Stable Singularity of the Euler Equations on $\mathbb{R}^3$

Adarsh Ganeshram* AGANESHRAM@BERKELEY.EDU

Valentin Duruisseaux* VDURUISS@CALTECH.EDU

Anima Anandkumar ANIMA@CALTECH.EDU

Department of Mathematics, University of California Berkeley, Berkeley, CA, USA

Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA

Abstract

We provide evidence of a stable finite-time singularity in the 3D Euler equations on the unbounded domain. Using a physics-informed neural network (PINN) with a self-similar ansatz, we find an approximate singular profile for the Euler system at the critical blowup rate of 0.5 and certify it using a spline representation. The transport field associated with the obtained profile suggests that linear damping can be established throughout the domain, providing strong evidence for the overall stability of the candidate profile. We also establish a complete framework for proving nonlinear stability of the approximate self-similar profile, reducing the analysis to a large but finite collection of explicit estimates and computable constants.

Keywords: Euler equations | fluid dynamics | singularity formation | finite-time blowup | physics-informed neural networks | computer-assisted proofs

1 Introduction

A central problem in the analysis of partial differential equations (PDEs) is whether all solutions arising from smooth initial data remain regular at all times or whether they may lose regularity in finite time. Such a loss of regularity is known as *singularity formation*, or *blowup*. In particular, establishing whether solutions of the 3D incompressible Navier–Stokes equations can develop a finite-time blowup is among the most prominent open problems in fluid dynamics, and is one of the six unsolved Millennium problems [1] and the 15^{th} Smale problem [2]. Closely related is the question of finite-time singularity formation in the Euler equations, which govern inviscid flows and differ from the Navier–Stokes equations by the absence of viscous dissipation.

Blowups have been established in related simpler problems or under modified assumptions. For instance, finite-time blowups have been established for the Euler equations in the presence of a boundary given smooth initial conditions [3], or under non-smooth initial conditions on the unbounded domain $\mathbb{R}^3$ [4–7]. However, it is not clear how these additional assumptions can be removed. For instance, boundary-driven blowup mechanisms exploit the confinement imposed by solid walls to generate rapid vorticity growth [3; 8; 9], which does not work when there is no boundary. In addition, there are non-existence results which rule out finite-time blowups in the Euler system on $\mathbb{R}^3$ under certain conditions [10; 11], which include strong decay and integrability assumptions on the vorticity profile.

A common strategy for establishing finite-time blowup is to construct an approximate solution numerically, and then to use stability theory to prove the existence of a singular solution to the governing PDEs that remains sufficiently close to this approximation. A self-similar ansatz offers a tractable path to both the numerical search for such an approximate profile and its stability analysis. It allows us to continuously zoom in or out at the specified rate, as the solution evolves, so that its overall "shape" is unchanged over time. More precisely, a blowup solution is *self-similar* if, under a suitable dynamic rescaling, it remains close to an approximate blowup profile all the way up to the blowup time. The dynamic rescaling equations thus serve as a starting point for stability analysis around a given approximate profile.

Figure 1: **From numerical discovery to certified stability.** An approximate self-similar profile is first discovered using PINNs and converted to a spline for analytic estimates and rigorous certification. Perturbations around this profile are introduced together with a parametrized weighted energy that captures the relevant singular behavior of the solution. The stability theory is then developed by linearizing around the spline profile and deriving analytic estimates to control the linear and nonlinear contributions. The analytic estimates depend on explicit numerical constants and quantitative properties of the approximate profile, which are rigorously enclosed using interval arithmetic together with the profile data, stability constants, and PDE residuals. Finally, the parametrized singular weights and other parameters are optimized so that the certified damping dominates the other terms and PDE residuals. This yields rigorous computer-assisted nonlinear stability, with the analytic and certified estimates also suitable for subsequent formalization in Lean.

Since it is challenging to manually construct approximate self-similar profiles as blowup candidates, Physics-Informed Neural Networks (PINNs) [12] have recently been adopted as a more flexible alternative. When the PINN computations are carried in high precision (FP64), they can discover approximate singular profiles [13–15], while (approximately) enforcing the governing equations together with symmetries, normalization, decay and other physical constraints. They have been successful in discovering profiles for simplified versions of the Euler equations, such as Burgers and Boussinesq equations. However, PINN formulations have not yet succeeded in discovering potential singularities in the Euler and Navier–Stokes equations on the unbounded domain. It remains unclear how to choose suitable constraints and ansatz for such a PINN formulation, and how to overcome the resulting optimization challenges.

Recent theoretical results also impose additional constraints on the regime where such finite-time blowups may occur, further reducing flexibility for the selection of ansatz and PINN formulations. Constantin et al. [16] recently establish that the critical self-similar scaling exponent of $\lambda = 0.5$, known to be natural for Navier–Stokes equations, is also the most promising (or "mathematically distinguished") regime for self-similar Euler blowups. They essentially rule out stable self-similar singularities for the Euler system when $\lambda < 0.5$, and $\lambda > 0.5$ for the commonly used axis-symmetric ansatz considered here. Despite sharing this critical scaling exponent of 0.5 for potential blowups, the Euler and Navier–Stokes equations cannot share a common singular profile, since the effect of viscosity cannot be ignored in this regime. Overall, establishing the existence of singularities in either system is highly non-trivial.

In addition to finding an approximate profile, establishing stability is crucial but complicated. This becomes simpler when a global outgoing property is satisfied [17], which intuitively pushes the self-similar flow away from the singular core. This expansive effect of the global outgoing flow is also closely connected to the favorable higher-order damping that underlies the stability analysis. However, for the Euler system with a self-similar ansatz, Constantin et al. [16] provides strong theoretical evidence that nontrivial fixed points of the meridional flow are likely to exist and have been "well known to cause tremendous headaches" in other related problems. Given that the global outgoing property cannot hold for the Euler system with self-similar ansatz, they instead consider a weaker local property that locally pushes the self-similar flow away from these fixed points, preventing perturbations from concentrating there and providing a favorable structure for the stability analysis. Establishing stability with this weaker local outgoing property however requires a substantially more delicate analysis.

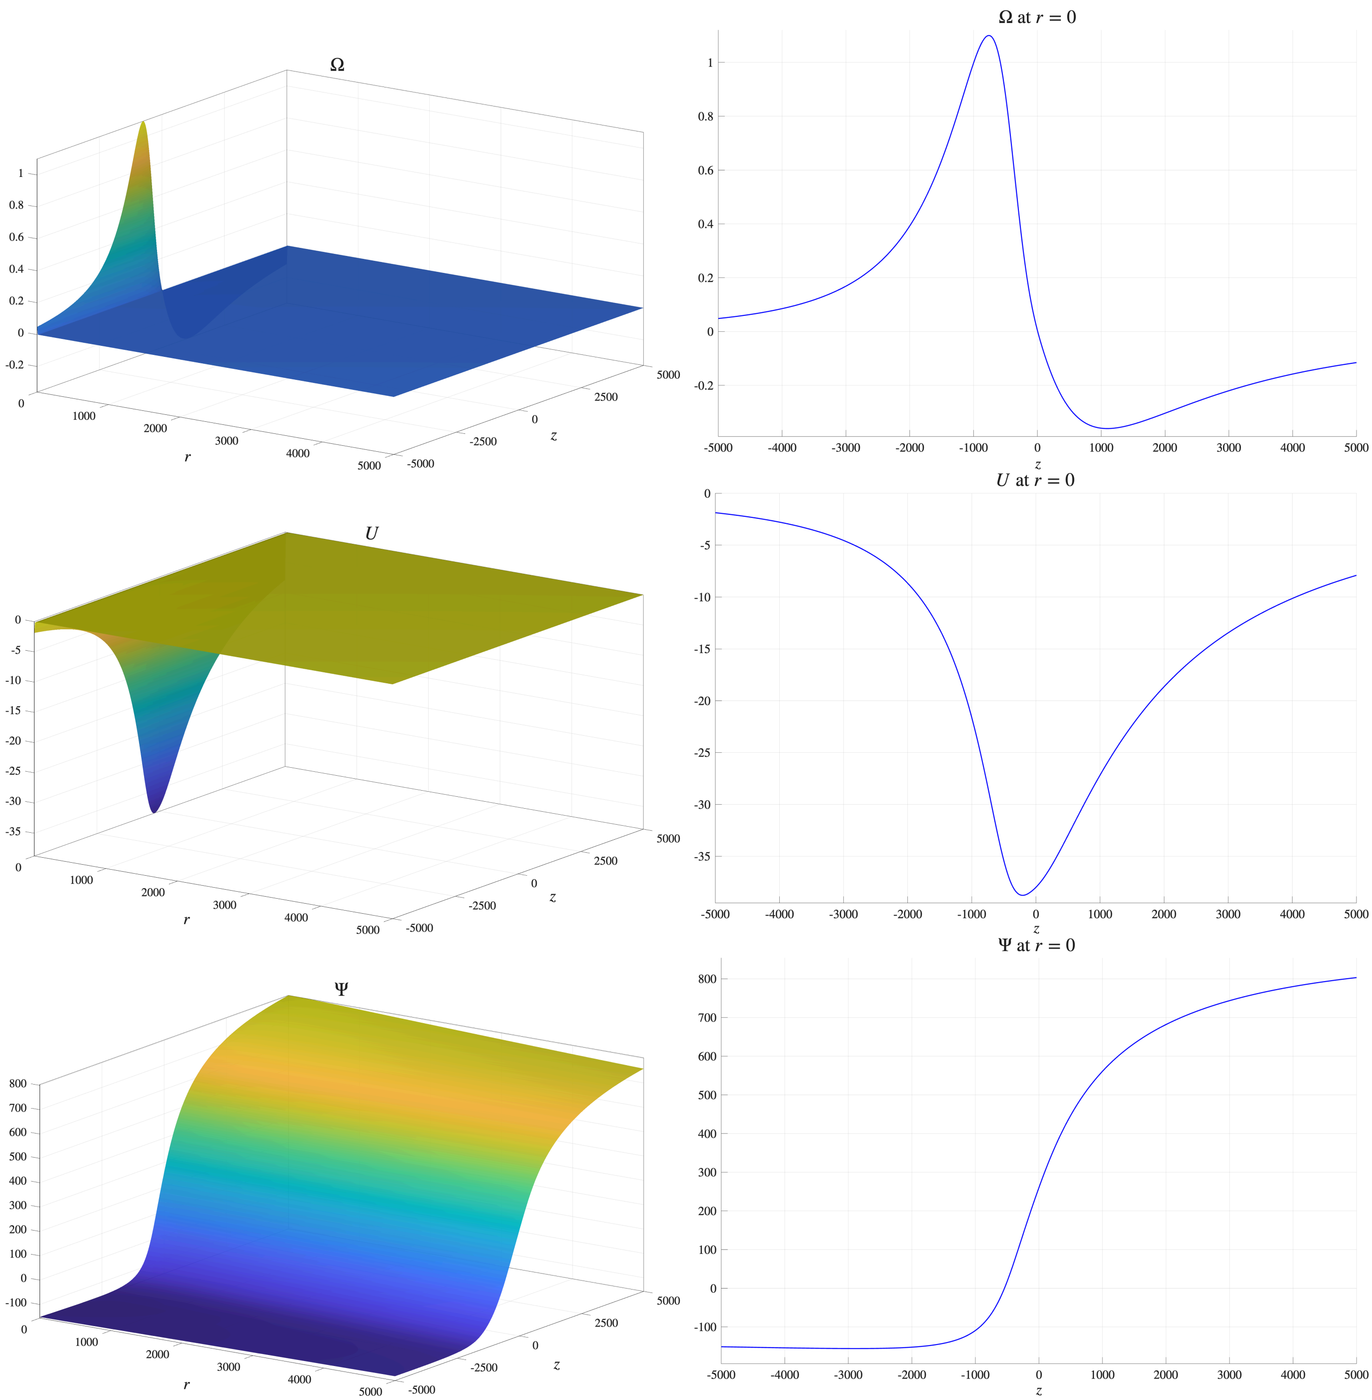

Figure 2: Visualization of the approximate profile Ω, U, Ψ as three-dimensional surface plots for the 3D axisymmetric Euler equations with their corresponding $r = 0$ cross-sections.

We provide the first evidence that the Euler equations admits a self-similar finite-time blowup on $\mathbb{R}^3$ at the critical scaling exponent $\lambda = 0.5$, referred to as "mathematically distinguished" by Constantin et al. [16]. We obtain certified linear damping for the Euler system, providing the key stabilizing mechanism. We also establish a complete proof framework for the nonlinear stability of the approximate blowup profiles in our companion stability paper [18], reducing the argument to a finite collection of explicit, rigorously verifiable estimates. The profile is made available on GitHub, together with the verification code demonstrating that the Euler equations are satisfied to high precision. Our overall methodology is shown in Figure 1.

Numerical Profiles. Using a PINN formulation, we obtain a highly accurate approximate self-similar profile for the Euler system. A novel combination of complementary tools is essential to achieving this level of accuracy, ensuring that the formulation captures the correct physics while converging to a high-precision solution. Imposing appropriate soft and hard constraints is necessary to prevent convergence to trivial or degenerate solutions, a failure mode observed with other PINN formulations.

Motivated by Hou et al. [19], we adopt a traveling-wave ansatz propagating along the z-axis in the axisymmetric setting. Although the traveling-wave and stationary formulations are equivalent in the Euler setting [16] since we can translate and scale the system appropriately, they induce substantially different optimization landscapes. In the traveling-wave formulation, the singularity propagation speed is treated as a free parameter, providing greater flexibility to the optimizer. Moreover, the traveling-wave ansatz does not impose the z-parity constraints that often accompany the stationary self-similar ansatz, and PINN optimization consistently fails to converge when these z-parity constraints are imposed. These features substantially alleviate the optimization difficulties encountered in previous PINN formulations and are essential to the success of our approach.

We employ optimization techniques to obtain a profile with high-accuracy. We use adaptive loss-weighting schedules to balance competing objectives, FP64 arithmetic to reduce round-off errors and improve gradient accuracy, PINN boosting [20; 21] to progressively correct residual errors, and adaptive collocation with residual-based resampling to concentrate training points in regions where the profile is most difficult to resolve. We also employ a sinh transformation to efficiently sample the unbounded domain. Finally, the choice of optimizer is also important, and we use efficient self-scaled curvature-aware optimizers. While Adam [22] typically plateaus at a profile-residual MSE of order 10^{-4}, the recently proposed SS-eSOAP [23] optimizer reaches the 10^{-6} regime with only a 15% increase in per-step cost. We finally employ the more computationally expensive SS-Broyden [24; 25] for high-accuracy fine-tuning, reducing the profile-residual MSE to order 10^{-10}.

The optimized PINN profile is converted into piecewise polynomial splines suitable for rigorous certification. The spline representation facilitates exact interval-wise differentiation and evaluation of profile-dependent quantities. We use `Arb` [26], an arbitrary-precision interval arithmetic library, to rigorously control numerical and rounding errors throughout the certification procedure. This allows in particular to bound the residual errors of the spline profile by $\mathcal{O}(10^{-5})$ in L^2 and $\mathcal{O}(10^{-3})$ in L^∞.

Connections to theory. We verify that our numerical profile does not violate non-existence results for singularities in the Euler system. In particular, our vorticity profile exhibits only polynomial spatial decay, whereas the non-existence theorems of Chae [10, 11] require super-algebraic decay. We further investigate the case in which the scaling exponent λ is treated as a free parameter and optimized jointly with the profile. Despite this added flexibility, the optimized scaling exponent converges to a value close to the critical exponent $\lambda = 0.5$ predicted by Constantin et al. [16]. Additionally, we consider the family of convection-weighted Euler equations proposed by Hou et al. [19], where a parameter ε allows convection to be progressively increased from the non-convective case at $\varepsilon = 0$ to the full Euler system at $\varepsilon = 1$. As ε approaches 1, the scaling exponent λ decreases toward the critical value $\lambda = 0.5$, providing further agreement with the theoretical prediction of Constantin et al. [16]. For the convection-varied equations, Hou et al. [19] found singular behavior for $\varepsilon \in [0, 0.3)$ using a numerical solver but not beyond $\varepsilon = 0.3$, and conjectured a transition from blowup to regularity near this value. Our PINN formulation disproves this conjecture, by discovering approximate singular profiles for all values of $\varepsilon \in [0, 1]$.

Looking at the transport field of the profiles, we identify two regions of low transport, which can be problematic for stability analysis. The first surrounds a fixed point of the meridional flow on the symmetry axis, while the second is centered at an off-axis meridional fixed point. The axial fixed point has a nearly degenerate locally outgoing property which agrees with the theory from [16] for $\lambda = 0.5$ (with $c_* = 0$ in their notation). In contrast, no such constraint is imposed on the local outgoing behavior of the off-axis fixed point in the theory of [16], which is consistent with our numerical results where the off-axis fixed point is strictly locally outgoing. We use these properties in our stability analysis.

Stability. Starting from the spline representation of the numerical profiles, we establish a complete stability framework, following the strategy from prior blowup stability works [3; 17]. We linearize the appropriate dynamically rescaled equations around the spline profiles and formulate a stability theorem in suitable weighted L^2 and Sobolev norms. We then derive complementary low-order and high-order energy estimates for both the linear and nonlinear terms, with explicit computable bounds. These estimates are incorporated into a single energy functional that yields the quantitative control required to close the full stability argument. We refer to our companion paper [18] for the details of the analysis. The proof is thereby reduced to the certification of a large finite collection of explicit estimates, each of which can be rigorously certified using interval arithmetic or `Arb` computations.

We first focus on linear damping for stability analysis in this paper. Establishing damping at the linear level is a key prerequisite for the complete stability proof, since insufficient linear damping cannot generally be overcome by the remaining nonlinear estimates. Its certification thus isolates a central stabilizing mechanism and provides strong quantitative evidence that the full nonlinear stability argument can be closed. In contrast to prior works, the singular weight functions defining the stability energy norms, together with other tunable parameters entering the stability argument, are optimized numerically rather than selected through a lengthy and often highly nontrivial process of manual trial and error. Here, we represent the weight functions as a combination of locally supported Gaussian basis functions, a singular component near the origin, asymptotically growing neural networks in the far field, and a global neural-network correction, with all components parametrized and jointly optimized.

After optimizing the weights, we obtain low-order damping throughout most of the domain, with the exception of two neighborhoods around the meridional fixed points. For the fixed point on the axis, we can use singular weights and local analysis to obtain low-order damping. At the off-axis fixed point, the flow is strictly locally outgoing, enabling higher-order damping to be achieved through appropriate tuning of the weighting functions. In particular, the relevant higher-order derivatives have a favorable sign and contribute positively to the damping estimate for adequate Sobolev derivative order. Overall, we obtain low-order damping on most of the computational domain, and high-order damping in the problematic off-axis low transport region, and damping outside the computational domain analytically. The combined energy exploits these complementary mechanisms. The resulting damping margin provides an important quantitative indication that the stability scheme can be closed.

Remaining Work to Complete the Proof. The remaining work is primarily computational for the full nonlinear stability in the Euler system: rigorously certifying the profile-dependent constants and margins required in the stability framework using spline representations, interval arithmetic, certified matrix bounds, and finite-dimensional optimization. Some estimates, presented in detail in our companion paper [18], may require further refinement to obtain sufficient margin, but the modular structure of the argument allows such adjustments without changing the underlying stability mechanism. Conditional on these certifications, nonlinear stability of the rescaled profile can be transferred through the dynamic rescaling to finite-time singularity formation in the original variables. At the critical scaling $\lambda = 1/2$, control of the modulation parameters near their self-similar values would yield stability of the candidate profile and, after reconstruction, an admissible solution that becomes singular in finite time.

Formalization in Lean. In the `LeanPDE` paper, we formalize key components of the argument in Lean [27], using `Mathlib` [28] to verify symbolic identities, derivations, and other algebraic steps that enter the proof.

Rigorous numerical certificates are generated separately in `Julia` using `Arb` interval arithmetic [26; 29], then replayed and verified in Lean using `TorchLean` [30] so that the verification pipeline assigns distinct roles to symbolic formalization and certified numerical computation. More broadly, this work is part of the `LeanPDE` effort to develop tools and infrastructure for machine-checked PDE analysis, including closer integration between proof assistants, computer algebra systems, and rigorous numerical methods.

Broader Implications. While much of the recent progress in AI-assisted mathematics has focused on the role of large language models (LLMs) in tasks such as theorem proving and conjecture testing, the central computational ingredient in this work is physics-informed optimization. The governing PDEs are incorporated directly into the optimization problem, allowing candidate singularity structures to be evaluated. Such physics-informed and physics-centric AI are critical ingredients across many areas of research involving physical systems [31], and LLMs lack such physical grounding.

In this paper, simple multilayer perceptron architectures are sufficient for PINNs to identify candidate singularity profiles. In other settings, however, more sophisticated architectures and frameworks, such as Neural Operators [32–34], may be needed. While the approach considered here seeks an individual solution for a specific PDE instance, neural operators are designed to learn solution operators across entire families of parameterized PDEs. Physics-informed neural operators (PINOs) further incorporate the governing PDEs and other physical constraints directly into the operator-learning framework [35; 36]. Such approaches could enable training across families of PDEs and support curriculum strategies that progress from simpler regimes toward the more challenging target problem. In the current context, for example, one could use a PINO to learn the solution operator across a family of Euler equations parameterized by convection strength, starting from weaker convection and progressively approaching the full-convection regime. Such progressive learning may be helpful to discover blowups in other open problems.

The central methodological components of this work, including the PINN framework, the choice of ansatz, and the optimization strategy, are developed by the human authors. LLMs served only in a limited supporting capacity, for tasks such as locating references, understanding relevant background material and definitions, checking intermediate derivations, and assisting with Lean formalization. In presenting the method, we make the reasoning behind the main methodological choices explicit, showing how they arise from the structure of the problem and why they are appropriate in the present setting, and how they agree with existing theoretical results. This is particularly important in AI-assisted mathematical research, where automated methods can yield successful results while leaving the underlying reasoning and mathematical structure opaque. Mathematical and scientific progress depends not only on obtaining successful results, but also on developing the reasoning and structural understanding that make those results meaningful. Making these choices explicit helps preserve that understanding and underscores the essential role of human judgment, interpretation, and scientific insight in mathematical research.

2 The Euler Equations

2.1 The 3D Axisymmetric Equations

Starting from the incompressible 3D Euler equations

$$\partial_t \tilde{u} + (\tilde{u}\cdot\nabla)\tilde{u} = -\nabla\tilde{p}, \qquad \nabla\cdot\tilde{u} = 0, \tag{1}$$

we solve for the velocity field $\tilde{u}(x,t) \in \mathbb{R}^3$ and the scalar pressure $\tilde{p}(x,t) \in \mathbb{R}$, for $x \in \mathbb{R}^3$ and $t \geq 0$. Taking the curl yields the vorticity formulation

$$\partial_t\tilde{\omega} + \tilde{u}\cdot\nabla\tilde{\omega} = \tilde{\omega}\cdot\nabla\tilde{u}, \qquad \tilde{\omega} := \nabla\times\tilde{u}. \tag{2}$$

For flows that are symmetric about a fixed spatial axis, such as the $\mathring{z}$-axis, we can rewrite (1) in cylindrical coordinates $x = (\mathring{r}\cos\theta, \mathring{r}\sin\theta, \mathring{z})$ with $\mathring{r} \geq 0$, $\theta \in [0, 2\pi)$, $\mathring{z} \in \mathbb{R}$. In these coordinates,

$$\tilde{u} = u^r(\mathring{r},\mathring{z},t)\,e_r + u^\theta(\mathring{r},\mathring{z},t)\,e_\theta + u^z(\mathring{r},\mathring{z},t)\,e_z,$$
$$\tilde{\omega} = \omega^r(\mathring{r},\mathring{z},t)\,e_r + \omega^\theta(\mathring{r},\mathring{z},t)\,e_\theta + \omega^z(\mathring{r},\mathring{z},t)\,e_z,$$

where e_r, e_θ, e_z are the cylindrical unit vectors. By definition of axisymmetry, we also have $\partial_\theta(\cdot) = 0$. Following Luo and Hou [8], we introduce the rescaled swirl and streamfunction variables

$$u_\bullet := \frac{u^\theta}{\mathring{r}}, \qquad \omega_\bullet := \frac{\omega^\theta}{\mathring{r}}, \qquad \psi_\bullet := \frac{\psi^\theta}{\mathring{r}},$$

where ψ^θ denotes the θ-component of the streamfunction ψ. This reformulation has the advantage of removing the $1/\mathring{r}$ singularity from the cylindrical coordinates. We also define the radial component u^r and axial component u^z of the velocity in cylindrical coordinates, which can be recovered from $\psi_\bullet$ via

$$u^r = -\,\mathring{r}\,\partial_{\mathring{z}}\psi_\bullet, \qquad u^z = 2\psi_\bullet + \mathring{r}\,\partial_{\mathring{r}}\psi_\bullet.$$

The axisymmetric 3D Euler equations are

$$\begin{aligned}
\partial_t u_\bullet + u^r\,\partial_{\mathring{r}} u_\bullet + u^z\,\partial_{\mathring{z}} u_\bullet &= 2u_\bullet\,\partial_{\mathring{z}}\psi_\bullet, \\
\partial_t \omega_\bullet + u^r\,\partial_{\mathring{r}} \omega_\bullet + u^z\,\partial_{\mathring{z}} \omega_\bullet &= 2u_\bullet\,\partial_{\mathring{z}} u_\bullet, \\
-\mathcal{E}_{\mathring{r},\mathring{z}}\psi_\bullet &= \omega_\bullet
\end{aligned}$$

where $\mathcal{E}_{\mathring{r},\mathring{z}} = \partial_{\mathring{r}}^2 + \frac{3}{\mathring{r}}\partial_{\mathring{r}} + \partial_{\mathring{z}}^2$. The 3D axisymmetric Euler equations can be viewed as the non-viscous counterpart of the 3D axisymmetric Navier–Stokes equations.

2.2 Traveling-Wave Ansatz

We consider a *traveling self-similar* ansatz. *Self-similar* refers to the fact that the geometry of the solution, after a proper rescaling of space, keeps an approximately fixed shape. *Traveling* means that the center of this singularity is not fixed. Instead, it moves along the symmetry ($\mathring{r} = 0$) axis while the solution converges to the self-similar geometry.

The translation is introduced only in the axial ($\mathring{z}$) variable. This is consistent with the axisymmetric setting, where $\mathring{r} = 0$ is the fixed symmetry axis. A radial ($\mathring{r}$) translation would move the singularity away from the symmetry axis and would generally destroy the parity and regularity structure imposed at $\mathring{r} = 0$. By contrast, an axial translation in $\mathring{z}$ preserves the cylindrical symmetry and provides a natural mechanism for the singular structure to drift while it concentrates.

The traveling-wave ansatz. Given the original axisymmetric variables as $(u_\bullet, \omega_\bullet, \psi_\bullet)$, we introduce the *rescaled profile variables* (u, ω, ψ) via

$$u_\bullet(\mathring{r}, \mathring{z}, t) = (T-t)^{c_u}\, u\left(\frac{\mathring{r}}{(T-t)^\lambda}, \frac{\mathring{z} - \mathring{z}_c(t)}{(T-t)^\lambda}\right), \quad (3)$$

$$\omega_\bullet(\mathring{r}, \mathring{z}, t) = (T-t)^{c_\omega}\, \omega\left(\frac{\mathring{r}}{(T-t)^\lambda}, \frac{\mathring{z} - \mathring{z}_c(t)}{(T-t)^\lambda}\right), \quad (4)$$

$$\psi_\bullet(\mathring{r}, \mathring{z}, t) = (T-t)^{c_\psi}\, \psi\left(\frac{\mathring{r}}{(T-t)^\lambda}, \frac{\mathring{z} - \mathring{z}_c(t)}{(T-t)^\lambda}\right). \quad (5)$$

In these equations, T is the *blowup time*, λ is the *spatial blowup rate*, $\mathring{z}_c(t)$ is the center of the traveling profile, and c_u, c_ω, c_ψ are *amplitude exponents*.

The rescaled coordinates are

$$r = \mathring{r}/(T-t)^\lambda,$$

and

$$z = (\mathring{z} - \mathring{z}_c(t))/(T-t)^\lambda,$$

and

$$\frac{d}{dt}\mathring{z}_c(t) = -C(T-t)^{\lambda-1}. \quad (6)$$

Thus C is the *traveling speed* of the profile in the rescaled axial coordinate. If $C = 0$, the singularity is centered at a fixed axial point in the rescaled frame. If $C \neq 0$, the singularity drifts in the original variables while remaining stationary in the moving rescaled frame.

We are considering a *single-scale traveling ansatz*. The same length scale $(T-t)^\lambda$ is used to measure the radial size of the profile, the axial size of the profile, and the axial displacement relative to the traveling center. The purpose is to isolate the traveling self-similar regime in which one spatial scale, one traveling speed, and the corresponding amplitude exponents determine the leading-order profile.

The exponents are fixed by balancing the powers of $T - t$ in the axisymmetric system, to enforce that time differentiation, transport, stretching, and the stream function relation contribute at the same order:

$$c_u = -1, \qquad c_\omega = -1 - \lambda, \qquad c_\psi = -1 + \lambda. \quad (7)$$

The traveling-wave ansatz provides more flexibility for the optimization than the more common *stationary self-similar* ansatz [3; 13–15], in which one exploits scaling invariance and seeks a time-independent backward self-similar profile

$$u_\bullet(x, t) = (T-t)^{c_u}\, u\left(\frac{x - x_0}{(T-t)^\lambda}\right).$$

In our paper, the singularity is allowed to *travel* along the symmetry axis. This introduces an additional axial center $\mathring{z}_c(t)$, and the profile is written in coordinates centered at $\mathring{z}_c(t)$.

2.3 Steady-State Profile Equations

We next present the profile equations generated by the traveling self-similar ansatz. This reduction transforms the original singular evolution into a time-independent nonlinear system for the rescaled variables (U, Ω, Ψ), together with the drift speed C.

Assume $(u_\bullet, \omega_\bullet, \psi_\bullet)$ is given by the ansatz (3)–(5), that the center satisfies (6), that the exponents are chosen as in (7). Then (U, Ω, Ψ) satisfies the profile equations

$$U + \big(\lambda r + U^r\big)\partial_r U + \big(C + \lambda z + U^z\big)\partial_z U = 2U\partial_z\Psi, \tag{8}$$

$$(1+\lambda)\Omega + \big(\lambda r + U^r\big)\partial_r \Omega + \big(C + \lambda z + U^z\big)\partial_z \Omega = 2U\partial_z U, \tag{9}$$

$$-\,\mathcal{E}\,\Psi = \Omega, \tag{10}$$

with $U^r = -r\partial_z\Psi$, $U^z = 2\Psi + r\partial_r\Psi$, and the elliptic operator $\mathcal{E} \coloneqq \partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2$.

Equations (8)–(10) are the steady-state equations for the traveling self-similar profile. Thus, the traveling-wave ansatz reduces the study of a possible finite-time singularity to an autonomous nonlinear profile problem for (U, Ω, Ψ) and the parameter C.

3 Discovery of a Finite-Time Blowup via Physics-Informed Optimization

3.1 Physics-Informed Optimization

Traditional numerical methods for PDE are typically developed and implemented on a case-by-case basis, with finite element methods among the most widely used approaches, but their computational cost can limit their practical applicability and they struggle to identify new self-similar blowup profiles.

Physics-informed optimization provides a more flexible and computationally practical alternative. Physics-Informed Neural Networks (PINNs) [12] are neural networks designed to approximate solutions to PDE. Their parameters are obtained by minimizing deviations from the governing PDE in an appropriate norm, taking a finite set of collocation points in the spatiotemporal domain as input and producing a parametrized approximation to the solution at each grid point. The PDE residual loss is often combined with additional penalty terms that encode boundary and initial conditions, smoothness priors, and other problem-specific constraints, resulting in a composite objective used for gradient-based optimization to train the neural network.

Despite their success, PINNs can be difficult to train due to poorly conditioned optimization problems and increasingly complex loss landscapes as the physics loss is emphasized [37]. They may also converge to trivial or undesired solutions [38], while imbalances among loss terms can lead to uneven gradient contributions during training [39]. Consequently, obtaining reliable PINN solutions requires careful design and tuning of the network architecture, sampling strategy, loss weights, hyperparameters, and optimization procedure. These challenges are particularly important in our setting, where making PINN solutions amenable to computer-assisted proofs requires very high numerical accuracy when solving the Euler equations.

To address these difficulties, we combine several complementary mechanisms designed to improve optimization, enforce the desired solution structure, and increase accuracy. In particular,

- **Soft/Hard constraints** to guide the optimization and encode known properties of the profile, together with smoothness losses.
- **Adaptive loss weighting schedules**, allowing the relative importance of the various objectives to evolve during training and helping mitigate imbalances between competing loss terms.
- **64-bit floating point precision** (FP64) to reduce round-off errors in autograd and achieve the highest fidelity available in PyTorch, improving gradient quality, convergence, and accuracy.
- **PINN boosting** [20; 21], where new networks are added sequentially after the current model reaches an optimization plateau, with each new network trained to correct the residual errors left by the previous ones. This turns one difficult global optimization problem into a sequence of simpler correction problems and allows later stages to use specialized architectures and localized corrections to target the remaining error more effectively.

Table 1: Residual errors for the U, Ω, Ψ profile equations for the Euler system, comparing PINN and spline representations. Note that the residuals are estimated empirically on an independent set of 1,000,000 randomly-sampled collocation points for the PINN representations, whereas they are evaluated analytically and exactly for the spline representations. Here, $\mathbb{D}$ denotes the full domain (i.e. the half plane), $B_3(0)$ the ball of radius 3 centered at the origin, axis the $r = 0$ axis, and FF the far-field region.

		PINNs		**Splines**	
	Domain	RMSE	ℓ^∞	L^2	L^∞
U-equation	$\mathbb{D}$	$4.1{\cdot}10^{-5}$	$3.6{\cdot}10^{-3}$	$4.1{\cdot}10^{-5}$	$3.6{\cdot}10^{-3}$
	$B_3(0)$	$1.6{\cdot}10^{-5}$	$5.7{\cdot}10^{-5}$	$1.7{\cdot}10^{-5}$	$5.7{\cdot}10^{-5}$
	axis	$9.1{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$	$9.1{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$
	FF	$6.2{\cdot}10^{-6}$	$2.2{\cdot}10^{-5}$	$6.2{\cdot}10^{-6}$	$2.2{\cdot}10^{-5}$
Ω-equation	$\mathbb{D}$	$1.7{\cdot}10^{-5}$	$9.3{\cdot}10^{-4}$	$1.7{\cdot}10^{-5}$	$9.3{\cdot}10^{-4}$
	$B_3(0)$	$4.3{\cdot}10^{-6}$	$1.2{\cdot}10^{-5}$	$4.3{\cdot}10^{-6}$	$1.2{\cdot}10^{-5}$
	axis	$4.8{\cdot}10^{-5}$	$3.7{\cdot}10^{-4}$	$4.8{\cdot}10^{-5}$	$3.7{\cdot}10^{-4}$
	FF	$3.1{\cdot}10^{-6}$	$1.4{\cdot}10^{-5}$	$3.1{\cdot}10^{-6}$	$1.4{\cdot}10^{-5}$
Ψ-equation	$\mathbb{D}$	$3.8{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$	$3.8{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$
	$B_3(0)$	$7.2{\cdot}10^{-6}$	$2.5{\cdot}10^{-5}$	$7.8{\cdot}10^{-6}$	$2.5{\cdot}10^{-5}$
	axis	$8.6{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$	$8.6{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$
	FF	$2.1{\cdot}10^{-6}$	$8.3{\cdot}10^{-6}$	$2.1{\cdot}10^{-6}$	$8.3{\cdot}10^{-6}$

- **Self-scaled curvature-aware optimizers**. We first use SS–eSOAP [23] as an efficient curvature-aware optimizer, followed by SS-Broyden [24; 25] for high-accuracy refinement. SS-Broyden approximates second-order curvature information without explicitly computing the Hessian, and its adaptive scaling can improve conditioning and convergence in challenging nonlinear optimization problems. In our case, this optimization strategy reduces the MSE from $\sim 10^{-4}$ with Adam [22] to $\sim 10^{-10}$.

- **A sinh transformation** to efficiently sample on the unbounded domain.

- **Collocation points resampling**, concentrated near the origin and regions where the solution varies rapidly, and residual-based adaptive resampling to target points with the largest PDE errors.

3.2 Numerical Results

We train PINNs for the Euler system, and visualize the resulting approximate profiles U, Ω, Ψ as 3D surfaces and show the corresponding $r = 0$ cross-sections in Figure 2. These are consistent with the profile geometry reported by Hou et al. [19], showing that the approximate self-similar profile remains robust under full convection. The approximate profiles are obtained to high accuracy, as indicated by Table 1.

The PINN serves as a numerical discovery tool that allows us to locate an approximate self-similar profile with high accuracy in a practical and well-conditioned way, especially when direct optimization over analytic ansatz would be difficult or unstable. For certification, we replace the computed profile by a fixed piecewise polynomial spline, which provides an explicit analytic representation on which derivatives and L^∞ bounds can be evaluated exactly. This avoids the more challenging task of certifying network evaluation, automatic differentiation, and L^∞ bounds. We choose splines because their low-degree local structure supports simple and exact differentiation and residual estimates while remaining sufficiently expressive to accurately capture the computed profile. Their locality also allows the fit to the PINN to be refined selectively in more important regions, e.g. near the profile center and the axis. The spline profile retains the small residuals of the PINN, with only a minor loss of accuracy.

4 Stability

A singularity is *stable* if the associated blowup mechanism persists under small perturbations of the initial data, in the sense that the same mechanism occurs for an open neighborhood of nearby initial conditions. Conversely, an *unstable* singularity requires infinitely precise initial conditions: small perturbations deflect the evolution away from the singularity mechanism.

The dynamic rescaling formulation presented in Section 4.1 is the natural starting point for such stability analysis. In the original physical variables, the solution may grow and concentrate as $t \to T^-$, so there is no fixed object around which to linearize. After rescaling, the candidate blowup profile becomes a steady state of the rescaled equations, while the blowup rate, translation rate, and amplitude normalization are recorded by modulation parameters. Linear stability can therefore be studied by perturbing this steady rescaled profile and analyzing the resulting linearized rescaled dynamics.

We follow a similar framework as Chen et al. [3; 17] and related works on stability of blowups. We first compute the blowup profile numerically, then establish linear stability for the dynamically rescaled equation in suitable weighted norms and prove nonlinear stability using Sobolev embeddings and additional weighted estimates. The main distinction is that the weight functions and other tunable parameters in the stability proof are optimized numerically, rather than being chosen and refined through a potentially long and highly nontrivial process of trial and error.

4.1 Dynamic Rescaling Equations

In the dynamic formulation, the fixed self-similar parameters in the profile ansatz are replaced by time-dependent *modulation parameters*. Rather than prescribing a blowup time, a spatial scale, and a traveling center in advance, it defines a long-time evolution in normalized variables, with the singularity kept at order-one size and near a fixed location in rescaled coordinates. This avoids the need to resolve a structure whose amplitude diverges and whose length scale collapses in the physical variables, while allowing the computation to identify am approximate self-similar profile and its associated scaling parameters.

More precisely, let $(\mathring{r}, \mathring{z}, \mathring{t})$ and (r, z, t) denote the physical and dynamically rescaled variables. We introduce a common spatial scale $\mathrm{s}_r(t) > 0$, amplitude scales $\mathrm{s}_u(t)$, $\mathrm{s}_\omega(t) > 0$, and an axial center $\mathring{z}_c(t)$ by

$$\mathring{r} = \mathrm{s}_r(t)\, r, \qquad \mathring{z} - \mathring{z}_c(t) = \mathrm{s}_r(t)\, z,$$
$$u_\bullet(\mathring{r}, \mathring{z}, \mathring{t}) = \mathrm{s}_u(t)\, u(r, z, t),$$
$$\omega_\bullet(\mathring{r}, \mathring{z}, \mathring{t}) = \mathrm{s}_\omega(t)\, \omega(r, z, t).$$

We then define the rescaling rates and axial drift by

$$\partial_t \mathrm{s}_r(t) = -\lambda\, \mathrm{s}_r(t), \quad \partial_t \mathrm{s}_u(t) = -c_u(t)\, \mathrm{s}_u(t),$$
$$\partial_t \mathrm{s}_\omega(t) = -c_\omega(t)\, \mathrm{s}_\omega(t)$$

and $\partial_t \mathring{z}_c(t) = -C(t)\, \mathrm{s}_r(t)$. Overall, this is encoded by the coordinate drift $(\partial_t r,\, \partial_t z)_{\mathring{r},\mathring{z}} = (\lambda r, \lambda z + C(t))$ where $C(t)$ is the *instantaneous traveling speed* of the axial center in the same rescaled frame. The scalar functions $c_u(t)$ and $c_\omega(t)$, governing the amplitude scaling of u and ω, maintain an order-one amplitude in the rescaled variables.

In the original variables, perturbations can change the blowup time, shift the center, or alter the amplitude normalization, producing apparent growth or drift even when the profile is stable. The normalization conditions remove these neutral directions by recentering, rescaling, and renormalizing the solution, so stability can be assessed directly in the rescaled variables. If the dynamic rescaling captures a stable self-similar blowup, the rescaled solution converges to a steady profile and the modulation parameters approach constants. The time derivatives then vanish, and for the traveling-wave ansatz the limiting amplitude rates are $c_u = -1$ and $c_\omega = -(1 + \lambda) = -3/2$. Substituting these values recovers the profile equations. Thus, self-similar profiles are steady states of the dynamically rescaled evolution, with finite-time blowup in physical variables corresponding to convergence toward a stable steady state.

4.2 Linearization

We derive the perturbation equations around a steady state of the dynamically-rescaled system, and separate linear and nonlinear terms in the perturbations. We denote perturbations via δ's in blue, and use the subscript $_0$ to denote evaluation at the origin.

We perturb a given approximate steady state $(\bar{u}, \bar{\omega}, \bar{\psi}, \bar{C}, \bar{c}_u, \bar{c}_\omega)$ of the dynamic rescaling equations, by writing

$$u = \bar{u} + \delta u, \qquad \omega = \bar{\omega} + \delta\omega, \qquad \psi = \bar{\psi} + \delta\psi,$$

$$C = \bar{C} + \delta C, \qquad c_u = \bar{c}_u + \delta c_u, \qquad c_\omega = \bar{c}_\omega + \delta c_\omega,$$

For notational convenience, we will use the symbol δ as shorthand for the full list of perturbation variables,

$$\delta := \{\delta u, \delta\omega, \delta\psi, \delta C, \delta c_u, \delta c_\omega\}.$$

The perturbation variables measure the difference between the evolving rescaled solution and the steady blowup profile. The perturbations must inherit the symmetries of the underlying steady state and therefore belong to the same functional class. In particular, the parity, far-field decay, and regularity conditions imposed on the rescaled variables are inherited by δu, $\delta\omega$, and $\delta\psi$, with the latter two coupled by the elliptic equation relating vorticity and streamfunction.

The reduced variables u and ω do not transform in the same way under rescaling. Instead, ω and ∇u are the natural pair of quantities in a scale-consistent linear stability analysis. We collect these perturbations into the scale-consistent perturbation vector

$$\delta v := (\delta\omega, \nabla\delta u) = (\delta\omega, \partial_r\delta u, \partial_z\delta u),$$

$$\delta v_\omega := \delta\omega, \quad \delta v_r := \partial_r\delta u, \quad \delta v_z := \partial_z\delta u.$$

Linearization consists of expanding the perturbation equation around the approximate blowup profile and retaining only the linear terms in the perturbation. The linearized equation then gives the leading-order dynamics of small perturbations near the profile. The discarded terms collected in nonlinear remainders, which record the higher-order interactions among the perturbations.

