avoid repeated calculations
- Re-ordering of arrays to maximise cache performance
- Use of NumPy array broadcasting for performance
- Transfer function spectra time-averaging parallelised in time
- Transfer into zonal flows as a function of poloidal angle, θ , parallelised in θ (time-averaged or a function of t as required)
- Bi-coherence spectra or summed bi-coherence is also available (where bi-coherence is a measure of the phase coupling, obtained by effectively dividing the transfer function's triple product by its magnitude). This was investigated but the results did not reveal anything beyond that available from the transfer functions, so were not reported.

References

- [1] Barnes, M *et al.* 2021 “GS2”.
- [2] Barnes, M *et al.* 2021 “Gyrokinetics / Gs2 — Bitbucket” URL: <https://bitbucket.org/gyrokinetics/gs2/src/master/> (visited on 04/07/2021).
- [3] Biggs-Fox, SN 2020 “Gs2_py_utils” GitLab URL: https://gitlab.com/stephen-biggs-fox/gs2_py_utils (visited on 08/15/2022).
- [4] Biggs-Fox, SN 2021 “Zonal Transfer Functions” GitLab URL: <https://gitlab.com/stephen-biggs-fox/zonal-transfer-functions> (visited on 05/07/2021).
- [5] Cook, J *et al.* 2016 *Environmental Research Letters*, **11** (4) 048002.
- [6] Smith, JB *et al.* 2009 *Proceedings of the National Academy of Sciences*, **106** (11) 4133.
- [7] Intergovernmental Panel on Climate Change 2013 “Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change” *Cambridge University Press*.
- [8] International Energy Agency 2016 “World Energy Outlook 2016” *Organisation for Economic Co-operation and Development*.
- [9] Bustreo, C *et al.* 2019 *Fusion Engineering and Design*, **146** 2189.
- [10] Gibney, E 2022 *Nature*, **602** (7897) 371.

- [11] Wesson, J 2011 "Tokamaks" (4th ed.) *Oxford University Press, International Series of Monographs on Physics* (149).
- [12] Chen, FF 1984 "Introduction to Plasma Physics and Controlled Fusion" Vol. 1: Plasma Physics (2nd ed.) *Springer Science+Business Media, LLC*.
- [13] Wagner, F *et al.* 1982 *Physical Review Letters*, **49** (19) 1408.
- [14] Snyder, P *et al.* 2011 *Nuclear Fusion*, **51** (10) 103016.
- [15] Saarelma, S *et al.* 2017 *Plasma Physics and Controlled Fusion*, **60** (1) 014042.
- [16] Saarelma, S *et al.* 2017 *Plasma Physics and Controlled Fusion*, **59** (6) 064001.
- [17] Martin Collar, JP *et al.* 2020 *Plasma Physics and Controlled Fusion*, **62** (9) 095005.
- [18] Diamond, PH *et al.* 2005 *Plasma Physics and Controlled Fusion*, **47** (5) R35.
- [19] Kim, Ej and Diamond, PH 2003 *Physical Review Letters*, **90** (18) 185006.
- [20] Shats, MG *et al.* 2004 *Fusion Science and Technology*, **46** (2) 279.
- [21] Manz, P *et al.* 2012 *Physics of Plasmas*, **19** (7) 072311.
- [22] Cziegler, I *et al.* 2014 *Plasma Physics and Controlled Fusion*, **56** (7) 075013.
- [23] Schmitz, L *et al.* 2014 *Nuclear Fusion*, **54** (7) 073012.
- [24] Cziegler, I *et al.* 2015 *Nuclear Fusion*, **55** (8) 083007.
- [25] Hillesheim, JC *et al.* 2016 *Physical Review Letters*, **116** (6) 065002.
- [26] Tynan, GR *et al.* 2016 *Plasma Physics and Controlled Fusion*, **58** (4) 044003.
- [27] Schmitz, L 2017 *Nuclear Fusion*, **57** (2) 025003.
- [28] Abdoul, PA *et al.* 2015 *Plasma Physics and Controlled Fusion*, **57** (6) 065004.

