



HAL
open science

Ensemble downscaling of a regional ocean model

Malek Ghantous, N.K. Ayoub, Pierre de Mey-Frémaux, Vassilios Vervatis,
Patrick Marsaleix

► **To cite this version:**

Malek Ghantous, N.K. Ayoub, Pierre de Mey-Frémaux, Vassilios Vervatis, Patrick Marsaleix.
Ensemble downscaling of a regional ocean model. *Ocean Modelling*, 2020, 145, pp.101511.
10.1016/j.ocemod.2019.101511 . hal-03052959

HAL Id: hal-03052959

<https://hal.science/hal-03052959v1>

Submitted on 11 Dec 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Ensemble downscaling of a regional ocean model*

Malek Ghantous^{†a}, Nadia Ayoub^a, Pierre De Mey-Frémaux^a, Vassilios Vervatis^b, Patrick Marsaleix^c,

^a*LEGOS, Université de Toulouse, CNES, CNRS, IRD, UPS*

^b*Faculty of Physics, University of Athens*

^c*LA, Université de Toulouse, CNRS*

Abstract

We downscaled a free ensemble of a regional, parent model to a high-resolution coastal, child ensemble in the Bay of Biscay. The child ensemble was forced at the open boundaries by the parent ensemble, and locally by perturbing the winds. By comparing ensembles generated by each of these forcings perturbations separately and combined we were able to consider the ensemble from either of two paradigms: (1) characterising high-resolution, coastal model errors using local and non-local forcing perturbations, or (2) downscaling regional model errors into the coastal domain. We found that most of the spread in the child ensembles was generated from the ensemble of open boundary conditions, with the local wind perturbations on their own generating substantially less ensemble spread. Together, the two sources of error increased the ensemble spread by only a small amount over the non-local perturbations alone. In general, the spread in sea surface height was greater in the child ensembles than in the parent ensemble, probably due to the more refined dynamics, while the spread in sea surface temperature was lower, likely due to the way the open boundary conditions were averaged. Deep below the surface, though, the child ensemble featured a large spread even where the parent model's spread was very weak. This enhanced error response is a promising result for an ensemble data assimilation system, as it could be exploited to correct the model deep below the surface.

*Manuscript accepted for publication in Ocean Modelling, doi: <https://doi.org/10.1016/j.ocemod.2019.101511>. This preprint is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

[†]Now at Météo France, Toulouse; Corresponding author

Keywords: Ensemble modelling, downscaling, model uncertainties, coastal modelling, Bay of Biscay, data assimilation

1. Introduction

We downscale geophysical models so that we can increase the resolution in a region of interest, without incurring the computational cost and the difficulty of increasing the resolution everywhere in the model domain. This modest aim leads to anything but a straightforward problem to solve, though, for the difficulties associated with nesting models are manifold, and the high-resolution child model cannot simply be regarded as a more detailed version of the parent model. A corollary of this is that the errors in the parent model cannot be expected to translate in a straightforward way into errors in the child model. It is this problem—the downscaling of model *errors*—which we concern ourselves with here, specifically in ocean circulation models.

The promise of higher resolution models to provide more detailed descriptions of the ocean balances against some inconvenient physics at small scales. As Sandery and Sakov (2017) show, while increased detail is possible, and useful, this does not necessarily mean increased accuracy, with the higher complexity of the small-scale flow even deteriorating the large-scale flow through an inverse energy cascade. There are further problems to address with the open boundary problem. Quite apart from the difficulty of matching solutions of different spatial and temporal scales, the internal variability of the downscaled model is decoupled from that of the parent model, which can lead to divergent solutions—an undesirable trait when the parent model outperforms the child. Sometimes this divergence can be redressed, as Katavouta and Thompson (2016) showed by applying a spectral nudging technique in the deeper regions. This maintained the parent model’s accuracy there, while giving free reign to the child model in the coastal region, where the parent model was unable to resolve many of the processes. With all this in mind, it would be fruitful then to explore the behaviour of errors in a downscaled model system.

In this study, we attempt to characterise a downscaled coastal model’s errors from ensembles of perturbed simulations. We do this principally by examining the ensemble’s statistics—above all the ensemble variance or, equivalently, its standard deviation, which is often referred to as the ensemble’s *spread*. There are many sources of error in a model, due to deficiencies in

34 the physical model, the numerical scheme and discretisation, and errors in
35 the initial and boundary conditions and the forcing fields. By narrowing our
36 focus onto a subspace of the full model error space, we can better under-
37 stand the sensitivity of the model to a specific set of error sources. For us,
38 this subset consists only of surface winds and the open boundary conditions.
39 Understanding the model error subspace will guide us in future efforts in
40 assimilating data, where corrections to the model will only be possible for
41 variables sensitive to these two sources.

42 Experimenting with forcing error perturbations in ocean models is not
43 new, of course, even for nested models. To name but a few studies, Echevin
44 et al. (2000) perturbed the initial conditions while leaving all other forcing
45 variables untouched; Auclair et al. (2003) examined the effect on a coastal
46 model of perturbing, in turn, the initial density field, a prominent slope
47 current in the domain, the wind field and river discharges; Jordà and De Mey
48 (2010) perturbed the wind and the initial, analytic velocity profile; Kim et al.
49 (2011) perturbed the winds, the initial conditions and the open boundary
50 conditions, the last of these by calculating EOFs for the boundary conditions
51 and, with random amplitudes, superimposing a linear combination of EOF
52 modes onto them (in fact, the same basic procedure we and others have
53 used to perturb the wind). The distinguishing feature of the work we are
54 presenting is thus the use of an ensemble of open boundary conditions. By
55 defining our open boundary perturbations from a pre-existing parent model
56 ensemble, we achieve something similar to Kim et al. (2011) as far as the
57 coastal model error space is concerned, but we gain the ability to compare our
58 coastal model errors with the parent model errors—in particular, our coastal
59 model’s error space may be thought of as representing a refinement of the
60 parent model’s error space. At this point, the problem may be viewed from
61 two equivalent perspectives: one is that of characterising coastal model errors
62 by generating a coastal model ensemble, and the other is that of downscaling
63 errors from a larger-scale (parent) model ensemble. Though philosophically
64 different, functionally these two paradigms amount to the same thing, as
65 long as the former is approached like the latter, using an ensemble of open
66 boundary conditions drawn from the parent model ensemble.

67 Since the parent ensemble was generated by perturbing solely the wind,
68 and we further perturb the local winds in the coastal ensemble, it is evident
69 that the only forcing variable directly perturbed is that of the wind. Thus the
70 open boundary conditions provided by the parent ensemble can be considered
71 to be a conduit through which the effects of non-local wind perturbations can

72 enter the coastal domain. It may seem limiting to use only one error source
73 to generate our coastal ensemble, but the implicit assumption underlying
74 this choice—that the model is perfect and that all the errors come from
75 the wind forcing—has the advantage that it enables us to concentrate on
76 understanding the consequences of one source of error. And indeed the wind
77 turns out to be a good choice, as Vervatis et al. (in review, 2019) found
78 that perturbing the wind in a regional model (the very same regional model
79 configuration as the parent model used in this study) had a much greater
80 effect than the other physical and biogeochemical perturbations they tested.
81 Being restricted to wind perturbations also allow us to compare like with like
82 when comparing the parent and child ensembles. To that end, we follow in the
83 footsteps of other studies where, in those too, only the wind was perturbed
84 (for example those by Le Hénaff et al. (2009) and Kourafalou et al. (2015)).

85 There are several ways one might choose to categorise the ensemble error
86 behaviour, the most obvious being with gross quantities like calculating the
87 mean ensemble spread over time. But it is informative to look beyond bulk
88 quantities and see where, as well as by how much, the spread changes. To
89 address this we subdivide the domain into deep-water, continental shelf and
90 intermediate zones. We also generate surface and cross-section maps to see
91 what the error pattern looks like at a given point in time and space, both
92 at the surface and below the surface. Different variables should be expected
93 to behave differently, too, and this should also depend on the regime; near
94 the river mouth we should see strong errors in salinity and temperature, but
95 sea surface height ought to vary most in deep water where the large eddies
96 dominate in the signal. We therefore consider these variables separately and
97 compare them between the parent and child ensembles.

98 After a brief and general description of the Bay of Biscay (section 2),
99 and of the regional parent and coastal child models (section 3), we discuss in
100 some detail how the wind perturbations were calculated and how we used the
101 open boundary ensemble to generate our ensembles (section 4). The bulk of
102 the paper is given over to an exploration of these ensembles (section 5). We
103 treat the surface extensively before diving beneath, showing how surface wind
104 perturbations can lead to large errors in specific zones under the surface, and
105 how this is also model-dependent. Because of its highly non-isotropic nature,
106 the behaviour of the river plume when subjected to wind perturbations holds
107 its own interest, so we treat this separately (section 5.5) before concluding
108 with a general discussion.

109 **2. The Bay of Biscay**

110 We shall present only a very brief summary of the major physical pro-
111 cesses in the Bay of Biscay, and in particular in the small domain at the
112 heart of this study. We invite the reader to consult the cited references for
113 more detail.

114 The Bay of Biscay is really a gulf situated in the eastern Atlantic Ocean,
115 with its land boundaries formed by Spain's northern coast and France's west-
116 ern coast. The Bay of Biscay and its bathymetry are shown in figure 1, which
117 also shows the domain boundaries for the regional parent model used in our
118 experiments, and the smaller coastal domain. The continental shelf is clearly
119 visible in the figure. It can be as narrow as about 20 km on Spain's north-
120 ern coast near Bilbao, but it broadens greatly as one travels north along
121 the French coast, and less so westward along the Spanish coast. For our
122 small coastal domain, shown in figure 2, the important submarine features
123 are the Gouf de Capbreton, a striking submarine canyon, and the Plateau
124 des Landes, which is really a gentle slope descending from the shelf into the
125 abyssal plain. Although the only river directly modelled in our domain is
126 the Adour, which meets the sea near Bayonne, the Gironde's plume enters
127 from the domain's northern boundary and is a major source of fresh, cool
128 water along the coast. The discharge from both rivers is significantly cooler
129 than the water in the Bay, so the river plumes can be identified from the sea
130 surface temperature maps, a fact we exploit in this study.

131 Even though our model domain covers only a small corner of the Bay of
132 Biscay, an array of different physical processes may be found therein. There
133 are the river plumes of cooler, fresh water, which are an important part of
134 the dynamics over the shelf; the river discharge is greatest during the win-
135 ter months (Puillat et al., 2006). An intermittent along-shore current, the
136 Iberian Poleward Current, extends along the Spanish coast in late Autumn
137 and Winter, and occasionally follows the slope as it turns northward (Ru-
138 bio et al., 2013; Kersalé et al., 2016). This current is also associated with
139 mesoscale, anticyclonic slope-water eddies (SWODDIES), which are provoked
140 by bathymetric features such as undersea canyons, which interrupt the cur-
141 rent's flow and send jets of slope water into deep water regions (Pingree and
142 Le Cann, 1992).

143 Of particular interest for this study, the Plateau des Landes plays host to
144 some intense mesoscale activity, with cyclones and anti-cyclones of between
145 50 and 80 km diameter (Le Cann and Serpette, 2009; Solabarrieta et al.,

146 2014). Beneath the surface, incursion of water masses such as the Mediter-
147 ranean Intermediate Water can lead to strong and deep sub-surface gradients
148 in temperature and salinity (van Aken, 2000).

149 The barotropic tidal amplitudes in the Bay of Biscay can be extremely
150 large, over 1.5 m in some places over the shelf, dominated by the semi-diurnal
151 lunar tide (M2) (Le Cann, 1990). As the tidal current crosses the shelf break,
152 quite complex baroclinic flows and internal tides can be generated, some of
153 whose energy dissipates by adding to the vertical mixing (Piraud et al.,
154 2010). In our model domain, because the shelf is narrower, the tidal currents
155 are not as strong as they are further north (Le Cann, 1990; Toubanc et al.,
156 2018).

157 **3. The ocean circulation models and their configurations**

158 Before we discuss ensemble generation, we shall say a few words about
159 the models themselves. We give only brief descriptions of them here as
160 they have been more thoroughly described elsewhere. We refer to the high-
161 resolution coastal model as the “child”, or coastal, model and the regional
162 model supplying the open boundary conditions as the “parent”.

163 The parent ensemble is the free ensemble produced by Vervatis et al.
164 (2016) and described fully in their paper. It was generated with the NEMO
165 model (Madec and the NEMO team, 2008) on a $1/36^\circ$ grid on a domain
166 encompassing the Bay of Biscay and the western part of the English Channel.
167 This configuration, named BISCAY36, was based on that of Maraldi et al.
168 (2013).

169 The coastal model is Symphonie, which is described in a series of papers
170 by Marsaleix et al. (2006, 2008, 2009a, 2011, 2012); Estournel et al. (2005,
171 2009) and Piraud et al. (2008). Over the last decade, the Symphonie model
172 has been used in several different configurations to model the Bay of Bis-
173 cay: see the studies of Piraud et al. (2008) and Piraud et al. (2010) for
174 the barotropic and internal tides; Le Hénaff et al. (2009) for the general cir-
175 culation; Herbert et al. (2011) on the Iberian Poleward Current along the
176 northern Spanish coasts; Herbert (2012) for the surface circulation response
177 to storm Klaus in January 2009 and, recently, Toubanc et al. (2018) for the
178 impact of tides on the open-boundary conditions. These studies have shown
179 that Symphonie was able to simulate the main circulation patterns in the
180 Bay of Biscay from daily to seasonal time scales.

181 The domain, whose limits can be seen in the bathymetry plots in fig-
182 ure 2, is much smaller than the parent model’s—it forms a box over the
183 southeastern corner of the Bay of Biscay extending northward up to just
184 beneath Arcachon on France’s southwest coast, and westward ending a lit-
185 tle before Santander on Spain’s northern coast (figure 1). The resolution
186 is set to 500 m between grid points in both horizontal directions, a sub-
187 stantial improvement over the parent model’s roughly 2.5 km resolution.
188 This enables it to better resolve mesoscale eddies—perhaps even allowing
189 the resolution of sub-mesoscale eddies—fronts and river plumes. The model
190 bathymetries are different, too: for the coastal model we used a new, high-
191 resolution bathymetry (Lyard, 2015), shown in figure 2. While the parent
192 model defines depth with z values, the coastal model uses σ layers. We refer
193 to the specific configuration of the Symphonie model in this high-resolution
194 coastal domain as BOBLAND.

195 The meteorological fields for both model configurations are the 3-hourly
196 operational analyses from the European Centre for Medium-Range Weather
197 Forecasts’ (ECMWF) high-resolution model, and are thus identical for the
198 parent’s and child’s unperturbed runs.

199 The child model is not two-way nested in the parent model, but rather
200 uses an independent run of the parent model to define its open boundary
201 conditions (that is, “offline”). The method for handling the open boundaries
202 is clearly of consequence for this study, but we defer this to the discussion
203 on downscaling in section 4.2.2.

204 A summary of some of the features and configuration parameters for both
205 models is given in table 1.

206 In this paper we perform sensitivity studies, so we have not sought to
207 verify in detail the realism of the model beyond what these prior studies
208 have done. We nevertheless thought it prudent to check that our results
209 were in line with observations. Unfortunately, very few observations were
210 available to us for our period of study in our domain. We have used the L3S
211 daily satellite SST fields (Orain, 2016) but the coverage is very poor due to
212 the frequently cloudy conditions in the Bay of Biscay in winter. Comparisons
213 of the unperturbed runs for both the parent and child with those days with
214 good coverage showed few differences of more than half a degree Celsius.
215 We have also checked that our results are qualitatively consistent with the
216 information found in the literature. For instance, the observations of SST
217 and HF radar presented by Solabarrieta et al. (2014) show a large eddy
218 similar to the one that persists throughout the study period. While the size

219 and location of the eddy are consistent between the coastal model and the
220 data, the parent model’s eddy properties seem less realistic. Also, we found
221 that the SST gradient over the eastern shelf and at the edge of the Adour’s
222 plume was in good agreement with Solabarrieta et al.’s 2014 observations;
223 along the Spanish shelf it was however too warm, maybe because of a local
224 recirculation (not shown).

225 4. Ensemble generation

226 One of our main assumptions is that the model errors result principally
227 from errors in the forcing variables. If we know the distribution of these
228 errors, or we can guess at it, we can generate an ensemble of perturbed
229 forcing variables which follow it. If we then apply this ensemble of perturbed
230 forcings to the model we can generate an ensemble of model realisations,
231 each member corresponding to a particular forcing perturbation. This model
232 ensemble can then provide a useful estimate of the model error covariances.
233 Strictly speaking, of course, the estimate is only of the subspace of error
234 covariances related to the forcing variables we perturb, but if these forcing
235 variables dominate the others then this is adequate. In this study, where
236 our focus is on downscaling an ocean circulation model, we only directly
237 perturb the winds and the open boundary conditions, although these latter
238 are themselves generated from perturbations to the parent model’s winds.
239 Vervatis et al. (in review, 2019) found that the wind was, and by quite some
240 way, the most important source of error among those that they tested.

