A&A, 710, A264 (2026)
https://doi.org/10.1051/0004-6361/202659496

Astronomy & Astrophysics

The Rayleigh–Taylor instability in partially ionised plasmas: Ambipolar diffusion effects in the non-linear phase

E. Callies[1,⋆], Z. Meliani[2], V. Guillet[1], and A. Marcowith[1]

[1] Laboratoire Univers et Particules de Montpellier (LUPM) Université Montpellier, CNRS/IN2P3, CC72, Place Eugène Bataillon, 34095 Montpellier Cedex5, France
[2] Laboratoire d'étude de l'Univers et des phénomènes eXtrêmes(LUX), Observatoire de Paris,Université PSL, Sorbonne Université, CNRS, 92190 Meudon, France

Received 18 February 2026 / Accepted 23 April 2026

ABSTRACT

Context. The Rayleigh–Taylor (RT) instability is a key mechanism that drives mixing and structure formation in stratified astrophysical media. In partially ionised environments such as the molecular interstellar medium, including irradiated H_2 regions associated with stellar clusters (e.g. the Pleiades), ion–neutral coupling, and ambipolar diffusion are expected to play a major role in shaping the instability evolution in the presence of gravity and magnetic fields.
Aims. Our aim was to determine how ion–neutral coupling and ambipolar diffusion affect the linear and the non-linear growth of the RT instability under astrophysically relevant conditions, and to identify the coupling regimes in which departures from the classical single-fluid picture become significant.
Methods. We performed high-resolution two-fluid numerical simulations using the `MPI-AMRVAC` code, spanning a wide range of perturbation wavelengths, coupling strengths, from uncoupled to strongly coupled, passing by intermediate or ambipolar diffusion-dominated regimes, and magnetic field configurations. The linear theory was revisited using a physically consistent formulation with different ion–neutral coupling strengths across the interface and validated against the simulations. We investigated the physics of the instability using morphology-based diagnostics of the mixing layer to compare simulations at equivalent non-linear stages, complemented by spectral, force, and energy budget analyses.
Results. In the linear regime, theoretical growth rates are recovered over a wide range of wavelengths, from the single-fluid limit to intermediate bi-fluid coupling. In the non-linear regime, ambipolar diffusion modifies the classical quadratic growth and introduces a coupling-dependent evolution. For multi-wavelength perturbations, the non-linear dynamics becomes strongly scale-dependent: intermediate coupling enhances fragmentation in hydrodynamic configurations, while magnetised cases exhibit a non-monotonic reorganisation of the interface, with the smoothest morphologies occurring at intermediate coupling. Spectral and energetic diagnostics indicate that these behaviours correlate with changes in the relative contributions of ion–neutral drift and magnetic stresses during the non-linear evolution.
Conclusions. Our results demonstrate that ambipolar diffusion does not merely rescale RT growth rates, but reshapes the non-linear, multi-scale dynamics by altering how gravity-driven kinetic energy is redistributed through magnetic tension and ion–neutral drift. In magnetised configurations, the magnetic field suppresses small-scale corrugations, while ambipolar diffusion weakens this constraint by allowing partial decoupling between ions and neutrals.

Key words. ISM: magnetic fields

1. Introduction

The Rayleigh–Taylor instability (RTI) is a fundamental hydrodynamic instability that occurs when a dense fluid is supported against gravity by a lighter one. First identified by Lord Rayleigh (Rayleigh 1883) and later confirmed experimentally by Taylor (Taylor 1950), it arises when an interface between two fluids of different densities becomes unstable under gravity, with the heavier fluid lying above the lighter one. Small perturbations of the interface are amplified as dense fluid sinks and light fluid rises, leading to the formation of characteristic spikes and bubbles in the non-linear regime. In its classical formulation, the RTI is a purely hydrodynamic, large-scale instability, independent at leading order of viscosity and compressibility, and is usually analysed for two immiscible, incompressible, inviscid fluids separated by a planar interface in a uniform gravitational field.

Beyond its idealised formulation, the Rayleigh–Taylor instability is of broad relevance in astrophysics because it naturally arises in any stratified system subjected to an effective acceleration, whether due to gravity, pressure gradients, or inertial forces. It is therefore expected to operate over a wide variety of macroscopic environments, ranging from accretion flows onto compact magnetised objects (Arons & Lea 1976; Wang & Nepveu 1983) to buoyant structures in galaxy clusters (Robinson et al. 2004; Ruszkowski et al. 2007), and in solar and stellar contexts where magnetic flux systems evolve under gravity (Isobe et al. 2005). The instability is also a key ingredient in the dynamics of supernova remnants and pulsar wind nebulae, where accelerated dense shells interact with lighter surrounding media (Jun & Norman 1996; Hester et al. 1996; Bucciantini et al. 2004). In all these systems, the Rayleigh–Taylor instability governs the growth of large-scale structures, promotes mixing between plasma components, and strongly influences the global morphology and evolution of the flow. This ubiquity highlights

⋆ Corresponding author: edouard.callies@umontpellier.fr

the RTI as a fundamental mechanism for momentum and energy redistribution in accelerated astrophysical plasmas, well beyond the idealised hydrodynamic setting.

A large body of analytical work has been devoted to understanding the linear phase of the instability. In hydrodynamics, the classical dispersion relation predicts an unbounded growth rate that increases with the square root of the wavenumber, implying the absence of any short-wavelength cutoff. In contrast, Chandrasekhar's seminal analysis of the magnetohydrodynamic (MHD) RTI (Chandrasekhar 1961) showed that a magnetic field tangential to the interface introduces a stabilising magnetic tension, leading to the existence of a critical wavenumber above which perturbations are stabilised. Later studies extended this framework by considering inclined magnetic fields. In particular, Vickers et al. (2020) demonstrated that when the magnetic field is oblique with respect to the interface, the stabilising effect is weakened and the cutoff can disappear, restoring unstable modes at all wavelengths. In partially ionised plasmas, a bi-fluid description is required to account for the relative drift between ions and neutrals and for non-ideal effects such as ambipolar diffusion. In this context, Díaz et al. (2012) and Díaz et al. (2014) investigated the impact of compressibility and ion–neutral coupling on the linear RTI, showing that partial ionisation can significantly modify growth rates and mode structure compared to single-fluid MHD predictions. In our recent linear analysis (Callies et al. 2025) (hereafter Paper I), we have combined these ingredients by considering the RTI in a bi-fluid framework with an inclined magnetic field. Our results demonstrate that ambipolar diffusion plays a non-negligible role in the linear development of the instability. In particular, in the parameter regime relevant to partially ionised plasmas, the growth rate exhibits a markedly different scaling, increasing approximately linearly with the wavenumber k, rather than following the standard hydrodynamic $k^{1/2}$ dependence.

While linear analytical studies have provided invaluable insight into the onset and early evolution of the RTI, they are inherently restricted to the initial exponential growth phase. Extending such approaches into the fully non-linear regime is notoriously difficult, as mode coupling, secondary instabilities, and turbulent mixing rapidly invalidate linear assumptions. For this reason, numerical simulations have become the primary tool for investigating the late-time development and saturation of the RTI. A qualitative criterion for the transition to non-linearity is that the interface displacement in the vertical direction becomes comparable to the inverse wavenumber, $\sim 1/k$, i.e. when the vertical and horizontal scales are of the same order (Fermi & von Neumann 1953) (see also the discussion in Hillier 2016). In the non-linear regime, the hydrodynamic RTI is observed to evolve in a self-similar manner; the thickness h of the mixing layer is governed by

$$h \simeq \alpha A g t^2, \tag{1}$$

where A is the Atwood number, g the acceleration, and α a dimensionless non-linear growth rate that is found to be only weakly sensitive to initial conditions (Ristorcelli & Clark 2004; Cook et al. 2004). This behaviour has been confirmed in numerous laboratory experiments and numerical studies.

In magnetised plasmas, the non-linear phase has been explored through two- and three-dimensional MHD simulations. Early 2D studies investigated how magnetic fields aligned either tangentially or normally to the interface modify the growth and morphology of the instability (Jun et al. 1995). Fully 3D simulations later demonstrated that magnetic tension can strongly affect bubble and spike structures and the efficiency of mixing, sometimes even enhancing the late-time growth relative to the purely hydrodynamic case by suppressing secondary Kelvin–Helmholtz roll-ups (Stone & Gardiner 2007b,c). Idealised 3D MHD simulations by Carlyle & Hillier (2017) systematically examined the effect of magnetic field strength in a parameter regime relevant to astrophysical plasmas, showing that stronger fields tend to reduce the non-linear growth rate of rising bubbles and introduce asymmetries between bubbles and spikes, while the overall evolution remains compatible with a self-similar $h \propto Agt^2$ scaling.

Recent numerical studies have significantly advanced the investigation of the non-linear RTI in solar prominence conditions. Using 2D two-fluid simulations, Popescu Braileanu et al. (2021a) explored the influence of magnetic field strength, shear, and mass loading on the non-linear evolution of the instability at a smooth prominence–corona interface, and showed that magnetic shear can substantially reduce or suppress the instability and that non-linear development is accompanied by coherent magnetic structuring and partial ion–neutral decoupling. In a follow-up study, Popescu Braileanu et al. (2021b) demonstrated that the ion–neutral collision frequency critically controls both linear growth rates and the emergence of small-scale structures in the non-linear regime. Popescu Braileanu et al. (2023) emphasised magnetic amplification and current-sheet formation driven by the RTI in partially ionised prominence plasmas, highlighting the feedback between instability-induced flows and magnetic topology evolution. These works provide a detailed description of the RTI in realistic solar prominence environments, including smooth interfaces and complex thermodynamics. In contrast, the present study adopts a deliberately idealised configuration with a sharp interface and controlled parameter variations, designed to isolate the dynamical and energetic role of ambipolar diffusion across distinct ion–neutral coupling regimes and to directly connect linear dispersion properties to non-linear mixing-layer evolution.

Beyond prominence-specific set-ups, Changmai et al. (2023) performed high-resolution 2.5D ideal MHD simulations to follow the RTI into a fully developed turbulent state, showing that RTI-driven dynamics can generate anisotropic MHD turbulence with coherent field-aligned flows and power-law spectra. At the same time, recent fully three-dimensional incompressible MHD simulations by Kalluri & Hillier (2025) demonstrated that the non-linear magnetic RTI can evolve differently in two and three dimensions, due to the presence in 3D of mixed and interchange modes that enhance small-scale mixing and modify energy partition. These results highlight the sensitivity of non-linear RTI dynamics to magnetic topology and dimensionality. In this context, controlled two-dimensional studies remain valuable for isolating specific physical mechanisms, particularly when the objective is not to reproduce fully developed turbulence, but to quantify how ion–neutral coupling and ambipolar diffusion modify growth laws, force balance, and energy redistribution relative to ideal MHD benchmarks.

More generally, as emphasised in the review by Soler & Ballester (2022), most non-linear studies of partially ionised plasmas rely either on ideal MHD or on single-fluid formulations including ambipolar diffusion, while fully non-linear two-fluid investigations remain comparatively scarce. In parallel, Khomenko et al. (2025) showed that ambipolar diffusion and ion–neutral drift can strongly influence small-scale structuring and energy dissipation in partially ionised solar plasmas when adequately resolved, underscoring the importance of multi-fluid effects for non-linear plasma dynamics.

In addition, recent high-resolution observations have started to probe the structure of the cold neutral medium (CNM) at unprecedented spatial scales. Using JWST/NIRCam data, Vigoureux et al. (2026) reported pronounced anisotropy in the two-dimensional Fourier statistics of dust emission in the Pleiades nebula, down to scales of order ~40 au. The observed scale-dependent anisotropy, aligned with the magnetic field, provides direct evidence that partially ionised, magnetised plasmas can sustain strongly anisotropic small-scale structuring. Although these observations do not identify a specific driving instability, they highlight the need to understand how magnetically mediated processes and ion–neutral coupling shape non-linear plasma dynamics across scales. In this context, buoyancy-driven instabilities such as the magnetic Rayleigh–Taylor instability, modified by ambipolar diffusion, offer a controlled framework to investigate mechanisms capable of generating anisotropic structuring and scale-dependent energy redistribution in partially ionised environments.

