Response to the reviewers

We thank the reviewers for their in-depth reviews of the paper, and their constructive comments. Below we give detailed point-by-point responses, with any numbers referring to equations, figures, or lines referring to the original and revised paper versions in the reviewer comments and replies, respectively.

Reviewer 1

General comments

The paper describes how the recently proposed two-scale momentum (2SM) theory can be coupled to a atmospheric perturbation model (APM) developed by the authors. The approach is novel and the paper well written, mathematically sound and mostly easy to follow, however the structure, validation and "fast" element of the paper need improvement.

The only part that seems to be "fast" is hidden in subsection 4.3.1, so either the title needs to be adapted, as it is currently misleading, or more work needs to be spend on timing the three different flavours of the coupling that are presented and show how the currently not so "fast" approaches can be sped up. Whilst going through section 4 and 5 the motivation of using one approach over the other remains somewhat obscure and currently just seems like a list of possible methods.

While the paper is mathematically rigorous, some decisions and sensitivities need to be explored further, especially how to choose the size of the control volume in the 2SM approach such that the error in determining the average wind farm wind speed, Uf, does not depend on the wind farm layout. The authors currently are able to avoid addressing this issue, as they use idealised square wind farms, the same idealised conditions also the 2SM theory is based on. At least one non-rectangular wind farm (or at least run one different wind direction) test case need to be added to showcase that the proposed approach has more general applicability and provides similar results as for an idealised setup. Additionally, one case above rated wind speed could be added, such that the farm has some CT, CP variation (potentially already the case but not explicitly discussed in the paper.

Reply: We would like to thank the reviewer for their in-depth review of the paper. We have tried to address them to the best of our abilities, and believe it significantly improved the paper.
We appreciate the reviewer's concern on the "fast" aspect of our work. We have added a discussion on the computational cost of the conventional APM at the end of section 3, and have made the effect of the derivation in section 4.2 on this cost more explicit. To differentiate the two model variants, we have added a paragraph in the conclusion on when one might be preferable over the other.
Finally, we agree with the reviewer that the current model and its validation is limited to idealized farm layouts and operating conditions. However, performing additional high-fidelity simulations is not feasible for us in the scope of the current paper. We have made modifications throughout the paper to emphasize that the used LES dataset represents idealized conditions, and have indicated how future work could extend the model derivation to arbitrary farm layouts. Finally, we have added an additional concluding paragraph explicitly state these limitations:
As powerful as the tools offered by the 2SM approach are, it remains unclear how valid the underlying assumptions are outside of the idealized conditions they have been tested in. So far, they have only been validated using LES of infinite farms (Kirby et al., 2022) and farms operating in horizontally homogeneous background conditions with uniform turbine layouts (Kirby et al., 2025). Future work should investigate

whether the farm coefficients remain constant under realistic and transient atmospheric conditions, and how to account for irregular turbine spacing and operating setpoints.

Specific comments

Title

Reviewer Point P 1.1 — The title is a little too general with regard to the content as there are fast blockage models out there, however not when gravity-waves are considered.

Reply: We have changed the title to
"Fast models for blockage and gravity-wave effects on wind-farm power output".

Abstract

Reviewer Point P 1.2 — The abstract should more clearly define to what problem the model should be applied to. I guess the aim is to provide accurate estimates of wind farm power for large wind farms when considering gravity wave effects at low computational costs? Please clarify.

Reply: We re-phrased the sentence describing the combination of 2SM and APM approaches to emphasize the goal of developing a faster version of the APM. It now reads:
To address this, this paper develops ways of speeding up atmospheric perturbation models (APMs) by an order of magnitude by applying insights from the recently proposed two-scale momentum (2SM) theory (Nishino and Dunstan, J. Fluid Mech., vol. 894, 2020, A2).

Reviewer Point P 1.3 — One might consider providing a reference to "two-scale momentum theory" as it is not that established yet.

Reply: We have added a reference to Nishino and Dunstan (2020).

Introduction

Reviewer Point P 1.4 — First paragraph: Whilst the numerical proof that large farms can trigger gravity waves is compelling and referenced here, this section of the text does not reflect that all existing studies related to gravity waves remain numerical - physical or operational validation of their significance remains missing. Please at least include a statement that these are numerical studies (Bleeg's paper is different but also does not discuss gravity waves or atmospheric interactions)

Reply: We have modified the paragraph to explicitly mention the cited studies being based on LES results. Furthermore, Bleeg (2018) is now cited in the first sentence, referring strictly to the blockage effect and not to gravity waves. The paragraph now reads:
As offshore wind turbines and farms continue to grow in both size and installed capacity, they reach the point where they form an obstacle to the entire atmospheric boundary layer (ABL) (Bleeg et al., 2018). Large Eddy Simulation (LES) studies have shown that in conventionally neutral boundary layers (CNBLs), where

a neutral ABL is topped by a capping inversion and a stratified upper atmosphere, wind farms can trigger gravity waves, creating upstream pressure gradients and blockage effects (Allaerts and Meyers, 2017, 2018; Centurelli et al., 2021; Lanzilao and Meyers, 2022, 2023a, 2024; Maas, 2023a, b; Stipa et al., 2024b). As these mesoscale effects affect turbine power output, they are becoming more important to farm operation and planning.

Reviewer Point P 1.5 — Second paragraph: It is not just the cost, but also the complexity of the inputs.

Reply: Thank you for this interesting point. While this is true for LES, where generating turbulent inflow conditions is a challenge, RANS and mesoscale models can directly assimilate weather predictions or reanalysis data. Overall, we do not believe input complexity is as big of a challenge to these methods as computational cost. Therefore, we make no change.

Reviewer Point P 1.6 — Third paragraph: One might want to mention what type of atmospheric dynamics it captures.

Reply: We have rephrased it to *"the generation of blockage and gravity waves"*.

Section 2

Reviewer Point P 1.7 — L44-51: The LES dataset is not only used for validation but also used in deriving model constants. This should be mentioned.

Reply: In response to later comments on the validation procedure (see below), we changed the fitting of the model constants to only use NBL cases, so that there is no overlap between the cases used for model development and the cases on which the model was tested. Therefore, we leave this sentence in the introduction unchanged.

Reviewer Point P 1.8 — L73: "...assumed to only depend on the internal problem." Is this an issue when coupling it to the APM, as this might change the turbine set-points?

Reply: This is a valid concern. Without additional simulations of farms in above-rated conditions, there is no way to say how big this issue is. We have modified the paragraph following this sentence to reference this limitation more explicitly (see next comment).

Reviewer Point P 1.9 — L74: It did work for their test cases, which are idealised square wind farms staggered and aligned configurations with a single turbine type and below rated operation (ie CT and CP not changing much inside farm). This limitation should be highlighted as it might not generally hold.

Reply: We re-wrote this paragraph to more explicitly reference these limitations. The first sentences now read:
Kirby et al. (2025) investigated this assumption using LES results of large wind farms with uniform turbine densities in below-rated conditions. They found that the internal power coefficient C_p^* as defined in equation (1) is a good heuristic for internal effects, and remained constant over cases with different mesoscale flow behavior.

Reviewer Point P 1.10 — Paragraph from L78: Please mention the operating point of the turbines that was used. I suppose below-rated or close to that? If so this does align with the assumption mentioned above. This should be made explicit.

Reply: We have added *"operating at below-rated conditions"* when describing the wind farm, and now explicitly mention that this same dataset was used by Kirby et al. (2025).