Starting from the dynamic rescaling equations, perturbing all variables, subtracting the steady-state equations, and separating first-order terms from higher-order terms gives

$$\begin{aligned}\partial_t\delta u &= \mathcal{L}_u(\delta) + \mathcal{N}_u(\delta) \; + \; \mathcal{R}_u,\\ \partial_t\delta\omega &= \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) \; + \; \mathcal{R}_\omega.\end{aligned}$$

Here $\mathcal{L}_u$ and $\mathcal{L}_\omega$ are linear in the perturbation variables, while $\mathcal{N}_u$ and $\mathcal{N}_\omega$ contain the nonlinear products of perturbations, and $\mathcal{R}_u$ and $\mathcal{R}_\omega$ are the PDE residuals accounting for the fact that the numerical profiles are not exact.

Since the scale-consistent variables are ω and ∇u, we differentiate the δu-equation:

$$\begin{aligned}\partial_t\partial_r\delta u &= \mathcal{L}_r(\delta) + \mathcal{N}_r(\delta) \; + \; \mathcal{R}_r,\\ \partial_t\partial_z\delta u &= \mathcal{L}_z(\delta) + \mathcal{N}_z(\delta) \; + \; \mathcal{R}_z,\end{aligned}$$

where we write $\mathcal{L}_r := \partial_r\mathcal{L}_u$, $\mathcal{N}_r := \partial_r\mathcal{N}_u$, $\mathcal{L}_z := \partial_z\mathcal{L}_u$, $\mathcal{N}_z := \partial_z\mathcal{N}_u$, $\mathcal{R}_z := \partial_z\mathcal{R}_u$, $\mathcal{R}_z := \partial_z\mathcal{R}_u$.

For stability, we use a two-head fit for U and Ψ, while Ω is obtained by applying the elliptic operator exactly to Ψ. Hence, the elliptic residual vanishes by construction, removing one source of certification error and simplifying substantially the stability analysis.

The modulation parameters $\delta C, \delta c_u, \delta c_\omega$ are fixed by normalization conditions, which express them in terms of the profile and its perturbations evaluated at the origin

$$C(t) + u_0^z(t) = 0, \qquad u_0(t) = \bar{u}_0.$$

For the perturbations, these conditions give

$$\delta C + 2\delta\psi_0 = 0, \qquad \delta u_0 = 0, \qquad (\partial_t\delta u)_0 = 0.$$

These normalization conditions remove the artificial freedom associated with axial translation, spatial rescaling, amplitude normalization. They are essential for formulating linear stability in the modulated variables, since otherwise the perturbation could drift along symmetry directions rather than measuring a genuine change in the profile.

4.3 Weighted Norms and Energy

The stability estimates are measured in weighted norms adapted to the blowup profile and damping structure of the linearized operator. We define positive weight functions

$$\Phi_\omega, \Phi_r, \Phi_z : \mathbb{D} \to (0, \infty)$$

on the half-plane

$$\mathbb{D} = \{(r, z) : r \geq 0, \; z \in \mathbb{R}\}.$$

The purpose of the weighted energy is to produce damping from selected linear terms, such as the linear transport terms. We choose weights that are singular at the origin since they strengthen the linear damping mechanism at the origin where the transport terms vanish, while still allowing the weighted energy to remain finite for the admissible perturbation class.

More explicitly, Φ_ω weights the vorticity perturbation $\delta\omega$, Φ_r weights $\partial_r \delta u$, and Φ_z weights $\partial_z \delta u$. We denote the collection of weights by $\Phi := (\Phi_\omega, \Phi_r, \Phi_z)$. For a positive scalar weight ϕ, define the associated inner product and norm by

$$\langle f, g\rangle_\phi := \int_{\mathbb{D}} f g \, \phi \, dr \, dz$$

and $\|f\|_\phi^2 := \langle f, f\rangle_\phi$. Given vector quantities $q := (q_\omega, q_r, q_z)$, and $\tilde{q} := (\tilde{q}_\omega, \tilde{q}_r, \tilde{q}_z)$, their Φ-weighted inner product is defined as

$$\langle q, \tilde{q}\rangle_\Phi := \langle q_\omega, \tilde{q}_\omega\rangle_{\Phi_\omega} + \langle q_r, \tilde{q}_r\rangle_{\Phi_r} + \langle q_z, \tilde{q}_z\rangle_{\Phi_z}.$$

The low-order energy is defined as the Φ-norm of δv:

$$\mathfrak{E}_0(t)^2 = \sum_{i \in \{\omega, r, z\}} \|\delta v_i(t)\|_{\Phi_i}^2.$$

The three weights Φ_ω, Φ_r, Φ_z are kept separate because the damping calculation acts on the three components of δv in different ways.

While the low-order energy remains the main source of damping, it is not sufficient by itself to close the estimates since the linear and nonlinear terms contain terms which cannot be controlled using a weighted L^2 norm, such as terms involving L^∞ norms or pointwise quantities evaluated at the origin. To control these terms, we supplement $\mathfrak{E}_0$ with a high-order weighted energy $\mathfrak{H}_k$, chosen so that the top derivatives provide the required pointwise, interpolation, elliptic, and modulation bounds.

The high-order part uses the adjusted weights

$$\Phi_{i,k} := 1 + \rho^k \Phi_i, \quad \rho(r, z) := r^2 + z^2, \quad i \in \{\omega, r, z\}.$$

We then define the *high-order weighted energy*

$$\mathfrak{H}_k^2 := \sum_{i \in \{\omega, r, z\}} \sum_{|\alpha| = k} \|D^\alpha \delta v_i\|_{\Phi_{i,k}}^2. \tag{11}$$

The full energy used for the stability argument is finally obtained as

$$\mathfrak{E}_k^2 := \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2, \qquad 0 < \mu_k \le 1,$$

where the parameter μ_k is tuned to allow both the low-order and high-order damping to absorb the low-order and high-order losses in the linear estimates.

The order of the top derivative k should be chosen large enough so that the energy controls all point values, products, commutators, and streamfunction terms used in the analytic estimates. With our proof construction, the lowest admissible choice is $k = 4$. Larger choices of k are possible, but increase combinatorially the number of derivatives and analytic constants that must be certified.

No intermediate derivatives are needed in the argument: every derivative level $0 < j < k$ is controlled by the low-order piece $\mathfrak{E}_0^2$, the top-order piece $\mathfrak{H}_k^2$, and explicit interpolation constants $I_{i,j,k}(\eta)$:

$$\|D^j f\|_{\Phi_{i,j}} \le \eta \|D^k f\|_{\Phi_{i,k}} + I_{i,j,k}(\eta) \|f\|_{\Phi_i}.$$

4.4 Stability Estimates

Linear stability asks whether perturbations of solutions of the linearized system remain bounded in a suitable norm. The desired estimate is

$$\|\delta v(t)\|_\Phi \le C e^{-\gamma t} \|\delta v(0)\|_\Phi, \qquad \gamma > 0.$$

A standard rigorous way to prove this is to establish the differential energy estimate

$$\frac{1}{2} \frac{d}{dt} \mathfrak{E}_0(t)^2 = \langle \mathcal{L} \delta v, \delta v\rangle_\Phi \le -\Lambda_{\mathcal{L}} \mathfrak{E}_0(t)^2.$$

For a blowup profile, this is the relevant notion of stability: the original physical solution may still blow up, but the normalized spatial profile is stable in the rescaled variables, i.e. the perturbations converge to zero under the linearized rescaled flow.

However, the low-order energy is not sufficient to close the linear estimates since the linear operators contain terms which cannot be controlled using a weighted L^2 norm. It is thus supplemented with the high-order energy $\mathfrak{H}_k$, to get an estimate of the form

$$\frac{1}{2} \frac{d}{dt} \mathfrak{E}_0(t)^2 \le -\Lambda_{\mathcal{L}}^{\mathrm{low}} \mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}} \mathfrak{H}_k^2.$$

The nonlinear stability argument upgrades this statement to the full perturbation equation, to prove that, if the initial perturbation is sufficiently small in a suitable energy, then the perturbation remains small for as long as the estimates apply. In the case of an approximate numerical profile, the solution is allowed to remain in a small neighborhood whose size also depends on the certified residual error.

4.5 Stability Theorem

The stability proof is reduced to an $\mathfrak{E}_k$ energy estimate. Differentiating $\mathfrak{E}_k^2$ in time and inserting the perturbation equations separates the argument into distinct components:

- *the residual error*, which records the defect of the approximate steady profile
- *low- and high-order linear parts* which contain the stabilizing mechanism: the low-order estimate gives the main source of damping, while the high-order estimate controls the differentiated equations, allowing only losses that can be absorbed by the low-order energy after choosing μ_k.
- *low- and high-order nonlinear parts*, bounded by cubic powers of the full energy $\mathfrak{E}_k$.

These estimates culminate in the stability theorem below, with the understanding that every certification constant entering the analytic bounds must be rigorously validated for the proof to be complete. The detailed estimates and derivations are given in our companion paper [18].

Theorem 1 (Euler Stability)**.** *Assume that the approximate profile has a sufficiently small residual and that the certified low- and high-order linear estimates can be combined, after a suitable choice of the high-order weight μ_k, to produce strict damping of the full energy $\mathfrak{E}_k$. Assume further that the nonlinear terms contribute only cubic powers of $\mathfrak{E}_k$.*

These estimates yield an inequality of the form

$$\frac{d}{dt}\mathfrak{E}_k^2 \leq -\Lambda_{\mathrm{stab}}\mathfrak{E}_k^2 + \Lambda_3\mathfrak{E}_k^3 + \Lambda_{\mathcal{R}}\mathfrak{E}_k,$$

where $\Lambda_{\mathrm{stab}} > 0$ represents the net linear damping, Λ_3 controls the nonlinear terms, and $\Lambda_{\mathcal{R}}$ measures the residual.

If there exists $\delta_* > 0$ such that

$$\Lambda_3\delta_* + \frac{\Lambda_{\mathcal{R}}}{\delta_*} < \Lambda_{\mathrm{stab}}, \tag{12}$$

then the damping dominates both the nonlinear effects and the residual forcing on the boundary $\mathfrak{E}_k = \delta_*$.

Consequently, the ball $\{\mathfrak{E}_k < \delta_*\}$ is forward invariant: every solution satisfying $\mathfrak{E}_k(0) < \delta_*$ continues to satisfy $\mathfrak{E}_k(t) < \delta_*$ for as long as the solution exists and the certified estimates remain valid. Thus, sufficiently small perturbations remain uniformly controlled in the full energy $\mathfrak{E}_k$.

The single-radius certificate can be relaxed by prescribing separate target bounds for the low-order and top-order energies, i.e. $\mathfrak{E}_0(t) < \delta_0$ and $\mathfrak{H}_k(t) < \delta_k$.

4.6 Rescaled Stability to Physical Blowup

The stability analysis is carried out in dynamically rescaled variables, so a reconstruction is needed to recover the physical behavior. If an exact admissible rescaled solution exists for all $t \geq 0$ and remains in the stability regime, then the amplitude scale $\mathrm{s}_u(t)$ grows exponentially whenever $c_u(t)$ stays negative. Since the physical clock satisfies $d\mathring{t}/dt = 1/\mathrm{s}_u(t)$, this growth makes the total remaining physical time finite. At the same time, the fixed normalization at the rescaled origin implies that the physical axial vorticity at the moving center grows proportionally to $\mathrm{s}_u(t)$ and thus becomes unbounded. The center itself converges to a finite physical location. Hence the rescaled stability regime reconstructs to a finite-time physical singularity.

This reconstruction step is conditional on the existence of the controlled rescaled trajectory for all $t \geq 0$, but does not require convergence to the numerical profile. An unconditional blowup result would additionally require a local well-posedness and continuation argument ruling out breakdown at finite rescaled time.

4.7 Computer-Assisted Proof Strategy

The stability theorems reduce stability to a finite-dimensional optimization problem. The aim is to choose the weights so that the linear damping margin is larger than the nonlinear estimates and residual error. For each candidate weight, we compute the certified constants, optimize the radius, and update the weights until the closing inequality becomes strict.

The optimization loop proceeds as follows:

1. *Parametrize the weight functions* with parameters ϑ, while ensuring that the weights have the required positivity, regularity, and symmetries. Choose an initial admissible candidate (ϑ, μ_k).

2. *Compute the linear constants* $\Lambda_{\mathcal{L}}^{\text{low}}$, $\Lambda_{\mathcal{L}}^{\text{high}}$, $\Lambda_{D^\alpha\mathcal{L}}^{\text{low}}$, $\Lambda_{D^\alpha\mathcal{L}}^{\text{high}}$ for the current candidate. If the damping margin available $\Lambda_{\text{stab}} \le 0$, this candidate cannot close the theorem, so we move to a new candidate.

3. *Compute the nonlinear and residual constants.* If $\Lambda_{\text{stab}} > 0$, compute Λ_3 and $\Lambda_{\mathcal{R}}$.

4. *Optimize the radius and evaluate the gap* $\mathcal{J}(\vartheta, \mu_k)$ which determines whether the candidate succeeds. When successful, choose an admissible radius δ_* such that (12) holds.

5. *Iterate* steps 2–4 by update (ϑ, μ_k) to decrease $\mathcal{J}(\vartheta, \mu_k)$, until the proof closes.

6. *Certify* rigorously and analytically that the proof closes with the candidate weights and radius.

4.8 Partial Results: Linear Damping

A central question in the stability argument is whether the leading linear terms provide a positive damping margin before the remaining terms are added. This matters because missing damping cannot generally be repaired by taking smaller perturbations. We evaluate the signed linearized operator directly in the weighted energies and obtain strongly favorable results. Low-order damping controls most of the half-plane, high-order damping supplies the localized repair near the off-axis low transport regions, singular weights can control the on-axis low transport region, the exterior region is closed analytically, and the profile residual is small in the weighted norms.

The low-order transport terms and selected couplings form a symmetric matrix $\mathbb{M}_{\mathcal{L}}(r, z)$. Pointwise damping is equivalent to $\lambda_{\max}(\mathbb{M}_{\mathcal{L}}(r, z)) < 0$. The weights Φ_i enter through the transport contribution and are optimized to decrease the largest eigenvalue of $\mathbb{M}_{\mathcal{L}}$. The largest eigenvalue is negative on the entire computational region, except in two small neighborhoods where the transport terms become very small. This defect is highly localized, and the high-order linear equations provide the natural repair off-axis because derivatives of the transport field need not be small when the transport itself vanishes. The numerical evaluation reveals higher-order damping in precisely the off-axis neighborhood where the low-order margin deteriorates. In addition, singular weights can be used to control the on-axis low transport region.

Outside the spline support (i.e. the finite computational box), the profile-dependent coefficients vanish exactly, the signed matrix reduces to an explicit diagonal tail form, and the weights can be continued analytically so that the tail inequalities close without additional numerical computation.

4.9 Remaining quantitative certification

The signed matrix provides damping, but the final linear estimate must also include the remaining linear terms. The relevant linear stability margin is

$$\Lambda_{\text{stab}} = \min\left\{ \Lambda_{\mathcal{L}}^{\text{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\text{low}},\ \Lambda_{D^\alpha\mathcal{L}}^{\text{high}} - \frac{\Lambda_{\mathcal{L}}^{\text{high}}}{\mu_k} \right\} > 0.$$

A positive Λ_{stab} gives a strict linear damping margin. The certified signed matrix bounds show that the central linear mechanism required for this margin is already present. The remaining task is to certify the terms treated outside the signed matrices and verify that their combined losses remain below the available damping. If additional margin is needed, the same pipeline can be improved by refining the profile in the regions that contribute most to the residuals, optimizing the weights and other parameters, or sharpening the analytic estimates.

The remaining task is quantitative rather than structural: rigorously certify and optimize the remaining terms and estimates of the proof. Some selected analytic estimates may still require sharpening to recover additional margin, but such changes can be made within the existing modular proof framework. If the remaining constants are rigorously certified with a positive closing margin, then the nonlinear stability argument is fully closed. Combined with the finite-time reconstruction from the dynamically rescaled variables, this produces an admissible solution converging to the approximate self-similar profile while developing a singularity in finite physical time. Thus, remaining work is rigorous quantitative certification and optimization within an already established stability and singularity framework.

Conclusion

We provided the first evidence that the Euler equations admit a self-similar blowup on $\mathbb{R}^3$ at the critical scaling exponent λ of 0.5, in agreement with prior theory. We obtained a highly accurate approximate profile using PINNs with a novel combination of complementary tools, including soft and hard constraints, adaptive loss weighting, FP64 arithmetic, boosting, adaptive resampling, and self-scaled curvature aware optimizers. The PINN profile was then converted into piecewise polynomial spline surrogates suitable for rigorous certification.

Complementary low- and high-order energy estimates for the linear and nonlinear terms are derived in the companion stability paper [18], with every constant and estimate made explicit or reduced to a rigorously computable quantity. These estimates are incorporated into a single energy functional that yields the quantitative control required to establish stability of the profiles. The stability proof is thereby reduced to the certification of a large finite collection of explicit inequalities, each of which can be rigorously certified using interval arithmetic or `Arb` computations. We initiated this computational program by focusing on linear damping, and overall obtained low-order damping on most of the domain and high-order damping in the problematic off-axis low-transport region. This damping provides an important quantitative indication that the stability scheme can close, and thus that the profiles are stable.

The remaining work is primarily computational: rigorously certifying the profile-dependent constants and margins in the stability framework using spline representations, interval arithmetic, certified matrix bounds, and finite-dimensional optimization. Some estimates may require refinement to achieve sufficient margin, but the modular structure of the argument permits such adjustments without altering the underlying stability mechanism. Conditional on these certifications, nonlinear stability of the rescaled profile can be transferred through dynamic rescaling to finite-time singularity formation in the original variables. When $\lambda = 1/2$, control of the modulation parameters near their self-similar values would yield stability of the candidate profile and, after reconstruction, an admissible solution developing a finite-time singularity.

In parallel, the `LeanPDE` paper develops a formalization in Lean [27] of key components of the argument, with particular emphasis on the symbolic derivations and proof steps. Numerical certificates are generated separately in `Julia` using `Arb` interval arithmetic [26; 29], while Lean verifies the algebraic identities and symbolic relations underlying the certified estimates. The resulting pipeline assigns complementary roles to its different components: physics-informed optimization is used for discovery, splines and interval arithmetic for rigorous numerical certification, and Lean for formal symbolic verification and, ultimately, machine-checked analysis.

Author Contributions

A.G. performed all numerical experiments and implementations, generated the figures, and contributed substantially to the development of the theoretical results in close collaboration with A.A. and V.D. A.G. also proofread the manuscript and provided suggestions on the exposition.

V.D. established most of the derivations and theoretical results, closely supervised the design, development, and execution of the numerical experiments, and carried out most of the writing and overall preparation of the manuscript.

A.A. led the project, closely supervised numerical experiments and theoretical developments, and made substantial contributions to the framing and presentation of the paper, including writing the abstract, introduction, and conclusion with V.D.

Acknowledgments

A. Anandkumar is supported in part by the Bren endowed chair, ONR (MURI grant N00014-18-12624), Darpa ExpMath HR0011 and by the AI2050 senior fellow program at Schmidt Sciences. V. Duruisseaux is supported in part by the AI2050 program and SciDAC P-240000887. A. Ganeshram was supported by the Caltech SURF program during the summer.

We are especially grateful to Robert J. George for his generous contributions to this work, including his help with the certification framework and his formal verification of the derivations in Lean as part of the `LeanPDE` paper.

We are also grateful to Peicong Song for his generous time, insightful discussions, and valuable suggestions. We also benefited from discussions with Changhe Yang and Yixuan Wang during the early stages of this work.

We acknowledge the use of LLMs for general brainstorming, verification of derivations, and Lean autoformalization. Importantly, LLMs were not used in the central component of this work, namely, the discovery of the approximate profile functions. The PINN formulation, including the choice of soft and hard constraints and efficient sampling and optimization strategies, were all developed solely by the authors.

Supporting Information for

Stable Singularity of the Euler Equations on $\mathbb{R}^3$

Adarsh Ganeshram* AGANESHRAM@BERKELEY.EDU

Valentin Duruisseaux* VDURUISS@CALTECH.EDU

Anima Anandkumar ANIMA@CALTECH.EDU

Department of Mathematics, University of California Berkeley, Berkeley, CA, USA

Department of Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA, USA

Contents

From Numerical Profile Discovery to Certified Nonlinear Stability

Soft Constraints

Hard Constraints

Boosting

Curature-Aware Optimizer

High Precision

Specialized Sampling

NUMERICS

Approximate Numerical Profile

Profile discovery. Starting from a parametrized traveling wave self-similar profile ansatz, obtain a blowup profile (u, ω, ψ) using physics-informed neural networks (PINNs).

Analytic representation. Fit the PINN profile with splines to obtain a representation suitable for rigorous certification, enabling exact estimates, and analytic expressions without introducing further numerical approximation errors.

linearize around the numerical profile

PERTURBATIONS

Perturbation Equations and Stability Energy

Linearize around the spline profile and introduce perturbations $\delta = (\delta\omega, \delta u)$, with $\delta v_\omega := \delta\omega$, $\delta v_r := \partial_r \delta u$, and $\delta v_z := \partial_z \delta u$. Schematically,

$$\partial_t \delta\omega = \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) + \mathcal{R}_\omega, \qquad \partial_t \partial_i \delta u = \partial_i \mathcal{L}_u(\delta) + \mathcal{N}_i(\delta) + \mathcal{R}_i, \quad i \in \{r, z\}.$$

Weights and energy. Define the full energy $\mathfrak{E}_k^2 := \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2$, where

$$\mathfrak{E}_0^2 := \|\delta\omega\|_{\Phi_\omega}^2 + \|\partial_r \delta u\|_{\Phi_r}^2 + \|\partial_z \delta u\|_{\Phi_z}^2, \qquad \mathfrak{H}_k^2 := \sum_{i \in \{\omega, r, z\}} \sum_{|\alpha| = k} \|D^\alpha \delta v_i\|_{\Phi_{i,k}}^2,$$

with $\Phi_{i,k} := 1 + \rho^k \Phi_i$ and $\rho := r^2 + z^2$.

derive analytic estimates and form the stability proof

STABILITY THEORY

Analytic Estimates and Conditional Stability Proof

Analytic results and estimates. Leverage Sobolev embeddings, weighted transport estimates, Leibniz estimates, elliptic bounds, and modulation estimates to derive the linear and nonlinear energy estimates required to close the stability argument, conditional on obtaining the necessary numerical constants in these estimates and suitable properties of the numerical profile.

The derivations, analytic estimates, and the bootstrap argument, can later be formalized in Lean.

certify the numerical inputs and explicit estimates

CERTIFICATION

Certified Constants and Estimates

Rigorous certification. Starting from the spline profile, interval arithmetic rigorously encloses the profile data, the operator and stability constants, and the PDE residual bound, so that every numerical ingredient is converted into an explicit, machine-verifiable bound.

Certified proof. The certified quantities can then be incorporated into the conditional stability argument, with the aim of closing the proof. Subsequently, these estimates can be encoded within a Lean formalization framework to obtain a fully computer-verified proof.

optimize the singular weights and parameters to conclude

CONCLUSION

Computer-Assisted Nonlinear Stability

Optimize the admissible singular weight functions Φ, μ_k, and stability radius δ_* so that the certified damping dominates the nonlinear and residual contributions.

Stability conclusion. If $\mathfrak{E}_k(0) < \delta_*$, then $\mathfrak{E}_k(t) < \delta_*$ for every later admissible time. Therefore $\{\mathfrak{E}_k < \delta_*\}$ is forward invariant, yielding the desired nonlinear stability.

1 The Euler Equations

1.1 The 3D Axisymmetric Equations

Starting from the incompressible 3D Euler equations written by the great Euler [40],

$$\partial_t \tilde{u} + (\tilde{u} \cdot \nabla)\tilde{u} = -\nabla \tilde{p}, \qquad \nabla \cdot \tilde{u} = 0, \tag{S1}$$

we solve for the velocity field $\tilde{u}(x,t) \in \mathbb{R}^3$ and the scalar pressure $\tilde{p}(x,t) \in \mathbb{R}$, for $x \in \mathbb{R}^3$ and $t \geq 0$. Taking the curl of the above equation yields the vorticity formulation

$$\partial_t \tilde{\omega} + \tilde{u} \cdot \nabla \tilde{\omega} = \tilde{\omega} \cdot \nabla \tilde{u}, \qquad \tilde{\omega} := \nabla \times \tilde{u}. \tag{S2}$$

In $\mathbb{R}^3$, one can recover $\tilde{u}$ from $\tilde{\omega}$ via the Biot–Savart law. Equivalently, one can introduce a vector streamfunction $\tilde{\psi}(x,t) \in \mathbb{R}^3$ satisfying

$$-\Delta \tilde{\psi} = \tilde{\omega}, \qquad \tilde{u} = \nabla \times \tilde{\psi}, \tag{S3}$$

together with a gauge condition (e.g. $\nabla \cdot \tilde{\psi} = 0$) to fix $\tilde{\psi}$.

For flows that are symmetric about a fixed spatial axis, such as the $\mathring{z}$-axis, it is convenient to rewrite equation (S1) in cylindrical coordinates

$$x = (\mathring{r}\cos\theta, \mathring{r}\sin\theta, \mathring{z}), \qquad \mathring{r} \geq 0, \quad \theta \in [0, 2\pi), \quad \mathring{z} \in \mathbb{R}. \tag{S4}$$

In these coordinates, the velocity and vorticity decompose as

$$\tilde{u} = u^r(\mathring{r}, \mathring{z}, t)\, e_r + u^\theta(\mathring{r}, \mathring{z}, t)\, e_\theta + u^z(\mathring{r}, \mathring{z}, t)\, e_z, \tag{S5}$$

$$\tilde{\omega} = \omega^r(\mathring{r}, \mathring{z}, t)\, e_r + \omega^\theta(\mathring{r}, \mathring{z}, t)\, e_\theta + \omega^z(\mathring{r}, \mathring{z}, t)\, e_z, \tag{S6}$$

where e_r, e_θ, e_z are the unit vectors in cylindrical coordinates. In addition, by definition of axisymmetry, we have that $\partial_\theta(\cdot) = 0$.

Following Luo and Hou [8], it is convenient to introduce the rescaled swirl and streamfunction variables

$$u_\bullet := \frac{u^\theta}{\mathring{r}}, \qquad \omega_\bullet := \frac{\omega^\theta}{\mathring{r}}, \qquad \psi_\bullet := \frac{\psi^\theta}{\mathring{r}}, \tag{S7}$$

where ψ^θ denotes the θ-component of the vector potential ψ in (S3). This reformulation has the advantage of removing the $1/\mathring{r}$ singularity from the cylindrical coordinates.

We also define the radial component u^r and axial component u^z of the velocity in cylindrical coordinates. These meridian velocity components (u^r, u^z) can be recovered from $\psi_\bullet$ via the (axisymmetric) streamfunction relations

$$u^r = -\,\mathring{r}\,\partial_{\mathring{z}}\psi_\bullet, \qquad u^z = 2\psi_\bullet + \mathring{r}\,\partial_{\mathring{r}}\psi_\bullet. \tag{S8}$$

3D Axisymmetric Euler Equations. The axisymmetric 3D Euler equations can be written as

$$\partial_t u_\bullet + u^r\,\partial_{\mathring{r}} u_\bullet + u^z\,\partial_{\mathring{z}} u_\bullet = 2u_\bullet\,\partial_{\mathring{z}}\psi_\bullet, \tag{S9}$$

$$\partial_t \omega_\bullet + u^r\,\partial_{\mathring{r}} \omega_\bullet + u^z\,\partial_{\mathring{z}} \omega_\bullet = 2u_\bullet\,\partial_{\mathring{z}} u_\bullet, \tag{S10}$$

$$-\Big[\partial_{\mathring{r}}^2 + \frac{3}{\mathring{r}}\partial_{\mathring{r}} + \partial_{\mathring{z}}^2\Big]\psi_\bullet = \omega_\bullet. \tag{S11}$$

Equations (S8), (S9)–(S11) form the reduced axisymmetric-with-swirl system in $(\mathring{r}, \mathring{z})$ for $(u_\bullet, \omega_\bullet, \psi_\bullet)$.

Axis interpretation. Every occurrence of $\partial_{\mathring{r}}^2 + 3\mathring{r}^{-1}\partial_{\mathring{r}} + \partial_{\mathring{z}}^2$ at $\mathring{r} = 0$ denotes its axis-regular continuation, not literal division by zero. For a C^2 profile even in $\mathring{r}$,

$$\left(\partial_{\mathring{r}}^2 + \frac{3}{\mathring{r}}\partial_{\mathring{r}} + \partial_{\mathring{z}}^2\right) f(0,\mathring{z}) = 4\,\partial_{\mathring{r}\mathring{r}} f(0,\mathring{z}) + \partial_{\mathring{z}\mathring{z}} f(0,\mathring{z}). \tag{S12}$$

1.2 Traveling-Wave Ansatz

We consider a *traveling self-similar* ansatz.

- *Self-similar* refers to the fact that the geometry of the solution, after a proper rescaling of space, keeps an approximately fixed shape.
- *Traveling* means that the center of this singularity is not fixed. Instead, it moves along the symmetry ($\mathring{r} = 0$) axis while the solution converges to the self-similar geometry.

Since the ansatz is written in terms of powers of $\tau = T - t$, it is also a *backward self-similar ansatz*: the possible singular profile is described by rescaling backward from the blowup time T.

The translation is introduced only in the axial ($\mathring{z}$) variable. This is consistent with the axisymmetric setting, where $\mathring{r} = 0$ is the fixed symmetry axis and $\mathring{r} \geq 0$ measures distance to that axis. A radial ($\mathring{r}$) translation would move the singularity away from the symmetry axis and would generally destroy the parity and regularity structure imposed at $\mathring{r} = 0$. By contrast, an axial translation in $\mathring{z}$ preserves the cylindrical symmetry and provides a natural mechanism for a coherent structure to drift while it concentrates.

The traveling-wave ansatz. Given the original axisymmetric variables as $(u_\bullet, \omega_\bullet, \psi_\bullet)$, we introduce the *rescaled profile variables* (u, ω, ψ) via

$$u_\bullet(\mathring{r},\mathring{z},t) = (T-t)^{c_u}\, u\left(\frac{\mathring{r}}{(T-t)^\lambda}, \frac{\mathring{z}-\mathring{z}_c(t)}{(T-t)^\lambda}\right), \tag{S13}$$

$$\omega_\bullet(\mathring{r},\mathring{z},t) = (T-t)^{c_\omega}\, \omega\left(\frac{\mathring{r}}{(T-t)^\lambda}, \frac{\mathring{z}-\mathring{z}_c(t)}{(T-t)^\lambda}\right), \tag{S14}$$

$$\psi_\bullet(\mathring{r},\mathring{z},t) = (T-t)^{c_\psi}\, \psi\left(\frac{\mathring{r}}{(T-t)^\lambda}, \frac{\mathring{z}-\mathring{z}_c(t)}{(T-t)^\lambda}\right). \tag{S15}$$

Here T is the *blowup time*, λ is the *spatial blowup rate*, $\mathring{z}_c(t)$ is the center of the traveling profile, and c_u, c_ω, c_ψ are *amplitude exponents*.

The rescaled coordinates are

$$r = \frac{\mathring{r}}{(T-t)^\lambda}, \qquad z = \frac{\mathring{z}-\mathring{z}_c(t)}{(T-t)^\lambda}, \tag{S16}$$

and the rescaled meridian velocity components u^r and u^z, can be recovered from ψ via

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi. \tag{S17}$$

The traveling speed is encoded by the constant C through

$$\frac{d}{dt}\mathring{z}_c(t) = -C(T-t)^{\lambda-1}. \tag{S18}$$

With this convention, the moving coordinate z satisfies

$$(\partial_t z)_{\mathring{r},\mathring{z}} = \frac{\lambda z + C}{T-t}. \tag{S19}$$

Thus C is the *traveling speed* of the profile in the rescaled axial coordinate. If $C = 0$, the singularity is centered at a fixed axial point in the rescaled frame. If $C \neq 0$, the singularity drifts in the original variables while remaining stationary in the moving rescaled frame.

The traveling-wave ansatz considered here provides more flexibility during the optimization than the more common *stationary self-similar* dynamics studied in [3; 13–15]. In the stationary case, one exploits scaling invariance and seeks a time-independent backward self-similar profile of the form

$$u_\bullet(x,t) = (T-t)^{c_u}\, u\!\left(\frac{x - x_0}{(T-t)^\lambda}\right). \tag{S20}$$

Here x denotes the relevant spatial variable, $(T-t)^\lambda$ is the concentrating length scale, and $(T-t)^{c_u}$ gives the amplitude growth.

The profile is *stationary* because its center is fixed in physical space: after rescaling, the same spatial point remains the center of the core. In the present axisymmetric setting, the singularity is allowed to *travel* along the symmetry axis. This introduces an additional axial center $\mathring{z}_c(t)$, and the profile is written in coordinates centered at $\mathring{z}_c(t)$ rather than at a fixed point.

Scaling exponents. The exponents are fixed by balancing the powers of $T-t$ in the axisymmetric system. This balance enforces that time differentiation, transport, stretching, and the elliptic streamfunction relation all contribute at the same leading order. It gives

$$c_u = -1, \qquad c_\omega = -1-\lambda, \qquad c_\psi = -1+\lambda. \tag{S21}$$

These exponents have a simple interpretation. The variable $u_\bullet$ grows like $(T-t)^{-1}$. The vorticity variable $\omega_\bullet$ contains one additional spatial derivative at scale $(T-t)^\lambda$, and therefore grows like $(T-t)^{-1-\lambda}$. The streamfunction is two spatial derivatives smoother than $\omega_\bullet$, which leads to the exponent $-1+\lambda$.

Energy scaling constraint. The traveling-wave ansatz also has consequences for energy quantities in the original variables. For the convection-varied Euler equations, the relevant energy identity is recalled from Theorem 4.2 of [41]. The scaling calculation below applies whenever the corresponding energy is finite and remains bounded as $t \to T^-$.

Written in the original variables $(\mathring{r}, \mathring{z})$, the relevant energy functional is

$$E = \int_{-\infty}^{\infty}\int_0^{\infty} \left(|u_\bullet|^2 + |\nabla_{\mathring{r},\mathring{z}}\psi_\bullet|^2\right)\, \mathring{r}^3\, d\mathring{r}\, d\mathring{z}. \tag{S22}$$

This identity links the self-similar scaling exponent λ to a quantity expressed in the original physical variables. The following proposition shows how E transforms under the traveling-wave ansatz and provides a simple but important consistency condition: a finite-energy self-similar profile cannot have an arbitrary concentration rate λ.

Proposition S1 (Energy scaling under the self-similar ansatz)**.** *Let $\tau = T - t$, and assume that*

$$u_{\bullet}(\mathring{r}, \mathring{z}, t) = \tau^{-1} u\left(\frac{\mathring{r}}{\tau^{\lambda}}, \frac{\mathring{z} - \mathring{z}_c(t)}{\tau^{\lambda}}\right), \qquad \psi_{\bullet}(\mathring{r}, \mathring{z}, t) = \tau^{\lambda-1} \psi\left(\frac{\mathring{r}}{\tau^{\lambda}}, \frac{\mathring{z} - \mathring{z}_c(t)}{\tau^{\lambda}}\right). \tag{S23}$$

Equivalently, introduce the rescaled variables

$$r = \frac{\mathring{r}}{\tau^{\lambda}}, \qquad z = \frac{\mathring{z} - \mathring{z}_c(t)}{\tau^{\lambda}}. \tag{S24}$$

If the rescaled profile energy is finite and nonzero, then boundedness of $E(t)$ as $t \to T^-$ requires

$$\lambda \geq \frac{2}{5}. \tag{S25}$$

Proof The proof is provided in Appendix E. ■

The proposition shows that the concentration rate λ is constrained by the scaling of the energy in the original variables. If the rescaled profile energy is finite and nonzero, boundedness of $E(t)$ as $t \to T^-$ requires $\lambda \geq 2/5$. When $\lambda > 2/5$, the contribution of an exact self-similar core to $E(t)$ tends to zero as $t \to T^-$. When $\lambda = 2/5$, the energy contribution is invariant under the self-similar scaling.

Remark. *Constantin et al. [16] investigate the admissible scaling exponents for a hypothetical self-similar blowup of the Euler equations. They establish the lower bound $\lambda \geq 2/5$ and further show that, under a local outflow condition essential for the corresponding linear stability analysis, one must in fact have $\lambda \geq 1/2$. Prior to this work, it was hoped that an Euler singularity with $2/5 \leq \lambda < 1/2$ might persist in the Navier–Stokes setting, since viscosity would be asymptotically lower order in this regime. In that case, a self-similar Euler blowup implies a Navier–Stokes singularity. The result of Constantin et al. [16], however, rules out this scenario for profiles satisfying the local outflow condition. The exponent $\lambda = 1/2$ is especially important in both Euler and Navier–Stokes flows. For Euler, its importance is tied to the invariance of circulation, and for Navier–Stokes, this is the unique exponent at which self-similar blowup can occur with viscosity present as the blowup time is approached.*

1.3 Steady-State Profile Equations

We next present the profile equations generated by the traveling self-similar ansatz. This reduction transforms the original singular evolution into a time-independent nonlinear system for the rescaled variables (U, Ω, Ψ), together with the scaling exponent λ and the drift speed C.

For notation simplicity, we define the elliptic operator

$$\mathcal{E} \coloneqq \partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2. \tag{S26}$$

Proposition S2 (Traveling-wave profile equations)**.** *Assume that* $(u_\bullet, \omega_\bullet, \psi_\bullet)$ *is given by the ansatz* (S13)–(S15)*, that the center satisfies* (S18)*, and that the exponents are chosen as in* (S21)*. Set*

$$\lambda = \frac{1}{2}. \tag{S27}$$

Then (U, Ω, Ψ) *satisfies*

$$U + \big(\lambda r + U^r\big)\partial_r U + \big(C + \lambda z + U^z\big)\partial_z U = 2U\partial_z\Psi, \tag{S28}$$
$$(1+\lambda)\Omega + \big(\lambda r + U^r\big)\partial_r \Omega + \big(C + \lambda z + U^z\big)\partial_z \Omega = 2U\partial_z U, \tag{S29}$$
$$-\mathcal{E}\,\Psi = \Omega, \tag{S30}$$

where

$$U^r = -r\partial_z\Psi, \qquad U^z = 2\Psi + r\partial_r\Psi. \tag{S31}$$

Equivalently,

$$U + r\big(\lambda - \partial_z\Psi\big)\partial_r U + \big(C + \lambda z + 2\Psi + r\partial_r\Psi\big)\partial_z U = 2U\partial_z\Psi, \tag{S32}$$
$$(1+\lambda)\Omega + r\big(\lambda - \partial_z\Psi\big)\partial_r \Omega + \big(C + \lambda z + 2\Psi + r\partial_r\Psi\big)\partial_z \Omega = 2U\partial_z U, \tag{S33}$$
$$-\mathcal{E}\,\Psi = \Omega. \tag{S34}$$

Proof The derivation is provided in Appendix C. ■

In these equations,

- The terms $\lambda r\partial_r$ and $\lambda z\partial_z$ come from observing the solution in a shrinking coordinate system.
- The term $C\partial_z$ comes from following the center of the profile as it travels along the symmetry axis.
- The terms involving U^r and U^z are the convection terms.
- The terms $2U\partial_z\Psi$ and $2U\partial_z U$ on the right-hand side are the stretching and coupling terms inherited from the axisymmetric formulation.

Equations (S28)–(S30) are the steady-state equations for the traveling self-similar profile. Thus, the traveling-wave ansatz reduces the study of a possible finite-time singularity to an autonomous nonlinear profile problem for (U, Ω, Ψ) and the parameters λ and C. The parameter λ determines the rate at which the spatial scale collapses, while C determines the speed of the axial drift in the rescaled frame.

2 Discovering a Finite-Time Blowup via Physics-Informed Optimization

Classical numerical methods for solving PDE typically rely on problem-specific discretizations and solver configurations, which can make their implementation and adaptation across different equations, geometries, and parameter regimes challenging. Among these approaches, finite element methods are among the most widely used due to their versatility and well-established mathematical foundation. However, high-fidelity finite element simulations can be computationally demanding, particularly for large-scale systems, problems characterized by rapidly varying dynamics, or applications requiring repeated PDE solutions. These limitations can become particularly acute for nonlinear PDE, where general-purpose finite element solvers may require weeks of CPU time and can struggle to discover previously unknown blowup profiles, especially when little prior information about their structure is available.

Physics-informed optimization offers a flexible and computationally tractable framework for solving PDE, with the potential to transfer across problem settings and facilitate the discovery of previously unknown solution profiles. Here, we focus on a particular class of such methods, physics-informed neural networks (PINNs), which recast the solution of a given PDE as an optimization problem whose objective is constructed directly from the governing equations and associated constraints. From this perspective, PINNs can be viewed as an alternative class of numerical methods for PDE, in which a parameterized representation of the solution is obtained through gradient-based optimization rather than through traditional algorithms applied to a prescribed discretization of the domain.

2.1 Methods

2.1.1 Physics-Informed Optimization

Physics-Informed Neural Networks (PINNs) [12] are neural networks designed to approximate solutions to PDE. Their parameters are obtained by minimizing deviations from the governing PDE in an appropriate norm, taking a finite set of collocation points in the spatiotemporal domain as input and producing a parametrized approximation $u_\theta(x,t)$ to the solution at each grid point (x,t). Physics-Informed Neural Operators (PINOs) [35; 36] extend the same principle to learning a *solution operator* mapping to an entire family of solutions [34; 42], for instance with Fourier Neural Operators [32; 33].

The PDE residual loss is often combined with additional penalty terms that encode boundary and initial conditions, smoothness priors, and other problem-specific constraints. This is necessary because minimizing the residual may not, in general, guarantee that the optimized solution satisfies the prescribed constraints or exhibits the desired regularity, especially when training data are sparse or the optimization landscape is ill-conditioned. A typical resulting composite objective is given by

$$\mathcal{L} = \mathcal{L}_{\text{PDE}} + \alpha\,\mathcal{L}_{\text{BC}} + \beta\,\mathcal{L}_{\text{Smoothness}} + \dots . \tag{S35}$$

This composite loss is then used as the objective for gradient-based optimization to train the neural network u_θ. Here, $\alpha, \beta, \ldots \in \mathbb{R}^+$ denote positive hyperparameters that weight the respective loss terms.

While PINNs have proven extremely successful in numerous practical applications, the corresponding optimization task can be challenging and prone to numerical issues. The training loss can have poor conditioning as it involves differential operators that can be ill-conditioned. In particular, Krishnapriyan et al. [37] showed that the loss landscape becomes increasingly complex and harder to optimize as the physics loss coefficient increases. The model could also converge to a trivial or non-desired solution that satisfies the physics laws on the set of points where the physics loss is computed [38]. There could also be conflicts between the multiple loss terms when their gradients point in opposite directions. Even without such conflicts, the losses can vary significantly in magnitude, leading to unbalanced gradients during training [39], and thus to the contribution and reduction of certain losses being relatively negligible during training. Consequently, obtaining reliable PINN solutions can require substantial care throughout the training process, including the choice of network architecture, collocation sampling strategy, loss weights, training hyperparameters, optimizer, and optimization strategy. In practice, these components must be tuned and monitored carefully to ensure that the optimized solution reflects the intended physical behavior rather than artifacts of the optimization procedure.

2.1.2 Achieving High-Precision with PINNs

To make PINN solutions amenable to computer-assisted proofs, we must achieve very high accuracy when solving the Euler and Navier–Stokes equations. In this paper, we leverage several mechanisms that substantially improve the precision attainable with PINNs and enable the construction of sufficiently accurate approximate solutions.

