Water Waves manuscript No.
(will be inserted by the editor)

The error in predicted phase velocity of surface waves atop a shear current with uncertainty

Peter Maxwell · Benjamin K. Smeltzer · Simen Å. Ellingsen

Received: date / Accepted: date

Abstract The effect of a depth-dependent shear current $U(z)$ on surface wave dispersion is conventionally calculated by assuming $U(z)$ to be an exactly known function, from which the resulting phase velocity $c(k)$ is determined. This, however, is not the situation in reality. Field measurements of the current profile are performed at a finite number of discrete depths and with nonzero experimental uncertainty. Here we analyse how imperfect knowledge of $U(z)$ affects estimates of $c(k)$. We performed a numerical experiment simulating a large number of "measurements" of three different shear currents: an exponential profile, a $1/7$-law profile, and a profile measured in the Columbia River delta. A number of measurement points were specified, the topmost of which at $z = -h_s$ (permitting simulation of measurement points which do not fully extend to the surface), and measurements taken from a normal distribution with standard deviation ΔU. Four different methods of reconstructing a continuous $U(z)$ from the measurements are compared with respect to mean value and variance of $c(k)$. We find that an ordinary least-squares polynomial fit seems robust against mispredicting mean values at the expense of relatively high variance. Its performance is similar for all profiles whereas a fit to an exponential form is excellent in one case and poor in another. A clear conclusion is the need for a measurement of the surface velocity $U(0)$ when there is significant shear near the surface. For the exponential and Columbia profiles alike, errors due to extrapolation of U from $z = -h_s$ to 0 dominate the resulting error of c, especially for shorter wavelengths. In contrast, the error in $c(k)$ decreases slowly with a higher density of measurement points, indicating that better, not more, velocity sensors should be invested in. A pseudospectral analysis of the linear operator corresponding to the three velocity profiles was performed. In all cases, the pseudospectrum shows strong asymmetry around the eigenvalue for c, indicating that a perturbation in the underlying current is more likely to push the value

P. Maxwell · S Å. Ellingsen · B. K. Smeltzer
Department of Energy and Process Engineering,
Norwegian University of Science and Technology,
N-7491 Trondheim, Norway
E-mail: peter.maxwell@ntnu.no

of c to higher, not lower, values. This is in tentative agreement with our observation that for sufficiently large ΔU, c is found to have predominantly positive skewness, although the direct relationship between the two is not altogether obvious.

Keywords Wave-current interaction · Dispersion relation · Waves with vorticity

1 Introduction

A sub-surface shear current can drastically alter the dispersion of surface waves, something that has been recognised for a long time [20]. A number of phenomena occur because of shear that are not present for a depth-uniform current. Some of these effects can be dramatic. Already 45 years ago Dalrymple, studying forces from waves during hurricanes, concluded that "it is obvious ... that rational offshore design must include the effect of [shear] currents" [5]. Only recently has the importance of including sub-surface shear been widely recognised. The effect of shear currents on wave dispersion has now been implemented in ocean models [14,8], and new theory has been developed aimed at modelling, in particular the all-important wave action conservation [1,21].

The strongly sheared tidal current in the Columbia River estuary has been a particular focus of attention for these purposes because the shear is unusually strong and excellent measurement data for the depth-varying velocity is available [12]. Every year thousands of ships traverse these waters, which have been dubbed "the Graveyard of the Pacific" for its dangerous conditions and many shipwrecks. Zippel and Thomson found wave steepness predictions to be off by up to 20% if not accounting for sub-surface shear [30]. In the same current some of us have predicted that wave resistance on small ships can vary by a factor 3 or more between upstream and downstream motion at the same velocity relative to the surface [16], an effect solely due to shear. Methods to reconstruct the sub-surface flow profile from measurements of the dispersion relation have recently been developed [18,4] and tested in the Columbia mouth [3].

What all of these studies have in common though is the study of change in wave behavour due to a horizontal shear current $\mathbf{U}(z)$ that is presumed to be a *known* analytical function. For linear waves, a range of numerical methods [29,6,24,15] and integral approximation schemes [25,22,13,9] exist for calculating the phase velocity $c(\mathbf{k})$ —the dispersion relation— given arbitrary $\mathbf{U}(z)$. However, this is not the situation in practice. Invariably what is available is not a smooth function $\mathbf{U}(z)$, but rather a set of measurements $\mathbf{U}_n$ at a set of discrete depths z_n, each measured with some nonzero uncertainty $\Delta\mathbf{U}_n$. The common procedure is to construct a function $\mathbf{U}_{\text{fit}}(z)$ by fitting some analytical function to these points, and assume this to be the exact velocity profile thereafter. It is pertinent, therefore, to ask how and to what extent imperfect knowledge of the current $\mathbf{U}(z)$ can result in errors in the predicted value of $c(\mathbf{k})$ propagating atop it.

There are several sources of error to account for. The most obvious is the measurement error $\Delta\mathbf{U}$ but often more grievous is the fact that measurements are typically not available for the top few meters of the water column, making it necessary to extrapolate $\mathbf{U}(z)$ from the uppermost measured point to the water surface. This is especially

common for measurements using Acoustic Doppler Current Profilometry (ADCP). Finally, there exists an infinity of options for how to fit a continuous function $\mathbf{U}_{\text{fit}}(z)$ to measurement data and it is in no way obvious to what extent the choice of fitting method can affect the predicted phase velocity.

Naturally, the set of cases and parameters relevant to our question is vast, and we must by necessity restrict ourselves to a few values of a few of the possible parameters. We have tested three different shear profiles typical of real-life applications: an exponential profile with shear restricted to a shear layer near the surface, a profile to represent a bottom boundary layer in a shallow water flow, and an exemplar velocity profile from the aforementioned Columbia River dataset. We consider only unidirectional profiles $\mathbf{U}(z) = U(z)\mathbf{e}_x$ and, reasoning that the measurement error doesn't change appreciably with depth, we make the simplification that the measurement error ΔU is Gaussian and is the same for all depths. We further assume that the depths of the measurement points are exactly known. Waves are presumed to propagate streamwise in the same direction as the current, but see Sec. 2.2 for how results for waves propagating counter to the current can be deduced. We compare some of the most common procedures for fitting curves $U_{\text{fit}}(z)$ to the set $\{z_n, U_n\}$ both for cases where the surface current velocity is and is not included in the data set.

In considering the linear dispersion relation, our analysis is necessarily concerned with linear waves. Clearly, many situations exist where nonlinear effects become imoportant. Whilst it is conceptually possible to extend the analysis to second or higher order in wave steepness, this would vastly complicate our efforts — indeed the necessary theory for Stokes expansions on arbitrary shear currents does not presently exist to our knowledge. Moreover, the question of dispersion becomes complicated, particularly so if a spectrum of wavelengths is present, since sums of eigenfrequencies come into play. Adding to this that our parameter space is already extensive and that such analysis would greatly increase the scope and length of our paper, no nonlinear effects will be considered.

1.1 Problem definition

The geometry of our problem is the same as in Refs. [13,9] and many others. We consider a shear current with profile $U(z)$ whose surface when undisturbed is at $z = 0$ and the bottom is at $z = -h$ where h can be finite or infinite. A surface wave of amplitude a and wave number k propagates on the free surface and is assumed to have small steepness, $ka \ll 1$, so that linear theory can be assumed. We assume inviscid, incompressible flow and neglect effects of surface tension. Due to the wave, the surface elevation is $z = \hat{\eta}(x,t)$ and the velocity field becomes $[U(z) + \hat{u}(x,z,t), \hat{v}(x,z,t), \hat{w}(x,z,t)]$. Hatted quantities are small, and governing equations as well as boundary conditions are linearised with respect to these. The solutions to the flow field are of form

$$[\hat{u}, \hat{v}, \hat{w}](x,t,z) = [u, v, w](z,k)\mathrm{e}^{\mathrm{i}k[x-c(k)]t} \tag{1}$$

where k is the wave number and $c(k)$ is the phase velocity. For reasons of manageability, we shall limit ourselves to unidirectional (2D) systems with waves propagating

along or against the flow direction. Now, let $U_0 = U(0)$ and

$$\tilde{U} = U/U_0;\ \ \tilde{w}(\tilde{z}) = w(z)/w(0);\ \ \tilde{z} = z/d;\ \ \tilde{c} = c/U_0;\ \ \tilde{k} = kd;\ \ \tilde{h} = h/d, \quad (2)$$

where d is a characteristic lengthscale of the velocity profile which we will define for each profile as they are considered. The strength of the current is now described by the Froude number

$$\mathrm{Fr} = U_0/\sqrt{gd} \quad (3)$$

with g being acceleration due to gravity.

The system is governed by the Rayleigh equation, which is readily derived from the Euler and continuity equations (see, e.g. [20]). Imposing the linearised free surface boundary condition and Dirichlet boundary condition at the fluid bottom gives the boundary value problem, in our units,

$$\tilde{w}'' - \frac{\tilde{U}''(\tilde{z})}{\tilde{U}(\tilde{z}) - \tilde{c}}\tilde{w} = \tilde{k}^2\tilde{w}; \quad (4a)$$

$$(1-\tilde{c})^2\tilde{w}' - [\mathrm{Fr}^{-2} + \tilde{U}'(0)(1-\tilde{c})]\tilde{w} = 0;\ \ \tilde{z} = 0; \quad (4b)$$

$$\tilde{w} = 0;\ \ \tilde{z} = -\tilde{h}, \quad (4c)$$

where a prime denotes differentiation with respect to $\tilde{z}$. Note that (2) implies $\tilde{U}(0) = \tilde{w}(0) = 1$.

When $\tilde{U}(z)$ is known, the eigenvalue $\tilde{c}(\tilde{k})$ and eigenfunction $\tilde{w}(\tilde{z};\tilde{k})$ follow from Eqs. (4) and can be calculated with different methods; see Sec. 2 below.

1.2 Outline

In Sec. 2 we outline the stochastic numerical experiments to obtain statistical properties of $\tilde{c}(\tilde{k})$ based on the distributions of measured $\tilde{U}_n$ then present the numerical results and discuss the implications of these in Sec. 3. Sec. 4 presents some further numerical experiment results concerning dependence of deviation on spacing of measurement points. Sec. 5 presents a pseudospectral analysis before conclusions in Sec. 6.

2 Numerical experiments with ensemble of perturbed shear profiles

In this section we describe a numerical experiment to investigate the effects of a number of limitations in realistic data from field measurements of the current profile.

One such limitation is that data is only available in a finite set of points $\{\tilde{z}_n\}$, $n = 1...N$. There is an infinite number of ways in which to reconstruct a continuous velocity profile from experimental measurements yet no obvious 'correct' choice. We seek to characterise the effect of employing different approximation schemes for constructing a continuous $\tilde{U}(\tilde{z})$ by using a heuristic analysis. We do not consider interpolation methods that force $\tilde{U}(\tilde{z})$ to pass exactly through the measured points since these have the unacceptable property of becoming worse for increasing number of measurement points, N.

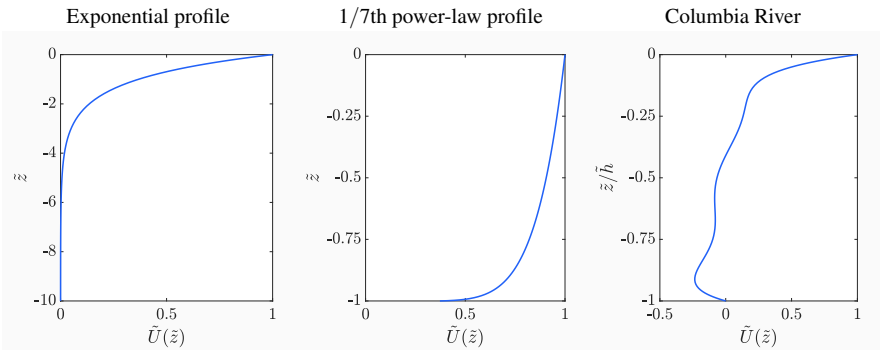

Fig. 1 The three shear profiles studied.

The approach used in this section is to take three candidate velocity profiles $\tilde{U}(z)$ which are known and select a vector of reasonable 'measurement' locations $\{\tilde{z}_n\}$ to represent a field measurement. The exact dispersion relation for a given profile is $\tilde{c}_e(\tilde{k})$. Next, we perform a series of simulations wherein Gaussian random noise is applied to the velocity profile at the measurement points to replicate what a single experimental measurement might return. For each run, we then apply our approximation methods to yield a continuous reconstructed velocity profile $\tilde{U}_{\text{fit}}(\tilde{z})$ and calculate the resulting dispersion relation $\tilde{c}_{\text{fit}}(\tilde{k})$ for a broad range of wave numbers $\tilde{k}$. After a suitable number of runs have completed, we can calculate the componentwise relative error $\Delta\tilde{c}/\tilde{c}_e$, which can in turn be used to characterise the statistical properties of the reconstruction methods. Here $\Delta\tilde{c} = \tilde{c}_{\text{fit}} - \tilde{c}_e$.

To calculate $\tilde{c}_{\text{fit}}(\tilde{k})$ for each $\tilde{k}$ from a function $\tilde{U}_{\text{fit}}(\tilde{z})$, we used a path-following scheme [19] for solving Eqs. (4). The choice of method is in no way essential: one can equally well use any of a range of available methods [6,24,15]. The accuracy of the numerical method has an error of circa 10^{-10}, so is far too small to contribute to any error in our analysis. Even using an analytical approximation such as Refs. [25, 22,9] gives results that are hard to discern from our high-accuracy ones in the graphs we present. Indeed, for a small proportion of the test runs, a critical layer can be encountered. The numerical schemes take an unacceptably long time to compute in this scenario and become less reliable. Therefore, for the small number of $\tilde{U}_{\text{fit}}(\tilde{z})$ fit functions that generate a critical layer, we substitute the $\tilde{c}_{\text{fit}}(\tilde{k})$ result using the integral approximation method from Ref. [9] (computed using Clenshaw–Curtis quadrature). This approach has been tested and any disparity is too small to influence the results in our figures.