Reviewer Point P 1.11 — Paragraph from L86: "... we believe ..." is not sufficient justification. Kirby showed for their test cases that 2.5zh worked, you found it is 2zh, but mostly because otherwise you were reaching into the upper atmosphere. So is this some implicit tuning parameter? We just need to find some height and averaging volume that ensures that it matches what we would like to see? Please add a sensitivity study with respect to Hf and the choice of averaging volume size, as otherwise this seems somewhat arbitrary. Please also add at least one test case for a non-square wind farm and describe how the volume should be calculated in that case.

Reply: We did not mean for this paragraph to represent a tuning of H_f, or a general recommendation for the optimal averaging volume in all atmospheric conditions. Therefore, we re-wrote this paragraph with an emphasis on justifying the choice of H_f as corresponding to the APM layer distribution, and removed the comparison with earlier literature. Furthermore, we added a brief reference on how complex farm layouts could be investigated by future work:
We calculate U_f as a straightforward average over the wind farm, so that

$$U_f = \frac{1}{H_f \Omega_f} \int_0^{H_f} \iint_{\Omega_f} u d\Omega_f dz,$$

where H_f is the height of the 2SM averaging volume, and Ω_f the planform area of the wind farm. We take Ω_f to be a rectangle around the farm edge, thereby including the outer edges of the turbine disks at the sides. We set the height H_f of the volume to twice the turbine hub height z_h, as this corresponds to the setup of the APM (see section 4.1). This height is above the turbine tip height but below the ABL height for all the cases in the dataset of Lanzilao and Meyers (2024). While some of the H150 cases have lower unperturbed ABL heights, strong vertical displacements of the capping inversions over the wind farm ensure that even there the ABL height is above the 2SM control volume. The exact dimensions of the resulting control volume are listed in table 1. This simple approach is possible because the analyzed farm has a convex and uniform layout. More realistic farms might require a more complex horizontal averaging method such as the one proposed by Van der Laan et al. (2023), where the wind speed is weighted by Gaussian distributions centered around the turbines. Alternatively, a farm could be split in two or more homogeneous regions. Further investigation of such approaches is beyond the scope of this study, and left as promising future work.
We agree that the sensitivity of the dependence of C_p^* on H_f requires further research. We speculate that in the CNBL conditions studied here, the dependence is minimal as long as H_f is located above turbine tip height and below the capping inversion. However, without more high-resolution CFD data we can not generalize this to other atmospheric conditions. Since this analysis constitutes a deeper investigation of 2SM theory in itself, we feel this is outside the scope of the current paper. Therefore, we do not comment on it further.

Reviewer Point P 1.12 — Paragraph from L95: Could you elaborate on the scaling? Could the good match between 2SM and LES not also stem from performing this scaling?

Reply: At finite grid resolutions, the actuator-disk method introduces an error in the calculation of the power output as the higher order moments of the wind speeds are taken over the disk footprint. For more information on this issue, see Shapiro, Carl R., Dennice F. Gayme, and Charles Meneveau. "Filtered actuator disks: Theory and application to wind turbine models in large eddy simulation." Wind Energy (2019). We have re-phrased the sentence to refer to this paper as well.
The scaling does not affect the good match at all, as the difference across cases between P_f is much larger than the difference in P_0. To show this, figure RR1 shows the same results as those reported in the manuscript, but without the scaling.

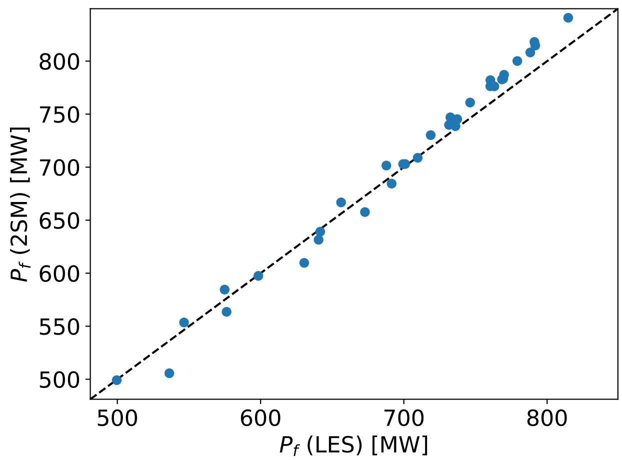

Figure RR1: A comparison of farm power P_f as calculated by LES (x-axis) and equation 1 (y-axis) for all 36 CNBL cases from the dataset by Lanzilao and Meyers (2024).

Section 3

Reviewer Point P 1.13 — Equation 3: I am missing a definition of x,y,z

Reply: Thank you for noticing this. We have added their definitions as the streamwise, spanwise, and vertical coordinates, respectively.

Reviewer Point P 1.14 — L119: Is this the same filter size as in Equation 11 and is it fixed?

Reply: Yes, it is. We have modified the text introducing the filter length ℓ to explicitly say it is fixed, and refer to this earlier filtering when introducing the turbine forces.

Reviewer Point P 1.15 — L160-: Does this mean that this tight coupling at individual turbine level is the issue for high computation costs? It could be nice if such a statement could be build in here and maybe already be included in the introduction to make it very clear. The 2SM model does not say that the forcing does not vary within the farm (at least the impression I got from a statement in the introduction), but only that it can all be related to some average farm wind speed, which is actually very similar to what the actuator wind farm models have shown, and that is what you are leveraging here, assuming that pressure perturbations will not change forces significantly or the turbine operating points.

Reply: Yes, this is correct. Thank you for the suggestions for making this more clear to readers. We have modified the paragraph in the introduction discussion the conventional APM, so that it explicitly mentions this tight coupling as the main cost of the model:

. . . Building on earlier work by Smith (2010), this model simulates the large-scale wind-farm atmosphere interaction explicitly, while coupling to conventional wake models to account for turbine-scale flow effects. In order to capture the local inter-farm feedback effects between the turbine forces and the mesoscale flow, the governing system of equations is solved repeatedly in a fixed-point iteration. This constitutes the main cost of the model, which is low compared to conventional CFD tools but relatively high compared to standard wake models (several hundred to a thousand times slower, or approximately one to five minutes on a standard laptop per simulation) (Devesse et al., 2024b).

We also agree that the paragraph in the introduction referring to "variations" was inaccurate. We have expanded this introductory paragraph to more clearly state our goal and why we work with the APM, so that it now reads:

Our goal is to develop fast models for predicting wind-farm power output that account for blockage and gravity wave effects. To that end, the 2SM framework can be useful, as it implies that flow solvers simulating the farm-scale flow do not have to resolve the turbine-scale flow. This would allow for coarser grid resolutions and reduced computational costs. However, even when using coarser grids, conventional CFD tools are still multiple orders of magnitude slower than engineering wake models (Ouro et al., 2025). In contrast, the APM already uses a coarse resolution, while its main cost lies in capturing the local feedback effects between the turbine forces and the mesoscale flow. Therefore, we focus on combining the 2SM and APM approaches, leveraging insights from the former to speed up and simplify the latter.

Since, according to the 2SM hypothesis, a farm's power output only depends on the average velocity in the farm, the APM may neglect local changes in the turbine force distribution within the farm induced by gravity wave effects. Instead, this force distribution can be scaled uniformly over the farm to match the total farm force as predicted by 2SM theory. Due to the linear nature of the APM, this transforms the model's governing equations into a system of scalar equations, drastically lowering its computational cost. We use this to develop . . .

Additionally, we have added the following paragraph at the end of section 3, in order to clearly state our goal:

The main numerical costs of the classical APM lie in solving equations 7-11 and re-calculating the wake model background velocity at each step of the fixed point iteration (Devesse et al., 2024b). These repeated calculations were needed to resolve the local coupling between the mesoscale APM state ($\vec{u}'_1, \eta_1, p'$) and the turbine forces. However, 2SM theory allows us to neglect this local coupling, and consider a farm-averaged approach instead. This is the focus of the following sections.