241 We describe in turn the procedure for perturbing the wind locally and
242 how we utilised the ensemble of open boundary conditions.

243 4.1. The wind perturbations

244 We followed Auclair et al. (2003) and Le Hénaff et al.’s 2009 method for
245 perturbing the wind field, which is also the method used by Vervatis et al.
246 (2016). Our basic assumption is that the uncertainties in the ECMWF wind
247 fields have the same spatial and temporal structure as the wind variabil-
248 ity. Consequently, the perturbations are built from the empirical orthogonal
249 functions (EOFs) of the ECMWF wind fields in the BOBLAND domain over
250 the period desired—for this study, this is December 2011 to March 2012.

251 Let \mathbf{U} be the matrix containing both zonal and meridional wind velocities.
252 It may be decomposed into its temporal mean and an anomaly matrix,

$$\mathbf{U} = \overline{\mathbf{U}} + \mathbf{U}'. \quad (1)$$

253 Now for a matrix \mathbf{U}' containing both zonal and meridional wind anomalies,
 254 and of dimension $k \times l$ where k and l are the temporal and spatial dimensions
 255 respectively, we may perform a singular value decomposition giving us the
 256 expression

$$\mathbf{U}' = \mathbf{X}\mathbf{\Lambda}\mathbf{Y}^T, \quad (2)$$

257 where $\mathbf{X} \in \mathbb{R}^{k \times j}$ and $\mathbf{Y} \in \mathbb{R}^{l \times j}$ are composed of orthogonal vectors and
 258 $\mathbf{\Lambda} \in \mathbb{R}^{j \times j}$ is diagonal, whose non-zero elements σ_j are the singular values of
 259 \mathbf{U}' . The dimension j corresponds to the number of modes, of which only
 260 the first few are computed, using the method described by Toumazou and
 261 Cretaux (2001). To decide how many to keep, we can use the singular values
 262 from $\mathbf{\Lambda}$ to determine what percentage P_j of the total variance each mode j
 263 represents:

$$P_j = \frac{\sigma_j^2}{\text{Tr } \mathbf{\Lambda}^2}. \quad (3)$$

264 where Tr indicates that the trace of the matrix is taken. For our EOFs
 265 the first two modes explained $\sim 64\%$ and $\sim 29\%$ of the variance respectively,
 266 and there was a steep drop to the third mode at only a little more than
 267 2%, after which the percentage for subsequent modes fell gradually. Figure
 268 3 shows the temporal standard deviation of the first five modes, where the
 269 dominance of the first zonal and second meridional modes is made clear,
 270 as is the much weaker contribution from higher-order modes. We therefore
 271 decided to keep only these first two modes, together explaining about 93%
 272 of the total variance.

273 The EOF modes may be recombined with random numbers $\epsilon_{i,j}$, one for
 274 each ensemble member i and each spatial mode j , to produce a wind per-
 275 turbation which is added to the original wind field. Following Lucas et al.
 276 (2008), we choose the random numbers from a Gaussian distribution, in pairs
 277 of positive and negative numbers of the same absolute value. The probability
 278 density function is then

$$p(\epsilon) = \frac{1}{\sqrt{2\pi s^2}} e^{-\frac{\epsilon^2}{2s^2}}, \quad (4)$$

279 where s is the standard deviation. The perturbation can be expressed as

$$\tilde{\mathbf{U}}_i = \sum_{j=1}^2 \epsilon_{i,j} \mathbf{x}_j \sigma_j \mathbf{y}_j^T, \quad (5)$$

280 and the wind field for ensemble member i is then given by

$$\mathbf{U}_i = \mathbf{U} + \tilde{\mathbf{U}}_i. \quad (6)$$

281 If our ocean model were linear, the pairs of Gaussian random numbers would
282 guarantee that the ensemble of simulations be Gaussian with zero mean. Fol-
283 lowing Vervatis et al. (2016), we assume an error standard deviation of $s = 0.3$
284 for equation (4), reflecting a presumed observational error of 30% for the
285 wind. (It's worth mentioning that Le Hénaff et al. (2009) assumed an error
286 of 0.2, but chose the value of 0.3 to deliberately inflate the ensemble disper-
287 sion. Since there are many other sources of error we have not perturbed,
288 and noting that our ensembles are not especially long either, we feel that the
289 choice is reasonable for our configuration, whatever the real error value.)

290 To prevent the wind perturbations from being correlated between mem-
291 bers, the random numbers must be replaced by a new set generated with
292 the same Gaussian and pairing criteria, and uncorrelated with previous sets.
293 All the same, the rate at which we change them should not be too high,
294 otherwise coherent weather patterns would not have enough time to develop.
295 Vervatis et al. (2016) and Le Hénaff et al. (2009) chose to update the random
296 numbers every five days. Their choice also corresponded to the decorrelation
297 time of the wind EOFs. Although their EOFs are on a much larger spatial
298 scale than ours, ultimately it is still an arbitrary decision, so we did not seek
299 to replicate this methodology for our reduced domain, preferring instead to
300 follow their example and use five days as well. To avoid shocks when up-
301 dating from one set of random numbers to the next we linearly interpolated
302 them in time, ensuring a smooth evolution from one perturbation to the next
303 for all times.

304 *4.2. An ensemble of open boundary conditions*

305 *4.2.1. Generating the parent ensemble*

306 Downscaling a regional model to a high resolution coastal model implies
307 using the regional, parent model to prescribe the open boundary conditions
308 for the coastal, child model. In our study we want to take this a step further:
309 we want to downscale a regional model *ensemble*, so for each member of our
310 coastal ensemble we use a different member of the parent ensemble to force
311 the open boundaries. From the point of view of the coastal ensemble, this
312 is equivalent to saying that its open boundaries are being perturbed. Other
313 ways of perturbing them could be envisaged, such as the methods tested by

314 Jordà and De Mey (2010), but by using the parent ensemble we maintain the
315 dual paradigm of ensemble downscaling and—equivalently—perturbation of
316 the open boundary conditions. For the case where we perturb the winds
317 locally as well, this adds extra degrees of freedom to our ensemble.

318 Our parent ensemble is the regional free ensemble described by Vervatis
319 et al. (2016), generated, as we have said, by perturbing the ECMWF winds
320 using the same EOF method we have just presented. However, their EOFs
321 were calculated over the whole BISCAY36 domain, for a different period from
322 ours and from daily averages of the wind fields as opposed to the 3-hourly
323 fields we used. The temporal standard deviations of the first five of their
324 EOF modes for the BISCAY36 configuration are shown in figure 4, where the
325 view has been restricted to the BOBLAND domain. If we compare them to the
326 EOFs we calculated for the BOBLAND configuration, shown in figure 3, we see
327 that although both EOFs are dominated by the first zonal mode, thereafter
328 they are markedly dissimilar. The BOBLAND EOFs were calculated only over
329 the BOBLAND domain, so we see a slightly more complex structure, and the
330 second meridional mode is stronger relative to its neighbouring modes—first
331 meridional and second zonal—than in BISCAY36. Vervatis et al. (2016) kept
332 the first six EOF modes to generate their wind perturbations, but even by
333 only keeping the first two modes in BOBLAND we might expect stronger and
334 more varied wind perturbations, both because of the stronger first zonal and
335 second meridional modes, and because of the more structured second zonal
336 and first meridional modes. We test this hypothesis in section 5.3.3.

337 There are two other differences between the way we and Vervatis et al.
338 (2016) handled the time-dependent component of the perturbations. Firstly,
339 they didn't use every index from the EOF time series like we did, but instead
340 every fifth one, interpolating them the same way as the random numbers.
341 Secondly, when we started the ensemble runs we set the first random number
342 to zero for all members, so that the first perturbation would ramp up from
343 zero. The omission of this step created a spike in the parent ensemble's
344 spread (discussed in section 5.3.1).

345 *4.2.2. The downscaling method*

346 Some thought was needed about the manner in and the extent to which
347 the regional model should influence the coastal model. The regional model's
348 fields were provided as daily averages, where the mean was taken over 25
349 instantaneous fields spaced an hour apart. The fields used to force the open
350 boundaries of the child model are temperature, velocity, and sea surface

351 height. The “25-hour” averaging for the parent ensemble’s fields does a fairly
 352 good job of removing the tides from the sea surface height and velocity fields,
 353 but some signals are still present, principally an artefact due to aliasing. This
 354 artefact has a period of around 14 to 15 days, with an amplitude of about
 355 a couple of centimetres in deeper water. It acts as a forcing on the coastal
 356 model, producing likewise a signal with a fortnightly oscillation of up to about
 357 4 cm over the shelf or 2 cm in deep water. A more thorough discussion of
 358 the problems associated with the way these open boundary conditions have
 359 been averaged was presented by Toubanc et al. (2018), though the context
 360 is that of a much larger scale coastal model.

361 The boundary condition in the Symphonie model is of the radiative
 362 Flather type. (For a complete description on the way Symphonie handles
 363 its open boundary conditions, we refer the reader to Marsaleix et al. (2006)
 364 and Marsaleix et al. (2009b).) We use a nudging layer of 30 grid points for
 365 the main experiments, corresponding to 15 km in the BOBLAND configura-
 366 tion. At the boundary the nudging time scales are at $\tau = 0.1$ days for the
 367 barotropic velocity and $\tau = 1$ day for the baroclinic velocity. Marchesiello
 368 et al. (2001) applied a cosine half-period to reduce these values to zero, which
 369 is the method Symphonie uses, to wit,

$$\tau_r(d) = \begin{cases} \frac{\tau}{2} \left(1 - \cos \pi \frac{s_L - d}{s_L} \right), & \text{for } d \leq s_L \\ 0, & \text{for } d > s_L, \end{cases} \quad (7)$$

370 where τ_r is the reduced nudging time scale at distance d from the domain
 371 boundary, for a nudging layer of width s_L .

372 We also considered using a narrower nudging layer of only 10 grid points
 373 before settling on 30. Comparing unperturbed model runs of both nudging
 374 layer widths with sea surface temperature data, we were unable to discern
 375 much difference between the two. However, a wider nudging layer could serve
 376 to constrain the coastal ensemble spread: for example if the nudging layer is
 377 too wide, the boundary conditions will impose themselves to the detriment of
 378 any other perturbed dynamics. On the other hand, a wider nudging layer, we
 379 reasoned, should cause the open boundary conditions to have more influence
 380 further into the domain. For our application this may have benefits, since
 381 it means, all other things being equal, that the coastal model ensemble will
 382 capture more of the parent ensemble’s spread. This was desirable too in
 383 subsequent work, where we wished to preserve the large scale circulation

384 of the parent model while assimilating data at the same time (to be presented
385 in a future paper). Thus the wider nudging layer was preferred, although in
386 the sections that follow we discuss the impact of the nudging layer width on
387 the ensemble where pertinent.

388 5. Ensemble behaviour and statistics

389 5.1. *Introducing the ensembles*

390 We generated several ensembles in the BOBLAND configuration, whose
391 defining characteristics are summarised in table 2. Each one consists of 50
392 members, from which the unperturbed run is excluded. Of these ensembles,
393 we concentrate on three, plus the parent ensemble, with the others used
394 mostly to answer specific questions raised by the first three and to test some
395 of our assumptions. The four principal ensembles on which the bulk of our
396 analysis is based are the parent, WIND, OBC (Open Boundary Condition)
397 and WIND×OBC ensembles, the names corresponding to the forcing being
398 perturbed. For WIND×OBC the two types of perturbations are applied at the
399 same time. When we perturb only the wind we use the same unperturbed
400 open boundary conditions for all members, and likewise when we perturb
401 only the open boundary conditions we use the same unperturbed winds for
402 all members. The coastal ensemble runs begin after a model spinup, starting
403 on September 12, 2011 and ending on November 30. At this date all of the
404 ensemble members are identical and have the exact same open boundary
405 conditions and winds as the unperturbed run, and they diverge from then
406 on.

407 In what follows, all ensemble statistics were computed over the 50-member
408 ensemble of perturbed runs (excluding the unperturbed member). To better
409 represent and compare the ensemble statistics of each ensemble proper, with-
410 out contamination from the parent ensemble in the nudging layer, the spatial
411 averages exclude points which would fall inside the 30-point nudging layer of
412 the BOBLAND domain. Similarly, to compare our ensembles with the parent
413 ensemble, statistics for the parent were only computed over the BOBLAND
414 domain, again excluding points which fall inside the nudging layer.

415 The statistical quantity which primarily interests us is the ensemble vari-
416 ance, or equivalently the ensemble standard deviation (which is also referred
417 to in this paper and in the literature as the ensemble spread). This tells
418 us to what extent the perturbations have led the ensemble to diverge, and
419 may be interpreted as a measure of the error in the model resulting from

420 an imperfectly known forcing. (This is of particular significance in a data
 421 assimilation context: we may think of it of as indicative of how amenable the
 422 model is to correction by data assimilation, since the more error there is in
 423 the model then the more weight the observations will carry in the analysis
 424 for any given observational error.)

425 We follow Vervatis et al. (2016) in treating the first month of the ensem-
 426 bles (i.e. December) as an “ensemble spin-up” period, where the ensemble
 427 members have still not diverged enough to give useful statistics. We hence
 428 exclude the first month of our ensembles from the temporal moments we ex-
 429 amine later in this paper, but it can still be useful to look in here: one reason
 430 is that it shows how the ensemble spread grows as the model perturbations
 431 take effect. But there are also some dynamics which only, or mostly, occurred
 432 in this period.

433 *5.2. Quantifying the interaction with tides*

434 Exploring the model error subspace will guide us in future choices for data
 435 assimilation experiments. In particular, when dealing with SSH observations
 436 (such as altimetric data), the question arises whether the tidal signal should
 437 be assimilated together with the residual circulation. We therefore need to
 438 understand if the tidal dynamics are represented in the model error subspace.
 439 Indeed, tides influence and are influenced by the circulation, so perturbations
 440 on wind and open boundary conditions could in principle perturb tides. Here
 441 we seek to quantify the effect of the perturbations on the tidal elevation
 442 signals. To do this we need to remove the tidal signal from the elevation. For
 443 accuracy we run a harmonic analysis on the surface elevation, which requires
 444 a time series with a high sampling rate. The analyses are therefore done at
 445 a few representative locations where we recorded these high-frequency time
 446 series with a sampling rate of $23'15.5''$; they are marked in figure 2.

447 We can estimate the effect of the perturbations on the tides by calculating
 448 the ratio of the ensemble spread of the tidal signal, to the ensemble spread
 449 of the residual circulation. For elevation η , this may be expressed as

$$\frac{\text{std}(\eta - \eta')}{\text{std}(\eta')}, \quad (8)$$

450 where the prime indicates detiding by harmonic analysis. Figure 5 shows
 451 this ratio for each of the ten locations marked in figure 2. The highest
 452 value is about 0.16 for a point on the shelf, but most of the time the values

453 are closer to 0.05, only breaking the 0.1 threshold on a few occasions. For
454 points over deep water, the ratios are even smaller. One sixth is not a
455 completely insignificant contribution to the ensemble spread, but it is an
456 absolute maximum, and other high values only occur occasionally and over
457 the shelf. Consequently, we consider that the perturbations' effect on tidal
458 elevations is weak enough when compared with the residual circulation that
459 it can be neglected. Thus in the following sections we consider only the
460 modification of the residual circulation, using detided daily fields of SSH,
461 where the detiding is done online by subtracting the tidal forcing from the
462 model output fields before computing the daily average.

463 *5.3. Ensemble spread of surface variables*

464 Since it is at the surface that our forcing perturbations are generated
465 (bearing in mind that the parent ensemble was generated by wind pertur-
466 bations too), so it is that we begin our examination of the ensembles at the
467 surface too.

468 *5.3.1. Time series of the spatial mean spread*

469 Figures 6 and 7 show time series for the spatial mean of the ensemble
470 spread of sea surface height (SSH) and sea surface temperature (SST) re-
471 spectively. We shall focus initially on the top-left plot in these figures, in
472 which the spatial means are taken across the whole of the BOBLAND domain.
473 For the moment we discuss only the four principal ensembles: the WIND,
474 OBC, WIND×OBC and parent ensembles, all drawn in black.

475 We first address the most prominent feature of the SSH curves. Except
476 for the WIND ensemble, there is an extremely high spike in the spread at the
477 beginning of the time series. This is related to the manner in which the parent
478 model's wind perturbations were abruptly introduced. The subsequent spikes
479 are, on the contrary, not artefacts but are related to a strong, synchronous
480 sea surface height change from the parent model, which is accentuated on the
481 shelf. The SSH curves for the WIND ensemble show no such spikes as this was
482 the only ensemble not to use the parent ensemble for its open boundaries.
483 Because we treat this first month as a spin-up, the anomalous initial spike is
484 excluded from any temporal statistical calculations we make.