Whatever numerical or approximation scheme used, the granularity of the computational grid must be small enough to capture any strong curvature or oscillatory behaviour in the fits. This is particularly important near the surface. We performed tests with both the path-following method and approximation scheme to ensure sufficiently many z computational grid points were used.

We choose three much-used velocity profiles: an exponential profile, a $1/7$th power law profile, and a profile which is based on a set of real current measurements taken in the Columbia River estuary [12]. These are shown in Fig. 1.

Additionally, we select a number of free parameters to characterise the problem:

- $\Delta\tilde{U}$ is the standard deviation of the normal distribution of velocity measurements. We assume the error is the same for all depths and the mean to be equal to the real value at the point of measurement.
- $\Delta\tilde{h}$ is the spacing of the equidistantly spaced measurement points along the $\tilde{z}$ axis.
- $\tilde{h}_s \geq 0$ is the depth at which our topmost measurement sits.

$\tilde{h}_s$ is a particularly relevant parameter because most practical measurement techniques such as ACDP cannot reliably measure current nearest the surface. When the shear is concentrated near the surface such as in wind driven (Ekman) flow, the top few meters of the water column have the greatest effect on surface wave dispersion.

The exponential profile we use is defined as

$$\tilde{U}_{\exp}(\tilde{z}) = \mathrm{e}^{\tilde{z}}. \tag{5}$$

In other words, the characteristic length d is the depth at which $\tilde{U}$ is reduced by $1/\mathrm{e}$. This profile is an oft-assumed model of a wind-driven surface shear layer in the upper ocean.

The power law, or "$1/7$", profile is characteristic of river flow over a smooth bed. To avoid numerical issues due to (unphysical) singular behaviour in the derivatives at $\tilde{z} = -\tilde{h}$, we use the slightly scaled form

$$\tilde{U}_{1/7}(\tilde{z}) = \left((1-0.001)\frac{\tilde{z}}{\tilde{h}} + 1\right)^{1/7}. \tag{6}$$

The third profile is created using a methodology similar to as detailed in Ref. [16]: a suitable exemplar profile is taken from the Columbia River dataset upon which we fit a low-order polynomial. Further detail on this is given in Sec. 3.3.

We compare several standard methods for reconstructing a global approximant from a discrete set of data points:

- Ordinary Least Squares *(OLS)*. A 7th degree polynomial in $\tilde{z}$ is fitted to the set $\{\tilde{z}_n, \tilde{U}_n\}$. We used a Chebyshev basis for numerical stability and performance reasons but in practical testing we found no difference in accuracy compared to a monomial basis of the same degree, e.g. MATLAB's `polyfit()` function.
- Tikhonov Regularised Least Squares *(TLS)*. In particular, we solve the problem in the canonical linear least squares form [10, eqn. 3], $\min\{\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \alpha\|\boldsymbol{L}\boldsymbol{x}\|_2^2\}$, using the SVD method. In our implementation, we chose matrix $\boldsymbol{L}$ as to constrain the first derivative and picked $\alpha = 1$; again this is over a Chebyshev basis of 7th degree.
- Nonlinear Least Squares Exponential Fit *(EXP)*. In our implementation, we used MATLAB's `lsqcurvefit()` with a model function $\tilde{U}(\tilde{z}) = a + be^{c\tilde{z}}$.
- Linear Extrapolation then OLS *(OLS+Extr)*. This method is only applicable when $\tilde{h}_s \neq 0$. We apply a linear extrapolation to the surface using the uppermost 6 measured points, add the extrapolated surface velocity to the set of measurement points, and then finally apply an OLS fit to this new set.

Note that for the OLS, TLS, and EXP schemes, the function fitted to data must be continued (extrapolated) to the surface when $\tilde{h}_s \neq 0$.

For the OLS+Extr method an alternative version would be to do a curve fit of the available points up to depth $\tilde{h}_s$, then continue the fit in a straight line to the surface with continuous derivative, add the resulting surface point to the set of measurements and fit again. This may improve results when the curvature is monotonic such as for the exponential profile. However, in general the curvature is not known *a priori*, and we opine that the chosen method is a fairer comparison.

We also tested a global Bezier curve fit using the canonical De Casteljau's algorithm. Despite obtaining a good fit on visual inspection, it performed very poorly in the simulations so is not further considered in our analysis. The reason for the poor performance is likely to be that the Bezier curve must terminate exactly on the given measurement points at each end of the interval. Therefore, it is highly sensitive to the error in the surface measurement. Also, if the error in the measurements near the surface impart some disparity then it forces strong curvature in the fit, which exacerbates the problem. There may be a more suitable way to use Bezier curves for this purpose but it is not immediately obvious.

An illustration of how the fitting schemes perform is shown in Fig. 2, for the exponential velocity profile and "measurement" points from one numerical run. Note already that the different approximation schemes exhibit different, and in some cases biased, properties.

2.1 Remarks on how a sub-surface shear flow affects $c(k)$

Before going on to the results of our numerical experiment, we will briefly consider some aspects of how the dispersion relation of $c(k)$ is affected by a shear flow $U(z)$ in order to facilitate our later analysis. We assume $k > 0$.

The k-dependent phase velocity can be written in the form

$$c(k) = c_0 + \bar{U}(k) \tag{7}$$

where $\bar{U}(k)$ is a k-dependent effective Doppler shift, and $c_0(k) = \sqrt{(g/k)\tanh kh}$. For flows that are not too strongly sheared (e.g. oceanographic currents), a good approximation is that of Stewart & Joy [25],

$$\bar{U}(k) \approx 2k \int_{-\infty}^{0} U(z) \mathrm{e}^{2kz}\, \mathrm{d}z \tag{8}$$

in infinitely deep water, or the corresponding finite-depth version [22]

$$\bar{U}(k) \approx 2k \int_{-h}^{0} U(z) \frac{\sinh 2k(z+h)}{\cosh 2kh}\, \mathrm{d}z. \tag{9}$$

In all the currents we consider, these approximations are good to within a few percent. In dimensionless units, $\tilde{c}_0 = \mathrm{Fr}^{-1}\tilde{k}^{-1/2}$ or $\tilde{c}_0 = \mathrm{Fr}^{-1}[\tilde{k}^{-1}\tanh(\tilde{k}\tilde{h})]^{1/2}$.

Considering the deep water case, the proportionality to $\exp(2kz)$ reflects the well-known rule-of-thumb that a wave is affected by the water beneath it only down to a

depth of about half a wavelength. Assuming, as for our exponential and Columbia River profiles, that the shear is concentrated near the surface penetrating a distance of order d into the depth, and that the wavelength is long compared to this, $kd \ll 1$, Eq. (8) is a weighted average of $U(z)$ over all of the water column where the shear is appreciable.

In contrast, short waves compared to the shear penetration d only feel a small part of the shear profile nearest the surface. In the limit $k \to \infty$, Eq. (8) and (9) both become

$$\bar{U}(k) \sim U(0) + \frac{U'(0)}{2k} + ... \tag{10}$$

In other words, with an approximate velocity profile $U_{\text{fit}}(z)$ based on a data set, the misprediction of c for very short waves will tend to

$$\Delta c \sim U_{\text{fit}}(0) - U(0), \tag{11}$$

equal to the misprediction of the surface velocity.

Thus, the error in the estimation of $c(k)$ for long waves (small k) is mitigated by the fact that the noise in the measurements $\{U_n\}$ is averaged over the water column. In contrast, the very shortest waves depend only on the accuracy of the measured flow at the surface. When the surface current itself is measured, $\tilde{c}_{\text{fit}}$ will have a standard deviation of order $\Delta\tilde{U}$ in this limit. When measurements start at a nonzero depth $\tilde{h}_s$, however, this error can be much greater, amplified by having to extrapolate the current near the surface.

We must expect therefore that the *absolute* error Δc grows for higher k until it reaches its asymptotic value (11).

2.2 Note on the relative error

We will be plotting the relative error $\Delta\tilde{c}/\tilde{c} = \Delta c/c$ for a number of cases, whose meaning requires some comment.

We constrain ourselves herein to considering positive values of $U(z)$, meaning we consider waves propagating downstream in our frame of reference, which we consistently take to be the 'lab' frame where, in particular, the current tends to zero at the bottom. One should note that while Δc is independent of both these choices, c, and consequently $\Delta c/c$, is not.

Were we to consider instead a wave propagating *against* the current, changing the sign of U turns Eq. (7) into

$$c(k) = c_0 - \bar{U}(k).$$

The relative error, as defined, is amplified in this case, by a factor $(c_0 + \bar{U})/(c_0 - \bar{U})$. To avoid doubling an already large parameter space we do not plot results for the counter-propagating scenario (which will differ only by this factor), but note that we have conservatively plotted the *lower* of the two relative errors.

3 Numerical results

In the following subsections, we perform a sufficiently large number of runs of the numerical experiment and collate the componentwise relative errors $\Delta\tilde{c}(\tilde{k})/\tilde{c}(\tilde{k})$ for $\tilde{k}$ ranging from "long" waves $\tilde{k} \ll 1$ which essentially feel a Doppler shift from a weighted average of $\tilde{U}(z)$ over the full depth-varying profile, to "short" waves $\tilde{k} \gg 1$ which are only affected by the current very close to the surface.

3.1 Exponential profile

The exponential profile, Eq. (5), is a common model for the wind-driven surface shear layer in oceans and lakes. It has its greatest shear near the surface, which will particularly affect the shorter waves (higher $\tilde{k}$). Having to use extrapolation from a depth $\tilde{h}_s$ to the surface can be expected to introduce serious errors in the regime of short waves of wavelength shorter than about $2\tilde{h}_s$, i.e. $\tilde{k}\tilde{h}_s \gtrsim \pi$.

Figs. 3 and 4 illustrate how the relative phase velocity error depends on $\tilde{k}$ for the different approximation schemes and physical parameter combinations. From 10,000 runs of the experiment, we plot the mean value as a function of $\tilde{k}$ as well as the 10% percentiles on either side (i.e. 80% of the area under the distribution curve lies between these lines) to illustrate the development of the width of the profile as well as give some indication of its skewness.

The current in Fig. 3 is weaker ($\mathrm{Fr}^2 = 0.01$) than that of Fig. 4 ($\mathrm{Fr}^2 = 0.05$). When the surface velocity is part of the set of measured points ($\tilde{h}_s = 0$, the two top rows) the results are very close for the two regimes. When $\tilde{h}_s > 0$, however, an extrapolation is necessary and the stronger current widens the distribution of $\tilde{c}$. On the other hand the goodness of the mean value prediction is hardly affected by the change in current strength.

We notice more generally the large difference in variance at high $\tilde{k}$ whether or not the surface velocity is measured. This must be expected since, according to Eq. (11), this simply reflects the error in determining the surface current velocity, an error which is clearly amplified by extrapolation. The OLS scheme consistently avoids bias of the mean itself although the results have the largest variance. The TLS scheme is arguably unfavourable in all cases: its mean value tends to underpredict at short wavelengths (corresponding to a systematic underprediction of the surface velocity), the resulting distribution is badly skewed as seen from the asymmetry of the percentiles, and the variance is also comparatively high.

The exponential fitting scheme performs reasonably well for the exponential current. Despite the expected somewhat unfair advantage of fitting data to a function of the correct functional form in this particular case, the method is no significant improvement over the simpler OLS. Moreover, this scheme has another problem namely that it very occasionally produces estimates for $\tilde{c}$ which are off by many orders of magnitude, so much so that a single outlier out of thousands could shift the mean so much as to lie outside of the 10% percentile. Some of this instability is still visible in the results (for example in panel l of Figs. 3 and 4 where the mean may lie outside the percentiles) despite a "sanity filter" being used in our tests: we discard results that are

clearly far too large. Specifically, we discard results outside an absolute relative error threshold of 3 against the dispersion relation for the known exact profile. In a practical situation, that approach is obviously not possible and some other sensible check would be necessary to mitigate the effect of these rare predictions, e.g. comparing phase velocity against known statistics in the context being measured.

We finally notice how the OLS+Extr scheme can avoid the increase in variance for short wavelengths suffered by the other methods. On the other hand, the method systematically underpredicts $\tilde{c}$ in this range, corresponding to an underprediction of surface velocity, expected due to the strong non-linearity of $\tilde{U}$ near the surface. The perceptive reader may pause upon the somewhat counter-intuitive reduction in deviation for OLS+Extr going from $\Delta h = 0.1$ to $\Delta h = 0.5$. This is likely an artefact of the method, i.e. of taking a linear extrapolation of the topmost 6 points to predict the surface current. For the sparse grid, these 6 points reach 5 times further into the deep and so is fitting over a longer portion of the curve. Furthermore, because $\tilde{h}_s$ is fixed, the gradient on the sparser grid will be shallower. This may reduce the deviation at the cost of worse systematic bias. Whether OLS+Extr is preferable to OLS, say, in practice depends on the context – whether it is more important to avoid very poor estimates (high variance) or a systematic bias (nonzero mean).

In Fig. 5 we plot the first four statistical moments of the distribution of $\tilde{c}(\tilde{k})$ for the strongest exponential current $\mathrm{Fr}^2 = 0.05$. For clarity, only the OLS fit method is used. Columns from left to right show increasing measurement uncertainty $\Delta\tilde{U}$. The first row shows the mean values (as distributed through the panels in Fig. 4). A clear and intuitively obvious observation is that the mean value is better predicted at high $\tilde{k}$ when the surface velocity has been measured. Perhaps more interestingly, the badness of the mean estimated $\tilde{c}$ does not get appreciably worse with growing $\Delta\tilde{U}$. This emphasises again that the error in the mean value of $\tilde{c}$ when $\tilde{h}_s > 0$ depends almost exclusively on the accuracy of $\tilde{U}_{\mathrm{fit}}(0)$ by extrapolation and is relatively insensitive to other parameters. In the present case, however, the misprediction of the mean of $\Delta\tilde{c}/\tilde{c}$ using the OLS approximation is less than 2% even for the most unfavourable cases, which is unlikely to be significant in practice (other approximation methods fare worse as previously discussed).

The standard deviation of the distributions of $\tilde{c}$ (second row of Fig. 5) increases approximately linearly as a function of $\Delta\tilde{U}$, as one expects for a near-normal distribution of $\tilde{c}$.