Section 4

Reviewer Point P 1.16 — L186: I disagree that this is "straightforward" at all and that's due to Ω_f. How do we define a wind farm planform area in a realistic setup? In you test cases and those of Kirby this might be "straightforward", but in any practical application it is not. There should be a discussion around how to determine Ω_f for real farms in a consistent manner, otherwise Uf could be biased depending on farm configuration. Actuator wind farm models also struggle defining Uf consistently, but some practical solutions exist and should be discussed.

Reply: We agree with the reviewer that the equations given here are inadequate for general wind farm

layouts. However, this is an issue with general 2SM theory, and not with the 2SM–APM methods we are developing here. Therefore, we have moved the definition of U_f and the discussion on the averaging volume to section 2 (cfr. point 11). There, we cite the averaging method proposed by Van der Laan et al. (2023) or a multiple-zone split of a farm as possible approaches to be investigated by future work.
We have added a short note below the equation re-stating the definition of U_f at the start of section 4 to reference this discussion:
As noted in section 2, modeling wind farms with irregular turbine layouts may require more complex horizontal averaging operators. However, this does not affect the following discussion conceptually.

Reviewer Point P 1.17 — Equation 14: Wind farm forcing is being applied with Gaussians, but forces are applied outside of the wind farm planform if it is defined as the outer turbine perimeter, however Ω_f is not defined so it remains unclear. The 2SM probably requires all forces to remain inside the CV.

Reply: The turbine footprints are indeed non-zero outside of the wind farm planform. However, to our knowledge 2SM theory only proposes that the total farm force scales with the average farm velocity, and does not prevent these forces from being smeared out. As long as the Gaussian filtering does not affect the calculation of the average farm velocity, there should be no issues.
We already clearly identify this discrepancy between filtered and unfiltered flow in the paragraph below equation 16, and clearly state that we neglect it in equation 17. Therefore, we have made no further changes.

Reviewer Point P 1.18 — L215: To have a tuning parameter that provides good results is nice, but it would be good to show how it varies between the different cases. Will it change if we change CV and wind farm size, this sensitivity should be explored.

Reply: The tuning parameter α is an attempt to capture the changes in the vertical wind profile above the wind farm, and will therefore depend on the relative strengths of the internal boundary layer growth and pressure gradients induced by gravity waves, among other factors. As was the case for determining the optimal H_f, the available dataset is too limited to adequately explore this dependence. Furthermore, following a later comment on our validation approach (P 1.22), we now only tune α on a neutral case, which prevents a further investigation of how this parameter changes with atmospheric conditions.
We have added a comment when introducing α to indicate the limits of this basic tuning approach:
The velocity in this region depends on the relative strength of the turbulent entrainment and gravity-wave-induced pressure gradients. Here, we take the simplest possible approach by assuming that the average velocity perturbation in the part of the APM layer above the 2SM volume scales with the velocity deficit within it.

Reviewer Point P 1.19 — Equation 20: Should this not be f_{wf} or am I missing something here?

Reply: Thank you for noticing this oversight. We have corrected it throughout the manuscript.

Reviewer Point P 1.20 — L235: Could make it a little more clear what you mean by “APM state”, as lot seems to be possible once we assume it is not changing.

Reply: By this, we mean the mesoscale velocities, layer thicknesses, and pressure perturbations that form the solution of equations 7-11 for a given turbine forcing, i.e. the variables $(\vec{u}_1', \eta_1, p')$. We have added these variables when mentioning the APM state to clarify the statement.

Reviewer Point P 1.21 — L247: This is a little unclear. So you are running an engineering wind farm simulation, excluding atmospheric effects etc to get the footprint? Also can it really be chosen "freely"? Is it not somewhat related to the CV used to compute Uf?

Reply: We agree that the phrasing did not clearly communicate how the footprint was chosen. We rewrote it as follows:
The closer the footprint $\vec{\Pi}_f$ is to the actual normalized force distribution, the more accurate we expect the coefficients a_u and a_η to be. However, this actual footprint is not known a priori. Instead, we use a standalone wake model to estimate the footprint cheaply. The wake model can predict the turbine disk velocities S_k, which can be combined with equations 14–15 to find the farm forcing, so that

$$\tilde{\vec{\Pi}}_f = \frac{\sum_k^{N_t} \left(C'_{T,k,WM} A_k S_{k,WM}^2 \vec{e}_x \right)}{\sum_k^{N_t} \left(C'_{T,k,WM} A_k S_{k,WM}^2 \right)}, \tag{1}$$

where the subscript WM indicates the variables are calculated using a standalone wake model. This approximates the footprint of a farm in neutral conditions at a low computational cost.

Reviewer Point P 1.22 — L264: This is not really a validation, it is only a verification, unless I misunderstood your approach. You have derived/trained your model constants using the LM LES cases and now you are running your model and comparing to them. Validation would be if you had another set of data that you could compare to that you did not use in training. Please modify and make this clear unless you have some other test cases using that data you did not use in deriving your model.

Reply: We appreciate the reviewer's concern. The only model constant we tune is the α parameter relating the 2SM averaging volume to the APM layer. In order to avoid using the same cases in the validation as in the model development, we changed the tuning of this α to now only use the NBL case H500-C0-G0, so that the CNBL cases we use to validate the model are never used in its derivation or during tuning.

Reviewer Point P 1.23 — L275-: What are the error metrics based on (total wind farm power)? That should be mentioned. Generally errors should be discussed in more detail, as there are some cases with large errors and also a general bias towards overestimating production. Also why are you not also comparing to your fully coupled APM model, maybe that would be a fairer comparison? You mention that errors are similar in the next paragraph, but maybe they are biased differently? There is no way to discern that here.

Reply: Thank you for these suggestions. We now state explicitly that the figures and table show the results for the total wind farm power P_f.
We also expanded the discussion on the model errors in multiple ways. Firstly, we added a figure that ranks the power output predicted by LES and the 2SM–APM per case. Secondly, we now show two additional cases in the figure of the centerline velocity, and added a sentence on the model's prediction of resonant lee waves:
*In general, accurate predictions of the average velocity are in part due to a cancellation of errors, as the model over- and underpredicts the velocity at the front and back of the farm, respectively. While the model predicts the wrong location for the resonant lee waves for some of the H*300 *cases, as can bee seen in figure 6 (right), this has no consistent impact on its prediction of P_f.*
Finally, we added a brief discussion on the model error behavior:

Looking at figure 5, there seems to be no consistent relation between the model error and the varying atmospheric parameters H, $\Delta\Theta$, or Γ. Instead, figure 4 shows that there is a small but consistent positive bias, and that the scatter on the model's predictions is larger at lower farm power outputs, i.e. for cases with stronger blockage effects.

Reviewer Point P 1.24 — Table 2: Explain what all the model cases are. This is not mentioned anywhere. In the text one could also go into a little more detail (best case, worst case scenario . . .)

Reply: We added a column "Section" that references the parts of the paper where each of the listed models is discussed or derived. We also expanded the discussion in the text (cfr. the above reply).

Reviewer Point P 1.25 — L287: It should be highlighted that this is closer to the real use case, as if we cannot run a LES simulation first to get Cpstar, Ctstar etc if we want a fast model.

Reply: We agree with this comment. We have rewritten the introductory paragraph of this section to better reflect this:
When using the 2SM–APM in practical settings, it is not feasible to rely on LES to calculate the farm thrust and power coefficients C_p^* and C_T^*, as these simulations are too expensive. Therefore, this section aims to briefly demonstrate how a basic wake model can be used to calculate these coefficients instead of LES. We merely wish to show that this approach is conceptually possible and point out some potential issues, and do not attempt to minimize its computational cost.