Despite these advances, a systematic exploration of how ambipolar diffusion modifies the non-linear self-similar growth of the magnetic RTI across ion–neutral coupling regimes remains lacking. In particular, it is still unclear whether the classical $h \propto Agt^2$ scaling persists, is delayed, or is fundamentally altered in magnetised, partially ionised plasmas. In this work, we address this issue using high-resolution 2D two-fluid simulations with an oblique magnetic field, combining a self-similar analysis based on the normalised mixing height h/λ with controlled perturbations spanning multiple wavelengths and coupling regimes. To enable meaningful comparisons, simulations are analysed at equivalent dynamical stages using the normalised mixing layer thickness $\Delta h/L_x$, rather than at equal physical times, thereby removing the biases associated with different growth rates across coupling regimes (see Sect. 4).

The paper is organised as follows. In Sect. 2 we describe the physical model and numerical set-up. In Sect. 3 we detail the diagnostics used to define the mixing layer as well as the verification of the theoretical analysis performed in (I). In Sect. 4 we discuss the dependence of the non-linear evolution on wavelength and coupling strength. Finally, in Sect. 5 we summarise our main conclusions and outline the perspectives for future work.

2. Method

The aim of this work was to quantify how ambipolar diffusion modifies the non-linear development of the Rayleigh–Taylor (RT) instability in a partially ionised plasma with an oblique magnetic field. Building on the linear analysis of Paper I, we extended the study beyond the exponential stage and investigated how ion–neutral drift reshapes the structure, growth, and self-similar evolution of the mixing layer once non-linear interactions dominate. We solved the two-fluid equations for charged and neutral species in a stratified configuration under gravity, exploring parameter regimes where magnetic tension, buoyancy, and ion–neutral coupling compete. The selected wavelengths and coupling strengths are guided by the linear predictions of Paper I. We first validated the numerical implementation against the theoretical growth rates before analysing the non-linear multi-scale evolution. Our simulations are not designed to reproduce a specific observational configuration, but to isolate the physical mechanisms controlling ion–neutral drift and ambipolar diffusion in non-linear buoyancy-driven mixing. Resolving these effects numerically requires working in parameter regimes where the relevant coupling and diffusion scales are accessible at finite resolution. Our approach was therefore mechanistic: we identified robust trends in how ambipolar diffusion modifies non-linear RT evolution, which can then inform astrophysical interpretations.

2.1. Governing two-fluid equations

We considered a two-fluid system composed of a charged component (c) and a neutral component (n). The magnetic field $\mathbf{B}$ is inclined by an angle θ with respect to the interface. Ionisation and recombination, Ohmic resistivity, the Hall effect, and explicit resistive heating are neglected. The governing equations correspond to the two-fluid formulation implemented in `MPI-AMRVAC` (Braileanu & Keppens 2022), with the specific choice adopted here that the gravitational acceleration acts only on the charged component. As a result, the neutral fluid does not experience gravity directly and is accelerated solely through ion–neutral collisional coupling. This asymmetric forcing is consistent with the configuration analysed in Paper I and allows us to isolate the role of ion–neutral momentum exchange in the buoyancy-driven evolution of the instability. The governing equations read:

$$\frac{\partial \rho_n}{\partial t} + \nabla \cdot (\rho_n \mathbf{v}_n) = 0, \tag{2}$$

$$\frac{\partial \rho_c}{\partial t} + \nabla \cdot (\rho_c \mathbf{v}_c) = 0, \tag{3}$$

$$\frac{\partial (\rho_n \mathbf{v}_n)}{\partial t} + \nabla \cdot (\rho_n \mathbf{v}_n \mathbf{v}_n + p_n \mathbf{I}) = \mathbf{R}_n, \tag{4}$$

$$\frac{\partial (\rho_c \mathbf{v}_c)}{\partial t} + \nabla \cdot \left[\rho_c \mathbf{v}_c \mathbf{v}_c + \left(p_c + \frac{1}{2}B^2\right)\mathbf{I} - \mathbf{B}\mathbf{B}\right] = \rho_c \mathbf{g}_c - \mathbf{R}_n, \tag{5}$$

$$\frac{\partial e_n^{\rm tot}}{\partial t} + \nabla \cdot \left[\mathbf{v}_n (e_n^{\rm tot} + p_n)\right] = M_n, \tag{6}$$

$$\frac{\partial e_c^{\rm tot}}{\partial t} + \nabla \cdot \left[\mathbf{v}_c (e_c^{\rm tot} + p_c + \tfrac{1}{2}B^2) - \mathbf{B}(\mathbf{v}_c \cdot \mathbf{B})\right] = \rho_c \mathbf{v}_c \cdot \mathbf{g}_c - M_n, \tag{7}$$

$$\frac{\partial \mathbf{B}}{\partial t} + \nabla \cdot \left[\mathbf{v}_c \mathbf{B} - \mathbf{B}\mathbf{v}_c\right] = 0, \tag{8}$$

with collisional momentum and energy exchange terms

$$\mathbf{R}_n = \rho_n \rho_c \alpha \, (\mathbf{v}_c - \mathbf{v}_n), \tag{9}$$

$$M_n = \frac{1}{2}\left(v_c^2 - v_n^2\right)\rho_n \rho_c \alpha + \frac{1}{\gamma - 1}(T_c - T_n)\,\rho_n \rho_c \alpha. \tag{10}$$

Here α is the constant collision coefficient, so that the collision frequencies satisfy $\nu_{\rm nc} = \alpha \rho_n$ and $\nu_{\rm cn} = \alpha \rho_c$ (for the charges and neutrals respectively). We focus on the asymmetric configuration where only the charged fluid experiences gravity, $\mathbf{g}_c = -g\,\hat{\mathbf{z}}$ and $g_n = 0$ as discussed in Paper I.

2.2. Initial equilibrium configuration

The domain consists of two uniform layers separated by a sharp interface at $z = 0$, with identical density contrasts for both fluids:

$$\mathcal{A} = \frac{\rho_{c,2} - \rho_{c,1}}{\rho_{c,2} + \rho_{c,1}} = \frac{\rho_{n,2} - \rho_{n,1}}{\rho_{n,2} + \rho_{n,1}} = \frac{3}{5}. \tag{11}$$

A constant gravity $\mathbf{g}_c = -g\,\hat{\mathbf{z}}$ is applied to the charged fluid. The pressure of the charged fluid satisfies hydrostatic equilibrium,

$$\nabla p_c = \rho_c \mathbf{g}_c, \tag{12}$$

with $p_c(z_{\max}) = 2n_c k_B T_c$, while the neutral pressure is uniform, $p_n = n_n k_B T_n$. We set $T_c = T_n = 100\,\mathrm{K}$ at the top of the box. The magnetic field strength is fixed from the interface pressure through

$$|\mathbf{B}|^2 = 2\frac{p_c(z=0)}{\beta}, \tag{13}$$

where β is the ratio between thermal and magnetic pressure, i.e. the beta plasma. Thereby fixing the magnetic field strength from the interface pressure and the chosen value of β.

2.3. Numerical set-up

The simulations are performed with the two-fluid module of the open-source code `MPI-AMRVAC`[1]. parallelised with MPI. We adopt a 2D configuration. The computational domain is discretised on a 64×64 base grid and refined dynamically through adaptive mesh refinement (AMR). Depending on the simulation, between five and six AMR levels are used, selected through a Löhner-type refinement criterion based on gradients of key physical variables. This allows the mesh to concentrate around the evolving interface and to resolve the fine-scale structures that develop during the instability growth. With six refinement levels, the effective resolution reaches 4096 cells per direction.

Time integration is carried out with the three-step IMEX–ARS3 scheme (Ascher et al. 1997), which treats stiff collisional terms implicitly and ideal MHD fluxes explicitly. Spatial fluxes are computed with the HLLD Riemann solver (Ruuth & Spiteri 2002; Miyoshi & Kusano 2005) and third-order PPM reconstruction (Colella & Woodward 1984). A Courant number of 0.8 is used for the explicit part of the scheme. The divergence of the magnetic field is controlled through a multi-grid cleaning method.

Finally, we impose periodic boundary conditions in the horizontal (x) direction. In the vertical direction (perpendicular to the interface), the upper and lower boundaries are treated with fixed-value (non-periodic) conditions for all primitive variables. This choice prevents the appearance of vanishing or numerically unstable pressures at the domain edges, which would otherwise result from a naïve extrapolation of the hydrostatic profiles. The vertical boundaries therefore act as reservoirs that maintain the prescribed density, pressure, and magnetic field values, while still allowing the instability to develop freely within the interior of the domain.

2.4. Initial perturbations and physical parameter space

The instability is triggered by imposing a sinusoidal perturbation of the interface,

$$\eta(x) = \eta_0 \sin\left(\frac{2\pi}{\lambda}x\right), \tag{14}$$

in a square domain ($L_z = L_x$). Multi-wavelength perturbations are introduced separately in Sect. 4.2. The smallest injected wavelength is $\lambda_{\min} = L_x/58$. At maximum refinement (4096 cells), it is resolved by ~71 grid points, ensuring that all injected modes lie well below the Nyquist limit.

To characterise the simulations and to explore the influence of ambipolar diffusion, magnetic tension, and coupling strength, we adopt a strategy in which a subset of physical parameters is fixed while a controlled set of others is varied systematically. The quantities that remain constant throughout the study are summarised in Table 1. They define the global normalisation of the problem, including the size of the computational domain, the gravitational acceleration applied to the charged component, and the masses and relative abundances of neutrals and charges. The density ratio $n_c/n_n = 10^{-4}$ places the system in a weakly ionised regime where ion–neutral drift and ambipolar diffusion can play a dynamically significant role. All spatial scales are expressed relative to the box size L_x, so that perturbation wavelengths are naturally given as fractions of L_x. In particular, the largest injected wavelength is fixed to $\lambda_{\max} = L_x/4$, which ensures that several unstable modes fit within the computational domain while avoiding direct interaction with the vertical boundaries during the linear stage.

The present set-up is inspired by CNM conditions, representative of cold HI clouds in the interstellar medium. It adopts a weakly ionised regime in which ambipolar diffusion operates in the strong-coupling limit, where ion–neutral drift plays a dynamically significant role. However, the absolute values of L_x and g_c are not chosen to reproduce a specific CNM object. Instead, L_x is set by numerical constraints, while g_c is adjusted so that magnetic tension acts on resolved scales. The dynamics is therefore primarily controlled by dimensionless parameters such as λ/L_x, $\nu_{\rm nc}/\omega_{\rm th}$, and β, rather than by absolute scales. The simulations should thus be interpreted as controlled experiments in a CNM-inspired regime. In practice, we fix the magnetic field strength through the plasma beta parameter, adopting $\beta = 3 \times 10^2$, which ensures that magnetic tension is dynamically important without fully suppressing the instability. The gravitational acceleration is then tuned so that the magnetic cutoff wavelength satisfies $L_{\rm cut} \simeq L_x/20$, allowing stabilised and unstable modes to coexist. This choice maximises the sensitivity of the non-linear evolution to ambipolar diffusion by enforcing a competition between buoyancy, magnetic tension, and ion–neutral coupling on comparable scales (Table 2).

Table 1. Constant physical parameters used in all simulations.

Name of the quantity	Value
Size of the box $L_x = L_z$	3×10^{10} cm
Gravity g_c	10^3 cm s^{-2}
Mass of neutrals m_n	$m_n = m_p$
Mass of charges m_c	$m_c = 12\,m_p$
Ratio of density n_c/n_n	10^{-4}

Table 2. Default values of the parameters varied in the simulations.