One interesting observation is that reducing the spacing between measurement points by a factor of 5, from $\Delta\tilde{h} = 0.5$ to 0.1, reduces the standard deviation of results by less than a factor of 2. The same observation holds also for the two other profiles considered. The indication is that for the purposes of reducing the variance of phase velocity estimates, one would expect to be better off buying better sensors rather than more of them.

The skewness of a variable X with standard deviation σ and mean μ is defined as

$$S = \left\langle \left(\frac{X - \mu}{\sigma} \right)^3 \right\rangle. \tag{12}$$

The angular brackets indicate averaging. A positive skewness indicates that the most likely values of X are higher than μ, and *vice versa*. The excess kurtosis of X is defined as

$$K = \left\langle \left(\frac{X-\mu}{\sigma} \right)^4 \right\rangle - 3. \tag{13}$$

A positive excess kurtosis implies a greater emphasis on the tails of the distribution, i.e. the likelihood of measuring outliers is higher than for the normal distribution, and *vice versa*. For the normal distribution $S = K = 0$.

As for the skewness (third row of Fig. 5) and excess kurtosis (fourth row of same) of the distribution of $\Delta\tilde{c}/\tilde{c}$ on the exponential profile, we are able to draw few definite conclusions; more interesting results appear for the other two profiles considered. The one observation we make for future reference is that for the highest measurement uncertainty, the skewness is predominantly positive, albeit small in magnitude.

3.2 $z^{1/7}$ profile

The '1/7' profile is a common model for a turbulent boundary layer flow over a flat bottom. Since the shear in this instance is due to a boundary layer near the bottom rather than at the surface, the shear is concentrated around depth $\tilde{h}$ and is relatively weak near the surface. The characteristic shear depth is thus $d = h$. An extrapolation from $\tilde{z} = -\tilde{h}_s$ to 0 is now expected to be less problematic since the velocity deviates much less from a straight line than was the case for the exponential profile. From the outset we thus expect this profile to be the "easiest" of the three on which to correctly predict $\tilde{c}$ based on imperfect data. This turns out broadly true when the surface current has been measured ($\tilde{h}_s = 0$), but the picture is richer once extrapolation is involved.

As for the exponential profile before, we plot in Figs. 8 and 9 the mean and 10% percentiles of the distribution of $\tilde{c}(\tilde{k})$ from 10,000 random "measurements" $\{\tilde{z}_n, \tilde{U}_n\}$. Compared to the exponential profile case we notice that the detrimental effects for high $\tilde{k}$ of not having measurements of the surface flow velocity is slightly less pronounced. For this profile, the shear is relatively small near the surface and so extrapolations are expected to be more accurate than when near-surface curvature is present.

The exponential approximation procedure (EXP) performs remarkably well in this case, but for a slight overprediction of the mean value. Given its very poor performance for the Columbia profile later, and reasonable performance for the exponential profile, its performance clearly depends very strongly on the shape of the profile and seems to us to be somewhat unpredictable.

It is noticeable that the OLS now performs among the worst in terms of variance, particularly when $\tilde{h}_s > 0$. It also shows a tendency to marginally mispredict the mean for short wavelengths. One can see the reason for this by noting that the shape of the profile (Fig. 1b) is hard to reproduce by only a small number of terms of a power series in powers of $\tilde{z}$: $\tilde{U}(\tilde{z})$ has large curvature but it is situated far from the surface, and in order to fit it, higher order terms in the polynomial must be very significant near $\tilde{z} = -\tilde{h}$, but cancel each other for the rest of the water column. Once a high order polynomial is then extended beyond the range of points to which it was fitted,

its behaviour becomes very volatile. Due to inherent similarities, TLS also suffers the same problem.

The difficulty in reproducing the $\tilde{z}^{1/7}$ profile with a low order polynomial is also apparent when considering skewness and excess kurtosis of $\Delta\tilde{c}/\tilde{c}$ for the largest $\Delta\tilde{U}$ and $\tilde{h}_s$, see Fig. 10. Both skewness and kurtosis are significant. Note that the skewness is perturbed strongly in the positive direction. For the strongest $\Delta\tilde{U}$, kurtosis is also perturbed strongly in a positive direction indicating that one should expect outlier results more often.

3.3 Columbia River test profile

Finally we turn to a velocity profile based on real measurements in the estuary of the Columbia River [12], where buoyant fresh water forms a surface jet atop the heavier salt water to produce strong shear near the surface. To do this, we investigated the dataset and picked one sample which was of both of good quality and had qualitatively 'interesting' features; we discarded the spanwise component, which is of small magnitude and irrelevant for our purposes. Precisely as in our models, this data consists of a number of discrete points of unknown accuracy with no data for the top 2m of the water column. In order to construct an exemplar profile with the same characteristics, we arbitrarily shift the profile upwards so that the topmost available measurement depth is now considered the surface (in terms of performance this should be conservative, since the shear is likely to further increase through the top 2 meters). A polynomial is then fitted to the points to give the profile shown in Fig. 1c, around which our simulated "measurements" will be made. With a reasonable choice of d for the characteristic length scale of the velocity variation with depth, we obtain $\mathrm{Fr}^2 = 0.05$. Measurement points are distributed equidistantly throughout the water column down to depth $z = -h$.

Compared to the two previously considered profiles, which were both idealised model flows, the Columbia profile has a more complicated structure, containing several inflection points and a degree of backflow near the bottom. Occurrences of particularly good (or poor) performance of particular fitting schemes which hold only for special cases with monotonic curvature may now be exposed.

From Fig. 13, perhaps the clearest trend is that all methods struggle to correctly predict the surface velocity, and consequently the value of $\tilde{c}(\tilde{k})$ for high $\tilde{k}$, except for OLS. The exponential fitting in particular can very badly underpredict the surface velocity, and combined with a small variance is almost guaranteed to give too low values of $\tilde{c}$ for sufficiently large $\tilde{k}$ in all the cases considered, especially when $\tilde{h}_s > 0$. It is surprising to what extent this underprediction of the surface velocity $\tilde{U}_0$ holds true also when a measurement point is present at the surface itself. Worse still, for Fig. 13g, Fig. 13h, and Fig. 13i, the predicted mean actually lies well outside the percentiles indicating that the fitting process is very unstable in this context.

For OLS the trend is as for the previous two profile, however: the mean value of the distribution of $\tilde{c}(\tilde{k})$ remains very close to correct but at the cost of higher variance.

The bottom row of Fig. 14 shows that the OLS fit on this profile has nontrivial positive excess kurtosis for medium-to-long wavelengths, again in the positive direc-

tion. As for the previous two profiles the skewness is observed to be predominantly positive. The indication is that for a profile with much structure and several inflection points the shape of the distribution of c, in particular the thickness of its tails and bias, becomes unpredictable.

4 Further numerical experiments to find optimal $\Delta\tilde{h}$ for reducing deviation

For all three velocity profiles, we have observed that increasing the number of measurement points — or, equivalently, reducing $\Delta\tilde{h}$ — appears to effect only a modest reduction in variance in estimated phase velocity. This may have implications when considering equipment choice and considering more vs more accurate sensors, for instance. To make an initial attempt at answering this question, we repeated some numerical experiments using the OLS fitting method, this time changing $\Delta\tilde{h}$ in a granular manner while fixing all other parameters ($\Delta\tilde{h}_s = 0$, $\Delta\tilde{U} = 0.1$, $\mathrm{Fr}^2 = 0.05$). Instead of using $\Delta\tilde{h}$, for these purposes it is easier to talk in terms of the number of measurement points, which we denote N_h.

In Figs. 6, 11, and 15 we calculate the standard deviation of componentwise relative error in $\tilde{c}(\tilde{k})$ for the exponential, $1/7$, and Columbia River velocity profiles, respectively, for various N_h values. The value of the standard deviation for each N_h at $\tilde{k} = \pi$ is also noted. It is clear from these figures that for all three cases, the standard deviation decreases slowly with increasing N_h. For instance, in the exponential profile, increasing N_h fivefold from 10 to 50 reduces the standard deviation by only circa 40%.

For particularly low N_h, it is evident that the standard deviation becomes sensitive to the exact placement of the measurement points and has some unpredictable $\tilde{k}$ dependency. We in addition seek a scalar measure that avoids this $\tilde{k}$ dependency. Given the standard deviation of the relative error in phase velocity evaluated at $\tilde{k}$ using N_h measurement points, $\sigma_{N_h}(\Delta\tilde{c}(\tilde{k})/\tilde{c}(\tilde{k}))|_{\tilde{k}}$, we define

$$\gamma(N_h) := \int_{\tilde{k}_{\min}}^{\tilde{k}_{\max}} \sigma_{N_h}(\Delta\tilde{c}(\tilde{k})/\tilde{c}(\tilde{k}))|_{\tilde{k}} \, \mathrm{d}\tilde{k} \tag{14}$$

as a scalar measure of the global deviation for a particular N_h.

In Fig. 7, 12, and 16, we plot normalised $\sigma_{N_h}(\Delta\tilde{c}/\tilde{c})$ evaluated at $\tilde{k} = \pi$ against N_h and also normalised γ against N_h. For $\gamma(N_h)$, in all three profiles, the trend is linear up to about $N_h \sim 100$ after which the trend is shallower. The slope in the linear region is approximately -0.004 for all profiles. In other words, to reduce $\gamma(N_h)$ by 20% requires 50 extra measurement points.

We therefore conclude that increasing the number of measurement points has only a modest effect on improving the accuracy of estimating $\tilde{c}(\tilde{k})$.

5 Pseudospectral analysis

The key question in this paper is how error in, or a perturbation of, a velocity profile can influence phase velocities of the waves propagating atop that flow. If the Rayleigh

equation is viewed as a linear operator then a perturbation of the velocity profile is a perturbation of that linear operator. This permits analysis using pseudospectra: for a perturbation of the linear operator of order ε, how far can the eigenvalue — the phase velocity — be perturbed as a consequence. In our context, noise or mismeasurement of the velocity profile can be considered a subset of such possible perturbations.

In this section we adopt slightly different notation in an attempt to adhere to the conventions in the literature. Particularly, in accordance with standard mathematical notation λ denotes an arbitrary eigenvalue (not a wavelength), and $\boldsymbol{x}$ is a corresponding eigenvector.

Preamble. The system described in Eq. (4) can be naturally viewed as a Sturm–Liouville eigenvalue problem with $\tilde{k}$ as the eigenvalue, and $\tilde{c}$ acting as a scalar parametrisation. However, akin to the example in [2, p. 135], we can equally well consider the system as an eigenproblem with the phase velocity, $\tilde{c}$, as the eigenvalue and $\tilde{k}$ acting as the parametrisation of the problem. Rewriting Eq. (4) in powers of $\tilde{c}$ for emphasis:

$$\left[\tilde{c}\left(\frac{\mathrm{d}^2}{\mathrm{d}z^2}-\tilde{k}^2\right)+\tilde{U}(\tilde{z})\left(\frac{\mathrm{d}^2}{\mathrm{d}z^2}-\tilde{k}^2\right)-\tilde{U}''(\tilde{z})\right]\tilde{w}=0; \tag{15a}$$

$$\left[\tilde{c}^2\frac{\mathrm{d}}{\mathrm{d}z}+\tilde{c}\left(\tilde{U}'(\tilde{z})-2\frac{\mathrm{d}}{\mathrm{d}z}\right)+\frac{\mathrm{d}}{\mathrm{d}z}-Fr^{-2}-\tilde{U}'(\tilde{z})\right]\tilde{w}=0;\ \ \tilde{z}=0; \tag{15b}$$

$$\tilde{w}=0;\ \ \tilde{z}=-\tilde{h}. \tag{15c}$$

The solution sets must obviously be the same in that any triplet $\{\tilde{k},\tilde{c},\tilde{w}\}$ which is a solution when $\tilde{k}$ is considered the eigenvalue is also a solution when $\tilde{c}$ is considered the eigenvalue. However, in practical scenarios, we are only interested in the subset containing real quantities and it is more useful to have the facility to calculate $\tilde{c}$ from $\tilde{k}$. Hence, the distinction becomes more important.

The system described by Eq. (15) can be discretised and solved numerically. Recapitulating from [19], a collocation pseudospectral method is used to discretise, yielding a quadratic eigenvalue problem of the form:

$$(\boldsymbol{A}_2\tilde{c}^2+\boldsymbol{A}_1\tilde{c}+\boldsymbol{A}_0)\boldsymbol{w}=0,\ \ \boldsymbol{w}\neq 0, \tag{16}$$

where a standard two boundary point row-replacement strategy [27] is used so that the free surface condition is in the topmost rows of $\boldsymbol{A}_2,\boldsymbol{A}_1,\boldsymbol{A}_0$ and similarly the bottom rows are zero for the bottom Dirichlet boundary condition (which are then removed). This implies that the only nonzero row of $\boldsymbol{A}_2$ is the first row. We also note that matrix $\boldsymbol{U}$ and its derivatives are diagonal. A QZ decomposition is then used to solve the quadratic eigenvalue problem along with appropriate mechanisms to determine which eigenvalue in the numerical solution set is the physical phase velocity.

Background and motivation of pseudospectra. What then of the behaviour of Eq. (15) under a perturbation? This question can be answered by using pseudospectral analysis. We first consider the standard eigenvalue problem,

$$\boldsymbol{A}\boldsymbol{x}=\lambda\boldsymbol{x},\boldsymbol{x}\neq 0, \tag{17}$$

for $\boldsymbol{A} \in \mathbb{C}^{n\times n}, \boldsymbol{x} \in \mathbb{C}^n, \lambda \in \mathbb{C}$. We denote the spectrum of $\boldsymbol{A}$ by $\Lambda(A)$, i.e. the set of eigenvalues of Eq. (17).

If we consider $\boldsymbol{A}$ to be a finite dimensional discretisation of some linear operator then it is natural to ask what happens to the spectrum when $\boldsymbol{A}$ is perturbed by some matrix $\boldsymbol{E}$ with $\|\boldsymbol{E}\| < \varepsilon$. This is directly analogous to the fundamental problem we are considering, viz. how do perturbations of a background velocity profile influence the phase velocity of the waves. This notion can be made formal [28, p. 14, eqn 2.3].