Reviewer Point P 1.26 — L293: Why are you mirroring the wake effects using the original Gaussian model? The expansion parameters of that wake model were derived from simulations of a turbine in ground effect, so implicitly it already accounts for that.

Reply: The original paper by Bastankhah and Porté-Agel (2020) finds that the ground does not significantly affect the turbine wake in their region of interest ($x/D < 20$), and their results are not impacted by adding mirrored turbines. For presumably similar reasons, Niayifar and Porté-Agel (2016) tuned the wake spreading parameters that we use without explicitly accounting for ground effects. In contrast, we find that mirroring the turbines drastically improves the wake model's results. To demonstrate this, figure RR2 shows the wake model predictions of the APM-level velocity field (i.e. horizontally filtered and height-averaged) and turbine power outputs, as compared to the LES results.
It's clear that mirroring the turbines is necessary for the wake model to provide good results. The conflicting findings with the original papers might be because we use a turbine type where $D/z_h \approx 1.66$, whereas Niayifar and Porté-Agel (2016) used turbines with $D/z_h \approx 1.14$. The rotor being relatively closer to the ground could explain the need for mirroring the turbines in our case.
We expanded the paragraph in the text to include this reasoning:
The turbines are mirrored to account for ground effects, as this significantly reduces the error on the predicted power output and farm-averaged velocity. This conflicts with the results from Bastankhah and Porté-Agel (2020), who found that mirroring the turbines had little effect on the WM's predictions. This difference with the original papers might be because we use a turbine type with a higher ratio of D/z_h, leading to the rotors being relatively closer to the ground.

Reviewer Point P 1.27 — L300-: This is a major issue. Without it we do not get a "fast" model after all so one might argue that this type of 2SM-APM coupling is actually not really fast nor really good. Could you comment on that and its relevance to practical applications?

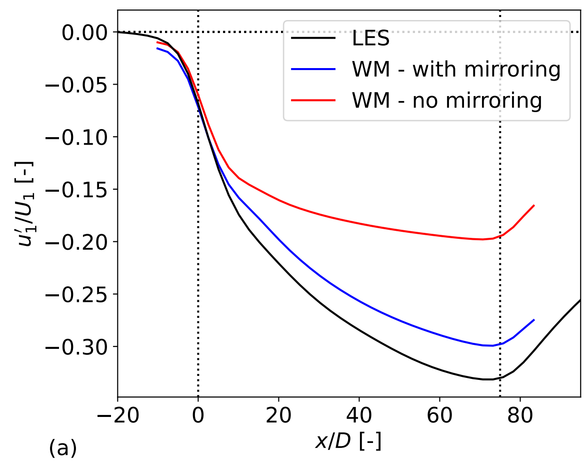

Figure RR2: Performance of the wake model for the case H500-$\Delta\theta$0-Γ0. (a): the meso-scale velocity deficit as found by LES (black) and the wake model, both with and without turbine mirroring (blue and red, respectively). The dotted lines denote the wind farm region. (b): the power output as found by LES (x-axis) and wake model (y-axis), using the same color code. The triangles and the circles denote the first two turbine rows and the remainder of the farm, respectively.

Reply: We agree that this is a potential issue, but deem it out of the scope of the current paper to solve this. We have modified the introductory paragraph of the section to reflect this (see Point P1.25). Furthermore, we have added a short paragraph discussing this issue at the end of the section:
The above discussion shows that in principle wake models can serve the role of a prior CFD simulation in determining C_T^* and C_P^*. The associated computational cost can be significantly reduced by future research, as we have made no attempts to minimize it. Furthermore, constructing lookup tables for C_T^* and C_P^* as a function of wind speed and wake model inputs would allow this upfront cost to be shared by many 2SM–APM cases when performing eg. multiyear AEP estimates. Finally, the coefficients could potentially also be inferred directly from operational data. Instead of investigating these approaches further, the following section develops a second version of the 2SM–APM which directly couples to a wake model, eliminating the need for an explicit calculation of C_T^* and C_P^*.

Section 5

Reviewer Point P 1.28 — L323: Please check again whether it is actually validation and not verification.

Reply: There is no parameter in the pressure-based 2SM–APM that is tuned using the dataset of Lanzilao and Meyers, making this a pure validation.

Reviewer Point P 1.29 — L341: It might be good to mention the "filtered, height-averaged" part already before in connection with U1 in L335 as all the different operators can be a little confusing (or forgotten by this stage).

Reply: We thank the reviewer for this suggestion. We have modified the introduction of the velocity matching equation to read

In contrast, when coupling a wake model to the APM, $\vec{U}_b$ has to ensure that horizontally filtering and height-averaging the turbine-scale velocity field results in the APM velocity, so that

$$\langle \widetilde{WM(\vec{U}_b)} \rangle_1 = \vec{U}_1 + \vec{u}'_1.$$

Reviewer Point P 1.30 — L346: Typo “new” ?

Reply: Thank you for noticing this oversight, we have replaced “now” with “then”.

Reviewer Point P 1.31 — L411: Typo “... simplify this solver, ...”?

Reply: Thank you for noticing this, we have modified it to “In order to simplify this solver, ..”

Reviewer Point P 1.32 — L435-: Maybe this will be discussed afterwards, but clearly this approach is also dropping a lot of the simplicity of the initial coupling and computational speed gains? Especially the latter needs to be discussed in relation to the promises made in the title.

Reply: We added two sentences discussing this potential disadvantage:
Its main disadvantage is that it requires an additional solver for the operator $\mathcal{A}^c_{u,P}$, corresponding to equations 38–39. However, since it can use a simple first-order marching scheme and a relatively small grid, this solver is straightforward to implement and cheap to run.

Reviewer Point P 1.33 — L443: Is it important where the inlet is placed?

Reply: Yes, it is. We added a short discussion on the reasoning behind the inlet placement when introducing the marching solver in section 5.1:
... a simple first-order explicit Euler marching scheme, with as inlet boundary conditions $\boldsymbol{u}'_{1,c}, \boldsymbol{u}'_{2,c}, \eta_{1,c}, \eta_{2,c} = 0$, which correspond to there being no pressure-induced change in background velocity at the inlet. For these conditions to be valid, the inlet should be placed far enough upstream of the farm so that $p' \approx 0$, which for realistic atmospheric conditions is typically less than 100 km (Smith, 2024). As a result, the operator $\mathcal{A}^c_{u,P}$ can use far smaller grid sizes than the periodic boundary conditions required of the conventional APM.

Reviewer Point P 1.34 — L445: The error is actually quite larger at times, one could argue that there is not much benefit in running with a correction for many of these cases. Could you provide some insights as to why?

Reply: As discussed in point P1.23, the variance in the error seems to mostly scale with the blockage losses, and not directly with the varying atmospheric state variables. We did not find any metrics that could be evaluated *a priori* that scale well with the eventual error of the model. Furthermore, any guidance on when the 2SM–APM models should be trusted above standalone wake models in practical applications should be based on validations with operational data. As such, we leave this to future work.

Reviewer Point P 1.35 — L457: Use “It is” instead of “It’s”

Reply: Unfortunately, we did not find any instance of “it’s” in this paragraph.

Reviewer Point P 1.36 — Figure 8: It clearly captures the trend but there is an offset.

Reply: We have modified the paragraph discussing this figure to mention this explicitly:
It's clear that the farm-average coupling performs decently in the role of a blockage-correction model, as it captures the overall trend despite a consistent bias. The error on this power difference is similar to the error on the farm powers shown in figure 8, which indicates that the farm-averaged PB model's bias is mostly due to an overprediction of blockage losses.

Conclusion

Reviewer Point P 1.37 — L457: Remove "interesting" as it is a very subjective statement.

Reply: Thank you for the suggestion, we have removed it from the new version of the manuscript.