Name of the quantity	Value
Magnetic β	3×10^2
Atwood number $\mathcal{A}$	3/5
Magnetic inclination in the xOz plane θ	10°
Largest injected wavelength $\lambda_{\max}$	$L_x/4$

[1] MPI-AMRVAC v3.0-289-g034bdf3d (commit `034bdf3d`), compiled and run with OpenMPI 5.0.8.

Table 3. Definition of the four ion–neutral coupling regimes used throughout this work.

Coupling regime	ν_{nc}/ρ_c [cm^3 g^{-1} s^{-1}]	Physical interpretation
NC (No Coupling)	0	Neutrals and charges evolve independently; no momentum exchange occurs between the two fluids.
LC (Low Coupling)	7.5×10^{19}	Weak coupling regime: ion–neutral interactions are present but insufficient to enforce coherent motion; partial decoupling persists.
IC (Intermediate Coupling)	7.5×10^{20}	Transitional regime: ion–neutral interactions are dynamically important but do not impose a single-fluid behaviour; ambipolar diffusion effects are maximal.
HC (High Coupling)	7.5×10^{21}	Strong coupling regime approaching the fully coupled asymptotic limit; charges and neutrals move nearly coherently as a single fluid.
HC-Lim (High Coupling Limit)	7.5×10^{22}	Totally coupled regime

Notes. The regimes are classified according to the effective collision frequency per unit charged density, ν_{nc}/ρ_c, which controls the strength of momentum exchange between ions and neutrals.

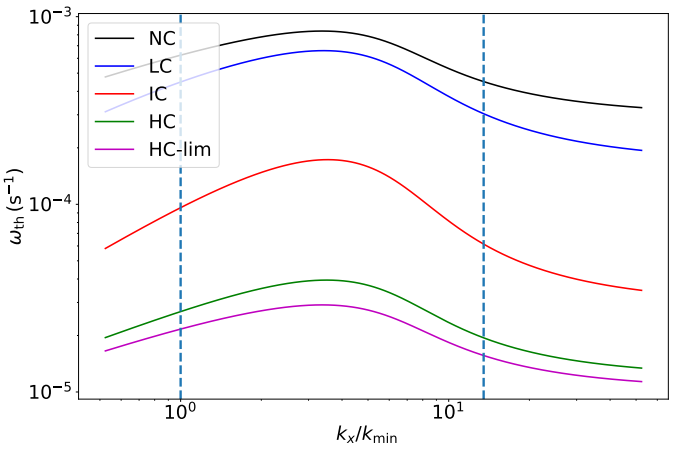

Fig. 1. Theoretical linear growth rate as a function of the horizontal wavenumber in the presence of a magnetic field ($\theta = 10°$). The black curves indicate the asymptotic limits corresponding to the fully uncoupled and fully coupled regimes. The colored curves show intermediate coupling cases for two different assumptions on the variation of the collision frequency across the interface. The dashed bands mark the ranges of wavelengths injected in the simulations, illustrating how the numerical experiments sample the different coupling regimes predicted by linear theory.

The role of ion–neutral coupling is characterised through the neutral–ion collision frequency, which is treated as a free parameter and varied over several orders of magnitude. For clarity, coupling regimes are classified according to the ratio ν_{nc}/ω_{th}, where ω_{th} (see Paper I) is the theoretical linear growth rate associated with the considered wavelength. This naturally defines four regimes (see Table 3 for a summary): No coupling (NC), Low coupling (LC), Intermediate coupling (IC), High coupling (HC), and High limit coupling (HC-Lim). Figure 1 illustrates the corresponding theoretical growth rates as functions of the wavenumber for different assumptions regarding the variation of collision frequencies across the interface. In particular, it highlights how adopting a physically consistent variation of ν_{nc} with density effectively shifts the system toward stronger coupling, without altering the asymptotic limits of the dispersion relation. The scales delimited byje vouthe vertical dashed lines in this figure indicate the ranges of wavelengths injected in the simulations, making explicit how the numerical experiments sample the different coupling regimes predicted by linear theory. The overall survey strategy is therefore twofold: first, to validate the linear behaviour derived in Paper I within this physically consistent parameterisation; and second, to explore how ambipolar diffusion reshapes the non-linear evolution of the instability across scales, from single-mode configurations to fully broadband perturbations.

3. Verification of the linear theory

The aim of this section is to verify that our numerical set-up correctly reproduces the predictions of the linear Rayleigh–Taylor theory in all configurations relevant to this work (hydrodynamic, single-fluid MHD, and bi-fluid). We proceed gradually, beginning with the definition and validation of the diagnostic used to measure the instability growth, and then increasing the physical complexity step by step. The linear theory developed in Paper I relied on the simplifying assumption that the ion–neutral and neutral–ion collision frequencies were identical on both sides of the density interface. While this choice allowed for a compact analytical formulation, it is not strictly consistent with the numerical implementation used in the present work, where the conserved quantity is the drag coefficient rather than the collision frequency itself. As a result, the collision frequencies naturally vary across the interface in proportion to the local density contrast. In order to ensure a physically consistent comparison between the analytical predictions and the numerical simulations, we relax this assumption and allow the collision frequencies to differ on either side of the interface. We show in Appendix A that this modification does not introduce any new qualitative regime of instability. Instead, it primarily results in a renormalisation of the effective coupling strength at intermediate wavelengths, while leaving the uncoupled and fully coupled asymptotic limits unchanged. The analytical framework of Paper I therefore remains fully applicable, provided that the coupling strength is interpreted in terms of an effective collision frequency.

3.1. Hydrodynamic reference: Compressibility check

The theoretical predictions used throughout this work are based on the incompressible linear analysis of Paper I, while the simulations solve the fully compressible equations. It is therefore necessary to assess whether compressibility may significantly alter the linear growth rates in the parameter regime explored

here. Following Díaz et al. (2012), a convenient estimate is provided by the dimensionless ratio gL/c_s^2, where L is the characteristic perturbation wavelength and c_s the sound speed. When $gL/c_s^2 \ll 1$, pressure adjusts rapidly compared to buoyancy-driven motions, and the instability behaves effectively incompressibly. Using the largest wavelength present in our simulations, we obtain the conservative estimate

$$\frac{g_c L}{c_s^2} \simeq 0.1, \tag{15}$$

indicating that compressibility effects should remain weak during the linear stage. This expectation is confirmed quantitatively by comparing incompressible and compressible hydrodynamic growth rates for a representative wavelength $L_x/4$.

The incompressible growth rate $\omega_{\rm th,incomp}$ is obtained from the dispersion relation derived in Paper I, while the compressible growth rate $\omega_{\rm th,comp}$ is computed using the analytical expression of Díaz et al. (2012) (their Eq. (30)), evaluated for the same physical parameters (density contrast, gravity, and wavelength): $\omega_{\rm th,comp} = 6.94 \times 10^{-4}$ s^{-1}, $\omega_{\rm th,incomp} = 7.02 \times 10^{-4}$ s^{-1} which differ by less than 2%. In the remainder of this work, the incompressible dispersion relation is therefore used as a reference framework to interpret the simulations. It is not intended to reproduce the detailed temporal evolution of the compressible system, but to capture the correct ordering of growth rates across wavelengths and coupling regimes. Small temporal offsets between theory and simulations are therefore expected. Given that the smallest injected modes are resolved by ∼71 grid points, numerical diffusivity is not expected to significantly affect the linear growth. In addition, the good agreement with the theoretical linear growth rates further indicates that numerical diffusion remains subdominant compared to physical effects. The remaining discrepancy is therefore attributed to weak compressibility and transient initialisation effects.

3.2. Diagnostic tool: Finger–bubble height in the linear and weakly non-linear regimes

In order to quantify the evolution of the Rayleigh–Taylor instability in a way that is robust and comparable across all physical configurations explored in this work (HD, MHD, and bi-fluid), we follow the vertical penetration of both the rising light-fluid bubbles and the descending heavy-fluid spikes (fingers). Our primary observable is the horizontally averaged heavy-fluid fraction as a function of height,

$$\langle f_h \rangle(z,t) = \frac{1}{L_x} \int_0^{L_x} f_h(x,z,t)\, \mathrm{d}x, \tag{16}$$

where f_h is the local mass fraction of heavy fluid. From $\langle f_h \rangle(z,t)$ we define two characteristic interface positions: the bubble position $z_b(t)$ as the lowest height in the upper half for which the heavy fluid is still present, and the finger position $z_F(t)$ as the greatest height in the lower half for which light fluid is still present, using fixed thresholds $f_{\rm low}$ and $f_{\rm high}$ (here $f_{\rm low} = 0.01$ and $f_{\rm high} = 0.99$). In practice, $z_f(t)$ is obtained from the crossings of $\langle f_h \rangle(z,t) = f_{\rm low}$ on the $z < 0$ side (finger penetration), while $z_b(t)$ is obtained from the crossings of $\langle f_h \rangle(z,t) = f_{\rm high}$ on the $z > 0$ side (bubble rise). We then define the bubble–finger height

$$h(t) \equiv z_b(t) - z_f(t), \tag{17}$$

which directly measures the thickness of the developing mixing region in a way that is insensitive to a drift of the mean interface position. This diagnostic is widely used in RT studies (e.g. Stone & Gardiner 2007a).

3.3. Linear validation with multiple injected wavelengths

We now validate the linear regime of the instability in a controlled multi-wavelength setting, both in the single-fluid MHD limit and in the bi-fluid framework, before turning to non-linear and multi-scale effects. We independently inject several wavelengths and analyse their early-time evolution with the diagnostic of Sect. 3.2. In all simulations, the initial interface displacement is prescribed with a controlled amplitude proportional to the injected wavelength, $h_0 = \varepsilon\lambda$ (here $\epsilon \simeq 1/3$), so that all runs start from the same dimensionless configuration. We analyse the normalised height $h(t)/h_0$ and normalise time by the theoretical linear growth rate $\omega_{\rm th}$ associated with each wavelength, defining $\bar{t} = \omega_{\rm th} t$. In the linear regime, the correct theory or diagnostic behaviour therefore implies $h/h_0 \propto \exp(\bar{t})$ and collapse of curves across different wavelengths when plotted as a function of the scaled time $\bar{t}$. The extent of the linear regime being not strictly identical across configurations, even after normalisation by $\omega_{\rm th}$, it is identified independently in each case as the interval over which the curves follow a clear exponential trend in semi-logarithmic representation.

We first consider the single-fluid MHD configuration obtained by setting the neutral–ion collision frequency to zero ($\nu_{\rm nc} = 0$), with a uniform magnetic field and plasma beta $\beta = 3 \times 10^2$. Three wavelengths are selected so as to sample distinct regions of the theoretical growth rate curve, and each wavelength is injected individually in otherwise identical runs. The resulting evolution of h/h_0 is shown in Fig. 2.

In the linear regime ($0.5 \lesssim \bar{t} \lesssim 2$), the three curves collapse and appear as straight lines in semi-logarithmic representation, indicating that the exponential growth rates measured in the simulations are consistent with the theoretical predictions. Normalising time by $\omega_{\rm th}$ does not modify the exponential behaviour, but rescales the time axis so that different wavelengths can be compared on the same dimensionless basis. A small systematic offset is nevertheless visible at early times ($\bar{t} \lesssim 0.5$), where the initial perturbation has not yet relaxed onto a pure growing eigenmode. This transient phase contains non-modal contributions, explaining the departure from the ideal exponential behaviour. This validates both the single-fluid dispersion relation and the consistency of the height diagnostic across spatial scales. At later times, the curves naturally deviate from a common exponential as non-linear effects develop, as expected when the interface displacement becomes comparable to the mode scale.