Definition 1 (Matrix definition of pseudospectral set) The $\varepsilon-$pseudospectral set of $\boldsymbol{A}$ for $\varepsilon > 0$ is defined as,

$$\Lambda_\varepsilon(\boldsymbol{A}) := \{\lambda \in \Lambda(\boldsymbol{A}+\boldsymbol{E}) : \text{for all } \boldsymbol{E} \in \mathbb{C}^{n\times n}, \|\boldsymbol{E}\| \leq \varepsilon\}. \tag{18}$$

Equivalently, we may approach this using the resolvent, $(\lambda \boldsymbol{I} - \boldsymbol{A})^{-1}$. The resolvent set is defined as $\rho(\boldsymbol{A}) := \{\lambda : (\lambda \boldsymbol{I} - \boldsymbol{A}) \text{ is invertible }\}$; the spectrum is then the compliment of the resolvent set in the complex plane. So we have that for λ to be an eigenvalue, the resolvent $(\lambda \boldsymbol{I} - \boldsymbol{A})^{-1}$ must be undefined. As noted in [28, p. 12], an eigenvalue is inherently fragile in the sense that upon an arbitrarily small perturbation, it will cease to be an eigenvalue. So, from a practical standpoint, it's perhaps better to ask for a set which is 'almost' an eigenvalue or equivalently when the resolvent norm, $\|(\lambda \boldsymbol{I} - \boldsymbol{A})^{-1}\|$ is suitably large. This leads us to the second definition of pseudospectra using the resolvent norm [28, p. 13, eqn 2.1].

Definition 2 (Resolvent definition of pseudospectral set) The $\varepsilon-$pseudospectral set of $\boldsymbol{A}$ for $\varepsilon > 0$ is defined as,

$$\Lambda_\varepsilon(\boldsymbol{A}) := \{\lambda \in \mathbb{C} : \|(\lambda \boldsymbol{I} - \boldsymbol{A})^{-1}\| \geq 1/\varepsilon\}. \tag{19}$$

Note that [28, thm 2.1, p. 16] establishes equivalence between the two definitions.

Thus ends our brief description of standard pseudospectra and we now move into slightly more interesting territory.

Pseudospectra for Rayleigh equations. We now turn to the Rayleigh equation, governing linear surface waves on a shear current. We are considering perturbations and eigenvalues which are by definition complex. Although phase velocities for propagating waves in inviscid fluid mechanics are real, this poses no problems as we can restrict ourselves to considering effects only along the real axis. Dissipation and/or interaction with critical layers can moreover move the eigenvalues off the real c axis, something which is used deliberately in systems with stationary and periodic time dependence to enforce the radiation condition (see e.g. Ref. [17]).

The discretised system, Eq. (16), is a quadratic eigenvalue problem and therefore the definitions of pseudospectra have to be modified to accommodate. Repeating the definitions from [26], a polynomial matrix can be written,

$$\boldsymbol{P}(\lambda) := \lambda^m \boldsymbol{A}_m + \lambda^{m-1}\boldsymbol{A}_{m-1} + \ldots + \boldsymbol{A}_0,\ \boldsymbol{A}_j \in \mathbb{C}^{n\times n}. \tag{20}$$

which leads naturally to the polynomial eigenvalue problem,

$$\boldsymbol{P}(\lambda)\boldsymbol{x} = 0,\ \boldsymbol{x} \neq 0. \tag{21}$$

The $\varepsilon-$pseudospectrum can then defined by [26, eqn 2.3],

Definition 3 (Matrix definition of pseudospectra for polynomial eigenproblem)

$$\Lambda_\varepsilon(\boldsymbol{P}) := \{\lambda \in \mathbb{C} : (\boldsymbol{P}(\lambda) + \Delta\boldsymbol{P}(\lambda))\boldsymbol{x} = 0, \boldsymbol{x} \neq 0, \\ \Delta\boldsymbol{P}(\lambda) := \lambda^m \Delta\boldsymbol{A}_m + \lambda^{m-1}\Delta\boldsymbol{A}_{m-1} + \ldots + \Delta\boldsymbol{A}_0, \|\Delta\boldsymbol{A}_j\| \leq \alpha_j\}, \quad (22)$$

where $\boldsymbol{\alpha} := (\alpha_m, \alpha_{m-1}, \ldots, \alpha_0)$ is a vector of scalar parameters which specify the relative magnitude of the matrix perturbations.

As may be expected, there is an equivalent definition using the resolvent established by [26, lemma 2.1],

Definition 4 (Resolvent definition of pseudospectra for polynomial eigenproblem)

$$\Lambda_\varepsilon(\boldsymbol{P}) := \{\lambda \in \mathbb{C} : \|\boldsymbol{P}(\lambda)^{-1}\| \geq (\varepsilon p(|\lambda|))^{-1}\} \quad (23)$$

where $p(x) := \sum_{j=0}^{m} \alpha_j x^k$.

For our purposes, we use Eq. (16) and let $\boldsymbol{P}(c) := (\boldsymbol{A}_2 c^2 + \boldsymbol{A}_1 c + \boldsymbol{A}_0)$ so that the relevant pseudospectral set is,

$$\Lambda_\varepsilon(\boldsymbol{P}) := \{c \in \mathbb{C} : \|(\boldsymbol{A}_2 c^2 + \boldsymbol{A}_1 c + \boldsymbol{A}_0)^{-1}\| \geq (\varepsilon p(|c|))^{-1}\}, \quad (24)$$

with $\alpha_0 = \alpha_1 = \alpha_2 = 1$. This choice of α_j preserves an absolute scale which is appropriate for our problem given the disparity of $\boldsymbol{A}_2$ compared to $\boldsymbol{A}_1$ and $\boldsymbol{A}_0$.

A note on homogeneous eigenvalues. We note that the recommendation in [11] is to address the pseudospectra of matrix polynomials by expressing the polynomial in homogeneous form $\boldsymbol{P}(\alpha, \beta) = \alpha^m \boldsymbol{A}_m + \alpha^{m-1}\beta\boldsymbol{A}_{m-1} + \ldots + \beta^m \boldsymbol{A}_0$. The eigenvalues become of the form $(\alpha, \beta) \in \mathbb{C}^2$ rather than a scalar, with $\lambda = \alpha/\beta$ for $\beta \neq 0$. This, in a sense, is more suited to handling infinite eigenvalues. However, it does significantly complicate plots with the suggestion being a stereographic projection onto the Riemann sphere. Whilst it solves the problem of infinite eigenvalues in relation to level set style plots, the spherical projection is far from intuitive. In our situation, we choose to accept difficulties with infinite eigenvalues because it retains a clearer representation overall and permits the reader to infer at least broad conclusions.

5.1 Pseudospectra plots for candidate velocity profiles

Fig. 17 shows the pseudospectra of the exemplar velocity profiles for $\tilde{k} = \pi$; the full spectrum returned by the collocation scheme is superimposed on top. The eigenvalue corresponding to the phase velocity is marked with a red asterisk (rightmost), the lower branch eigenvalue can also be seen in a blue asterisk (leftmost). The collocation scheme also returns a cluster of values from the essential spectrum Ref. [7, ch. 4], in this case approximately around the origin, such that $\tilde{U}(z) - \tilde{c} = 0$ for some $\tilde{z} \in [-\tilde{h}, 0]$; in other words, $\tilde{c}$ which satisfy the condition for a critical layer at some $\tilde{z}$. The number of these values returned depends on the order of the collocation scheme: as the order increases, it returns more and more values from this interval. We shall

consider only the discrete spectrum: the leftmost and rightmost eigenvalues, which represent propagating wave modes.

Whilst the plots for the three shear profiles are clearly different, they share some key qualitative characteristics.

- The level set curves are not isotropic around the phase velocity $\tilde{c}$, which suggests the underlying operator is not normal.
- The arrangement of the level set curves appear to be influenced by infinite eigenvalues, which is a consequence of A_2 being singular. This was mentioned above and was anticipated by [11].
- For a perturbation of order ε, a real phase velocity can increase far more than it can decrease. This applies in all cases.
- Proximity of the phase velocity to the critical layer region appears to determine the tightness of the level curves on the negative side of the phase velocity.

5.2 Remarks and conclusions

It is clear that perturbations of the underlying linear operator have an asymmetric influence on the phase velocity, although the exactly how this relates to a given perturbation of the (actual/measured/reconstructed) velocity profile is not obvious. That the phase velocity is likely to deviate by a larger amount in the positive real direction compared to the negative direction seems consistent with the skewness results in the statistical moments as discussed in Sec. 3. Whenever the uncertainty of measurements, ΔU, was sufficiently large, skewness was found to be predominantly positive in all cases, meaning the estimated value of $\tilde{c}$ was more likely to lie above than below the mean value. Correspondingly we notice that a sufficiently large perturbation ε is able to shift the eigenvalue $\tilde{c}$ further along the real axis in the positive direction than in the negative (although, note carefully that the shift in the mean value observed for high $\tilde{k}$ in a number of cases is due to misprediction of the surface velocity by the curve fitting schemes and bears no relation to the pseudospectral analysis in which consider a single value of $\tilde{k}$).

It is anticipated that structured polynomial pseudospectra arranged in such a manner as to match the diagonal nature of the shear matrices may provide further insight. However, this presents both some philosophical and practical difficulties and would be the subject of future research effort.

6 Conclusions

We have considered the effect of imperfect knowledge of a sub-surface shear current on the error in the calculated phase velocity of linear surface waves, i.e. in the dispersion relation. This corresponds to realistic situations where the current is only measured at a finite set of depths, with nonzero uncertainty, and often not including the surface but beginning at some depth $\tilde{h}_s$, typically a few meters.

Four different methods for constructing a continuous velocity profile from a discrete set of points were tested: an Ordinary Least Squares (OLS) fitting of the data to a

polynomial, an exponential fit (EXP), a Tikhonov Regularised Least Squares method (TLS) and when the topmost measurement point is not at the surface, a method of first linear-extrapolating the profile to obtain an estimate of the surface velocity then performing an OLS (OLS+Extr).

Based on our three rather different velocity profiles — an exponential profile modelling a near-surface shear layer, a profile varying as $z^{1/7}$ modelling a turbulent bottom boundary layer, and an exemplar of a measured flow in the Columbia River delta — it seems indicated that the OLS method is the most robust against systematic bias of the mean but at the expense, in some but not all cases, of a higher variance of results. The TLS typically has similar variance as OLS, but with a tendency to systematically bias results for short waves, and is therefore hardly ever the most favourable option. The performance of EXP varies very greatly based on the exact shape of the velocity profile: for the $z^{1/7}$ profile it performs excellently, for the exponential profile it is among the best, but for the Columbia River profile it fails badly. When the shape of the profile is not known *a priori* the performance of EXP is therefore unpredictable. Likewise, the OLS+Extr method, relevant when extrapolation is necessary, performs poorly for profiles which are strongly curved near the surface but better for the $z^{1/7}$ profile whose large curvature is situated at greater depth.

The clearest and unsurprising conclusion is doubtlessly the importance of performing a surface current measurement when the profile is expected to be most strongly sheared near the top of the water column. Both for the exponential and Columbia River profiles the necessity to extrapolate from the topmost measurement point to the surface greatly increases the variance of the resulting distribution of $\tilde{c}(\tilde{k})$. In both of these cases the error from extrapolation quickly outweighs that from other sources. Conventional current measurements using Acoustic Doppler profilometry cannot capture the portion of the water column closest to the surface. In contrast, measurements based on measuring the wave dispersion and inferring the current from this (e.g. [18], [3], [4], [23]), have the *best* accuracy in the near-surface region, an example of the power of this technique (naturally, $c(k)$ is directly measured for a set of k values in this case, posing somewhat of a complementary problem to that which we consider).

Studying the standard deviation of the distributions of $\Delta\tilde{c}/\tilde{c}$ estimated with OLS approximation, the trend is the same for all profiles: not having a surface measurement (topmost measuring point sits at $\tilde{z} = -\tilde{h}_s$) increases the deviation manyfold. In the statistical moments tests, the standard deviation is roughly proportional to that of the measurement uncertainty, ΔU, but a fivefold increase in the density of sensors only reduces the deviation by a factor of 2. This result is replicated when the dependence of standard deviation on the number of points N_h is examined. All else being equal it seems indicated that, after a certain threshold, it is preferable to invest in better measurements in each point rather than more measurement points.

The OLS fitting method shows the least tendency to mispredict mean value and also the least skewness. Nevertheless, both skewness and excess kurtosis can be significant, particularly once extrapolation to the water surface is necessary. The sign of the excess kurtosis seems to be predominantly positive but depends in a complicated manner on the interplay between velocity profile and the method whereby $U(z)$ is reconstructed. There is also a clear tendency for skewness to be predominantly positive

in all cases when the uncertainty ΔU is sufficiently large, indicating that even though the average value from many measurements is accurate, too-high values are likely to be further off than mispredictions which are far too low.

A pseudospectral analysis, considering the Rayleigh equation as a linear operator, can shed light on the question: If the velocity profile is somehow perturbed by an amount ε –e.g. a measurement is performed giving a different result than the real profile– how far can the consequent eigenvalue c shift from its true value?

While there is no simple relation between ΔU and ε, our pseudospectral analysis nevertheless produces some interesting observations. While the pseudospectra differ significantly for the different profiles, it holds for all three that a perturbation of magnitude ε can shift the value of c further in the positive direction than it can in the negative. When ε is large enough this asymmetry becomes marked. This observation appears to concord with our observation from the numerical experiments that when ΔU is sufficiently large, skewness was found to be predominantly positive. We regard this conclusion as somewhat tentative at the present stage, worthy of further investigation.