Higher Floating Point Precision. 64-bit floating point precision (FP64) with PINNs helps preserve accuracy when computing automatic-differentiation derivatives and PDE residuals, where FP32 round-off can accumulate and destabilize training. With cleaner gradients and more reliable physics losses, FP64 often improves convergence and yields higher-precision solutions that better satisfy the governing equations and physics constraints [43]. We carry out PINN training in FP64 to ensure our solution losses are evaluated with the highest numerical fidelity that is possible in PyTorch.

Boosting. Our methodology is inspired by the existing literature on PINN boosting [20; 21]. We write our solution approximation as a stage-wise additive model,

$$u(x,t) \approx f_{\theta_0} \tag{S36}$$

where each newly added PINN f_{θ_k} is introduced only after the current surrogate has reached an optimization plateau. The stage-k PINN f_{θ_k} is trained *conditioned* on the previously optimized sum: the parameters of the earlier PINNs $f_{\theta_0}, \ldots, f_{\theta_{k-1}}$ are frozen, and only f_{θ_k} is optimized so that it serves as a correction targeting the residual structure left. In effect, boosting replaces a single, difficult global fitting task with a sequence of easier residual-fitting subproblems. The successive correction networks can continue to reduce the PDE residual beyond the saturation point of a single PINN, yielding higher-accuracy solutions. Because each stage solves a distinct residual-correction problem, one can vary the PINN architecture and training strategy across stages, tailoring f_{θ_k} and its optimization to the specific residual features it is intended to capture. This flexibility also allows later-stage residual networks to be localized to small subdomains or patches where errors are concentrated, or to be combined naturally with domain-decomposition strategies that assign different correction networks to different regions of the computational domain.

Self-Scaled Curvature-Aware Optimizers. Adam [22] is widely used for neural-network training, but its reliance on first-order gradient information can limit convergence in PINNs, whose loss landscapes are often stiff and poorly conditioned. Curvature-aware optimization can alleviate these difficulties and continue reducing the objective after first-order methods begin to stagnate, often leading to substantially higher solution accuracy, though this comes at an increased computational cost. Quasi-Newton methods achieve this by constructing an approximation to the inverse Hessian from successive parameter and gradient updates. However, these approximations may themselves become poorly conditioned, degrading the quality of the resulting search directions.

Self-scaling addresses this issue by adaptively rescaling the inverse-Hessian approximation to better capture the local geometry of the objective landscape. In PINNs, such self-scaled quasi-Newton updates have been shown to improve conditioning and accelerate convergence [25].

More recently, structured curvature-aware optimizers such as Shampoo [44] and SOAP [45] have provided a more scalable means of exploiting curvature information. Rather than maintaining a dense inverse-Hessian approximation, they construct structured preconditioners that capture important correlations in the gradient while remaining computationally practical for larger neural networks. In particular, SOAP performs adaptive optimization in the eigenbasis of the Shampoo preconditioner, yielding updates that are better aligned with the local geometry of the objective. Standard Shampoo and SOAP, however, do not have the self-scaling mechanism.

Recently, SS–eSOAP [23] was introduced to augment SOAP with self-scaling and an adaptive eigenbasis update. This retains much of the computational efficiency of first-order optimization while providing a more effective curvature-aware update. In our experiments, SS–eSOAP incurs only a 15% increase in computational cost relative to Adam, while substantially improving the PDE MSE residual from $\mathcal{O}(10^{-4})$ with Adam to $\mathcal{O}(10^{-6})$. For further refinement, we employ the self-scaled Broyden (SS-Broyden) method [24; 25]. SS-Broyden maintains a more expressive approximation of the inverse Hessian and updates both its scaling and Broyden parameters using local curvature information. This richer curvature model comes at a substantially higher computational and memory cost than SS–eSOAP, making SS-Broyden less attractive as a general-purpose optimizer throughout training. However, once the solution has entered a sufficiently accurate regime, the improved curvature approximation can continue reducing the objective beyond the level at which the cheaper optimizer begins to saturate. We use SS-Broyden as a final high-accuracy refinement stage, reducing the PDE MSE residual in our experiments from $\mathcal{O}(10^{-6})$ with SS–eSOAP to $\mathcal{O}(10^{-10})$.

2.1.3 Dealing with Unbounded Domains

Sampling on an unbounded domain is a fundamental challenge for numerical methods. Similar to [13; 15], we use the sinh transformation (and the chain rule for derivatives), to sample over the half-plane

$$\mathbb{D} = \{(r, z) : r \geq 0,\ z \in \mathbb{R}\}. \tag{S37}$$

More precisely, we sample the computational variables on the finite box

$$\big(r_{\mathtt{comp}}, z_{\mathtt{comp}}\big) \in [0, 30] \times [-30, 30], \tag{S38}$$

and map them to the physical variables by

$$(r, z) = \big(\sinh r_{\mathtt{comp}},\ \sinh z_{\mathtt{comp}}\big). \tag{S39}$$

Under the sinh mapping, the transformation is approximately linear near the origin, since $\sinh(x) \sim x$ as $x \to 0$. Uniform sampling in the computational variables therefore produces nearly uniform, finely resolved sampling in the physical domain near the origin, where the blowup profile is expected to vary most rapidly. Farther from the origin, the mapping progressively stretches the sampling points, enabling the same computational grid to cover a large portion of the physical domain without unnecessarily high resolution in the far field. This distribution is well suited to the anticipated solution structure, which exhibits sharp variation near the blowup region and a much slower decay toward zero in the far field.

Under the sinh mapping, since $\sinh(30) \approx 5.34 \times 10^{12}$, the computational box $[0, 30] \times [-30, 30]$ maps roughly to $[0,\, 5.34 \times 10^{12}] \times [-5.34 \times 10^{12},\, 5.34 \times 10^{12}]$, which is effectively unbounded for practical purposes. The restriction to nonnegative values of r is consistent with the assumed even parity in r, so the solution can be mirrored across the axis $r = 0$ (or enforced as a hard constraint in the PINN architecture). The odd symmetry of the sinh map is particularly convenient here, as it preserves the underlying parity structure when expressed in the computational variables.

The choice of sinh also avoids some of the drawbacks associated with alternative transformations. For example, a compactifying map such as $r = \operatorname{arctanh} \rho$, with $0 \le \rho < 1$, maps the infinite physical domain to the finite interval $\rho \in [0, 1)$, but becomes singular as $\rho \to 1$, since $dr/d\rho = 1/(1 - \rho^2)$. Polynomial maps avoid this singular behavior, but their algebraic growth is typically too slow to sample the physical far field efficiently. The sinh map therefore provides a natural compromise: it reaches large physical scales rapidly while remaining smooth and preserving parity through its odd symmetry.

We represent the PINN in the computational variables and obtain derivatives with respect to the physical variables via the chain rule. When the computational and physical coordinates are related by the sinh map introduced above, we use the same notation u_Θ for the induced representation in physical space:

$$u_\Theta(r, z) := u_\Theta\big(r_{\mathsf{comp}}, z_{\mathsf{comp}}\big). \tag{S40}$$

The first derivatives in physical coordinates are then

$$\frac{\partial u_\Theta}{\partial r} = \frac{\partial u_\Theta}{\partial r_{\mathsf{comp}}} \operatorname{sech}\big(r_{\mathsf{comp}}\big), \qquad \frac{\partial u_\Theta}{\partial z} = \frac{\partial u_\Theta}{\partial z_{\mathsf{comp}}} \operatorname{sech}\big(z_{\mathsf{comp}}\big). \tag{S41}$$

The same principle applies to higher-order derivatives: the transformed expressions involve higher-order chain-rule coefficients, built from smooth bounded functions such as sech and tanh on the computational box, and are evaluated using automatic differentiation [46], as is standard in PINN implementations [47]. We give the full transformed PDE residual operator in Appendix B.

2.1.4 PINN Objectives and Constraints

Let u_Θ be a neural-network approximation to the PDE solution with parameters Θ, being the collection of all the network parameters to be optimized and the traveling wave speed C. To optimize the PINN representation, we minimize a weighted combination of loss terms

$$\mathcal{L}(\Theta) \;=\; \mathcal{L}_{\text{PDE}}(\Theta) \;+\; \alpha\, \mathcal{L}_{\mathsf{FF}}(\Theta) \;+\; \beta\, \mathcal{L}_{\text{Smoothness}}(\Theta) \;+\; \gamma\, \mathcal{L}_{\text{Non-triviality}}(\Theta), \tag{S42}$$

where

- $\mathcal{L}_{\text{PDE}}(\Theta) = \int_{\mathbb{D}} \big\|\mathcal{R}\big(u_\Theta\big)(x)\big\|^2 \, dx$ denotes the PDE residual loss, where $\mathcal{R}$ is the PDE residual operator.
- $\mathcal{L}_{\mathsf{FF}}(\Theta)$ denotes the far field loss conditions that $u,\, \omega, \psi \to 0$ as $r, z \to \infty$. We emphasize that these are decay conditions at infinity rather than classical boundary conditions (i.e Dirichlet, Neumann, Robin) at a finite boundary.
- $\mathcal{L}_{\text{Smoothness}}(\Theta) = \int_{\mathbb{D}} \big\|\nabla \mathcal{R}\big(u_\Theta\big)(x)\big\|^2 \, dx$ penalizes the gradient of the PDE residual. This promotes smoothness of the optimized representation. Heuristically, controlling $\nabla \mathcal{R}$ discourages highly oscillatory residuals and biases the optimized solution toward higher regularity on the region where the penalty is enforced. The gradient is computed via automatic differentiation.
- $\mathcal{L}_{\text{Non-triviality}}(\Theta)$ is an additional loss term designed to discourage the optimization from converging to trivial solutions.

In contrast to the preceding constraints, which are encouraged softly through the loss function, we enforce the following symmetry exactly as a hard constraint. Because our self-similar ansatz travels along the z-axis, no symmetry condition is imposed in z. Instead, we impose even parity in r:

$$u(r,z) = u(-r,z), \quad \omega(r,z) = \omega(-r,z), \quad \psi(r,z) = \psi(-r,z). \tag{S43}$$

This parity condition is physically motivated by Luo and Hou [8], where, in the presence of a boundary, an analogous symmetry serves as a trapping mechanism that promotes vorticity growth toward singularity formation.

2.1.5 Sampling Collocation Points

To improve generalization and avoid overfitting to a fixed set of collocation points, we periodically resample training points and bias sampling toward the origin, where the solution varies most rapidly. Since the solution is even in r, we restrict sampling to $r \geq 0$. We also use residual-based adaptive sampling [48], progressively adding points in regions where the PDE residual is largest. Separate samples are used for the PDE, smoothness, and far-field losses, with additional points placed along the symmetry axis $r = 0$, where the formally singular radial term is evaluated using its regular limiting form.

For SS-Broyden, we use the same overall strategy but resample less frequently, since full-batch second-order methods are more sensitive to changes in the training set through their use of past iterates in Hessian approximations. We also find that sufficiently large batches are important for good generalization across the computational domain.

2.2 Numerical PINN Profiles

Leveraging the methods described in Section 2.1, we parametrize and train a PINN solution for the Euler self-similar profile system. We visualize the optimized approximate profiles u, ω, and ψ as three-dimensional surfaces and display the corresponding $r = 0$ cross-sections in Figure S1. The optimized solution reproduces the qualitative profile structure reported by Hou et al. [19], showing that the self-similar profile remains robust even when convection is fully amplified. These profiles were obtained to high accuracy, as can be seen from the profile equations residuals reported in Table S1.

Remark. *We also investigate the family of self-similar profiles for the convection-varied Euler system introduced in Appendix A, obtained by varying the convection amplitude $\varepsilon \in [0, 1]$. We observe that the value of λ decreases steadily towards $\lambda = 1/2$ as ε increases towards $\varepsilon = 1$, as expected.*

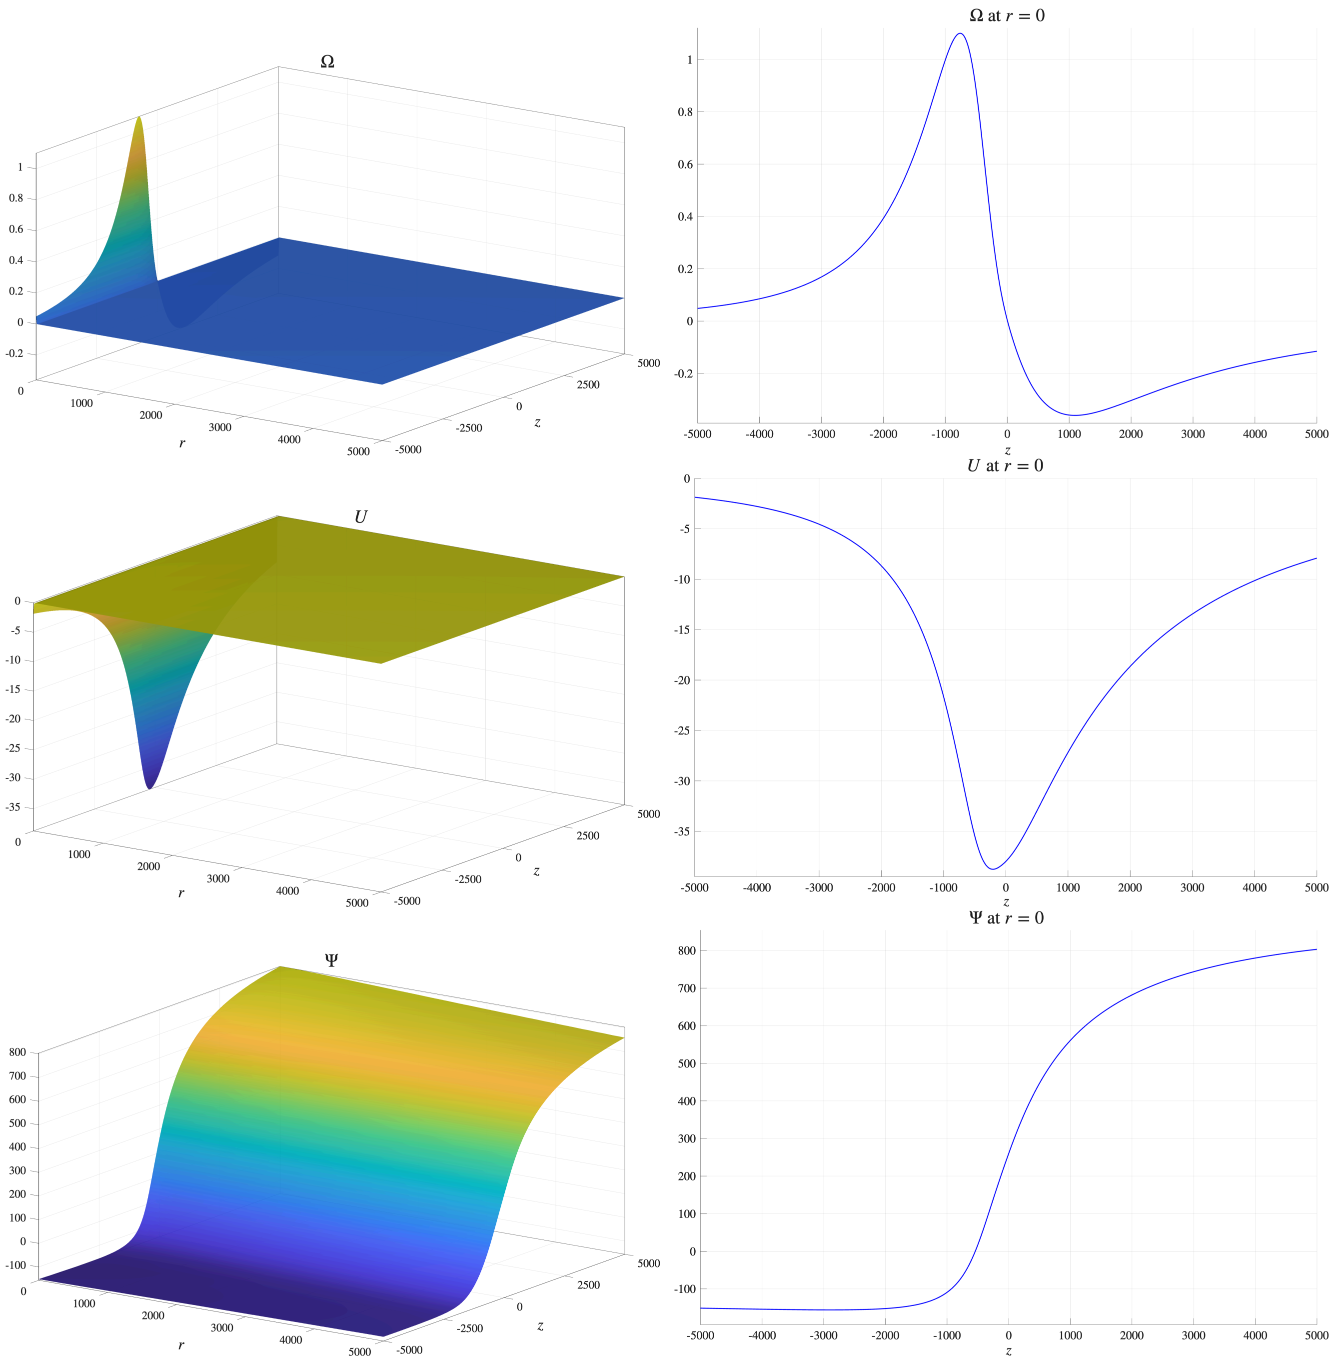

Figure S1: Visualization of the approximate profiles Ω, U, Ψ as surface plots for the 3D axisymmetric Euler system. The corresponding $r = 0$ cross-sections.

2.3 Analytic Spline Representation of the Steady State

The PINN computation serves as a numerical discovery tool, providing a practical and well-conditioned means of obtaining an accurate approximate self-similar profile when direct optimization over analytic ansatzes or alternative basis expansions would be considerably more difficult. We denote steady-state quantities by an overhead bar, and let

$$(\bar{u}_{\text{PINN}}, \bar{\psi}_{\text{PINN}}, \bar{\omega}_{\text{PINN}})$$

denote the profile produced by the PINN.

Table S1: Residual errors for the U, Ω, Ψ profile equations for the Euler system, comparing PINN and spline representations. Note that the residuals are estimated empirically on an independent set of 1,000,000 randomly-sampled collocation points for the PINN representations, whereas they are evaluated analytically and exactly for the spline representations. Here, $\mathbb{D}$ denotes the full domain (i.e. the half plane), $B_3(0)$ the ball of radius 3 centered at the origin, axis the $r = 0$ axis, and FF the far-field region.

		PINNs		**Splines**	
	Domain	RMSE	ℓ^∞	L^2	L^∞
U-equation	$\mathbb{D}$	$4.1{\cdot}10^{-5}$	$3.6{\cdot}10^{-3}$	$4.1{\cdot}10^{-5}$	$3.6{\cdot}10^{-3}$
	$B_3(0)$	$1.6{\cdot}10^{-5}$	$5.7{\cdot}10^{-5}$	$1.7{\cdot}10^{-5}$	$5.7{\cdot}10^{-5}$
	axis	$9.1{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$	$9.1{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$
	FF	$6.2{\cdot}10^{-6}$	$2.2{\cdot}10^{-5}$	$6.2{\cdot}10^{-6}$	$2.2{\cdot}10^{-5}$
Ω-equation	$\mathbb{D}$	$1.7{\cdot}10^{-5}$	$9.3{\cdot}10^{-4}$	$1.7{\cdot}10^{-5}$	$9.3{\cdot}10^{-4}$
	$B_3(0)$	$4.3{\cdot}10^{-6}$	$1.2{\cdot}10^{-5}$	$4.3{\cdot}10^{-6}$	$1.2{\cdot}10^{-5}$
	axis	$4.8{\cdot}10^{-5}$	$3.7{\cdot}10^{-4}$	$4.8{\cdot}10^{-5}$	$3.7{\cdot}10^{-4}$
	FF	$3.1{\cdot}10^{-6}$	$1.4{\cdot}10^{-5}$	$3.1{\cdot}10^{-6}$	$1.4{\cdot}10^{-5}$
Ψ-equation	$\mathbb{D}$	$3.8{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$	$3.8{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$
	$B_3(0)$	$7.2{\cdot}10^{-6}$	$2.5{\cdot}10^{-5}$	$7.8{\cdot}10^{-6}$	$2.5{\cdot}10^{-5}$
	axis	$8.6{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$	$8.6{\cdot}10^{-5}$	$1.4{\cdot}10^{-3}$
	FF	$2.1{\cdot}10^{-6}$	$8.3{\cdot}10^{-6}$	$2.1{\cdot}10^{-6}$	$8.3{\cdot}10^{-6}$

In practice, however, the PINN solution is not the steady state we work over for our stability analysis and numerical certification. This is due to the fact that the basis optimized by the PINN can be highly nonlinear and oscillatory, making it difficult to interpret. We instead freeze an explicit analytic profile $(\bar{u}, \bar{\psi}, \bar{\omega})$, together with certified intervals for the fitted parameters. All residuals, energy estimates, and nonlinear constants are then evaluated using this fitted profile rather than the original neural network. This different representation allows for certified residuals and norms, which is essential in our stability argument as it involves the repeatedly differentiating, calculating weighted perturbations, and finding L^∞ bounds of the numerical profile. These operations can be certified for explicit analytic representations (via interval arithmetic), whereas certifying them directly for the PINN would require a separate verification of the network evaluation, automatic differentiation procedure, and L^∞ bounds.

We must now choose an analytic representation that is sufficiently expressive to reproduce the computed profile with high accuracy, while remaining sufficiently explicit for rigorous certification. For this purpose, we use *piecewise polynomial splines*. This choice is common in numerical analysis and is particularly well suited to verification because, on each subinterval, the approximation is given by an explicit low-degree polynomial, so that its derivatives, extrema, and residuals can be evaluated and bounded using standard rigorous tools, including interval arithmetic. Moreover, the local nature of the representation allows the mesh to be refined only where the profile exhibits rapid variation without unnecessarily increasing the complexity of the approximation elsewhere. In contrast, global polynomial or spectral representations have to find a tradeoff between global and local behavior which may require a relatively high degree of polynomials, leading to larger expressions and less convenient rigorous bounds. Piecewise splines avoid this, achieving high local accuracy with moderate polynomial degree, preserve smoothness across subintervals, and yield a representation whose errors and derivatives can be certified in an more interpretable and computationally efficient manner.

Overall, the certification pipeline has four separate layers, as described in the diagram below, where only the first arrow is a numerical approximation step:

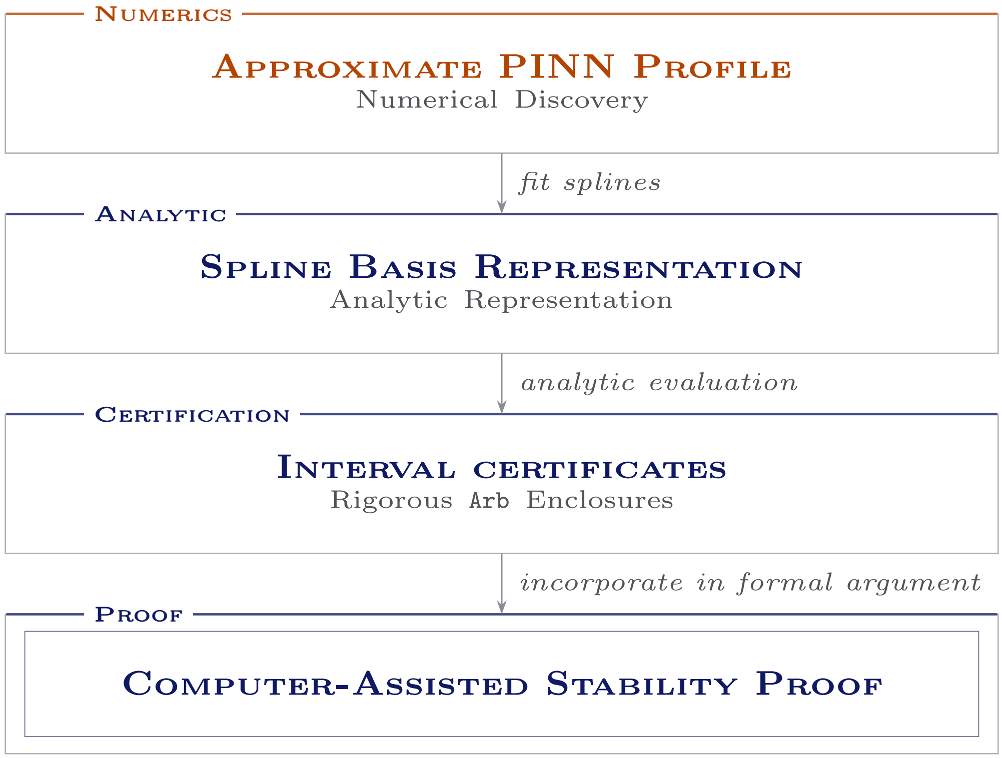

We define $\mathcal{R}_u^{\mathsf{raw}}$ and $\mathcal{R}_\omega^{\mathsf{raw}}$ to be the PDE residuals accounting for the fact that the numerical profiles do not satisfy exactly the Euler equations:

$$\mathcal{R}_u^{\mathsf{raw}} := -(\lambda r + \bar{u}^r)\partial_r \bar{u} - (\lambda z + \bar{C} + \bar{u}^z)\partial_z \bar{u} + 2\bar{u}\partial_z \bar{\psi} + \bar{c}_u \bar{u}, \tag{S44}$$

$$\mathcal{R}_\omega^{\mathsf{raw}} := -(\lambda r + \bar{u}^r)\partial_r \bar{\omega} - (\lambda z + \bar{C} + \bar{u}^z)\partial_z \bar{\omega} + 2\bar{u}\partial_z \bar{u} + (\bar{c}_u - \lambda)\bar{\omega}. \tag{S45}$$

The spline fit preserves the small profile-equation residuals observed in the numerical solution, so the passage to an analytic representation incurs only a minor loss of accuracy. For the stability analysis, we further construct the surrogate through a two-head fit: $(\bar{u}, \bar{\psi})$ are fitted separately, while $\bar{\omega}$ is defined by applying the elliptic operator exactly to $\bar{\psi}$. Consequently, the residual in the elliptic equation is exactly 0, eliminating one source of certification error and significantly simplifying the subsequent stability analysis.

3 Stability

A singularity is *stable* if the associated blowup mechanism persists under small perturbations of the initial data, in the sense that the same mechanism occurs for an open neighborhood of nearby initial conditions. Conversely, an *unstable* singularity requires infinitely precise initial conditions: small perturbations deflect the evolution away from the singularity mechanism.

The dynamic rescaling formulation presented in Section 3.1 is the natural starting point for such a stability analysis. In the original physical variables, the solution may grow and concentrate as $t \to T^-$, so there is no fixed object around which to linearize. After rescaling, the approximate blowup profile becomes a steady state of the rescaled equations, while the blowup rate, translation rate, and amplitude normalization are recorded by modulation parameters. Linear stability can therefore be studied by perturbing this steady rescaled profile and analyzing the resulting linearized rescaled dynamics.

We follow the same overall structure as [3; 17] and related works on stability of blowups. We first compute the blowup profile numerically, then establish linear stability for the dynamically rescaled equation in suitable weighted norms and prove nonlinear stability using Sobolev embeddings and weighted estimates. The main distinction is that the weight functions and other tunable parameters in the stability proof are optimized numerically to ensure that the argument closes, rather than being chosen and refined through expert intuition and a potentially long and technically challenging process of trial and error. The overall proof architecture is summarized in the diagram on page 50, while the optimization strategy used to close the proof is described in detail in Section 3.9.

Remark. *An alternative sufficient criterion is to prove invertibility of the full linearized profile operator [49]. In this case, the self-similar profile equation is viewed perturbatively around an approximate profile, and the goal is to invert the linearized operator after removing the neutral scaling direction. This is stronger than the criterion used here. In contrast, the numerical eigenvalue analysis in [41, Section 4.4] gives a more limited, finite-dimensional test of the same stability mechanism: it consists of discretizing the dynamically rescaled linearized equation around the computed profile and inspecting the spectrum of the resulting Jacobian. This provides useful evidence for stability, but is not by itself a rigorous substitute for the required linear estimate. It only concerns a finite-dimensional discretization and therefore controls, at best, resolved low-frequency perturbations. As noted in [50, Section 1.3.1], the proof must control the infinite-dimensional operator, including high-frequency modes, nonlocal velocity terms, possible neutral directions, and transient growth caused by non-normality. Thus, negative real parts of the computed eigenvalues do not yield the uniform estimate required for nonlinear stability.*

3.1 Dynamic Rescaling Equations

3.1.1 Introduction

The steady profile equations above describe an exact traveling self-similar solution. In computations and stability arguments, however, one usually does not know the correct profile, or traveling speed in advance. The *dynamic rescaling formulation* addresses this by letting the rescaled solution evolve while continuously adjusting its location, scale, and amplitude. The key idea is to replace the finite-time singular behavior in physical variables by a long-time evolution in normalized variables. The singularity is kept at order-one size and near a fixed location in the rescaled coordinates. In this frame, blowup no longer appears primarily as unbounded growth or shrinking length scale. Instead, a stable blowup scenario should appear as convergence toward a steady state. That steady-state is precisely the traveling-wave profile described in Proposition S2.

Modulated coordinates and amplitudes. In the dynamic formulation, the fixed self-similar parameters in the ansatz are replaced by time-dependent *modulation parameters*. Rather than prescribing a blowup time, a spatial scale, and a traveling center in advance, the rescaled frame is adjusted as the solution evolves. More precisely, let $(\mathring{r}, \mathring{z}, \mathring{t})$ denote the physical variables and let (r, z, t) denote the dynamically rescaled variables. We introduce a common spatial scale $\mathrm{s}_r(t) > 0$, amplitude scales $\mathrm{s}_u(t)$, $\mathrm{s}_\omega(t) > 0$, and an axial center $\mathring{z}_c(t)$ by

$$\mathring{r} = \mathrm{s}_r(t)\, r, \qquad \mathring{z} - \mathring{z}_c(t) = \mathrm{s}_r(t)\, z, \tag{S46}$$

and

$$u_\bullet(\mathring{r}, \mathring{z}, \mathring{t}) = \mathrm{s}_u(t)\, u(r, z, t), \qquad \omega_\bullet(\mathring{r}, \mathring{z}, \mathring{t}) = \mathrm{s}_\omega(t)\, \omega(r, z, t). \tag{S47}$$

Then,

$$\psi_\bullet(\mathring{r}, \mathring{z}, \mathring{t}) = \mathrm{s}_\omega(t)\, \mathrm{s}_r(t)^2\, \psi(r, z, t). \tag{S48}$$

The normalization of the transport, stretching, and time-derivative terms is fixed by

$$\mathrm{s}_u(t) = \mathrm{s}_\omega(t)\, \mathrm{s}_r(t), \qquad \frac{dt}{d\mathring{t}} = \mathrm{s}_\omega(t)\, \mathrm{s}_r(t). \tag{S49}$$

We then define the rescaling rates by

$$\partial_t \mathrm{s}_r(t) = -\lambda\, \mathrm{s}_r(t), \qquad \partial_t \mathrm{s}_u(t) = -c_u(t)\, \mathrm{s}_u(t), \qquad \partial_t \mathrm{s}_\omega(t) = -c_\omega(t)\, \mathrm{s}_\omega(t). \tag{S50}$$

and define the axial drift via

$$\partial_t \mathring{z}_c(t) = -C(t)\, \mathrm{s}_r(t). \tag{S51}$$

Overall, this adjustment is encoded by the coordinate drift

$$(\partial_t r)_{\mathring{r},\mathring{z}} = \lambda r, \qquad (\partial_t z)_{\mathring{r},\mathring{z}} = \lambda z + C(t). \tag{S52}$$

where

- λ is the *spatial rescaling rate*. It measures, in the current rescaled time variable, how quickly the computational frame expands
- the function $C(t)$ is the *instantaneous traveling speed* of the axial center in the same rescaled frame.
- the term λz is the axial part of the dilation, while $C(t)$ is the additional axial translation.

The amplitude of the solution must also be normalized. This is achieved via two scalar functions $c_u(t)$ and $c_\omega(t)$, which are the logarithmic amplitude rates for u and ω. The $c_u(t)u$ and $c_\omega(t)\omega$ terms appearing in the rescaled evolution equations do not represent new forcing, they only record the rate at which the normalization of the rescaled variables is changed in order to keep the profile at order-one amplitude.

Use and connection with stability. The profile formulation and the dynamic rescaling formulation play complementary roles. The profile equations describe the limiting singularity once the blowup rate, traveling speed, and amplitude rates have settled to constants, while the dynamic equations describe the evolution toward that limiting object. Numerically, this is essential: instead of resolving a structure whose amplitude diverges and whose length scale collapses in physical variables, one follows a normalized solution whose main profile remains order one in a fixed computational frame. The modulation parameters provide direct diagnostics for the blowup scenario: if $C(t)$, $c_u(t)$, and $c_\omega(t)$ converge, their limits give the axial traveling speed, and the amplitude growth rates. If the rescaled profiles converge at the same time, the computation gives more than evidence of singular growth. It gives a candidate self-similar profile, the values of the scaling parameters, and the profile equations that the limiting object should satisfy.

The dynamic rescaling formulation is the natural setting for stability. In the original physical variables, a perturbation of a blowup solution can change the blowup time, shift the center location, or alter the amplitude normalization. Such changes may appear as growth or drift even when the underlying profile is stable. The normalization conditions remove these neutral directions by continuously recentering, rescaling, and renormalizing the solution. After this gauge freedom has been fixed, stability becomes a question about the evolution of genuine perturbations in the dynamically rescaled variables. A stable traveling self-similar singularity is represented by a stable steady state of the dynamic rescaling equations: solutions starting close to the profile should remain controlled, the modulation parameters should approach limiting constants, and the rescaled solution should converge to the steady profile. In this formulation, finite-time blowup in the original variables is linked to convergence toward a stable fixed point of the evolution.

3.1.2 Dynamic Rescaling Equations for the Euler Equations

As before, define the elliptic operator

$$\mathcal{E} := \partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2. \tag{S53}$$

The proposition below provides the axisymmetric Euler equations in the dynamically rescaled variables.

Proposition S3 (Dynamic rescaling equations for the Euler)**.** *In the dynamically rescaled variables, the axisymmetric system takes the form*

$$\partial_t u + \big(\lambda r + u^r\big)\partial_r u + \big(\lambda z + C(t) + u^z\big)\partial_z u = 2u\partial_z\psi + c_u(t)u, \tag{S54}$$
$$\partial_t \omega + \big(\lambda r + u^r\big)\partial_r \omega + \big(\lambda z + C(t) + u^z\big)\partial_z \omega = 2u\partial_z u + c_\omega(t)\omega, \tag{S55}$$
$$-\mathcal{E}\psi = \omega, \tag{S56}$$

with

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi, \tag{S57}$$

and $\lambda = 1/2$.

Proof The derivation is provided in Appendix D. ■

Normalization conditions. When $\lambda = 1/2$ is fixed, the independent modulation parameters are $C(t)$ and $c_u(t)$. The vorticity amplitude rate is determined by

$$c_\omega(t) = c_u(t) - \frac{1}{2}. \tag{S58}$$

Before the stability analysis, write the radial and axial transport coefficients of the numerically obtained profile as

$$G_r(r,z) = \lambda r + u^r_{\mathsf{num}}(r,z), \qquad G_z(r,z) = \lambda z + C_{\mathsf{num}} + u^z_{\mathsf{num}}(r,z). \tag{S59}$$

We choose an axial shift z_{shift} so that the point $(0, z_{\mathsf{shift}})$ of the numerical profile becomes the analysis origin. For this point to be a fixed point of the meridional flow, we require

$$G_r(0, z_{\mathsf{shift}}) = 0, \qquad G_z(0, z_{\mathsf{shift}}) = 0. \tag{S60}$$

The radial condition is automatic because $u^r_{\mathsf{num}}(0,z)=0$ on the axis. Using $u^z_{\mathsf{num}}(0,z)=2\psi_{\mathsf{num}}(0,z)$, the axial condition $G_z(0,z_{\mathsf{shift}})=0$ becomes the scalar equation

$$C_{\mathsf{num}}+\lambda z_{\mathsf{shift}}+2\psi_{\mathsf{num}}(0,z_{\mathsf{shift}})=0. \tag{S61}$$

We impose this condition and choose z_{shift} as a numerically certified root. Having fixed z_{shift}, we recenter each numerical profile field by

$$\bar f(r,z):=f_{\mathsf{num}}(r,z+z_{\mathsf{shift}}). \tag{S62}$$

Denoting the constant term in the recentered axial transport by $\bar C$, we continue to write the reference transport coefficients as

$$G_r(r,z)=\lambda r+\bar u^r(r,z),\qquad G_z(r,z)=\lambda z+\bar C+\bar u^z(r,z). \tag{S63}$$

With this notation, the imposed root condition gives

$$G_r(0,0)=0,\qquad G_z(0,0)=\bar C+2\bar\psi_0=0. \tag{S64}$$

After applying this shift, set $\bar u_0:=\bar u(0,0)$. We preserve the meridional fixed point and the amplitude at the origin by imposing

$$C(t)+u^z(0,0,t)=0,\qquad u(0,0,t)=\bar u_0. \tag{S65}$$

The first condition keeps the reference meridional fixed point at the origin. The second keeps the amplitude fixed there and therefore gives $\partial_t u(0,0,t)=0$.

Limiting behaviour. If the dynamic rescaling captures a stable traveling self-similar blowup, the rescaled solution should converge as rescaled time advances to a time-independent profile,

$$u(r,z,t)\to u_\infty(r,z),\qquad \omega(r,z,t)\to\omega_\infty(r,z),\qquad \psi(r,z,t)\to\psi_\infty(r,z), \tag{S66}$$

and the modulation parameters should converge to constants,

$$C(t)\to C_\infty,\qquad c_u(t)\to c_{u,\infty},\qquad c_\omega(t)\to c_{\omega,\infty}. \tag{S67}$$

In that limit, the time derivatives in the dynamic rescaling equations disappear. The steady dynamic equations are

$$\Big(\lambda\, r+u^r_\infty\Big)\partial_r u_\infty+\Big(\lambda\, z+C_\infty+u^z_\infty\Big)\partial_z u_\infty=2u_\infty\,\partial_z\psi_\infty+c_{u,\infty}\,u_\infty, \tag{S68}$$
$$\Big(\lambda\, r+u^r_\infty\Big)\partial_r \omega_\infty+\Big(\lambda\, z+C_\infty+u^z_\infty\Big)\partial_z \omega_\infty=2u_\infty\,\partial_z u_\infty+c_{\omega,\infty}\,\omega_\infty, \tag{S69}$$

with $u^r_\infty=-r\partial_z\psi_\infty$ and $u^z_\infty=2\psi_\infty+r\partial_r\psi_\infty$.

For the traveling-wave ansatz in (S13)–(S15), the limiting amplitude rates are

$$c_{u,\infty}=-1,\qquad c_{\omega,\infty}=-(1+\lambda)=-3/2. \tag{S70}$$

By writing

$$\big(u_\infty,\,\omega_\infty\,,\psi_\infty,\,C_\infty\big)=\big(U,\,\Omega,\,\Psi,\,C\big), \tag{S71}$$

and substituting (S70) into (S68)–(S69), we see that the limiting dynamic equations agree with the corresponding profile equations. In this sense, the traveling-wave profiles are fixed points of the appropriate dynamically rescaled evolution.

3.2 Linearization

In the present section, we derive the perturbation equations around a steady state of the dynamically rescaled system. We then separate the terms that are linear in the perturbation from the higher-order remainders. The linearized equations identify the operator whose stability determines whether the candidate blowup profile is attracting to first order in the rescaled variables.

3.2.1 Perturbations

We denote perturbations using δ's in blue. We use the subscript $_0$ to denote evaluation at the origin $(r, z) = (0, 0)$. We perturb a given steady state

$$\bar{u}, \bar{\omega}, \bar{\psi}, \bar{C}, \bar{c}_u, \bar{c}_\omega \tag{S72}$$

of the dynamic rescaling equations, by writing

$$u = \bar{u} + \delta u, \qquad \omega = \bar{\omega} + \delta\omega, \qquad \psi = \bar{\psi} + \delta\psi, \tag{S73}$$

$$C = \bar{C} + \delta C, \qquad c_u = \bar{c}_u + \delta c_u, \qquad c_\omega = \bar{c}_\omega + \delta c_\omega. \tag{S74}$$

The perturbation variables measure the difference between the evolving rescaled solution and the steady blowup profile. The goal of linear stability analysis is to determine whether these perturbations remain bounded, or preferably decay, under the dynamics obtained by linearizing around the steady state.

The perturbations must inherit the symmetries of the underlying steady state and therefore belong to the same functional class. In particular, the parity and regularity conditions imposed on the rescaled variables are inherited by δu, $\delta\omega$, and $\delta\psi$, with the latter two coupled by the elliptic equation relating vorticity and streamfunction. The perturbations must also decay at the far-field. The admissible perturbation class includes the weighted regularity, integrability, trace, and boundary-flux assumptions used in the energy estimates. In particular, the cutoff boundary-flux terms vanish as the cutoffs are removed. This ensures that the perturbation equations remain in the same functional class as the steady profile.

For notational convenience, we will use δ as shorthand for the full list of perturbation variables,

$$\delta := \{\delta u, \delta\omega, \delta\psi, \delta C, \delta c_u, \delta c_\omega\}. \tag{S75}$$

We also write the meridional velocity as

$$u^r = \bar{u}^r + \delta u^r, \qquad u^z = \bar{u}^z + \delta u^z, \tag{S76}$$

where

$$\delta u^r = -r\,\partial_z \delta\psi, \qquad \delta u^z = 2\,\delta\psi + r\,\partial_r \delta\psi, \tag{S77}$$

and also have

$$\bar{c}_\omega = \bar{c}_u - \lambda, \qquad \delta c_\omega = \delta c_u. \tag{S78}$$

Modulations and Normalization Conditions. The modulation parameters are fixed by normalization conditions at the origin. We impose

$$C(t) + u^z_0(t) = 0, \qquad u_0(t) = \bar{u}_0. \tag{S79}$$

For the perturbations, these conditions give

$$\delta C + 2\delta\psi_0 = 0, \qquad \delta u_0 = 0, \qquad (\partial_t \delta u)_0 = 0. \tag{S80}$$

These normalization conditions fix the axial translation and amplitude normalization associated with $C(t)$ and $c_u(t)$. The spatial rate $\lambda = 1/2$ has already been fixed. They are essential for formulating linear stability in the modulated variables, since otherwise the perturbation could drift along symmetry directions rather than measuring a genuine change in the profile.

3.2.2 Choice of Variables for the Linear Stability Analysis

The reduced variables u and ω do not transform in the same way under rescaling (see Appendix F). Indeed, the scaling of the full three-dimensional fields induces the following scaling of the reduced variables:

$$u_\eta(r,z,t) = \eta^{\alpha+1} u(\eta r, \eta z, \eta^{\alpha+1} t), \qquad \omega_\eta(r,z,t) = \eta^{\alpha+2} \omega(\eta r, \eta z, \eta^{\alpha+1} t). \tag{S81}$$

The relative scaling of u, ω, and ∇u is unchanged. Thus ω carries one extra factor of η relative to u, reflecting the one-derivative relation between the underlying full velocity and vorticity fields. By contrast,

$$\omega_\eta(r,z,t) = \eta^{\alpha+2} \omega(\eta r, \eta z, \eta^{\alpha+1} t), \qquad \nabla u_\eta(r,z,t) = \eta^{\alpha+2} (\nabla u)(\eta r, \eta z, \eta^{\alpha+1} t). \tag{S82}$$

Therefore ω and ∇u are the natural pair of quantities in a scale-consistent linear stability analysis. Quantities with different scaling correspond to different differential orders, so comparing them directly mixes objects of different analytic strength. For this reason, instead of using the equations for $(\partial_t \delta\omega, \partial_t \delta u)$, the stability analysis is formulated using the equations for $\partial_t \delta\omega$, $\partial_t \partial_z \delta u$, $\partial_t \partial_r \delta u$. We collect the associated perturbations into the scale-consistent perturbation vector

$$\delta v := (\delta\omega, \nabla \delta u) = (\delta\omega, \partial_r \delta u, \partial_z \delta u), \tag{S83}$$

and for compact notation, we also set

$$\delta v_\omega := \delta\omega, \qquad \delta v_r := \partial_r \delta u, \qquad \delta v_z := \partial_z \delta u. \tag{S84}$$

3.2.3 Perturbation Equations

Linearization consists of expanding the perturbation equation around the candidate blowup profile and retaining only the terms that are first order in the perturbation. The resulting linearized equation gives the leading-order dynamics of small disturbances near the profile. The discarded terms are collected in nonlinear remainders, which record the higher-order interactions among the perturbations. The full derivations of the linearized and modulation equations are given in Appendix G.