We then perform the same validation in the bi-fluid model in a representative intermediate-coupling regime (IC), where ion–neutral interactions are significant but do not trivially enforce a single effective fluid. The same three wavelengths are injected with the same dimensionless initial amplitude h_0/λ, and time is normalised by the theoretical bi-fluid growth rate $\omega_{\rm th}$ for each wavelength and coupling strength. Figure 2b shows that the curves again collapse tightly in the linear regime, indicating that the bi-fluid dispersion relation correctly predicts the scale dependence of the linear growth and that the numerical coupling implementation reproduces the expected exponential stage. As in the single-fluid case, a slight early-time deviation from the theoretical exponential curve is observed, reflecting the finite-amplitude initialisation and the adjustment toward the linear eigenmode. The exponential slopes, however, remain consistent with theory over the identified linear interval. For completeness, we note that in strictly uncoupled bi-fluid runs ($\nu_{\rm nc} = 0$) one may

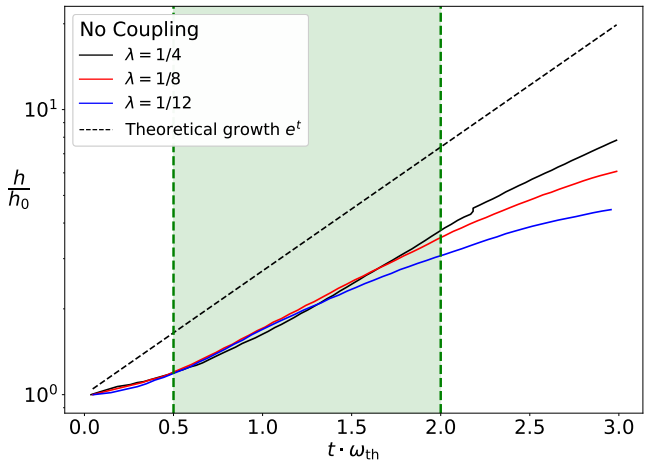

(a) NC

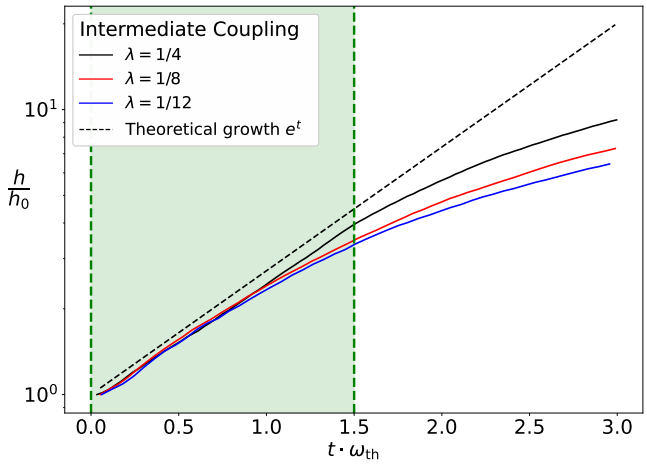

(b) IC

Fig. 2. Linear stage. Normalised finger–bubble height h/h_0 as a function of the reduced time $\bar{t} = \omega_{th} t$ for three injected wavelengths. The curves show the numerical results, and the solid green line corresponds to the theoretical exponential growth. The shaded green region indicates the time interval identified as the linear growth regime. Top panel: no-coupling (NC) case. Bottom panel: intermediate-coupling (IC) case.

observe a small systematic early-time offset even after normalisation by ω_{th}; since the exponential slopes remain correct, this effect is best interpreted as a transient adjustment of the initial conditions rather than as an error in the growth rate itself. Importantly, this delay is absent in the IC configuration shown here, suggesting that even moderate coupling suppresses such early relaxation effects. Overall, the multi-wavelength collapse in both figures provides a stringent validation of the linear theory and of the diagnostic pipeline used throughout the remainder of the paper.

The results show that the linear growth in our simulations is well reproduced by the incompressible dispersion relation of Paper I, within the expected limitations from compressibility and finite-amplitude initialisation. Growth rates and mode ordering are consistently recovered across wavelengths and coupling regimes. Small early-time deviations are due to transient adjustment effects rather than inaccuracies in the theoretical growth rates. This validates the numerical set-up and ensures that the deviations discussed later arise from non-linear and ambipolar diffusion effects.

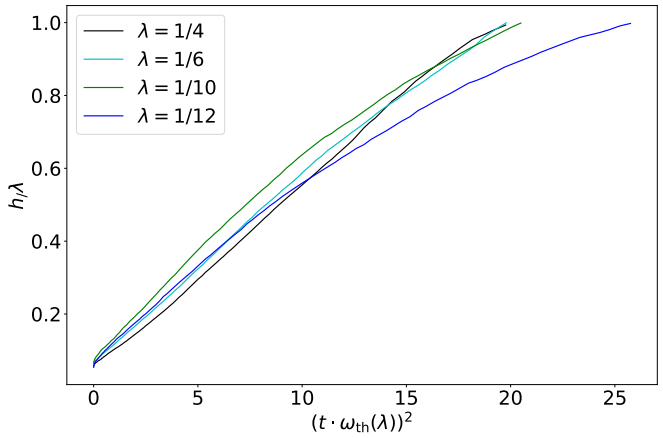

Fig. 3. Quadratic scaling of the mixing height in the single-fluid case ($\nu_{nc} = 0$). Each colour represents a different value of the perturbation wavelength.

4. Non-linear analysis

We now examine the non-linear evolution of the Rayleigh–Taylor instability and the role of ambipolar diffusion beyond the linear regime. The analysis starts from single-mode perturbations to isolate non-linear effects, and is then extended to multi-wavelength configurations in both hydrodynamic and magnetised regimes.

4.1. One wavelength

In the non-linear stage of the Rayleigh–Taylor instability, the growth of the mixing layer is often described by the quadratic law (Sharp 1984; Dimonte et al. 2004)

$$h(t) \approx \alpha \mathcal{A} g t^2, \tag{18}$$

where $\mathcal{A}$ is the Atwood number, g the gravitational acceleration, and α a dimensionless coefficient characterising the effective non-linear growth. This scaling is commonly associated with a self-similar regime in which the dynamics is governed by a single macroscopic length scale $h(t)$ and a characteristic velocity $U \sim \dot{h}$, while the detailed flow physics is absorbed into the value of α. In the following, this quadratic scaling is therefore used as a reference framework to assess how ambipolar diffusion modifies the non-linear evolution, rather than as an a priori universal growth law.

We first consider the uncoupled case ($\nu_{nc} = 0$), which provides a reference single-fluid behaviour. Figure 3 shows the dimensionless mixing height h/λ as a function of the rescaled time $t\,\omega_{th}$, where ω_{th} is the linear growth rate predicted by the two-fluid linear analysis. For all considered wavelengths, the curves are approximately linear, confirming the quadratic growth of the mixing layer, and their slopes remain comparable over a broad range of wavelengths, in agreement with previous numerical studies. At the shortest wavelengths, a tendency toward saturation is observed at small scales, which can be attributed to enhanced dissipation and finite resolution effects. Overall, in the absence of coupling, the quadratic model provides an accurate and robust description of the non-linear growth in our 2D simulations. Before turning to the coupled regimes, it is important to stress that the quadratic model in Eq. (18) should be regarded as an effective description rather than a strict law, especially in two-dimensional configurations. As emphasised by

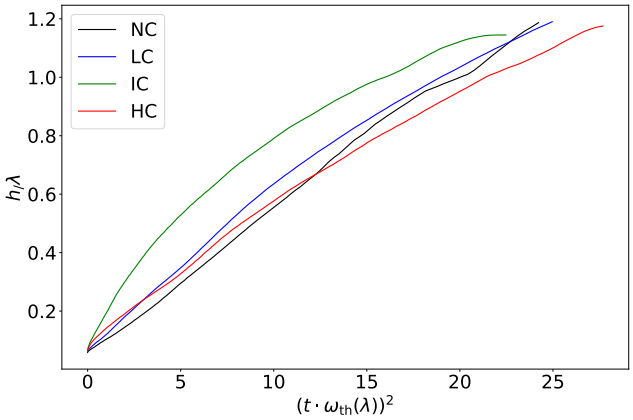

(a) Simulation

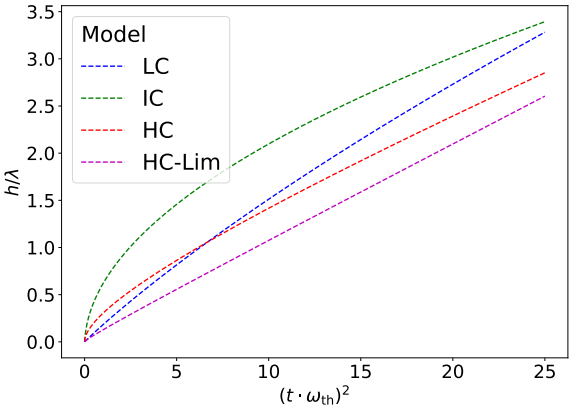

(b) Model

Fig. 4. Quadratic scaling of the mixing height for the No-coupling, Low-coupling, Intermediate-coupling, and High-coupling cases (NC, LC, IC, HC) for $\lambda = 1/4L_x$. Top Panel: simulation result. Bottom panel: theoretical evolution, predicted by the local two-fluid simplified model.

Kalluri & Hillier (2025), self-similarity in Rayleigh–Taylor mixing does not necessarily imply an exactly quadratic growth at all times, but rather the emergence of a regime controlled by a single macroscopic length scale, in which suitably defined dimensionless coefficients remain approximately constant. In 2D, the coexistence of direct and inverse cascades and the persistence of large-scale coherent structures make this regime intrinsically less robust than in three dimensions, so that temporal modulations of the effective growth rate are expected even in the absence of additional physics. In this context, the parameter α should be interpreted as a time-averaged measure of the non-linear growth efficiency, whose applicability may be limited when additional mechanisms interfere with the inertial-buoyancy balance.

When ambipolar coupling is introduced, the non-linear evolution departs more profoundly from the quadratic reference. As shown in Fig. 4, the ambipolar regime is characterised by a pronounced curvature of the growth curves, indicating a sub-quadratic evolution over an extended time interval. The mixing height therefore grows more slowly than predicted by Eq. (18), reflecting a strong reduction of the instantaneous growth rate compared to the uncoupled case. At later times, the curves tend to straighten again, suggesting a gradual recovery of a quasi-quadratic behaviour, but with a significantly reduced effective coefficient. This behaviour shows that ambipolar diffusion does not simply renormalise the quadratic law through a smaller constant, but instead induces a genuinely time-dependent growth process. The failure of the quadratic description in the ambipolar regime can be understood from a simple physical argument. In the classical single-fluid picture, the buoyancy power injected into the flow is assumed to be directly converted into bulk kinetic energy at the scale of the mixing layer, so that the non-linear dynamics can be closed in terms of the single macroscopic length scale $h(t)$. In a partially ionised medium, however, buoyancy initially accelerates only the charged component, while neutrals are entrained progressively through ion–neutral collisions. During this phase, a finite slip velocity develops and a non-negligible fraction of the injected buoyancy power is transferred into relative ion–neutral motion and dissipated by drag on a characteristic coupling timescale. As a result, the instantaneous growth rate of the mixing layer depends not only on $h(t)$, but also on the time-dependent state of coupling between the two fluids, preventing the closure of the growth law in terms of h alone.

To interpret the departure from a purely quadratic growth observed in Fig. 4a, we consider a minimal two-fluid model in which buoyancy acts only on the charged component, while ions and neutrals exchange momentum through linear ion–neutral drag. The model assumes spatially averaged velocities, a constant density ratio $\chi \equiv \rho_c/\rho_n$, and neglects pressure gradients, turbulent transport, and spatial intermittency. Its purpose is not to provide a quantitative description of the non-linear Rayleigh–Taylor instability, but to isolate the dynamical consequences of time-dependent ion–neutral coupling.