485 If we examine the full time series, we see that the OBC and WIND×OBC
486 ensembles have a greater overall spread in SSH than the parent ensemble, a
487 trend which is largely inverted for the SST. Meanwhile, the WIND ensemble

488 almost always has the least spread of all the ensembles, except for SST dur-
489 ing the ensemble spin-up phase. This is because so much of the SST spread
490 is generated by the cold river plume over the shelf, and during the ensem-
491 ble spin-up the parent ensemble's perturbations have not had the time to
492 propagate from the open boundaries to the near-coast, leaving the interior
493 initially unaffected unless the coastal ensemble's winds are perturbed. In a
494 similar fashion, the parent ensemble's spread is initially lower than the WIND
495 ensemble's until non-local perturbations have had time to reach the shelf.

496 That the SST and SSH should respond differently to the forcing perturba-
497 tions is not in itself remarkable (we discuss it further below), however it may
498 seem surprising at first that the parent ensemble, in which only the winds are
499 perturbed, should have a greater spread in SST than the higher-resolution
500 WIND×OBC ensemble, where the perturbations from the boundary are su-
501 perposed on the wind perturbations. There are several mechanisms that
502 might contribute to this. Firstly, the high-resolution coastal model is able to
503 produce a more complex velocity field, with more coherent eddies at smaller
504 scales. Even if the SSH fields are more distinct from member to member, this
505 could result in more horizontal diffusion of scalars like temperature, leading
506 the ensemble members to look more alike in those variables. But there is
507 more to it than that, because it also depends on how sub-grid scale diffu-
508 sion is handled in each model. Even though the tracer diffusion scheme is the
509 same for both models, the momentum diffusion is different, and the numerical
510 diffusion, resolution, time step and other features, which are specific to each
511 model, will influence horizontal diffusion. Because the differences between
512 the models run so deep, a thorough comparison is impractical. A simple
513 test would be to see the effect of increasing horizontal diffusion of tracers
514 in the coastal model alone, which we did by doubling it. The result partly
515 confirmed our hypothesis about diffusion and spread, in that the ensemble
516 spread decreased as horizontal diffusion increased. However, increasing the
517 horizontal diffusion parameter also lead to more diffusion of the velocity field,
518 in turn reducing the SSH spread. This implies that scalar diffusion cannot
519 be the sole mechanism responsible for the inverted spread behaviour.

520 To further get a sense of the relationship between the complexity of the
521 eddy field and horizontal diffusion, we looked at time series of spatial stan-
522 dard deviations for scalars and SSH. Calculating the temporal correlation
523 coefficients of these time series for the unperturbed member of both parent
524 and child model configurations was inconclusive, however: they do turn out
525 to be negative, but they are very weak, with absolute values less than 0.3 (ex-

526 cept for SST with SSH for the parent model, whose value is -0.59). Despite
527 the inconclusiveness of these findings, and the equivocal role that diffusion
528 appears to play in the different behaviour of the parent and child, in the fol-
529 lowing section we examine evidence suggesting that indeed the parent model
530 diffuses tracers rather less quickly than the coastal model.

531 The second reason for the inverted response of the SSH and SST spreads
532 in the two models has to do with the way we prescribe the open boundary
533 conditions themselves, which is as daily means. These are interpolated by
534 the coastal model for the current time step, but higher frequency information
535 is irretrievably lost. The coastal ensemble is thus being constrained into
536 a lower frequency error space than the parent ensemble by means of the
537 lower frequency forcing at the open boundaries. This latter explanation
538 leads to a contradiction, for the loss of high frequency information should
539 lead to a reduction in spread of all variables, yet it doesn't. We refer back
540 to the first reason to explain this discrepancy: the coastal model's higher
541 resolution and therefore greater ability to resolve eddies. The more defined
542 eddies are stronger and more localised in space, so their displacement must
543 impose a more marked difference between ensemble members than the weakly
544 resolved eddies of the regional model. This justifies the ensemble downscaling
545 paradigm, since despite the mechanism by which we may expect the ensemble
546 spread to be reduced—the loss of information at the open boundaries—the
547 spread of SSH nonetheless increases.

548 Finally, a big contributor to the reversed trend of SSH and SST spreads
549 appears to be directly attributable to the Adour's plume, which we consider
550 in the next section.

551 *5.3.2. Subdividing the domain*

552 In a previous ensemble study at the basin scale with the Symphonie
553 model, Kourafalou et al. (2015) emphasised the existence of distinct regimes
554 of model error over the shelf and in deep water. To find out how these regimes
555 differed in our ensembles we divided the BOBLAND domain into three sub-
556 domains according to depth. The subdivisions are as follows: deep water,
557 which we define as more than a thousand metres depth, and which encom-
558 passes the abyssal plain and much of the gently sloping Landes Plateau (top
559 right plots in the figures); an intermediate region where depths vary between
560 a thousand metres and 150 metres, which we may associate with the con-
561 tinental slope (bottom left plots); and the continental shelf with depths up
562 to 150 metres (bottom right plots). These regions are defined based on the

563 smoothed bathymetries seen by the respective models, which although not
564 precisely the same as each other are close enough (as illustrated in figure 2).
565 Time series of the spatial means of ensemble spread for each subdivision are
566 shown in the remaining plots in figures 6 and 7.

567 The SST spread is greatest over the shallow waters of the shelf, where the
568 river plumes are still coherent and can create very big and sharp differences
569 in SST. By contrast, the SSH spread is smaller here, and greater in deep
570 water, where large, spatially coherent eddies and steric effects from a deeper
571 water column have more influence. But there are some peaks in SSH spread
572 over the shelf for all ensembles except for the WIND ensemble. These peaks
573 are driven by wind-related processes in the parent model, and are discussed
574 by Vervatis et al. (2016). They manifest as large, synchronous increases in
575 the SSH everywhere, but especially over the shelf.

576 The SSH spread over the slope is most of the time somewhere in between
577 the shelf's and the deep water's. The SST shows something different, though.
578 From mid-February, the spread of all the ensembles is lower here than over
579 either the shelf or the abyssal plain. Meanwhile, at the end of December, and
580 from mid-January through to mid-February, the parent ensemble's spread is
581 far greater over the slope than that of the coastal ensembles. The great dif-
582 ference between the ensembles may be explained by looking at the surface
583 maps in figure 8, where we see that the parent model produces a more spa-
584 tially coherent river plume, and during those periods of large spread in the
585 parent ensemble it extends out some distance from the coast, being entrained
586 by the deep-water eddies. The coastal model diffuses the plume earlier, so
587 it does not extend as far out. Naturally, the further out the plume remains
588 undiffused, the more it can be deviated by the circulation, and the higher
589 the overall spread will be there than if it stayed near the coast and diffused
590 quickly. From mid-February, the plume doesn't extend as far out, with the
591 cooler river water tending to hug the coast, where the spread remains quite
592 large. The mesoscale eddies are the other major source of ensemble spread
593 at the surface, but, as is hinted in the SST maps and becomes clear when
594 looking at the SSH, these circulate mainly in deep water and do not greatly
595 influence what is happening over the slope. The slope is, in a sense, situated
596 in between two highly active regimes, insofar as SST spread goes.

597 This leaves us again to consider the greater spread in SST of the parent
598 ensemble compared with the child ensembles. Looking again at the surface
599 maps for SST and SSH in figure 8, we can see that coinciding with the large
600 eddy's centre (high SSH) are higher temperatures. These temperatures are

601 not only higher for the parent model, but the gradient is also steeper; in
602 other words, while the parent model’s eddy has a less distinct signature in
603 SSH than the coastal model, it has a conversely more distinct signature in
604 SST. This distinct signature naturally leads to a higher spread in SST for
605 the parent ensemble, and once more suggests that the parent model diffuses
606 tracers more weakly than the coastal model.

607 We also produced plots for the sea surface salinity (SSS) but they do
608 not add much to the argument, so we do not present them. The two salient
609 points are these: the pattern of the SSS spread follows quite closely that of
610 the SST, including the enhanced spread in the parent model over the slope;
611 and the spread of the SSS is reduced by an order of magnitude in deep water.
612 The former observation is easily explained since both tracers are controlled
613 by fairly similar mechanisms: the river plumes on the shelf and the eddies
614 over the abyssal plain. The latter is simply because, unlike the temperature,
615 the biggest source for changes in salinity is the fresh water discharge from
616 rivers, which is most pronounced near the coast.

617 *5.3.3. Experimenting with the wind perturbation*

618 The manner in which we perturbed the winds involved some choices based
619 more on convention than on physical law. To determine how much these
620 choices influenced the results we generated a few extra ensembles, described
621 in table 2. Note that for want of computer resources these experiments were
622 only performed for part of the period under examination, some not exceeding
623 the ensemble spin-up period. The spatial means of SST and SSH spread for
624 these ensembles are shown in figures 6 and 7 along with the four principal
625 ensembles.

626 It should be expected that increasing the wind perturbations should lead
627 to greater ensemble spread in response. The ensemble S05 bears out this
628 intuitive hypothesis. It is the exact same configuration as for our WIND en-
629 semble, with the same distribution of random numbers $\epsilon_{i,j}$ used to construct
630 the wind perturbations (equations (4) and (5)), except with a higher error
631 standard deviation of $s = 0.5$, instead of 0.3. The conclusion holds true for
632 both the SSH and SST and all subdivisions of the domain, although near the
633 end of the ensemble spin-up period the difference is sometimes very small.

634 In section 4.2, we discussed how our EOFs were calculated over a smaller
635 domain and a shorter time period than those of Vervatis et al. (2016), and
636 that we kept only the first two modes. The smaller domain could serve to
637 increase the structural detail for the same number of modes by virtue of

638 excluding contributions from outside the domain, but the shorter time frame
639 could have the opposite effect because it doesn't capture long-term changes;
640 similarly, restricting ourselves to only two modes precludes some of the small-
641 scale structure. The two sets of EOFs are, in short, not directly comparable.
642 This prompted two questions: firstly, what do we lose by only keeping the
643 first two modes, and secondly, what if we were to use exactly the same
644 wind perturbations as those which generated the parent ensemble? For the
645 second question, one intriguing possibility is that for any particular member,
646 the open boundary conditions would set up an ocean circulation pattern
647 in the child model which could be reinforced by matching winds—a kind of
648 resonance effect which would maintain a larger spread. Or, looked at another
649 way, if we used different wind perturbations than those used to generate the
650 parent ensemble, those perturbations might dampen the differences generated
651 by the ensemble of open boundary conditions.

652 To answer the first question we generated the EOF10 ensemble, which is
653 the same as our WIND ensemble except that we kept the first ten EOF modes
654 instead of only the first two. To be consistent, the random numbers associ-
655 ated with the first two modes are the same. Keeping ten modes instead of
656 two means that the total variance explained rises from $\sim 92.73\%$ to $\sim 98.97\%$.
657 The result is that, while the ensemble spread is enlarged in places, notably
658 for the SSH in deep water in January, most of the time it is about the same
659 as when only two modes are kept, and quite often it decreases. It seems
660 that by adding the extra complexity to the wind perturbations we are can-
661 celling out some of the variations we might otherwise produce. The higher
662 order modes might just be adding statistical noise which doesn't represent
663 the spatial pattern of the wind uncertainties, eventually leading to a decrease
664 in spread.

665 With the ensemble EOF||OBC we tested the second question by using
666 the same wind perturbations that Vervatis et al. (2016) applied to generate
667 the parent ensemble. We ensured for each member in the coastal ensemble
668 that we matched its wind perturbations to the corresponding parent mem-
669 ber which provided its open boundary conditions. Though the difference is
670 small, the spread for this configuration is actually lower than that produced
671 by our newly generated EOFs in the WIND \times OBC ensemble. We generated
672 another ensemble, EOFBISC, to measure the strength of these perturbations.
673 This ensemble uses the same wind perturbations as EOF||OBC but like the
674 WIND ensemble keeps the unperturbed open boundaries for all members. The
675 spread in SST is also almost always less than for the WIND ensemble, but

676 the spread in SSH increases from the end of the ensemble spin-up period.
677 This somewhat contradictory result demonstrates the difficulty in drawing
678 general conclusions from these experiments—clearly a different set of EOFs
679 will lead to a different set of wind perturbations which may be more or less
680 successful at perturbing the ocean circulation model at different times.

681 As a final test, we configured yet another ensemble using the wind per-
682 turbations from the parent ensemble. The EOF \perp OBC ensemble was config-
683 ured the same as EOF \parallel OBC except that the wind perturbations are no longer
684 matched with the open boundary conditions from the equivalent member
685 from the parent model. Instead, we forced each member’s wind perturbation
686 with the open boundary conditions from the member number two beneath it,
687 effectively randomising them. The spread is higher than that of EOF \parallel OBC for
688 the whole period simulated, though still lower than that of the WIND \times OBC
689 ensemble. If anything, we must conclude that any resonance effect, should it
690 exist, must be small, and that further perturbations (in this case generated
691 by simply dissociating the perturbations from the open boundary conditions)
692 can be expected to generate a greater spread rather than a lower one.

693 There are an infinite number of ways of generating perturbations, so we
694 needed to make some choices. We made them based on the simple physical
695 and intuitive assumption of generating wind uncertainties with a variability
696 comparable to that of the wind itself. Despite other interesting questions,
697 perhaps the most significant result for us from this series of experiments is
698 that small variations to this method do not significantly change the pattern
699 of model sensitivity, and hence do not change our conclusions.

700 5.3.4. *Experimenting with the nudging layer width*

701 A wider nudging layer ought to increase the influence of the parent model
702 on the child model, preserving more of its behaviour in the downscaled sys-
703 tem. But for an ensemble we expect the dynamics to be less impacted by the
704 local perturbations (the wind, here). To learn something about the way the
705 nudging layer width affects the ensemble spread we prepared three ensembles
706 with a 10-point nudging layer instead of 30; these are WIND-NL10, OBC-NL10
707 and WIND \times OBC-NL10. The narrower nudging layer was imposed after the en-
708 semble spinup period, so the ensemble statistics are directly comparable with
709 the other experiments. The time series of the ensemble spreads are shown in
710 figures 6 and 7. Changing the nudging layer width has several repercussions,
711 so we shall take the variables one at a time, starting with the SSH ensem-
712 ble spreads. Overall, except for the first few days over the shelf and slope

713 where the SSH is controlled mainly by non-local effects (the domain-wide,
714 synchronous oscillations previously mentioned), the narrower nudging layer
715 leads to greater spread for all three coastal ensembles. This is true even for
716 the OBC run, despite there being no wind perturbations, doubtless a result
717 of the greater freedom for the eddy dynamics. The WIND ensemble's spread
718 sees the greatest relative increase from this modification, especially over the
719 shelf and slope, where it even exceeds all the other ensembles for a few days
720 in mid-March. The OBC and WIND×OBC ensembles react in a more compli-
721 cated way, with the former's ensemble spread exceeding the latter's in deep
722 water for most of the 3-month period. The reason for this isn't clear, but one
723 possibility might be that the 5-day interval between random numbers (which
724 sets up the time correlation of the wind perturbations) is insufficient, leading
725 the wind perturbations to erode rather than strengthen the eddies. Over
726 the slope and shelf, where local effects are more important, the WIND×OBC
727 ensemble retains its supremacy.

728 For the SST things are again quite different. The WIND ensemble's spread
729 is almost unchanged over the shelf, sometimes slightly lower, but in deep
730 water it is higher everywhere with the narrower nudging layer. Here too,
731 the reduced influence of the parent model allows the wind perturbations
732 freer reign to spread the ensemble. But the OBC and WIND×OBC ensemble
733 spreads are often weaker with the narrower nudging layer. Again, why this
734 should be the case is not entirely clear, but one possibility is that while the
735 WIND ensemble is less constrained by the parent model's unperturbed run,
736 the other ensembles rely on the ensemble of open boundary conditions to
737 generate their ensemble spread. Why this should have the opposite effect
738 for SSH and SST may be explained by higher scalar diffusion in the coastal
739 model as discussed in sections 5.3.1 and 5.3.2. With already smaller spread
740 in SST because of the narrower nudging layer, the homogenisation of the
741 ensemble's SST fields by the eddies might be more readily achieved.

742 These nudging layer experiments reinforce our observations from sections
743 5.3.1 and 5.3.2, that increasing SSH spread often comes at the cost of reducing
744 the spread in SST. This is a trade-off which will have consequences for data
745 assimilation schemes.

746 *5.3.5. The spatial structure of the ensemble spread of surface variables*

747 The plots in figures 6 and 7 give us useful mean quantities but to really
748 understand how the perturbations affect the circulation it helps to look at
749 some surface maps. Shown in figure 9 are, for the four principal ensembles,

750 the time means of SSH ensemble spread for the three months following the
751 ensemble spin-up phase. This compacted view of the ensembles over the
752 period summarises their behaviour and major differences:

- 753 1. In all the ensembles the greatest contributor to the spread in SSH is
754 from the mesoscale eddies. This follows from observing that the region
755 where SSH ensemble spread is concentrated is also where the eddies are
756 present.
- 757 2. These eddies are concentrated in deep water, with the continental slope
758 (represented by the 200 m isobath) neatly circumscribing their range.
- 759 3. The WIND ensemble has the smallest spread, though in a concentrated
760 area in the middle of the domain its ensemble spread is greater than
761 that of the parent ensemble.
- 762 4. Over the shelf and at the domain's edge the WIND ensemble has the
763 lowest spread. For the other ensembles, non-local processes contribute
764 to the spread in these areas while the WIND ensemble is limited by only
765 having one set of open boundary conditions.
- 766 5. The OBC and WIND×OBC ensembles are very similar—non-local errors
767 dominate.
- 768 6. The ensemble spread of SSH is concentrated in the centre of the domain
769 of the child ensembles, a feature not shared by the parent ensemble.
770 One reason for this is the higher spatial coherence of the mesoscale
771 eddies than in the parent model. Another possibility is that the eddies
772 might be trapped in the coastal model's domain.