7 Figures

7.1 Examples of fits on noisy measurement data (Exponential profile as 'exact' exemplar)

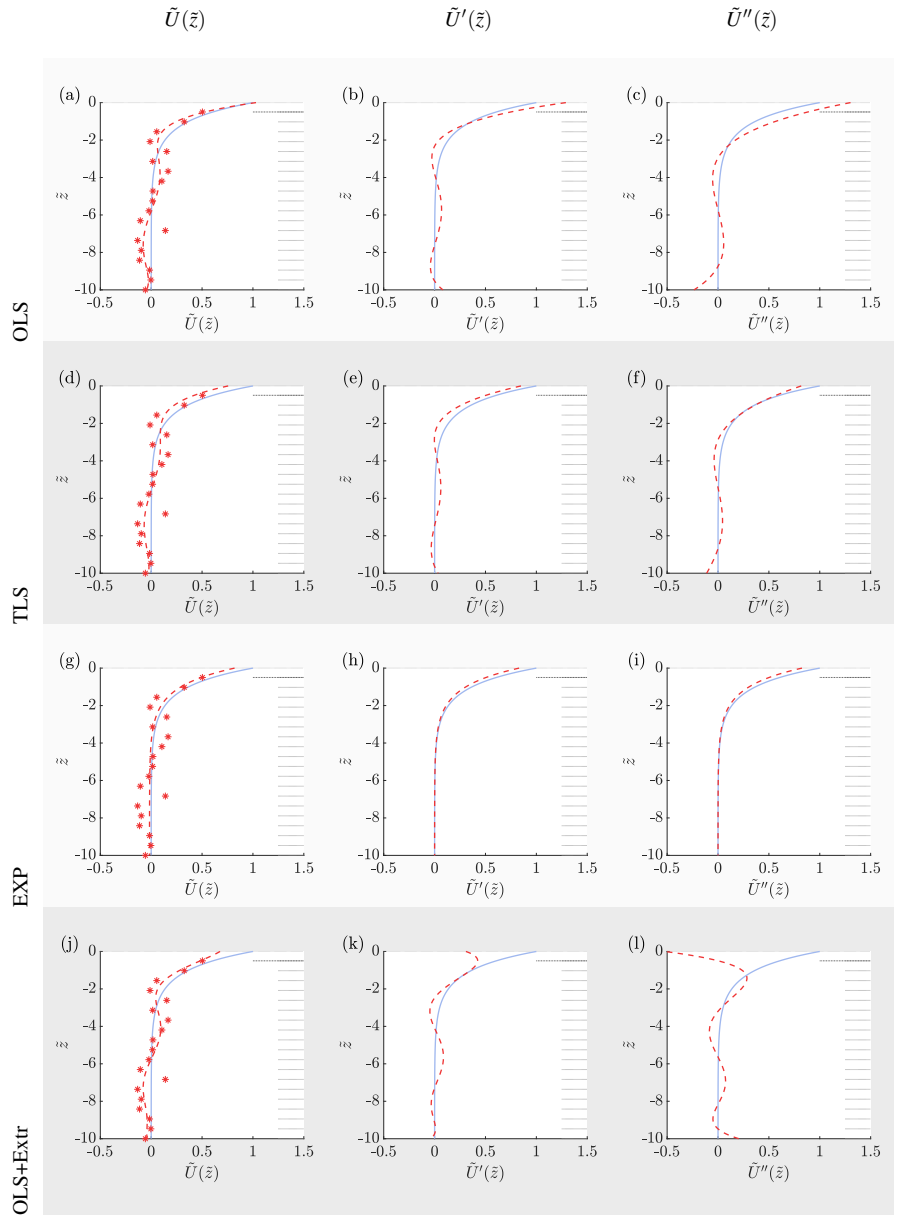

Fig. 2 Reconstruction of noisy measurements on an exponential shear profile using various fitting methods. The velocity profile has a shear of $Fr^2 = 0.01$. Light blue curve is the exact velocity profile, the red asterisks are the noisy measurements, the faint grey lines indicate the measurement positions, the topmost heavy feint grey line is the location of $\tilde{h}_s$, and the dashed red is the reconstructed profile.

7.2 Exponential profile simulations

7.2.1 Exponential profile: percentiles for $\mathrm{Fr}^2 = 0.01$

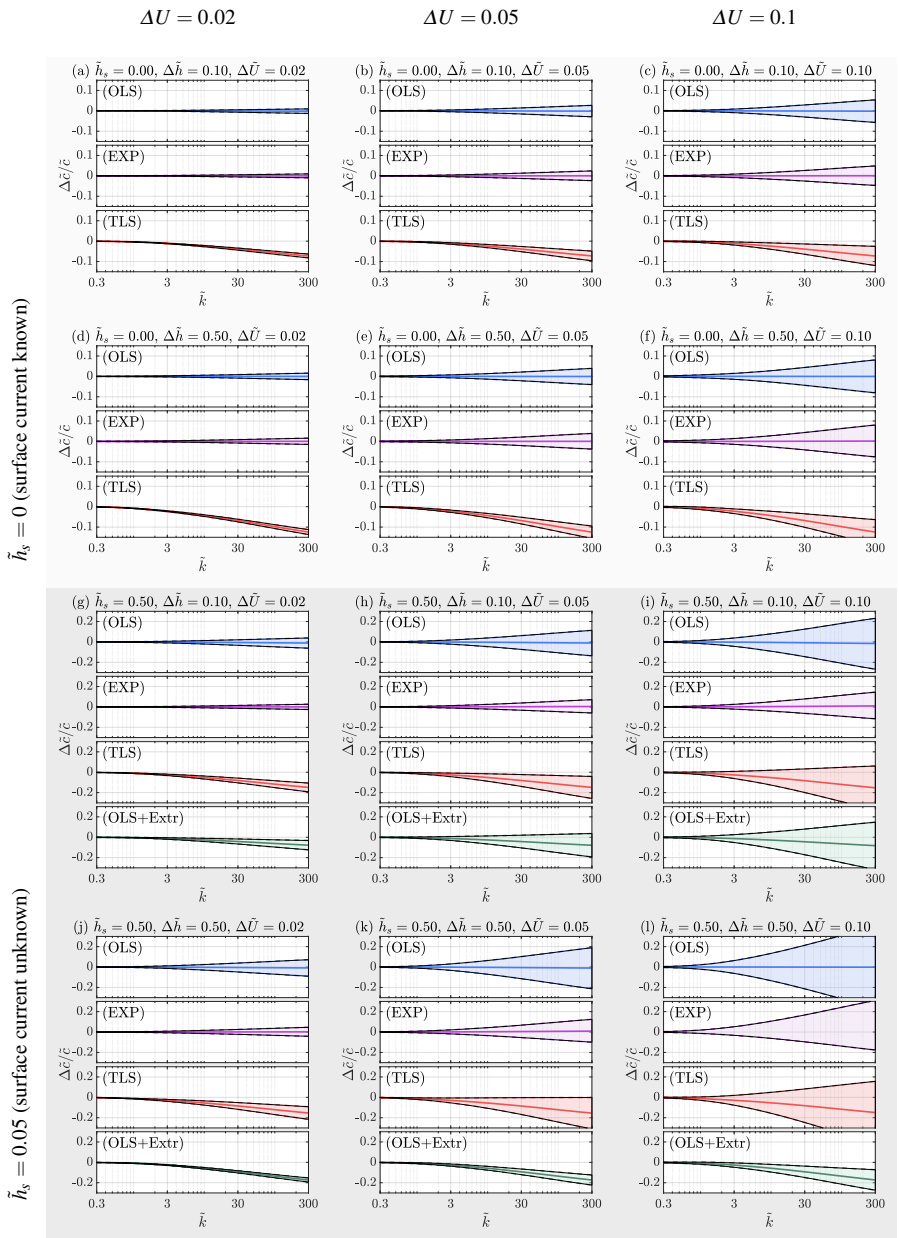

Fig. 3 10th percentile region with mean of the relative error in phase velocity on an ensemble of noisy measurements of an exponential velocity profile reconstructed using various fitting methods. The velocity profile has shear $Fr^2 = 0.01$. The top two rows assume a measurement point located at the surface whereas the bottom two rows assume the topmost measurement point is at $\tilde{h}_s$. Colomnwise represents greater magnitude Gaussian error. The alternate rows are due to the two different $\Delta\tilde{h}$ used.

7.2.2 *Exponential percentiles for* $\mathrm{Fr}^2 = 0.05$

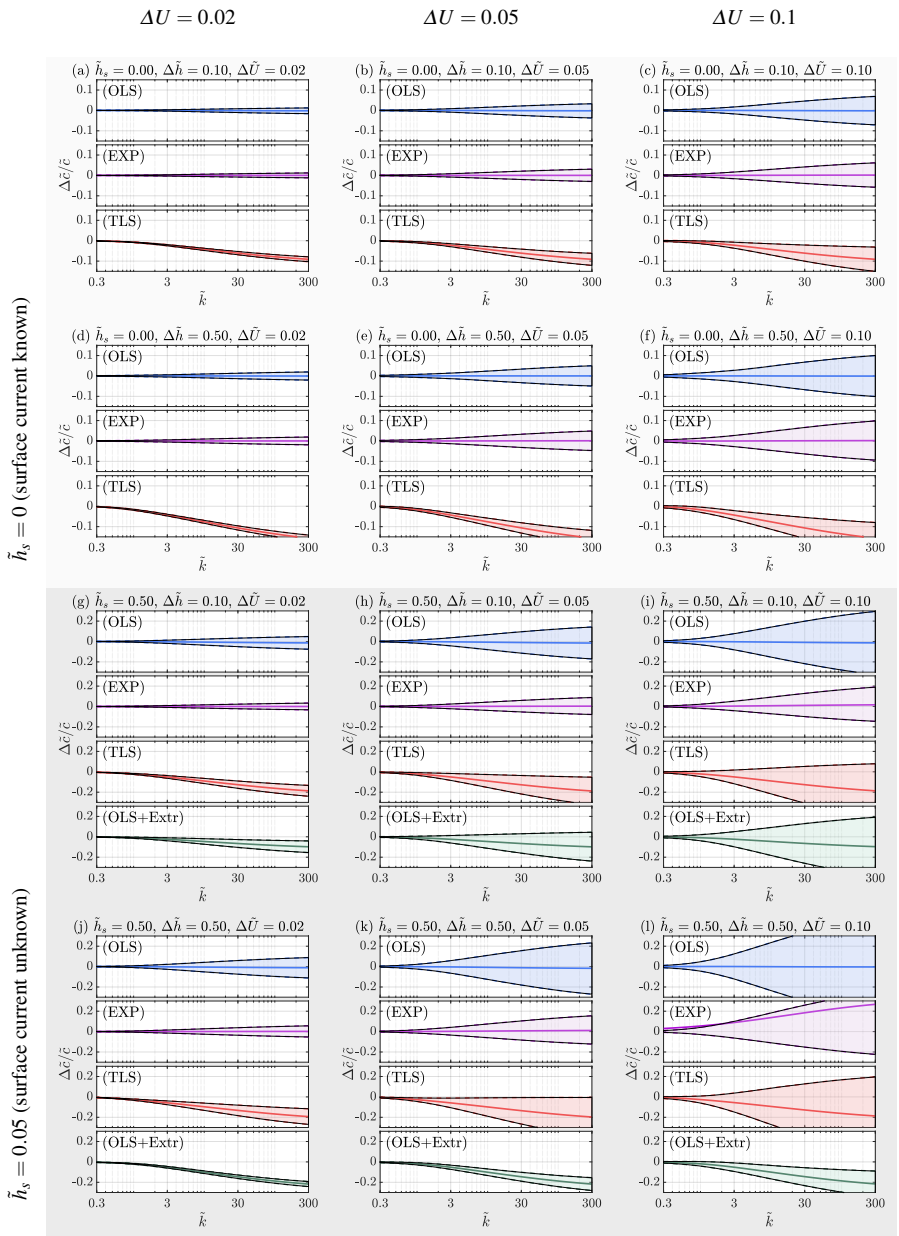

Fig. 4 Similar to Fig. 3 but the exponential profile has shear of $Fr^2 = 0.05$.

7.2.3 *Exponential profile statistical moments for* $\mathrm{Fr}^2 = 0.05$ *(for OLS fit)*