Reviewer Point P 1.38 — L458: What are the core hypothesis?

Reply: We have modified the sentence to make this more explicit:
Using results from Lanzilao and Meyers (2024), its core hypothesis of scale separation between the turbine- and farm-scale flows was validated for steady-state idealized CNBL conditions.

Reviewer Point P 1.39 — L468-469: But how does it perform when CT and CP are unknown. Was this on purpose not mentioned as it did not perform as "well"? Well is also very subjective, try to quantify your claims and also state the max error. The same goes for the next paragraph.

Reply: The inclusion of the "for known C_T^* and C_p^*" was meant to emphasize the model's dependence on these values, and not to hide its performance in other situations. We have rephrased the paragraph in order to address the reviewer's concerns:
In a validation campaign using the LES results from Lanzilao and Meyers (2024), its RMSE on farm power was only 3% of the installed capacity when using LES-based values for C_T^* and C_p^*. Furthermore, we addressed the model's reliance on accurately knowing these coefficients by showing how they could in principle be calculated using a standard wake model.

Reviewer Point P 1.40 — Last paragraph: Some nice points here, but what about computational cost of the different model variants? Also if adoption in industry is wanted than also model simplicity is key and the second variant seem quite a bit more involved to implement. Future work should for sure also try to validate using datasets that are not using square wind farms and conditions different from those used here.

Reply: Thank you for raising these important considerations. We have added a short paragraph comparing the two model variants:
Both the 2SM–APM model and the farm-averaged pressure-based approach have a similar accuracy and computational cost. Which of them is preferable largely depends on the trade-off between efficiently calculating the farm coefficients C_T^* and C_p^* for the former, and solving additional equations for the latter. This investigation is left as future work.
Furthermore, as discussed in our introductory reply, we have added a final paragraph emphasizing the limitations of the 2SM hypothesis, indicating it as a prime area for further research:
As powerful as the tools offered by the 2SM approach are, it remains unclear how valid the underlying assumptions are outside of the idealized conditions they have been tested in. So far, they have only been

validated using LES of infinite farms (Kirby et al., 2022) and farms operating in steady-state, horizontally homogeneous background conditions with uniform turbine layouts (Kirby et al., 2025). Future work should investigate whether the farm coefficients remain constant under realistic and transient atmospheric conditions, and for farms with irregular turbine spacing and operating setpoints.

Reviewer 2

This well written and simple article presents two fast reduced-order models to predict wind farm power including the impacts of hydrostatic blockage effect. The derivation of both models is presented and the models are novel. This article can be published after the following questions and concerns are addressed and changes are made accordingly.

Reply: We thank the reviewer for the kind words, and the insightful comments, which we did our best to address. We believe this has substantially improved the paper.

Reviewer Point P 2.1 — The authors aim at improving the speed of reduced-order models. However, it was not established why it is desirable to have faster models? How fast is fast? Also, is this the only way to build a faster model or is it that the authors are improving their previous models? I think the authors should motivate the article in a clear way; I think they are improving their existing models and that should reflect in the title and the introduction. Otherwise, convince the reader that this is the only approach to producing faster models.

Reply: We agree that we did not make clear why we see the APM as a promising starting point for developing the faster models. Therefore, we have added the following paragraph in the introduction:
Our goal is to develop fast models for predicting wind-farm power output that account for blockage and gravity wave effects. To that end, the 2SM framework can be useful, as it implies that flow solvers simulating the farm-scale flow do not have to resolve the turbine-scale flow. This would allow for coarser grid resolutions and reduced computational costs. However, even when using coarser grids, conventional CFD tools are still multiple orders of magnitude slower than engineering wake models (Ouro et al., 2025). In contrast, the APM already uses a coarse resolution, while its main cost lies in capturing the local feedback effects between the turbine forces and the mesoscale flow. Therefore, we focus on combining the 2SM and APM approaches, leveraging insights from the former to speed up and simplify the latter.

Reviewer Point P 2.2 — Line 89-91: I understand that limiting the height from reaching into the free atmosphere is a valid reason, but the justification given for using a lower height is rather vague. If using a shorter domain is closer to turbine-scale flow, why not consider the wind farm layer (D)? Later in this paragraph, the authors talk about including the actuator disks in the lateral direction. This contradicts with the justification given earlier. It is desirable to elaborate a generic approach to choose the height of the control volume.

Reply: We did not mean for this paragraph to represent a tuning of H_f, or a general recommendation for the optimal averaging volume in all atmospheric conditions. Therefore, we re-wrote this paragraph with an emphasis on justifying the choice of H_f as corresponding to the APM layer distribution, and removed the comparison with earlier literature. While we do not set up a general method for choosing H_f, we now discuss the limitations of our approach in-depth, and indicate possibilities for future work.

We agree that the sensitivity of the dependence of C_p^* on H_f requires further research. We speculate that in the CNBL conditions studied here, the dependence is minimal as long as H_f is located above turbine tip height and below the capping inversion. However, without more high-resolution CFD data we can not generalize this to other atmospheric conditions. Since this analysis constitutes a deeper investigation of 2SM theory in itself, we feel this is outside the scope of the current paper. Therefore, we do not comment on it further.

Reviewer Point P 2.3 — Line 92-93: Kirby's observation sounds better representative of the wind farm aerodynamics in shallow boundary layers as recent LES studies have shown how wind farm yield is affected by low boundary layers. I think, the authors should calculate C*p based on the CNBL LES Uf and they might see variations. Basing C*p on a truly neutral simulation inherently excludes the stratification effects that are captured by the CNBL LES cases.

Reply: The point of this section was to demonstrate that C_p^* remains constant as long as the internal problem stays the same, even when the farm-scale flow goes from truly neutral to a CNBL flow. If lower CNBLs would change C_p^*, we would not find such low errors when reconstructing the farm power output in figure 1. To make this more explicit, we modify the conclusion of section 2 to read:
The root-mean square error (RMSE) on P_f is 1% of the farm installed capacity, which shows that a C_p^* value calculated from a neutral simulation remains applicable over a wide range of CNBL flows. We conclude that the 2SM assumption holds under the idealized conditions tested here.

Reviewer Point P 2.4 — Equation 9: It is interesting that the contribution to p' of the interfacial and internal waves is disentangled. Is it possible or is this an assumption? I would expect that the total displacement of the boundary layer is a coupled effect of the wind farm perturbation and all kinds of gravity waves.

Reply: The factor g' follows from treating the capping inversion as an infinitely thin layer with a finite jump in potential temperature across it (Gill, 1982). This assumption is made often in linear gravity wave models. Examples from literature include Klemp and Lilly (1975, *Dynamics of Wave-Induced Downslope Winds*), Durran (1990, *Mountain Waves and Downslope Winds*) Vosper (2004, *Inversion effects on mountain lee waves*), and many more. We recommend Teixeira (2014, *The physics of orographic gravity wave drag*) for an overview.
In the text, we have added a citation referencing Gill (1982), where readers can find a didactical derivation.

Reviewer Point P 2.5 — Line 172: Linked to first comment, is this the only way to build a fast model using 2SM? How about wind farm canopy models in a CFD setup that could exploit the essence of the 2SM to project a reasonably correct wind farm drag to the flow in an iterative manner? The model can run on a mesoscale grid and exclude the assumptions (linearization and gravity wave feedback) made by APM. When the authors talk about "fast" they should justify how is APM their chosen fastest yet reasonable accurate approach? Or is it that the authors want to exploit the availability of an APM model?

Reply: We took this helpful comment into account when formulating our response to the first comment. As such, we refer to our response there.