773 What this view doesn't show is the pronounced ensemble spread, seen as
774 peaks in the time series in figure 6, which results from the domain-wide pul-
775 sations generated in the parent model. It also doesn't show the comparable
776 spread of the WIND ensemble to the others early in the period, before it
777 falls off from mid-January. This is clearly visible in the downward trend of
778 the WIND ensemble's deep-water SSH spread (the top-right graph in figure
779 6). It's an almost suspiciously smooth signal, but it seems to be a product
780 of a particularity of the unperturbed run of the parent model (recall that
781 all members of the WIND ensemble are forced at the open boundaries by
782 the parent's unperturbed run). The large mesoscale eddy in that simulation
783 shifts towards the northern and western boundaries of the BOBLAND domain
784 during the simulation period. Being closer to the boundary means that it is
785 more likely to be controlled by the prescribed open boundary conditions and
786 is more resilient to being displaced by the wind.

787 This behaviour is echoed in the ensemble means and spreads of the surface
788 velocity fields in figures 10 and 11, respectively for the 20th of January and
789 the 19th of February (representing the beginning and end of the downward
790 trend in SSH spread). Comparing the two we see that the WIND ensemble’s
791 spread on the 20th of January is greater than a month later, and this reduc-
792 tion in ensemble spread manifests as a more spatially coherent eddy signal
793 in the ensemble mean for the WIND ensemble in the February plot.

794 To lend support to this hypothesis we can compare with runs where we
795 reduced the nudging layer from 30 points to 10 points (figure 6, black and
796 green dashed curves respectively). Sure enough, the effect of the wind pertur-
797 bations on the SSH in this ensemble is less constrained by the open boundary
798 conditions and the overall ensemble spread is higher. It still falls in February,
799 but by proportionally less, and the recovery in March is stronger. Figure 12
800 shows surface maps for individual ensemble members for the 9th of March
801 for the WIND and WIND-NL10 ensembles, where the latter is identical to the
802 former except with the narrower 10-point nudging layer. By gathering all
803 the member plots into one figure we are able to rapidly make a qualitative
804 comparison (ignoring fine details) of the two ensembles. With the 30-point
805 nudging layer, the large eddy barely changes position from member to mem-
806 ber, whereas with a 10-point nudging layer the members are more distinct,
807 with the eddy’s position and strength varying much more. The 10-point
808 nudging layer also allows the eddy’s centre to get closer to the edge of the
809 domain.

810 Let us return briefly to the velocity ensemble means and spreads in figures
811 10 and 11. Conspicuously absent from the spreads for the 20th of January
812 (figure 10) is a signal corresponding to the slope current clearly visible in
813 the zonal means. As this is true for the parent ensemble as well as for the
814 coastal ensembles, this is an aspect of the parent model’s dynamics that
815 the coastal ensembles faithfully reproduce. Although we were able to find
816 a weak spread at certain times, most of the time the slope current—when
817 present—remained unmoved by the perturbations.

818 As with the SSH, we also produced plots for the time mean of surface
819 temperature spread; these are shown in figure 13. The combination of differ-
820 ent perturbation sources affects the SST’s ensemble spread in similar ways
821 to the SSH, though with some particularities. The salient points are:

- 822 1. Again the domain is divided into two zones by the continental slope,
823 but here the ensemble spread is greatest over the shelf, with the spread

- 824 produced by the river plumes dominating the signal.
- 825 2. In deep water, the WIND ensemble again has the lowest spread, and
826 at the edges, where the spread is controlled by the open boundary
827 conditions, the spread drops off; the parent ensemble has the highest
828 spread.
- 829 3. Over the shelf, however, the WIND ensemble's spread is comparable to
830 the other ensembles, with the river plume being the dominant source
831 of error there.
- 832 4. As a consequence of this, the WIND×OBC ensemble's spread over the
833 shelf is visibly greater than the OBC ensemble's, a finding reflected in
834 the time series of spatial means of ensemble spread in figure 7.
- 835 5. Meanwhile, it is the parent which has the largest spread over the shelf
836 as it does overall, though with a different pattern which extends further
837 south and west from the Adour's mouth and is weaker further north
838 along the shelf.
- 839 6. The ensemble spread in the parent has localised peaks in contrast to
840 the smoother pattern in the child model's ensembles.

841 5.4. *Beneath the surface*

842 In the preceding sections we've concentrated on the ensemble spread of
843 surface variables, a natural place to look considering that the wind pertur-
844 bations affect the surface directly and the ocean interior only indirectly. But
845 there is a significant response to these perturbations deep beneath the sur-
846 face too, in some places much stronger. Attributing an error response of the
847 oceanic interior to surface forcing only is interesting enough from the stand-
848 point of model and dynamic variability, but for data assimilation systems it
849 suggests the possibility of correcting the model well away from the surface
850 when surface data is assimilated. So with that in mind, in this section we
851 explore the ensemble response in the coastal ocean interior.

852 5.4.1. *A general view of the ensembles beneath the surface*

853 We begin, once again, by concentrating on the four principal ensembles:
854 WIND, OBC, WIND×OBC, and the parent ensemble. In all the ensembles the
855 maxima of ensemble spread for salinity and temperature are either below the
856 surface or over the continental shelf, under the influence of the river plume's
857 wandering tendrils of cool, fresh water. Looking at time series of mean
858 temperature spread near the base of the mixed layer (taken to be roughly

185 m below the surface), shown in figure 14, we note immediately that it is much greater than at the surface. We see that the spread of the parent is sometimes slightly higher than the other ensembles, but often considerably lower than all but the WIND ensemble. The slightly greater spread of parent relative to the coastal ensembles in the latter half of the period is mainly due to the deepening of the mixed layer, which puts the highly variable base beyond the depth at which we measured the statistics. A look at the vertical cross-sections in figure 15 shows a very high spread at the bottom of the mixed layer, which comes about as the depth of the thermocline changes in response to changes in mixing depth induced by the modified wind stress. This effect was also noted by Andreu-Burillo et al. (2002) following their sensitivity experiments. There are a couple of other conclusions we can draw immediately from figure 15: one is that from the surface to near the bottom of the mixed layer, but before reaching the thermocline, the parent ensemble's spread reduces more than the coastal ensembles', which remain fairly stable. This is consistent with our earlier conclusion of higher diffusion in the coastal model, but now it is manifested in the vertical direction. The second is that at the thermocline the spread of the coastal ensembles is greater than the parent's. Recalling that the coastal ensembles' higher spreads in SSH are reflected in their more coherent and stronger eddies, these would be expected to erode the surface stratification and deepen the mixed layer rather more than in the parent. This explains at least part of both observations, though it must be said that differences in the way that vertical mixing is parameterised would also be pertinent this close to the surface.

Once again, the equivalent plots for salinity do not add enough new information to warrant their inclusion here. Suffice it to say that deep below the surface the effect of the river plumes is very small.

At the increased depths of 500 m and 1000 m, shown in figure 16, the parent ensemble spread is still lower again than the coastal ensembles, though it still overtakes the WIND ensemble later in the time series. This behaviour too is visible in figure 15. Below the surface mixed layer, the coastal ensembles feature a strong spread which coincides with a second, deeper thermocline. At the same depth the parent ensemble's spread almost disappears. The vertical gradients of temperature and salinity are very similar in both parent and coastal models, so we may once more attribute this difference to the more detailed small-scale dynamics of the coastal model.

Arguably it is here, below the mixed layer, that the coastal model distinguishes itself the most from the regional parent model. The enhanced hori-

897 zontal resolution allows it to resolve small-scale eddies (as small as ~ 10 km
898 in diameter), and the sigma-level vertical discretisation allows for a more
899 detailed description of the circulation's interaction with bottom relief. As
900 a result, not only are the eddies stronger and penetrate deeper, but from
901 about a thousand metres below the surface the coastal model's horizontal
902 velocity fields become wonderfully complex, almost belying any connection
903 to the parent model's dynamics. Figure 17 illustrates this for the meridional
904 velocity component for unperturbed runs of the parent and coastal models
905 at the surface, and three horizontal sections at different depths. The velocity
906 fields are stronger in the coastal model at all depths, but that's not sufficient
907 to explain the fall at depth in the parent's ensemble spread relative to the
908 coastal model. In deep water, the surface velocity fields are dominated by
909 the large, anticyclonic mesoscale eddy, with a smaller cyclonic one adjoining
910 it on its south-eastern edge (clear in the coastal model, and hinted at in the
911 parent model). At 500 m these eddies still dominate the velocity field, but
912 at 1000 m they fracture into smaller, less coherent structures. This change
913 in regime isn't unique to the coastal model, but what differentiates it from
914 the parent model is the high density of variation. While the large structures
915 dominate at the surface for both models, in deeper water it is the smaller
916 structures, which only the coastal model resolves, which dominate. Some of
917 this complexity arises from small eddies generated by the flow around the
918 underwater gorges on the continental slopes, which shows the importance not
919 only of the model's higher resolution but of the bathymetry's too, as well as
920 the terrain-following vertical sigma coordinates.

921 *5.4.2. The reduced nudging layer's effects at depth*

922 As we did for section 5.4.1, we plotted curves for the coastal ensembles
923 with a reduced nudging layer at depth (figures 14 and 16 green curves).
924 The results below the surface are less ambiguous than for the surface, if
925 one remembers to take into account that in the latter part of the period the
926 mixed layer deepens, and the ensemble statistics start to resemble the surface
927 more than the thermocline. Apart from this, the reduced nudging layer leads
928 to a greater ensemble spread, consistent with the coastal model's stronger
929 response to perturbations at depth. Freed somewhat from the constraints of
930 the parent model, it develops even more differences between members as we
931 descend the water column.

932 *5.4.3. Interaction of eddies with bottom near a strong gradient*

933 We noticed a particularly strong spread at a specific time and place which
934 we thought worthy of closer examination. In figure 18 we show cross-sections
935 of the WIND ensemble centred on an eddy. The unperturbed run's SSH is
936 depicted in figure 18a, in which an eddy dipole can be made out.

937 Figure 18b shows a cross-section of ensemble spread of temperature, and
938 several features stand out. Firstly, there is a large spread at the base of the
939 surface mixed layer, as discussed in section 5.4.1. There is also a small re-
940 gion of larger spread at the coast, a result of the wind pushing the cool, fresh
941 water plume of the Adour in different directions, and it being mixed through
942 the water column. But most interestingly, there is a region between 1000 and
943 1500 m depth and -3 and -2.5° longitude with high ensemble spread. It
944 is particularly strong along the gently sloping Landes Plateau. This feature
945 is all the more remarkable because of the region of reduced spread between
946 it and the surface mixed layer, and at these depths the vertical mixing pro-
947 duced by surface stresses cannot have much influence—neither, in fact, can
948 the open boundary conditions, since only local wind perturbations have been
949 applied. The unperturbed temperature profile in figure 18c shows the sec-
950 ond, deep thermocline mentioned in section 5.4.1. It lies beneath a region
951 of much gentler temperature gradient which extends from the near-surface
952 thermocline to a depth of around 1200 m. The zonal velocity field at this
953 depth, plotted in figure 18e, shows that the eddy visible in the SSH signal in
954 figure 18a extends down to just about reach the slope.

955 The final plot in the figure, 18f, shows the spread of the zonal velocity.
956 While it decreases with depth, it is still significant near the second ther-
957 mocline. This and the anomalously high values of vertical velocity spread
958 shown in figure 18d suggest that interactions between the bottom and the
959 eddy are responsible for the enhanced spread in temperature. The verti-
960 cal velocity spread is especially high at the bottom where it coincides with
961 the high spread in temperature. There is meanwhile a very small ensem-
962 ble spread in vertical velocity coinciding with the near-surface thermocline.
963 This is telling: the thermocline is a region of strong temperature gradient,
964 and therefore high stability, where we would expect vertical movement to be
965 suppressed. We can hence differentiate the mechanism which displaces the
966 near-surface thermocline from that which produces the strong temperature
967 spread near the bottom, where strong vertical movements are present within
968 the stable layer. Oey and Zhang (2004) showed that warm eddies interacting

969 with a slope could lead to big modifications to the deep circulation, includ-
970 ing cyclones and jets. While in this particular case we do not see substantial
971 alterations to the horizontal velocities, the vertical velocities do appear to be
972 affected by such an interaction.

973 At other times not shown, vertical motion (and consequently its spread) is
974 suppressed in the deep thermocline as well even when there is a large spread of
975 temperature. And there are doubtless other complex processes leading to en-
976 hanced spread that a more exhaustive study of the ensembles could identify.
977 We can say with certainty that the presence of the eddy–slope interaction is
978 not necessary to generate a large spread deep beneath the surface. Over the
979 deep abyssal plain of the Bay of Biscay, Vervatis et al. (2016) demonstrated
980 this by assimilating SSH data, which were able correct the strong errors be-
981 tween 1000 and 1500 m. Ayoub et al. (in prep.) also observed a high spread
982 at these depths far from the slope. But the eddy–slope interaction provides
983 another mechanism by which model errors can be generated at depth and,
984 hence, by which the model can be corrected at depth by data assimilation at
985 the surface. The use of the coastal model is decisive here, because although
986 we could not verify whether or not the vertical velocities were affected in the
987 parent model, there is no such corresponding enhanced spread in tempera-
988 ture. It is possible that the eddy–slope interaction is present, but the result
989 is barely, if at all, discernible. Interestingly, it is the WIND ensemble which
990 shows the strongest spread here, and by some way. Even the WIND×OBC en-
991 semble doesn’t come close, resembling more closely the OBC ensemble. This
992 isn’t then a question of pure strength of perturbation, but of a confluence
993 of factors: the right perturbation leading to the right deviations of the eddy
994 with the right temperature gradient in the right place so that the condi-
995 tions are favourable for the eddy–slope interaction to generate the spread in
996 temperature.

997 *5.5. Extent and diffusion of river plumes*

998 Only the Adour river runoff is explicitly modelled in the coastal model,
999 but the Gironde’s runoff is also modelled in the parent model. Its plume
1000 enters the coastal model from the northern boundary, forming as it does,
1001 with the Adour’s plume, a significant body of cool, less saline water. As we
1002 noted in section 5.3.5, the behaviour of the Adour plume is quite different in
1003 the two models, leading to different ensemble statistics.

1004 Direct comparison between the models is complicated, because even
1005 though the prescribed outflow volumes of the Adour are the same for most of

1006 the time period in the BISCAY36 and BOBLAND configurations, the way that
1007 river water is injected into the Bay of Biscay is different. Both use a single
1008 grid point as the source but this grid point is much smaller in BOBLAND than
1009 in BISCAY36. The depth at the river mouth is also different, with BISCAY36
1010 drawing its depth from its coarse bathymetry, about 10.5 m, and BOBLAND
1011 using a prescribed depth of 5 m. The areas of the outward faces of the grid
1012 cells are then about 22950 m² and 2500 m² respectively, so the discharge ve-
1013 locity is more than 9 times greater in BOBLAND than in BISCAY36. A higher
1014 velocity could produce a longer plume, but could also generate more local
1015 turbulence and generate more mixing, dissipating the plume earlier. The
1016 salinity is treated in a similar way in the two river models: it is prescribed
1017 and insignificant compared to the salinity levels in the open ocean (0 PSU
1018 in BOBLAND and 0.1 PSU in BISCAY36). Temperature, on the other hand, is
1019 handled differently, with only minimum and maximum values prescribed in
1020 BOBLAND (10° and 22°) with the seasonal temperature cycle set with a sinu-
1021 soid. In BISCAY36 the temperature is determined from a Neumann boundary
1022 condition, with the river discharge grid point treated as an open boundary.
1023 With these differences, only a gross evaluation of the river plume behaviour
1024 between the two configurations would be appropriate, but between the three
1025 different BOBLAND ensembles we can be more concrete.

1026 Our analysis is based on using the surface salinity as a proxy for the
1027 river plume. For each member of an ensemble, we calculated the area of
1028 each plume, whose extents were defined by one of three thresholds. Taking
1029 the ensemble mean, we may determine a typical plume area for each model.
1030 These areas are shown plotted over time in figure 19.

