

The updated version of the manuscript is a clear improvement, showcasing a rigorous aeroelastic study of an IEA 15MW RWT. The authors compare their inhouse aeroelastic solver (“CFD-CSD”) with the OpenFAST modules BeamDyn and ElastoDyn. Comparing different setups (velocities only (“CFD-CSD:OV”) versus modelled torsion DOF (“CFD-CSD:T”); rotor only (“CFD-CSD:OV/T:RO”) versus modelled tower and nacelle (“CFD-CSD:OV/T:TN”)), they find the biggest discrepancies between the inhouse-code predictions and the ones from OpenFAST when the blade passes the tower. Regarding deformations, aerodynamic loads, and power coefficient, the torsional DOF appears to have a strong effect; differences between the structural-model formulations appear weaker in the presented comparisons.

The presented work lays a solid foundation for detailed aeroelastic studies, with the potential for an impactful contribution to modelling multi-megawatt class turbines. However, I have made the decision of rejecting the present manuscript for the following key reasons

1. **Incomplete examination of the apparent slower wake recovery due to tower presence (Fig. 4):** The authors compare the rotor-averaged streamwise velocity

along the centerline, normalized by the undisturbed velocity at hub height, between the present work and Santoni et al. (2017). They observe a trend opposite in sign to Santoni et al. (2017), namely a slower wake recovery when tower and nacelle are included.

First, this comparison is of limited interpretability because Santoni et al. (2017) investigated a 5 MW turbine under uniform inflow (i.e., different inflow and turbine class compared to the present case). While the authors acknowledge this limitation, it remains unclear what insight is gained by comparing the two datasets given the confounding differences.

Second, the manuscript does not provide a quantitative investigation to establish whether the observed slower recovery is a robust physical effect or a sensitivity to modeling/setup choices (e.g., tower/rotor scaling, inflow shear setup, actuator or boundary-condition choices). The authors offer plausible hypotheses and acknowledge the lack of validation, but even though make physical conclusions based on the observed results (“indicating that wake recovery is slightly hindered by the presence of the tower” in line 330 of the manuscript).

2. **Mixing fidelities for validation:** the present work compares CFD-CSD/OV with OpenFAST:ElastoDyn and CFD-CSD/T with OpenFAST:BeamDyn. This is fair insofar as the same aeroelastic DOFs are compared. However, the compared numerical models are the sum of combining different submodels that hold different fidelities, limiting the interpretability of the presented comparison:

		Structural model	
		Linear	Nonlinear
Flow model	BEM	Elastodyn	Beamdyn
	FVM	CFD-CSD (OV, TN)	

The authors recognize that the cross-solver comparisons are not directly attributable to a single modeling component; however, the manuscript nevertheless bases its central conclusions on these confounded comparisons. Without controlled sensitivity studies or a comparable reference, the resulting insights remain limited and largely descriptive.

3. **Unclear storyline and vague overall goal:** the presented work is a thorough and valid starting point for good aeroelastic benchmarks, however the reviewers struggled to understand the intent of the work. Is it to validate the setup of a 15 MW turbine for your inhouse solver? Or is it to examine the effects of modelling the presence of towers in aeroelastic CFD? Or is it to test the limits of BEM for 15MW turbine aeroelasticity? As a result of the many aspects the manuscript covers, the conclusion is too long (about 50 lines) and unclear. The 36-page length of the paper makes it exhausting to read.

I'm very sorry for my decision and acknowledge the rigor and diligence of the authors. I appreciate the willingness the authors showed for incorporating the reviewers' comments. However, the manuscript appears like the concatenation of two good papers that are not finished yet. I suggest rebuilding the paper from the ground up and split into two possible publications: (i) one on validating the effects of the tower and nacelle e.g., on a stiff rotor (isolating effects) and compare the observed results with another (established) LES solver; this would increase the credibility of related findings and this could be followed by (ii) one publication on comparing these simulations with the same setup but added structural response. Then you would only vary one submodel instead of varying fidelity of both, flow and structural model. In case you want to compare with BEM, try to use the same structural solver for both, FVM-LES and BEM.