Reviewer Point P 2.6 — Equation 16: I am confused about this assumption and how it holds when the tuning parameter is zero. No perturbation above the 2SM essentially means the zero-mean displacement of the wind farm layer in the APM. Then equation 16 becomes: $H_1 * u'_1 = H_1 * u'_f = H_1 * (U_f - U_1)$. Thus, $H_1 = H_f$. So where is the difference in U_f and U_1 coming from if there is no perturbation/displacement? This assumption is the core of this derivation, and it is unclear to me how it holds. This requires further discussion and clarity.

Reply: We thank the reviewer for drawing our attention to the potential confusion caused by the formulation of this derivation. We have done our best to improve the text.
The 2SM volume height is always set to $H_f = H_1$, as this minimizes the discrepancy between the 2SM controle volume and the APM layer. This reasoning is given below equation 17.
The tuning parameter α is introduced to deal with the remaining discrepancy between the two volumes. This discrepancy consists of the region $H_1 < z < H_1 + \eta_1$. We propose that the velocity deficit in this region above the 2SM volume scales with the velocity perturbation within the 2SM volume through the parameter α. To make this more clear, we have added an equation explicitly introducing this scaling. The derivation now reads as follows:
... Here, we take the simplest possible approach by assuming that the average velocity perturbation in the part of the APM layer above the 2SM volume scales with the velocity deficit within it, so that

$$\frac{1}{\overline{\eta}_1}\int_{H_1}^{H_1+\overline{\eta}_1}(\overline{u}-U_0)dz \approx \frac{\alpha}{H_1}\int_0^{H_1}(\overline{u}-U_0)dz = \alpha u'_f,$$

where U_0 is the unperturbed velocity profile, α is a tuning parameter between 0 (no velocity deficit in the region $H_1 < z < H_1 + \overline{\eta}_1$) and 1 (perfectly mixed APM layer), and u'_f is the 2SM-farm-averaged velocity deficit:

$$\begin{aligned} u'_f &= U_f - \int_0^{H_1} U_0 dz \\ &= U_f - U_1. \end{aligned}$$

Taking the vertical integral up to $H_1 + \overline{\eta}_1$, we then find

$$(H_1 + \overline{\eta}_1)\overline{u}'_1 \approx (H_1 + \alpha\overline{\eta}_1)u'_f.$$

The difference between U_f and U_1 is due to the perturbation caused by the wind farm, since U_f is the farm-averaged velocity, while U_1 is the unperturbed velocity in the absence of the wind farm. To remind readers of this definition, we have added a line explicitly identifying U_1 to equation 19 (second equation in the above excerpt).

Reviewer Point P 2.7 — Line 212: Is tuning of α generic or specific to the CNBL LES cases in the dataset? Note that CNBL itself isn't fully representative of all atmospheric conditions and the idea of using a fast model should not rely on availability of higher fidelity solutions. Also, tracing a particle through a turbulent field resolved in LES to find the height of the plaint surface seems overly simple and a straightforward assumption that this depicts the authors' wishful interpretation of a distinct wind farm layer and an upper layer. This requires further explanation and justification to support the method used to acquire APM states.

Reply: Following the comments of another reviewer, this tuning now only uses an NBL LES case. We agree with the reviewer that this is not representative of all atmospheric conditions. Therefore, we have added the following sentence to the end of the paragraph:

Note that this basic approach does not capture the significant differences in velocity profiles between NBL, CNBL, and realistic flows, and is instead chosen for its simplicity.
We maintain that our methodology for setting up APM states is valid. It is important to consider that h_1 is defined as a surface of streamlines. To emphasize this, we have added its mathematical definition to section 3:
The pliant surface h_1 is defined as following the streamlines of the horizontally filtered flow, so that (Allaerts and Meyers, 2019)

$$\widetilde{w}(x, y, z_1) = \widetilde{\vec{u}}(x, y, z_1) \cdot \nabla_h h_1,$$

where w is the vertical velocity component, and ∇_h is the horizontal del operator.
Therefore, advecting a set of points is an appropriate way of reproducing h_1 from LES data. Furthermore, the turbulent flow field is time-averaged and horizontally filtered with a filter length of 1 km before the points are advected, so small-scale turbulent flow features will not affect the outcome. To emphasize this, we now explicitly mention that the flow field is time-averaged when describing the methodology.

Reviewer Point P 2.8 — Line 288-290: Presumably the wake model will give C_p^* and C_T^* for a neutral or at best a specific ABL stability case. Then using the 2SM-APM for arbitrary boundary conditions means that the wind farm thrust and power is independent of the atmospheric state to simulate. Since the flow conditions (i.e. U_{inflow}) depend on atmospheric state, the predicted power will inherit the state simulated by the wake model. Thus, a workflow coupling a wake model with the 2SM-APM, for consistent atmospheric conditions across the models, is desirable. Can the authors comment on this or propose a better strategy? There is a typo too: "These can then be used by".

Reply: The core idea of 2SM theory is that the farm coefficients are not dependent on the mesoscale atmospheric state, as shown in section 2. We have expanded the paragraph to re-state this:
In essence, the process is the same as before, i.e. the standalone wake model is used to estimate U_f, P_f, and F for a given farm, so that C_p^* and C_T^* can be calculated using equations 1 and 2. As shown in section 2, the resulting coefficients capture the turbine-scale flow characteristics, and are valid across a range of atmospheric conditions as long as the turbine setpoints and wake shapes remain the same. Once they are calculated, the 2SM–APM can use them to predict the farm power output in different mesoscale flow conditions.
If one wants to use a coupling between the APM and a wake model that does not rely on 2SM theory, various options have been developed (see Devesse et al., 2024b for an overview).

Reviewer Point P 2.9 — Section 5: At this point, the formulation is a bit hard to follow as the equations are repeatedly related to expressions in previous sections. I suggest the authors should give a flow chart with steps clearly showing the relations among the main assumptions made. This will also help the reader follow the structure of this paper. I also suggest the authors include the formulations of operators in the appendix as the derivations deeply rely on them.

Reply: We thank the reviewer for this helpful suggestion. We added both a conceptual sketch of the pressure-based coupling method and its derivation (figure 8), and a table with an overview of all the different operators (table 3).

Reviewer Point P 2.10 — Equations 27-29: On what basis was the contribution of velocity deficit dropped here? I think a small perturbation or displacement would cause little blockage. More importantly, how is perturbation related to velocity deficit? Please establish this assumption, preferably with some expressions. I also don't follow that the wake component is included back in 28 and later estimated in 29.

Reply: We agree that the original formulation was not clear. We have improved the text surrounding equations 29-31. Combined with the new flowchart visualizing the relations between the different variables, this should make the derivation easier to follow. The new text now reads:
In order to find a computationally cheaper way to estimate $\vec{U}_b$, we decompose the filtered, height-averaged wake model operator into a wake and a correction component, so that

$$\langle\widetilde{WM(\vec{U}_b)}\rangle_1 = \vec{U}_1 + \vec{u}'_{1,w} + \vec{u}'_{1,c}.$$

where the wake component $\vec{u}'_{1,w}$ is the filtered and height-averaged velocity deficit corresponding to the turbine wakes, as predicted by the wake model, and the correction component $\vec{u}'_{1,c}$ accounts for the remainder of the total velocity perturbation $\vec{u}'_1$. This splitting is shown conceptually on the right-hand side of figure 8. For sufficiently small flow perturbations, we can then assume that subtracting the wake component $\vec{u}'_{1,w}$ from the above equation corresponds to not applying the wake model operator WM, so that

$$\langle\widetilde{\vec{U}_b}\rangle_1 = \vec{U}_1 + \vec{u}'_{1,c}.$$

We can therefore interpret the correction component $\vec{u}'_{1,c}$ as representing the changes in the background velocity which we are trying to estimate.