1031 The parent model clearly has the most expansive plume by all three of
1032 the threshold definitions, with only occasional exceptions, mainly at the be-
1033 ginning of the time series. Interestingly, the plumes from the three coastal
1034 ensembles are very similar; to the extent that they do differ from one another,
1035 they do not follow the pattern of ensemble spreads of surface temperature
1036 and salinity. As described in section 5.3 and figure 7 (recall that the surface
1037 salinity and temperature behave almost identically), these spreads are by far
1038 the greatest in the parent ensemble, followed by the WIND×OBC and OBC
1039 ensembles (these two are very close) and then the WIND ensemble, with by
1040 far the least spread. However although the parent ensemble's plume area
1041 is on average greater, the next greatest is, more often than not, the OBC
1042 ensemble, with the WIND×OBC ensemble usually found between it and the
1043 WIND ensemble. The greater spread of the parent ensemble seems to reflect

1044 its broader average plume, but between the coastal ensembles things are a
1045 bit subtler. A greater spread in SST or SSS in one coastal ensemble than
1046 in another could mean that since the plume is displaced further, it is also
1047 diffused more effectively. This threshold value is hence attained closer to the
1048 source. This is just one factor in a complicated process.

1049 To shed more light on the plume’s behaviour, we examine the salinity
1050 and its spatial gradients for the parent and child models, shown in figure 20.
1051 The example shown is for the first perturbed members of the WIND×OBC
1052 and parent ensembles. It appears as though we can distinguish three salinity
1053 regimes in both models: (1) the inner, fresh zone near the river mouth, with a
1054 strong gradient as the fresh water mixes quickly with the surrounding saline
1055 water; (2) an intermediate zone of mixed water distinct from the main, outer
1056 circulation; and (3) the saline, outer circulation. The inner zone of the coastal
1057 model has a very strong gradient at the river mouth, representing a rapid
1058 mixing of the plume here, while the parent’s mixing is slower. By the time
1059 we get to the edge of this zone, however, the parent model’s plume stabilises
1060 and remains coherent for longer, and the coastal model’s intermediate zone
1061 is patchy and diffuse. Even though a direct comparison is impossible, the
1062 ocean colour observations of Petus et al. (2014) do favour the coastal model’s
1063 representation. The very sharp boundary in the observations between the
1064 river plume and the surrounding water is situated very close to the river
1065 mouth, closer in character to the more localised and more quickly diffused
1066 plume of the coastal model.

1067 **6. Discussion and conclusions**

1068 We sought to downscale a region model ensemble by forcing a coastal
1069 model ensemble by an ensemble of open boundary conditions, as well as
1070 by local wind perturbations. We found that, overall, the ensemble of open
1071 boundary conditions coupled with local wind perturbations generated the
1072 greatest spread in the coastal model, but that the contribution from the
1073 open boundary ensemble was greater than the local wind perturbations, and
1074 the combination of the two led to only a slightly greater spread. It seems as
1075 though the non-local errors advected into the model were more important in
1076 generating local errors than the local wind perturbations. It was only during
1077 the early stages of the ensemble spinup, before the non-local perturbations
1078 had time to penetrate into the domain, did we see the local wind perturba-
1079 tions generate a greater spread over the shelf. Another contributing factor

1080 could be that although the OBC ensemble was generated without local wind
1081 perturbations, the perturbed surface characteristics resulting from the non-
1082 locally produced errors creates a difference in each member’s wind stress.
1083 This might partly explain why an additional wind perturbation doesn’t add
1084 too much to the coastal ensemble’s spread.

1085 There was a distinct difference between the signals over the shelf and over
1086 deep water for all the ensembles. The SSH spread was usually greater over
1087 deep water, where the forcing perturbations displaced the large mesoscale
1088 eddy; the SST spread was greater over the shelf, where the river plume dom-
1089 inated the signature. Conversely, the SST spread was lower over deep water
1090 and the SSH spread was for most of the period lower over the shelf, the ex-
1091 ception being the large, non-locally produced synchronous events which had
1092 an outsized effect over shallow waters. The behaviour is different too, with
1093 the spread of both variables over the shelf varying spatially and temporally
1094 more so than over deep water.

1095 We experimented with different methods of generating the wind pertur-
1096 bations and combining them with the ensemble of open boundary conditions.
1097 Interestingly, the exact form of the wind perturbations did not have a huge
1098 impact on the shape of the time series—the magnitude changed significantly,
1099 but periods of large spread and low spread were mostly conserved. This sug-
1100 gests that the intrinsic variability of the model dominates over any specific
1101 perturbation scheme. This may also explain the saturation in spread that ap-
1102 peared to occur near the end of the ensemble spinup period with the different
1103 wind perturbations.

1104 Of course this is not quite as true when we switch from an ensemble with
1105 only local wind perturbations to one forced by the parent ensemble at the
1106 open boundaries. In this case, the non-local errors can be related to much
1107 larger scale processes that the coastal model alone cannot represent, whatever
1108 the local perturbation scheme. The dominance of intrinsic variability and the
1109 insensitivity to the precise form of errors in boundary and initial conditions
1110 has been noted in geophysical modelling before, for example by Giorgi and
1111 Bi (2000) and Rinke et al. (2004). We must be careful not to draw too
1112 many similarities with our own experiments—for one thing, not all of their
1113 conclusions (which include insensitivity to perturbation magnitude) apply to
1114 our ensembles. But it seems safe to conclude that the finer details in a wind
1115 perturbation scheme are of relative unimportance.

1116 The choice of nudging layer had consequences for the ensemble spread,
1117 critically so for the WIND ensemble. With a broad nudging layer, the large

1118 mesoscale eddy was present in the same area near the edge of the domain for
1119 all members of the WIND ensemble. This led to a sustained drop in ensemble
1120 spread, an effect which disappeared almost entirely with application of the
1121 narrower nudging layer. The narrower nudging layer also led to a bigger
1122 spread in SSH for the other coastal ensembles, even though their spreads in
1123 SST saw only modest and inconsistent increases.

1124 One of the more intriguing results in the present work is how deep subsur-
1125 face errors can be generated by surface perturbations. A large spread at the
1126 base of the seasonal thermocline was anticipated, having been previously de-
1127 scribed (Andreu-Burillo et al., 2002). But for there to be such a strong signal
1128 at much greater depths, beyond 1000 m, was less expected. To be sure, the
1129 parent model had a measurable spread there too, but while the signals at the
1130 near-surface thermocline were similar for the parent and coastal ensembles,
1131 the deep water spread was in places at least twice as strong in the coastal
1132 ensembles than in the parent ensemble. This suggests a valuable contribu-
1133 tion that the downscaling paradigm can make to the study of model errors.
1134 It is especially appealing from a data assimilation perspective, for if we can
1135 generate deep, subsurface errors then we might hope to be able to correct
1136 these regions with only surface data. The vertical temperature gradients for
1137 the two models' unperturbed runs, we noted, were almost the same, but the
1138 model and bathymetry resolutions were much higher in the coastal model.
1139 A major consequence of this is the more detailed dynamics, of course, but
1140 more specifically it is the ability of the coastal model to resolve more coher-
1141 ent mesoscale eddies with stronger currents. These can reach deeper in the
1142 coastal model than in the parent model, whose less coherent eddy has a much
1143 weaker effect at these depths. We also saw that eddy-slope interactions could
1144 be an additional mechanism for ensemble divergence, something we did not
1145 see in the parent ensemble. The use of sigma layer depth coordinates in the
1146 coastal model may have contributed to this behaviour.

1147 The river plume's behaviour was very different between the parent and
1148 child models, and this contributed to the parent model's mostly higher SST
1149 spread. Among the coastal ensembles the spread was usually greatest for the
1150 OBC ensemble, though the difference was small. According to our threshold
1151 criterion, the mean plume area does not appear to be directly correlated with
1152 the ensemble spread. We did find that the plume's behaviour supported our
1153 hypothesis that the parent model diffuses scalars less than the higher reso-
1154 lution coastal model, which accounted for some of the differences in spread
1155 between the parent ensemble and the child ensembles.

1156 A recurring feature of the ensembles is that the WIND×OBC ensemble
1157 did not always have a greater spread than the OBC ensemble. Logically,
1158 more sources of perturbation should lead to more spread, as we can explore
1159 more regions of the error space. But with a limited number of ensemble
1160 members we cannot always guarantee an increase in spread as we pile on
1161 perturbations. This limitation could probably be somewhat alleviated by
1162 increasing the magnitude of the perturbations, but our aim is to not simply
1163 increase the spread at any cost—we are interested in the model’s response to
1164 forcing errors, and we would prefer the error space to be densely populated
1165 by our ensemble. Despite the sometimes lower spread, the benefit of adding
1166 an extra forcing perturbation is that it opens the error spaces which may
1167 be left untouched, or weakly affected, by any one forcing perturbation. We
1168 can see this in our ensembles. Figure 13 shows almost no spread at the edge
1169 of the continental shelf for the coastal ensembles, but the spread, while still
1170 weak, is slightly enhanced when wind perturbations are added.

1171 The next phase of this work will be to use what we have learned in a
1172 data assimilation context, in which we interpret the ensemble spread as the
1173 model’s uncertainty. Of course a successful data assimilation scheme requires
1174 that the ensemble’s spread span the data—Vervatis et al. (in review, 2019)
1175 propose empirical techniques to determine this. We hope to present these
1176 results in a future paper.

1177 **Acknowledgements**

1178 We would like to thank Cyril Nguyen and Thomas Duhaut of the Lab-
1179 oratoire d’Aérodologie in Toulouse for their help with running and adapting
1180 the Symphonie model, and Florent Lyard of LEGOS for helping us with
1181 the tidal analysis and for supplying the high-resolution bathymetry. We
1182 would also like to thank Patrick Marchesiello for useful discussions on model
1183 dissipation. This project was granted access to the HPC resources of the
1184 CALMIP supercomputing centre through grants P1119 and P1325 for the
1185 2016–2017 allocations, and thanks go to Emmanuel Courcel, Nicolas Renon
1186 and Pierrette Barbaresco for their support. Support from Didier Gazen and
1187 Laurent Cabanas of the Laboratoire d’Aérodologie, Toulouse, for the use of the
1188 Nuwa cluster is also gratefully acknowledged. This research was undertaken
1189 as part the AMICO project and was supported by the GMES–MDD joint
1190 program of the French Ministère de la Transition écologique et solidaire and
1191 CNRS/INSU.

1192 **References**

- 1193 van Aken, H.M., 2000. The hydrography of the mid-latitude Northeast At-
1194 lantic Ocean: II: The intermediate water masses. *Deep Sea Research*
1195 Part I: Oceanographic Research Papers 47, 789–824. doi:10.1016/S0967-
1196 0637(99)00112-0.
- 1197 Andreu-Burillo, I., Caniaux, G., Gavart, M., De Mey, P., Baraille, R., 2002.
1198 Assessing ocean-model sensitivity to wind forcing uncertainties. *Geophys-*
1199 *ical Research Letters* 29, 5.1–5.4. doi:10.1029/2001GL014473.
- 1200 Auclair, F., Marsaleix, P., De Mey, P., 2003. Space-time structure and
1201 dynamics of the forecast error in a coastal circulation model of the
1202 Gulf of Lions. *Dynamics of Atmospheres and Oceans* 36, 309–346.
1203 doi:doi:10.1016/S0377-0265(02)00068-4.
- 1204 Ayoub, N., De Mey, P., Marsaleix, P., in prep. Constraining a coastal ocean
1205 model by SST and SSH observations using an Ensemble Kalman Filter in
1206 the Bay of Biscay (North-East Atlantic) .
- 1207 Echevin, V., De Mey, P., Evensen, G., 2000. Horizontal and vertical struc-
1208 ture of the representer functions for sea surface measurements in a coastal
1209 circulation model. *Journal of Physical Oceanography* 30, 2627–2635.
- 1210 Estournel, C., Auclair, F., Lux, M., Nguyen, C., Marsaleix, P., 2009. “Scale
1211 oriented” embedded modeling of the North-Western Mediterranean in the
1212 frame of MFSTEP. *Ocean Science* , 73–90.
- 1213 Estournel, C., Zervakis, V., Marsaleix, P., Papadopoulos, A., Auclair, F.,
1214 Perivoliotis, L., Tragou, E., 2005. Dense water formation and cascading in
1215 the Gulf of Thermaikos (North Aegean) from observations and modelling.
1216 *Continental Shelf Research* , 2366–2386doi:doi:10.1016/j.csr.2005.08.014.
- 1217 Giorgi, F., Bi, X., 2000. A study of internal variability of a regional climate
1218 model. *Journal of Geophysical Research: Atmospheres* 105, 29503–29521.
- 1219 Griffies, S.M., Hallberg, R.W., 2000. Biharmonic friction with a
1220 Smagorinsky-like viscosity for use in large-scale eddy-permitting ocean
1221 models. *Monthly Weather Review* 128, 2935–2946.

- 1222 Herbert, G., 2012. Modélisation et observation de la dynamique haute
1223 fréquence de la circulation du Golfe de Gascogne. Ph.D. thesis. Université
1224 de Toulouse.
- 1225 Herbert, G., Ayoub, N., Marsaleix, P., Lyard, F., 2011. Signature of the
1226 coastal circulation variability in altimetric data in the southern Bay of
1227 Biscay during winter and fall 2004. *Journal of Marine Systems* 88, 139–
1228 158.
- 1229 Jordà, G., De Mey, P., 2010. Characterization of error dynamics in a 3D
1230 coastal model of the Catalan sea using stochastic modelling 30, 419–441.
- 1231 Katavouta, A., Thompson, K.R., 2016. Downscaling ocean conditions with
1232 application to the Gulf of Maine, Scotian Shelf and adjacent deep ocean.
1233 *Ocean Modelling* 104, 54–72. doi:10.1016/j.ocemod.2016.05.007.
- 1234 Kersalé, M., Marié, L., Cann, B.L., Serpette, A., Lathuilière, C., Le Boyer,
1235 A., Rubio, A., Lazure, P., 2016. Poleward along-shore current pulses on
1236 the inner shelf of the Bay of Biscay. *Estuarine, Coastal and Shelf Science*
1237 179, 155–171. doi:https://doi.org/10.1016/j.ecss.2015.11.018. special Issue:
1238 Functioning and dysfunctioning of Marine and Brackish Ecosystems.
- 1239 Kim, S., Samelson, R.M., Snyder, C., 2011. Towards an uncertainty budget
1240 for a coastal ocean model. *Monthly Weather Review* 139, 866–884.
- 1241 Kourafalou, V.H., De Mey, P., Staneva, J., Ayoub, N., Barth, A., Chao, Y.,
1242 Cirano, M., Fiechter, J., Herzfeld, M., Kurapov, A., Moore, A.M., Oddo,
1243 P., Pullen, J., van der Westhuysen, A., Weisberg, R.H., 2015. Coastal
1244 ocean forecasting: science foundation and user benefits. *Journal of Oper-
1245 ational Oceanography* 8, s147–s167. doi:10.1080/1755876X.2015.1022348,
1246 arXiv:https://doi.org/10.1080/1755876X.2015.1022348.
- 1247 Large, W.G., Yeager, S.G., 2004. Diurnal to Decadal Global Forcing for
1248 Ocean and Sea-Ice Models: The Data Sets and Flux Climatologies. Tech-
1249 nical Report. NCAR Technical Note NCAR/TN-460+STR.
- 1250 Le Cann, B., 1990. Barotropic tidal dynamics of the Bay of Biscay shelf:
1251 observations, numerical modelling and physical interpretation. *Continental
1252 Shelf Research* 10, 723–758.

- 1253 Le Cann, B., Serpette, A., 2009. Intense warm and saline upper ocean inflow
1254 in the southern Bay of Biscay in autumn–winter 2006–2007. *Continental*
1255 *Shelf Research* 29, 1014–1025.
- 1256 Le Hénaff, M., De Mey, P., Marsaleix, P., 2009. Assessment of observa-
1257 tional networks with the Representer Matrix Spectra method—application
1258 to a 3D coastal model of the Bay of Biscay. *Ocean Dynamics* , 3–
1259 20doi:10.1007/s10236-008-0144-7.
- 1260 Le Sommer, J., Penduff, T., Theetten, S., Madec, G., Barnier, B., 2009. How
1261 momentum advection schemes influence current-topography interactions at
1262 eddy permitting resolution. *Ocean Modelling* 29, 1–14.
- 1263 Lucas, M.A., Ayoub, N., Barnier, B., Penduff, T., De Mey, P., 2008. Stochas-
1264 tic study of the temperature response of the upper ocean to uncertainties
1265 in the atmospheric forcing in an Atlantic OGCM. *Ocean Modelling* 20,
1266 90–113. doi:10.1016/j.ocemod.2007.07.006.
- 1267 Lyard, F., 2015. private communication.
- 1268 Madec, G., the NEMO team, 2008. NEMO ocean engine. Note du Pôle de
1269 modélisation 27. Institut Pierre-Simon Laplace (IPSL), France. ISSN No
1270 1288-1619.
- 1271 Maraldi, C., Chanut, J., Levier, B., Ayoub, N., De Mey, P., Refray,
1272 G., Lyard, F., Cailleau, S., Drevillon, M., Fanjul, E.A., Sotillo, M.G.,
1273 Marsaleix, P., the Mercator Research, Team, D., 2013. NEMO on the
1274 shelf: assessment of the Iberia–Biscay–Ireland configuration. *Ocean Sci-*
1275 *ence* , 745–771doi:10.5194/os-9-745-2013.
- 1276 Marchesiello, P., McWilliams, J.C., Shchepetkin, A., 2001. Open boundary
1277 conditions for long-term integration of regional oceanic models. *Ocean*
1278 *Modelling* 3, 1–20.
- 1279 Marsaleix, P., Auclair, F., Estournel, C., 2006. Considerations
1280 on open boundary conditions for regional and coastal ocean mod-
1281 els. *Journal of Atmospheric and Oceanic Technology* , 1604–
1282 1613doi:http://dx.doi.org/10.1175/JTECH1930.1.