We start from the dynamic rescaling equations (S54)–(S56):

$$\partial_t u + \big(\lambda r + u^r\big)\partial_r u + \big(\lambda z + C(t) + u^z\big)\partial_z u = 2u\partial_z \psi + c_u(t) u, \tag{S85}$$
$$\partial_t \omega + \big(\lambda r + u^r\big)\partial_r \omega + \big(\lambda z + C(t) + u^z\big)\partial_z \omega = 2u\partial_z u + c_\omega(t) \omega, \tag{S86}$$
$$-\mathcal{E}\psi = \omega. \tag{S87}$$

Here,

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi, \tag{S88}$$

and $\mathcal{E}$ is the elliptic operator

$$\mathcal{E} := \partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2. \tag{S89}$$

Substituting (S73)–(S74), subtracting the steady-state equations, re-centering the modulation coefficients as specified below, and separating linear terms from nonlinear terms in the perturbations gives

$$\partial_t\delta u = \mathcal{L}_u(\delta) + \mathcal{N}_u(\delta) \; + \; \mathcal{R}_u, \qquad \partial_t\delta\omega = \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) \; + \; \mathcal{R}_\omega. \tag{S90}$$

The raw profile residuals $\mathcal{R}_u^{\mathsf{raw}}$ and $\mathcal{R}_\omega^{\mathsf{raw}}$ account for the fact that the numerical profiles are not exact:

$$\mathcal{R}_u^{\mathsf{raw}} = -(\lambda r + \bar{u}^r)\partial_r\bar{u} - (\lambda z + \bar{C} + \bar{u}^z)\partial_z\bar{u} + 2\bar{u}\partial_z\bar{\psi} + \bar{c}_u\bar{u}, \tag{S91}$$

$$\mathcal{R}_\omega^{\mathsf{raw}} = -(\lambda r + \bar{u}^r)\partial_r\bar{\omega} - (\lambda z + \bar{C} + \bar{u}^z)\partial_z\bar{\omega} + 2\bar{u}\partial_z\bar{u} + (\bar{c}_u - \lambda)\bar{\omega}. \tag{S92}$$

We also define the residual amplitude offset

$$c_{u,\mathrm{res}} := -\frac{\mathcal{R}_u^{\mathsf{raw}}(0)}{\bar{u}_0}, \tag{S93}$$

under the nondegeneracy condition $\bar{u}_0 \neq 0$.

We absorb this fixed offset into the reference amplitude rate. From this point onward, δC and δc_u denote the variable modulation coefficients, so that

$$C = \bar{C} + \delta C, \qquad c_u = \bar{c}_u + c_{u,\mathrm{res}} + \delta c_u, \qquad c_\omega = \bar{c}_\omega + c_{u,\mathrm{res}} + \delta c_u. \tag{S94}$$

The residuals used in the perturbation equations are

$$\mathcal{R}_u := \mathcal{R}_u^{\mathsf{raw}} + c_{u,\mathrm{res}}\,\bar{u}, \tag{S95}$$

$$\mathcal{R}_\omega := \mathcal{R}_\omega^{\mathsf{raw}} + c_{u,\mathrm{res}}\,\bar{\omega}. \tag{S96}$$

The linearized u-operator and ω-operator are

$$\begin{aligned}\mathcal{L}_u(\delta) := &\big(2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}}\big)\,\delta u - \big(\lambda r + \bar{u}^r\big)\,\partial_r\delta u - \big(\lambda z + \bar{C} + \bar{u}^z\big)\,\partial_z\delta u + \big(2\bar{u} + r\,\partial_r\bar{u}\big)\partial_z\delta\psi \\ &- \big(r\,\partial_z\bar{u}\big)\partial_r\delta\psi - \big(2\partial_z\bar{u}\big)\delta\psi - (\partial_z\bar{u})\,\delta C + \bar{u}\,\delta c_u,\end{aligned} \tag{S97}$$

$$\begin{aligned}\mathcal{L}_\omega(\delta) := &(\bar{c}_u + c_{u,\mathrm{res}} - \lambda)\,\delta\omega - \big(\lambda r + \bar{u}^r\big)\partial_r\delta\omega - \big(\lambda z + \bar{C} + \bar{u}^z\big)\,\partial_z\delta\omega + 2(\partial_z\bar{u})\,\delta u + 2\bar{u}\,\partial_z\delta u \\ &- 2(\partial_z\bar{\omega})\,\delta\psi + r\,(\partial_r\bar{\omega})\,\partial_z\delta\psi - r\,(\partial_z\bar{\omega})\,\partial_r\delta\psi - (\partial_z\bar{\omega})\,\delta C + \bar{\omega}\,\delta c_u.\end{aligned} \tag{S98}$$

The nonlinear remainders are

$$\mathcal{L}_\omega(\delta) = r\,\partial_z\delta\psi\,\partial_r\delta u - (\delta C + (2\delta\psi + r\partial_r\delta\psi))\,\partial_z\delta u + 2\,\delta u\,\partial_z\delta\psi + \delta c_u\,\delta u, \tag{S99}$$

$$\mathcal{N}_\omega(\delta) = -\,\delta u^r\partial_r\delta\omega - (\delta C + \delta u^z)\,\partial_z\delta\omega + 2\,\delta u\,\partial_z\delta u + \delta c_u\,\delta\omega, \tag{S100}$$

The elliptic relation (S300) is already linear, so

$$-\mathcal{E}\delta\psi := -\left(\partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2\right)\delta\psi = \delta\omega. \tag{S101}$$

Consequently, $\delta\psi$ is not an independent dynamical unknown: it is recovered from $\delta\omega$ through this elliptic equation, and then δu^r and δu^z are recovered from (S77).

3.2.4 Differentiated Equations for the Scale-Consistent Variables

Since the scale-consistent variables are ω and ∇u, we differentiate the δu-equation:

$$\partial_t\partial_r\delta u = \mathcal{L}_r(\delta) + \mathcal{N}_r(\delta) \;+\; \mathcal{R}_r, \qquad \partial_t\partial_z\delta u = \mathcal{L}_z(\delta) + \mathcal{N}_z(\delta) \;+\; \mathcal{R}_z, \tag{S102}$$

where we denote

$$\mathcal{L}_r(\delta) := \partial_r\mathcal{L}_u(\delta), \qquad \mathcal{N}_r(\delta) := \partial_r\mathcal{N}_u(\delta), \qquad \mathcal{L}_z(\delta) := \partial_z\mathcal{L}_u(\delta), \qquad \mathcal{N}_z(\delta) := \partial_z\mathcal{N}_u(\delta), \tag{S103}$$

and

$$\mathcal{R}_z := \partial_z\mathcal{R}_u, \qquad \mathcal{R}_r := \partial_r\mathcal{R}_u. \tag{S104}$$

Differentiating the $\mathcal{L}_u$ expressions gives

$$\begin{aligned}\mathcal{L}_r(\delta) = &-(\lambda r + \bar{u}^r)\,\partial_{rr}\delta u - (\lambda z + \bar{C} + \bar{u}^z)\,\partial_{zr}\delta u + \big(2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}} - \lambda - \partial_r\bar{u}^r\big)\partial_r\delta u\\ &-(\partial_r\bar{u}^z)\,\partial_z\delta u + 2\partial_{rz}\bar{\psi}\,\delta u + (2\bar{u} + r\partial_r\bar{u})\,\partial_{zr}\delta\psi + \big(3\partial_r\bar{u} + r\partial_{rr}\bar{u}\big)\partial_z\delta\psi - r(\partial_z\bar{u})\,\partial_{rr}\delta\psi\\ &-\big(3\partial_z\bar{u} + r\partial_{rz}\bar{u}\big)\partial_r\delta\psi - 2\partial_{rz}\bar{u}\,\delta\psi - (\partial_{rz}\bar{u})\,\delta C + (\partial_r\bar{u})\,\delta c_u,\end{aligned} \tag{S105}$$

$$\begin{aligned}\mathcal{L}_z(\delta) = &-(\lambda r + \bar{u}^r)\,\partial_{rz}\delta u - (\lambda z + \bar{C} + \bar{u}^z)\,\partial_{zz}\delta u + \big(2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}} - \lambda - \partial_z\bar{u}^z\big)\partial_z\delta u\\ &-(\partial_z\bar{u}^r)\,\partial_r\delta u + 2\partial_{zz}\bar{\psi}\,\delta u + (2\bar{u} + r\partial_r\bar{u})\,\partial_{zz}\delta\psi + \big(r\partial_{rz}\bar{u}\big)\partial_z\delta\psi - r(\partial_z\bar{u})\,\partial_{rz}\delta\psi\\ &-r(\partial_{zz}\bar{u})\,\partial_r\delta\psi - 2\partial_{zz}\bar{u}\,\delta\psi - (\partial_{zz}\bar{u})\,\delta C + (\partial_z\bar{u})\,\delta c_u.\end{aligned} \tag{S106}$$

The differentiated nonlinear remainders are

$$\begin{aligned}\mathcal{N}_r(\delta) = &\;(\partial_z\delta\psi + r\partial_{rz}\delta\psi)\,\partial_r\delta u + r\partial_z\delta\psi\,\partial_{rr}\delta u - (3\partial_r\delta\psi + r\partial_{rr}\delta\psi)\,\partial_z\delta u\\ &-(\delta C + (2\delta\psi + r\partial_r\delta\psi))\,\partial_{zr}\delta u + 2(\partial_r\delta u)(\partial_z\delta\psi) + 2\delta u\,\partial_{zr}\delta\psi + \delta c_u\,\partial_r\delta u,\end{aligned} \tag{S107}$$

$$\begin{aligned}\mathcal{N}_z(\delta) = &\;r(\partial_{zz}\delta\psi)\,\partial_r\delta u + r(\partial_z\delta\psi)\,\partial_{rz}\delta u - (2\partial_z\delta\psi + r\partial_{rz}\delta\psi)\,\partial_z\delta u\\ &-(\delta C + (2\delta\psi + r\partial_r\delta\psi))\,\partial_{zz}\delta u + 2(\partial_z\delta u)(\partial_z\delta\psi) + 2\delta u\,\partial_{zz}\delta\psi + \delta c_u\,\partial_z\delta u,\end{aligned} \tag{S108}$$

3.2.5 Modulation Equations

The modulation equations are obtained in Appendix G by imposing the normalization conditions (S80). They express the perturbations of the rescaling parameters in terms of the profile and its perturbations evaluated at the origin. The translation formula is the perturbation of the stagnation condition $C + u_0^z = 0$. For the amplitude formula, we differentiate the normalization $u_0 = \bar{u}_0$ in time, evaluate the u-equation at the meridional fixed point, and solve the resulting identity for δc_u.

With the fixed residual amplitude offset already absorbed into the reference coefficient, the modulation coefficients are

$$\delta C = -2\delta\psi_0, \qquad\qquad \delta c_u = -2(\partial_z \delta\psi)_0. \tag{S109}$$

Both δC and δc_u are linear in the perturbation. Products such as $\delta C\, \partial_z \delta u$ and $\delta c_u\, \delta u$ remain ordinary quadratic nonlinear terms. These equations determine the two modulation variables from the perturbations.

3.2.6 Closed Linearized Variables

In terms of δv, the closed linearized system may be written schematically as

$$\partial_t \delta v = \mathcal{L}_\nabla \delta v, \tag{S110}$$

or equivalently in components:

$$\partial_t \begin{pmatrix} \delta\omega \\ \partial_r \delta u \\ \partial_z \delta u \end{pmatrix} = \mathcal{L}_\nabla \begin{pmatrix} \delta\omega \\ \partial_r \delta u \\ \partial_z \delta u \end{pmatrix}. \tag{S111}$$

The operator $\mathcal{L}_\nabla$ includes transport by the steady rescaled flow, coupling between δu and $\delta\omega$, elliptic recovery of $\delta\psi$, velocity reconstruction, and the linearized modulation terms.

3.3 Weighted Norms and Energy

3.3.1 Weighted Norms

The stability estimates are measured in weighted norms adapted to the candidate blowup profile and to the damping structure of the linearized operator. For this purpose, define

$$\Phi_\omega,\ \Phi_r,\ \Phi_z : \mathbb{D} \to (0, \infty) \tag{S112}$$

to be the positive weight functions assigned respectively to the three components of δv, on the half-plane

$$\mathbb{D} = \{(r, z) : r \geq 0,\ z \in \mathbb{R}\}. \tag{S113}$$

More explicitly, Φ_ω weights the vorticity perturbation $\delta\omega$, Φ_r weights $\partial_r \delta u$, and Φ_z weights $\partial_z \delta u$. We denote the collection of weights by

$$\Phi := (\Phi_\omega, \Phi_r, \Phi_z). \tag{S114}$$

For a positive scalar weight ϕ, define the associated inner product and norm by

$$\|f\|_\phi^2 := \int_{\mathbb{D}} |f|^2 \phi \, dr\, dz, \qquad \langle f, g\rangle_\phi := \int_{\mathbb{D}} f g \, \phi \, dr\, dz. \tag{S115}$$

Given vector quantities $q := (q_\omega, q_r, q_z)$, and $\tilde{q} := (\tilde{q}_\omega, \tilde{q}_r, \tilde{q}_z)$, we define the corresponding Φ-weighted norm and Φ-weighted inner product by

$$\|q\|_\Phi^2 := \|q_\omega\|_{\Phi_\omega}^2 + \|q_r\|_{\Phi_r}^2 + \|q_z\|_{\Phi_z}^2. \tag{S116}$$

$$\langle q, \tilde{q}\rangle_\Phi := \langle q_\omega, \tilde{q}_\omega\rangle_{\Phi_\omega} + \langle q_r, \tilde{q}_r\rangle_{\Phi_r} + \langle q_z, \tilde{q}_z\rangle_{\Phi_z}. \tag{S117}$$

3.3.2 Low-Order Weighted Energy.

The low-order energy uses the three singular weights Φ_ω, Φ_r, Φ_z, and is defined as the Φ-norm of δv:

$$\mathfrak{E}_0(t) := \|\delta v(t)\|_\Phi, \qquad \text{i.e.} \quad \mathfrak{E}_0(t)^2 = \|\delta\omega(t)\|_{\Phi_\omega}^2 + \|\partial_r \delta u(t)\|_{\Phi_r}^2 + \|\partial_z \delta u(t)\|_{\Phi_z}^2. \tag{S118}$$

The three weights Φ_ω, Φ_r, Φ_z are kept separate because the damping calculation acts on the three components of δv in different ways. The purpose of the weighted energy is to make the coercive part of the linear part produce a positive damping constant $\Lambda_{\mathcal{L}} > 0$. A discussion of the constraints that the weight functions Φ_ω, Φ_r, and Φ_z must satisfy in our setting is provided in Section 3.3.5.

While the low-order energy remains the main source of damping, it is not sufficient by itself to close the estimates since the linear and nonlinear terms contain terms which cannot be controlled using a weighted L^2 norm, such as terms that must be controlled in L^∞, as well as modulation factors and pointwise quantities evaluated at the origin, such as $(\partial_z \delta u)_0$ and $(\partial_z \delta\omega)_0$.

3.3.3 High-Order Weighted Energy.

To control the linear and nonlinear terms which cannot be controlled in the low-order weighted energy $\mathfrak{E}_0$, we supplement $\mathfrak{E}_0$ with a high-order weighted energy $\mathfrak{H}_k$, chosen so that the top derivatives provide the required pointwise, interpolation, elliptic, and modulation bounds.

The high-order part uses the *higher-order weight functions*

$$\Phi_{i,k} := 1 + \rho^k \Phi_i, \qquad \rho(r,z) := r^2 + z^2, \qquad i \in \{\omega, r, z\}. \tag{S119}$$

The factor $\rho^k = |(r,z)|^{2k}$ compensates for the loss of k powers under differentiation near the origin. If f has local order $|(r,z)|^m$, then $D^k f$ has local order $|(r,z)|^{m-k}$; consequently, the factors $|D^k f|^2 \rho^k \Phi_i$ and $|f|^2 \Phi_i$ have the same local power of $|(r,z)|$. The added 1 prevents the high-order weight from degenerating at the origin and gives the direct bound

$$\Phi_{i,k} = 1 + \rho^k \Phi_i \geq 1, \qquad \|D^k f\|_{L^2(\mathbb{D})} \leq \|D^k f\|_{\Phi_{i,k}}. \tag{S120}$$

We then define the *high-order weighted energy*

$$\mathfrak{H}_k^2 := \sum_{|\alpha|=k} \|D^\alpha \delta\omega\|_{\Phi_{\omega,k}}^2 + \sum_{|\alpha|=k} \|D^\alpha \partial_r \delta u\|_{\Phi_{r,k}}^2 + \sum_{|\alpha|=k} \|D^\alpha \partial_z \delta u\|_{\Phi_{z,k}}^2. \tag{S121}$$

3.3.4 Full Energy.

The full energy used for the stability argument is obtained as

$$\mathfrak{E}_k^2 := \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2, \qquad 0 < \mu_k \le 1. \tag{S122}$$

The parameter μ_k is chosen only small enough so the low-order damping absorbs the low-order loss coming from the differentiated linear estimate, but also large enough so that the differentiated damping absorbs the top-order loss left by the low-order linear estimate.

Choosing k. The order of the top derivative k should be chosen large enough so that the energy controls all point values, products, commutators, and elliptic or streamfunction terms used in the estimates. The modulation formulas require the origin values

$$\delta\psi_0, \qquad (\partial_z \delta\psi)_0, \qquad (\mathcal{E}\delta u)_0. \tag{S123}$$

In two dimensions, a point value of a derivative of order m is controlled by Sobolev norms through order $m+2$. The largest origin derivative in (S123) has order one, so the modulation formulas require $k \ge 3$. Independently, the reverse allocation in the top-order transport commutators places a factor of derivative order at most two in L^∞, and the corresponding H^2-to-L^∞ estimate again uses derivatives through order four. **The lowest admissible choice is thus $\boldsymbol{k = 4}$**. Larger choices of k are possible, but they increase combinatorially the number of profile derivatives and analytic constants that must be certified.

Remark on intermediate derivative orders. Intermediate derivative orders $0 < j < k$ do not need to be included explicitly in the definition of the full energy. Indeed, for each such j, we derive in our companion paper [18] the certified interpolation bound

$$\sum_{|\alpha|=j} \|D^\alpha f\|_{\Phi_{i,j}} \le \eta_{i,j} \sum_{|\alpha|=k} \|D^\alpha f\|_{\Phi_{i,k}} + I_{i,j,k}(\eta_{i,j}) \|f\|_{\Phi_i}. \tag{S124}$$

Thus, every intermediate-order term is controlled by a combination of the top-order and low-order energy levels. In each product estimate, these two contributions are kept separate until Young's inequality is applied: the top-order term is absorbed by the available $\mathfrak{H}_k^2$ damping, while the low-order term is controlled by $\mathfrak{E}_0^2$. The finite family of parameters $\eta_{i,j}$ is chosen so that the total top-order interpolation contribution remains strictly below the available differentiated damping.

3.3.5 Constraints on the Weight Functions

Each weight function Φ_i for $i \in \{\omega, r, z\}$, is required to satisfy the following conditions.

1. **Positivity and coercivity.** There is a constant $C_{\text{floor}} > 0$ such that

$$\Phi_i(r, z) \ge C_{\text{floor}} \qquad \text{for almost every } (r, z) \in \mathbb{D}. \tag{S125}$$

   This prevents the weighted norm from becoming weak in a region where the perturbation must still be controlled.

2. **Near-origin admissibility.** Each weight function Φ_i may have its own singularity rate at the origin. We require only local integrability, without imposing a common exponent or a comparison between different components. For certification, one convenient sufficient condition is

$$\Phi_i(r,z) \le C_i(r^2+z^2)^{-p_i}, \qquad 0 < r^2+z^2 \le \rho_i, \qquad 0 \le p_i < 1, \tag{S126}$$

where C_i, ρ_i, and p_i are chosen separately for each component. The condition $p_i < 1$ guarantees local integrability in two dimensions. This power-law model is only a sufficient criterion, and a direct local-integrability certificate may be used instead.

3. **Regularity.** Away from the origin, C^1 regularity is enough for the basic weighted integrations by parts. When a later estimate differentiates Φ_i or $\Phi_{i,k}$ further, the required derivatives are assumed to exist on that region and to satisfy the corresponding certified bounds. In the computer-assisted implementation, we use smooth parametrizations with C^∞ activation functions, which meet these requirements and are convenient for gradient-based optimization.

4. **Far-field growth needed for velocity control.** There are constants $\rho_{\mathsf{growth}}, \mathsf{c}_{\mathsf{growth}} > 0$ such that

$$\min\{\Phi_r, \Phi_z\} \ge \mathsf{c}_{\mathsf{growth}}\,\rho \qquad \text{whenever } \rho \ge \rho_{\mathsf{growth}}, \qquad \rho = r^2+z^2. \tag{S127}$$

This lower growth condition is used to recover the unweighted L^2 control of δu needed in the global pointwise estimate.

5. (Optional) **Normalization.** When the weight functions are optimized via gradient-based optimization, it can be useful to fix a normalization for each weight function to remove the irrelevant scaling freedom.

6. (Optional) **Evenness in r.** Since the pointwise estimates use reflection across the axis, it is convenient to choose the weights with the same even symmetry in r:

$$\Phi_i(r,z) = \Phi_i(-r,z). \tag{S128}$$

This symmetry can be built directly into the weight parametrization.

3.4 Linear Stability Estimates

3.4.1 Low-Order Energy Linear Stability and the Need for Higher-Order Energy

Linear stability asks whether perturbations of solutions of the closed linearized system remain bounded or decay in a suitable norm. In terms of δv, the desired estimate is

$$\|\delta v(t)\|_\Phi \le Ce^{-\gamma t}\|\delta v(0)\|_\Phi, \qquad \gamma > 0. \tag{S129}$$

A standard rigorous way to prove (S129) is to establish the differential energy estimate

$$\frac{1}{2}\frac{d}{dt}\mathfrak{E}_0(t)^2 = \langle \mathcal{L}_\nabla \delta v, \delta v\rangle_\Phi \le -\Lambda_{\mathcal{L}}\mathfrak{E}_0(t)^2. \tag{S130}$$

This estimate gives a quantitative estimate for every solution of the closed linearized perturbation equation in the specified weighted space. Integrating (S130) gives exponential decay of $\mathfrak{E}_0(t)$, and hence of the scale-consistent perturbation δv, rather than merely a numerical or visual indication of convergence.

For a blowup profile, this is the relevant notion of stability: the original physical solution may still blow up, but the normalized spatial profile is stable in the rescaled variables. Linear stability means that $\delta u(t) \to 0$ and $\delta\omega(t) \to 0$ under the linearized rescaled flow. Note that the blowup time, blowup location, and scaling parameters may change under perturbation, but these changes are absorbed by the modulation parameters.

In spectral language, the idea is that the relevant spectrum of the closed linearized operator $\mathcal{L}_\nabla$ should lie in the stable half-plane. In a numerical stability calculation, one discretizes $\mathcal{L}_\nabla$ and studies the eigenvalues of the resulting matrix. Eigenvalues with negative real part correspond to decaying linear modes, while eigenvalues with positive real part correspond to growing directions. Neutral modes may arise from translation or amplitude normalization. The modulation conditions are designed to fix these directions, so that the stability question concerns perturbations transverse to the artificial symmetry directions. Thus, the role of the linearization is to produce the operator $\mathcal{L}_\nabla$ governing the first-order evolution of the scale-consistent perturbation δv. The role of the linear stability analysis is to show that this operator has no growing directions in the weighted energy $\mathfrak{E}_0$, once the elliptic constraint, the velocity reconstruction, and the modulation conditions are imposed. If such an estimate is proved, then the proposed rescaled profile is linearly stable as a blowup profile.

However, the low-order energy is not sufficient by itself to close the linear estimates since the linear operators contain terms which cannot be controlled using a weighted L^2 norm. It is thus supplemented with the high-order weighted energy $\mathfrak{H}_k$, to get an estimate of the form

$$\frac{1}{2}\frac{d}{dt}\mathfrak{E}_0(t)^2 = \langle \mathcal{L}_\nabla \delta v, \delta v\rangle_\Phi \le -\Lambda_{\mathcal{L}}^{\mathrm{low}}\mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}}\mathfrak{H}_k^2. \tag{S131}$$

3.4.2 Low-Order Linear Estimates

The low-order linear estimate to certify in the Euler setting is

$$\langle \mathcal{L}_\omega, \delta\omega\rangle_{\Phi_\omega} + \langle \partial_r\mathcal{L}_u, \partial_r\delta u\rangle_{\Phi_r} + \langle \partial_z\mathcal{L}_u, \partial_z\delta u\rangle_{\Phi_z} \le -\Lambda_{\mathcal{L}}^{\mathrm{low}}\,\mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}}\,\mathfrak{H}_k^2. \tag{S132}$$

Here the constant $\Lambda_{\mathcal{L}}^{\mathrm{low}} > 0$ is the net low-order linear damping and $\Lambda_{\mathcal{L}}^{\mathrm{high}} \ge 0$ is the top-order loss from direct estimates that require $\mathfrak{H}_k$.

3.4.3 High-Order Linear Estimates

We need to certify a top-order damping constant $\Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} > 0$ and a low-order constant $\Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}} \ge 0$ such that

$$\begin{aligned}\sum_{|\alpha|=k}\langle D^\alpha\mathcal{L}_\omega, D^\alpha\delta\omega\rangle_{\Phi_{\omega,k}} + \sum_{|\alpha|=k}\langle D^\alpha\partial_r\mathcal{L}_u, D^\alpha\partial_r\delta u\rangle_{\Phi_{r,k}} + \sum_{|\alpha|=k}\langle D^\alpha\partial_z\mathcal{L}_u, D^\alpha\partial_z\delta u\rangle_{\Phi_{z,k}}\\ \le \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}}\,\mathfrak{E}_0^2 - \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}}\,\mathfrak{H}_k^2.\end{aligned} \tag{S133}$$

3.5 Nonlinear Stability Estimates

The linear stability analysis shows that the rescaled profile is stable under the first-order perturbation dynamics, measured in the weighted low-order norm associated with Φ_ω, Φ_r, and Φ_z. The nonlinear stability argument upgrades this statement to the full perturbation equation. Its goal is to prove that, if the initial perturbation is sufficiently small in a suitable energy, then the perturbation remains small for as long as the estimates apply. In the case of an approximate numerical profile, the solution is allowed to remain in a small neighborhood whose size also depends on the certified residual error.

The nonlinear stability proof is reduced to an energy estimate for the full energy $\mathfrak{E}_k$. Differentiating $\mathfrak{E}_k^2$ in time and inserting the perturbation equations separates the argument into five distinct components: the residual error, the low-order linear part, the high-order linear part, the low-order nonlinear part, and the high-order nonlinear part. The residual records the defect of the approximate steady profile. The linear terms contain the stabilizing mechanism: the low-order estimate gives the principal damping, while the high-order estimate controls the differentiated equations, allowing only losses that can be absorbed by the low-order energy after choosing μ_k. The nonlinear terms are then estimated perturbatively, with the low-order and high-order contributions bounded by cubic and quartic powers of the full energy.

3.5.1 Low-Order Nonlinear Estimates

We start from the expressions (S100), (S107), (S108), for the nonlinear parts $\mathcal{N}_i$ for $i \in \{\omega, r, z\}$.

We substitute the exact modulation formulas before separating linear and nonlinear terms:

$$\delta C = -2\delta\psi_0, \qquad \delta c_u = -2(\partial_z \delta\psi)_0. \tag{S134}$$

The precise low-order nonlinear estimate to prove is

$$|\langle \mathcal{N}_\omega, \delta\omega\rangle_{\Phi_\omega} + \langle \mathcal{N}_r, \partial_r \delta u\rangle_{\Phi_r} + \langle \mathcal{N}_z, \partial_z \delta u\rangle_{\Phi_z}| \le \Lambda_{\mathcal{N},3}\, \mathfrak{E}_k^3. \tag{S135}$$

3.5.2 High-Order Nonlinear Estimates

The precise high-order nonlinear estimate to prove is

$$\mu_k \left| \sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{N}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \right| \le \Lambda_{D^\alpha\mathcal{N},3}\, \mathfrak{E}_k^3. \tag{S136}$$

3.6 PDE Residuals

If the $(\bar u, \bar\omega, \bar\psi, \bar\lambda, \bar C, \bar c_u)$ profile is exact, the raw residuals below vanish. If the profile is numerical, these are the defects obtained by substituting the profile into the steady equations:

$$\mathcal{R}_u^{\mathsf{raw}} := -(\lambda r + \bar u^r)\partial_r \bar u - (\lambda z + \bar C + \bar u^z)\partial_z \bar u + 2\bar u \partial_z \bar\psi + \bar c_u \bar u, \tag{S137}$$

$$\mathcal{R}_\omega^{\mathsf{raw}} := -(\lambda r + \bar u^r)\partial_r \bar\omega - (\lambda z + \bar C + \bar u^z)\partial_z \bar\omega + 2\bar u \partial_z \bar u + (\bar c_u - \lambda)\bar\omega. \tag{S138}$$

The residuals entering the perturbation equations are

$$\mathcal{R}_u := \mathcal{R}_u^{\mathsf{raw}} + c_{u,\mathrm{res}}\,\bar{u}, \tag{S139}$$
$$\mathcal{R}_\omega := \mathcal{R}_\omega^{\mathsf{raw}} + c_{u,\mathrm{res}}\,\bar{\omega}. \tag{S140}$$

The elliptic residual is not listed here because the spline representation of $\bar{\omega}$ is obtained by applying the elliptic operator to $\bar{\psi}$ exactly, as discussed in Section 2.3. As a result, the steady-state profile satisfies the elliptic equation exactly, and the elliptic residual is exactly 0.

Define

$$\mathcal{R}_r := \partial_r \mathcal{R}_u, \qquad \mathcal{R}_z := \partial_z \mathcal{R}_u. \tag{S141}$$

Then the low-order perturbation equations are

$$\partial_t \delta\omega = \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) + \mathcal{R}_\omega, \tag{S142}$$
$$\partial_t \partial_r \delta u = \partial_r \mathcal{L}_u(\delta) + \mathcal{N}_r(\delta) + \mathcal{R}_r, \tag{S143}$$
$$\partial_t \partial_z \delta u = \partial_z \mathcal{L}_u(\delta) + \mathcal{N}_z(\delta) + \mathcal{R}_z. \tag{S144}$$

Define the residual constants by separating the low-order and top-order pieces:

$$\Lambda_{\mathcal{R},\Phi}{}^2 := \|\mathcal{R}_\omega\|_{\Phi_\omega}^2 + \|\mathcal{R}_r\|_{\Phi_r}^2 + \|\mathcal{R}_z\|_{\Phi_z}^2, \tag{S145}$$
$$\Lambda_{D^\alpha\mathcal{R}}{}^2 := \sum_{|\alpha|=k} \Big(\|D^\alpha \mathcal{R}_\omega\|_{\Phi_{\omega,k}}^2 + \|D^\alpha \mathcal{R}_r\|_{\Phi_{r,k}}^2 + \|D^\alpha \mathcal{R}_z\|_{\Phi_{z,k}}^2 \Big), \tag{S146}$$
$$\Lambda_{\mathcal{R}}{}^2 := \Lambda_{\mathcal{R},\Phi}{}^2 + \mu_k \Lambda_{D^\alpha\mathcal{R}}{}^2. \tag{S147}$$

Applying the Cauchy–Schwarz inequality gives

$$\left| \sum_{i\in\{\omega,r,z\}} \langle \mathcal{R}_i, \delta v_i \rangle_{\Phi_i} + \mu_k \sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{R}_i, D^\alpha \delta v_i \rangle_{\Phi_{i,k}} \right| \le \Lambda_{\mathcal{R}}\, \mathfrak{E}_k. \tag{S148}$$

3.7 Stability Theorems

We are now ready to state the stability results for the Euler equations. Importantly, every named certification constant attached to an analytic estimate is understood to be a rigorously validated finite numerical value. Conceptually, the analytic argument identifies the constants that are needed, while the numerical step certifies them once for the chosen profiles, weights, and parameters. The stability proofs are not complete until appropriate constants have been certified numerically.

3.7.1 Single-Radius Stability Theorems

We now state the full stability theorem.

Theorem S1 (Single-radius stability)**.** *Assume that the following perturbation equations hold for admissible perturbations:*

$$\partial_t \delta\omega = \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) + \mathcal{R}_\omega, \tag{S149}$$

$$\partial_t \partial_r \delta u = \partial_r \mathcal{L}_u(\delta) + \mathcal{N}_r(\delta) + \mathcal{R}_r, \tag{S150}$$

$$\partial_t \partial_z \delta u = \partial_z \mathcal{L}_u(\delta) + \mathcal{N}_z(\delta) + \mathcal{R}_z. \tag{S151}$$

Assume that the following estimates have been certified with finite constants:

$$\langle \mathcal{L}_\omega, \delta\omega \rangle_{\Phi_\omega} + \langle \partial_r \mathcal{L}_u, \partial_r \delta u \rangle_{\Phi_r} + \langle \partial_z \mathcal{L}_u, \partial_z \delta u \rangle_{\Phi_z} \le -\Lambda_{\mathcal{L}}^{\mathrm{low}}\, \mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}}\, \mathfrak{H}_k^2. \tag{S152}$$

$$\sum_{i \in \{\omega, r, z\}} \sum_{|\alpha| = k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i \rangle_{\Phi_{i,k}} \le \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}}\, \mathfrak{E}_0^2 - \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{high}}\, \mathfrak{H}_k^2. \tag{S153}$$

$$|\langle \mathcal{N}_\omega, \delta\omega \rangle_{\Phi_\omega} + \langle \mathcal{N}_r, \partial_r \delta u \rangle_{\Phi_r} + \langle \mathcal{N}_z, \partial_z \delta u \rangle_{\Phi_z}| \le \Lambda_{\mathcal{N},3}\, \mathfrak{E}_k^3. \tag{S154}$$

$$\mu_k \left| \sum_{i \in \{\omega, r, z\}} \sum_{|\alpha| = k} \langle D^\alpha \mathcal{N}_i, D^\alpha \delta v_i \rangle_{\Phi_{i,k}} \right| \le \Lambda_{D^\alpha \mathcal{N},3}\, \mathfrak{E}_k^3. \tag{S155}$$

$$\left| \sum_{i \in \{\omega, r, z\}} \langle \mathcal{R}_i, \delta v_i \rangle_{\Phi_i} + \mu_k \sum_{i \in \{\omega, r, z\}} \sum_{|\alpha| = k} \langle D^\alpha \mathcal{R}_i, D^\alpha \delta v_i \rangle_{\Phi_{i,k}} \right| \le \Lambda_{\mathcal{R}}\, \mathfrak{E}_k. \tag{S156}$$

Choose μ_k so that $0 < \mu_k \le 1$ and

$$\Lambda_{\mathrm{stab}}(\mu_k) := \min \left\{ \Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}},\ \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k} \right\} > 0. \tag{S157}$$

After this choice of μ_k is fixed, write Λ_{stab} for the positive number $\Lambda_{\mathrm{stab}}(\mu_k)$. Define the combined nonlinear constants by

$$\Lambda_3 := \Lambda_{\mathcal{N},3} + \Lambda_{D^\alpha \mathcal{N},3}. \tag{S158}$$

Assume that $\delta_* > 0$ satisfies

$$\Lambda_3 \delta_* + \frac{\Lambda_{\mathcal{R}}}{\delta_*} < \Lambda_{\mathrm{stab}}. \tag{S159}$$

Then every solution with $\mathfrak{E}_k(0) < \delta_*$ remains in the ball $\mathfrak{E}_k(t) < \delta_*$ for as long as the solution exists and the certified estimates apply.

Proof The proof is presented in Appendix I. ■

The diagram below summarizes the strategy used for proving nonlinear stability.

PERTURBATIONS

PERTURBATION EQUATIONS

$$\partial_t \delta\omega = \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) + \mathcal{R}_\omega, \quad \partial_t \partial_r \delta u = \partial_r \mathcal{L}_u(\delta) + \mathcal{N}_r(\delta) + \mathcal{R}_r, \quad \partial_t \partial_z \delta u = \partial_z \mathcal{L}_u(\delta) + \mathcal{N}_z(\delta) + \mathcal{R}_z.$$

$-\mathcal{E}\delta\psi = \delta\omega$, with modulation coefficients δC and δc_u

build the weighted energy

WEIGHTS AND ENERGY

STABILITY ENERGY

$$\|f\|_\phi^2 := \int_{\mathbb{D}} |f|^2 \phi \, dr \, dz, \quad \langle f, g\rangle_\phi := \int_{\mathbb{D}} f g \, \phi \, dr \, dz, \qquad \Phi_{i,k} := 1 + \rho^k \Phi_i, \quad \rho := r^2 + z^2.$$

The singular base weights are $\Phi_\omega, \Phi_r, \Phi_z$, while $\Phi_{i,k}$ are the compensated top-order weights.

Low-Order Energy $\quad \mathfrak{E}_0^2 := \|\delta\omega\|_{\Phi_\omega}^2 + \|\partial_r \delta u\|_{\Phi_r}^2 + \|\partial_z \delta u\|_{\Phi_z}^2,$

Top-Order Energy $\quad \mathfrak{H}_k^2 := \sum_{|\alpha|=k} \|D^\alpha \delta\omega\|_{\Phi_{\omega,k}}^2 + \sum_{|\alpha|=k} \|D^\alpha \partial_r \delta u\|_{\Phi_{r,k}}^2 + \sum_{|\alpha|=k} \|D^\alpha \partial_z \delta u\|_{\Phi_{z,k}}^2,$

Full Energy $\quad \mathfrak{E}_k^2 := \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2, \qquad 0 < \mu_k \le 1$

derive the analytic ingredients

THEORY

ANALYTIC RESULTS AND ESTIMATES

Sobolev extension and reflection theorems, Sobolev embeddings and interpolation, weighted transport, and Leibniz estimates control lower-order factors and intermediate derivatives.

Elliptic and streamfunction PDE bounds, modulation estimates, and certified operator and residual constants reduce all pairings to explicit bounds.

obtain linear, nonlinear, and residual bounds

LINEAR

LOW-ORDER AND HIGH-ORDER LINEAR ESTIMATES

$$\langle \mathcal{L}_\omega, \delta\omega\rangle_{\Phi_\omega} + \langle \partial_r \mathcal{L}_u, \partial_r \delta u\rangle_{\Phi_r} + \langle \partial_z \mathcal{L}_u, \partial_z \delta u\rangle_{\Phi_z} \le -\Lambda_{\mathcal{L}}^{\mathrm{low}} \mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}} \mathfrak{H}_k^2,$$

$$\sum_{i \in \{\omega, r, z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \le \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}} \mathfrak{E}_0^2 - \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{high}} \mathfrak{H}_k^2$$

NONLINEAR

LOW-ORDER AND HIGH-ORDER NONLINEAR ESTIMATES

$$\left| \langle \mathcal{N}_\omega, \delta\omega\rangle_{\Phi_\omega} + \langle \mathcal{N}_r, \partial_r \delta u\rangle_{\Phi_r} + \langle \mathcal{N}_z, \partial_z \delta u\rangle_{\Phi_z} \right| \le \Lambda_{\mathcal{N},3} \mathfrak{E}_k^3$$

$$\mu_k \left| \sum_{i \in \{\omega, r, z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{N}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \right| \le \Lambda_{D^\alpha \mathcal{N},3} \mathfrak{E}_k^3$$

PDE RESIDUAL

PROFILE PDE RESIDUAL ESTIMATES

$$\left| \sum_{i \in \{\omega, r, z\}} \langle \mathcal{R}_i, \delta v_i\rangle_{\Phi_i} + \mu_k \sum_{i \in \{\omega, r, z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{R}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \right| \le \Lambda_{\mathcal{R}} \mathfrak{E}_k$$

choose coupling and bootstrap radius

CLOSURE

CHOOSE μ_k AND THE STABILITY RADIUS δ_*

$$\Lambda_{\mathrm{stab}}(\mu_k) := \min\left\{ \Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}}, \; \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k} \right\} > 0,$$

$$\text{choose} \quad \delta_* > 0 \quad \text{with} \quad \left(\Lambda_{\mathcal{N},3} + \Lambda_{D^\alpha \mathcal{N},3}\right) \delta_* + \frac{\Lambda_{\mathcal{R}}}{\delta_*} < \Lambda_{\mathrm{stab}}$$

conclude the stability proof

STABILITY

STABILITY

If $\mathfrak{E}_k(0) < \delta_*$, then $\mathfrak{E}_k(t) < \delta_*$ for every later admissible time.
Thus the ball $\{\mathfrak{E}_k < \delta_*\}$ is forward invariant, which is the desired nonlinear stability statement

3.7.2 Two-Radii Stability Theorems

The single-radius certificate can be relaxed by prescribing separate target bounds for the low-order and top-order energies. The same weighted energy argument still applies, provided the weighted ball is chosen small enough to lie inside the desired two-radii region,

$$\mathfrak{E}_0(t) < \delta_0, \qquad \mathfrak{H}_k(t) < \delta_k, \tag{S160}$$

with $\delta_0 > 0$ and $\delta_k > 0$. These bounds are obtained from the same full energy $\mathfrak{E}_k^2 = \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2$.

For a chosen weight $\mu_k > 0$, the invariant set is the weighted ball

$$\{\mathfrak{E}_k(t) < \delta\}. \tag{S161}$$

This ball is contained in the target set (S160) provided that

$$0 < \delta \leq \min\{\delta_0, \sqrt{\mu_k}\,\delta_k\}, \tag{S162}$$

since this implies

$$\mathfrak{E}_0(t) \leq \mathfrak{E}_k(t) < \delta \leq \delta_0, \qquad \mathfrak{H}_k(t) \leq \frac{\mathfrak{E}_k(t)}{\sqrt{\mu_k}} < \frac{\delta}{\sqrt{\mu_k}} \leq \delta_k. \tag{S163}$$

Theorem S2 (Two-radii nonlinear stability)**.** *Assume that the perturbation equations* (S149)–(S151) *hold for admissible perturbations. Fix $\delta_0 > 0$ and $\delta_k > 0$. Choose $\mu_k > 0$ and $\delta > 0$ satisfying*

$$0 < \delta \leq \min\{\delta_0, \sqrt{\mu_k}\,\delta_k\}. \tag{S164}$$

Assume that the same certified estimates (S152)–(S156) *of Theorem S1 hold, with the weighted energy*

$$\mathfrak{E}_k^2 = \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2. \tag{S165}$$

Assume that μ_k satisfies the stability condition

$$\Lambda_{\mathrm{stab}}(\mu_k) := \min\left\{\Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}},\ \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k}\right\} > 0. \tag{S166}$$

After this choice of μ_k is fixed, write Λ_{stab} for the positive number $\Lambda_{\mathrm{stab}}(\mu_k)$. Define

$$\Lambda_3 := \Lambda_{\mathcal{N},3} + \Lambda_{D^\alpha\mathcal{N},3}. \tag{S167}$$

Assume that

$$\Lambda_3\delta + \frac{\Lambda_{\mathcal{R}}}{\delta} < \Lambda_{\mathrm{stab}}. \tag{S168}$$

Then every solution with

$$\mathfrak{E}_0(0)^2 + \mu_k \mathfrak{H}_k(0)^2 < \delta^2 \tag{S169}$$

satisfies

$$\mathfrak{E}_0(t) < \delta_0, \qquad \mathfrak{H}_k(t) < \delta_k \tag{S170}$$

for as long as the solution exists and the certified estimates apply.