- [29] Biggs, SN 2016 “Gyrokinetic Simulations of Kinetic Ballooning Modes” MSc Fusion Energy, Project Dissertation: *University of York*.
- [30] Abdoul, PA *et al.* 2018 *Plasma Physics and Controlled Fusion*, **60** (2) 025011.
- [31] Xia, H *et al.* 2006 *Physical Review Letters*, **97** (25) 255003.
- [32] Cziegler, I *et al.* 2013 *Physics of Plasmas*, **20** (5) 055904.
- [33] Holland, C *et al.* 2007 *Physics of Plasmas*, **14** (5) 056112.
- [34] Nakata, M *et al.* 2012 *Physics of Plasmas*, **19** (2) 022303.
- [35] Maeyama, S *et al.* 2015 *Physical Review Letters*, **114** (25) 255002.
- [36] Dimits, AM *et al.* 2000 *Physics of Plasmas*, **7** (3) 969.
- [37] Hazeltine, R and Meiss, J 2003 “Plasma Confinement” *Dover Publications, Dover Books on Physics*.
- [38] Dickinson, D 2012 “Effects of Profiles on Microinstabilities in Tokamaks” PhD thesis: *University of York*.
- [39] Abdoul, PA 2015 “Generalised 2D Ballooning Theory in Tokamaks: Obtaining the Global Properties of Linear Microinstabilities From the Local Ballooning Equations” PhD thesis: *University of York*.
- [40] Di Siena, A *et al.* 2018 *Physics of Plasmas*, **25** (4) 042304.
- [41] Fitzpatrick, R 2016 “Plasma Physics” *The University of Texas at Austin*.
- [42] Lin, Y *et al.* 2005 *Plasma Physics and Controlled Fusion*, **47** (4) 657.
- [43] Groelj, D *et al.* 2017 *The Astrophysical Journal*, **847** (1) 28.
- [44] Groelj, D *et al.* 2018 *Physical Review Letters*, **120** (10) 105101.
- [45] Kwon, JM *et al.* 2015 *Journal of Computational Physics*, **283** 518.
- [46] Grandgirard, V *et al.* 2016 *Computer Physics Communications*, **207** 35.
- [47] Mandell, NR *et al.* 2020 *Journal of Plasma Physics*, **86** (1) .

- [48] Jolliet, S *et al.* 2007 *Computer Physics Communications*, **177** (5) 409.
- [49] Görler, T *et al.* 2011 *Journal of Computational Physics*, **230** (18) 7053.
- [50] Mishchenko, A *et al.* 2015 *Nuclear Fusion*, **55** (5) 053006.
- [51] Hastie, RJ *et al.* 1979 *Nuclear Fusion*, **19** (9) 1223.
- [52] Connor, JW *et al.* 1978 *Physical Review Letters*, **40** (6) 396.
- [53] Connor, JW *et al.* 1979 *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **365** (1720) 1.
- [54] Riley, KF *et al.* 2004 “Mathematical Methods for Physics and Engineering: A Comprehensive Guide” *Cambridge University Press*.
- [55] Taylor, JB and Wilson, HR 1996 *Plasma Physics and Controlled Fusion*, **38** (11) 1999.
- [56] Highcock, E 2012 “The Zero-Turbulence Manifold in Fusion Plasmas” PhD thesis: *University of Oxford*.
- [57] Dickinson, D *et al.* 2011 *Plasma Physics and Controlled Fusion*, **53** (11) 115010.
- [58] Diallo, A *et al.* 2013 *Nuclear Fusion*, **53** (9) 093026.
- [59] Hassan, E *et al.* 2021 *Physics of Plasmas*, **28** (6) 062505.
- [60] Greene, JM and Chance, MS 1981 *Nuclear Fusion*, **21** (4) 453.
- [61] Hatch, DR *et al.* 2019 *Nuclear Fusion*, **59** (8) 086056.
- [62] Bowman, C *et al.* 2017 *Nuclear Fusion*, **58** (1) 016021.
- [63] Huysmans, GTA *et al.* 1991 *International Journal of Modern Physics C*, **2** (1) 371.
- [64] Koh, S *et al.* 2012 *Physics of Plasmas*, **19** (7) 072505.
- [65] Miller, RL *et al.* 1998 *Physics of Plasmas*, **5** (4) 973.
- [66] Peeters, AG 2000 *Plasma Physics and Controlled Fusion*, **42** (12B) B231.