Instead of deriving an expression for $\vec{u}'_{1,c}$ directly, we will derive it as the difference between the APM state $\vec{u}'_1$ and the wake component $\vec{u}'_{1,w}$. Using equation 28, we find

$$\vec{u}'_{1,c} = \vec{u}'_1 - \vec{u}'_{1,w}.$$

Therefore, the next step is to estimate $\vec{u}'_{1,w}$, . . .

Reviewer Point P 2.11 — Line 351: Clearly, averaging over velocity is dropped because it's a slower operation. Yet the operator suggested depends on the turbine force given by the wake model, as mentioned on line 394-395. Can the authors establish if this estimation is more accurate and there isn't a trade-off between speed and accuracy?

Reply: This assumption of using the derived operator instead of averaging comes down to the difference between the velocity-matching and pressure-based coupling techniques for the conventional APM. The corresponding tradeoff has been thoroughly analyzed by Devesse et al., Comparing methods for coupling wake models to an atmospheric perturbation model in WAYVE, 2026.

Minor issues

Reviewer Point P 2.12 — Statements on Line 37-39 are not entirely correct. The AFMs resolve the turbine scales on a coarser grid than required by ADMs. More importantly, these are CFD models which inherently couple the resolved scales. Thus, there is no separation of scales, consider rephrasing these lines and maybe there is no need to classify them as 2SM approaches; drawing the analogy is fine.

Reply: We agree that classifying these models as 2SM approaches goes further than what the original papers argue, and have changed the line from "same idea" to "similar ideas".

Reviewer Point P 2.13 — Line 75: I don't see why C_p^* cannot be calculated for stable or unstable flow simulations. Perhaps, the authors are confined by the CNBL dataset they have for validation. This doesn't mean the same approach is not applicable to SBLs and CBLs.

Reply: We agree with the reviewer. It was not our intention to imply that this was a limitation of the approach. We have changed the line from "neutral flow simulations" to "small-scale simulations".

Reviewer Point P 2.14 — Line 86: "farm averaged velocity", I assume?

Reply: Correct. This entire paragraph has been rewritten in response to other comments.

Reviewer Point P 2.15 — Figure 1: I don't understand why the authors did not use CNBL data to calculate C_p^*?

Reply: Our goal was to demonstrate that even when the flow goes from being truly neutral to a CNBL, C_p^* remains constant. As discussed in our reply to comment P2.3, we have added a sentence to emphasize this.

Reviewer Point P 2.16 — Line 96: Which other cases? You didn't refer to a specific case or cases in the previous sentences of this paragraph.

Reply: Thank you for noticing this oversight. We have modified the line from "the other cases" to "the 36 CNBL cases".

Reviewer Point P 2.17 — Line 111: Maybe it is useful to mention a few tools or a class of tools. I assume the authors mean to compare APM with engineering tools and not CFD tools.

Reply: Thank you for the suggestion. Our intention was to compare it against the relatively high vertical resolution of conventional CFD and NWP models. We have made this explicit by modifying the line from "conventional simulation tools" to "conventional CFD or mesoscale simulations".

Reviewer Point P 2.18 — Line 153: For completeness, include the expression for vertical wavenumber.

Reply: Thank you for the suggestion, we have done so.

Reviewer Point P 2.19 — Line 165: I guess APM assumes, through equ 9, a steady-state impact of all types of atmospheric gravity waves and not only internal waves.

Reply: We meant to say that the closure equation for the internal waves requires the use of a Fourier method, which is not true for the inversion waves. We have made this more explicit by rephrasing the sentence:
Since the pressure feedback of the internal gravity waves is mostly easily expressed in Fourier space, the system of equations 7-11 is solved with a Fourier–Galerkin method.

Reviewer Point P 2.20 — Line 279: Please also specify computational load in quantifiable units instead of saying standard laptop, e.g. core hours or flops etc.

Reply: Thank you for the suggestion. We have repeated our tests on a single CPU, and now list the results in the text as follows:

. . . each calculation taking roughly 1.3 CPU seconds on a Intel Xeon Platinum 8360Y 2.4GHz CPU. This is only approximately 35 times longer than a wake model evaluation for the same farm, . . .

Reviewer Point P 2.21 — Table 2: Some of the short forms are not defined at first use.

Reply: Thank you for noticing this. We have added a column specifying the sections of the paper that describe the different models, which should clarify any confusion.

Reviewer Point P 2.22 — Line 303: The authors aim to demonstrate that wake models can provide Uf and the remaining routine follows. Since wake models are bottom-up approaches and not ideal for wind farm interactions, I am wondering if there is an opportunity to establish a method to estimate Uf for an arbitrary wind farm? If so, why use wake models to estimate Uf?

Reply: We have added a larger discussion on how to extend the 2SM averaging operators to arbitrary farm layouts in section 2, as discussed in P2.2.

Reviewer Point P 2.23 — Line 345: "we can now assume"?

Reply: Thank you for noticing this. We have replaced "now" with "then".

Reviewer Point P 2.24 — Line 411: How do these assumptions affect the solutions produced by the solver?

Reply: Thank you for the question. Upon investigating, we found that including Coriolis effects and wind veer significantly improved the results, and we no longer neglect the corresponding terms. We still leave out the horizontal eddy viscosity in order to avoid having to discretize second-order derivatives in our marching scheme. The sentence now reads:
Finally, in order to simplify the marching solver, we neglect the horizontal eddy viscosity, so that there are no second-order derivatives in the equations.

Reviewer 3

This article has to do with the reduced-order modelling of wind farm flows. Specifically, it focuses on an effort to improve the efficiency of an atmospheric perturbation model (APM) developed by the authors and described in earlier publications. To achieve efficiency gains, the authors leverage insights from two-scale momentum theory. The authors offer two different approaches to speed up the APM and compare the results with a dataset of large-eddy simulations over a range of different inflow conditions.

I found the article to be well written, and for the most part I found the model descriptions to be clear. The validation indicates that the two approaches used to improve calculation efficiency hold promise. Some potentially important terms were left out of the APM to simplify this study. The terms will likely be restored in a future study; it will be interesting to learn how this affects the results.

There is a need for better engineering/reduced-order models for wind farm flows. This paper is a helpful contribution towards that end, and I recommend its publication. I have a number of minor comments. I hope at least some of them can lead to improvements in the paper.

Reply: We thank the reviewer for their kind words and helpful suggestions. We have incorporated them to the best of our ability, which has substantially improved the text. Below we give a detailed point-by-point response, with any numbers referring to equations, figures, or lines referring to the original and revised paper versions in the reviewer comments and replies, respectively.

Reviewer Point P 3.1 — Line 30 "one to five minutes on a standard laptop": Is that one to five minutes for each atmospheric condition simulated? Or is it one to five minutes to cover the full range of conditions, including wind speed and wind direction, that would be assessed during a typical energy yield analysis? This is probably clarified in the reference cited at the end of the sentence, but it is probably worth clarifying here.

Reply: We have modified the line so that it now explicitly states that this cost is one to five minutes for each simulated case.

Reviewer Point P 3.2 — Lines 68-72, where $C'_{p,k}$ and $C'_{T,k}$ are defined as disk-based turbine power and turbine thrust coefficients. While reading this paper, I was wondering where these values come from. I thought they might be simply inputs to the farm-scale model, but then line 250 indicates that $C'_{T,k}$ is "calculated using a standalone wake model". This makes me think that I may not fully understand what $C'_{T,k}$ is. Some clarification would help. I am not confident I understand how the values are determined for the different farm-scale models (wake model with no induction model, wake model with induction model, and truly neutral LES).