Proof The weighted-energy part of the argument is identical to the single-radius argument in Theorem S1. Thus, $\mathfrak{E}_k(t) < \delta$ for all times under consideration. It remains only to translate this invariant weighted ball into the two separate target bounds. By the containment condition (S162),

$$\mathfrak{E}_0(t) \leq \mathfrak{E}_k(t) < \delta \leq \delta_0, \qquad \mathfrak{H}_k(t) \leq \frac{\mathfrak{E}_k(t)}{\sqrt{\mu_k}} < \frac{\delta}{\sqrt{\mu_k}} \leq \delta_k. \tag{S171}$$

■

3.8 From Rescaled Stability to Physical Blowup

The nonlinear stability theorems provide control of an exact solution in dynamically rescaled variables, where the concentrating core is kept at order-one scale and amplitude. By the exact dynamic-rescaling change of variables, every finite rescaled time corresponds to an equivalent physical solution.

To obtain finite-time blowup from a trajectory that exists for all rescaled times, two facts remain to be checked: the physical clock must have a finite limit, and undoing the normalization must force a physical quantity to become unbounded. This subsection provides exactly this reconstruction.

Both checks are controlled by the amplitude scale $\mathrm{s}_u(t)$. The normalization fixes $u(0,0,t) = \bar{u}_0 \neq 0$, where $\bar{u}_0$ is a constant, while

$$\frac{d\mathring{t}}{dt} = \frac{1}{\mathrm{s}_u(t)}. \tag{S172}$$

Thus growth of s_u makes the remaining physical time small and, after undoing the amplitude normalization, forces growth of the physical axial vorticity at the moving center. A uniform negative upper bound on $c_u(t)$ is enough to produce this growth.

As defined in Section 1, $\mathring{t}$ denotes physical time, $\mathring{z}_c(t)$ the physical axial center, $\tilde{\omega}$ the full physical vorticity, and ω^z its axial component.

Proposition S4 (Finite-time physical blowup from a signed amplitude rate)**.** *Consider the Euler equations in the dynamic rescaling formulation above, with $\lambda = 1/2$. Assume that the exact admissible rescaled solution exists for every $t \geq 0$, remains in the corresponding stability ball, and satisfies the modulation bounds throughout. Assume in addition that*

$$c_u(t) \leq -\gamma_u < 0, \qquad t \geq 0, \tag{S173}$$

for some $\gamma_u > 0$. Then the physical time $\mathring{t}(t)$ converges to a finite limit T, with

$$0 < T - \mathring{t}(t) \leq \frac{e^{-\gamma_u t}}{\gamma_u\, \mathrm{s}_u(0)}. \tag{S174}$$

The traveling center $\mathring{z}_c(t)$ converges to a finite physical location. Moreover, at the moving center the physical axial vorticity satisfies

$$\left|\omega^z\big(0, \mathring{z}_c(t), \mathring{t}(t)\big)\right| = 2|\bar{u}_0|\, \mathrm{s}_u(t) \longrightarrow \infty. \tag{S175}$$

Consequently, since ω^z is a component of the full physical vorticity $\tilde{\omega}$,

$$\|\tilde{\omega}(\cdot, \mathring{t}(t))\|_{L^\infty(\mathbb{R}^3)} \longrightarrow \infty, \tag{S176}$$

and the reconstructed physical solution cannot be continued smoothly through T.

Proof The proof is given in Appendix J. ■

Certification of the sign condition. Since $\bar{c}_u = -1$ and $c_u = \bar{c}_u + c_{u,\mathrm{res}} + \delta c_u$, the modulation bounds reduce (S173) to scalar checks. It is enough to verify

$$1 - c_{u,\mathrm{res}} - M^{(1)}_{c_u}\delta_* > 0, \tag{S177}$$

where $M^{(1)}_{c_u}$ are constants that measure the sensitivity of the amplitude modulation parameter c_u to the perturbation. Its precise definition is given in the companion stability paper [18].

Remark. *This assumes that an exact admissible rescaled solution exists for every $t \geq 0$ and remains in the stability ball. The stability estimates provide control of such a solution only for as long as it exists. It does not by itself prove all-time existence in rescaled time. To obtain an unconditional blowup theorem, the stability estimate must be combined with a local well-posedness and continuation result showing that the controlled rescaled solution cannot break down at any finite rescaled time. The residual of the numerical profile does not cause any issue because its certified residual is already included in the perturbation equations.*

3.9 Computer-Assisted Strategy for Proving Stability

Theorems S1 and S2 reduce the nonlinear stability problem to a finite-dimensional optimization problem. The aim is to choose the weights so that the linear damping margin is larger than the optimized nonlinear and residual error. For each candidate choice of weights, we compute the certified constants, optimize the radius, and then update the weights until the closing inequality becomes strict. Once the weights are fixed, the remaining optimization is only over the scalar radius δ. Decreasing δ reduces the cubic contributions, while increasing δ reduces the residual contribution $\Lambda_{\mathcal{R}}/\delta$. Thus the best radius balances these two effects, and the outer search changes the weights to improve this balance.

The optimization loop proceeds as follows.

Step 1. Choose the class of admissible weight functions.

Start with a finite-dimensional parametrization of a class of admissible weight functions, chosen so that the weights have the required positivity, regularity, and symmetries.

For each parameter value ϑ, write the low-order weights as

$$\Phi_\vartheta = \left(\Phi_{\omega,\vartheta}, \Phi_{r,\vartheta}, \Phi_{z,\vartheta}\right), \tag{S178}$$

and define the order-k weights by

$$\Phi_{i,k,\vartheta} = 1 + \rho^k \Phi_{i,\vartheta}, \qquad \rho(r,z) = r^2 + z^2, \qquad i \in \{\omega, r, z\}. \tag{S179}$$

The optimization variables are the weight parameter ϑ and the relative top-order weight μ_k. In the single-radius case of Theorem S1, one imposes $0 < \mu_k \leq 1$. In the two-radii case of Theorem S2, one only imposes $\mu_k > 0$. In both cases the weighted energy is $\mathfrak{E}_k^2 = \mathfrak{E}_0^2 + \mu_k \mathfrak{H}_k^2$.

Choose an initial admissible candidate (ϑ, μ_k).

Step 2. Compute the linear damping for the current candidate.

For the current values of (ϑ, μ_k), compute

$$\Lambda_{\mathcal{L}}^{\text{low}}, \qquad \Lambda_{\mathcal{L}}^{\text{high}}, \qquad \Lambda_{D^\alpha \mathcal{L}}^{\text{low}}, \qquad \Lambda_{D^\alpha \mathcal{L}}^{\text{high}}. \tag{S180}$$

The damping margin available for the combined energy is

$$\Lambda_{\text{stab}} := \min\left\{ \Lambda_{\mathcal{L}}^{\text{low}} - \mu_k \Lambda_{D^\alpha \mathcal{L}}^{\text{low}},\ \Lambda_{D^\alpha \mathcal{L}}^{\text{high}} - \frac{\Lambda_{\mathcal{L}}^{\text{high}}}{\mu_k} \right\}. \tag{S181}$$

If $\Lambda_{\text{stab}} \leq 0$, the candidate cannot close the theorem, so we penalize this term and move to a new candidate.

Step 3. Compute the nonlinear and residual constants.

If $\Lambda_{\mathrm{stab}} > 0$, compute

$$\Lambda_3 = \Lambda_{\mathcal{N},3} + \Lambda_{D^\alpha \mathcal{N},3}. \tag{S182}$$

For this candidate, the error that must be dominated by the damping is

$$\Lambda_3 \delta + \frac{\Lambda_{\mathcal{R}}}{\delta}. \tag{S183}$$

Step 4. Optimize the radius and evaluate the gap.

For the current candidate, let $\mathcal{A}_\delta(\vartheta, \mu_k) \subset (0, \infty)$ denote the nonempty set of radii for which every radius-dependent hypothesis used in the proof is valid.

In the single-radius case, the admissible condition is

$$\delta \in \mathcal{A}_\delta(\vartheta, \mu_k). \tag{S184}$$

In the two-radii case with prescribed target bounds $\delta_0 > 0$ and $\delta_k > 0$, the admissible conditions are

$$\delta \in \mathcal{A}_\delta(\vartheta, \mu_k), \qquad 0 < \delta \le \min\{\delta_0, \sqrt{\mu_k}\, \delta_k\}. \tag{S185}$$

Define the optimized gap

$$\mathcal{J}(\vartheta, \mu_k) := \inf_{\delta \text{ admissible}} \left(\Lambda_3 \delta + \frac{\Lambda_{\mathcal{R}}}{\delta} \right) - \Lambda_{\mathrm{stab}}. \tag{S186}$$

The candidate succeeds exactly when $\mathcal{J}(\vartheta, \mu_k) < 0$. In this case, choose an admissible radius δ_* such that

$$\Lambda_3 \delta_* + \frac{\Lambda_{\mathcal{R}}}{\delta_*} < \Lambda_{\mathrm{stab}}. \tag{S187}$$

Step 5. Iterate the outer optimization.

The outer objective is the gap $\mathcal{J}(\vartheta, \mu_k)$.

One evaluation of this objective consists of Steps 2–4: compute the damping margin, compute the error constants, optimize δ, and subtract the margin.

If $\mathcal{J}(\vartheta, \mu_k) < 0$, the proof closes.

If $\mathcal{J}(\vartheta, \mu_k) \ge 0$, update (ϑ, μ_k) to decrease $\mathcal{J}$, and repeat Steps 2–4.

Step 6. Record the tentative certificate.

Once $\mathcal{J} < 0$, record the weights and a selected admissible radius δ_* satisfying (S187). The single-radius certificate gives

$$\mathfrak{E}_k(t) < \delta_*. \tag{S188}$$

Provided that $0 < \delta_* \le \min\{\delta_0, \sqrt{\mu_k}\delta_k\}$, the two-radii certificate gives

$$\mathfrak{E}_0(t) < \delta_0, \qquad \mathfrak{H}_k(t) < \delta_k. \tag{S189}$$

Step 7. Validate rigorously the tentative certificate.

Validate rigorously and analytically that the proof closes with the candidate weights. The final certificate also verifies every radius-dependent admissibility guard at the selected radius δ_*.

We summarize the numerical strategy used to prove the stability theorems below

---

Algorithm 1. Loop for optimizing the nonlinear stability certificate

1. Initialize an admissible candidate (ϑ, μ_k).

2. Repeat the following loop.

a. Compute

$$\Lambda_{\mathrm{stab}} = \min\left\{ \Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}},\ \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k} \right\}. \tag{S190}$$

If $\Lambda_{\mathrm{stab}} \le 0$, penalize the candidate and update (ϑ, μ_k).

b. If $\Lambda_{\mathrm{stab}} > 0$, compute $\Lambda_{\mathcal{R}}$ and $\Lambda_3 = \Lambda_{\mathcal{N},3} + \Lambda_{D^\alpha\mathcal{N},3}$.

c. Determine the nonempty admissible set of radii for the current candidate, and compute

$$\mathcal{J}(\vartheta, \mu_k) = \inf_{\delta \text{ admissible}} \left(\Lambda_3 \delta + \frac{\Lambda_{\mathcal{R}}}{\delta} \right) - \Lambda_{\mathrm{stab}}. \tag{S191}$$

d. If $\mathcal{J} < 0$, select an admissible radius δ_* satisfying (S187) and stop, the argument closes and stability is proven.

e. If $\mathcal{J} \ge 0$, update (ϑ, μ_k) to decrease $\mathcal{J}$, and repeat the loop.

3. Record the final weights, the selected admissible radius δ_*, and the resulting bounds.

4. Validate rigorously the tentative proof using analytic certificates, including every radius-dependent admissibility guard at δ_*.

---

4 Partial Results: Linear Damping

4.1 Objective and Scope

We now examine the signed linear damping mechanism that drives the stability argument. The objective is to obtain a coercive linear contribution to the full weighted energy after the low-order and high-order estimates are combined and the remaining linear terms are included.

The complete linear estimates we aim to prove have the two-level form

$$\sum_{i\in\{\omega,r,z\}} \langle \mathcal{L}_i, \delta v_i\rangle_{\Phi_i} \le -\Lambda_{\mathcal{L}}^{\mathrm{low}} \mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}} \mathfrak{H}_k^2, \tag{S192}$$

$$\sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \le \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}} \mathfrak{E}_0^2 - \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} \mathfrak{H}_k^2. \tag{S193}$$

The four coefficients have complementary roles. The constant $\Lambda_{\mathcal{L}}^{\mathrm{low}}$ is the damping margin supplied by the low-order estimate, while $\Lambda_{\mathcal{L}}^{\mathrm{high}}$ measures the accompanying loss at top order. Conversely, $\Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}}$ is the damping margin supplied by the high-order estimate, while $\Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}}$ measures its accompanying low-order loss. The purpose of the remainder of this section is to construct and certify these four constants.

The reason the two estimates must be considered together is already visible from their form. Multiplying equation (S193) by μ_k and adding it to equation (S192) gives

$$\begin{aligned}\sum_{i\in\{\omega,r,z\}} \langle \mathcal{L}_i, \delta v_i\rangle_{\Phi_i} &+ \mu_k \sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \\ &\le -\left(\Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}}\right) \mathfrak{E}_0^2 - \mu_k \left(\Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k}\right) \mathfrak{H}_k^2.\end{aligned} \tag{S194}$$

Thus neither estimate needs to be coercive in both energy levels by itself. A positive $\mathfrak{H}_k^2$ loss in the low-order estimate can be absorbed by the negative high-order contribution, while a positive $\mathfrak{E}_0^2$ loss in the high-order estimate can be absorbed by the low-order damping. This is the basic role of the high-order estimate in the argument.

It is therefore natural to define the net linear damping margin by

$$\Lambda_{\mathrm{stab}}(\mu_k) = \min\left\{ \Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}},\ \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k} \right\} > 0. \tag{S195}$$

The two quantities inside the minimum are precisely the damping margins for $\mathfrak{E}_0^2$ and $\mu_k \mathfrak{H}_k^2$ in (S194). A positive value of Λ_{stab} gives a genuinely coercive linear estimate before the nonlinear and residual contributions are added. Unlike those later contributions, a missing linear sign cannot in general be repaired simply by restricting to a smaller perturbation neighborhood.

The rest of the section is organized as follows. First, we construct the low-order signed matrix and identify the pointwise criterion that yields low-order damping. Second, we choose admissible weights and certify the low-order matrix on the finite computational box and in the analytic far field. Third, once this direct low-order bound is established, we analyze any exceptional region where the same pointwise margin is not available. Fourth, we construct the complete high-order signed matrix and combine its damping with that low-order estimate. Finally, we incorporate the linear terms kept outside the two matrices, together with the nonlinear and residual contributions, to obtain the closing stability condition.

4.2 Low-Order Signed Damping Matrix

4.2.1 Transport Field and Retained Linear Terms

At low order, the question is whether the transport terms and a selected set of local linear couplings form a negative quadratic form at almost every point. The selection is made so that every retained term acts pointwise on the three variables already present in the low-order weighted energy and has a controlled coefficient after the corresponding weight conversion. Terms with a different structure are kept outside the pointwise matrix and estimated separately by specialized mechanisms adapted to them.

Recall that the transport field is

$$G_r = \lambda r + \bar{u}^r, \qquad G_z = \lambda z + \bar{C} + \bar{u}^z, \tag{S196}$$

where

$$\bar{u}^r = -r\partial_z\bar{\psi}, \qquad \bar{u}^z = 2\bar{\psi} + r\partial_r\bar{\psi}. \tag{S197}$$

Besides transport, the signed matrix retains the local diagonal coefficients multiplying $\delta\omega$, $\partial_r\delta u$, and $\partial_z\delta u$, together with the two local r–z couplings between $\partial_r\delta u$ and $\partial_z\delta u$. These are exactly the local terms that can be represented in the same three-component pointwise quadratic form.

Note that the coupling term $2\bar{u}\,\partial_z\delta u$ is kept outside the low-order signed matrix, because after passing to the weighted variables its coefficient would contain $2\bar{u}\sqrt{\Phi_\omega/\Phi_z}$, and the independent low-order singularity rates of Φ_ω and Φ_z do not by themselves guarantee that this multiplier is bounded. The complete low-order term list and the precise partition between matrix and non-matrix terms are recorded in Appendix H.1.1.

4.2.2 Weighted Transport Contribution

The role of the weights is most transparent in the transport contribution. Assuming the boundary terms vanish under the boundary and decay conditions used in the weighted energy estimate, integration by parts gives

$$\begin{aligned}
\langle -G_r\partial_r f - G_z\partial_z f, f\rangle_{\Phi_i} &= -\frac{1}{2}\int_{\mathbb{D}} \left(G_r\partial_r(f^2) + G_z\partial_z(f^2)\right)\Phi_i\,dr\,dz \\
&= \frac{1}{2}\int_{\mathbb{D}} |f|^2\left[\partial_r(G_r\Phi_i) + \partial_z(G_z\Phi_i)\right]dr\,dz \\
&= \frac{1}{2}\int_{\mathbb{D}} |f|^2\Phi_i\left(\partial_r G_r + \partial_z G_z + G_r\partial_r\log\Phi_i + G_z\partial_z\log\Phi_i\right)dr\,dz.
\end{aligned} \tag{S198}$$

Thus the coefficient multiplying $|f|^2\Phi_i$ is

$$\frac{1}{2}\left(\partial_r G_r + \partial_z G_z\right) + \frac{1}{2}\left(G_r\partial_r\log\Phi_i + G_z\partial_z\log\Phi_i\right). \tag{S199}$$

The first term is fixed by the divergence of the transport field. The second depends on the variation of the weight along the transport direction and is thus the part that can be adjusted through the choice of Φ_i. This is the basic mechanism by which the weights can improve the signed transport contribution.

4.2.3 Explicit Low-Order Matrix

We introduce the weighted low-order state

$$d\xi(r,z) = \left(\Phi_\omega^{1/2}\delta\omega, \Phi_r^{1/2}\partial_r\delta u, \Phi_z^{1/2}\partial_z\delta u\right)^\top, \qquad \int_{\mathbb{D}} |d\xi|^2\, dr\, dz = \mathfrak{E}_0^2. \tag{S200}$$

With the component order (ω, r, z), the retained low-order contribution is represented by a symmetric 3×3 matrix $\mathbb{M}_{\mathcal{L}}(r,z)$. We now define its entries from the transport contribution and the retained local coefficients.

Transport coefficients. The derivatives of the transport field that enter the matrix are

$$\partial_r\bar{u}^r = -\partial_z\bar{\psi} - r\partial_{rz}\bar{\psi}, \qquad \partial_z\bar{u}^r = -r\partial_{zz}\bar{\psi}, \tag{S201}$$

$$\partial_r\bar{u}^z = 3\partial_r\bar{\psi} + r\partial_{rr}\bar{\psi}, \qquad \partial_z\bar{u}^z = 2\partial_z\bar{\psi} + r\partial_{rz}\bar{\psi}, \tag{S202}$$

so that

$$\partial_r G_r = \lambda + \partial_r\bar{u}^r, \qquad \partial_z G_z = \lambda + \partial_z\bar{u}^z. \tag{S203}$$

For each $i \in \{\omega, r, z\}$, we denote the transport contribution to the corresponding diagonal entry by

$$D_i := \frac{1}{2}\left(\partial_r G_r + \partial_z G_z + G_r\partial_r \log\Phi_i + G_z\partial_z\log\Phi_i\right). \tag{S204}$$

Diagonal entries. The constant modulation offset $c_{u,\mathrm{res}}$ is absorbed into the reference amplitude coefficient and therefore contributes to the zeroth-order coefficients added to the transport contribution, i.e.,

$$(\mathbb{M}_{\mathcal{L}})_{\omega\omega} = D_\omega + \bar{c}_u + c_{u,\mathrm{res}} - \lambda, \tag{S205}$$

$$(\mathbb{M}_{\mathcal{L}})_{rr} = D_r + 2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}} - \lambda - \partial_r\bar{u}^r, \tag{S206}$$

$$(\mathbb{M}_{\mathcal{L}})_{zz} = D_z + 2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}} - \lambda - \partial_z\bar{u}^z. \tag{S207}$$

Cross entry. The two retained r–z couplings contribute to the same symmetric off-diagonal entry. After converting each term to the weighted variables and symmetrizing the quadratic form, we obtain

$$(\mathbb{M}_{\mathcal{L}})_{rz} = (\mathbb{M}_{\mathcal{L}})_{zr} = -\frac{1}{2}\left[\partial_r\bar{u}^z\left(\frac{\Phi_r}{\Phi_z}\right)^{1/2} + \partial_z\bar{u}^r\left(\frac{\Phi_z}{\Phi_r}\right)^{1/2}\right]. \tag{S208}$$

There are no retained ω–r or ω–z entries at low order. In particular, the ω–z coupling term $2\bar{u}\,\partial_z\delta u$ discussed earlier remains outside the matrix because its weight conversion is not controlled by the low-order weight assumptions.

Altogether, these entries form the symmetric matrix $\mathbb{M}_{\mathcal{L}}$. By construction, the retained transport and selected couplings contribute to the weighted energy through

$$\int_{\mathbb{D}} d\xi^\top \mathbb{M}_{\mathcal{L}}(r,z) d\xi\, dr\, dz. \tag{S209}$$

The pointwise sign of this quadratic form is determined by the largest eigenvalue of $\mathbb{M}_{\mathcal{L}}$:

$$\lambda_{\max}(\mathbb{M}_{\mathcal{L}}(r,z)) < 0 \qquad \Longleftrightarrow \qquad \eta^\top\mathbb{M}_{\mathcal{L}}(r,z)\eta < 0 \quad \text{for every } \eta \neq 0. \tag{S210}$$

Thus $-\lambda_{\max}(\mathbb{M}_{\mathcal{L}})$ is the local raw damping margin whenever the largest eigenvalue is negative.

4.3 Low-order Damping Certification

4.3.1 Weight Optimization and Finite-Box Evaluation

Weight optimization. In the transport diagonal, the weight enters only through

$$G_r \partial_r \log \Phi_i + G_z \partial_z \log \Phi_i, \qquad i \in \{\omega, r, z\}. \tag{S211}$$

The free parameters in the admissible low-order weights are therefore chosen to improve the pointwise sign of the complete matrix while preserving the admissibility conditions imposed on the low-order weights. Because the r–z cross entry also depends on weight ratios, it is not sufficient to optimize the transport diagonal component by component. The numerical objective is instead based on the largest eigenvalue of the full symmetric matrix $\mathbb{M}_{\mathcal{L}}$, with the aim of making its worst value as negative as possible and reducing the size of any region where strict negativity fails.

For the optimized candidate, we evaluate $\lambda_{\max}(\mathbb{M}_{\mathcal{L}}(r,z))$ throughout the finite half-plane box on which the reference profile is represented by splines. The rigorous finite-box step is then to enclose the largest eigenvalue with interval arithmetic, requiring a uniform negative upper bound on the region treated directly and only a finite upper bound on any localized pieces introduced below. Numerically, the largest eigenvalue is negative on almost all of the box, and the weakest margin is confined to small neighborhoods where the meridional transport field is small. These locations are analyzed in Subsection 4.4. The numerical evaluation identifies two low-transport neighborhoods: one around the normalized origin and one around a single off-axis fixed point. Their treatment is determined by the sign of the full low-order matrix, not by the smallness of the transport alone. Around the origin the full matrix remains in the directly damped regime, whereas the off-axis neighborhood is the single region on which the uniform negative matrix margin is not imposed and is therefore treated by localization.

The optimization step is used to identify a promising admissible set of weights. It is not itself the rigorous damping certificate. Once a candidate set of weights is fixed, the largest eigenvalue of the resulting matrix is evaluated on the computational box and then enclosed with interval arithmetic. The exterior region is treated separately by the analytic tail argument discussed in the subsequent section.

4.3.2 Analytic Far-Field Closure

Define the numerical enclosure box by

$$\mathbb{D}_{\text{box}} := [0, R_{\text{box}}] \times [-R_{\text{box}}, R_{\text{box}}]. \tag{S212}$$

On $\mathbb{D}_{\text{box}}$, the finite-box certification step consists of interval upper bounds for the largest eigenvalue of the pointwise matrix as described in the previous subsection. To extend these bounds to the full half-plane, it remains to control the exterior region $\mathbb{D} \setminus \mathbb{D}_{\text{box}}$ analytically.

Throughout the tail calculation,

$$\lambda = \frac{1}{2}, \qquad \bar{c}_u = -1. \tag{S213}$$

The profile-dependent constant $\bar{C}$ remains symbolic until the numerical certificate is evaluated.

In our spline representation, the steady profile and every profile derivative entering $\mathbb{M}_{\mathcal{L}}$ vanish exactly on $\mathbb{D} \setminus \mathbb{D}_{\text{box}}$. Then

$$\bar{u} = \bar{\omega} = \bar{\psi} = 0, \qquad \bar{u}^r = \bar{u}^z = 0, \qquad G_r = \frac{r}{2}, \qquad G_z = \frac{z}{2} + \bar{C}. \tag{S214}$$

All retained cross couplings vanish in this region. As a result, the tail damping matrix is diagonal, and its i-th diagonal entry is

$$-1 + c_{u,\mathrm{res}} + \frac{1}{4}(r\partial_r + z\partial_z)\log\Phi_i + \frac{1}{2}\bar{C}\,\partial_z \log\Phi_i, \qquad i \in \{\omega, r, z\}. \tag{S215}$$

Scalar tail condition. It is sufficient to find $\gamma_{\mathrm{tail}} > 0$ such that

$$\sup_{\mathbb{D}\setminus\mathbb{D}_{\mathrm{box}}} \left[(r\partial_r + z\partial_z)\log\Phi_i + 2\bar{C}\,\partial_z\log\Phi_i\right] \le 4(1 - c_{u,\mathrm{res}} - \gamma_{\mathrm{tail}}) \tag{S216}$$

for every $i \in \{\omega, r, z\}$. Substituting this bound into (S215) yields

$$\lambda_{\max}\big(\mathbb{M}_{\mathcal{L}}(r,z)\big) \le -\gamma_{\mathrm{tail}} \qquad \text{on } \mathbb{D}\setminus\mathbb{D}_{\mathrm{box}}. \tag{S217}$$

The weights can be continued analytically in the tail. For the polynomial-growth continuation,

$$\Phi_i(r,z) = C_{\mathrm{floor}} + a_i\rho, \qquad \rho = r^2 + z^2, \qquad a_i > 0, \tag{S218}$$

with $p = 1$, the calculation in Appendix H.1.2 gives the sufficient condition

$$\left(1 + \frac{2|\bar{C}|}{R_{\mathrm{box}}}\right) < 2(1 - c_{u,\mathrm{res}}). \tag{S219}$$

Whenever this inequality holds, the tail matrix has a uniform strictly negative largest eigenvalue. Once the finite-box interval enclosure is certified, combining it with this tail bound gives a pointwise low-order matrix estimate on the entire half-plane except for any problematic small region of low transport on which the uniform negative margin is not obtained.

4.4 Meridional Fixed Points and Localized Low-Order Control

4.4.1 Fixed Points and Low-Transport Regions

We found earlier that the low-order damping is weakest near locations where the meridional transport field becomes small. We refer to $(r_j^{\odot}, z_j^{\odot})$ as a *meridional fixed point* if the meridional transport field vanishes there, i.e.,

$$G_r(r_j^{\odot}, z_j^{\odot}) = 0, \qquad G_z(r_j^{\odot}, z_j^{\odot}) = 0. \tag{S220}$$

We also refer to a neighborhood in which

$$\sqrt{|G_r(r,z)|^2 + |G_z(r,z)|^2} \ll 1 \tag{S221}$$

as a *low-transport region*.

The weight-dependent part of the low-order transport coefficient is

$$G_r\partial_r\log\Phi_i + G_z\partial_z\log\Phi_i. \tag{S222}$$

At a meridional fixed point this contribution vanishes, and it becomes small in a low-transport neighborhood when the logarithmic derivatives of the admissible weights remain controlled. This weakens the main tunable contributions to the low-order diagonal entries, although it does not by itself determine the sign of the full damping matrix, because the divergence term and the retained local coefficients remain present.

Numerically, the meridional transport field has two fixed points in the computational domain. One lies on the symmetry axis and is shifted to the origin by the normalization condition (S80), while the second is off-axis. These two points are treated differently because $G_r = G_z = 0$ does not determine the sign of the full low-order matrix. For the optimized candidate, the numerical evaluation indicates that the remaining matrix terms are sufficient to keep a negative largest eigenvalue around the normalized origin, so the origin remains in the region treated by the direct finite-box bound. Around the off-axis fixed point, the same uniform negative margin is not imposed on a small neighborhood whose contribution is controlled instead. Finally, the first derivatives of G_r and G_z may remain nonzero at both fixed points, so the high-order transport commutators can provide a nontrivial damping mechanism even where the low-order weight-dependent transport contribution vanishes.

4.4.2 Localized Matrix Estimate

To quantify the off-axis low-transport region without imposing the same pointwise margin there as on the rest of the domain, we use a localized estimate. A uniform negative upper bound for $\lambda_{\max}(\mathbb{M}_{\mathcal{L}})$ on the entire half-plane would be sufficient for the integrated energy estimate, but it is stronger than necessary. A failure of the pointwise margin at an isolated point has no effect on the integral, because a set of measure zero carries no L^2 mass. What must be controlled is a positive-measure neighborhood on which the uniform negative margin is unavailable. The localized estimate permits such neighborhoods provided that the weighted energy that can concentrate in them is quantitatively controlled.

Small measure of the neighborhood alone is not sufficient, since an L^2 perturbation can concentrate most of its mass in a small region. We thus combine two inputs: a uniform negative matrix bound on the complement of the localized regions and a concentration estimate that controls the weighted energy inside each localized piece by the global low-order and top-order energies. The low-order part of this concentration estimate reduces the effective coefficient of $\mathfrak{E}_0^2$, while the top-order part produces an additional $\mathfrak{H}_k^2$ term that must be absorbed by the negative high-order matrix contribution after the high-order estimate is weighted by μ_k.

Localized regions. Let the region on which the uniform low-order matrix margin is not obtained be a finite union of pairwise disjoint measurable pieces

$$\mathbb{D}_{\mathcal{L}}^{\mathrm{loc}} = \bigcup_{\ell \in \mathcal{I}_{\mathcal{L}}^{\mathrm{loc}}} B_\ell, \qquad \mathbb{D} = \mathbb{D}_{\mathcal{L}}^{\mathrm{damp}} \cup \mathbb{D}_{\mathcal{L}}^{\mathrm{loc}}, \qquad \mathbb{D}_{\mathcal{L}}^{\mathrm{damp}} \cap \mathbb{D}_{\mathcal{L}}^{\mathrm{loc}} = \emptyset. \tag{S223}$$

We keep the notation above in a general finite-union form because the same argument would apply if several disjoint low-transport regions had to be localized. In the present computation, however, $\mathcal{I}_{\mathcal{L}}^{\mathrm{loc}} = \{\ell_{\mathrm{off}}\}$ and $B_{\ell_{\mathrm{off}}}$ is the single neighborhood of the unique off-axis fixed point. In particular, $B_{\ell_{\mathrm{off}}}$ stays a positive distance from the origin, which is useful for the concentration estimate because the low-order weights are singular only at the origin and are therefore bounded on $B_{\ell_{\mathrm{off}}}$.

Inputs for the localized certificate. Once the finite-box enclosure is validated, the direct-region constant $\Lambda_{\mathcal{L}}^{\mathrm{damp}}$ can be chosen as the minimum of the certified finite-box margin on the directly treated portion of $\mathbb{D}_{\mathrm{box}}$ and the analytic tail margin γ_{tail}. On $\mathbb{D}_{\mathcal{L}}^{\mathrm{damp}}$, we then require

$$\lambda_{\max}\big(\mathbb{M}_{\mathcal{L}}(r,z)\big) \le -\Lambda_{\mathcal{L}}^{\mathrm{damp}} \qquad (r,z) \in \mathbb{D}_{\mathcal{L}}^{\mathrm{damp}}, \qquad \Lambda_{\mathcal{L}}^{\mathrm{damp}} > 0. \tag{S224}$$

On each localized piece, we require the finite upper bound

$$\lambda_{\max}\big(\mathbb{M}_{\mathcal{L}}(r,z)\big) \le \Lambda_{\mathcal{L}}^{\mathrm{loc},\ell} \qquad (r,z) \in B_\ell, \qquad \Lambda_{\mathcal{L}}^{\mathrm{loc},\ell} \ge 0. \tag{S225}$$

Thus the largest eigenvalue is uniformly negative on $\mathbb{D}_{\mathcal{L}}^{\mathrm{damp}}$, while on each B_ℓ only a finite upper bound is required.

The second input is a concentration estimate of the form

$$\int_{B_\ell} |d\xi|^2 \, dr \, dz \le m_{\mathcal{L}}^{0,\mathrm{loc},\ell} \mathfrak{E}_0^2 + m_{\mathcal{L}}^{k,\mathrm{loc},\ell} \mathfrak{H}_k^2, \qquad \ell \in \mathcal{I}_{\mathcal{L}}^{\mathrm{loc}}, \tag{S226}$$

with nonnegative constants $m_{\mathcal{L}}^{0,\mathrm{loc},\ell}$ and $m_{\mathcal{L}}^{k,\mathrm{loc},\ell}$. The first coefficient measures the part of the localized weighted mass controlled at low order. The second measures the part that must be controlled through the top-order energy.

Effective matrix estimate. Define

$$\gamma_\ell := \Lambda_{\mathcal{L}}^{\mathrm{damp}} + \Lambda_{\mathcal{L}}^{\mathrm{loc},\ell}. \tag{S227}$$

Using the domain split and $\int_{\mathbb{D}} |d\xi|^2 = \mathfrak{E}_0^2$,

$$\int_{\mathbb{D}} d\xi^\top \mathbb{M}_{\mathcal{L}} d\xi \, dr \, dz \le -\Lambda_{\mathcal{L}}^{\mathrm{damp}} \mathfrak{E}_0^2 + \sum_{\ell \in \mathcal{I}_{\mathcal{L}}^{\mathrm{loc}}} \gamma_\ell \int_{B_\ell} |d\xi|^2 \, dr \, dz. \tag{S228}$$

Define

$$\Lambda_{\mathcal{L}}^{\mathrm{mat,loc}} := \Lambda_{\mathcal{L}}^{\mathrm{damp}} - \sum_{\ell \in \mathcal{I}_{\mathcal{L}}^{\mathrm{loc}}} \gamma_\ell m_{\mathcal{L}}^{0,\mathrm{loc},\ell}, \tag{S229}$$

$$A_{\mathcal{L}}^{k,\mathrm{loc}} := \sum_{\ell \in \mathcal{I}_{\mathcal{L}}^{\mathrm{loc}}} \gamma_\ell m_{\mathcal{L}}^{k,\mathrm{loc},\ell}. \tag{S230}$$

Then the concentration estimate gives

$$\int_{\mathbb{D}} d\xi^\top \mathbb{M}_{\mathcal{L}} d\xi \, dr \, dz \le -\Lambda_{\mathcal{L}}^{\mathrm{mat,loc}} \mathfrak{E}_0^2 + A_{\mathcal{L}}^{k,\mathrm{loc}} \mathfrak{H}_k^2. \tag{S231}$$

Thus localization reduces the effective low-order matrix margin from $\Lambda_{\mathcal{L}}^{\mathrm{damp}}$ to $\Lambda_{\mathcal{L}}^{\mathrm{mat,loc}}$ and introduces the top-order term $A_{\mathcal{L}}^{k,\mathrm{loc}} \mathfrak{H}_k^2$. To retain a genuine low-order damping term, the localized certificate must in particular satisfy $\Lambda_{\mathcal{L}}^{\mathrm{mat,loc}} > 0$. The next subsection constructs the high-order matrix and provides the negative $\mathfrak{H}_k^2$ contribution that is compared directly with $A_{\mathcal{L}}^{k,\mathrm{loc}} \mathfrak{H}_k^2$.

4.5 High-Order Signed Damping Matrix

4.5.1 High-Order Transport Commutators and Meridional Fixed Points

The first transport commutators depend on the full Jacobian of the meridional transport field,

$$\partial_r G_r = \lambda + \partial_r \bar{u}^r, \qquad \partial_z G_r = \partial_z \bar{u}^r, \qquad \partial_r G_z = \partial_r \bar{u}^z, \qquad \partial_z G_z = \lambda + \partial_z \bar{u}^z. \tag{S232}$$

We consider $\alpha = (q, k-q)$ with $0 \le q \le k$ and apply D^α to $G_r \partial_r f + G_z \partial_z f$. The top-order commutator terms are exactly those in which one derivative lands on G_r or G_z. Each such term still contains k derivatives of f, so it belongs to the top-order signed matrix rather than to a lower-order remainder. Terms in which at least two derivatives land on the transport coefficients contain at most $k-1$ derivatives of the perturbation.

Top-order derivative variables. For $0 \le q \le k$, set

$$V_{i,q} := \partial_r^q \partial_z^{k-q} \delta v_i, \qquad i \in \{\omega, r, z\}. \tag{S233}$$

Here q is the number of r-derivatives and $k-q$ is the number of z-derivatives. The neighboring terms in the formula below occur only when the corresponding derivative is available: the $V_{i,q-1}$ term carries a factor q and therefore disappears when $q = 0$, while the $V_{i,q+1}$ term carries a factor $k-q$ and disappears when $q = k$.

For each (i, q), the complete top-order contribution coming from high-order transport and its first commutators is

$$-G_r \partial_r V_{i,q} - G_z \partial_z V_{i,q} - \big(q\, \partial_r G_r + (k-q)\, \partial_z G_z\big) V_{i,q} - q\, \partial_r G_z\, V_{i,q-1} - (k-q)\, \partial_z G_r\, V_{i,q+1}. \tag{S234}$$

The first two terms transport the top derivative itself. The coefficient of $V_{i,q}$ contains the two diagonal commutator contributions $q\partial_r G_r$ and $(k-q)\partial_z G_z$. The index shifts can be read directly from the Leibniz rule. In the term $G_z \partial_z f$, if one of the q r-derivatives falls on G_z, then only $q-1$ r-derivatives remain on f, while the original ∂_z supplies one additional z-derivative. This gives $V_{i,q-1}$ with coefficient $q\, \partial_r G_z$. Similarly, in $G_r \partial_r f$, if one of the $k-q$ z-derivatives falls on G_r, the perturbation factor gains one r-derivative and loses one z-derivative, giving $V_{i,q+1}$ with coefficient $(k-q)\, \partial_z G_r$.

These are all of the transport contributions that remain at top perturbation order.

Behavior at a meridional fixed point. This is precisely where the high-order estimate differs from the low-order one. At low order, the tunable weight contribution is proportional to G_r and G_z and therefore vanishes at a meridional fixed point. At high order, the first commutators contain derivatives of the transport field, such as $\partial_r G_r$, $\partial_z G_z$, $\partial_r G_z$, and $\partial_z G_r$. These quantities need not vanish when $G_r = G_z = 0$, so the high-order transport matrix can retain a nontrivial signed contribution at the fixed point.

After weighted integration by parts, the transport contribution to the diagonal (i, q) entry contains

$$D_{i,q}^{(k),\mathrm{tr}} = \frac{1}{2} \left(\partial_r G_r + \partial_z G_z + G_r \partial_r \log \Phi_{i,k} + G_z \partial_z \log \Phi_{i,k}\right) - q\, \partial_r G_r - (k-q)\, \partial_z G_z. \tag{S235}$$

At a meridional fixed point, the terms multiplied by G_r and G_z vanish, while the derivatives of the transport field remain, so

$$D_{i,q}^{(k),\mathrm{tr}}(r_j^{\odot}, z_j^{\odot}) = \left(\frac{1}{2} - q\right) \partial_r G_r(r_j^{\odot}, z_j^{\odot}) + \left(\frac{1}{2} - (k-q)\right) \partial_z G_z(r_j^{\odot}, z_j^{\odot}). \tag{S236}$$

The neighboring derivative couplings also retain the cross derivatives $\partial_z G_r$ and $\partial_r G_z$. Thus the high-order transport contribution can remain nontrivial even where the low-order weight-dependent transport contribution vanishes or becomes weak. This observation explains why high-order damping is available as a complementary mechanism near the fixed points, although its sign is determined only after the full high-order matrix is assembled.

4.5.2 Complete High-Order Signed Matrix

The high-order transport contribution described in the previous subsection determines the transport part of the top-order matrix. It remains to add the linear terms for which all k derivatives stay on the perturbation.

These terms are, component by component,

$$\omega\text{-term}: \quad (\bar{c}_u - \lambda + c_{u,\mathrm{res}})V_{\omega,q} + 2\bar{u}\, V_{z,q}, \tag{S237}$$

$$r\text{-term}: \quad \left(2\partial_z\bar{\psi} + \bar{c}_u - \lambda - \partial_r\bar{u}^r + c_{u,\mathrm{res}}\right) V_{r,q} - \partial_r\bar{u}^z\, V_{z,q}, \tag{S238}$$

$$z\text{-term}: \quad -\partial_z\bar{u}^r\, V_{r,q} + \left(2\partial_z\bar{\psi} + \bar{c}_u - \lambda - \partial_z\bar{u}^z + c_{u,\mathrm{res}}\right) V_{z,q}. \tag{S239}$$

These are exactly the local terms for which all k derivatives remain on the perturbation. If a derivative lands on one of their spatially varying coefficients, the resulting perturbation factor has order at most $k-1$ and is treated as a remainder.

As in the low-order construction, it is convenient to write the signed quadratic form directly as a symmetric matrix. Order the weighted top derivatives by component and derivative index:

$$\mathbf{V}_k := \left(\Phi_{\omega,k}^{1/2}V_{\omega,0}, \ldots, \Phi_{\omega,k}^{1/2}V_{\omega,k}, \Phi_{r,k}^{1/2}V_{r,0}, \ldots, \Phi_{r,k}^{1/2}V_{r,k}, \Phi_{z,k}^{1/2}V_{z,0}, \ldots, \Phi_{z,k}^{1/2}V_{z,k}\right)^\top. \tag{S240}$$

With this ordering, the matrix separates into diagonal (i,q) terms, neighboring q-couplings within each component, and same-q cross-component couplings.