- 1283 Marsaleix, P., Auclair, F., Estournel, C., 2009a. Low-order pressure gradient
1284 schemes in sigma coordinate models: The seamount test revisited. *Ocean*
1285 *Modelling* , 169–177doi:http://dx.doi.org/10.1016/j.ocemod.2009.06.011.
- 1286 Marsaleix, P., Auclair, F., Estournel, C., Nguyen, C., Ulses, C.,
1287 2011. An accurate implementation of the compressibility terms
1288 in the equation of state in a low order pressure gradient scheme
1289 for sigma coordinate ocean models. *Ocean Modelling* , 1–
1290 13doi:http://dx.doi.org/10.1016/j.ocemod.2011.07.004.
- 1291 Marsaleix, P., Auclair, F., Estournel, C., Nguyen, C., Ulses, C., 2012.
1292 Alternatives to the Robert-Asselin filter. *Ocean Modelling* , 53–
1293 66doi:http://dx.doi.org/10.1016/j.ocemod.2011.11.002.
- 1294 Marsaleix, P., Auclair, F., Floor, J.W., Herrmann, M.J., Estournel,
1295 C., Pairaud, I., Ulses, C., 2008. Energy conservation issues in
1296 sigma-coordinate free-surface ocean models. *Ocean Modelling* , 61–
1297 89doi:http://dx.doi.org/10.1016/j.ocemod.2007.07.005.
- 1298 Marsaleix, P., Ulses, C., Pairud, I., Herrmann, M.J., Floor, J.W., Estour-
1299 nel, C., F., A., 2009b. Open boundary conditions for internal grav-
1300 ity wave modelling using polarization relations. *Ocean Modelling* , 27–
1301 42doi:http://dx.doi.org/10.1016/j.ocemod.2009.02.010.
- 1302 Oey, L.Y., Zhang, H.C., 2004. The generation of subsurface cyclones and jets
1303 through eddy–slope interaction. *Continental Shelf Research* 24, 2109–2131.
1304 doi:10.1016/j.csr.2004.07.007. recent *Developments in Physical Oceanog-*
1305 *raphic Modelling: Part I*.
- 1306 Orain, F., 2016. Product user manual for Level 3 SST products over European
1307 Seas. Technical Report.
- 1308 Pairaud, I.L., Auclair, F., Marsaleix, P., Lyard, F., Pichon, A., 2010. Dy-
1309 namics of the semi-diurnal and quarter-diurnal internal tides in the Bay
1310 of Biscay. part 2: Baroclinic tides. *Continental Shelf Research* , 253–
1311 269doi:10.1016/j.csr.2009.10.008.
- 1312 Pairaud, I.L., Lyard, F., Auclair, F., Letellier, T., Marsaleix, P., 2008. Dy-
1313 namics of the semi- diurnal and quarter-diurnal internal tides in the bay
1314 of biscay. part 1: Barotropic tides. *Continental Shelf Research* , 1294–
1315 1315doi:doi:10.1016/j.csr.2008.03.004.

- 1316 Petus, C., Marieu, V., Novoa, S., Chust, G., Bruneau, N., Froidefond, J.M.,
1317 2014. Monitoring spatio-temporal variability of the adour river turbid
1318 plume (bay of biscay, france) with modis 250-m imagery. *Continental*
1319 *Shelf Research* 74, 35–49.
- 1320 Pingree, R.D., Le Cann, B., 1992. Three anticyclonic Slope Water Oceanic
1321 eDDIES (SWODDIES) in the Southern Bay of Biscay in 1990. *Deep-Sea*
1322 *Research* 39, 1147–1175.
- 1323 Puillat, I., Lazure, P., Jégou, A.M., Lampert, L., Miller, P., 2006. Mesoscale
1324 hydrological variability induced by northwesterly wind on the French con-
1325 tinental shelf of the Bay of Biscay. *Scientia Marina* , 15–26.
- 1326 Rinke, A., Marbaix, P., Dethloff, K., 2004. Internal variability in Arctic
1327 regional climate simulations: case study for the SHEBA year. *Climate*
1328 *research* 27, 197–209.
- 1329 Rubio, A., Fontán, A., Lazure, P., González, M., Valencia, V., Ferrer, L.,
1330 Mader, J., C., H., 2013. Seasonal to tidal variability of currents and tem-
1331 perature in waters of the continental slope, southeastern Bay of Biscay.
1332 *Journal of Marine Systems* , S121–S133doi:10.1016/j.jmarsys.2012.01.004.
- 1333 Sandery, P.A., Sakov, P., 2017. Ocean forecasting of mesoscale features can
1334 deteriorate by increasing model resolution towards the submesoscale. *Nature*
1335 *Communications* 8. doi:10.1038/s41467-017-01595-0.
- 1336 Solabarrieta, L., Rubio, A., Castanedo, S., Medina, R., Charria, G.,
1337 Hernández, C., 2014. Surface water circulation patterns in the south-
1338 eastern Bay of Biscay: New evidences from HF radar data. *Continental*
1339 *Shelf Research* , 60–76.
- 1340 Toublanc, F., Ayoub, N.K., Lyard, F., Marsaleix, P., Allain, D.J., 2018.
1341 Tidal downscaling from the open ocean to the coast: a new ap-
1342 proach applied to the Bay of Biscay. *Ocean Modelling* 214, 16–32.
1343 doi:https://doi.org/10.1016/j.ocemod.2018.02.001.
- 1344 Toumazou, V., Cretaux, J.F., 2001. Using a Lanczos eigensolver in the
1345 computation of empirical orthogonal functions. *Monthly Weather Review*
1346 129, 1243–1250. doi:10.1175/1520-0493(2001)129;1243:UAEIT;2.0.CO;2.

1347 Vervatis, V., Testut, C.E., De Mey, P., Ayoub, N., Chanut, J., Quat-
1348 trocchi, G., 2016. Data assimilative twin-experiment in a high-
1349 resolution Bay of Biscay configuration: 4DEnOI based on stochas-
1350 tic modeling of the wind forcing. *Ocean Modelling* 100, 1–19.
1351 doi:<http://dx.doi.org/10.1016/j.ocemod.2016.01.003>.

1352 Vervatis, V.D., De Mey-Frémaux, P., Ayoub, N., Sofianos, S., Testut, C.E.,
1353 Kailas, M., Karagiorgos, J., Ghantous, M., in review, 2019. Assessment
1354 of a NEMO-based physical-biogeochemical stochastic model in a high-
1355 resolution Bay of Biscay configuration. *Geoscientific Model Development*
1356 .

	NEMO v2.3	Symphonie 66
horizontal discretisation		Arakawa C
horizontal resolution	1/36°	500 m
vertical coordinates	50 z levels	43 σ levels
vertical resolution	0.5 m to 400 m	even if depth < 100 m, variable with depth thereafter
domain	11–0.26°W, 41–51.5°N	3.5–1.23°W, 43.28–44.49°N
horizontal tracer advection		QUICKEST
horizontal momentum advection	vector form (Le Sommer et al., 2009)	flux-divergence form (Griffies and Hallberg, 2000)
vertical mixing scheme		k- ϵ
atmospheric forcing	ECMWF HRES, 3-hourly 10 m wind velocity, 2 m dew point temperature, surface pressure, air temperature, radiative heat flux and precipitation; 1/8° CORE bulk formulae (Large and Yeager, 2004)	
heat, freshwater and momentum fluxes	Loire, Gironde and Adour	Adour (daily flux from hydro.eaufrance.fr)
rivers		9 tidal harmonics from FES2012
tides at open boundaries	11 tidal harmonics from a barotropic version of the IBI system forced by TPXO 7.1	
ensemble initial conditions	Dec 1st 00:00	Dec 1st 00:00; coastal model initialised from parent on Sep 1st 12:00

Table 1: Model configuration parameters

Experiment name	Description
WIND	Perturb wind only, using the first two EOF modes, with $s = 0.3$.
OBC	Same wind for all members, but each member takes its open boundary conditions from a different member of the parent ensemble.
WIND×OBC	Perturb winds as for WIND, but also vary open boundaries as for OBC.
EOF10	As for WIND, but use the first ten EOF modes instead of just two.
S05	As for WIND, but $s = 0.5$ instead of 0.3.
EOF OBC	As for WIND×OBC but replace our wind perturbations with perturbations reconstructed using the same parameters and random numbers as Vervatis et al. (2016), and preserving the relationship between the wind perturbations and the open boundaries for each member.
EOFBISC	Same as EOF OBC but without varying the open boundaries.
EOF⊥OBC	Same as EOF OBC but the open boundaries and wind perturbations are dissociated.
WIND-NL10	Same as WIND but with a narrower 10-point nudging layer.
OBC-NL10	Same as OBC but with a narrower 10-point nudging layer.
WIND×OBC-NL10	Same as WIND×OBC but with a narrower 10-point nudging layer.

Table 2: Wind perturbation experiments




Figure 1: Bathymetry of the regional parent model. The black square shows the coastal domain.

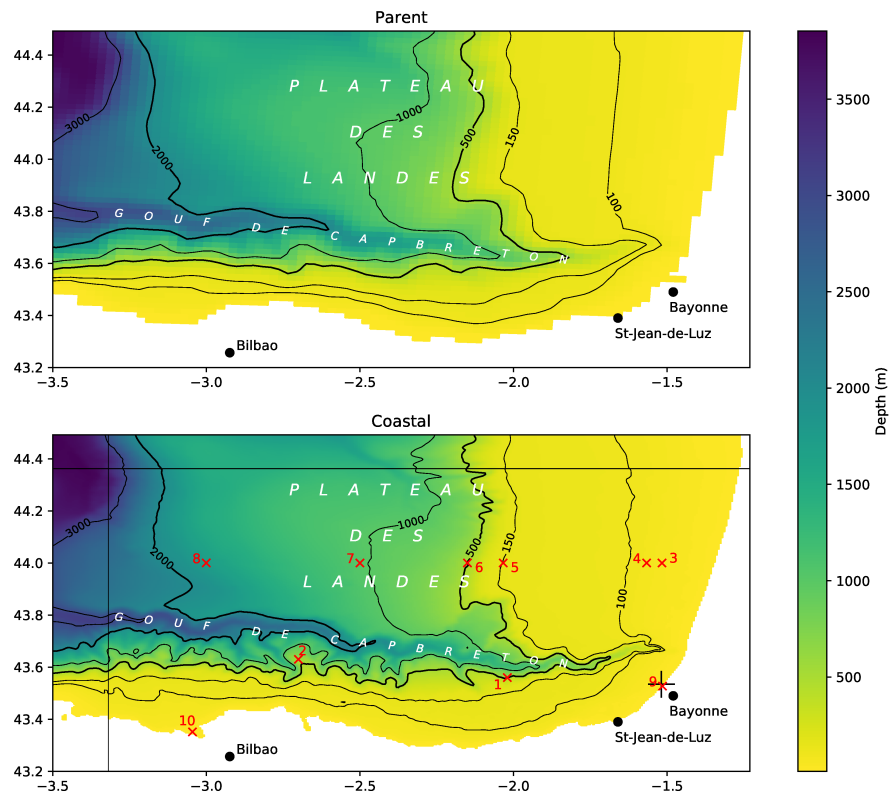


Figure 2: Bathymetries used in the regional parent model (top) and the coastal child model (bottom) shown as they are interpolated onto the model grids in the coastal domain. The numbered crosses show the locations of the high frequency time series used in section 5.2 and figure 5. The vertical and horizontal lines show the edge of the coastal model's 30-point nudging layer. The large, black cross indicates the location of the Adour's discharge in the child model; the notch in the land mask above Bayonne in the parent locates the river discharge there.

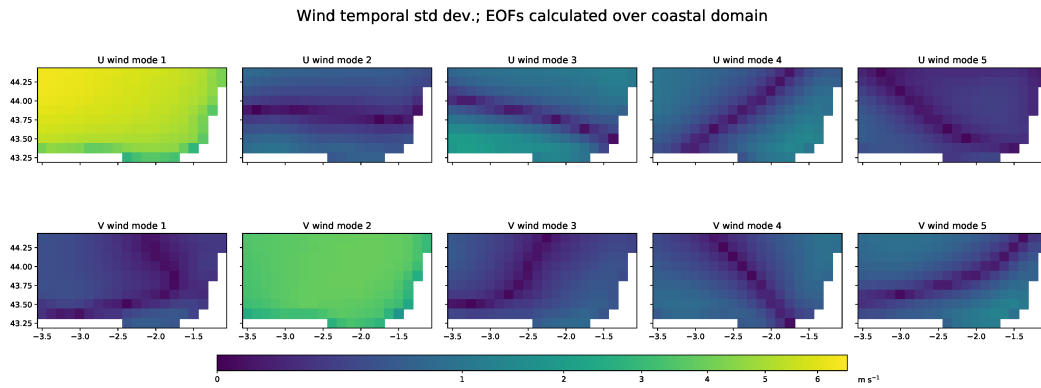


Figure 3: Temporal standard deviations of the first five modes of the wind EOFs for the BOBLAND configurations over the four month period of the ensemble runs. Except for the EOF10ensemble, only the first two modes were used to generate coastal ensembles.

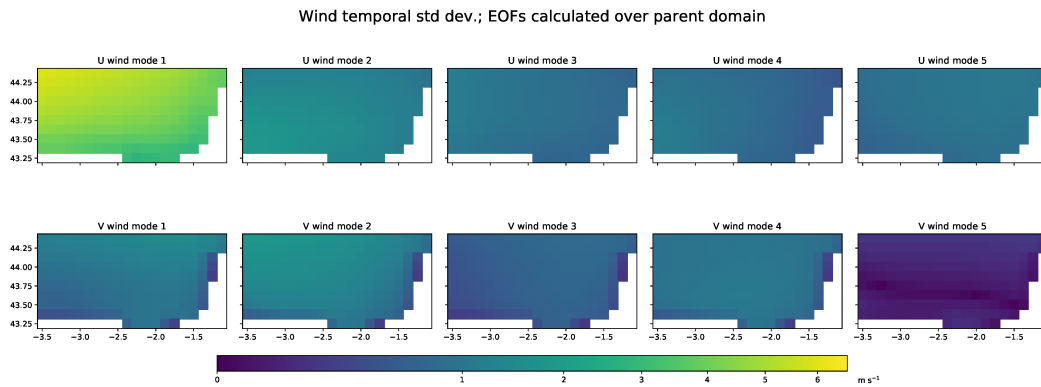


Figure 4: Temporal standard deviations of the first five modes of the reconstructed wind EOFs for the BISCAY36 model configuration, over the four month period of the ensemble runs.




Figure 5: Plot of the ratio (8) for ten point locations indicated by the corresponding numbered crosses in the lower plot of figure 2. The elevations from which the ratio is calculated are high frequency (sampled every ~ 23 minutes).