Reply: These coefficients are known from the turbine specifications, which typically include curves that link the power and thrust generated by the turbine to a reference velocity. This velocity is often the inflow velocity upstream of the turbine. However, in numerical studies using actuator disk methods, the disk-average is also commonly used. Since this section uses LES data, we followed that convention.
To clarify any confusion, we re-wrote the paragraph discussion the wake model's role in the APM underneath equation 15 to explicitly mention the turbine coefficient curves:
where S_k and $\vec{e}_k$ are the turbine disk velocities and orientations, respectively, and $C'_{T,k}$ can be evaluated using S_k from the turbine thrust coefficient curve. Because the APM's resolution is too low to resolve the turbine-scale flow, these quantities are calculated using a wake model. To capture the dependence of the turbine forces on the local mesoscale flow, past studies developed ways to calculate the background velocity used by the wake model based on the APM state (Allaerts and Meyers, 2019; Stipa et al., 2024a; Devesse et al., 2024a, b).

Reviewer Point P 3.3 — Figure 1: For the y-axis results, is P_0 the power of a standalone turbine simulated in truly neutral conditions? Might be worth clarifying in the paper.

Reply: Thank you for the suggestion. We have added the capping inversion strength and atmospheric lapse rate for the single-turbine simulations when introducing the LES dataset:
For each boundary layer height, the dataset also includes simulations of a wind-farm in truly neutral flow and of a stand-alone turbine with $\Delta\theta = 5K$ and $\Gamma = 4K/km$.

Reviewer Point P 3.4 — Line 143 "For the sake of simplicity, the latter two are not included in this study": This comes up a few times in the paper. Are the authors able to reasonably speculate on the impact of this simplification on the results? I cannot and am left wondering.

Reply: Thank you for the interesting question. Presumably, the effects will somewhat balance out, as dispersive stresses slow down the flow coming into the farm, while the turbulent entrainment will add streamwise momentum within the farm. Which of the two is dominant might depend on the farm layout

and the atmospheric conditions. For a detailed discussion of these effects, we refer to Devesse et al., 2024a, A meso–micro atmospheric perturbation model for wind farm blockage, section 2.4.
One major reason for leaving out these effects is that the current available parametrizations are limited, and are not able to account for realistic atmospheric conditions and general farm layouts. This is an area of active research for us, and we hope to be able to include these terms in the 2SM-APM in the coming years.

Reviewer Point P 3.5 — Line 207 "(no perturbation above the 2SM volume)": I originally misread this to mean that the flow is not perturbed upwards, but I am pretty sure that is not what you mean. This line instead refers to a perturbation in wind speed. Perhaps I should have understood the meaning right away, but you may want to clarify to avoid others having this experience.

Reply: Thank you for the suggestion. We have modified the sentence to read:
. . . α is a tuning parameter between 0 (no velocity deficit in the region $H_1 < z < H_1 + \overline{\eta}_1$) and 1 (perfectly mixed APM layer) . . .

Reviewer Point P 3.6 — Line 232 "we will assume that the distribution of the force in space, normalized by the total force $F/\overline{h}_1$, is independent of the APM state": Meyers has published papers indicating that this is not a good assumption (i.e. the papers show that the details of the stratification, which directly influence the APM state, can have a large impact on the distribution of the force across the wind farm). Any thoughts on the consequences of this simplification?

Reply: As discussed in that paragraph, the stratification effects will impact the spatial distribution of force. When investigating turbine-level metrics such as turbine-specific power output, capturing this will be relevant. However, the 2SM approach argues that when trying to predict the total farm power output, these local variations are no longer relevant.
To make this more clear, we expanded the preceding sentences to emphasize that this simplification is valid when predicting the farm-averaged velocity:
. . . While these local values impact individual turbine performance, they are not relevant to the 2SM approach, as it only requires the average velocity over the farm to calculate the farm power output. Since the goal of the 2SM–APM is to compute this farm-averaged velocity as cheaply as possible, including local variations is an unnecessary cost. . . .

Reviewer Point P 3.7 — Line 268 "We leave out the H150 cases. . . ": On line 86 and 87, the authors write, "This height (H_f) is above the turbine tip height but below the ABL height for all the cases in the dataset of Lanzilao and Meyers (2024)." When I first read this line, I was confused because I know the Lanzilao and Meyers paper included cases with ABL heights below H_f. It was not until I reached line 268 that the matter was clarified for me. This does not necessarily mean that a change needs to be made, but I wanted to let the authors know that at least this reader was confused by the sentence starting on line 86.

Reply: Thank you for noticing this. It seems that while some of the cases in the dataset indeed have a precursor ABL height that is below the turbine tips, vertical displacement of the capping inversion over the farm causes the capping inversion to always be above the turbines. We have made this more clear by adding the following sentence to the text:
. . . This height is above the turbine tip height but below the ABL height for all the cases in the dataset of Lanzilao and Meyers (2024). While some of the H150 cases have lower unperturbed ABL heights, strong

vertical displacements of the capping inversions over the wind farm ensure that even there the ABL height is above the 2SM control volume. . . .

Reviewer Point P 3.8 — Line 318 "... calculating these coefficients with a wake model places requirements on said wake model's performance. . . ": I think I missed what those requirements are. Was it the need to include a turbine induction model? Or perhaps something else related to calculating U_f?

Reply: That is indeed one of the requirements. To emphasize this, we added a concluding paragraph at the end of section 4.3.1:
The above discussion shows that in principle wake models can serve the role of a prior CFD simulation in determining C_T^* and C_p^*, but that this requires the wake model to accurately represent the flow in the induction and near-wake regions. . . .

Reviewer Point P 3.9 — Line 337 "This condition gives in an implicit equation for U_b. . . ": Should the word "in" be there?

Reply: Thank you for noticing this oversight, we have removed the "in" from the sentence.

Reviewer Point P 3.10 — Line 346 "For sufficiently small perturbations, we can new assume": I think "new" should be "now" or perhaps it should be simply dropped.

Reply: Thank you for noticing this. We have changed it to "then".

Reviewer Point P 3.11 — Equation 41: H should be H_1, right?

Reply: Thank you for noticing this oversight, we have corrected the mistake.

Reviewer Point P 3.12 — Lines 476-477 "Tuning the underlying wake models could address this, as both the wake spreading coefficients and the farm-averaged power and thrust coefficients can be optimized": How does this align with the earlier statement that the bias mostly has to with an overprediction of blockage loss (line 454)? This question mostly relates to the tuning of wake spreading coefficients. Separately, I'm not familiar with the tuning of the thrust coefficients used in wake models. How would that be done and how would it be justified? Perhaps this relates back to my question about $C'_{T,k}$? I am also not clear on what would be done to tune the "farm-averaged power" of the wake model. Is this just a consequence of tuning the wake spreading coefficients? To be clear, I have no problem with the idea of tuning engineering wake models, which are highly simplified representations of wind farm flow – I just want to better understand what the authors are suggesting here and why they are suggesting it.

Reply: Thank you for this insightful question. In our validation, the bias indeed comes mostly from the overprediction of blockage losses, which tuning wake model coefficients or C_T^* and C_p^* would not address. However, tuning these values could cause the wake model to have its own counteracting bias, which would compensate the error in the 2SM–APM, so that the combined model has a lower error. We have re-written the sentence to make this more explicit:
Both models achieve a lower variance in the error on their power predictions compared to a standalone wake model, at the cost of introducing a larger bias. Tuning the underlying wake models could compensate for this, as both the wake spreading coefficients and the farm power and thrust coefficients C_T^* and C_p^* can be modified. . . .

Finally, the farm coefficients C_T^* and C_p^* can be tuned to match operational data, without relying on underlying wake models. We have added a sentence in the concluding paragraph of section 4.3.1 to suggest such alternative approaches:
. . . *Finally, the coefficients could potentially also be inferred directly from operational data.* . . .