The matrix $\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}$ is indexed by the pairs (i,q), with $i \in \{\omega, r, z\}$ and $0 \le q \le k$:

$$\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}} = \left[\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(i,q),(j,p)}\right]_{\substack{i,j\in\{\omega,r,z\}\\ 0\le q,p\le k}}, \qquad \mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}} = \left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)^\top. \tag{S241}$$

Diagonal entries. The diagonal entries are obtained exactly as in the low-order matrix: weighted integration by parts contributes the transport divergence and weight-gradient terms, the first commutators contribute the terms proportional to q and $k-q$, and equations (S237)–(S239) contribute the local diagonal coefficients. Explicitly, for $0 \le q \le k$,

$$\begin{aligned}\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(\omega,q),(\omega,q)} &= \frac{1}{2}\left(\partial_r G_r + \partial_z G_z + G_r\partial_r \log\Phi_{\omega,k} + G_z\partial_z\log\Phi_{\omega,k}\right)\\ &\quad + \bar{c}_u - \lambda + c_{u,\mathrm{res}} - q\,\partial_r G_r - (k-q)\,\partial_z G_z,\end{aligned} \tag{S242}$$

$$\begin{aligned}\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(r,q),(r,q)} &= \frac{1}{2}\left(\partial_r G_r + \partial_z G_z + G_r\partial_r \log\Phi_{r,k} + G_z\partial_z\log\Phi_{r,k}\right)\\ &\quad + 2\partial_z\bar{\psi} + \bar{c}_u - \lambda - \partial_r\bar{u}^r + c_{u,\mathrm{res}} - q\,\partial_r G_r - (k-q)\,\partial_z G_z,\end{aligned} \tag{S243}$$

$$\begin{aligned}\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(z,q),(z,q)} &= \frac{1}{2}\left(\partial_r G_r + \partial_z G_z + G_r\partial_r \log\Phi_{z,k} + G_z\partial_z\log\Phi_{z,k}\right)\\ &\quad + 2\partial_z\bar{\psi} + \bar{c}_u - \lambda - \partial_z\bar{u}^z + c_{u,\mathrm{res}} - q\,\partial_r G_r - (k-q)\,\partial_z G_z.\end{aligned} \tag{S244}$$

Same-component off-diagonal entries. For a fixed component i, equation (S234) couples $V_{i,q}$ only to the neighboring derivative splits $V_{i,q-1}$ and $V_{i,q+1}$. After symmetrization, each same-component block is therefore tridiagonal in the derivative index q: its only off-diagonal entries connect adjacent indices (i,q) and $(i,q+1)$. For $0 \le q < k$,

$$\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(i,q),(i,q+1)} = \left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(i,q+1),(i,q)} = -\frac{1}{2}\left[(k-q)\,\partial_z\bar{u}^r + (q+1)\,\partial_r\bar{u}^z\right]. \tag{S245}$$

Cross-component entries. The cross-component entries are obtained as in the low-order matrix by keeping each coefficient together with its required weight conversion and then symmetrizing. For $0 \le q \le k$,

$$\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(\omega,q),(z,q)} = \left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(z,q),(\omega,q)} = \bar{u}\left(\frac{\Phi_{\omega,k}}{\Phi_{z,k}}\right)^{1/2}, \tag{S246}$$

$$\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(r,q),(z,q)} = \left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\right)_{(z,q),(r,q)} = -\frac{1}{2}\left[\partial_r\bar{u}^z\left(\frac{\Phi_{r,k}}{\Phi_{z,k}}\right)^{1/2} + \partial_z\bar{u}^r\left(\frac{\Phi_{z,k}}{\Phi_{r,k}}\right)^{1/2}\right]. \tag{S247}$$

For the ω–z entry, the quantity to certify is the complete multiplier $\bar{u}\sqrt{\Phi_{\omega,k}/\Phi_{z,k}}$. Near the origin, $\Phi_{i,k} = 1 + \rho^k\Phi_i$ and $\rho^k\Phi_i \to 0$, so both high-order weights tend to 1 and this multiplier has no singular weight-conversion obstruction there. On the rest of the finite box and on the tail, the complete product is enclosed directly.

There are no ω-r entries. There are also no cross-component entries with different derivative indices $q \ne p$.

Together with symmetry, equations (S242)–(S247) therefore specify every entry of $\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}$.

High-order matrix certificate. The high-order sign is tested on the complete symmetric matrix, because the off-diagonal derivative and component couplings can affect the largest eigenvalue even when all diagonal entries are negative. The required certificate is

$$\operatorname*{ess\,sup}_{(r,z)\in\mathbb{D}} \lambda_{\max}\left(\mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}(r,z)\right) \le -\Lambda_{D^\alpha L}^{\mathrm{full}} < 0. \tag{S248}$$

The current numerical evaluation indicates a favorable high-order margin in the same neighborhoods where the low-order margin is weakest. A rigorous certificate requires a validated enclosure of the largest eigenvalue over the full domain. A successful full-domain enclosure defines the high-order matrix margin $\Lambda_{D^\alpha L}^{\mathrm{full}}$ carried into the compatibility condition below.

4.6 Compatibility of the Low-Order and High-Order Damping

Assuming the localized low-order bound and the high-order matrix bound above, the two estimates have complementary roles. The low-order estimate supplies the negative $\mathfrak{E}_0^2$ term, but localization introduces the positive contribution $A_{\mathcal{L}}^{k,\mathrm{loc}}\mathfrak{H}_k^2$. The high-order matrix supplies a negative $\mathfrak{H}_k^2$ contribution with margin $\Lambda_{D^\alpha L}^{\mathrm{full}}$. Since the full energy uses the high-order estimate with weight μ_k, these two top-order contributions must be compared after multiplying the high-order matrix estimate by μ_k.

Using (S231) and (S248) gives

$$\int_{\mathbb{D}} d\xi^\top \mathbb{M}_{\mathcal{L}} d\xi\, dr\, dz + \mu_k\int_{\mathbb{D}} \mathbf{V}_k^\top \mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}}\mathbf{V}_k\, dr\, dz \le -\Lambda_{\mathcal{L}}^{\mathrm{mat,loc}}\mathfrak{E}_0^2 - \left(\mu_k\Lambda_{D^\alpha L}^{\mathrm{full}} - A_{\mathcal{L}}^{k,\mathrm{loc}}\right)\mathfrak{H}_k^2. \tag{S249}$$

This inequality is the precise sense in which high-order damping compensates for a localized failure of low-order pointwise damping. The high-order estimate does not change $\lambda_{\max}(\mathbb{M}_{\mathcal{L}})$ inside the localized region. Instead, the concentration estimate converts the low-order contribution from that region into the top-order cost $A_{\mathcal{L}}^{k,\mathrm{loc}}\mathfrak{H}_k^2$, and the negative high-order term absorbs that cost after multiplication by μ_k.

Therefore the localized low-order loss is compatible with the high-order damping provided

$$A_{\mathcal{L}}^{k,\mathrm{loc}} < \mu_k \Lambda_{D^\alpha L}^{\mathrm{full}}. \tag{S250}$$

Together with $\Lambda_{\mathcal{L}}^{\mathrm{mat,loc}} > 0$, this condition gives strict negativity of both matrix-level energy coefficients. This comparison is still at the matrix level. The linear terms kept outside the two signed matrices are incorporated in Subsection 4.8 when the full stability margin is assembled.

4.7 Weighted PDE Residuals

The same weights used in the damping estimate also enter the norm in which the PDE residual is measured. Improving the signed linear matrix could therefore, in principle, amplify the weighted residual. It is important to check this after the damping weights have been selected rather than treating the two calculations independently.

For the present optimized weights, the current numerical evaluation indicates that the corresponding value of $\Lambda_{\mathcal{R}}$ remains small. Thus the weight choice that improves the linear damping does not appear to introduce a competing residual penalty.

The size of this residual matters differently from the nonlinear constants. In the final closing inequality, the residual enters through $\Lambda_{\mathcal{R}}/\delta_*$, whereas the nonlinear losses decrease as the stability radius δ_* is reduced. Consequently, a large residual cannot be compensated for simply by taking a smaller perturbation neighborhood. The weighted residual must therefore be evaluated after the damping weights have been fixed and carried explicitly into the closing inequality. If it is too large relative to the available margin, the profile itself must be refined in the weighted norms selected by the damping calculation.

4.8 Remaining Certification, Practical Next Steps, and Formalization

The preceding subsections reduce the required linear sign to explicit low-order and high-order matrix conditions. The low-order matrix is responsible for the principal $\mathfrak{E}_0^2$ damping, the localized estimate accounts for the small off-axis region where the same uniform pointwise margin is not imposed, and the high-order matrix provides the complementary top-order mechanism. The remaining task is to pass from these matrix-level conditions to the full stability inequality by incorporating the linear terms kept outside the matrices, the nonlinear losses, and the weighted residual.

From matrix damping to the full stability margin. To pass from the matrix-level bounds to the linear coefficients in the full energy estimate, we now include the linear terms that were intentionally kept outside the signed matrices. In the localized low-order estimate, the resulting constants are

$$\Lambda_{\mathcal{L}}^{\mathrm{low}} = \Lambda_{\mathcal{L}}^{\mathrm{mat,loc}} - A_{\mathcal{L}}^{0}, \qquad \Lambda_{\mathcal{L}}^{\mathrm{high}} = A_{\mathcal{L}}^{k} + A_{\mathcal{L}}^{k,\mathrm{loc}}. \tag{S251}$$

At high order, the corresponding constants are

$$\Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} = \Lambda_{D^\alpha L}^{\mathrm{full}} - A_{\mathrm{lin}}^{k,\mathrm{rem}}, \qquad \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}} = A_{\mathrm{lin}}^{0,\mathrm{rem}}. \tag{S252}$$

After choosing μ_k, the actual linear coercive margin is

$$\Lambda_{\mathrm{stab}}(\mu_k) = \min\left\{\Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}},\ \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k}\right\} > 0. \tag{S253}$$

A positive value of Λ_{stab} establishes coercivity before the nonlinear terms are invoked. On the boundary of a ball of radius δ_*, the cubic contribution is $\Lambda_3\delta_*$. These nonlinear contributions decrease when δ_* is reduced. Consequently, large but finite nonlinear constants usually shrink the stability neighborhood that can be certified rather than destroy an already positive linear sign. This makes the damping calculation particularly informative, because it isolates the linear sign condition that cannot be recovered merely by shrinking the perturbation neighborhood. The full closing condition is

$$\Lambda_3\delta_* + \frac{\Lambda_{\mathcal{R}}}{\delta_*} < \Lambda_{\mathrm{stab}}. \tag{S254}$$

Remaining quantitative certification. Closing the proof requires rigorous certification of the explicit constants entering the energy argument, which are presented in full detail in our companion stability paper [18]. These include the direct linear remainders, elliptic and interpolation constants, point-evaluation and modulation bounds, weighted multiplier constants, nonlinear product constants, and the parameters introduced by Cauchy–Schwarz and Young inequalities. This leads to a large but finite collection of scalar inequalities. The task is computationally extensive, but each required quantity is explicitly defined and has a prescribed route to rigorous certification.

The spline representation, interval arithmetic, matrix eigenvalue enclosures, and finite-dimensional optimization provide a direct route to this remaining verification with rigorous error control. If a particular contribution limits the final margin, the corresponding part of the construction can be refined without changing the overall proof mechanism. For example, the profile can be refined preferentially where the weighted residual is largest, while the weights, spline resolution, localization parameters, or analytic estimates can be adjusted when they are the limiting factors. The modular organization of the proof is intended precisely to make these quantitative refinements local to the relevant estimate.

If the complete collection of remaining constants can be rigorously certified with a positive closing margin, then the nonlinear stability argument developed here is fully closed. Together with the finite-time reconstruction from the dynamically rescaled variables, this yields an admissible solution that remains asymptotic to the candidate self-similar profile and becomes singular in finite physical time. Thus the principal remaining work along this route is the rigorous certification and quantitative optimization of an already specified proof framework, rather than the discovery of a new stability mechanism or a fundamentally different argument.

Appendix A. Convection–Varied Euler System

A.1 Convection-varied equations

Starting from the axisymmetric equations, consider the following variation on the convection terms:

$$\partial_t u_\bullet + \varepsilon\, u^r\, \partial_{\mathring{r}} u_\bullet + \varepsilon\, u^z\, \partial_{\mathring{z}} u_\bullet = 2u_\bullet\, \partial_{\mathring{z}} \psi_\bullet, \tag{S255}$$

$$\partial_t \omega_\bullet + \varepsilon\, u^r\, \partial_{\mathring{r}} \omega_\bullet + \varepsilon\, u^z\, \partial_{\mathring{z}} \omega_\bullet = 2u_\bullet\, \partial_{\mathring{z}} u_\bullet, \tag{S256}$$

$$-\Big[\partial_{\mathring{r}}^2 + \frac{3}{\mathring{r}}\partial_{\mathring{r}} + \partial_{\mathring{z}}^2\Big]\psi_\bullet = \omega_\bullet, \tag{S257}$$

$$u^r = -\,\mathring{r}\, \partial_{\mathring{z}}\psi_\bullet, \qquad u^z = 2\psi_\bullet + \mathring{r}\, \partial_{\mathring{r}}\psi_\bullet, \tag{S258}$$

with amplification factor ε, which is used to vary the strength of convection that exist within the Euler equations, with $\varepsilon = 1$ corresponding to the full Euler system.

When $\varepsilon = 0$, this corresponds to the non-convective blowup scenario proposed by Hou and Lei [51]. However, when convection is included, this finite-time blowup scenario can be destroyed. Hou et al. [19] studied the amplification of this ε term, with their numerical investigation suggesting that for $\varepsilon \le 0.3$ the equations exhibit singularity formation, while for larger ε the solution does not develop a singularity.

A.2 Traveling-wave profile equations

Proposition S5 (Traveling-wave profile equations for the convection-varied system)**.** *Assume that* $(u_\bullet, \omega_\bullet, \psi_\bullet)$ *is given by the ansatz* (S13)–(S15)*, that the center satisfies* (S18)*, and that the exponents are chosen via* (S21)*. Then the transformed profile* (U, Ω, Ψ) *satisfies*

$$U + \big(\lambda r + \varepsilon U^r\big)\partial_r U + \big(C + \lambda z + \varepsilon U^z\big)\partial_z U = 2U\partial_z\Psi, \tag{S259}$$

$$(1+\lambda)\Omega + \big(\lambda r + \varepsilon U^r\big)\partial_r \Omega + \big(C + \lambda z + \varepsilon U^z\big)\partial_z \Omega = 2U\partial_z U, \tag{S260}$$

$$-\,\mathcal{E}\,\Psi = \Omega, \tag{S261}$$

where

$$U^r = -r\partial_z\Psi, \qquad U^z = 2\Psi + r\partial_r\Psi. \tag{S262}$$

Equivalently, the system can be written in terms of U, Ω, Ψ *only as*

$$U + r\big(\lambda - \varepsilon\partial_z\Psi\big)\partial_r U + \big(C + \lambda z + \varepsilon(2\Psi + r\partial_r\Psi)\big)\partial_z U = 2U\partial_z\Psi, \tag{S263}$$

$$(1+\lambda)\Omega + r\big(\lambda - \varepsilon\partial_z\Psi\big)\partial_r \Omega + \big(C + \lambda z + \varepsilon(2\Psi + r\partial_r\Psi)\big)\partial_z \Omega = 2U\partial_z U, \tag{S264}$$

$$-\,\mathcal{E}\,\Psi = \Omega. \tag{S265}$$

Proof (✓ Formalized and verified in Lean)

Roadmap. The factors produced by time differentiation, spatial differentiation, convection, stretching, and the elliptic streamfunction relation are balanced to determine the amplitude exponents in (S21). After the common power of $T - t$ is removed, the Euler equations yield a time-independent profile problem.

Set

$$\tau = T - t, \qquad r = \frac{\mathring{r}}{\tau^{\lambda}}, \qquad z = \frac{\mathring{z} - \mathring{z}_c(t)}{\tau^{\lambda}}. \tag{S266}$$

The traveling-center condition (S18) is

$$\frac{d}{dt}\mathring{z}_c(t) = -C\tau^{\lambda-1}. \tag{S267}$$

Differentiating at fixed physical coordinates $(\mathring{r}, \mathring{z})$ gives

$$\partial_t r = \frac{\lambda r}{\tau}, \qquad \partial_t z = \frac{\lambda z + C}{\tau}. \tag{S268}$$

Moreover,

$$\partial_{\mathring{r}} = \tau^{-\lambda}\partial_r, \qquad \partial_{\mathring{z}} = \tau^{-\lambda}\partial_z. \tag{S269}$$

From the streamfunction relations in the physical variables,

$$u^r_{\text{phys}} = -\mathring{r}\,\partial_{\mathring{z}}\psi_{\bullet}, \qquad u^z_{\text{phys}} = 2\psi_{\bullet} + \mathring{r}\,\partial_{\mathring{r}}\psi_{\bullet}. \tag{S270}$$

Therefore,

$$\psi_{\bullet}(\mathring{r}, \mathring{z}, t) = \tau^{c_\psi}\Psi(r, z), \tag{S271}$$

gives

$$u^r_{\text{phys}} = -\mathring{r}\,\tau^{c_\psi-\lambda}\partial_z\Psi = \tau^{c_\psi}\big(-r\partial_z\Psi\big), \tag{S272}$$

$$u^z_{\text{phys}} = 2\tau^{c_\psi}\Psi + \mathring{r}\,\tau^{c_\psi-\lambda}\partial_r\Psi = \tau^{c_\psi}\big(2\Psi + r\partial_r\Psi\big). \tag{S273}$$

Thus

$$u^r_{\text{phys}} = \tau^{c_\psi}U^r, \qquad u^z_{\text{phys}} = \tau^{c_\psi}U^z, \qquad U^r = -r\partial_z\Psi, \qquad U^z = 2\Psi + r\partial_r\Psi. \tag{S274}$$

We first derive the equation for U. The ansatz

$$u_{\bullet}(\mathring{r}, \mathring{z}, t) = \tau^{c_u}U(r, z) \tag{S275}$$

and (S268) give

$$\partial_t u_{\bullet} = -c_u\tau^{c_u-1}U + \tau^{c_u}\left(\partial_r U\,\partial_t r + \partial_z U\,\partial_t z\right) = \tau^{c_u-1}\left[-c_u U + \lambda r\partial_r U + (\lambda z + C)\partial_z U\right]. \tag{S276}$$

The spatial derivatives are

$$\partial_{\mathring{r}}u_{\bullet} = \tau^{c_u-\lambda}\,\partial_r U, \qquad \partial_{\mathring{z}}u_{\bullet} = \tau^{c_u-\lambda}\,\partial_z U. \tag{S277}$$

Using (S274), the convection term is

$$\varepsilon\,u^r_{\text{phys}}\,\partial_{\mathring{r}}u_{\bullet} + \varepsilon\,u^z_{\text{phys}}\,\partial_{\mathring{z}}u_{\bullet} = \tau^{c_u+c_\psi-\lambda}\,\varepsilon\left(U^r\partial_r U + U^z\partial_z U\right). \tag{S278}$$

The stretching term is

$$2u_\bullet\,\partial_{\mathring{z}}\psi_\bullet = 2\tau^{c_u+c_\psi-\lambda}\,U\,\partial_z\Psi. \tag{S279}$$

Substitution into the physical $u_\bullet$ equation gives

$$\tau^{c_u-1}\left[-c_uU+\lambda r\partial_rU+(\lambda z+C)\partial_zU\right]+\tau^{c_u+c_\psi-\lambda}\,\varepsilon\,\left(U^r\partial_rU+U^z\partial_zU\right)=2\,\tau^{c_u+c_\psi-\lambda}\,U\partial_z\Psi. \tag{S280}$$

The time derivative, convection, and stretching terms have the same power of τ when

$$c_u-1=c_u+c_\psi-\lambda, \qquad \text{hence} \qquad c_\psi=\lambda-1. \tag{S281}$$

After division by τ^{c_u-1}, the equation becomes

$$-c_uU+\big(\lambda r+\varepsilon U^r\big)\partial_rU+\big(C+\lambda z+\varepsilon U^z\big)\partial_zU=2U\partial_z\Psi. \tag{S282}$$

Using $c_u=-1$ gives (S259).

We next derive the equation for Ω.

From

$$\omega_\bullet(\mathring{r},\mathring{z},t)=\tau^{c_\omega}\Omega(r,z) \tag{S283}$$

we obtain

$$\partial_t\omega_\bullet=\tau^{c_\omega-1}\left[-c_\omega\Omega+\lambda r\partial_r\Omega+(\lambda z+C)\partial_z\Omega\right], \tag{S284}$$

and

$$\partial_{\mathring{r}}\omega_\bullet=\tau^{c_\omega-\lambda}\,\partial_r\Omega, \qquad\qquad \partial_{\mathring{z}}\omega_\bullet=\tau^{c_\omega-\lambda}\,\partial_z\Omega. \tag{S285}$$

The convection term is

$$\varepsilon\,u^r_{\mathrm{phys}}\,\partial_{\mathring{r}}\omega_\bullet+\varepsilon\,u^z_{\mathrm{phys}}\,\partial_{\mathring{z}}\omega_\bullet=\tau^{c_\omega+c_\psi-\lambda}\,\varepsilon\,\left(U^r\partial_r\Omega+U^z\partial_z\Omega\right). \tag{S286}$$

Using (S281), this exponent is $c_\omega-1$.

The stretching term is

$$2u_\bullet\,\partial_{\mathring{z}}u_\bullet=2\,\tau^{2c_u-\lambda}\,U\partial_zU. \tag{S287}$$

Consequently,

$$\tau^{c_\omega-1}\left[-c_\omega\Omega+\lambda r\partial_r\Omega+(\lambda z+C)\partial_z\Omega\right]+\tau^{c_\omega-1}\,\varepsilon\,(U^r\partial_r\Omega+U^z\partial_z\Omega)=2\,\tau^{2c_u-\lambda}\,U\partial_zU. \tag{S288}$$

Matching the remaining powers gives

$$c_\omega-1=2c_u-\lambda. \tag{S289}$$

After division by $\tau^{c_\omega - 1}$, we obtain

$$-c_\omega \Omega + \big(\lambda r + \varepsilon U^r\big)\partial_r \Omega + \big(C + \lambda z + \varepsilon U^z\big)\partial_z \Omega = 2U\partial_z U. \tag{S290}$$

Using $c_\omega = -(1+\lambda)$ gives (S260).

It remains to check the elliptic streamfunction equation.

Since

$$\partial_{\mathring{r}\mathring{r}}\psi_\bullet = \tau^{c_\psi - 2\lambda}\,\partial_{rr}\Psi, \qquad \partial_{\mathring{z}\mathring{z}}\psi_\bullet = \tau^{c_\psi - 2\lambda}\,\partial_{zz}\Psi, \tag{S291}$$

and since $\mathring{r} = r\tau^\lambda$, we also have

$$\frac{3}{\mathring{r}}\,\partial_{\mathring{r}}\psi_\bullet = \frac{3}{r\tau^\lambda}\,\tau^{c_\psi - \lambda}\,\partial_r \Psi = \tau^{c_\psi - 2\lambda}\,\frac{3}{r}\,\partial_r \Psi. \tag{S292}$$

Therefore,

$$-\mathcal{E}_{\mathring{r},\mathring{z}}\psi_\bullet = -\tau^{c_\psi - 2\lambda}\mathcal{E}\Psi. \tag{S293}$$

The right-hand side of the physical elliptic equation is

$$\omega_\bullet = \tau^{c_\omega}\Omega. \tag{S294}$$

The common power cancels when

$$c_\omega = c_\psi - 2\lambda. \tag{S295}$$

Together with (S281), this gives $c_\omega = -1 - \lambda$. Combining this identity with (S289) gives $c_u = -1$.

After cancelling the common power of τ, the elliptic equation is

$$-\mathcal{E}\Psi = \Omega. \tag{S296}$$

■

In these profile equations,

- The terms $\lambda r\partial_r$ and $\lambda z\partial_z$ come from observing the solution in a shrinking coordinate system.
- The term $C\partial_z$ comes from following the center of the profile as it travels along the symmetry axis.
- The terms involving εU^r and εU^z are the convection terms from the convection-varied Euler system.
- The terms $2U\partial_z\Psi$ and $2U\partial_z U$ on the right-hand side are the stretching and coupling terms inherited from the axisymmetric formulation.

For the convection-varied Euler equations, these profile equations are the steady-state equations for the traveling self-similar profile. Thus, for the convection-varied Euler equations, the traveling-wave ansatz reduces the study of a possible finite-time singularity to an autonomous nonlinear profile problem for (U, Ω, Ψ) and the parameters λ and C. The parameter λ determines the rate at which the spatial scale collapses, while C determines the speed of the axial drift in the rescaled frame.

A.3 Dynamic rescaling equations for the convection-varied systems

As before, define the elliptic operator

$$\mathcal{E} := \partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2. \tag{S297}$$

The proposition below provides the convection-varied axisymmetric Euler equations in the dynamically rescaled variables.

Proposition S6 (Dynamic rescaling equations)**.** *In the dynamically rescaled variables, the convection-varied axisymmetric system takes the form*

$$\partial_t u + \big(\lambda(t)r + \varepsilon u^r\big)\partial_r u + \big(\lambda(t)z + C(t) + \varepsilon u^z\big)\partial_z u = 2u\partial_z\psi + c_u(t)u, \tag{S298}$$

$$\partial_t \omega + \big(\lambda(t)r + \varepsilon u^r\big)\partial_r \omega + \big(\lambda(t)z + C(t) + \varepsilon u^z\big)\partial_z \omega = 2u\partial_z u + c_\omega(t)\omega, \tag{S299}$$

$$-\mathcal{E}\psi = \omega, \tag{S300}$$

with

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi, \tag{S301}$$

and

$$c_\omega(t) = c_u(t) - \lambda(t). \tag{S302}$$

Proof The proof is provided in Appendix D. ■

Normalization conditions. The independent *modulation parameters* are $\lambda(t)$, $C(t)$, and $c_u(t)$. The vorticity amplitude rate is determined by the scaling relation

$$c_\omega(t) = c_u(t) - \lambda(t). \tag{S303}$$

Choose the axial translation of the reference profile so that

$$\bar{C} + \varepsilon \bar{u}_0^z = \bar{C} + 2\varepsilon\bar{\psi}_0 = 0. \tag{S304}$$

We impose the normalization conditions

$$\begin{aligned} &C(t) + \varepsilon u^z(0,0,t) = 0, \\ &\partial_t\big(\partial_z\omega\big)(0,0,t) = 0, \qquad u(0,0,t) = \bar{u}_0. \end{aligned} \tag{S305}$$

The first condition fixes the meridional fixed point at the origin. The second preserves the initially prescribed value of $\partial_z\omega(0,0,t)$; evaluating the z-differentiated vorticity equation at the origin determines $\lambda(t)$. The third fixes the amplitude at the origin and therefore implies $\partial_t u(0,0,t) = 0$. Together these conditions determine $\lambda(t)$, $C(t)$, and $c_u(t)$. In perturbation variables, the stagnation and amplitude conditions are

$$\delta C + 2\varepsilon\delta\psi_0 = 0, \qquad \delta u_0 = 0, \tag{S306}$$

so that $\delta C = -2\varepsilon\delta\psi_0$.

Limiting behaviour. If the dynamic rescaling captures a stable traveling self-similar blowup, the rescaled solution should converge as rescaled time advances to a time-independent profile,

$$u(r,z,t) \to u_\infty(r,z), \qquad \omega(r,z,t) \to \omega_\infty(r,z), \qquad \psi(r,z,t) \to \psi_\infty(r,z), \tag{S307}$$

and the modulation parameters should converge to constants,

$$\lambda(t) \to \lambda_\infty, \qquad C(t) \to C_\infty, \qquad c_u(t) \to c_{u,\infty}, \qquad c_\omega(t) \to c_{\omega,\infty},. \tag{S308}$$

In that limit, the time derivatives in (S298)–(S299) disappear. The steady dynamic equations are

$$\Big(\lambda_\infty\, r + \varepsilon u^r_\infty\Big)\partial_r u_\infty + \Big(\lambda_\infty\, z + C_\infty + \varepsilon u^z_\infty\Big)\partial_z u_\infty = 2u_\infty\, \partial_z \psi_\infty + c_{u,\infty}\, u_\infty, \tag{S309}$$

$$\Big(\lambda_\infty\, r + \varepsilon u^r_\infty\Big)\partial_r \omega_\infty + \Big(\lambda_\infty\, z + C_\infty + \varepsilon u^z_\infty\Big)\partial_z \omega_\infty = 2u_\infty\, \partial_z u_\infty + c_{\omega,\infty}\, \omega_\infty, \tag{S310}$$

with $u^r_\infty = -r\partial_z\psi_\infty$ and $u^z_\infty = 2\psi_\infty + r\partial_r\psi_\infty$.

For the traveling-wave ansatz in (S13)–(S15), the limiting amplitude rates are

$$c_{u,\infty} = -1, \qquad c_{\omega,\infty} = -\big(1 + \lambda_\infty\big). \tag{S311}$$

In the Euler setting, substituting (S311) into (S309)–(S310) gives the traveling-wave profile equations for general λ_∞.

Therefore, the steady states of the Euler dynamic rescaling system reproduce the Euler traveling-wave profile equations for general λ_∞. After writing

$$\big(u_\infty, \omega_\infty, \psi_\infty, C_\infty\big) = \big(U, \Omega, \Psi, C\big), \tag{S312}$$

the limiting dynamic equations agree with the corresponding profile equations. In this sense, the traveling-wave profiles are fixed points of the appropriate dynamically rescaled evolution.

Appendix B. PDE in the Computational Domain

We use the coordinate transformation

$$r = \sinh(r_{\mathtt{comp}})\,, \qquad z = \sinh(z_{\mathtt{comp}})\,, \tag{S313}$$

and write the corresponding profiles in computational coordinates as

$$(u, \omega, \psi)(r_{\mathtt{comp}}, z_{\mathtt{comp}}) = (U, \Omega, \Psi)(\sinh(r_{\mathtt{comp}}), \sinh(z_{\mathtt{comp}}))\,. \tag{S314}$$

Since

$$\frac{dr}{dr_{\mathtt{comp}}} = \cosh(r_{\mathtt{comp}})\,, \qquad \frac{dz}{dz_{\mathtt{comp}}} = \cosh(z_{\mathtt{comp}})\,, \tag{S315}$$

the chain rule gives

$$\partial_r = \frac{1}{\cosh(r_{\mathtt{comp}})}\partial_{r_{\mathtt{comp}}}\,, \qquad \partial_z = \frac{1}{\cosh(z_{\mathtt{comp}})}\partial_{z_{\mathtt{comp}}}\,. \tag{S316}$$

The derivatives are related by

$$f_r = \frac{f_{r_{\mathtt{comp}}}}{\cosh(r_{\mathtt{comp}})}, \qquad f_z = \frac{f_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})}. \tag{S317}$$

Applying the chain rule once more yields

$$f_{rr} = \frac{f_{r_{\mathtt{comp}} r_{\mathtt{comp}}}}{\cosh^2(r_{\mathtt{comp}})} - \frac{\sinh(r_{\mathtt{comp}})}{\cosh^3(r_{\mathtt{comp}})} f_{r_{\mathtt{comp}}}\,, \qquad f_{zz} = \frac{f_{z_{\mathtt{comp}} z_{\mathtt{comp}}}}{\cosh^2(z_{\mathtt{comp}})} - \frac{\sinh(z_{\mathtt{comp}})}{\cosh^3(z_{\mathtt{comp}})} f_{z_{\mathtt{comp}}}\,. \tag{S318}$$

The elliptic equation involves $f_{rr} + \frac{3}{r} f_r + f_{zz}$. In computational coordinates, we denote this expression by $\mathcal{E}_{\mathtt{comp}} f$. Using (S317)–(S318),

$$\begin{aligned} \mathcal{E}_{\mathtt{comp}} f := \frac{f_{r_{\mathtt{comp}} r_{\mathtt{comp}}}}{\cosh^2(r_{\mathtt{comp}})} - \frac{\sinh(r_{\mathtt{comp}})}{\cosh^3(r_{\mathtt{comp}})} f_{r_{\mathtt{comp}}} + \frac{3 f_{r_{\mathtt{comp}}}}{\sinh(r_{\mathtt{comp}}) \cosh(r_{\mathtt{comp}})} \\ + \frac{f_{z_{\mathtt{comp}} z_{\mathtt{comp}}}}{\cosh^2(z_{\mathtt{comp}})} - \frac{\sinh(z_{\mathtt{comp}})}{\cosh^3(z_{\mathtt{comp}})} f_{z_{\mathtt{comp}}}\,. \end{aligned} \tag{S319}$$

The factor $1/\sinh(r_{\mathtt{comp}})$ in (S319) appears singular at the axis. If the profile is smooth and even in r, then it is also even in $r_{\mathtt{comp}}$. Near the axis,

$$f_{r_{\mathtt{comp}}} = f_{r_{\mathtt{comp}} r_{\mathtt{comp}}}(0, z_{\mathtt{comp}})\, r_{\mathtt{comp}} + \mathcal{O}(r_{\mathtt{comp}}), \qquad \sinh(r_{\mathtt{comp}}) \cosh(r_{\mathtt{comp}}) = r_{\mathtt{comp}} + \mathcal{O}(r_{\mathtt{comp}}). \tag{S320}$$

Therefore,

$$\lim_{r_{\mathtt{comp}} \to 0} \frac{f_{r_{\mathtt{comp}}}}{\sinh(r_{\mathtt{comp}}) \cosh(r_{\mathtt{comp}})} = f_{r_{\mathtt{comp}} r_{\mathtt{comp}}}(0, z_{\mathtt{comp}}), \tag{S321}$$

The apparent singularity is therefore removable. At the axis, the radial part of $\mathcal{E}_{\mathtt{comp}} f$ extends continuously to $4 f_{r_{\mathtt{comp}} r_{\mathtt{comp}}}(0, z_{\mathtt{comp}})$.

The remaining factors in the transport terms are

$$r f_r = \frac{\sinh(r_{\mathtt{comp}})}{\cosh(r_{\mathtt{comp}})} f_{r_{\mathtt{comp}}}, \qquad z = \sinh(z_{\mathtt{comp}}). \tag{S322}$$

Substituting these identities into the profile equations gives the Euler system in computational coordinates

$$\begin{aligned} 0 = u - \frac{2u\,\psi_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})} + \frac{\sinh(r_{\mathtt{comp}})}{\cosh(r_{\mathtt{comp}})}\, u_{r_{\mathtt{comp}}} \left(\lambda - \varepsilon\, \frac{\psi_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})}\right) \\ + \frac{u_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})} \left(C + \lambda \sinh(z_{\mathtt{comp}}) + \varepsilon \left(2\psi + \frac{\sinh(r_{\mathtt{comp}})}{\cosh(r_{\mathtt{comp}})} \psi_{r_{\mathtt{comp}}}\right)\right), \end{aligned} \tag{S323}$$

$$\begin{aligned} 0 = (1+\lambda)\omega - \frac{2u\,u_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})} + \frac{\sinh(r_{\mathtt{comp}})}{\cosh(r_{\mathtt{comp}})}\, \omega_{r_{\mathtt{comp}}} \left(\lambda - \varepsilon\, \frac{\psi_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})}\right) \\ + \frac{\omega_{z_{\mathtt{comp}}}}{\cosh(z_{\mathtt{comp}})} \left(C + \lambda \sinh(z_{\mathtt{comp}}) + \varepsilon \left(2\psi + \frac{\sinh(r_{\mathtt{comp}})}{\cosh(r_{\mathtt{comp}})} \psi_{r_{\mathtt{comp}}}\right)\right), \end{aligned} \tag{S324}$$

$$0 = \omega + \mathcal{E}_{\mathtt{comp}}\, \psi, \tag{S325}$$

where ε, as before, denotes the convection amplification factor, and

$$\begin{aligned} \mathcal{E}_{\mathtt{comp}} f := \frac{f_{r_{\mathtt{comp}} r_{\mathtt{comp}}}}{\cosh^2(r_{\mathtt{comp}})} - \frac{\sinh(r_{\mathtt{comp}})}{\cosh^3(r_{\mathtt{comp}})} f_{r_{\mathtt{comp}}} + \frac{3 f_{r_{\mathtt{comp}}}}{\sinh(r_{\mathtt{comp}})\cosh(r_{\mathtt{comp}})} \\ + \frac{f_{z_{\mathtt{comp}} z_{\mathtt{comp}}}}{\cosh^2(z_{\mathtt{comp}})} - \frac{\sinh(z_{\mathtt{comp}})}{\cosh^3(z_{\mathtt{comp}})} f_{z_{\mathtt{comp}}}. \end{aligned} \tag{S326}$$

As before, the full Euler system is obtained by setting $\varepsilon = 1$.

✓ *These formulas have been formalized and verified in Lean.*

Appendix C. Derivations of the Profile Equations

Proposition S7 (Traveling-wave profile equations). *Assume that* $(u_\bullet, \omega_\bullet, \psi_\bullet)$ *is given by the ansatz* (S13)–(S15)*, that the center satisfies* (S18)*, and that the exponents are chosen as in* (S21)*. Set*

$$\lambda = \frac{1}{2}. \tag{S327}$$

Then (U, Ω, Ψ) *satisfies*

$$U + \big(\lambda r + U^r\big)\partial_r U + \big(C + \lambda z + U^z\big)\partial_z U = 2U\partial_z\Psi, \tag{S328}$$
$$(1+\lambda)\Omega + \big(\lambda r + U^r\big)\partial_r \Omega + \big(C + \lambda z + U^z\big)\partial_z \Omega = 2U\partial_z U, \tag{S329}$$
$$-\mathcal{E}\,\Psi = \Omega, \tag{S330}$$

where

$$U^r = -r\partial_z\Psi, \qquad U^z = 2\Psi + r\partial_r\Psi. \tag{S331}$$

Equivalently,

$$U + r\big(\lambda - \partial_z\Psi\big)\partial_r U + \big(C + \lambda z + 2\Psi + r\partial_r\Psi\big)\partial_z U = 2U\partial_z\Psi, \tag{S332}$$
$$(1+\lambda)\Omega + r\big(\lambda - \partial_z\Psi\big)\partial_r \Omega + \big(C + \lambda z + 2\Psi + r\partial_r\Psi\big)\partial_z \Omega = 2U\partial_z U, \tag{S333}$$
$$-\mathcal{E}\,\Psi = \Omega. \tag{S334}$$

Proof (✓ Formalized and verified in Lean)

We specialize the convection-varied calculation in the proof of Proposition S5 to $\varepsilon = 1$ and $\lambda = 1/2$.

The scaling exponents in (S21) become

$$c_u = -1, \qquad c_\omega = -\frac{3}{2}, \qquad c_\psi = -\frac{1}{2}. \tag{S335}$$

The coordinate identities in (S268) and (S269) therefore reduce to

$$\partial_t r = \frac{r}{2\tau}, \qquad \partial_t z = \frac{z/2 + C}{\tau}, \qquad \partial_{\mathring{r}} = \tau^{-1/2}\partial_r, \qquad \partial_{\mathring{z}} = \tau^{-1/2}\partial_z. \tag{S336}$$

For the $u_\bullet$ equation, the time derivative is

$$\partial_t u_\bullet = \tau^{-2}\left[U + \frac{1}{2}r\partial_r U + \left(C + \frac{1}{2}z\right)\partial_z U\right]. \tag{S337}$$

Because $c_\psi = -\frac{1}{2}$, the physical meridian velocity components satisfy

$$u^r_{\text{phys}} = \tau^{-1/2}\,U^r, \qquad u^z_{\text{phys}} = \tau^{-1/2}\,U^z, \qquad U^r = -r\partial_z\Psi, \qquad U^z = 2\Psi + r\partial_r\Psi. \tag{S338}$$

Also,

$$\partial_{\mathring{r}} u_\bullet = \tau^{-3/2}\partial_r U, \qquad \partial_{\mathring{z}} u_\bullet = \tau^{-3/2}\partial_z U. \tag{S339}$$

Hence the convection terms and the stretching term $2u_\bullet\,\partial_{\mathring{z}}\psi_\bullet$ all carry the same factor τ^{-2} as the time derivative.

For the vorticity equation,

$$\partial_t \omega_\bullet = \tau^{-5/2} \left[\frac{3}{2}\Omega + \frac{1}{2} r \partial_r \Omega + \left(C + \frac{1}{2} z \right) \partial_z \Omega \right]. \tag{S340}$$

The convection terms scale like $\tau^{-1/2}\tau^{-2} = \tau^{-5/2}$.

The stretching term $2u_\bullet \, \partial_{\mathring{z}} u_\bullet$ scales like $\tau^{-1}\tau^{-3/2} = \tau^{-5/2}$.

Canceling the common factor gives the desired equation.

Finally,

$$c_\psi - 2\lambda = -\frac{1}{2} - 1 = -\frac{3}{2} = c_\omega, \tag{S341}$$

so both sides of the physical elliptic equation have the factor $\tau^{-3/2}$. Canceling it gives the desired equation.
∎

Appendix D. Derivation of the Dynamic Rescaling Equations

In this Appendix, we derive the dynamic rescaling equations for the convection-varied system in the more general setting where the spatial rescaling rate λ is time-dependent, since in the convection-varied case one does not necessarily have $\lambda = 1/2$. The corresponding equations for the Euler system are recovered as the special case $\varepsilon = 1$ and $\lambda = 1/2$.

Proposition S8 (Dynamic rescaling equations)**.** *In the dynamically rescaled variables, the convection-varied axisymmetric system takes the form*

$$\partial_t u + \big(\lambda(t) r + \varepsilon u^r\big)\partial_r u + \big(\lambda(t) z + C(t) + \varepsilon u^z\big)\partial_z u = 2u\partial_z\psi + c_u(t)u \tag{S342}$$
$$\partial_t \omega + \big(\lambda(t) r + \varepsilon u^r\big)\partial_r \omega + \big(\lambda(t) z + C(t) + \varepsilon u^z\big)\partial_z \omega = 2u\partial_z u + c_\omega(t)\omega \tag{S343}$$
$$-\mathcal{E}\psi = \omega, \tag{S344}$$

with

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi, \tag{S345}$$

and

$$c_\omega(t) = c_u(t) - \lambda(t). \tag{S346}$$

Proof (✓ Formalized and verified in Lean)

Coordinate and amplitude scaling. At fixed physical coordinates, differentiating equation (S46) and using (S50)–(S51) gives

$$(\partial_t r)_{\mathring{r},\mathring{z}} = \lambda(t) r, \qquad (\partial_t z)_{\mathring{r},\mathring{z}} = \lambda(t) z + C(t). \tag{S347}$$

For any rescaled scalar field $f(r, z, t)$,

$$(\partial_t f)_{\mathring{r},\mathring{z}} = \partial_t f + \lambda(t) r\,\partial_r f + \big(\lambda(t) z + C(t)\big)\,\partial_z f. \tag{S348}$$

Applying this identity to the amplitudes in (S47)–(S48) gives

$$\partial_{\mathring{t}} u_\bullet = \mathrm{s}_u \frac{dt}{d\mathring{t}} \left[\partial_t u + \lambda r\,\partial_r u + (\lambda z + C)\partial_z u - c_u u\right], \tag{S349}$$
$$\partial_{\mathring{t}} \omega_\bullet = \mathrm{s}_\omega \frac{dt}{d\mathring{t}} \left[\partial_t \omega + \lambda r\,\partial_r \omega + (\lambda z + C)\partial_z \omega - c_\omega \omega\right]. \tag{S350}$$

The streamfunction scaling gives the physical meridian velocities

$$u^r_{\text{phys}} = \mathrm{s}_\omega\,\mathrm{s}_r^2\,u^r, \qquad u^z_{\text{phys}} = \mathrm{s}_\omega\,\mathrm{s}_r^2\,u^z, \tag{S351}$$

where

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi. \tag{S352}$$

Since $\partial_{\mathring{r}} = \mathrm{s}_r^{-1}\partial_r$ and $\partial_{\mathring{z}} = \mathrm{s}_r^{-1}\partial_z$, the physical convection and stretching terms in the $u_\bullet$ equation have the common coefficient

$$\mathrm{s}_\omega\,\mathrm{s}_r^2\,\frac{\mathrm{s}_u}{\mathrm{s}_r} = \mathrm{s}_u\,\mathrm{s}_\omega\,\mathrm{s}_r = \mathrm{s}_u\,\frac{dt}{d\mathring{t}}. \tag{S353}$$

In the vorticity equation, the convection coefficient is $\mathrm{s}_\omega^2\, \mathrm{s}_r$, while the stretching coefficient is

$$\frac{\mathrm{s}_u^2}{\mathrm{s}_r} = \mathrm{s}_\omega^2\, \mathrm{s}_r = \mathrm{s}_\omega\, \frac{dt}{d\mathring{t}}. \tag{S354}$$

Thus the relations in (S49) normalize all terms in both equations. Taking the logarithmic derivative of $\mathrm{s}_u = \mathrm{s}_\omega\, \mathrm{s}_r$ gives

$$c_\omega = c_u - \lambda. \tag{S355}$$

The rescaled u equation. Applying the coordinate chain rule (S348), adding the physical convection terms, and including the amplitude normalization gives

$$\partial_t u + \lambda(t) r \partial_r u + \big(\lambda(t) z + C(t)\big)\partial_z u + \varepsilon u^r \partial_r u + \varepsilon u^z \partial_z u - c_u(t) u = 2u\partial_z \psi. \tag{S356}$$

The first two transport terms come from the moving rescaled frame. The two terms multiplied by ε are the physical convection terms.