This model does not describe the detailed evolution of the interface, but rather a spatially averaged bulk dynamics of the two fluids. It should be understood as a minimal Lagrangian description, in which the characteristic velocity is identified with the large-scale growth of the mixing layer:

$$\rho_c \frac{\mathrm{d}U_c}{\mathrm{d}t} = \mathcal{A}g\,\rho_c - \rho_c \nu_{\mathrm{nc}}(U_c - U_n), \tag{19}$$

$$\rho_n \frac{\mathrm{d}U_n}{\mathrm{d}t} = \rho_c \nu_{\mathrm{nc}}(U_c - U_n), \tag{20}$$

where ν_{nc} is the neutral-on-charge collision frequency and $\Delta u \equiv U_c - U_n$ is the slip velocity. Subtracting the two equations yields a closed evolution equation for the drift:

$$\frac{\mathrm{d}\Delta u}{\mathrm{d}t} = \mathcal{A}g - \nu_{\mathrm{nc}}(1+\chi)\Delta u, \tag{21}$$

which relaxes exponentially on the coupling timescale $\tau_c = [\nu_{\mathrm{nc}}(1+\chi)]^{-1}$. Summing the momentum equations gives the total momentum balance,

$$\frac{\mathrm{d}}{\mathrm{d}t}(\rho_c U_c + \rho_n U_n) = \mathcal{A}g\,\rho_c \tag{22}$$

from which the center-of-mass velocity grows linearly as $U = \mathcal{A}g(\rho_c/\rho_{\mathrm{tot}})\,t$.

Combining these two results, the charged velocity can be written explicitly as the sum of a center-of-mass contribution and a transient drift term. Identifying the charged bulk velocity with the growth rate of the mixing layer, $U_c \simeq \dot{h}$, one obtains an analytical expression for the mixing height:

$$h(t) - h_0 = \frac{\mathcal{A}g\chi}{2(1+\chi)}\,t^2 + \frac{\mathcal{A}g}{\nu_{\mathrm{nc}}^2(1+\chi)^3}\left[(1+\chi)\nu_{\mathrm{nc}}t - (1 - e^{-(1+\chi)\nu_{\mathrm{nc}}t})\right]. \tag{23}$$

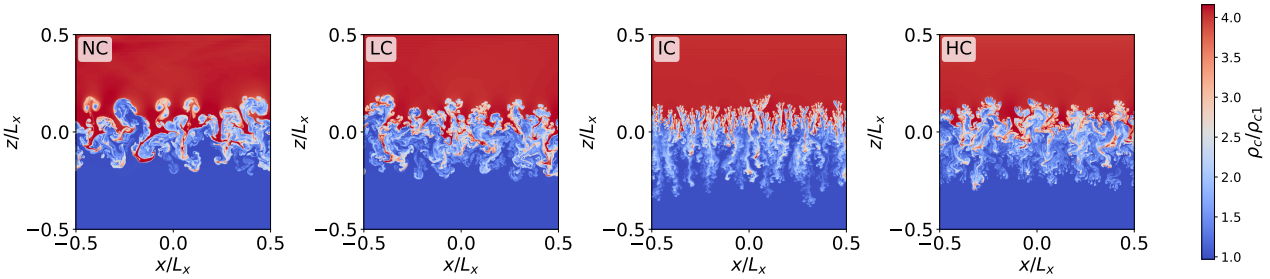

Fig. 5. Snapshots of the charge density ρ_c in hydrodynamic simulations with multi-wavelength initial perturbations, extracted at an equivalent non-linear stage defined by a fixed normalised mixing layer thickness, $\Delta h/L_x \simeq 0.35$. From left to right, the panels show increasing coupling strength ($\alpha = 0$, weak, intermediate, and strong coupling). All snapshots are shown using the same spatial window and colour scale. The density of the charge fluid ρ_c is given in units of the ρ_{c1}, the lighter fluid.

Introducing the reduced time $\tilde{t} \equiv \nu_{\rm nc} t$, the mixing height can be rewritten in dimensionless form as

$$h\left(\frac{\tilde{t}}{\nu_{\rm nc}}\right) - h_0 = \frac{\mathcal{A}g}{\nu_{\rm nc}^2}\left[\frac{\chi}{2(1+\chi)}\,\tilde{t}^2 + \frac{(1+\chi)\tilde{t} - (1-e^{-(1+\chi)\tilde{t}})}{(1+\chi)^3}\right]. \tag{24}$$

In this form, the growth law depends explicitly only on the reduced elapsed time $\nu_{\rm nc} t$, making clear that the non-linear evolution is governed by the relative ordering of the buoyancy timescale and the ion–neutral coupling time.

At early times, $\nu_{\rm nc} t \ll 1$, ions and neutrals are effectively decoupled and the charged component accelerates freely, recovering the classical quadratic growth $h \propto \mathcal{A}g\, t^2$. In the opposite limit, $\nu_{\rm nc} t \gg 1$, ions and neutrals are locked together and the system behaves as a single fluid of total inertia, yielding $h \propto \mathcal{A}g(\rho_c/\rho_{\rm tot})\, t^2$. The intermediate regime $\nu_{\rm nc} t \sim 1$ corresponds to a transient phase during which neutrals are progressively entrained by drag, leading to a time-dependent increase of the effective inertia and a reduced instantaneous growth rate.

This behaviour is illustrated in Fig. 4b, where the theoretical prediction for h/λ is plotted as a function of $(t \cdot \omega_{\rm th})^2$ for weakly (LC), intermediately (IC), and strongly (HC) coupled cases. The model reproduces the pronounced curvature of the IC curve, as well as the reduced asymptotic prefactor in the HC limit, in close qualitative agreement with the simulations. The remaining offset in the normalisation of h/λ reflects the fact that the present model does not include the phenomenological efficiency factor usually introduced in self-similar RT growth, such that $h = \alpha\, \mathcal{A}g\, t^2$ with $\alpha < 1$. As a result, the model correctly captures the temporal evolution but not the absolute normalisation of the mixing height. Importantly, the intermediate-coupling regime does not represent a breakdown of self-similar behaviour, but rather a modified form of self-similarity in which the growth law contains both quadratic and linear contributions. The linear term reflects the finite time required for momentum transfer between the two fluids and the associated evolution of the effective inertia of the mixing layer. The observed sub-quadratic growth therefore arises naturally from the intrinsic two-fluid dynamics, without invoking additional non-linear saturation mechanisms.

4.2. Multi-wavelength regime: Hydrodynamic reference

In the multi-wavelength configurations, the interface perturbation is constructed as a broadband superposition of modes,

$$\eta(x) \propto \eta_0 \sum_{i=n_{\rm osc}}^{N+n_{\rm osc}} c_i \sin\left(\frac{2\pi i}{L_x} x + \phi_i\right), \tag{25}$$

where N is the number of excited modes and $n_{\rm osc}$ sets the largest injected wavelength. The coefficients c_i and phases ϕ_i are drawn from uniform random distributions, with $c_i \in [-1, 1]$ and $\phi_i \in [0, 2\pi[$, and are fixed across all simulations to ensure reproducibility. We first consider the hydrodynamic configurations as a reference for the non-linear evolution under multi-wavelength perturbations. All comparisons are performed at an equivalent non-linear stage defined by a fixed normalised mixing layer thickness, $\Delta h/L_x \simeq 0.35$.

Figure 5 illustrates the non-linear morphology of the mixing layer in the hydrodynamic multi-wavelength simulations at a fixed normalised thickness, $\Delta h/L_x \simeq 0.35$, it shows snapshots of the charge density ρ_c for increasing coupling strength. In the uncoupled case ($\alpha = 0$), the interface is dominated by a small number of large-scale structures. The initially broadband perturbation rapidly evolves toward coherent bubbles and fingers that undergo efficient lateral merging, leading to extended plumes and a relatively smooth mixing layer. As the coupling strength increases, the morphology is progressively modified. For weak coupling, the global organisation remains close to the uncoupled reference, although small-scale corrugations and secondary structures become more visible along the sides of the plumes. In contrast, the intermediate-coupling regime exhibits a markedly different behaviour: at the same mixing height, the interface is significantly more fragmented, with numerous thin fingers and strongly reduced lateral coalescence. Large-scale plumes fail to emerge, and the mixing region retains a highly structured appearance down to smaller spatial scales. In the strongly coupled regime, the interface recovers a smoother and more coherent morphology, with reduced small-scale fragmentation and a renewed tendency toward large-scale organisation. To quantify these morphological differences, we compute one-dimensional power spectra of the charge density ρ_c along the horizontal direction x, evaluated at the same normalised mixing layer thickness. The spectra are obtained using a Welch-averaged estimator (Welch 1967) and normalised by their maximum value to highlight the relative distribution of power across scales.

Figure 6 shows that the redistribution of structures observed in real space is accompanied by clear, systematic changes in spectral content. In the uncoupled case, the power is predominantly concentrated at small wavenumbers, consistent with the dominance of large-scale plumes. As the coupling strength increases, power is progressively shifted toward larger horizontal wavenumbers (i.e. k_x), indicating an enhanced contribution from smaller spatial scales at the same mixing height. This redistribution is most pronounced in the intermediate-coupling regime, which exhibits the strongest relative enhancement of high-k_x power. In the strongly coupled case, the spectrum departs less markedly from the uncoupled reference, consistent with the partial recovery of large-scale organisation seen in the snapshots.

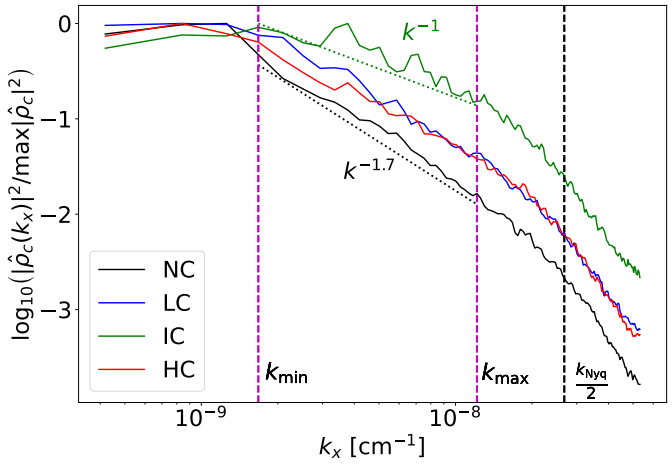

Fig. 6. One-dimensional power spectrum of the charge density ρ_c along the horizontal direction x, computed for the hydrodynamic multi-wavelength simulations at a fixed non-linear stage defined by $\Delta h/L_x \simeq 0.35$. The spectra correspond to $P_{\rho_c}(k_x) = \langle|\hat{\rho}_c(k_x)|^2\rangle$ and are normalised by their maximum value. The vertical dashed lines indicate the minimum and maximum injected wavenumbers, as well as half the Nyquist limit.

Taken together, the snapshots and density spectra demonstrate that ion–neutral coupling has a non-monotonic impact on the non-linear, multi-scale organisation of the mixing layer. At fixed $\Delta h/L_x$, intermediate coupling is the most effective at inhibiting large-scale coalescence and maintaining power at smaller horizontal scales, while both weak and strong coupling favor the emergence of larger-scale structures. This hydrodynamic reference establishes that ambipolar coupling can fundamentally reorganise the non-linear cascade of structures even in the absence of magnetic forces, providing a baseline for interpreting the magnetised configurations discussed below.

4.3. Multi-wavelength MHD regime

We now shift to 2D MHD simulations with multi-wavelengths. Figure 7 shows snapshots of the charge density ρ_c for the magnetised simulations at an equivalent non-linear stage, $\Delta h/L_x \simeq 0.35$, for increasing values of the coupling parameter α.