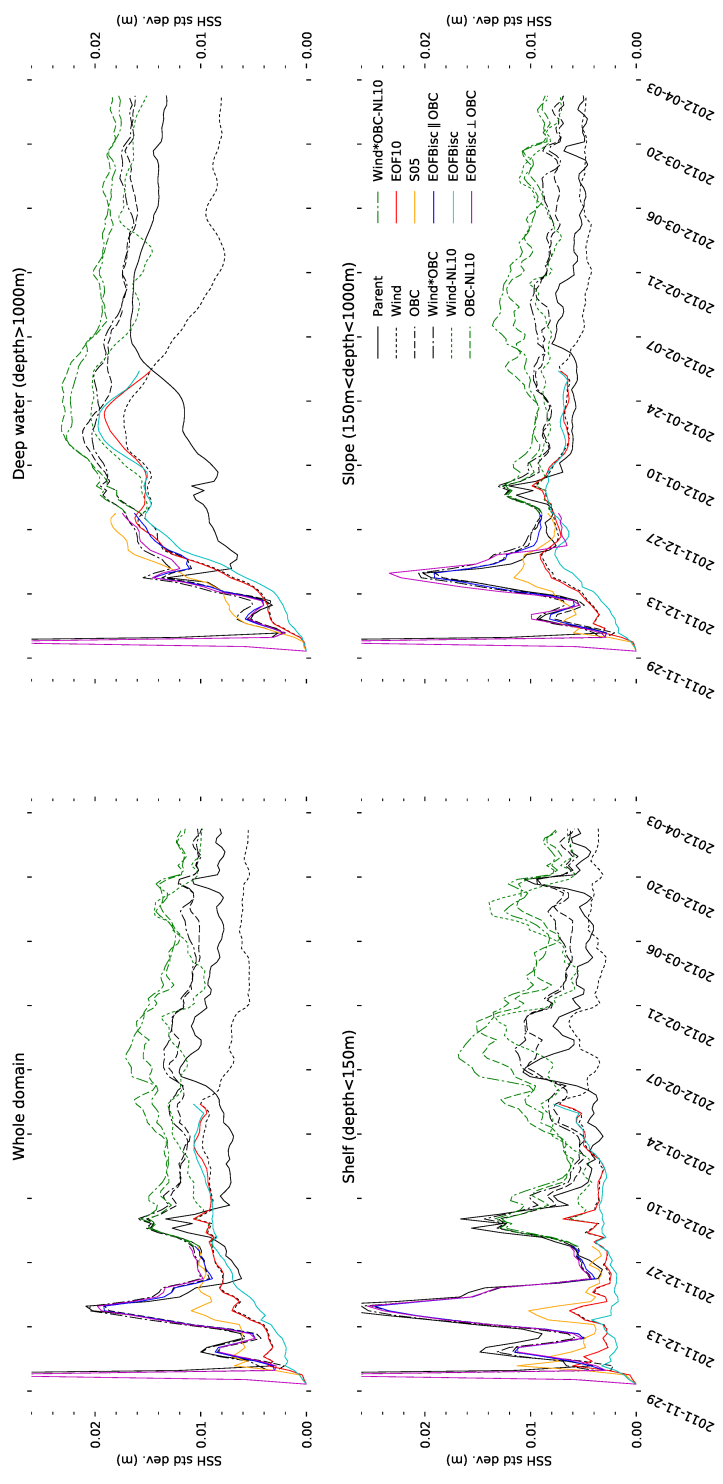


Figure 6: Spatial means of ensemble standard deviations of SSH for the free ensembles and the parent ensemble. The plot at top left is for the whole domain (excluding the nudging layer) while subsequent plots divide the regions by ocean depth. Note that the spikes at the start of the time series reflect spikes in the parent ensemble which come about from the sudden application of the wind perturbations.

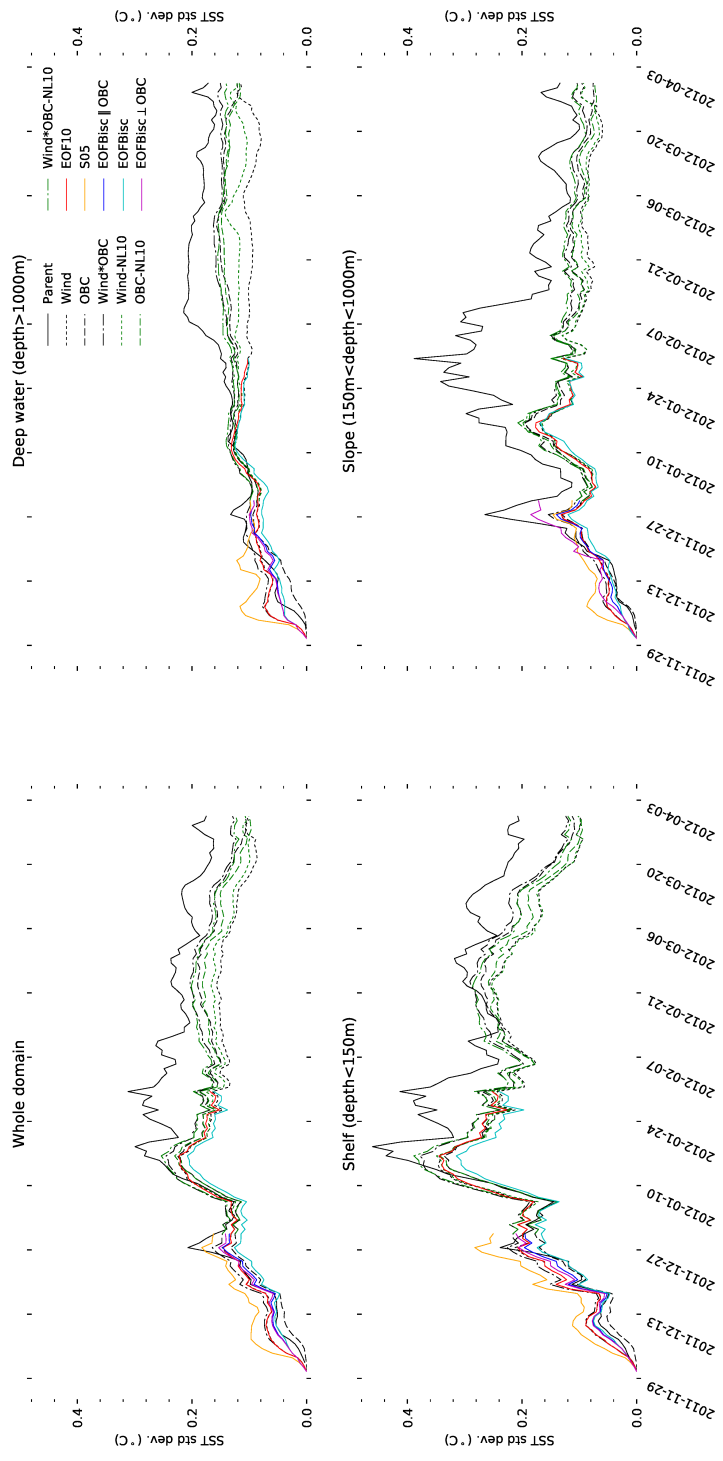


Figure 7: Spatial means of ensemble standard deviations of SST for the free ensembles and the parent ensemble. The plot at top left is for the whole domain (excluding the nudging layer) while subsequent plots divide the regions by ocean depth.

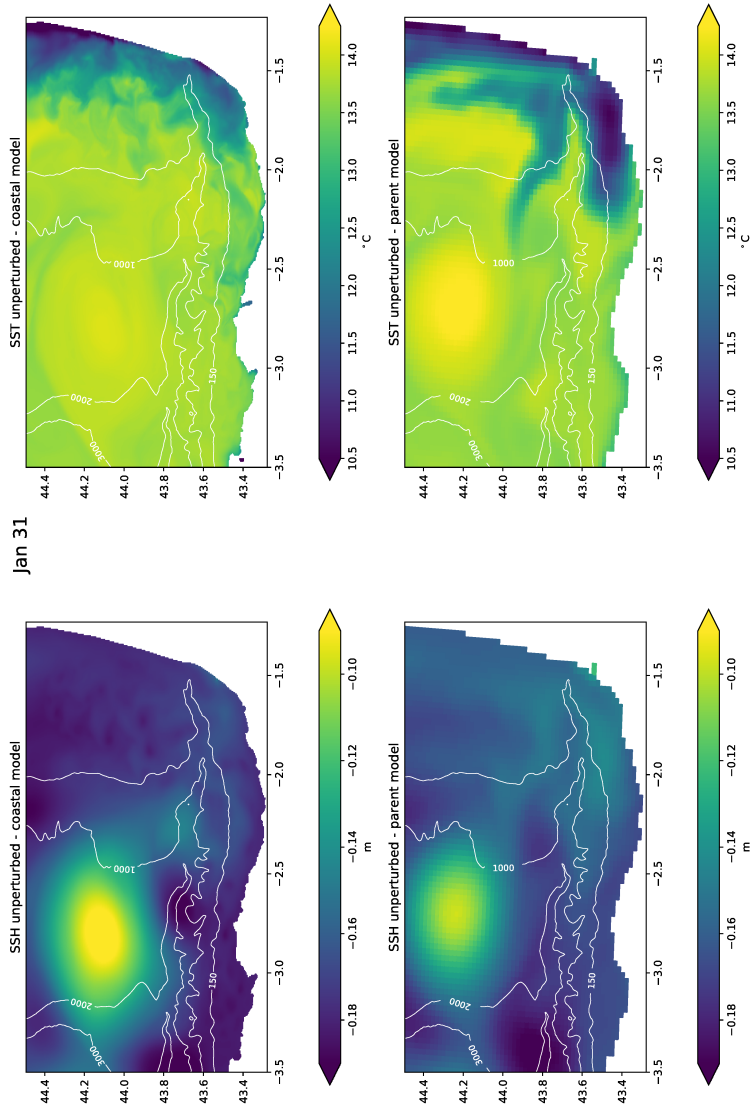


Figure 8: The unperturbed runs for the coastal (top) and regional models (bottom) for January 31st. On the left is SSH, while on the right is SST. The inverted barometer effect of the atmosphere has been removed from the SSH signal. (The Mundaka estuary has been masked from the SSH plot for the coastal model.)




Figure 9: Time mean over the three months of January, February and March of ensemble standard deviations of SSH for the four free ensembles.

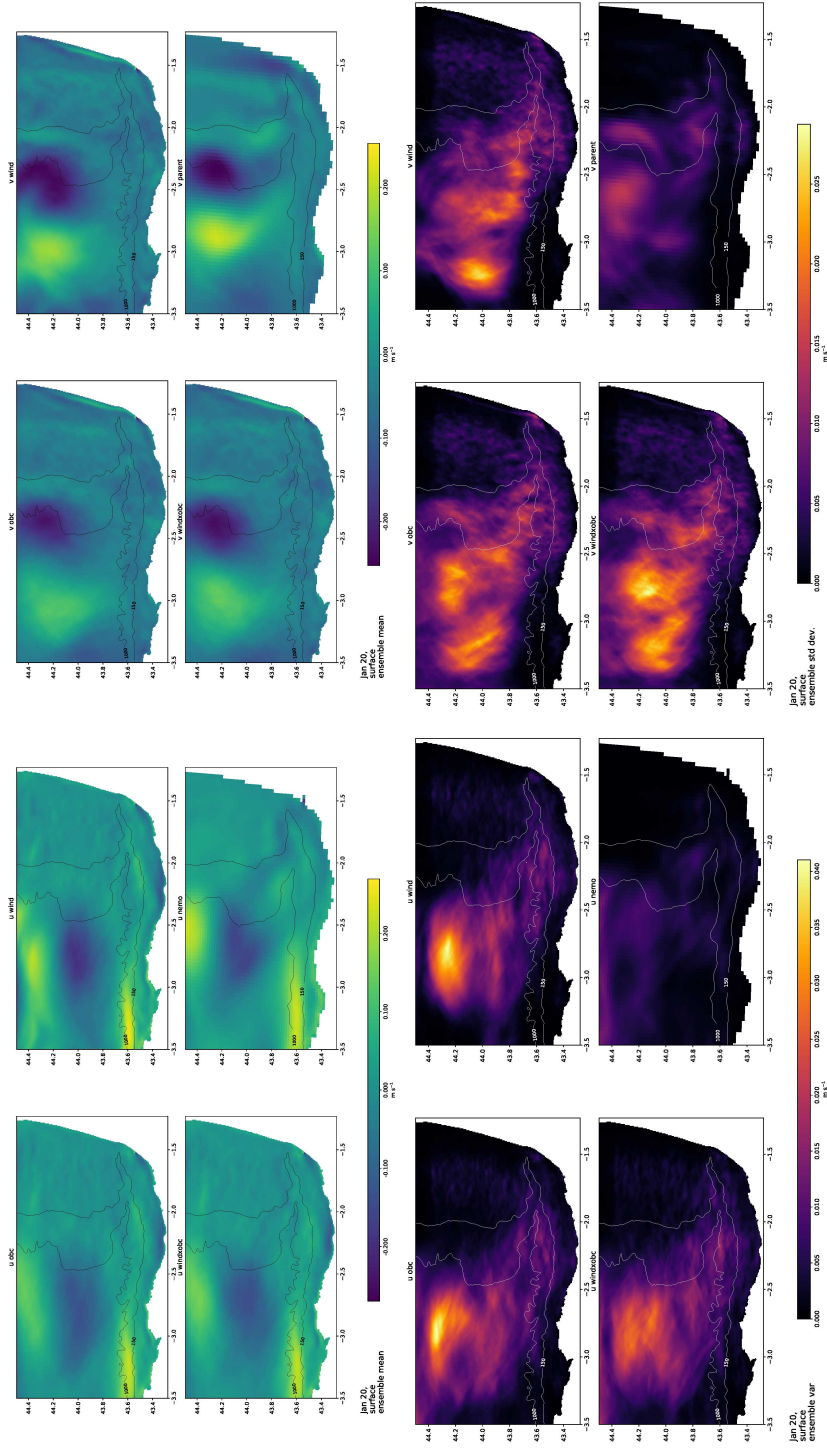


Figure 10: Velocity ensemble means and ensemble standard deviations of the daily mean on the 20th of January for all four ensembles, as labelled in the figure. The left columns are of zonal velocity, the right columns of meridional; means are on top, standard deviations on the bottom.




Figure 11: Velocity ensemble means and ensemble standard deviations of zonal velocity, the right columns of meridional; means are on top, standard deviations on the bottom.

SSH daily means for the Wind ensemble members for Mar 9

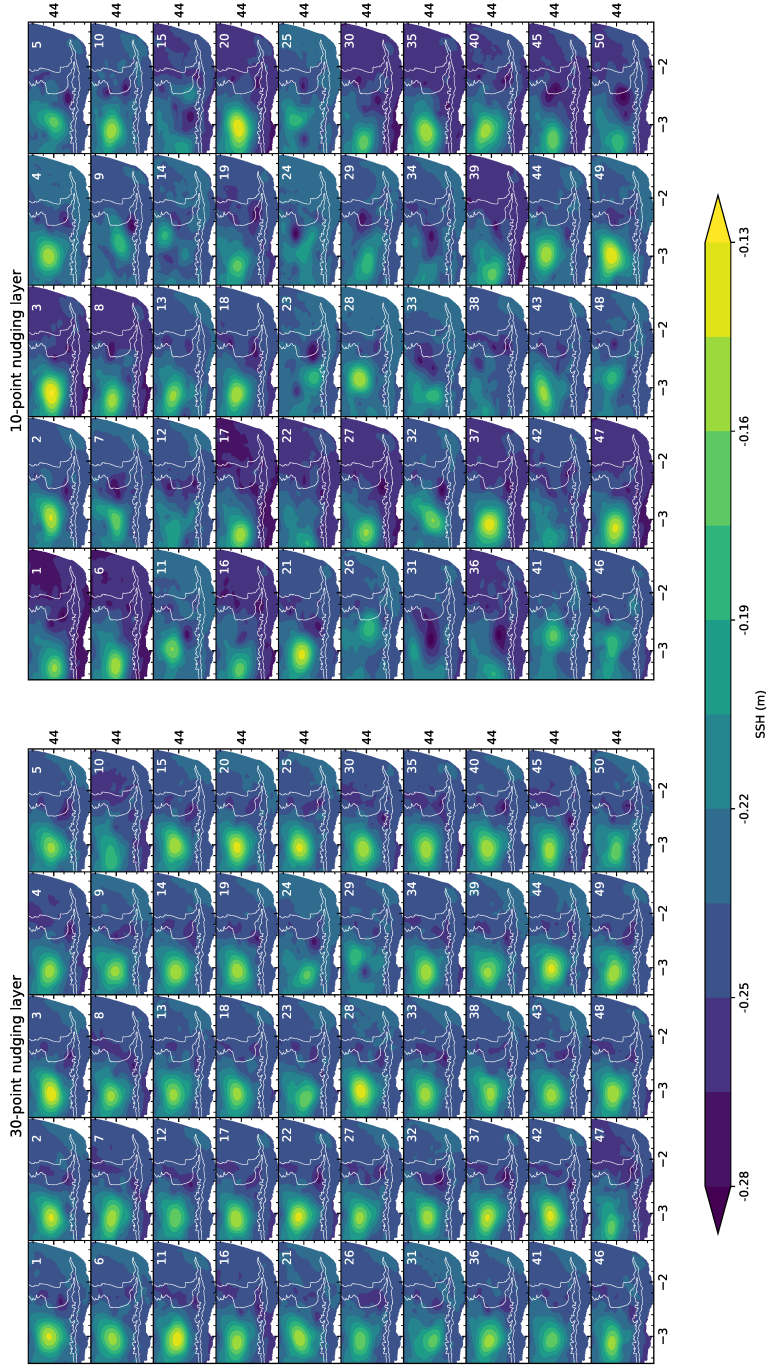


Figure 12: Illustration of the effect of different nudging layer widths on the WIND ensemble's SSH spread. Except for the unperturbed run, each member's SSH is shown, with the number in the top-right corner of each plot identifying the ensemble member. With the narrower nudging layer we immediately see a bigger difference in the location and amplitude of the anticyclonic eddy from member to member across the ensemble. The atmospheric inverted barometer effect has been removed from the SSH signals.




Figure 13: Time mean over the three months of January, February and March of ensemble standard deviations of SST for the four free ensembles.