Grouping the radial and axial derivatives gives

$$\partial_t u + \big(\lambda(t) r + \varepsilon u^r\big)\partial_r u + \big(\lambda(t) z + C(t) + \varepsilon u^z\big)\partial_z u - c_u(t) u = 2u\partial_z \psi. \tag{S357}$$

Moving the amplitude term to the right gives

$$\partial_t u + \big(\lambda(t) r + \varepsilon u^r\big)\partial_r u + \big(\lambda(t) z + C(t) + \varepsilon u^z\big)\partial_z u = 2u\partial_z \psi + c_u(t) u. \tag{S358}$$

The rescaled vorticity equation. The moving frame contributes

$$\partial_t \omega + \lambda(t) r \partial_r \omega + \big(\lambda(t) z + C(t)\big)\partial_z \omega, \tag{S359}$$

while the physical convection contributes

$$\varepsilon u^r \partial_r \omega + \varepsilon u^z \partial_z \omega. \tag{S360}$$

The stretching term is $2u\partial_z u$, and the amplitude normalization contributes $-c_\omega(t)\omega$ on the left.

Thus

$$\partial_t \omega + \lambda(t) r \partial_r \omega + \big(\lambda(t) z + C(t)\big)\partial_z \omega + \varepsilon u^r \partial_r \omega + \varepsilon u^z \partial_z \omega - c_\omega(t)\omega = 2u\partial_z u. \tag{S361}$$

Combining the frame transport with the physical convection gives

$$\partial_t \omega + \big(\lambda(t) r + \varepsilon u^r\big)\partial_r \omega + \big(\lambda(t) z + C(t) + \varepsilon u^z\big)\partial_z \omega - c_\omega(t)\omega = 2u\partial_z u. \tag{S362}$$

Moving the amplitude term to the right gives

$$\partial_t \omega + \big(\lambda(t) r + \varepsilon u^r\big)\partial_r \omega + \big(\lambda(t) z + C(t) + \varepsilon u^z\big)\partial_z \omega = 2u\partial_z u + c_\omega(t)\omega. \tag{S363}$$

The elliptic equation and velocity recovery. The streamfunction scaling gives

$$-\mathcal{E}_{\mathring{r},\mathring{z}}\psi_\bullet = -\mathrm{s}_\omega(t)\,\mathcal{E}\,\psi, \qquad \omega_\bullet = \mathrm{s}_\omega(t)\,\omega. \tag{S364}$$

Canceling the common factor $\mathrm{s}_\omega(t)$ gives

$$-\mathcal{E}\psi = \omega. \tag{S365}$$

Starting from the physical recovery formulas

$$u^r_{\text{phys}} = -\mathring{r}\partial_{\mathring{z}}\psi_\bullet, \qquad u^z_{\text{phys}} = 2\psi_\bullet + \mathring{r}\partial_{\mathring{r}}\psi_\bullet, \tag{S366}$$

and using (S46)–(S48), the common factor $\mathrm{s}_\omega(t)\,\mathrm{s}_r(t)^2$ cancels. This gives

$$u^r = -r\partial_z\psi, \qquad u^z = 2\psi + r\partial_r\psi. \tag{S367}$$

■

Appendix E. Energy Identities

In this appendix, we consider the energy functional for the convection-varied system,

$$E_\varepsilon = \int_{-\infty}^{\infty} \int_0^{\infty} \left(|u_\bullet|^2 + (2-\varepsilon) |\nabla_{\mathring{r},\mathring{z}} \psi_\bullet|^2 \right) \mathring{r}^3 \, d\mathring{r} \, d\mathring{z}. \tag{S368}$$

As before, the energy functional for the Euler system can be obtained as the special case where $\varepsilon = 1$, i.e.

$$E = \int_{-\infty}^{\infty} \int_0^{\infty} \left(|u_\bullet|^2 + |\nabla_{\mathring{r},\mathring{z}} \psi_\bullet|^2 \right) \mathring{r}^3 \, d\mathring{r} \, d\mathring{z}. \tag{S369}$$

Proposition S9 (Energy scaling under the self-similar ansatz)**.** *Let $\tau = T - t$, and assume that*

$$u_\bullet(\mathring{r}, \mathring{z}, t) = \tau^{-1} u\left(\frac{\mathring{r}}{\tau^\lambda}, \frac{\mathring{z} - \mathring{z}(t)}{\tau^\lambda} \right), \qquad \psi_\bullet(\mathring{r}, \mathring{z}, t) = \tau^{\lambda-1} \psi\left(\frac{\mathring{r}}{\tau^\lambda}, \frac{\mathring{z} - \mathring{z}(t)}{\tau^\lambda} \right). \tag{S370}$$

Equivalently, introduce the rescaled variables

$$r = \frac{\mathring{r}}{\tau^\lambda}, \qquad z = \frac{\mathring{z} - \mathring{z}(t)}{\tau^\lambda}. \tag{S371}$$

Then

$$E_\varepsilon(t) = \tau^{5\lambda-2} \int_{-\infty}^{\infty} \int_0^{\infty} \left(|u|^2 + (2-\varepsilon) |\nabla_{r,z} \psi|^2 \right) r^3 \, dr \, dz. \tag{S372}$$

If the rescaled profile energy is finite and nonzero, then boundedness of $E_\varepsilon(t)$ as $t \to T^-$ requires $\lambda \geq 2/5$.

Proof (✓ Formalized and verified in Lean)

We have

$$\mathring{r} = r\tau^\lambda, \qquad \mathring{z} = z\tau^\lambda + \mathring{z}(t). \tag{S373}$$

Therefore

$$d\mathring{r} \, d\mathring{z} = \tau^{2\lambda} \, dr \, dz, \qquad \mathring{r}^3 = r^3 \tau^{3\lambda}. \tag{S374}$$

Hence the weighted measure transforms as

$$\mathring{r}^3 \, d\mathring{r} \, d\mathring{z} = r^3 \tau^{5\lambda} \, dr \, dz. \tag{S375}$$

Next, since $\mathring{z}(t)$ is independent of z, we have

$$\partial_{\mathring{r}} = \tau^{-\lambda} \partial_r, \qquad \partial_{\mathring{z}} = \tau^{-\lambda} \partial_z. \tag{S376}$$

Thus

$$\nabla_{\mathring{r},\mathring{z}} = \tau^{-\lambda} \nabla_{r,z}. \tag{S377}$$

Using $\psi_\bullet = \tau^{\lambda-1}\psi$, it follows that

$$\nabla_{\mathring{r},\mathring{z}} \psi_\bullet = \tau^{\lambda-1} \nabla_{\mathring{r},\mathring{z}} \psi = \tau^{\lambda-1} \tau^{-\lambda} \nabla_{r,z} \psi = \tau^{-1} \nabla_{r,z} \psi. \tag{S378}$$

Also, by the ansatz,

$$u_\bullet = \tau^{-1} u. \tag{S379}$$

Substituting these identities into the definition of E_ε gives

$$E_\varepsilon(t) = \int_{-\infty}^{\infty}\int_0^{\infty} \left(\left|\tau^{-1}u\right|^2 + (2-\varepsilon)\left|\tau^{-1}\nabla_{r,z}\psi\right|^2\right) r^3 \tau^{5\lambda}\, dr\, dz \tag{S380}$$

$$= \tau^{5\lambda-2}\int_{-\infty}^{\infty}\int_0^{\infty} \left(|u|^2 + (2-\varepsilon)|\nabla_{r,z}\psi|^2\right) r^3\, dr\, dz. \tag{S381}$$

This proves the claimed scaling relation.

Assume now that the rescaled profile energy

$$\int_{-\infty}^{\infty}\int_0^{\infty} \left(|u|^2 + (2-\varepsilon)|\nabla_{r,z}\psi|^2\right) r^3\, dr\, dz \tag{S382}$$

is finite and nonzero. Since $\tau \to 0^+$ as $t \to T^-$, boundedness of $E_\varepsilon(t)$ requires

$$5\lambda - 2 \geq 0. \tag{S383}$$

Therefore

$$\lambda \geq \frac{2}{5}. \tag{S384}$$

■

Appendix F. Scaling Symmetry and the Choice of Variables for Linear Stability

We consider the three-dimensional incompressible Euler equations

$$\partial_t u + (u \cdot \nabla)u + \nabla p = 0, \qquad \nabla \cdot u = 0, \tag{S385}$$

where $u = u(x,t) \in \mathbb{R}^3$ is the velocity field and $p = p(x,t) \in \mathbb{R}$ is the pressure.

Scaling of the velocity. *Let (u,p) be a solution of (S385). For any $\eta > 0$ and any $\alpha \in \mathbb{R}$, define*

$$u_\eta(x,t) := \eta^\alpha u(\eta x, \eta^{\alpha+1} t), \qquad p_\eta(x,t) := \eta^{2\alpha} p(\eta x, \eta^{\alpha+1} t). \tag{S386}$$

Then (u_η, p_η) is again a solution of (S385).

Proof (✓ Formalized and verified in Lean)

Set

$$X = \eta x, \qquad T = \eta^{\alpha+1} t. \tag{S387}$$

Then

$$u_\eta(x,t) = \eta^\alpha u(X,T), \qquad p_\eta(x,t) = \eta^{2\alpha} p(X,T). \tag{S388}$$

We first compute the time derivative. By the chain rule,

$$\partial_t u_\eta(x,t) = \partial_t\big(\eta^\alpha u(X,T)\big) = \eta^\alpha (\partial_t u)(X,T)\,\partial_t T. \tag{S389}$$

Since $\partial_t T = \eta^{\alpha+1}$, this gives

$$\partial_t u_\eta(x,t) = \eta^\alpha \eta^{\alpha+1} (\partial_t u)(X,T) = \eta^{2\alpha+1} (\partial_t u)(X,T). \tag{S390}$$

Next we compute the spatial derivatives. Since $X = \eta x$, each derivative with respect to x_j contributes one factor of η. More precisely,

$$\partial_{x_j} u_\eta(x,t) = \partial_{x_j}\big(\eta^\alpha u(X,T)\big) = \eta^\alpha \sum_{k=1}^{3} (\partial_{X_k} u)(X,T)\,\partial_{x_j} X_k. \tag{S391}$$

Because $\partial_{x_j} X_k = \eta \delta_{jk}$, we obtain

$$\partial_{x_j} u_\eta(x,t) = \eta^\alpha \eta (\partial_{X_j} u)(X,T). \tag{S392}$$

Therefore

$$\nabla u_\eta(x,t) = \eta^{\alpha+1} (\nabla u)(X,T). \tag{S393}$$

Using (S393), we can now compute the transport term:

$$\begin{aligned}
(u_\eta \cdot \nabla) u_\eta &= \big(\eta^\alpha u(X,T)\big) \cdot \nabla\big(\eta^\alpha u(X,T)\big) \\
&= \big(\eta^\alpha u(X,T)\big) \cdot \big(\eta^{\alpha+1} (\nabla u)(X,T)\big) \\
&= \eta^{2\alpha+1} \big((u \cdot \nabla) u\big)(X,T).
\end{aligned} \tag{S394}$$

We now compute the pressure gradient. Again by the chain rule,

$$\partial_{x_j} p_\eta(x,t) = \partial_{x_j}\big(\eta^{2\alpha} p(X,T)\big) = \eta^{2\alpha} \sum_{k=1}^{3} (\partial_{X_k} p)(X,T)\, \partial_{x_j} X_k. \tag{S395}$$

Since $\partial_{x_j} X_k = \eta \delta_{jk}$, it follows that

$$\partial_{x_j} p_\eta(x,t) = \eta^{2\alpha} \eta (\partial_{X_j} p)(X,T), \tag{S396}$$

and therefore

$$\nabla p_\eta(x,t) = \eta^{2\alpha+1} (\nabla p)(X,T). \tag{S397}$$

Combining (S390), (S394), and (S397), we obtain

$$\begin{aligned}\partial_t u_\eta + (u_\eta \cdot \nabla) u_\eta + \nabla p_\eta &= \eta^{2\alpha+1} (\partial_t u)(X,T) + \eta^{2\alpha+1} \big((u \cdot \nabla) u\big)(X,T) + \eta^{2\alpha+1} (\nabla p)(X,T) \\ &= \eta^{2\alpha+1} \Big[\partial_t u + (u \cdot \nabla) u + \nabla p\Big](X,T).\end{aligned} \tag{S398}$$

Since (u,p) solves (S385), the bracketed term vanishes. Hence

$$\partial_t u_\eta + (u_\eta \cdot \nabla) u_\eta + \nabla p_\eta = 0. \tag{S399}$$

It remains to verify the divergence-free condition. Using the same computation as above,

$$\begin{aligned}\nabla \cdot u_\eta(x,t) &= \sum_{j=1}^{3} \partial_{x_j}\big(\eta^\alpha u_j(X,T)\big) = \eta^\alpha \sum_{j=1}^{3}\sum_{k=1}^{3} (\partial_{X_k} u_j)(X,T)\, \partial_{x_j} X_k \\ &= \eta^\alpha \eta \sum_{j=1}^{3} (\partial_{X_j} u_j)(X,T) = \eta^{\alpha+1} (\nabla \cdot u)(X,T).\end{aligned} \tag{S400}$$

Since $\nabla \cdot u = 0$, it follows that $\nabla \cdot u_\eta = 0$ as well. Thus (u_η, p_η) satisfies (S385). ■

We now turn to the vorticity $\omega = \nabla \times u$ and the velocity gradient ∇u.

Scaling of the vorticity. *Let $\omega = \nabla \times u$. If u_η is defined by (S386), then*

$$\omega_\eta(x,t) := \nabla \times u_\eta(x,t) = \eta^{\alpha+1} \omega(\eta x, \eta^{\alpha+1} t). \tag{S401}$$

Proof (✓ Formalized and verified in Lean)

As before, we write

$$X = \eta x, \qquad T = \eta^{\alpha+1} t. \tag{S402}$$

Then

$$u_\eta(x,t) = \eta^\alpha u(X,T). \tag{S403}$$

We compute the curl componentwise. For example, the first component is

$$(\omega_\eta)_1 = \partial_{x_2}(u_\eta)_3 - \partial_{x_3}(u_\eta)_2 = \partial_{x_2}\big(\eta^\alpha u_3(X,T)\big) - \partial_{x_3}\big(\eta^\alpha u_2(X,T)\big). \tag{S404}$$

Using the chain rule,

$$\partial_{x_2}\big(\eta^\alpha u_3(X,T)\big) = \eta^\alpha \eta (\partial_{X_2} u_3)(X,T), \qquad \partial_{x_3}\big(\eta^\alpha u_2(X,T)\big) = \eta^\alpha \eta (\partial_{X_3} u_2)(X,T). \tag{S405}$$

Therefore

$$(\omega_\eta)_1 = \eta^{\alpha+1}\big(\partial_{X_2} u_3 - \partial_{X_3} u_2\big)(X,T) = \eta^{\alpha+1}\omega_1(X,T). \tag{S406}$$

The same argument gives

$$(\omega_\eta)_2 = \eta^{\alpha+1}\omega_2(X,T), \qquad (\omega_\eta)_3 = \eta^{\alpha+1}\omega_3(X,T). \tag{S407}$$

Hence

$$\omega_\eta(x,t) = \eta^{\alpha+1}\omega(X,T), \tag{S408}$$

that is,

$$\omega_\eta(x,t) = \eta^{\alpha+1}\omega(\eta x, \eta^{\alpha+1} t). \tag{S409}$$

■

Scaling of the velocity gradient. *If u_η is defined by (S386), then*

$$\nabla u_\eta(x,t) = \eta^{\alpha+1}(\nabla u)(\eta x, \eta^{\alpha+1} t). \tag{S410}$$

Proof (✓ Formalized and verified in Lean)

As before, let $X = \eta x$ and $T = \eta^{\alpha+1} t$. For each pair of indices i, j,

$$\partial_{x_j}(u_\eta)_i(x,t) = \partial_{x_j}\big(\eta^\alpha u_i(X,T)\big). \tag{S411}$$

Applying the chain rule gives

$$\partial_{x_j}(u_\eta)_i(x,t) = \eta^\alpha \sum_{k=1}^{3} (\partial_{X_k} u_i)(X,T)\, \partial_{x_j} X_k. \tag{S412}$$

Since $\partial_{x_j} X_k = \eta \delta_{jk}$, we obtain

$$\partial_{x_j}(u_\eta)_i(x,t) = \eta^\alpha \eta (\partial_{X_j} u_i)(X,T) = \eta^{\alpha+1}(\partial_{X_j} u_i)(X,T). \tag{S413}$$

Because this holds for every i and j, it follows that

$$\nabla u_\eta(x,t) = \eta^{\alpha+1}(\nabla u)(X,T), \tag{S414}$$

which is exactly (S410). ■

The variables u, ω, and ∇u do not all transform in the same way. Indeed, from (S386) and (S401),

$$u_\eta(x,t) = \eta^{\alpha} u(\eta x, \eta^{\alpha+1} t), \qquad \omega_\eta(x,t) = \eta^{\alpha+1}\omega(\eta x, \eta^{\alpha+1} t). \tag{S415}$$

Thus the full vorticity carries one extra factor of η relative to the full velocity. This is expected since $\omega = \nabla \times u$ contains one spatial derivative of u. By contrast, (S401) and (S410) show that

$$\omega_\eta(x,t) = \eta^{\alpha+1}\omega(\eta x, \eta^{\alpha+1} t), \qquad \nabla u_\eta(x,t) = \eta^{\alpha+1}(\nabla u)(\eta x, \eta^{\alpha+1} t). \tag{S416}$$

The preceding scaling laws concern the full three-dimensional velocity and vorticity fields. In the stability analysis, however, u and ω denote the reduced variables

$$u = \frac{u^\theta}{r}, \qquad \omega = \frac{\omega^\theta}{r}. \tag{S417}$$

Since division by r contributes one additional factor of η under the rescaling $(r,z) \mapsto (\eta r, \eta z)$, the reduced variables satisfy

$$u_\eta(r,z,t) = \eta^{\alpha+1} u(\eta r, \eta z, \eta^{\alpha+1} t), \qquad \omega_\eta(r,z,t) = \eta^{\alpha+2}\omega(\eta r, \eta z, \eta^{\alpha+1} t). \tag{S418}$$

Consequently,

$$\nabla u_\eta(r,z,t) = \eta^{\alpha+2}(\nabla u)(\eta r, \eta z, \eta^{\alpha+1} t), \tag{S419}$$

where ∇ acts on the reduced spatial variables (r,z).

Thus the reduced variables ω and ∇u transform in the same way. This is the reason why they are the natural pair of quantities in a scale-consistent linear stability analysis. Quantities with different scaling correspond to different differential orders, so comparing them directly mixes objects of different analytic strength.

Appendix G. Linearization and Modulation Equations

G.1 Setup

We now derive the linearized Euler perturbation equations, together with the modulation equations used to close the perturbation system. We use blue for perturbation quantities, red highlighting for nonlinear perturbation products. A subscript $_0$ denotes evaluation at $(r, z) = (0, 0)$.

✓ *All the derivations in Appendix G have been formalized and verified in Lean.*

We start from the dynamically rescaled equations

$$\partial_t u + \big(\lambda r + u^r\big)\partial_r u + \big(\lambda z + C(t) + u^z\big)\partial_z u = 2u\partial_z\psi + c_u(t)u, \tag{S420}$$

$$\partial_t \omega + \big(\lambda r + u^r\big)\partial_r \omega + \big(\lambda z + C(t) + u^z\big)\partial_z \omega = 2u\partial_z u + c_\omega(t)\omega, \tag{S421}$$

$$-\mathcal{E}\psi = \omega, \tag{S422}$$

where

$$u^r = -r\partial_z\psi, \tag{S423}$$

$$u^z = 2\psi + r\partial_r\psi, \tag{S424}$$

$$c_\omega = c_u - \lambda, \tag{S425}$$

and

$$\mathcal{E}f := \partial_r^2 f + \frac{3}{r}\partial_r f + \partial_z^2 f. \tag{S426}$$

We denote steady-state values using an overhead bar. Thus $\bar{u}, \bar{\omega}, \bar{\psi}, \bar{C}, \bar{c}_u, \bar{c}_\omega$ form an Euler steady state. The profile may be approximate, so the steady state equations are

$$\partial_t \bar{u} + (\lambda r + \bar{u}^r)\partial_r \bar{u} + (\lambda z + \bar{C} + \bar{u}^z)\partial_z \bar{u} = 2\bar{u}\partial_z\bar{\psi} + \bar{c}_u\bar{u} \; - \; \mathcal{R}_u^{\mathsf{raw}}, \tag{S427}$$

$$\partial_t \bar{\omega} + (\lambda r + \bar{u}^r)\partial_r \bar{\omega} + (\lambda z + \bar{C} + \bar{u}^z)\partial_z \bar{\omega} = 2\bar{u}\partial_z\bar{u} + \bar{c}_\omega\bar{\omega} \; - \; \mathcal{R}_\omega^{\mathsf{raw}}. \tag{S428}$$

We write the perturbation ansatz as

$$u = \bar{u} + \delta u, \qquad \omega = \bar{\omega} + \delta\omega, \qquad \psi = \bar{\psi} + \delta\psi, \tag{S429}$$

$$C = \bar{C} + \delta C, \qquad c_u = \bar{c}_u + c_{u,\mathrm{res}} + \delta c_u, \qquad c_\omega = \bar{c}_\omega + c_{u,\mathrm{res}} + \delta c_\omega. \tag{S430}$$

The induced meridional velocity perturbations are

$$u^r = \bar{u}^r + \delta u^r, \qquad u^z = \bar{u}^z + \delta u^z, \tag{S431}$$

with

$$\delta u^r = -r\partial_z\delta\psi, \qquad \delta u^z = 2\delta\psi + r\partial_r\delta\psi, \qquad \bar{c}_\omega = \bar{c}_u - \lambda, \qquad \delta c_\omega = \delta c_u. \tag{S432}$$

Before linearizing, write the radial and axial transport coefficients of the numerically obtained profile as

$$G_r(r,z) = \lambda r + u^r_{\mathsf{num}}(r,z), \qquad G_z(r,z) = \lambda z + C_{\mathsf{num}} + u^z_{\mathsf{num}}(r,z). \tag{S433}$$

We choose z_{shift} so that $(0, z_{\mathsf{shift}})$ becomes the analysis origin. For this point to be a meridional fixed point, we require

$$G_r(0, z_{\mathsf{shift}}) = 0, \qquad G_z(0, z_{\mathsf{shift}}) = 0. \tag{S434}$$

The radial condition is automatic because $u^r_{\mathsf{num}}(0, z) = 0$ on the axis. Using $u^z_{\mathsf{num}}(0, z) = 2\psi_{\mathsf{num}}(0, z)$, the axial condition $G_z(0, z_{\mathsf{shift}}) = 0$ becomes

$$C_{\mathsf{num}} + \lambda z_{\mathsf{shift}} + 2\psi_{\mathsf{num}}(0, z_{\mathsf{shift}}) = 0. \tag{S435}$$

We therefore impose this scalar condition and choose z_{shift} as a numerically certified root. We then recenter each numerical profile field by

$$\bar{f}(r, z) := f_{\mathsf{num}}(r, z + z_{\mathsf{shift}}). \tag{S436}$$

Denoting the constant term in the recentered axial transport by $\bar{C}$, we continue to write the reference transport coefficients as

$$G_r(r, z) = \lambda r + \bar{u}^r(r, z), \qquad G_z(r, z) = \lambda z + \bar{C} + \bar{u}^z(r, z). \tag{S437}$$

With this notation, the imposed root condition gives

$$G_r(0, 0) = 0, \qquad G_z(0, 0) = \bar{C} + 2\bar{\psi}_0 = 0. \tag{S438}$$

Having fixed the shift, set

$$\bar{u}_0 := \bar{u}(0, 0), \qquad \bar{u}_0 \neq 0. \tag{S439}$$

and certify a positive lower bound for $|\bar{u}_0|$. Thus the recentered numerical profile used in the analysis has a fixed point of the meridional flow at the origin.

The two scalar modulation conditions used below preserve this choice:

$$C(t) + u^z_0(t) = 0, \qquad u_0(t) = \bar{u}_0. \tag{S440}$$

Since $u^r_0(t) = 0$ identically, these conditions imply

$$\big(\lambda r + u^r\big)_0 = 0, \qquad \big(\lambda z + C + u^z\big)_0 = 0 \tag{S441}$$

for every rescaled time.

Equivalently, relative to the reference profile,

$$\delta C + 2\delta\psi_0 = 0, \qquad \delta u_0 = 0, \tag{S442}$$

and the time-differentiated normalization condition used in the second modulation equation is

$$\partial_t \delta u_0 = 0. \tag{S443}$$

G.2 Linearization of the Euler equations

G.2.1 Linearization of the first equation

We consider the u-equation

$$\partial_t u + (\lambda r + u^r)\,\partial_r u + (\lambda z + C + u^z)\,\partial_z u = 2u\,\partial_z\psi + c_u(t)u. \tag{S444}$$

Substituting (S429)–(S432) into (S444) gives

$$\begin{aligned}\partial_t(\bar u + \delta u) + \big(\lambda r + (\bar u^r + \delta u^r)\big)\,\partial_r(\bar u + \delta u) + \big(\lambda z + (\bar C + \delta C) + (\bar u^z + \delta u^z)\big)\,\partial_z(\bar u + \delta u)\\ = 2(\bar u + \delta u)\,\partial_z(\bar\psi + \delta\psi) + (\bar c_u + c_{u,\mathrm{res}} + \delta c_u)(\bar u + \delta u).\end{aligned} \tag{S445}$$

Since $\bar u, \bar\psi, \bar C, \bar c_u$ satisfy the steady-state equation, we have

$$\partial_t\bar u + (\lambda r + \bar u^r)\,\partial_r\bar u + (\lambda z + \bar C + \bar u^z)\,\partial_z\bar u = 2\bar u\,\partial_z\bar\psi + \bar c_u\,\bar u \;-\; \mathcal{R}_u^{\mathsf{raw}}. \tag{S446}$$

We subtract (S446) from (S445) and group the nonlinear perturbation products in red to obtain

$$\partial_t\delta u + (\lambda r + \bar u^r)\,\partial_r\delta u + (\lambda z + \bar C + \bar u^z)\,\partial_z\delta u + (\delta u^r)\,\partial_r\bar u + (\delta C + \delta u^z)\,\partial_z\bar u \tag{S447}$$

$$\begin{aligned}&= (2\,\partial_z\bar\psi + \bar c_u + c_{u,\mathrm{res}})\,\delta u + 2\,\bar u\,\partial_z\delta\psi + \delta c_u\,\bar u \;+\; c_{u,\mathrm{res}}\,\bar u \;+\; \mathcal{R}_u^{\mathsf{raw}}\\ &\quad - \Big(\delta u^r\Big)\,\partial_r\delta u - \Big(\delta C + \delta u^z\Big)\,\partial_z\delta u + 2\,\delta u\,\partial_z\delta\psi + \delta c_u\,\delta u\end{aligned} \tag{S448}$$

Using (S432), we rewrite (S447) as

$$\begin{aligned}\partial_t\delta u + (\lambda r + \bar u^r)\,\partial_r\delta u + (\lambda z + \bar C + \bar u^z)\,\partial_z\delta u - r\,\partial_z\delta\psi\,\partial_r\bar u + \big(\delta C + (2\delta\psi + r\partial_r\delta\psi)\big)\,\partial_z\bar u\\ = (2\,\partial_z\bar\psi + \bar c_u + c_{u,\mathrm{res}})\,\delta u + 2\,\bar u\,\partial_z\delta\psi + \delta c_u\,\bar u \;+\; c_{u,\mathrm{res}}\,\bar u \;+\; \mathcal{R}_u^{\mathsf{raw}}\\ + r\,\partial_z\delta\psi\,\partial_r\delta u - \Big(\delta C + (2\delta\psi + r\partial_r\delta\psi)\Big)\,\partial_z\delta u + 2\,\delta u\,\partial_z\delta\psi + \delta c_u\,\delta u\,.\end{aligned} \tag{S449}$$

Collecting and reorganizing terms gives

$$\begin{aligned}\partial_t\delta u = -(\lambda r + \bar u^r)\,\partial_r\delta u - (\lambda z + \bar C + \bar u^z)\,\partial_z\delta u + (2\,\partial_z\bar\psi + \bar c_u + c_{u,\mathrm{res}})\,\delta u + 2\,\bar u\,\partial_z\delta\psi + \bar u\,\delta c_u\\ + r\,(\partial_z\delta\psi)\,\partial_r\bar u - \big(2\delta\psi + r\partial_r\delta\psi\big)\,\partial_z\bar u - (\partial_z\bar u)\,\delta C \;+\; c_{u,\mathrm{res}}\,\bar u \;+\; \mathcal{R}_u^{\mathsf{raw}}\\ + r\,\partial_z\delta\psi\,\partial_r\delta u - \Big(\delta C + (2\delta\psi + r\partial_r\delta\psi)\Big)\,\partial_z\delta u + 2\,\delta u\,\partial_z\delta\psi + \delta c_u\,\delta u\,.\end{aligned} \tag{S450}$$

Define the residual amplitude offset by

$$c_{u,\mathrm{res}} := -\frac{\mathcal{R}_u^{\mathsf{raw}}(0)}{\bar u_0} \tag{S451}$$

under the standing nondegeneracy condition $\bar u_0 \neq 0$.

We absorb the residual amplitude offset into the reference amplitude rate and use the modulation coefficients directly:

$$C = \bar{C} + \delta C, \qquad c_u = \bar{c}_u + c_{u,\mathrm{res}} + \delta c_u, \qquad c_\omega = \bar{c}_\omega + c_{u,\mathrm{res}} + \delta c_u. \tag{S452}$$

The residuals associated with this decomposition are

$$\mathcal{R}_u := \mathcal{R}_u^{\mathsf{raw}} + c_{u,\mathrm{res}}\,\bar{u}, \tag{S453}$$

$$\mathcal{R}_\omega := \mathcal{R}_\omega^{\mathsf{raw}} + c_{u,\mathrm{res}}\,\bar{\omega}. \tag{S454}$$

The linear operator is

$$\begin{aligned}\mathcal{L}_u(\delta) := \big(2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}}\big)\,\delta u - \big(\lambda r + \bar{u}^r\big)\,\partial_r\delta u - \big(\lambda z + \bar{C} + \bar{u}^z\big)\,\partial_z\delta u + \big(2\bar{u} + r\,\partial_r\bar{u}\big)\partial_z\delta\psi \\ - \big(r\,\partial_z\bar{u}\big)\partial_r\delta\psi - \big(2\partial_z\bar{u}\big)\delta\psi - (\partial_z\bar{u})\,\delta C + \bar{u}\,\delta c_u.\end{aligned} \tag{S455}$$

The nonlinear operator is

$$\mathcal{L}_\omega(\delta) = r\,\partial_z\delta\psi\,\partial_r\delta u - (\,\delta C + (2\delta\psi + r\partial_r\delta\psi))\,\partial_z\delta u + 2\,\delta u\,\partial_z\delta\psi + \delta c_u\,\delta u. \tag{S456}$$

G.2.2 Linearization of the second equation

We consider the ω-equation

$$\partial_t\omega + (\lambda r + u^r)\,\partial_r\omega + (\lambda z + C + u^z)\,\partial_z\omega = 2u\,\partial_z u + c_\omega(t)\,\omega. \tag{S457}$$

Substituting (S429)–(S432) into (S457) gives

$$\begin{aligned}\partial_t(\bar{\omega} + \delta\omega) + \big(\lambda r + (\bar{u}^r + \delta u^r)\big)\,\partial_r(\bar{\omega} + \delta\omega) + \big(\lambda z + (\bar{C} + \delta C) + (\bar{u}^z + \delta u^z)\big)\,\partial_z(\bar{\omega} + \delta\omega) \\ = 2(\bar{u} + \delta u)\,\partial_z(\bar{u} + \delta u) + (\bar{c}_\omega + c_{u,\mathrm{res}} + \delta c_\omega)(\bar{\omega} + \delta\omega).\end{aligned} \tag{S458}$$

Now, $\bar{u}, \bar{\omega}, \bar{C}, \bar{c}_\omega$ satisfy the steady-state equation, so

$$\partial_t\bar{\omega} + (\lambda r + \bar{u}^r)\,\partial_r\bar{\omega} + (\lambda z + \bar{C} + \bar{u}^z)\,\partial_z\bar{\omega} = 2\bar{u}\,\partial_z\bar{u} + \bar{c}_\omega\,\bar{\omega} \;-\; \mathcal{R}_\omega^{\mathsf{raw}}. \tag{S459}$$

We subtract (S459) from (S458) and group the nonlinear perturbation products in red to obtain

$$\begin{aligned}\partial_t\delta\omega + (\lambda r{+}\bar{u}^r)\,\partial_r\delta\omega + (\lambda z + \bar{C} + \bar{u}^z)\,\partial_z\delta\omega + (\delta u^r)\,\partial_r\bar{\omega} + (\delta C + \delta u^z)\,\partial_z\bar{\omega} \\ = 2\bar{u}\,\partial_z\delta u + 2(\partial_z\bar{u})\,\delta u + (\bar{c}_\omega + c_{u,\mathrm{res}})\,\delta\omega + \bar{\omega}\,\delta c_\omega \;+\; c_{u,\mathrm{res}}\,\bar{\omega} \;+\; \mathcal{R}_\omega^{\mathsf{raw}} \\ + 2\,\delta u\,\partial_z\delta u - \Big(\delta u^r\Big)\,\partial_r\delta\omega - \Big(\delta C + \delta u^z\Big)\,\partial_z\delta\omega + \delta c_u\,\delta\omega\,.\end{aligned} \tag{S460}$$

We collect and reorganize terms

$$\begin{aligned}\partial_t\delta\omega = -(\lambda r + \bar{u}^r)\,\partial_r\delta\omega - (\lambda z + \bar{C} + \bar{u}^z)\,\partial_z\delta\omega + (\bar{c}_\omega + c_{u,\mathrm{res}})\,\delta\omega + 2(\partial_z\bar{u})\,\delta u + 2\bar{u}\,\partial_z\delta u + \bar{\omega}\,\delta c_\omega \\ + r\,(\partial_z\delta\psi)\,\partial_r\bar{\omega} - \big(2\delta\psi + r\partial_r\delta\psi\big)\,\partial_z\bar{\omega} - (\partial_z\bar{\omega})\,\delta C \;+\; c_{u,\mathrm{res}}\,\bar{\omega} \;+\; \mathcal{R}_\omega^{\mathsf{raw}} \\ -\Big(\delta u^r\Big)\,\partial_r\delta\omega - \Big(\delta C + \delta u^z\Big)\,\partial_z\delta\omega + 2\,\delta u\,\partial_z\delta u + \delta c_u\,\delta\omega\,.\end{aligned} \tag{S461}$$

Then, substituting $\delta c_\omega = \delta c_u$, we obtain

$$\begin{aligned}\partial_t \delta\omega = &-(\lambda r + \bar{u}^r)\,\partial_r \delta\omega - (\lambda z + \bar{C} + \bar{u}^z)\,\partial_z \delta\omega + (\bar{c}_\omega + c_{u,\mathrm{res}})\,\delta\omega + 2(\partial_z \bar{u})\,\delta u + 2\bar{u}\,\partial_z \delta u \\ &+ \bar{\omega}\,(\delta c_u) + r\,(\partial_z \delta\psi)\,\partial_r \bar{\omega} - \big(2\delta\psi + r\partial_r \delta\psi\big)\,\partial_z \bar{\omega} - (\partial_z \bar{\omega})\,\delta C \;+\; c_{u,\mathrm{res}}\,\bar{\omega} \;+\; \mathcal{R}_\omega^{\mathsf{raw}} \\ &-\Big(\delta u^r\Big)\,\partial_r \delta\omega - \Big(\delta C + \delta u^z\Big)\,\partial_z \delta\omega + 2\,\delta u\,\partial_z \delta u + \delta c_u\,\delta\omega\,. \end{aligned} \tag{S462}$$

The linear operator is

$$\begin{aligned}\mathcal{L}_\omega(\delta) := &(\bar{c}_u + c_{u,\mathrm{res}} - \lambda)\,\delta\omega - \big(\lambda r + \bar{u}^r\big)\partial_r \delta\omega - \big(\lambda z + \bar{C} + \bar{u}^z\big)\,\partial_z \delta\omega + 2(\partial_z \bar{u})\,\delta u + 2\bar{u}\,\partial_z \delta u \\ &- 2(\partial_z \bar{\omega})\,\delta\psi + r\,(\partial_r \bar{\omega})\,\partial_z \delta\psi - r\,(\partial_z \bar{\omega})\,\partial_r \delta\psi - (\partial_z \bar{\omega})\;\delta C + \bar{\omega}\;\delta c_u. \end{aligned} \tag{S463}$$

The nonlinear operator is

$$\mathcal{N}_\omega(\delta) = \; -\delta u^r \partial_r \delta\omega - (\;\delta C + \delta u^z)\,\partial_z \delta\omega + 2\,\delta u\,\partial_z \delta u + \;\delta c_u\,\delta\omega. \tag{S464}$$

G.2.3 Linearization of the elliptic equation

We consider

$$-\Big[\partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2\Big]\psi = \omega. \tag{S465}$$

Then

$$-\Big[\partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2\Big]\big(\bar{\psi} + \delta\psi\big) = \bar{\omega} + \delta\omega. \tag{S466}$$

Since $(\bar{\psi}, \bar{\omega})$ form a steady-state, they satisfy

$$-\Big[\partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2\Big]\bar{\psi} = \bar{\omega}. \tag{S467}$$

There is no elliptic residual $\mathcal{R}_\psi$ here because the spline representation of $\bar{\omega}$ is obtained by applying the elliptic operator to $\bar{\psi}$ exactly, as discussed in Section 2.3.

We subtract (S467) from (S466) and obtain only linear terms:

$$-\Big[\partial_r^2 + \frac{3}{r}\partial_r + \partial_z^2\Big]\delta\psi = \delta\omega. \tag{S468}$$

G.2.4 z-differentiated form of the linearized first equation

Taking one z-derivative of the full linearized first equation, we use the notation

$$\mathcal{L}_z(\delta) := \partial_z \mathcal{L}_u(\delta), \tag{S469}$$
$$\mathcal{N}_z(\delta) := \partial_z \mathcal{L}_\omega(\delta). \tag{S470}$$

The main product-rule contributions are

$$\partial_z\Big[-(\lambda r + \bar{u}^r)\partial_r \delta u\Big] = -(\lambda r + \bar{u}^r)\partial_{rz}\delta u - (\partial_z \bar{u}^r)\partial_r \delta u, \tag{S471}$$
$$\partial_z\Big[-(\lambda z + \bar{C} + \bar{u}^z)\partial_z \delta u\Big] = -(\lambda z + \bar{C} + \bar{u}^z)\partial_{zz}\delta u - (\lambda + \partial_z \bar{u}^z)\partial_z \delta u, \tag{S472}$$
$$\partial_z\Big[(2\partial_z\bar{\psi} + \bar{c}_u)\delta u\Big] = (2\partial_z\bar{\psi} + \bar{c}_u)\partial_z \delta u + 2\partial_{zz}\bar{\psi}\,\delta u. \tag{S473}$$

These identities, together with term-by-term differentiation of the $\delta\psi$, δC, and δc_u terms, give the following explicit formulas.

The linear part is

$$\begin{aligned}\mathcal{L}_z(\delta) = \; & -(\lambda r + \bar{u}^r)\,\partial_{rz}\delta u - (\lambda z + \bar{C} + \bar{u}^z)\,\partial_{zz}\delta u \\ & + \big(2\partial_z\bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}} - \lambda - \partial_z \bar{u}^z\big)\partial_z \delta u - (\partial_z \bar{u}^r)\,\partial_r \delta u + 2\partial_{zz}\bar{\psi}\,\delta u \\ & + (2\bar{u} + r\partial_r \bar{u})\,\partial_{zz}\delta\psi + \big(r\partial_{rz}\bar{u}\big)\partial_z \delta\psi \\ & - r(\partial_z \bar{u})\,\partial_{rz}\delta\psi - r(\partial_{zz}\bar{u})\,\partial_r \delta\psi - 2\partial_{zz}\bar{u}\,\delta\psi \\ & - (\partial_{zz}\bar{u})\;\delta C + (\partial_z \bar{u})\;\delta c_u.\end{aligned} \tag{S474}$$

The nonlinear operator is

$$\begin{aligned}\mathcal{N}_z(\delta) = \; & r(\partial_{zz}\delta\psi)\,\partial_r \delta u + r(\partial_z \delta\psi)\,\partial_{rz}\delta u \\ & - (2\partial_z \delta\psi + r\partial_{rz}\delta\psi)\,\partial_z \delta u - (\;\delta C + (2\delta\psi + r\partial_r \delta\psi))\,\partial_{zz}\delta u \\ & + 2(\partial_z \delta u)(\partial_z \delta\psi) + 2\delta u\,\partial_{zz}\delta\psi + \;\delta c_u\,\partial_z \delta u.\end{aligned} \tag{S475}$$

G.2.5 r-differentiated form of the linearized first equation

We now differentiate the full linearized first equation with respect to r. We write

$$\mathcal{L}_r(\delta) := \partial_r \mathcal{L}_u(\delta), \tag{S476}$$

$$\mathcal{N}_r(\delta) := \partial_r \mathcal{L}_\omega(\delta), . \tag{S477}$$

The main product-rule contributions are

$$\partial_r\Big[-(\lambda r + \bar{u}^r)\partial_r \delta u\Big] = -(\lambda r + \bar{u}^r)\partial_{rr}\delta u - (\lambda + \partial_r \bar{u}^r)\partial_r \delta u, \tag{S478}$$

$$\partial_r\Big[-(\lambda z + \bar{C} + \bar{u}^z)\partial_z \delta u\Big] = -(\lambda z + \bar{C} + \bar{u}^z)\partial_{zr}\delta u - (\partial_r \bar{u}^z)\partial_z \delta u, \tag{S479}$$

$$\partial_r\Big[(2\partial_z \bar{\psi} + \bar{c}_u)\delta u\Big] = (2\partial_z \bar{\psi} + \bar{c}_u)\partial_r \delta u + 2\partial_{rz}\bar{\psi}\,\delta u. \tag{S480}$$

These identities, together with term-by-term differentiation of the $\delta\psi$, δC, and δc_u terms, give the following explicit formulas.