At $\Delta h/L_x \simeq 0.35$, all magnetised simulations exhibit a fully developed mixing layer whose overall morphology differs markedly from the hydrodynamic reference at the same non-linear stage. In all MHD cases, the interface appears globally smoother, with a reduced level of small-scale corrugations and fewer secondary structures, consistent with the suppression of short-wavelength features by magnetic tension. In the uncoupled MHD case ($\alpha = 0$), the mixing layer is organised into well-defined bubbles and fingers, but the interface retains a noticeable degree of fine structure. Small protrusions and lateral modulations are present along the sides of the main plumes, indicating that, in the absence of ion–neutral coupling, the charged component can still sustain a significant level of small-scale variability despite the presence of the magnetic field. For weak coupling, the overall appearance of the mixing layer remains similar, but the interface becomes slightly more regular. The fingers are more uniform in shape, and the smallest-scale features observed in the uncoupled case are less frequent, suggesting a partial damping of secondary structures as the coupling increases. In the intermediate-coupling regime ($\alpha = 7.5 \times 10^{20}$), the morphology becomes noticeably more coherent. The dominant bubbles are comparatively straight and elongated, with smooth lateral boundaries and very limited internal substructure. The interface is dominated by a small number of clean, well-defined plumes, and fine-scale corrugations are strongly reduced. In the strongly coupled case, the mixing layer remains globally similar in extent, but the interface recovers a degree of complexity. While large-scale bubbles persist, thin fingers and localised distortions reappear along the interface, leading to a more intricate and less uniform morphology than in the intermediate case. This indicates that the influence of coupling on the non-linear organisation of the flow is not monotonic with α.

To further quantify the spatial organisation of the flow, we compute the one-dimensional power spectra of the charge density along the horizontal direction at the same non-linear stage, $\Delta h/L_x \simeq 0.35$. The spectra, shown in Fig. 8, are obtained using a Welch-averaged estimator applied within the mixing layer. Despite the morphological differences observed in real space, the spectra exhibit no significant dependence on the coupling strength over the range of resolved scales. In particular, both the spectral slope and the relative distribution of power across horizontal wavenumbers remain remarkably similar between the different MHD runs. This indicates that, at this advanced stage of the non-linear evolution, the global statistical distribution of power in the charge density is largely insensitive to the value of α, even though the geometry of individual bubbles and fingers continues to vary.

To identify the mechanisms governing the non-linear reorganisation of the mixing layer in the magnetised simulations, we examine spatial maps of the dominant force densities at a fixed non-linear stage, $\Delta h/L_x \simeq 0.35$. Figure 9 shows, for each coupling regime, the Lorentz force acting on the charged component and the ion–neutral drag force acting on the neutrals, computed from the instantaneous fields.

The top row displays the magnitude of the Lorentz force density,

$$\mathbf{f}_L = \frac{(\nabla \times \mathbf{B}) \times \mathbf{B}}{4\pi},$$

together with magnetic field lines (cgs units). The Lorentz force is strongly localised along the mixing interface, with enhanced amplitudes near bubble and finger edges and in regions of strong field-line curvature. This localisation indicates a tension-dominated response, acting through thin layers that redistribute momentum along the interface rather than throughout the volume. The bottom row shows the magnitude of the ion–neutral drag force, $\mathbf{R}_n = \alpha\, \rho_c \rho_n (\mathbf{v}_c - \mathbf{v}_n)$, with streamlines of the relative velocity $\Delta\mathbf{v} = \mathbf{v}_c - \mathbf{v}_n$. In contrast to $\mathbf{f}_L$, the drag force extends over a larger fraction of the mixing region and highlights zones of strong inter-fluid slip, particularly in compressive and shear regions around plume edges. These maps reveal a clear separation of roles: magnetic stresses are localised and associated with field-line curvature, while ion–neutral drag is more spatially extended and governs dissipation through drift. This does not imply a local balance between the two forces, as pressure gradients and inertial terms remain dominant in the momentum equation. The Lorentz force also contributes to energy exchange via the work term $\mathbf{v}_c \cdot \mathbf{f}_L$, enabling local conversion of magnetic energy into kinetic energy.

To complement the local force maps, we now examine the global energy pathways that govern the non-linear evolution of the mixing layer. Figure 10 shows the cumulative gravitational energy injection $E_g(t) = \int_0^t \langle P_g \rangle \, \mathrm{d}t$, with $P_g = \rho_c v_{c,z} g$, together with the cumulative ion–neutral drag dissipation

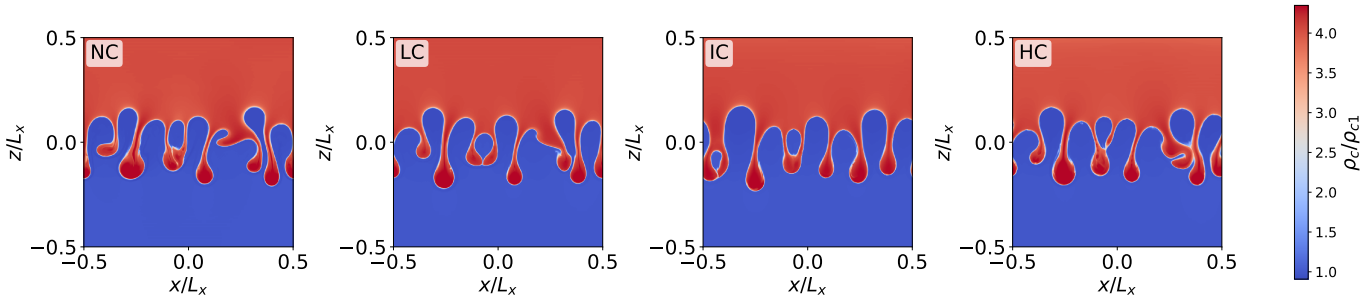

Fig. 7. Snapshots of the charge density ρ_c in MHD simulations with multi-wavelength initial perturbations, extracted at an equivalent non-linear stage defined by a fixed normalised mixing layer thickness, $\Delta h/L_x \simeq 0.35$. From left to right, the panels correspond to increasing coupling strength: $\alpha = 0$, weak coupling, intermediate coupling, and strong coupling. All snapshots are shown using the same spatial window and colour scale. The density of the charge fluid ρ_c is given in units of the ρ_{c1}, the lighter fluid.

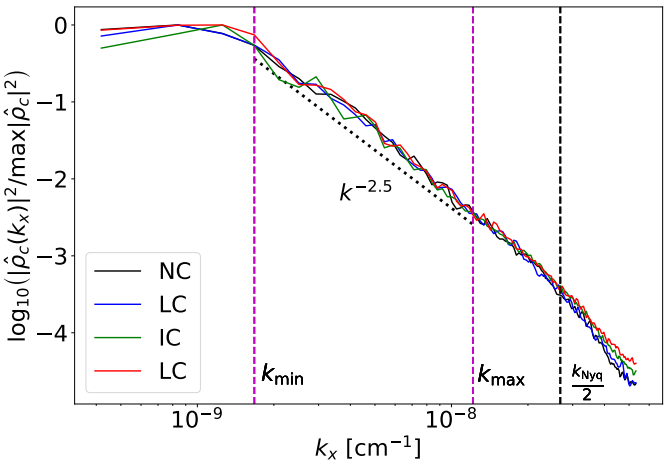

Fig. 8. One-dimensional power spectrum of the charge density ρ_c along the horizontal direction x, computed for the magnetised multi-wavelength simulations at an equivalent non-linear stage defined by $\Delta h/L_x \simeq 0.35$. All spectra were obtained using the same spectral procedure and are normalised by their maximum value. The different curves correspond to increasing ion–neutral coupling strength and are labelled NC, LC, IC, and HC (see Sect. 3). The vertical dashed lines indicate the minimum and maximum injected wavenumbers, as well as half the Nyquist limit.

$E_{\rm kin}(t) = \int_0^t \langle D_{\rm kin}\rangle \, {\rm d}t$, where $D_{\rm kin} = \alpha \rho_c \rho_n |\mathbf{v}_c - \mathbf{v}_n|^2$, both represented as functions of the normalised mixing layer thickness $\Delta h/L_x$. A robust result is that the Intermediate-coupling (IC) regime stands out energetically: at a given non-linear stage, IC simultaneously reaches a maximum of drag-related dissipation and a minimum of gravitational energy content. This behaviour is physically expected, since the gravitational source term scales with the charged velocity, $P_g \propto v_{c,z}$, so that efficient momentum exchange with neutrals reduces the charged flow speed and limits the gravitational energy retained in the charged component, while maximising the relative slip between the two fluids and therefore the drag dissipation. As a result, the gravitational energy injected into the system is converted most efficiently into ion–neutral drift in the IC regime. In contrast, the weakly coupled (LC) and strongly coupled (HC) regimes display comparable global energetic behaviour, both characterised by lower dissipation efficiencies. In the LC case, the charged fluid evolves almost independently of the neutrals, limiting the development of significant inter-fluid drift, whereas in the HC limit the two fluids are nearly locked together, suppressing relative motion and reducing drag dissipation. The HC-limit case further exhibits the largest gravitational energy content, consistent with the persistence of higher charged velocities when slip is strongly inhibited. It is important to emphasise that these simulations are performed in an open vertical configuration, so that the total energy budget is not strictly closed and energy can enter or leave the domain through the boundaries. For this reason, the present analysis does not aim at establishing a strict conservation law, but rather at comparing how gravitational energy is redistributed between physical channels across coupling regimes. In addition, although not shown explicitly in the cumulative diagnostics, the magnetic channel contributes to the energy redistribution through the work of the Lorentz force, $\mathbf{v}_c \cdot \mathbf{f}_L$, which transfers energy between magnetic stresses and bulk kinetic motions. This contribution is spatially localised, as discussed in Fig. 9, and complements the dominant drag-mediated dissipation channel. From this perspective, the global trends reinforce the interpretation suggested by the local force maps: the non-linear organisation of the mixing layer is controlled not by the total amount of injected gravitational energy, but by how efficiently this energy is redirected into ion–neutral drift versus magnetic stresses.

We deliberately limit the analysis to this stage of the non-linear evolution and do not attempt to characterise a late-time saturated regime. In magnetised Rayleigh–Taylor configurations, saturation does not necessarily coincide with the onset of fully developed turbulence, as is often assumed in purely hydrodynamic studies, but may instead result from a geometry-driven or force balance-driven arrest that depends on the wavelength content and on the horizontal magnetic component B_x. Moreover, because the vertical boundaries act as open reservoirs, the global energy budget is not strictly closed, so that extending the energy-based analysis to later times would require a dedicated control-volume approach tied to the mixing layer thickness or fully three-dimensional simulations with carefully designed boundary conditions. For these reasons, we focus here on the physically robust regime where the non-linear organisation of the mixing layer can be unambiguously related to the partition of injected gravitational power between ion–neutral drift and magnetic stresses.

5. Conclusion

In this work, we have investigated the non-linear development of the Rayleigh–Taylor instability in a partially ionised plasma, with a particular emphasis on the role of ambipolar diffusion in the presence of an oblique magnetic field. Building on

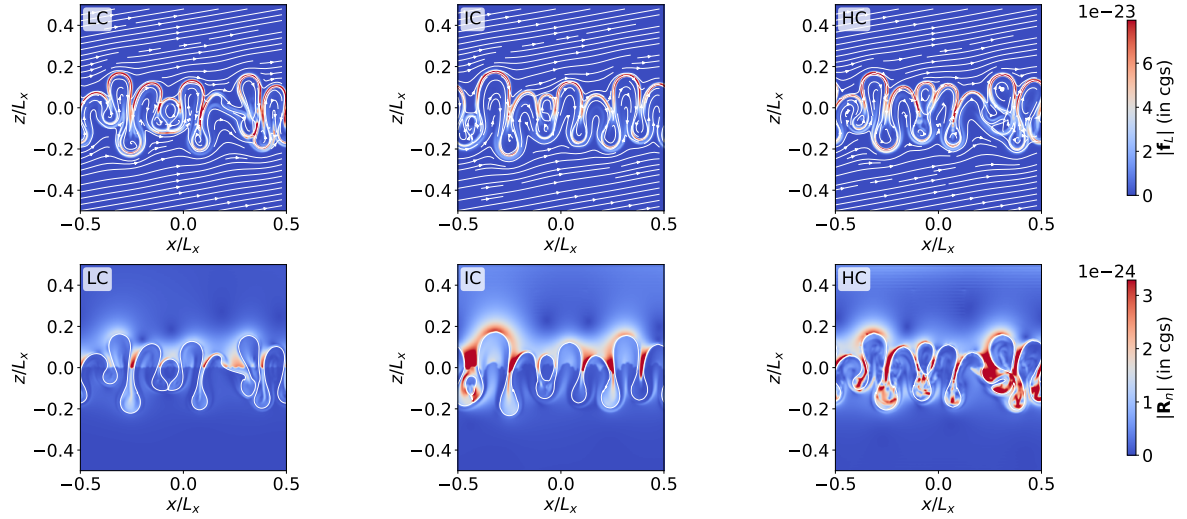

Fig. 9. Spatial organisation of force densities in the magnetised multi-wavelength simulations at a fixed non-linear stage $\Delta h/L_x \simeq 0.35$. Top row: magnitude of the Lorentz force acting on the charged fluid, $|\mathbf{f}_L| = |(\nabla \times \mathbf{B}) \times \mathbf{B}|/4\pi$, with magnetic field lines overlaid. Bottom row: magnitude of the ion–neutral drag force, $|\mathbf{R}_n| = \alpha \rho_c \rho_n |\mathbf{v}_c - \mathbf{v}_n|$, with streamlines of the relative velocity $\Delta\mathbf{v} = \mathbf{v}_c - \mathbf{v}_n$. The white contours delineate the mixing layer. From left to right, the coupling strength increases while the spatial window and colour scales are kept identical within each row. Both forces are expressed in cgs units (dyn cm^{-3}).