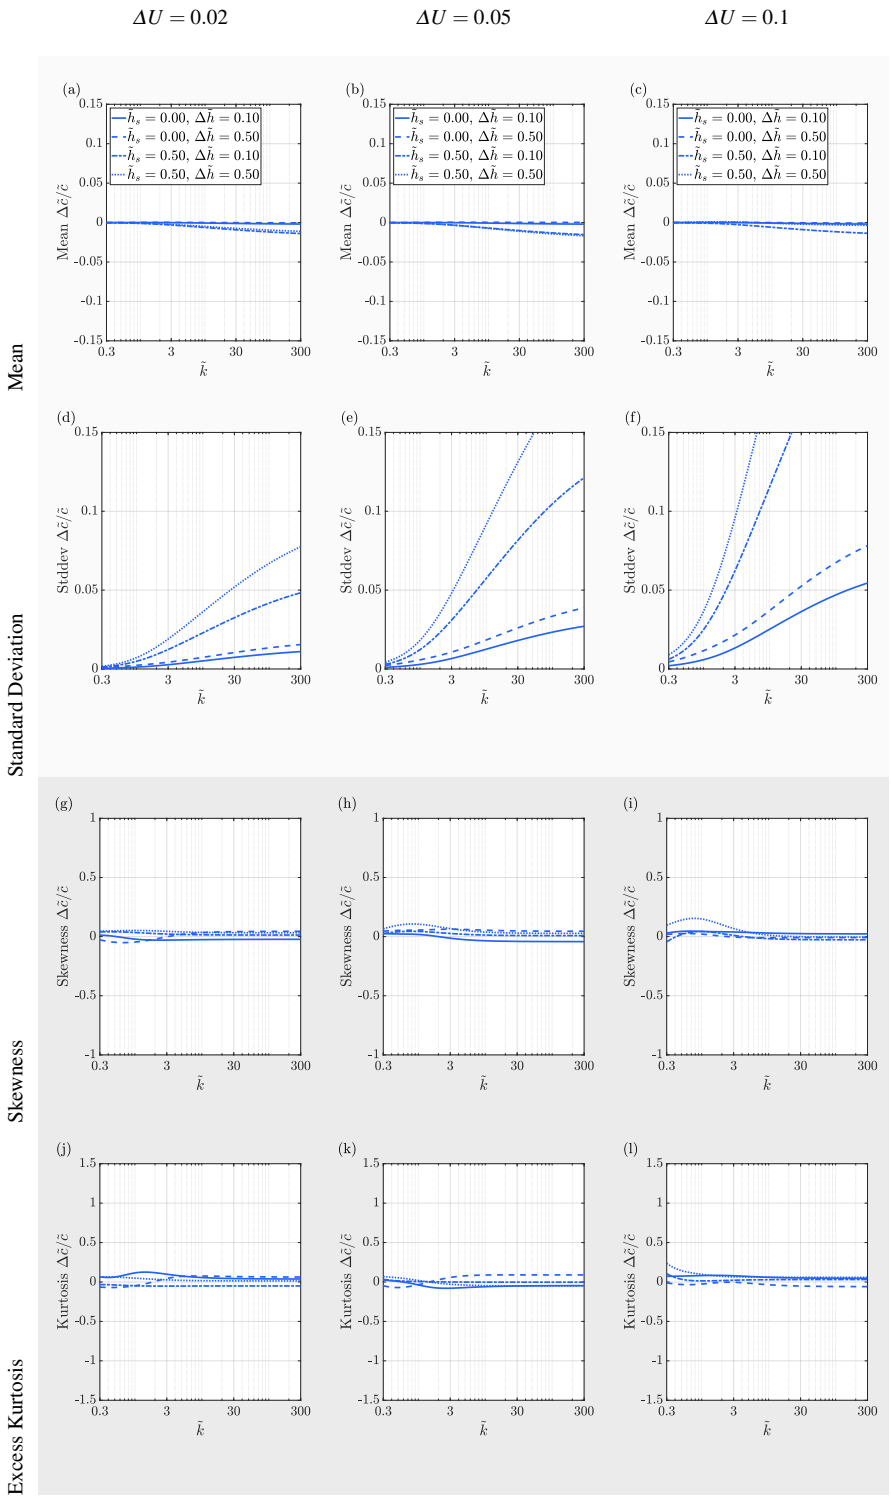

Fig. 5 Here we plot the statistical moments –mean, standard deviation, skewness, and kurtosis– of the relative error in phase velocity on an ensemble of noisy measurements of an exponential velocity profile reconstructed using the OLS method. The velocity profile has shear $Fr^2 = 0.05$. The standard deviation is chosen instead of the variance to keep the same scale as the mean.

7.2.4 Exponential profile deviation dependence on $\Delta\tilde{h}$

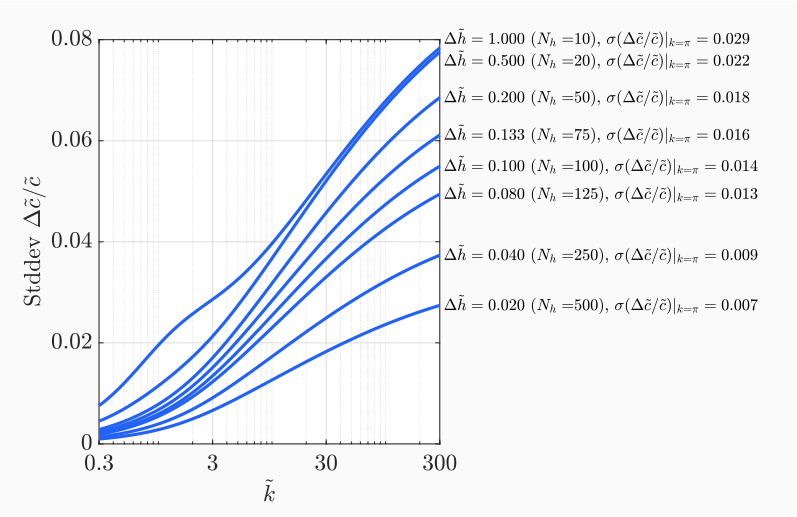

Fig. 6 The standard deviation in $\Delta\tilde{c}/\tilde{c}$ against k for several $\Delta\tilde{h}$ values for the exponential profile. The other parameters are kept constant: $\mathrm{Fr}^2 = 0.05$, $\Delta U = 0.1$, $\tilde{h}_s = 0$.

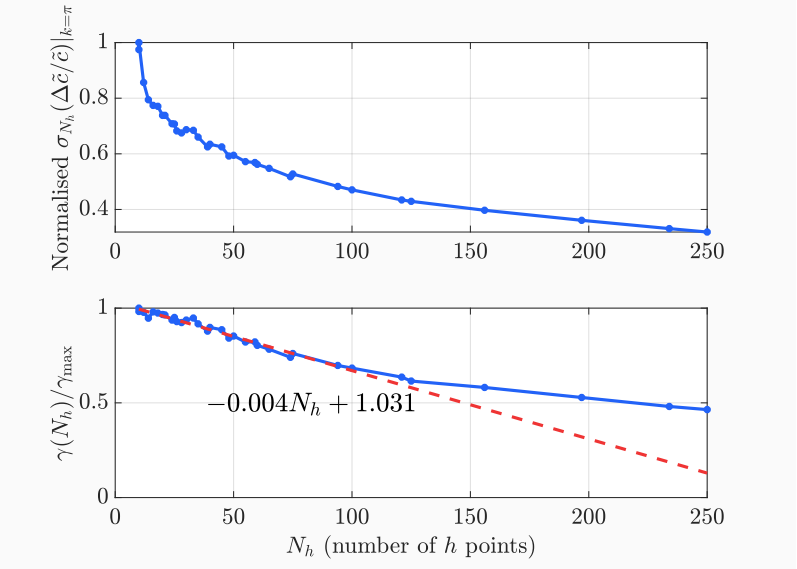

Fig. 7 The integral of the standard deviation of $\Delta\tilde{c}/\tilde{c}$ with respect to k for a range of $\Delta\tilde{h}$ values, shown here instead as the number of h points used for clarity. The other parameters are kept constant: $\mathrm{Fr}^2 = 0.05$, $\Delta U = 0.1$, $\tilde{h}_s = 0$.

7.3 $z^{1/7}$ profile simulations

7.3.1 $z^{1/7}$ profile percentiles for $\mathrm{Fr}^2 = 0.01$

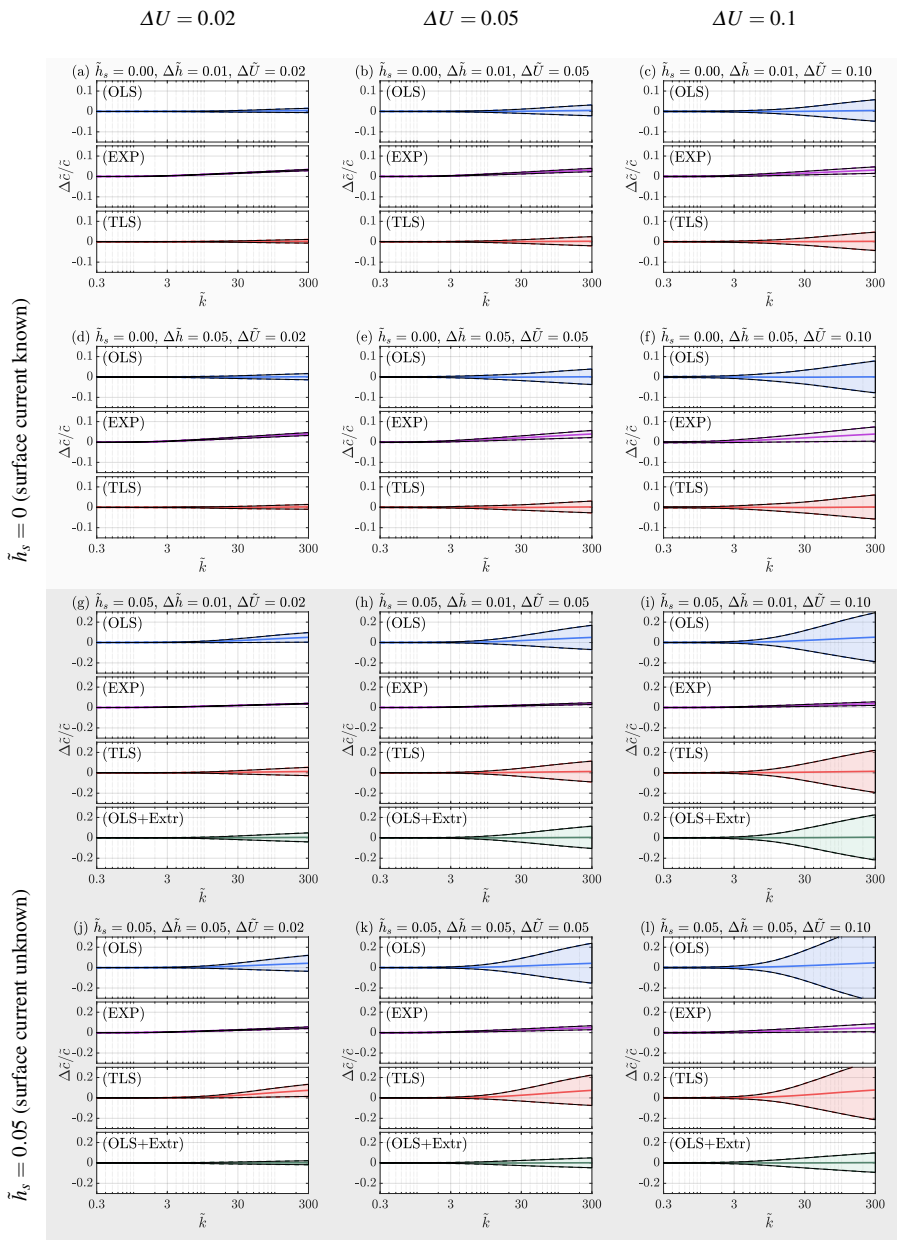

Fig. 8 Similar to Fig. 3 but for the '1/7' profile (6) with shear $Fr^2 = 0.01$

7.3.2 $z^{1/7}$ *profile percentiles for* $\mathrm{Fr}^2 = 0.05$

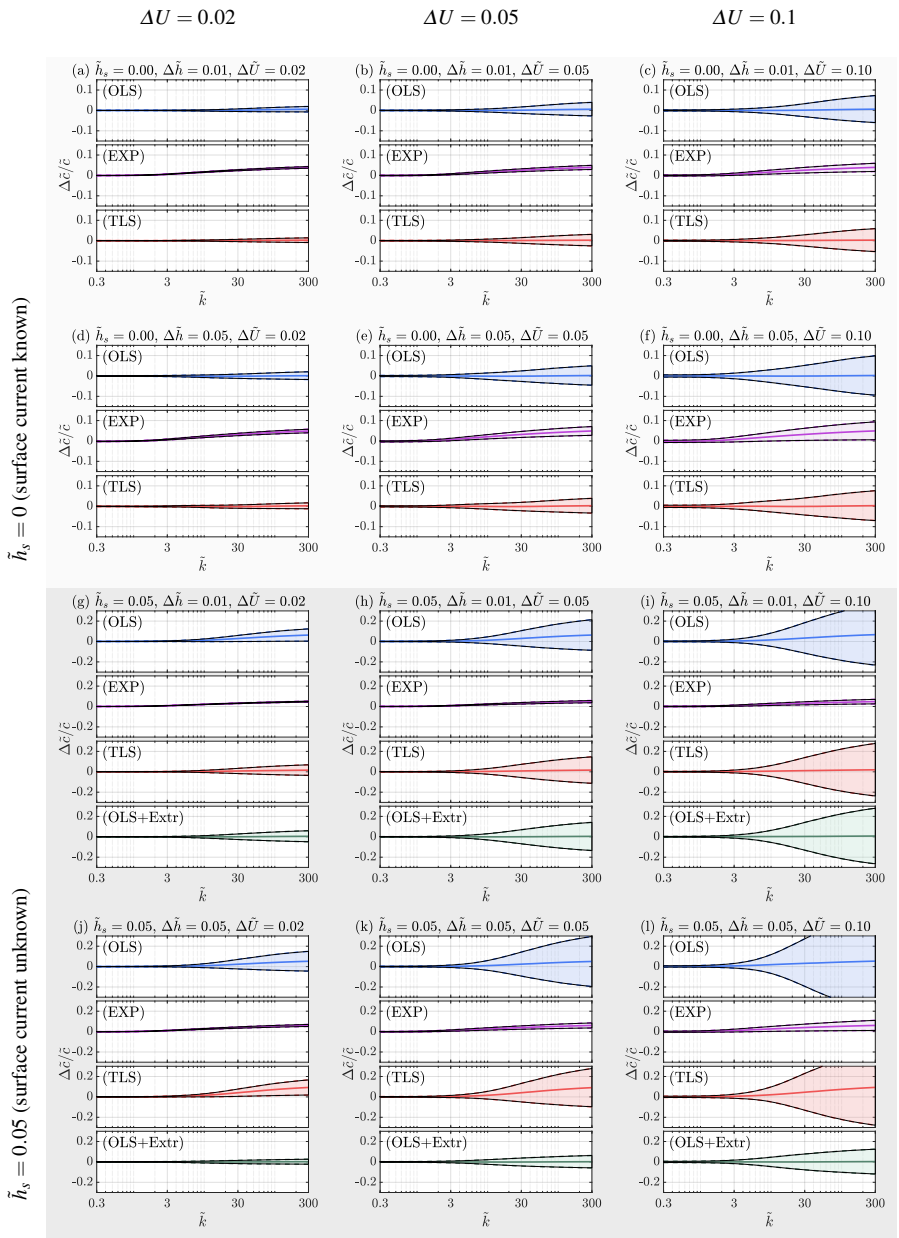

Fig. 9 Similar to Fig. 3 but for the '$1/7$' profile (6) with shear $Fr^2 = 0.05$

7.3.3 $z^{1/7}$ profile statistical moments for $\mathrm{Fr}^2 = 0.05$ *(for OLS fit)*