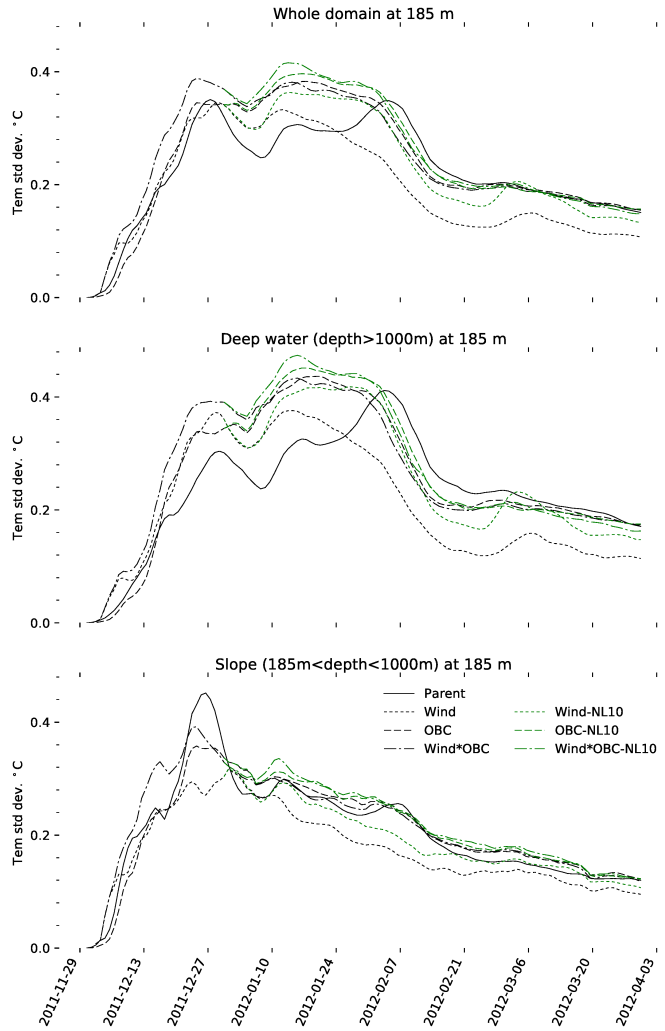


Figure 14: Spatial means of ensemble standard deviations of temperature on a horizontal slice at 185 m depth for the free ensembles and the parent ensemble. The plot at top left is for the whole domain (excluding the nudging layer) while subsequent plots divide the regions by ocean depth.

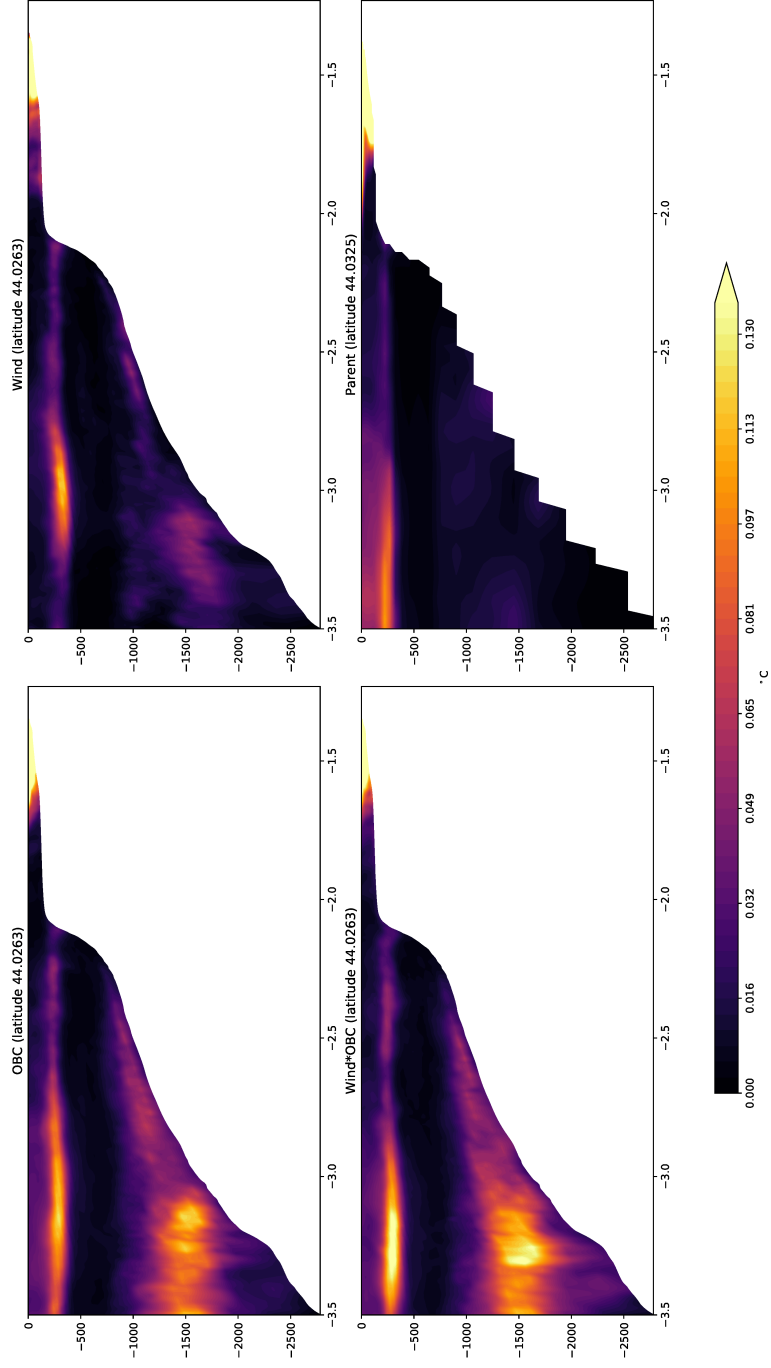


Figure 15: Cross sections of standard deviation of the daily mean temperature for March 11 for the four ensembles at 44.0263° latitude.

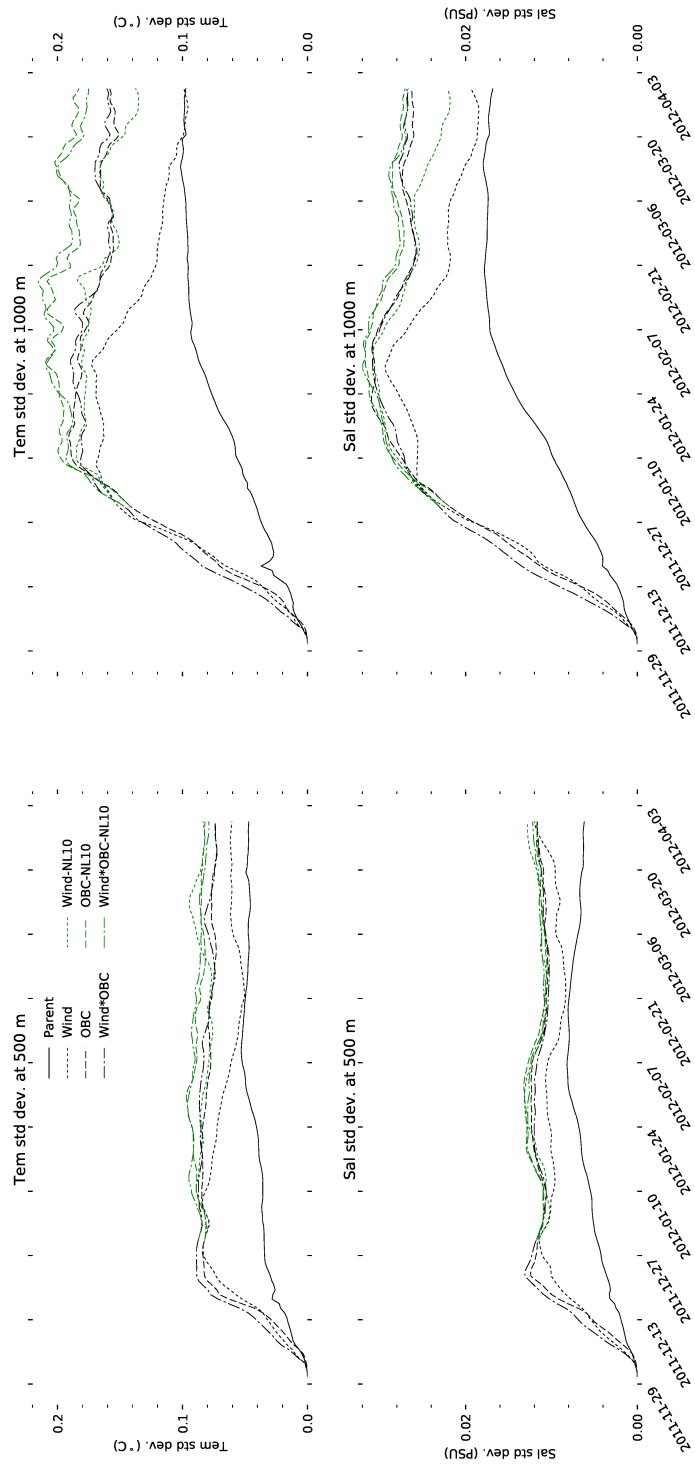
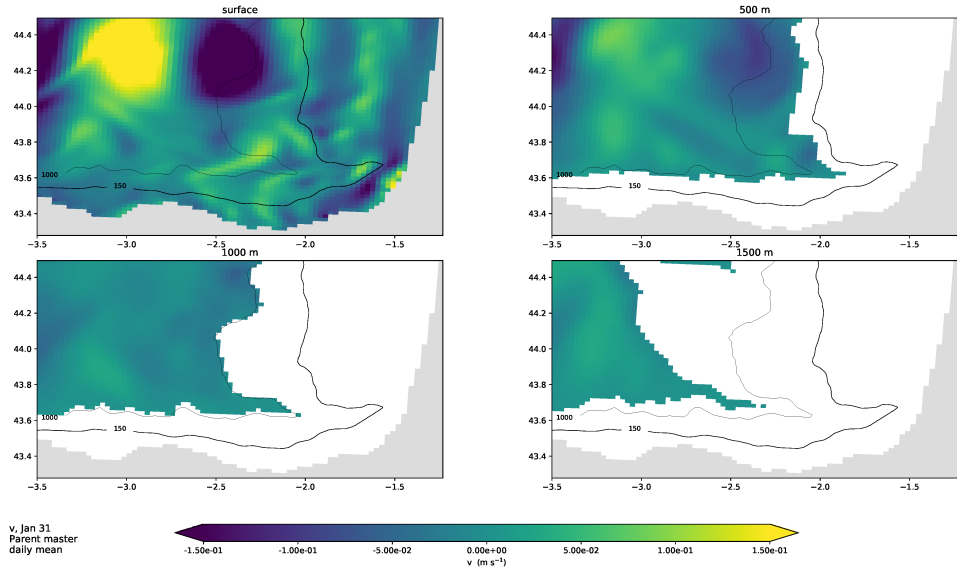
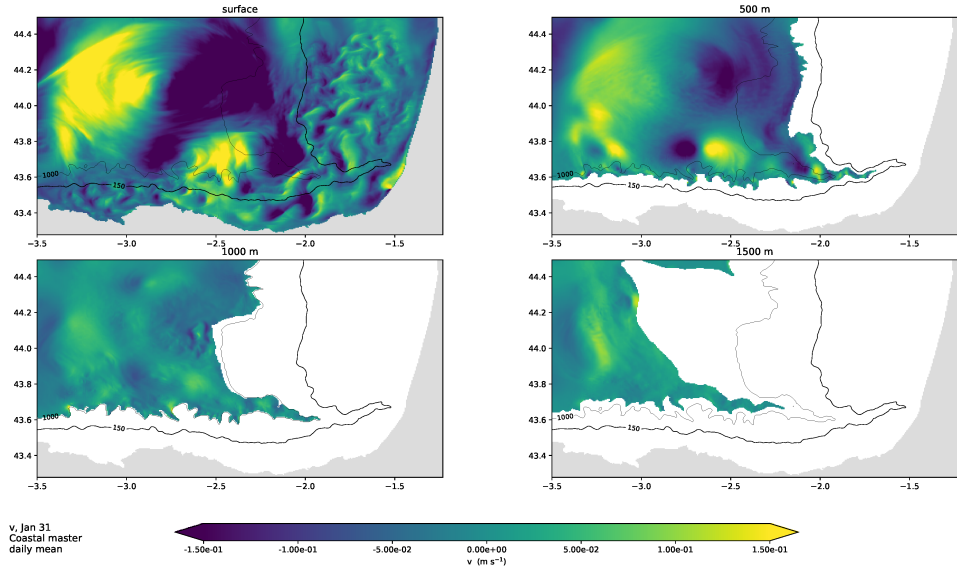


Figure 16: Spatial means over the whole domain (excluding the nudging layer) of ensemble standard deviations of temperature (top row) and salinity (bottom) on horizontal slices at 500 m (left column) and 1000 m depth for the free ensembles and the parent ensemble.



(a)



(b)

Figure 17: Meridional velocity at different depths for the unperturbed run of the parent (a) and coastal (b) models. Daily means for 31st January.

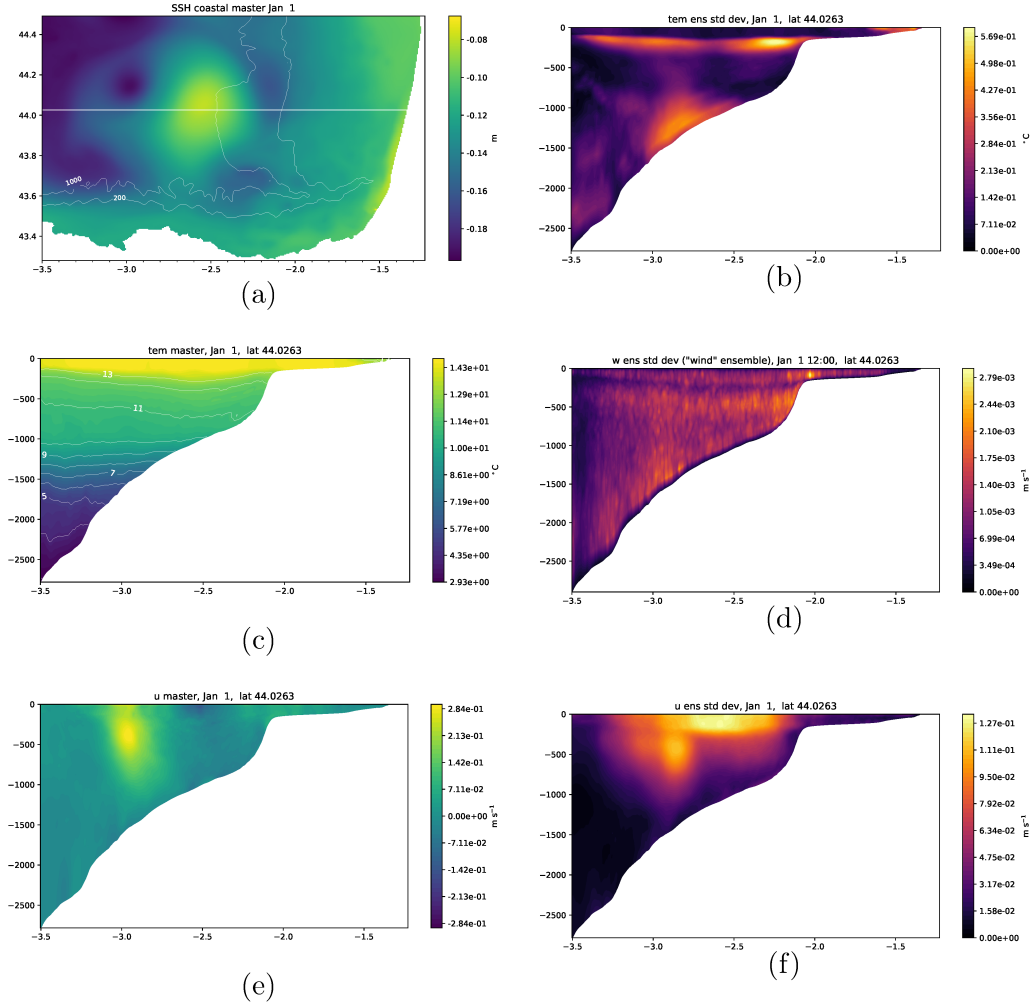


Figure 18: Illustration of how forcing perturbations at the surface alone can affect the model response at depth through eddy interaction with the bottom—plots for the 1st of January, 2012. In (a) the horizontal line shows where the cross section in subsequent plots is taken, and 200 and 1000 m depth contours are shown. (a) shows the SSH for the unperturbed model (with the atmospheric inverted barometer effect removed from the SSH signal); (b) the standard deviation of temperature; (c) the temperature for the unperturbed model; (d) is the standard deviation of vertical velocity; (e) the zonal velocity from the unperturbed model; (f) the standard deviation of zonal velocity. The vertical velocity plot is an instantaneous field, whereas the others are daily means.




Figure 19: Graph of ensemble mean for river plume areas defined by different salinity thresholds for the parent ensemble and the three coastal ensembles. The threshold values are 32, 34.5 and 35 PSU (low, mid and high, respectively).




Figure 20: Salinity (left) and norm of salinity gradients (right) for the WIND×OBC (upper plots) and parent ensembles' member 1; daily mean for 19th January. Note the geometric colour scales.