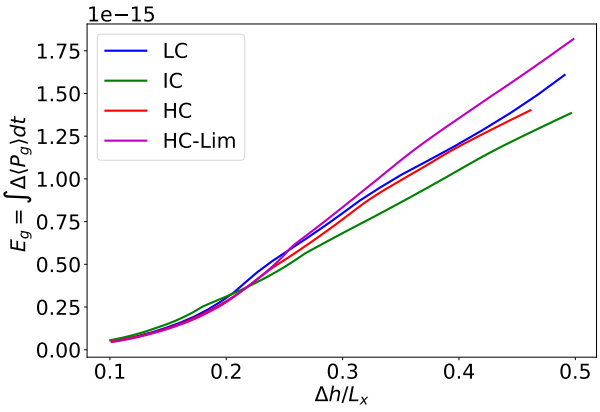

(a) Cumulative gravitational energy injection $E_g(t) = \int_0^t \langle P_g \rangle \, dt$.

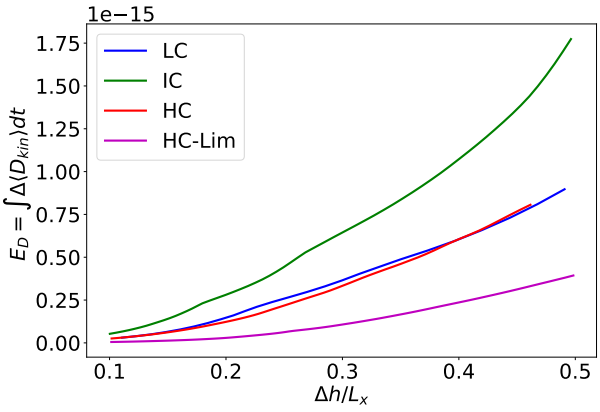

(b) Cumulative ion–neutral drag dissipation $E_{\rm kin}(t) = \int_0^t \langle D_{\rm kin} \rangle \, dt$.

Fig. 10. Global energy pathways as functions of the normalised mixing layer thickness $\Delta h/L_x$ for different ion–neutral coupling strengths. The intermediate-coupling regime maximises the conversion of gravitational energy into drag-related dissipation, while both weak and strong coupling lead to reduced dissipation efficiency.

the linear analysis presented in Paper I, we combined high-resolution two-fluid simulations with self-consistent diagnostics in order to quantify how ion–neutral coupling modifies both the growth and the morphology of the mixing layer beyond the linear regime. We first demonstrated that our numerical set-up accurately reproduces the predictions of the linear theory across hydrodynamic, single-fluid MHD, and bi-fluid configurations. Using a finger–bubble diagnostic, we confirmed that the theoretical growth rates are recovered for a wide range of wavelengths and coupling strengths, and that compressibility effects remain negligible in the parameter regime considered. This validation step provides a robust foundation for the non-linear analysis. In the non-linear regime, we showed that while the classical quadratic scaling of the mixing height, $h(t) \propto Agt^2$, remains a useful global description, it fails to capture important time-dependent features induced by ambipolar diffusion. By introducing a morphology-based alignment procedure based on a fixed normalised mixing height h/λ, we were able to perform meaningful comparisons between simulations with vastly different linear growth rates. This approach revealed a characteristic inflected behaviour of the mixing-layer evolution when ambipolar diffusion becomes dynamically important. In particular, we identified a specific coupling regime in which ambipolar diffusion produces a non-trivial modification of the non-linear dynamics. After synchronisation at an equivalent morphological stage, the evolution exhibits a transient phase of enhanced growth relative to the uncoupled reference, followed by a pronounced slowdown that leads to a reduced mixing efficiency at later times. The ratio of mixing heights between coupled and uncoupled runs asymptotically approaches a value significantly below unity, of order ∼0.7 for the most affected cases. This behaviour demonstrates that ambipolar diffusion does not simply rescale the non-linear growth coefficient, but instead reshapes the temporal evolution of the instability in a genuinely time-dependent manner.

We further showed that these effects persist when the interface is perturbed over a broad range of wavelengths. In the multi-mode regime, ambipolar diffusion strongly modifies the

morphology of the mixing layer, inhibiting large-scale plume coalescence and promoting a more fragmented interface at intermediate coupling. Although one-dimensional power spectra of the charge density at fixed non-linear stage do not reveal a clear change in global scaling laws, they nonetheless confirm that the statistical distribution of power across horizontal wavenumbers remains largely insensitive to coupling. This absence of a strong spectral signature indicates that the observed morphological differences are not associated with a simple spectral cutoff, but rather reflect a local and geometrical reorganisation of the flow. A more direct physical interpretation emerges from the analysis of energy conversion channels. By comparing the relative contributions of gravitational injection, Lorentz work, and ion–neutral drag dissipation at equivalent non-linear stages, we showed that ambipolar diffusion primarily acts by redistributing energy locally within the mixing layer. In particular, the fraction of gravitational power converted into ion–neutral drift reaches a maximum at intermediate coupling, while the global contribution of magnetic stresses remains comparatively modest and weakly dependent on the coupling strength. This non-monotonic behaviour provides a natural explanation for the enhanced smoothness and coherence observed in the intermediate regime, where the drift efficiently extracts energy from the flow without completely suppressing the development of large-scale structures.

Taken together, these results demonstrate that ambipolar diffusion introduces a bounded regime of non-linear behaviour in which the Rayleigh–Taylor instability departs qualitatively from the standard self-similar picture. Rather than acting as a simple effective diffusivity, ambipolar coupling produces an early acceleration followed by a sustained non-linear braking, accompanied by a redistribution of energy between bulk motions, magnetic stresses, and differential ion–neutral drift. The resulting morphology reflects a delicate balance between magnetic tension, which suppresses short-wavelength perturbations, and ambipolar diffusion, which locally reorganises the flow through slip-mediated dissipation. From a methodological standpoint, this work highlights the importance of morphology-based synchronisation when comparing non-linear evolutions across different physical regimes. Aligning simulations at a fixed normalised mixing height provides a robust framework to disentangle genuine physical effects from trivial timing differences associated with distinct linear growth rates. This approach is particularly well suited to multi-fluid systems, where additional timescales associated with coupling can strongly affect the apparent non-linear evolution. Finally, while the present study focuses on idealised two-dimensional configurations, the mechanisms identified here are expected to play an important role in a wide range of astrophysical environments where partially ionised plasmas and magnetic fields coexist, such as molecular clouds, supernova remnants, or the interfaces of expanding H II regions. Extensions to three dimensions, as well as to regimes including additional physical ingredients such as cooling, stratification, or cosmic-ray feedback, constitute natural directions for future work. The results presented here provide a physically grounded reference for such studies and emphasise that ambipolar diffusion can qualitatively alter the non-linear outcome of classical hydromagnetic instabilities.

Acknowledgements. The authors are grateful to K. Ferrière, P. Lesaffre, F. Boulanger, O. Berné for fruitful discussions. We also thank A. Cogez for his contribution during his internship. This work was supported by the Thematic Action "Programme National Physique Stellaire" (PNPS) of INSU Programme National "Astro", with contributions from CNRS Physique & CNRS Chimie, CEA, and CNES.

References

Arons, J., & Lea, S. M. 1976, ApJ, 207, 914
Ascher, U. M., Ruuth, S. J., & Spiteri, R. J. 1997, Appl. Numer. Math., 25, 151
Braileanu, B. P., & Keppens, R. 2022, A&A, 664, A55
Bucciantini, N., Amato, E., & Del Zanna, L. 2004, A&A, 423, 253
Callies, E., Guillet, V., Marcowith, A., Meliani, Z., & Lesaffre, P. 2025, A&A, 698, A147
Carlyle, J., & Hillier, A. 2017, A&A, 605, A21
Chandrasekhar, S. 1961, Hydrodynamic and Hydromagnetic Stability (Oxford: Oxford University Press)
Changmai, M., Jenkins, J. M., Durrive, J. B., & Keppens, R. 2023, A&A, 672, A152
Colella, P., & Woodward, P. R. 1984, J. Comput. Phys., 54, 174
Cook, A. W., Cabot, W. H., & Miller, P. L. 2004, J. Fluid Mech., 511, 333
Díaz, A. J., Soler, R., & Ballester, J. L. 2012, ApJ, 754, 41
Díaz, A. J., Khomenko, E., & Collados, M. 2014, A&A, 564, A97
Dimonte, G., Youngs, D., Dimits, A., et al. 2004, Phys. Fluids, 16, 1668
Fermi, E., & von Neumann, J. 1953, Los Alamos Scientific Laboratory Report
Hester, J. J., et al. 1996, ApJ, 456, 225
Hillier, A. 2016, PhD thesis, University of St Andrews, UK
Isobe, H., Miyagoshi, T., Shibata, K., & Yokoyama, T. 2005, Nature, 434, 478
Jun, B.-I., & Norman, M. L. 1996, ApJ, 465, 800
Jun, B.-I., Norman, M. L., & Stone, J. M. 1995, ApJ, 453, 332
Kalluri, M. T., & Hillier, A. 2025 [arXiv:2510.27053]
Khomenko, E., Vitas, N., Collados, M., & Modestov, M. 2025, A&A, 697, A29
Miyoshi, T., & Kusano, K. 2005, J. Comput. Phys., 208, 315
Popescu Braileanu, B., Lukin, V. S., Khomenko, E., & de Vicente, Á. 2021a, A&A, 646, A93
Popescu Braileanu, B., Lukin, V. S., Khomenko, E., & de Vicente, Á. 2021b, A&A, 650, A181
Popescu Braileanu, B., Lukin, V. S., & Khomenko, E. 2023, A&A, 670, A31
Rayleigh, L. 1883, Proc. London Math. Soc., 14, 170
Ristorcelli, J. R., & Clark, T. T. 2004, J. Fluid Mech., 507, 213
Robinson, K., Dursi, L. J., Ricker, P. M., et al. 2004, ApJ, 601, 621
Ruszkowski, M., Enßlin, T. A., Brüggen, M., & Heinz, S. 2007, MNRAS, 378, 662
Ruuth, S. J., & Spiteri, R. J. 2002, J. Sci. Comput., 17, 211
Sharp, D. H. 1984, Phys. Rep., 101, 221
Soler, R., & Ballester, J. L. 2022, Front. Astron. Space Sci., 9, 789083
Stone, J. M., & Gardiner, T. 2007a, ApJ, 671, 1726
Stone, J. M., & Gardiner, T. A. 2007b, ApJ, 671, 1726
Stone, J. M., & Gardiner, T. A. 2007c, Phys. Fluids, 19, 094104
Taylor, G. I. 1950, Proc. R. Soc. London. Ser. A, 201, 192
Vickers, E., Ballai, I., & Erdélyi, R. 2020, A&A, 634, A96
Vigoureux, G., Flagey, N., Boulanger, F., et al. 2026, A&A, 707, A374
Wang, Y.-M., & Nepveu, M. 1983, A&A, 118, 267
Welch, P. 1967, IEEE Trans. Audio Electroacoustics, 15, 70