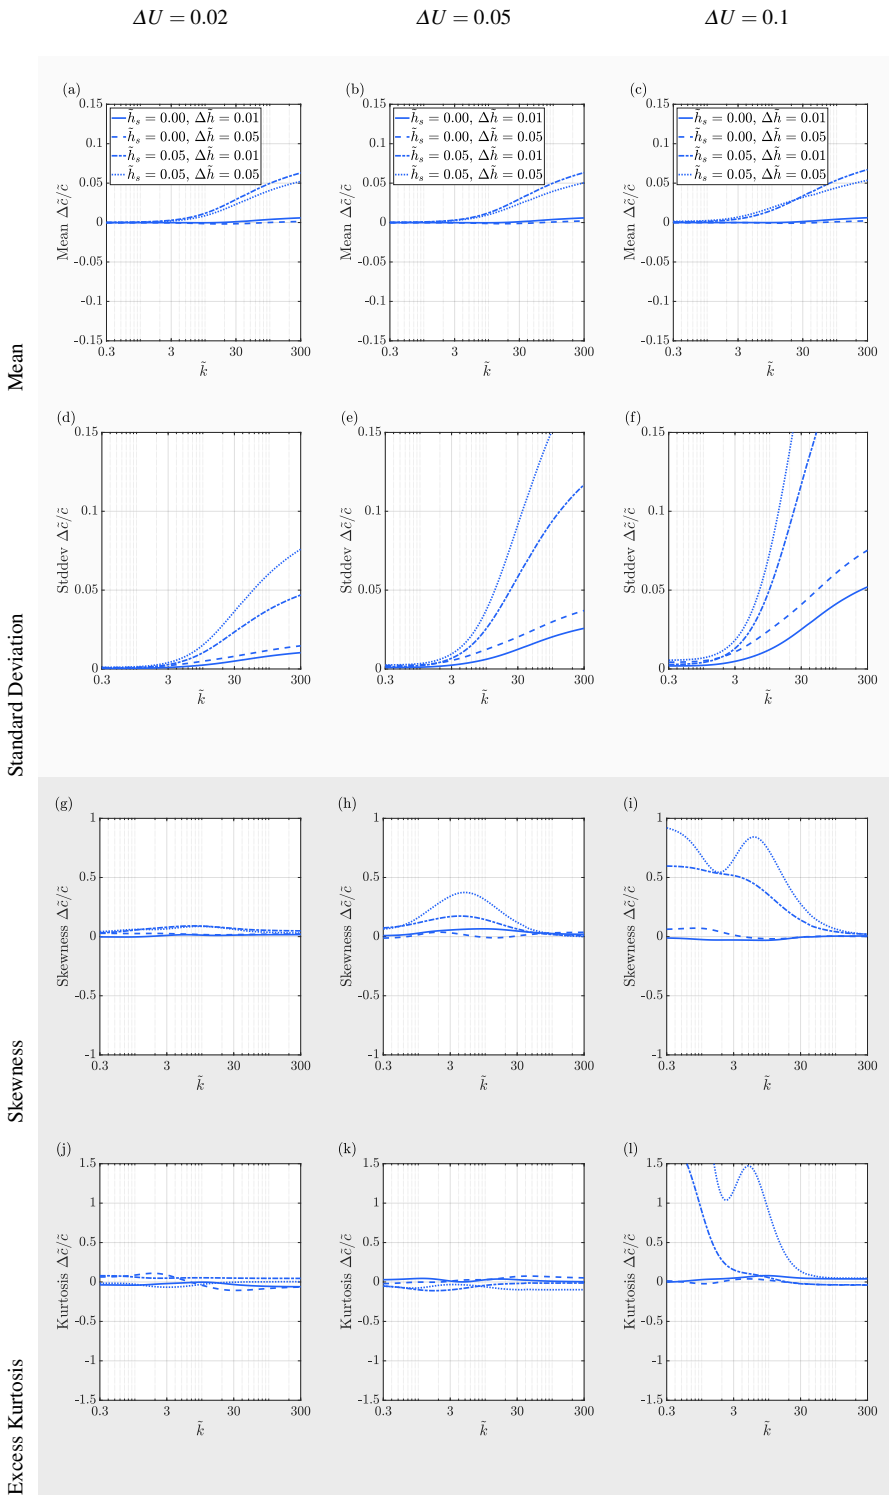

Fig. 10 Similar to statistical moments plots Fig. 5 but for the '1/7' profile (6) with $Fr^2 = 0.05$

7.3.4 $z^{1/7}$ profile deviation dependence on $\Delta\tilde{h}$

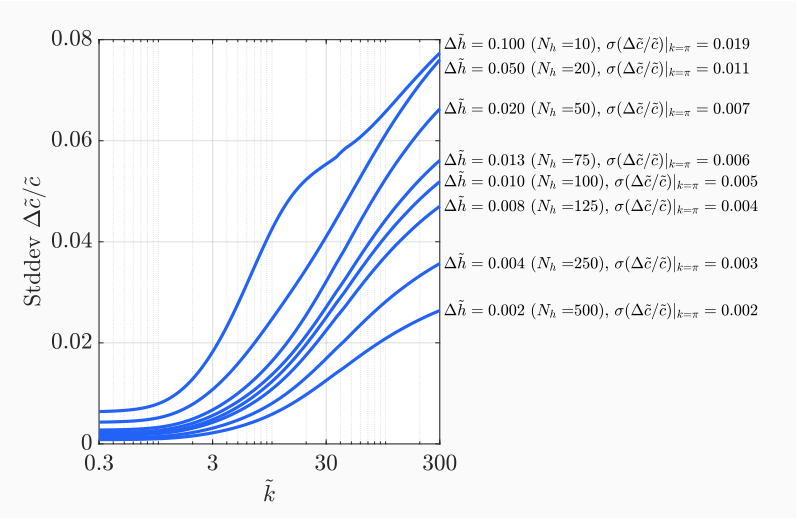

Fig. 11 The standard deviation in $\Delta\tilde{c}/\tilde{c}$ against k for several $\Delta\tilde{h}$ values for the '1/7' profile. The other parameters are kept constant: $\mathrm{Fr}^2 = 0.05$, $\Delta U = 0.1$, $\tilde{h}_s = 0$.

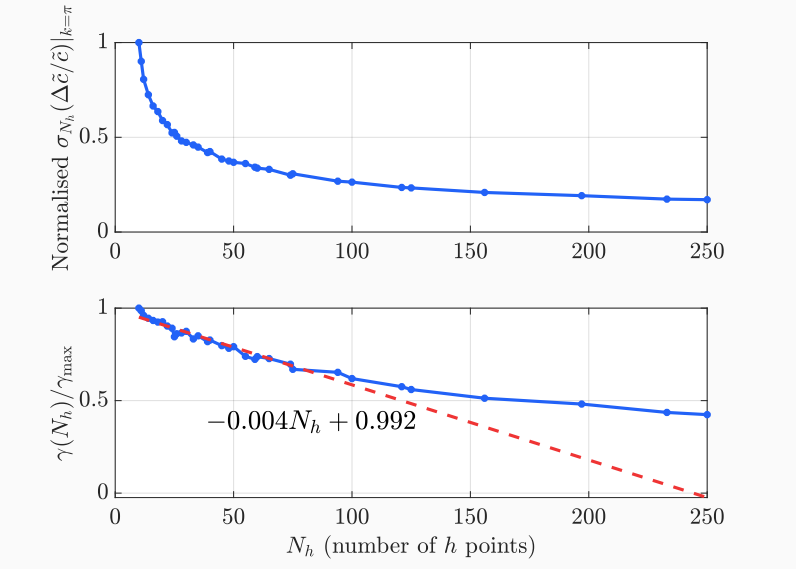

Fig. 12 The integral of the standard deviation of $\Delta\tilde{c}/\tilde{c}$ with respect to k for a range of $\Delta\tilde{h}$ values, shown here instead as the number of h points used for clarity. The other parameters are kept constant: $\mathrm{Fr}^2 = 0.05$, $\Delta U = 0.1$, $\tilde{h}_s = 0$.

7.4 Columbia River profile simulations

7.4.1 Columbia River profile percentiles

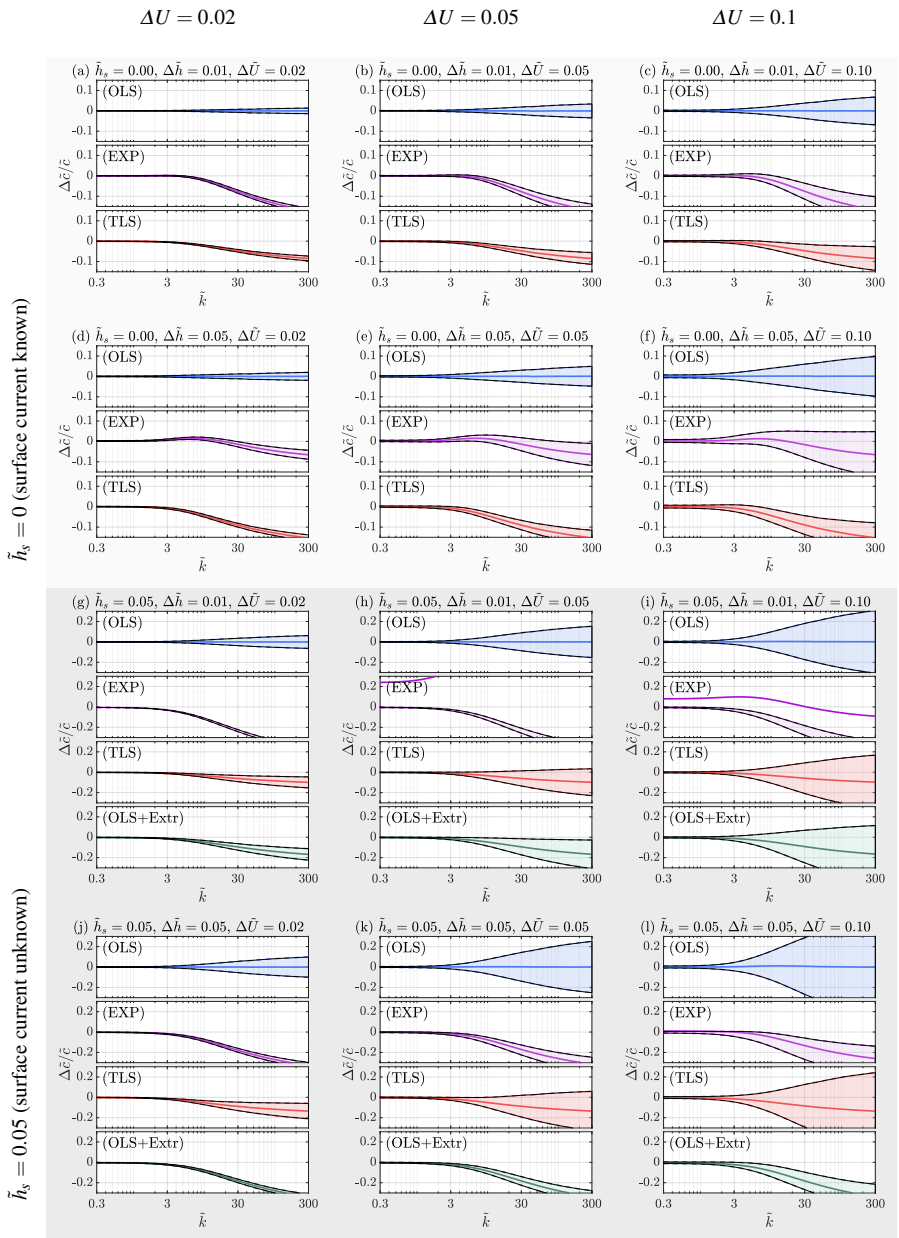

Fig. 13 Similar to Fig. 3 but using the Columbia River indicative velocity profile which has been scaled to have $Fr^2 = 0.05$

7.4.2 Columbia River profile statistical moments (for OLS fit)

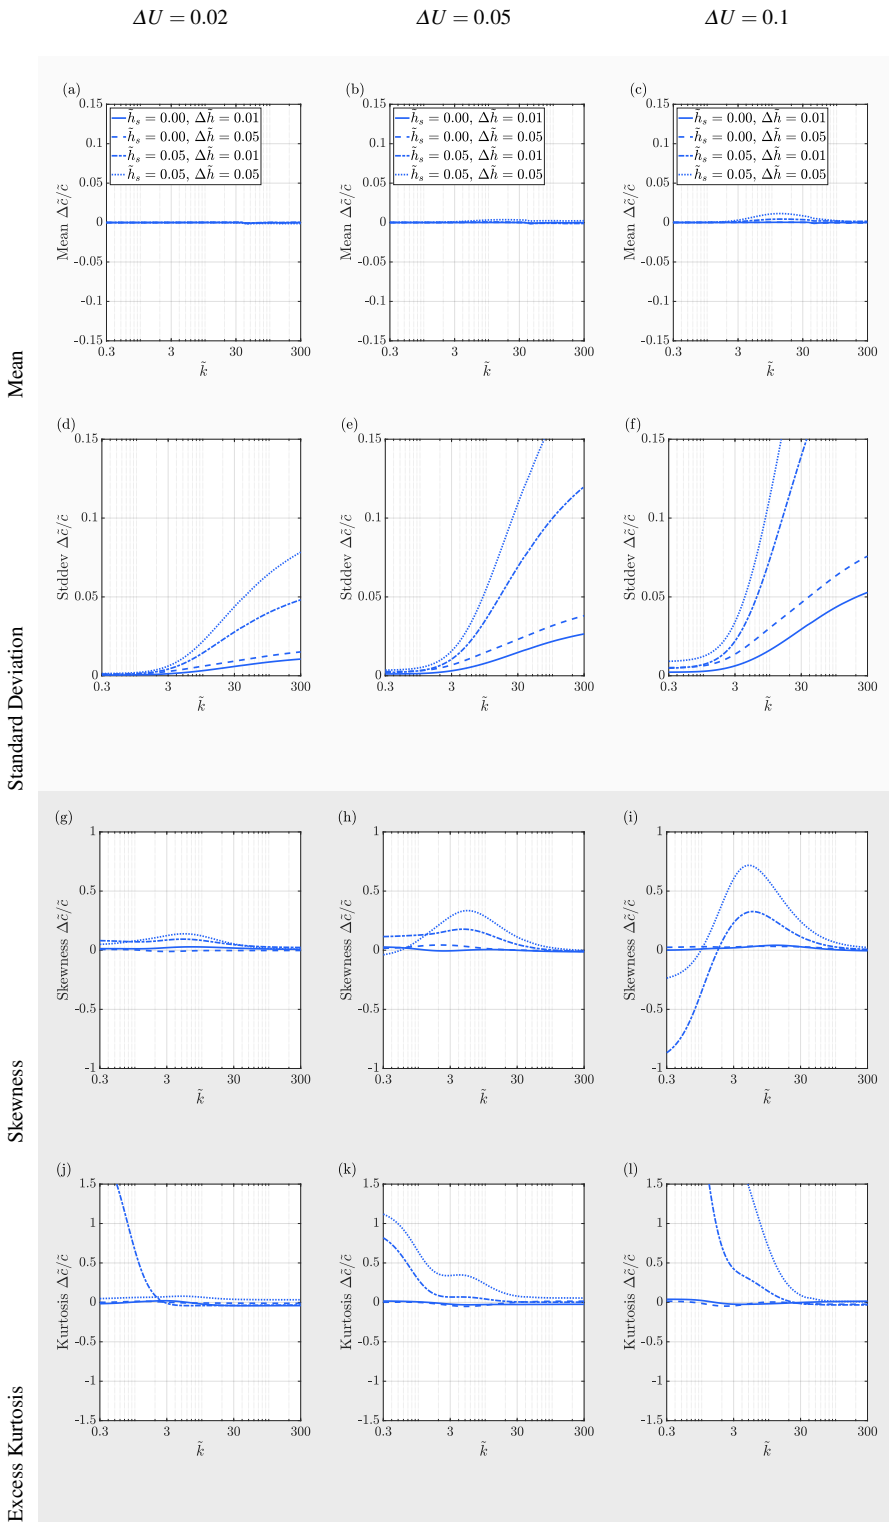

Fig. 14 Similar to statistical moments plots in Fig. 5 but for the indicative Columbia River velocity profile that has shear $Fr^2 \approx 0.04$.