The linear part is

$$\begin{aligned}\mathcal{L}_r(\delta) = &-(\lambda r + \bar{u}^r)\,\partial_{rr}\delta u - (\lambda z + \bar{C} + \bar{u}^z)\,\partial_{zr}\delta u \\ &+ \big(2\partial_z \bar{\psi} + \bar{c}_u + c_{u,\mathrm{res}} - \lambda - \partial_r \bar{u}^r\big)\partial_r \delta u - (\partial_r \bar{u}^z)\,\partial_z \delta u + 2\partial_{rz}\bar{\psi}\,\delta u \\ &+ (2\bar{u} + r\partial_r \bar{u})\,\partial_{zr}\delta\psi + \big(3\partial_r \bar{u} + r\partial_{rr}\bar{u}\big)\partial_z \delta\psi \\ &- r(\partial_z \bar{u})\,\partial_{rr}\delta\psi - \big(3\partial_z \bar{u} + r\partial_{rz}\bar{u}\big)\partial_r \delta\psi - 2\partial_{rz}\bar{u}\,\delta\psi \\ &- (\partial_{rz}\bar{u})\;\delta C + (\partial_r \bar{u})\;\delta c_u.\end{aligned} \tag{S481}$$

The nonlinear part is

$$\begin{aligned}\mathcal{N}_r(\delta) = &\ (\partial_z \delta\psi + r\partial_{rz}\delta\psi)\,\partial_r \delta u + r\partial_z \delta\psi\,\partial_{rr}\delta u \\ &- (3\partial_r \delta\psi + r\partial_{rr}\delta\psi)\,\partial_z \delta u - (\;\delta C + (2\delta\psi + r\partial_r \delta\psi))\,\partial_{zr}\delta u \\ &+ 2(\partial_r \delta u)(\partial_z \delta\psi) + 2\delta u\,\partial_{zr}\delta\psi + \;\delta c_u\,\partial_r \delta u.\end{aligned} \tag{S482}$$

G.3 Modulation equations

The modulation conditions determine δC and δc_u.

G.3.1 Modulation equation for δC

The first normalization condition is

$$C + u_0^z = 0. \tag{S483}$$

Since $u_0^z = 2\psi_0$, substituting $C = \bar{C} + \delta C$ and $\psi = \bar{\psi} + \delta\psi$ gives

$$0 = (\bar{C} + 2\bar{\psi}_0) + (\delta C + 2\delta\psi_0). \tag{S484}$$

The reference stagnation condition $\bar{C} + 2\bar{\psi}_0 = 0$ therefore yields $\delta C + 2\delta\psi_0 = 0$, and thus

$$\delta C = -2\delta\psi_0. \tag{S485}$$

G.3.2 Modulation equation for δc_u

The second normalization condition fixes $u_0(t) = \bar{u}_0$, hence $\partial_t u_0 = 0$.

Evaluating the u-equation at the origin gives

$$0 = -(\lambda r + u^r)_0(\partial_r u)_0 - (\lambda z + C + u^z)_0(\partial_z u)_0 + 2u_0(\partial_z\psi)_0 + c_u u_0. \tag{S486}$$

Since $u_0^r = 0$ and $C + u_0^z = 0$, both transport terms vanish.

Using $u_0 = \bar{u}_0$, the origin equation becomes

$$0 = 2\bar{u}_0(\partial_z\psi)_0 + c_u\bar{u}_0. \tag{S487}$$

We introduce the re-centered residual

$$\mathcal{R}_u := \mathcal{R}_u^{\mathsf{raw}} + c_{u,\mathrm{res}}\bar{u}, \qquad c_{u,\mathrm{res}} = -\frac{\mathcal{R}_u^{\mathsf{raw}}(0)}{\bar{u}_0}. \tag{S488}$$

By construction,

$$\mathcal{R}_u(0) = \mathcal{R}_u^{\mathsf{raw}}(0) + c_{u,\mathrm{res}}\bar{u}_0 = 0. \tag{S489}$$

Substituting $\psi = \bar{\psi} + \delta\psi$ and $c_u = \bar{c}_u + c_{u,\mathrm{res}} + \delta c_u$ separates the zero reference contribution from the perturbative terms:

$$0 = \underbrace{\big(2(\partial_z\bar{\psi})_0 + \bar{c}_u + c_{u,\mathrm{res}}\big)\bar{u}_0}_{\mathcal{R}_u(0)=0} + 2\bar{u}_0(\partial_z\delta\psi)_0 + \bar{u}_0\delta c_u. \tag{S490}$$

The underbraced term vanishes by (S489), leaving

$$\bar{u}_0\delta c_u = -2\bar{u}_0(\partial_z\delta\psi)_0. \tag{S491}$$

Since $\bar{u}_0 \neq 0$,

$$\delta c_u = -2(\partial_z\delta\psi)_0. \tag{S492}$$

Appendix H. Additional Details for the Linear Damping Estimate

H.1 Additional details for the low-order linear estimate

Here we record the complete low-order operator decomposition and identifies which terms are retained in the signed matrix and which are kept outside it in Section 4. The separate mechanisms used to estimate the terms outside the matrix are developed in our companion stability paper [18].

H.1.1 Complete low-order operator decomposition

The low-order linear pairing is

$$\mathcal{Q}_{\mathcal{L}}^{\mathrm{low}}(\delta) := \langle \mathcal{L}_\omega(\delta), \delta v_\omega \rangle_{\Phi_\omega} + \langle \mathcal{L}_r(\delta), \delta v_r \rangle_{\Phi_r} + \langle \mathcal{L}_z(\delta), \delta v_z \rangle_{\Phi_z}. \tag{S493}$$

Recall that $\mathcal{Q}_{\mathcal{L}}^{\mathrm{low}}$ also contains the linear terms obtained when the residual amplitude offset from the modulation equations multiplies a perturbation. Since δC and δc_u are exactly linear and therefore enter the linear operator only.

For each $i \in \{\omega, r, z\}$, separate the transport part and non-transport terms as

$$\mathcal{L}_i(\delta) = -G_r \partial_r \delta v_i - G_z \partial_z \delta v_i + \sum_m B_{i,m}(\delta). \tag{S494}$$

The complete list of non-transport terms is

$$\begin{aligned}
B_{\omega,1} &= (\bar{c}_u - \lambda + c_{u,\mathrm{res}})\delta v_\omega, & B_{\omega,2} &= 2\partial_z \bar{u}\, \delta u, \\
B_{\omega,3} &= 2\bar{u}\, \delta v_z, & B_{\omega,4} &= -2\partial_z \bar{\omega}\, \delta\psi, \\
B_{\omega,5} &= r\partial_r \bar{\omega}\, \partial_z \delta\psi, & B_{\omega,6} &= -r\partial_z \bar{\omega}\, \partial_r \delta\psi, \\
B_{\omega,7} &= -\partial_z \bar{\omega}\, \delta C, & B_{\omega,8} &= \bar{\omega}\, \delta c_u,
\end{aligned} \tag{S495}$$

$$\begin{aligned}
B_{r,1} &= (2\partial_z \bar{\psi} + \bar{c}_u - \lambda - \partial_r \bar{u}^r + c_{u,\mathrm{res}})\delta v_r, & B_{r,2} &= -\partial_r \bar{u}^z\, \delta v_z, \\
B_{r,3} &= 2\partial_{rz} \bar{\psi}\, \delta u, & B_{r,4} &= (2\bar{u} + r\partial_r \bar{u})\partial_{rz} \delta\psi, \\
B_{r,5} &= (3\partial_r \bar{u} + r\partial_{rr} \bar{u})\partial_z \delta\psi, & B_{r,6} &= -r\partial_z \bar{u}\, \partial_{rr} \delta\psi, \\
B_{r,7} &= -(3\partial_z \bar{u} + r\partial_{rz} \bar{u})\partial_r \delta\psi, & B_{r,8} &= -2\partial_{rz} \bar{u}\, \delta\psi, \\
B_{r,9} &= -\partial_{rz} \bar{u}\, \delta C, & B_{r,10} &= \partial_r \bar{u}\, \delta c_u,
\end{aligned} \tag{S496}$$

$$\begin{aligned}
B_{z,1} &= -\partial_z \bar{u}^r\, \delta v_r, & B_{z,2} &= (2\partial_z \bar{\psi} + \bar{c}_u - \lambda - \partial_z \bar{u}^z + c_{u,\mathrm{res}})\delta v_z, \\
B_{z,3} &= 2\partial_{zz} \bar{\psi}\, \delta u, & B_{z,4} &= (2\bar{u} + r\partial_r \bar{u})\partial_{zz} \delta\psi, \\
B_{z,5} &= (r\partial_{rz} \bar{u})\partial_z \delta\psi, & B_{z,6} &= -r\partial_z \bar{u}\, \partial_{rz} \delta\psi, \\
B_{z,7} &= -r\partial_{zz} \bar{u}\, \partial_r \delta\psi, & B_{z,8} &= -2\partial_{zz} \bar{u}\, \delta\psi, \\
B_{z,9} &= -\partial_{zz} \bar{u}\, \delta C, & B_{z,10} &= \partial_z \bar{u}\, \delta c_u.
\end{aligned} \tag{S497}$$

Among the non-transport terms above, the low-order signed matrix retains

$$\mathcal{S}_{\mathcal{L}} = \{B_{\omega,1}, B_{r,1}, B_{r,2}, B_{z,1}, B_{z,2}\}. \tag{S498}$$

The three diagonal labels contribute to the diagonal entries, while $B_{r,2}$ and $B_{z,1}$ contribute to the symmetric r–z entry. All other labels in the lists above are kept outside the signed matrix.

The terms outside the signed matrix are precisely the terms in the complete lists above that are not contained in $\mathcal{S}_{\mathcal{L}}$. We denote by $A^0_{\mathcal{L}}$ and $A^k_{\mathcal{L}}$ the aggregate low-order and top-order losses obtained after these terms are estimated by the separate mechanisms developed in the stability argument and presented in our companion stability paper [18].

H.1.2 Alternative weight continuations outside the computational box

Section 4.3.2 uses the tail continuation employed in the numerical construction. For completeness, other radial tail models can be analyzed from the same scalar condition. Suppose

$$\Phi_i(r,z) = \phi_i(\rho), \qquad \rho = r^2 + z^2. \tag{S499}$$

Then

$$\partial_r \log \Phi_i = 2r\frac{d}{d\rho}\log\phi_i, \qquad \partial_z \log \Phi_i = 2z\frac{d}{d\rho}\log\phi_i, \tag{S500}$$

and the left side of (S216) becomes

$$\left(1 + \frac{2z}{\rho}\bar{C}\right)\left(2\rho\frac{d}{d\rho}\log\phi_i(\rho)\right). \tag{S501}$$

Since $\sqrt{\rho} \geq R_{\text{box}}$ and $|z| \leq \sqrt{\rho}$ on the tail,

$$1 - \frac{2|\bar{C}|}{R_{\text{box}}} \leq 1 + \frac{2z}{\rho}\bar{C} \leq 1 + \frac{2|\bar{C}|}{R_{\text{box}}}. \tag{S502}$$

The convenient sufficient assumptions

$$c_{u,\text{res}} < 1, \qquad R_{\text{box}} > 2|\bar{C}| \tag{S503}$$

leave a positive zeroth-order margin and keep the correction factor positive throughout the tail.

The main asymptotic guide is

$$\limsup_{\rho\to\infty} 2\rho\frac{d}{d\rho}\log\phi_i(\rho) < 4(1 - c_{u,\text{res}}). \tag{S504}$$

We consider different continuation families below:

- **Exponential decay.** If
$$\phi_i(\rho) = C_{\text{floor}} + a\exp(-b\rho^p), \qquad a,b,p > 0, \tag{S505}$$
then $d(\log\phi_i)/d\rho \leq 0$. Under (S503), the left side of (S216) is nonpositive. Any $0 < \gamma_{\text{tail}} < 1 - c_{u,\text{res}}$ therefore works.

- **Algebraic decay.** If
$$\phi_i(\rho) = C_{\text{floor}} + a\rho^{-p}, \qquad a,p > 0, \tag{S506}$$
In this case,
$$2\rho\frac{d}{d\rho}\log\phi_i(\rho) = -2p\frac{a\rho^{-p}}{C_{\text{floor}} + a\rho^{-p}} \leq 0.$$
Under (S503), the factor $1 + 2z\bar{C}/\rho$ in (S501) is positive throughout the tail. Hence the entire expression in (S501) is nonpositive, so the scalar condition (S216) holds for any $0 < \gamma_{\text{tail}} < 1 - c_{u,\text{res}}$.

- **Algebraic growth.** If

$$\phi_i(\rho) = C_{\text{floor}} + a\rho^p, \qquad a, p > 0, \tag{S507}$$

then

$$2\rho \frac{d}{d\rho} \log \phi_i(\rho) = 2p \frac{a\rho^p}{C_{\text{floor}} + a\rho^p} \le 2p. \tag{S508}$$

A sufficient finite-box condition is

$$p\left(1 + \frac{2|\bar{C}|}{R_{\text{box}}}\right) < 2(1 - c_{u,\text{res}}). \tag{S509}$$

- **Positive exponential growth.** If

$$\phi_i(\rho) = C_{\text{floor}} + a\exp(b\rho^p), \qquad a, b, p > 0, \tag{S510}$$

then

$$2\rho \frac{d}{d\rho} \log \phi_i(\rho) = 2bp\rho^p \frac{a\exp(b\rho^p)}{C_{\text{floor}} + a\exp(b\rho^p)}, \tag{S511}$$

which is unbounded as $\rho \to \infty$. Hence the scalar tail certificate cannot hold.

The four continuation families above are summarized in the following table, separating the asymptotic rate restriction from the finite-tail condition used in the certificate.

Summary of radial model rates.

Rate	**Model** $\phi_i(\rho)$	**Asymptotic condition**	**Finite-tail condition**
Exponential decay	$C_{\text{floor}} + a\exp(-b\rho^p)$	$a, b, p > 0$, with $c_{u,\text{res}} < 1$	(S503)
Algebraic decay	$C_{\text{floor}} + a\rho^{-p}$	$a, p > 0$, with $c_{u,\text{res}} < 1$	(S503)
Algebraic growth	$C_{\text{floor}} + a\rho^p$	$p < 2(1 - c_{u,\text{res}})$	(S509)
Exponential growth	$C_{\text{floor}} + a\exp(b\rho^p)$	Not possible	Not possible

For the r, z weights, compatibility with the low-order weighted energy also requires $p \ge 1$. Thus an algebraically growing tail is compatible when one can choose

$$1 \le p < \frac{2(1 - c_{u,\text{res}})}{1 + 2|\bar{C}|/R_{\text{box}}}. \tag{S512}$$

This interval is nonempty if and only if

$$1 + \frac{2|\bar{C}|}{R_{\text{box}}} < 2(1 - c_{u,\text{res}}). \tag{S513}$$

H.1.3 Local concentration estimates

To apply the localized certificate, it is enough to prove (S226) on each piece B_ℓ. Since every B_ℓ is bounded and stays away from the origin, and the weights are singular only at the origin, $\sup_{B_\ell} \Phi_i < \infty$ for every $i \in \{\omega, r, z\}$. The weighted energy can thus be controlled by ordinary local estimates for the components δv_i.

Pointwise estimate. One option is to use pointwise estimates. If

$$\|\delta v_i\|_{L^\infty(B_\ell)} \leq P^0_{i,\ell}\mathfrak{E}_0 + P^k_{i,\ell}\mathfrak{H}_k, \qquad i \in \{\omega, r, z\}, \tag{S514}$$

then $|d\xi|^2 = \sum_i \Phi_i |\delta v_i|^2$ and the identity $(a+b)^2 \leq 2a^2 + 2b^2$ give (S226), i.e.,

$$\int_{B_\ell} |d\xi|^2 \, dr\, dz \leq m^{0,\mathrm{loc},\ell}_{\mathcal{L}} \mathfrak{E}_0^2 + m^{k,\mathrm{loc},\ell}_{\mathcal{L}} \mathfrak{H}_k^2, \qquad \ell \in \mathcal{I}^{\mathrm{loc}}_{\mathcal{L}}, \tag{S515}$$

with

$$m^{0,\mathrm{loc},\ell}_{\mathcal{L}} := 2|B_\ell| \sum_{i \in \{\omega,r,z\}} \Big(\sup_{B_\ell} \Phi_i\Big)(P^0_{i,\ell})^2, \tag{S516}$$

$$m^{k,\mathrm{loc},\ell}_{\mathcal{L}} := 2|B_\ell| \sum_{i \in \{\omega,r,z\}} \Big(\sup_{B_\ell} \Phi_i\Big)(P^k_{i,\ell})^2. \tag{S517}$$

Thus shrinking B_ℓ improves the constants, provided that the weight suprema and the pointwise constants remain uniformly controlled.

Local L^2 estimate. Alternatively, suppose one has a local L^2 estimate of the form

$$\int_{B_\ell} |\delta v_i|^2 \, dr\, dz \leq Q^0_{i,\ell}\mathfrak{E}_0^2 + Q^k_{i,\ell}\mathfrak{H}_k^2, \qquad i \in \{\omega, r, z\}, \tag{S518}$$

then (S226) holds with

$$m^{0,\mathrm{loc},\ell}_{\mathcal{L}} := \sum_{i \in \{\omega,r,z\}} \Big(\sup_{B_\ell} \Phi_i\Big) Q^0_{i,\ell}, \qquad m^{k,\mathrm{loc},\ell}_{\mathcal{L}} := \sum_{i \in \{\omega,r,z\}} \Big(\sup_{B_\ell} \Phi_i\Big) Q^k_{i,\ell}. \tag{S519}$$

H.2 Additional details for the high-order linear estimate

Remainder structure. Section 4.5.2 defines the complete top-order signed matrix, including the transport contribution, every first transport commutator, and the local top-order component couplings. The high-order remainder begins exactly where that complete signed matrix ends.

To see the transport split directly, write $\alpha = (q, k-q)$. In the Leibniz expansion of

$$D^\alpha \left(G_r \partial_r f + G_z \partial_z f\right), \tag{S520}$$

the terms in which exactly one derivative lands on G_r or G_z are the first transport commutators retained in the matrix. Equivalently, these are the $|\beta| = 1$ terms when β denotes the multi-index of derivatives falling on the transport coefficient. Each leaves exactly k derivatives on f. Terms with $|\beta| \geq 2$ contain at most $k-1$ derivatives of the perturbation and are therefore estimated as remainders.

The same principle applies to the local top-order coefficients. The undifferentiated coefficient multiplying a top derivative is retained in the signed matrix. Once a derivative lands on a spatially varying coefficient, the perturbation factor has order at most $k-1$ and is treated outside the matrix. The nonlocal streamfunction terms and the remaining high-order modulation terms are also kept outside the pointwise signed matrix and are controlled by the weighted elliptic, interpolation, and lower-order estimates.

Aggregate Euler estimate. All high-order linear terms not represented by the complete signed matrix are required to satisfy

$$\left| \sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i \rangle_{\Phi_{i,k}} - \int_{\mathbb{D}} \mathbf{V}_k^\top \mathbb{M}_{\mathcal{L}}^{(k),\mathrm{full}} \mathbf{V}_k \, dr\, dz \right| \le A_{\mathrm{lin}}^{k,\mathrm{rem}} \mathfrak{H}_k^2 + A_{\mathrm{lin}}^{0,\mathrm{rem}} \mathfrak{E}_0^2. \tag{S521}$$

Provided

$$A_{\mathrm{lin}}^{k,\mathrm{rem}} < \Lambda_{D^\alpha L}^{\mathrm{full}}, \tag{S522}$$

define the net high-order constants by

$$\Lambda_{D^\alpha \mathcal{L}}^{\mathrm{high}} := \Lambda_{D^\alpha L}^{\mathrm{full}} - A_{\mathrm{lin}}^{k,\mathrm{rem}}, \qquad \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}} := A_{\mathrm{lin}}^{0,\mathrm{rem}}. \tag{S523}$$

Combining the remainder bound with the matrix certificate (S248) gives the complete Euler high-order estimate

$$\sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i \rangle_{\Phi_{i,k}} \le \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}} \, \mathfrak{E}_0^2 - \Lambda_{D^\alpha \mathcal{L}}^{\mathrm{high}} \, \mathfrak{H}_k^2. \tag{S524}$$

Thus the complete high-order matrix margin is reduced only by the top-order part of the remainder bound, while its low-order part becomes the coefficient $\Lambda_{D^\alpha \mathcal{L}}^{\mathrm{low}}$. This is the high-order linear estimate used together with the low-order result when the final stability margin is assembled.

Possible refinements. The coupled low-order and high-order mechanism above is the first target. If its certified margin is positive but smaller than desired, the natural next step is to improve the constants already present rather than change the proof strategy. In particular, one can refine the steady profile, re-optimize the weights, shrink the localized sets B_ℓ, sharpen the concentration and interpolation bounds, retain additional favorable linear terms with their sign, and optimize k using the complete high-order matrix together with its remainder. These are quantitative improvements of the same damping mechanism.

Only if the complete top-order signed matrix remained too weak on a substantial region after these refinements would a genuinely stronger top-order estimate be needed. The low-order localization argument cannot simply be repeated for concentration of the top derivatives. In that case one would refine the top-order energy or exploit additional forms of damping.

Appendix I. Proofs of the Stability Theorems

We now prove the single-radius stability Theorem S1.

Theorem S3 (Euler single-radius stability)**.** *Assume that the following perturbation equations hold for admissible perturbations:*

$$\partial_t \delta\omega = \mathcal{L}_\omega(\delta) + \mathcal{N}_\omega(\delta) + \mathcal{R}_\omega, \tag{S525}$$
$$\partial_t \partial_r \delta u = \partial_r \mathcal{L}_u(\delta) + \mathcal{N}_r(\delta) + \mathcal{R}_r, \tag{S526}$$
$$\partial_t \partial_z \delta u = \partial_z \mathcal{L}_u(\delta) + \mathcal{N}_z(\delta) + \mathcal{R}_z. \tag{S527}$$

Assume that the following estimates have been certified with finite constants:

$$\langle \mathcal{L}_\omega, \delta\omega\rangle_{\Phi_\omega} + \langle \partial_r \mathcal{L}_u, \partial_r \delta u\rangle_{\Phi_r} + \langle \partial_z \mathcal{L}_u, \partial_z \delta u\rangle_{\Phi_z} \le -\Lambda_{\mathcal{L}}^{\mathrm{low}}\, \mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\mathrm{high}}\, \mathfrak{H}_k^2. \tag{S528}$$

$$\sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{L}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \le \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}}\, \mathfrak{E}_0^2 - \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}}\, \mathfrak{H}_k^2. \tag{S529}$$

$$|\langle \mathcal{N}_\omega, \delta\omega\rangle_{\Phi_\omega} + \langle \mathcal{N}_r, \partial_r \delta u\rangle_{\Phi_r} + \langle \mathcal{N}_z, \partial_z \delta u\rangle_{\Phi_z}| \le \Lambda_{\mathcal{N},3}\, \mathfrak{E}_k^3. \tag{S530}$$

$$\mu_k \left| \sum_{|\alpha|=k} \langle D^\alpha \mathcal{N}_\omega, D^\alpha \delta\omega\rangle_{\Phi_{\omega,k}} + \sum_{|\alpha|=k} \langle D^\alpha \mathcal{N}_r, D^\alpha \partial_r \delta u\rangle_{\Phi_{r,k}} + \sum_{|\alpha|=k} \langle D^\alpha \mathcal{N}_z, D^\alpha \partial_z \delta u\rangle_{\Phi_{z,k}} \right| \le \Lambda_{D^\alpha\mathcal{N},3}\, \mathfrak{E}_k^3. \tag{S531}$$

$$\left| \sum_{i\in\{\omega,r,z\}} \langle \mathcal{R}_i, \delta v_i\rangle_{\Phi_i} + \mu_k \sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha \mathcal{R}_i, D^\alpha \delta v_i\rangle_{\Phi_{i,k}} \right| \le \Lambda_{\mathcal{R}}\, \mathfrak{E}_k. \tag{S532}$$

Choose $\mu_k \in (0,1]$ and

$$\Lambda_{\mathrm{stab}}(\mu_k) := \min\left\{ \Lambda_{\mathcal{L}}^{\mathrm{low}} - \mu_k \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{low}},\ \Lambda_{D^\alpha\mathcal{L}}^{\mathrm{high}} - \frac{\Lambda_{\mathcal{L}}^{\mathrm{high}}}{\mu_k} \right\} > 0. \tag{S533}$$

After this choice of μ_k is fixed, write Λ_{stab} for the positive number $\Lambda_{\mathrm{stab}}(\mu_k)$.

Define the combined nonlinear constant by

$$\Lambda_3 := \Lambda_{\mathcal{N},3} + \Lambda_{D^\alpha\mathcal{N},3}. \tag{S534}$$

Assume that $\delta_* > 0$ satisfies

$$\Lambda_3 \delta_* + \frac{\Lambda_{\mathcal{R}}}{\delta_*} < \Lambda_{\mathrm{stab}}. \tag{S535}$$

Then every solution with $\mathfrak{E}_k(0) < \delta_*$ remains in the ball $\mathfrak{E}_k(t) < \delta_*$ for as long as the solution exists and the certified estimates apply.

Proof (✓ Formalized and verified in Lean)

We pair the three perturbation equations with $(\delta\omega, \partial_r\delta u, \partial_z\delta u)$ in the low-order weights and add the kth-order pairings with prefactor μ_k.

The two linear estimates in the theorem, with the high-order estimate multiplied by μ_k, give

$$\begin{aligned} &-\Lambda_{\mathcal{L}}^{\text{low}}\,\mathfrak{E}_0^2 + \Lambda_{\mathcal{L}}^{\text{high}}\,\mathfrak{H}_k^2 + \mu_k\big(-\Lambda_{D^\alpha\mathcal{L}}^{\text{high}}\,\mathfrak{H}_k^2 + \Lambda_{D^\alpha\mathcal{L}}^{\text{low}}\,\mathfrak{E}_0^2\big) \\ &\qquad = -(\Lambda_{\mathcal{L}}^{\text{low}} - \mu_k\Lambda_{D^\alpha\mathcal{L}}^{\text{low}})\mathfrak{E}_0^2 - (\mu_k\Lambda_{D^\alpha\mathcal{L}}^{\text{high}} - \Lambda_{\mathcal{L}}^{\text{high}})\mathfrak{H}_k^2. \end{aligned} \tag{S536}$$

By (S533),

$$\Lambda_{\mathcal{L}}^{\text{low}} - \mu_k\Lambda_{D^\alpha\mathcal{L}}^{\text{low}} \geq \Lambda_{\text{stab}}, \qquad \mu_k\Lambda_{D^\alpha\mathcal{L}}^{\text{high}} - \Lambda_{\mathcal{L}}^{\text{high}} \geq \mu_k\Lambda_{\text{stab}}. \tag{S537}$$

Therefore

$$-(\Lambda_{\mathcal{L}}^{\text{low}} - \mu_k\Lambda_{D^\alpha\mathcal{L}}^{\text{low}})\mathfrak{E}_0^2 - (\mu_k\Lambda_{D^\alpha\mathcal{L}}^{\text{high}} - \Lambda_{\mathcal{L}}^{\text{high}})\mathfrak{H}_k^2 \leq -\Lambda_{\text{stab}}\mathfrak{E}_0^2 - \mu_k\Lambda_{\text{stab}}\mathfrak{H}_k^2 = -\Lambda_{\text{stab}}\mathfrak{E}_k^2, \tag{S538}$$

since $\mathfrak{E}_k^2 = \mathfrak{E}_0^2 + \mu_k\mathfrak{H}_k^2$.

The low-order nonlinear estimate (S530), the high-order nonlinear estimate (S531), and (S534) give

$$\left| \sum_{i\in\{\omega,r,z\}} \langle \mathcal{N}_i, \delta v_i\rangle_{\Phi_i} + \mu_k \sum_{i\in\{\omega,r,z\}} \sum_{|\alpha|=k} \langle D^\alpha\mathcal{N}_i, D^\alpha\delta v_i\rangle_{\Phi_{i,k}} \right| \leq \Lambda_3\mathfrak{E}_k^3. \tag{S539}$$

Combining (S538), (S539), and the residual bound (S532) yields

$$\frac{1}{2}\frac{d}{dt}\mathfrak{E}_k^2 \leq -\Lambda_{\text{stab}}\mathfrak{E}_k^2 + \Lambda_3\mathfrak{E}_k^3 + \Lambda_{\mathcal{R}}\mathfrak{E}_k. \tag{S540}$$

It remains to show that the energy ball is forward invariant. We suppose that $\mathfrak{E}_k(0) < \delta_*$ and argue by contradiction. Define the first boundary-contact time by

$$t_* := \inf\{t > 0 : \mathfrak{E}_k(t) = \delta_*\} \tag{S541}$$

Set $F(t) := \mathfrak{E}_k(t)^2$. Then $F(t) < \delta_*^2$ for $t < t_*$ and $F(t_*) = \delta_*^2$, so the left derivative satisfies

$$F'(t_*) = \lim_{h\downarrow 0} \frac{F(t_*) - F(t_* - h)}{h} \geq 0. \tag{S542}$$

On the other hand, evaluating (S540) at t_* gives

$$\frac{1}{2}F'(t_*) \leq -\Lambda_{\text{stab}}\delta_*^2 + \Lambda_3\delta_*^3 + \Lambda_{\mathcal{R}}\delta_* = -\delta_*^2\left(\Lambda_{\text{stab}} - \Lambda_3\delta_* - \frac{\Lambda_{\mathcal{R}}}{\delta_*}\right) < 0 \tag{S543}$$

by (S535), which contradicts (S542).

Therefore the ball $\mathfrak{E}_k < \delta_*$ is forward invariant, which proves the stated nonlinear stability. ■

Appendix J. Proof of the Physical Reconstruction Proposition

This appendix provides the proof of the reconstruction proposition stated in Section 3.8.

Proposition S10 (Finite-time physical blowup from a signed amplitude rate)**.** *Consider the Euler equations in the dynamic rescaling formulation above, with $\lambda = 1/2$. Assume that the exact admissible rescaled solution exists for every $t \geq 0$, remains in the corresponding stability ball, and satisfies the modulation bounds throughout. Assume in addition that*

$$c_u(t) \leq -\gamma_u < 0, \qquad t \geq 0, \tag{S544}$$

for some $\gamma_u > 0$. Then the physical time $\mathring{t}(t)$ converges to a finite limit T, with

$$0 < T - \mathring{t}(t) \leq \frac{e^{-\gamma_u t}}{\gamma_u \, \mathrm{s}_u(0)}. \tag{S545}$$

The traveling center $\mathring{z}_c(t)$ converges to a finite physical location. Moreover,

$$\left|\omega^z\big(0, \mathring{z}_c(t), \mathring{t}(t)\big)\right| = 2|\bar{u}_0| \, \mathrm{s}_u(t) \longrightarrow \infty. \tag{S546}$$

Consequently,

$$\|\tilde{\omega}(\cdot, \mathring{t}(t))\|_{L^\infty(\mathbb{R}^3)} \longrightarrow \infty, \tag{S547}$$

and the reconstructed physical solution cannot be continued smoothly through T.

Proof (✓ Formalized and verified in Lean)

The amplitude scaling law and (S544) give

$$\partial_t \mathrm{s}_u = -c_u \mathrm{s}_u, \qquad \mathrm{s}_u(t) \geq \mathrm{s}_u(0) \, e^{\gamma_u t}. \tag{S548}$$

Using (S172), define

$$T = \mathring{t}(0) + \int_0^\infty \frac{d\tau}{\mathrm{s}_u(\tau)}. \tag{S549}$$

The lower bound in (S548) makes this integral finite and gives

$$0 < T - \mathring{t}(t) = \int_t^\infty \frac{d\tau}{\mathrm{s}_u(\tau)} \leq \frac{e^{-\gamma_u t}}{\gamma_u \mathrm{s}_u(0)}, \tag{S550}$$

which proves (S545).

Since $\mathrm{s}_u(t) > 0$, the physical clock is strictly increasing and maps $0 \leq t < \infty$ onto $\mathring{t}(0) \leq \mathring{t} < T$.

At the moving physical center, the rescaled coordinates are $(r, z) = (0, 0)$. The normalization condition $u(0, 0, t) = \bar{u}_0$ and the inverse amplitude rescaling give

$$u_\bullet\big(0, \mathring{z}_c(t), \mathring{t}(t)\big) = \mathrm{s}_u(t)\bar{u}_0. \tag{S551}$$

Since $u^\theta = \mathring{r} \, u_\bullet$, axisymmetry gives

$$\omega^z = \frac{1}{\mathring{r}} \partial_{\mathring{r}}(\mathring{r} u^\theta) = 2u_\bullet + \mathring{r} \, \partial_{\mathring{r}} u_\bullet. \tag{S552}$$

Axis regularity then yields

$$\omega^z(0, \mathring{z}, \mathring{t}) = 2u_\bullet(0, \mathring{z}, \mathring{t}). \tag{S553}$$

Combining this identity with (S551) and (S548) gives

$$\left|\omega^z\big(0, \mathring{z}_c(t), \mathring{t}(t)\big)\right| = 2|\bar{u}_0|\mathrm{s}_u(t) \longrightarrow \infty. \tag{S554}$$

It remains only to locate the singularity in physical space.

The modulation bounds give $|C(t)| \le C_*$ for some finite C_*.

Since $\lambda = 1/2$,

$$\mathrm{s}_r(t) = \mathrm{s}_r(0)\, e^{-t/2}. \tag{S555}$$

The center equation therefore gives

$$\int_0^\infty |\partial_t \mathring{z}_c(t)|\, dt \le C_* \mathrm{s}_r(0) \int_0^\infty e^{-t/2} dt < \infty. \tag{S556}$$

Hence $\mathring{z}_c(t)$ converges to a finite limit $\mathring{z}_*$.

For each finite rescaled time, the exact change of variables gives an equivalent smooth physical solution. A smooth continuation through T would keep the vorticity finite near $(0, \mathring{z}_*, T)$, contradicting (S546). ■

References

[1] Charles L. Fefferman. Existence and smoothness of the Navier–Stokes equation. In James Carlson, Arthur Jaffe, and Andrew Wiles, editors, *The Millennium Prize Problems*, pages 57–67. Clay Mathematics Institute, Cambridge, MA, 2006.

[2] Steve Smale. Mathematical problems for the next century. *The mathematical intelligencer*, 20(2): 7–15, 1998.

[3] Jiajie Chen and Thomas Y Hou. Singularity formation in 3d Euler equations with smooth initial data and boundary. *Proceedings of the National Academy of Sciences*, 122(27):e2500940122, 2025.

[4] Tarek M Elgindi. Finite-time singularity formation for cˆ1,α solutions to the incompressible Euler equations on rˆ3. *Annals of Mathematics*, 194(3):647–727, 2021.

[5] Steve Shkoller. Incompressible Euler Blowup at the $C^{1,\frac{1}{3}}$ Threshold. *arXiv preprint arXiv:2603.10945*, 2026.

[6] Jiajie Chen. Asymptotically self-similar blowup for 3d incompressible Euler with $c^{1,1/3-}$ velocity i: c^{∞} 1d limiting profiles, 2026. URL https://arxiv.org/abs/2605.15149.

[7] Jiajie Chen. Asymptotically self-similar blowup for 3d incompressible Euler with $c^{1,1/3-}$ velocity ii: 3d profiles, blowup, and limiting behavior, 2026. URL https://arxiv.org/abs/2605.15130.

[8] Guo Luo and Thomas Y. Hou. Potentially singular solutions of the 3d axisymmetric Euler equations. *Proceedings of the National Academy of Sciences*, 111(36):12968–12973, 2014. doi: 10.1073/pnas.1405238111.

[9] Guo Luo and Thomas Y. Hou. Toward the finite-time blowup of the 3d axisymmetric Euler equations: A numerical investigation. *Multiscale Modeling & Simulation*, 12(4):1722–1776, 2014. doi: 10.1137/140966411.

[10] Dongho Chae. Nonexistence of asymptotically self-similar singularities in the Euler and the Navier–Stokes equations. *Mathematische Annalen*, 338(2):435–449, 2007.

[11] Dongho Chae. On the blow-up problem for the Euler equations and the liouville type results in the fluid equations. *Discrete & Continuous Dynamical Systems-S*, 6(5):1139–1150, 2013. doi: 10.3934/dcdss.2013.6.1139.

[12] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

[13] Yongji Wang, C-Y Lai, Javier Gómez-Serrano, and Tristan Buckmaster. Asymptotic self-similar blow-up profile for three-dimensional axisymmetric Euler equations using neural networks. *Physical Review Letters*, 130(24):244002, 2023.

[14] Yongji Wang, Mehdi Bennani, James Martens, Sébastien Racanière, Sam Blackwell, Alex Matthews, Stanislav Nikolov, Gonzalo Cao-Labora, Daniel S. Park, Martin Arjovsky, Daniel Worrall, Chongli Qin, Ferran Alet, Borislav Kozlovskii, Nenad Tomašev, Alex Davies, Pushmeet Kohli, Tristan Buckmaster, Bogdan Georgiev, Javier Gómez-Serrano, Ray Jiang, and Ching-Yao Lai. Discovery of Unstable Singularities, 2025.

[15] Yixuan Wang, Ziming Liu, Zongyi Li, Anima Anandkumar, and Thomas Y. Hou. High precision PINNs in unbounded domains: application to singularity formulation in PDEs. *arXiv preprint arXiv:2506.19243*, 2025.

[16] Peter Constantin, Mihaela Ignatova, and Vlad Vicol. On putative self-similarity for incompressible 3d Euler. *arXiv preprint arXiv:2602.17570*, 2026.

[17] Jiajie Chen, Thomas Y. Hou, and De Huang. On the finite time blowup of the De Gregorio model for the 3d Euler equations. *Communications on Pure and Applied Mathematics*, 74(6):1282–1350, 2021. doi: https://doi.org/10.1002/cpa.21991. URL https://onlinelibrary.wiley.com/doi/abs/10.1002/cpa.21991.

[18] Valentin Duruisseaux, Adarsh Ganeshram, Robert J. George, and Anima Anandkumar. Stability framework for the singularity of the Euler equations on $\mathbb{R}^3$. *arXiv preprint*, 2026.

[19] Thomas Y. Hou, Tianling Jin, and Pengfei Liu. Potential singularity for a family of models of the axisymmetric incompressible flow. *Journal of Nonlinear Science*, 28(6):2217–2247, 12 2018. ISSN 1432-1467. doi: 10.1007/s00332-017-9370-9. URL https://doi.org/10.1007/s00332-017-9370-9.

[20] Zhiwei Fang, Sifan Wang, and Paris Perdikaris. Ensemble learning for physics informed neural networks: a gradient boosting approach. In *ICLR 2024 Workshop on AI4DifferentialEquations In Science*, 2024.

[21] Yongji Wang and Ching-Yao Lai. Multi-stage neural networks: Function approximator of machine precision. *Journal of Computational Physics*, 504:112865, 2024.

[22] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

[23] Guangyuan Wang, Mads Toftrup, Sebastian Loeschcke, Yixuan Wang, and Anima Anandkumar. SS-ESOAP: Self-scaled adaptive preconditioning with applications to matrix optimization and physics-informed machine learning. *ArXiv*, 2026.

[24] M. Al-Baali. Numerical experience with a class of self-scaling quasi-newton algorithms. *Journal of Optimization Theory and Applications*, 96(3):533–553, March 1998. ISSN 0022-3239. doi: 10.1023/A:1022608410710.

[25] Jorge F. Urbán, Petros Stefanou, and José A. Pons. Unveiling the optimization process of physics informed neural networks: How accurate and competitive can pinns be? *Journal of Computational Physics*, 523:113656, 2025. ISSN 0021-9991. doi: https://doi.org/10.1016/j.jcp.2024.113656.

[26] Fredrik Johansson. Arb: Efficient arbitrary-precision midpoint-radius interval arithmetic. *IEEE Transactions on Computers*, 66(8):1281–1292, 2017. doi: 10.1109/TC.2017.2690633.

[27] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.

[28] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs*, CPP 2020, pages 367–381, New York, NY, USA, 1 2020. ACM. doi: 10.1145/3372885.3373824.

[29] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B. Shah. Julia: A fresh approach to numerical computing. *SIAM Review*, 59(1):65–98, 2017. doi: 10.1137/141000671.

[30] Robert Joseph George, Jennifer Cruden, Will Adkisson, Xiangru Zhong, Huan Zhang, and Anima Anandkumar. TorchLean: Formalizing neural networks in lean, 2026.

[31] Anima Anandkumar. How do we build AI to push the frontiers of scientific discovery? *Dædalus*, 155 (1–2):216–228, 2026. doi: 10.1162/daed.a.983.

[32] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

[33] Valentin Duruisseaux, Jean Kossaifi, and Anima Anandkumar. Fourier neural operators explained: A practical perspective, 2025.

[34] Kamyar Azizzadenesheli, Nikola Kovachki, Zongyi Li, Miguel Liu-Schiaffini, Jean Kossaifi, and Anima Anandkumar. Neural operators for accelerating scientific simulations and design. *Nature Reviews Physics*, 6:320–328, 2024. doi: 10.1038/s42254-024-00712-5.

[35] Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM/IMS Journal of Data Science*, 1(3):1–27, 2024. doi: 10.1145/3648506.

[36] Adarsh Ganeshram, Haydn Maust, Valentin Duruisseaux, Zongyi Li, Yixuan Wang, Daniel Leibovici, Oscar Bruno, Thomas Hou, and Anima Anandkumar. FC-PINO: High Precision Physics-Informed Neural Operators via Fourier Continuation, 2025.

[37] Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, Robert M. Kirby, and Michael W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 34, pages 26548–26560, 2021.

[38] Raphael Leiteritz and Dirk Pflüger. How to avoid trivial solutions in physics-informed neural networks, 2021.

[39] Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021. doi: 10.1137/20M1318043.

[40] Leonhard Euler. Principes généraux de l'état d'équilibre des fluides. *Mémoires de l'Académie des sciences de Berlin*, 11:217–273, 1757.

[41] Pengfei Liu. *Spatial Profiles in the Singular Solutions of the 3D Euler Equations and Simplified Models*. PhD thesis, California Institute of Technology, 2017.

[42] Julius Berner, Miguel Liu-Schiaffini, Jean Kossaifi, Valentin Duruisseaux, Boris Bonev, Kamyar Azizzadenesheli, and Anima Anandkumar. Principled approaches for extending neural architectures to function spaces for operator learning. *Nature Machine Intelligence*, 8:1173–1181, 2026. doi: 10.1038/s42256-026-01267-z.

[43] Chenhui Xu, Dancheng Liu, Amir Nassereldine, and Jinjun Xiong. Fp64 is all you need: Rethinking failure modes in physics-informed neural networks. In *Advances in Neural Information Processing Systems*, volume 38, pages 158613–158634, 2025. doi: 10.52202/085713-4785.

[44] Vineet Gupta, Tomer Koren, and Yoram Singer. Shampoo: Preconditioned stochastic tensor optimization. In *International Conference on Machine Learning*, pages 1842–1850, 2018.

[45] Nikhil Vyas, Depen Morwani, Rosie Zhao, Itai Shapira, David Brandfonbrener, Lucas Janson, and Sham M. Kakade. SOAP: Improving and stabilizing Shampoo using Adam for language modeling. In *The Thirteenth International Conference on Learning Representations*, 2025.

[46] Dougal Maclaurin, David Duvenaud, and Ryan P Adams. Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, 2015.

[47] Atılım Günes Baydin, Barak A. Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*, 18 (153):1–43, 2018. ISSN 1532-4435.

[48] Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023. ISSN 0045-7825. doi: https://doi.org/10.1016/j.cma.2022.115671. URL https://www.sciencedirect.com/science/article/pii/S0045782522006260.

[49] Tarek M. Elgindi. Dynamics of ideal fluid flows. In *Proceedings of the International Congress of Mathematicians 2026 - Volume 5: Invited Lectures (Sections 9–11)*, pages 285–306. Society for Industrial and Applied Mathematics, 2026. doi: 10.1137/25M1807137.

[50] Jiajie Chen. *Singularity Formation in Incompressible Fluids and Related Models*. PhD thesis, California Institute of Technology, Pasadena, California, 5 2022.

[51] Thomas Hou and Zhen Lei. On the partial regularity of a 3d model of the Navier–Stokes equations. *Communications in Mathematical Physics*, 287:589–612, 2009. doi: 10.1007/s00220-008-0689-9.