Appendix A: Discussion of Paper I

In Paper I, the analytical dispersion relation was derived under the simplifying assumption that the ion–neutral and neutral–ion collision frequencies were identical on both sides of the density interface. While this assumption facilitates the algebraic treatment of the matching conditions, it is not fully consistent with the numerical prescription adopted in MPI–AMRVAC, where the drag coefficient α is conserved and the collision frequencies scale with the local density. In the numerical model, the collision frequencies are defined as

$$\nu_{\rm nc} = \alpha\rho_{\rm n}, \qquad \nu_{\rm cn} = \alpha\rho_{\rm c}, \tag{A.1}$$

so that they vary proportionally to the density of the interacting species. Across an interface separating two homogeneous layers of densities ρ_1 and ρ_2, we define the density contrast:

$$d \equiv \frac{\rho_1}{\rho_2}. \tag{A.2}$$

Because the drag coefficient α is constant, the collision frequencies on either side of the interface scale in the same way, yielding

$$\nu_{\rm nc,2} = d\,\nu_{\rm nc,1}, \qquad \nu_{\rm cn,2} = d\,\nu_{\rm cn,1}. \tag{A.3}$$

To account for this asymmetry in the analytical formulation, we introduce the parameter

$$\xi = \frac{\nu_{\rm nc,1}}{\nu_{\rm nc,2}} = \frac{\nu_{\rm cn,1}}{\nu_{\rm cn,2}}. \tag{A.4}$$

In the physically consistent case corresponding to a constant drag coefficient, this parameter is therefore not free, but directly related to the density contrast as

$$\xi = d. \tag{A.5}$$

The case $\xi = 1$ corresponds to the simplifying assumption adopted in Paper I, where the ion–neutral collision frequency was taken to be uniform across the interface. In that framework, the linear system reduces to the matrix given in Eq. (54), which further simplifies to Eq. (81) when $g_n = 0$, due to additional symmetries arising from the assumption $\rho_c = \rho_n$. In contrast, the case $\xi = d$ considered here reflects the physically consistent variation of the collision frequency with density, as implemented in the numerical model. Because $\rho_c \neq \rho_n$ across the interface, the collision frequencies differ between the two layers, and the simplifications leading to Eq. (81) in Paper I can no longer be applied. As a result, the full matrix structure must be retained, including the additional terms associated with D_1 and D_2. The matrix should therefore be directly compared to Eq. (54) of Paper I, rather than to its simplified form. :

$$M = \begin{pmatrix}
1 & 1 & 0 & -1 & -1 & 0 & 0 & 0 \\
ik_x/k & 0 & 1 & ik_x/k & 0 & -1 & 0 & 0 \\
k & m_{1-} & 0 & k & -m_{2+} & 0 & 0 & 0 \\
k^2 & m_{1-}^2 & 0 & -k^2 & -m_{2+}^2 & 0 & 0 & 0 \\
ik_x & 0 & m_{1-} & -ik_x & 0 & -m_{2+} & 0 & 0 \\
\alpha_{1c} & \beta_{1c} & \gamma_{1c} & -\alpha_{2c} & -\beta_{2c} & -\gamma_{2c} & \delta_{1c} & -\delta_{2c} \\
0 & \psi_{n1} & 0 & 0 & -\psi_{n2} & 0 & 1 & -1 \\
\alpha_{1n} & \beta_{1n} & 0 & -\alpha_{2n} & -\beta_{2n} & 0 & \delta_{1n} & -\delta_{2n}
\end{pmatrix}, \tag{A.6}$$

with

$$\phi_{c1} = \sqrt{\frac{\Omega + \nu_{\rm cn1}}{\Omega + \nu_{\rm cn1} + \nu_{\rm nc1}}} \tag{A.7}$$

$$\phi_{c2} = \sqrt{\frac{\xi\Omega + \nu_{\rm cn1}}{\xi\Omega + \nu_{\rm cn1} + \nu_{\rm nc1}}} \tag{A.8}$$

$$\psi_{n1} = \frac{\nu_{\rm cn1}}{\nu_{\rm cn1} + \Omega} \tag{A.9}$$

$$\psi_{n2} = \frac{\nu_{\rm cn1}}{\nu_{\rm cn1} + \xi\Omega} \tag{A.10}$$

$$m_{1-} = -\left[i\left(\frac{k_x}{\tan\theta}\right) - \frac{\Omega}{\sin\theta\, c_{\rm Ac}\phi_{c1}} \right] \tag{A.11}$$

$$m_{2+} = -\left[i\left(\frac{k_x}{\tan\theta}\right) + \frac{\Omega}{\sin\theta\, \sqrt{d} c_{\rm Ac}\phi_{c2}} \right] \tag{A.12}$$

$$\alpha_{1c} \equiv k_x \left[g_c - \left(c_{\rm Ac1}^2 k \sin^2\theta - \frac{\Omega\,(\nu_{\rm nc1} + \Omega)}{k} \right) \right], \tag{A.13}$$

$$\alpha_{2c} \equiv d^{-1} k_x \left[g_c + \left(d c_{\rm Ac1}^2 k \sin^2\theta - \frac{\Omega\,(\nu_{\rm nc1}/\xi + \Omega)}{k} \right) \right], \tag{A.14}$$

$$\beta_{1c} \equiv k_x g_c\,, \qquad \beta_{2c} \equiv d^{-1} k_x g_c\,, \tag{A.15}$$

$$\gamma_{1c} \equiv i\left(c_{\rm Ac1}^2 m_{1-}^2 \sin^2\theta - \phi_{c1}^{-2}\Omega^2 \right), \tag{A.16}$$

$$\gamma_{2c} \equiv d^{-1} i\left(d c_{\rm Ac1}^2 m_{2+}^2 \sin^2\theta - \phi_{c2}^{-2}\Omega^2 \right), \tag{A.17}$$

$$\delta_{1c} \equiv -k_x \frac{\Omega\nu_{\rm nc1}}{k}\,, \qquad \delta_{2c} \equiv d^{-1} k_x \frac{\Omega\nu_{\rm nc1}/\xi}{k}\,, \tag{A.18}$$

$$\alpha_{1n} \equiv -k_x \frac{\Omega\nu_{\rm cn1}}{k}\,, \tag{A.19}$$

$$\alpha_{2n} \equiv d^{-1} k_x \frac{\Omega\nu_{\rm cn1}/\xi}{k}\,, \tag{A.20}$$

$$\beta_{1n} \equiv k_x \psi_{n1} g_n\,, \qquad \beta_{2n} \equiv d^{-1} k_x \psi_{n2} g_n\,, \tag{A.21}$$

$$\delta_{1n} \equiv k_x \left(g_n + \frac{\Omega(\nu_{\rm cn1} + \Omega)}{k} \right), \tag{A.22}$$

$$\delta_{2n} \equiv d^{-1} k_x \left(g_n - \frac{\Omega(\nu_{\rm cn1}/\xi + \Omega)}{k} \right). \tag{A.23}$$

Allowing the collision frequencies to vary across the interface modifies the analytical formulation in a controlled but non-trivial way. When $\xi \neq 1$, the neutral–ion and ion–neutral collision frequencies explicitly differ on each side of the interface, so that the matching conditions involve distinct dynamical responses in the two layers. As a result, several auxiliary quantities introduced in Paper I, such as the parameters $\phi_{1,2}$ and $\psi_{1,2}$, acquire an explicit dependence on ξ, and the coefficients entering the interface conditions—in particular α_{ic}, β_{ic}, γ_{ic}, δ_{ic} and their neutral counterparts—must be evaluated using different collision frequencies on either side of the interface. While the overall block structure of the matching matrix is preserved, this asymmetry breaks part of the algebraic simplifications exploited in Paper I, where the coupling could be represented by a single effective parameter. In the present formulation, the coupling enters through several coefficients simultaneously, which makes the dispersion relation less amenable to a compact analytical reduction and renders the interpretation of the coupling regimes less direct from the mathematical expression alone. Nevertheless, the physical structure of the problem remains unchanged: the dispersion relation still connects continuously the limits of

vanishing coupling and perfect coupling, and the same asymptotic behaviours are recovered when the collision frequency becomes either negligible or dominant compared to the dynamical timescale.

The practical impact of this modification is illustrated in Fig. A.1, where the linear growth rate is plotted as a function of the wavenumber k for different coupling strengths, comparing the cases $\xi = 1$ (as assumed in Paper I) and $\xi = d$, corresponding to a collision frequency variation driven solely by the density contrast across the interface. The uncoupled and fully coupled asymptotic limits are shown for reference and are identical in both cases, confirming that the modification does not alter the fundamental limiting behaviour of the instability. Differences arise primarily at intermediate wavenumbers, where the growth rate obtained with $\xi = d$ is systematically reduced compared to the $\xi = 1$ case for the same nominal collision frequency. This reduction reflects the asymmetric collision-frequency distribution across the interface: when $\xi = d$, the coupling strength differs between the two layers, which modifies the relative contribution of each fluid to the linear dynamics and enhances the effective damping at intermediate scales. For practical comparison with the formulation of Paper I, it is convenient to introduce an effective collision frequency $\nu_{\mathrm{nc,eq}}$ defined operationally as the value that minimises the mismatch between the $\xi = d$ dispersion curve and the $\xi = 1$ curve over a chosen intermediate-k interval. With this definition, the two formulations can be brought into close quantitative agreement without affecting the uncoupled and perfectly coupled limits. The ratio $\nu_{\mathrm{nc,eq}}/\nu_{\mathrm{nc}}$ is not universal: it depends on the density contrast d and, to a lesser extent, on the coupling regime considered. Therefore, the transition from $\xi = 1$ to $\xi = d$ does not introduce a new qualitative regime of instability, but rather corresponds to a renormalisation of the effective coupling strength when the collision frequencies are allowed to scale consistently with density. The theoretical framework developed in Paper I thus remains fully applicable, provided that collision frequencies are interpreted in terms of an effective coupling when comparing analytical predictions with physically consistent numerical prescriptions.

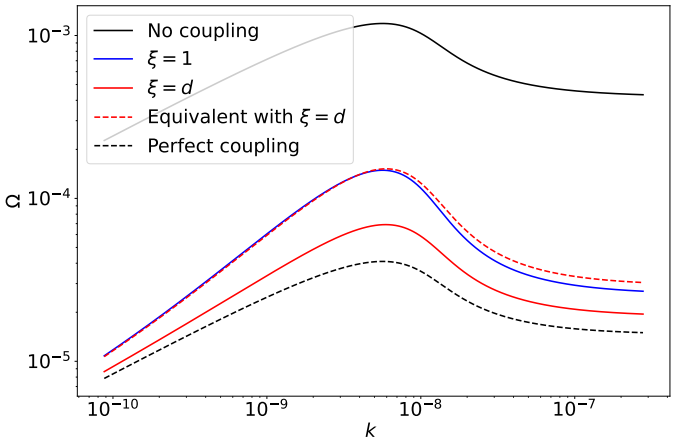

Fig. A.1: Linear growth rate as a function of the wavenumber k. The black curves show the two asymptotic limits: the uncoupled case (No coupling) and the perfectly coupled case (Perfect coupling). The blue curve corresponds to the formulation with $\xi = 1$ used in Paper I, while the red curve shows the physically consistent case $\xi = d$, where the collision frequencies scale with the density contrast across the interface (constant drag coefficient α). The dashed red curve labelled 'Equivalent with $\xi = d$' shows the $\xi = 1$ formulation evaluated with an effective collision frequency $\nu_{\mathrm{nc,eq}}$ chosen so as to reproduce the $\xi = d$ dispersion curve over the intermediate-k range.