7.4.3 Columbia River profile deviation dependence on $\Delta\tilde{h}$

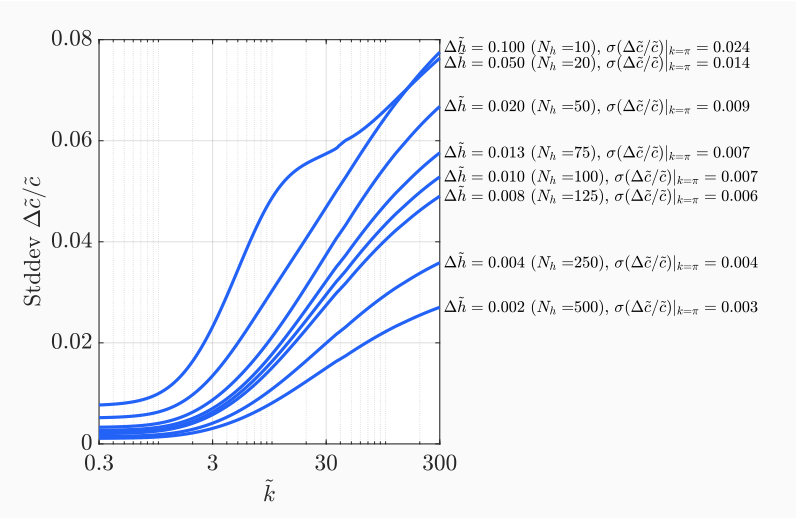

Fig. 15 The standard deviation in $\Delta\tilde{c}/\tilde{c}$ against k for several $\Delta\tilde{h}$ values for the Columbia River profile. The other parameters are kept constant: $\Delta U = 0.1$, $\tilde{h}_s = 0$.

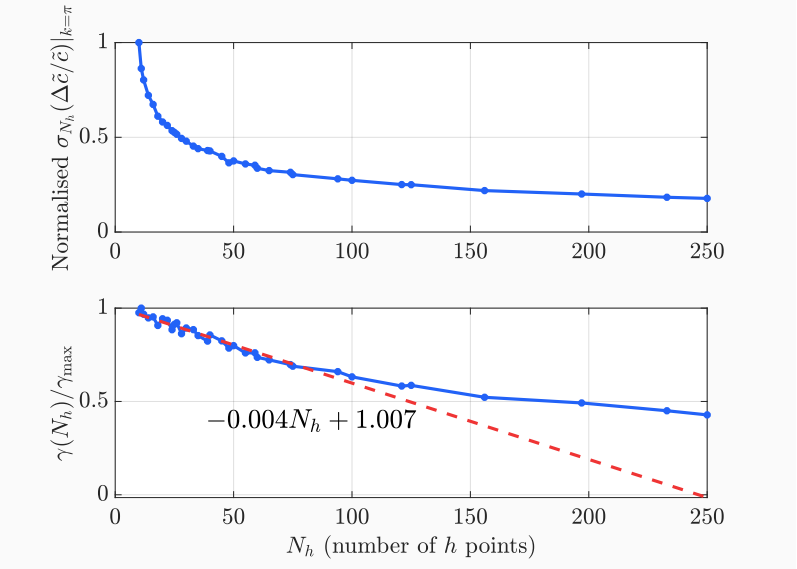

Fig. 16 The integral of the standard deviation of $\Delta\tilde{c}/\tilde{c}$ with respect to k for a range of $\Delta\tilde{h}$ values, shown here instead as the number of h points used for clarity. The other parameters are kept constant: $\Delta U = 0.1$, $\tilde{h}_s = 0$.

7.5 Pseudospectra plots

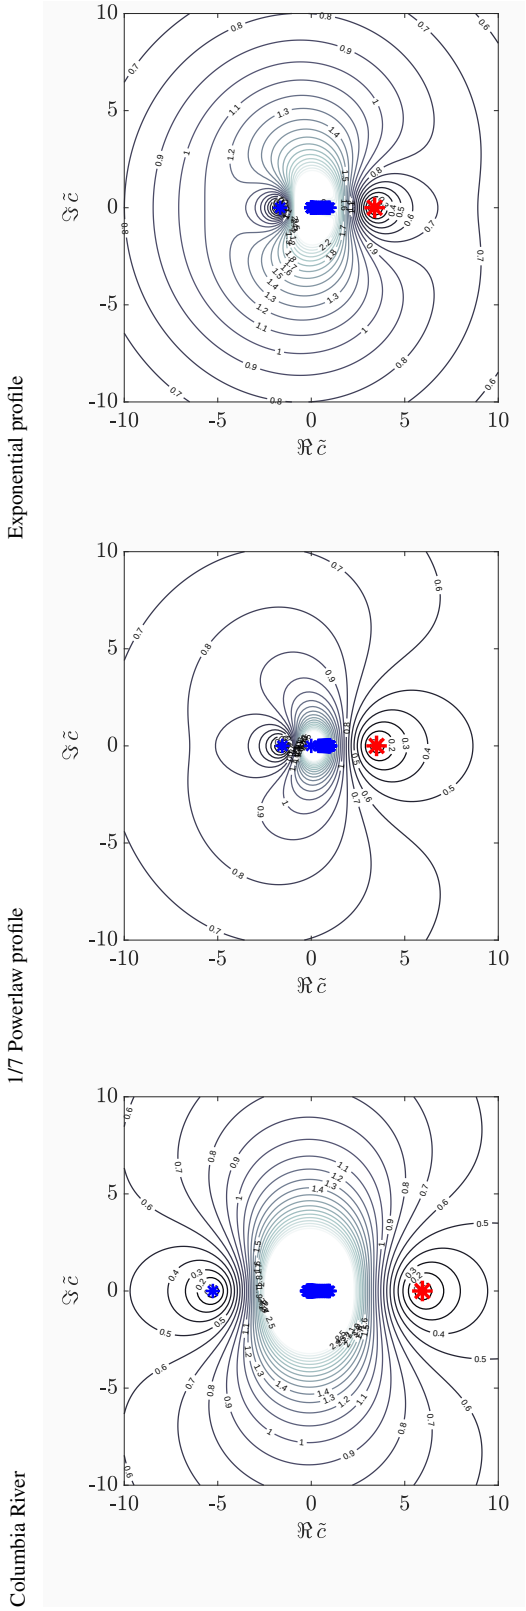

Fig. 17 Pseudospectra for our exemplar velocity profiles at $\tilde{k} = \pi$.

Acknowledgements P.M. and S.Å.E. are funded by the Research Council of Norway, grant number 249740. We are grateful to Prof James T. Kirby for suggesting this research question to us in the first place.

References

1. Banihashemi, S., Kirby, J.T., Dong, Z.: Approximation of wave action flux velocity in strongly sheared mean flows. Ocean Modelling **116**, 33–47 (2017)
2. Boyd, J.: Chebyshev and Fourier Spectral Methods: Second Revised Edition. Dover Books on Mathematics. Dover Publications (2001)
3. Campana, J., Terrill, E.J., de Paolo, T.: The development of an inversion technique to extract vertical current profiles from x-band radar observations. Journal of Atmospheric and Oceanic Technology **33**(9), 2015–2028 (2016)
4. Campana, J., Terrill, E.J., de Paolo, T.: A new inversion method to obtain upper-ocean current-depth profiles using x-band observations of deep-water waves. Journal of Atmospheric and Oceanic Technology **34**(5), 957–970 (2017)
5. Dalrymple, R.A.: Water wave models and wave forces with shear currents. Tech. Rep. 20, Coastal and Oceanographic Engineering Lab, University of Florida Gainesvilie (1973)
6. Dong, Z., Kirby, J.T.: Theoretical and numerical study of wave-current interaction in strongly-sheared flows. In: P. Lynett, J. Smith (eds.) Proc. 32nd Coastal on Coastal Engineering, 33 (2012)
7. Drazin, P.G., Reid, W.H.: Hydrodynamic Stability, 2 edn. Cambridge Mathematical Library. Cambridge University Press (2004). DOI 10.1017/CBO9780511616938
8. Elias, E.P., Gelfenbaum, G., Van der Westhuysen, A.J.: Validation of a coupled wave-flow model in a high-energy setting: The mouth of the columbia river. Journal of Geophysical Research: Oceans **117**(C9) (2012)
9. Ellingsen, S.Å., Li, Y.: Approximate dispersion relations for waves on arbitrary shear flows. Journal of Geophysical Research: Oceans **122**(12), 9889–9905 (2017)
10. Golub, G., Hansen, P., O'Leary, D.: Tikhonov regularization and total least squares. SIAM Journal on Matrix Analysis and Applications **21**(1), 185–194 (1999). DOI 10.1137/S0895479897326432
11. Higham, N.J., Tisseur, F.: More on pseudospectra for polynomial eigenvalue problems and applications in control theory. Linear Algebra and its Applications **351-352**, 435 – 453 (2002). DOI https://doi.org/10.1016/S0024-3795(01)00542-0. Fourth Special Issue on Linear Systems and Control
12. Kilcher, L.F., Nash, J.D.: Structure and dynamics of the columbia river tidal plume front. Journal of Geophysical Research: Oceans **115**(C5) (2010)
13. Kirby, J.T., Chen, T.M.: Surface waves on vertically sheared flows: approximate dispersion relations. Journal of Geophysical Research: Oceans **94**(C1), 1013–1027 (1989)
14. Kumar, N., Voulgaris, G., Warner, J.C.: Implementation and modification of a three-dimensional radiation stress formulation for surf zone and rip-current applications. Coastal Engineering **58**(12), 1097–1117 (2011)
15. Li, Y., Ellingsen, S.Å.: A framework for modelling linear surface waves on shear currents and varying bathymetry (2018). Submitted manuscript
16. Li, Y., Smeltzer, B.K., Ellingsen, S.Å.: Transient wave resistance upon a real shear current. Eur. J. Mech.-B/Fluids (in press, DOI:10.1016/j.euromechflu.2017.08.012) (2017)
17. Lighthill, J.: Waves in fluids. Cambridge University Press (1978)
18. Lund, B., Graber, H.C., Tamura, H., Collins III, C., Varlamov, S.: A new technique for the retrieval of near-surface vertical current shear from marine x-band radar images. Journal of Geophysical Research: Oceans **120**(12), 8466–8486 (2015)
19. Maxwell, P., Ellingsen, S.Å.: Path-following methods for calculating linear surface wave dispersion relations on vertical shear flows. arXiv e-prints arXiv:1905.03187 (2019). (Preprint)
20. Peregrine, D.: Interaction of water waves and currents. Advances in applied mechanics. **16**, 9–117 (1976)
21. Quinn, B., Toledo, Y., Shrira, V.: Explicit wave action conservation for water waves on vertically sheared flows. Ocean Modelling **112**, 33–47 (2017)
22. Skop, R.A.: Approximate dispersion relation for wave-current interactions. Journal of Waterway, Port, Coastal, and Ocean Engineering **113**(2), 187–195 (1987)
23. Smeltzer, B.K., Æsøy, E., Ådnøy, A., Ellingsen, S.Å.: An improved method for determining near-surface currents from wave dispersion measurements. arXiv e-prints arXiv:1904.11575 (2019)

24. Smeltzer, B.K., Ellingsen, S.Å.: Surface waves on currents with arbitrary vertical shear. Physics of Fluids **29**(4), 047102 (2017)
25. Stewart, R.H., Joy, J.W.: Hf radio measurements of surface currents. Deep Sea Research and Oceanographic Abstracts **21**(12), 1039–1049 (1974)
26. Tisseur, F., Higham, N.: Structured pseudospectra for polynomial eigenvalue problems, with applications. SIAM Journal on Matrix Analysis and Applications **23**(1), 187–208 (2001). DOI 10.1137/S0895479800371451
27. Trefethen, L.: Spectral Methods in MATLAB. Software, Environments, and Tools. Society for Industrial and Applied Mathematics (2000). DOI 10.1137/1.9780898719598
28. Trefethen, L.N., Embree, M.: Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press (2005). URL http://press.princeton.edu/titles/8113.html
29. Zhang, X.: Short surface waves on surface shear. Journal of Fluid Mechanics **541**, 345–370 (2005)
30. Zippel, S., Thomson, J.: Surface wave breaking over sheared currents: Observations from the mouth of the columbia river. Journal of Geophysical Research: Oceans **122**(4), 3311–3328 (2017)