- [67] Gormezano, C *et al.* 2007 *Nuclear Fusion*, **47** (6) S285.
- [68] Davies, B *et al.* 2022 *In Preparation*.
- [69] Aleynikova, K and Zocco, A 2017 *Physics of Plasmas*, **24** (9) 092106.
- [70] Lütjens, H *et al.* 1996 *Computer Physics Communications*, **97** (3) 219.
- [71] Kotschenreuther, M 1986 *The Physics of Fluids*, **29** (9) 2898.
- [72] Belli, EA and Candy, J 2010 *Physics of Plasmas*, **17** (11) 112314.
- [73] Graves, JP *et al.* 2019 *Plasma Physics and Controlled Fusion*, **61** (10) 104003.
- [74] Tang, WM *et al.* 1980 *Nuclear Fusion*, **20** (11) 1439.
- [75] Snyder, PB *et al.* 2004 *Plasma Physics and Controlled Fusion*, **46** (5A) A131.
- [76] Snyder, P *et al.* 2009 *Nuclear Fusion*, **49** (8) 085035.
- [77] Wan, W *et al.* 2012 *Physical Review Letters*, **109** (18) 185004.
- [78] Wan, W *et al.* 2013 *Physics of Plasmas*, **20** (5) 055902.
- [79] Fujisawa, A 2008 *Nuclear Fusion*, **49** (1) 013001.
- [80] Ritz, CP and Powers, EJ 1986 *Physica D: Nonlinear Phenomena*, **20** (2) 320.
- [81] Ritz, CP *et al.* 1989 *Physics of Fluids B: Plasma Physics*, **1** (1) 153.
- [82] Camargo, SJ *et al.* 1995 *Physics of Plasmas*, **2** (1) 48.
- [83] Kim, JS *et al.* 1996 *Physics of Plasmas*, **3** (11) 3998.
- [84] Kim, JS *et al.* 1997 *Physical Review Letters*, **79** (5) 841.
- [85] Holland, C *et al.* 2003 *Nuclear Fusion*, **43** (8) 761.
- [86] Xu, M *et al.* 2009 *Physics of Plasmas*, **16** (4) 042312.
- [87] Xu, M *et al.* 2012 *Physical Review Letters*, **108** (24) 245001.
- [88] Xu, Y *et al.* 2014 *Plasma Physics and Controlled Fusion*, **57** (1) 014028.

- [89] Kobayashi, S *et al.* 2015 *Physics of Plasmas*, **22** (9) 090702.
- [90] Chang, CS *et al.* 2017 *Physical Review Letters*, **118** (17) 175001.
- [91] Maeyama, S *et al.* 2021 *New Journal of Physics*, **23** (4) 043049.
- [92] Holland, C *et al.* 2004 *Review of Scientific Instruments*, **75** (10) 4278.
- [93] Kriete, DM *et al.* 2018 *Review of Scientific Instruments*, **89** (10) 10E107.
- [94] Plunk, GG and Tatsuno, T 2011 *Physical Review Letters*, **106** (16) 165003.
- [95] Plunk, GG *et al.* 2012 *New Journal of Physics*, **14** (10) 103030.