We are grateful to the two reviewers for their constructive and thoughtful comments (Stéphane Vannitsem and the anonymous reviewer) on our manuscript "New insights into decadal climate variability in the North Atlantic revealed by data-driven dynamical models". We also thank Andrey Gritsun their support as editor.

Our responses to reviewer comments are listed below. The requests are in black text and our replies in blue font. Corrections to text in the manuscript are indicated in *italic font*.

Anonymous Reviewer #1

I thoroughly enjoyed reading this paper. The authors use quadratic regression in the spirit of Kravtsov et al. (2005) and a slew of model selection criteria to construct a three-variable representation of decadal coupled dynamics over the North Atlantic. They end up with the model that produces realistic continuous spectra and is able to forecast out-of-sample data at multi-year lead times. The algebraic structure of the model provides essential clues as to the dynamics of the observed climate variability over the North Atlantic region and guides the targeted data analysis to support the emerging hypotheses.

> We are grateful that the manuscript is regarded as an interesting read and offers promising insights.

A couple of personal first-reading impressions on presentation and content:

(1) section 3.1.2 and Fig. 3 are a bit in the way of paper's presentation flow, IMO (with the only message in ll. 308-310 about the key role of precipitation). I'd remove this section or put it in the appendix.

> Thank you for this comment. We do agree, although a novel part of all this analysis is how precipitation acts as a mediating variable, fully coupled and feeding back onto both AMV and NAO. The precipitation terms found by the equation discovery is robust and present in all potential model configurations as shown in Fig. 3. We accept that this does not come across well, as it has been reduced to the last lines 308-310.
>
> We amended this section by:
> 1) Including introductory sentences explaining why we do the analysis shown in Fig. 3
> 2) Explaining the importance of Fig. 3 and its features which were not discussed, including the sudden occurrence of precipitation equations.
> 3) Removing text which was not focused enough on the main conclusions of the section.

The new text is included below:

"Examining this ensemble as a whole provides further insights into the robustness of each term in the three models, which, in turn, provides information on the dominant interactions between variables. This analysis helps to distinguish terms related to noise from those representing true dynamics in the more complex low-order models."

"The persistence of cross-terms involving precipitation across all models suggests that precipitation holds important dynamical information for AMV-NAO coupling. Moreover, it may act as a state-dependent coupling parameter, modifying the effective strength and sign of interactions between the AMV and NAO. Additionally, there is a delayed appearance of the dP/dt equations, which occurs when models have five terms or more. In the most sparse models, precipitation serves as a diagnostic variable, where its current state provides useful information about the coupled ocean-atmosphere system. However, its own evolution may be too noisy or weak to resolve compared to the leading-order AMV-NAO interactions. Only when the model includes sufficient coupling terms can precipitation become dynamically active, allowing the lower-frequency deterministic component of precipitation to be separated from the unresolved noise. The emergence of the dP/dt equations suggests an important role for active precipitation feedbacks on the atmosphere–ocean system, involving latent heating, freshwater fluxes, and atmospheric stability."

(2) section 3.2 and Fig. 5 suggest skill, but would benefit from including a few concrete numbers in text - and, maybe, comparison with a benchmark forecast (say, persistence, or three-variable LIM model).

Thank you for this comment. We agree, and we have now included exact numbers of correlations in the text. We have also included a comparison with a persistence benchmark.

The new text is included below:

"For the training data, the correlations at LYR[1-1] for the AMV, NAO, and precipitation are 0.86, 0.54, 0.11, respectively. For the test data, the correlations are 0.93, 0.52, 0.7, respectively. We compare the skill of the model for LYR[1-1] to a baseline forecast, which we choose here to be persistence. To calculate the persistence correlation, we correlate observations with the initial forecast states for LYR[1-1]. We find no significant differences between model forecast skill and persistence correlation for this lead year range across all variables and datasets. This suggests that, for early lead times, the model does not provide much additional predictive information beyond the memory of the initial state. However, the low-order model exhibits dominant variability on timescales of approximately 20 years and longer (see Fig. 4 d,e,f). As such, the skill diagnosed here, which is limited to lead times of up to 10 years, targets timescales shorter

> *than the intrinsic timescales of the model. Therefore, the finding that forecast skill is comparable to persistence in this range does not necessarily imply the absence of predictive skill at longer lead times and averaging windows. We also note that correctly initialising an ensemble of forecasts greatly improves the model's predictive performance at longer lead times, while also beating the persistence benchmark (not shown here)."*

(3) Section 3.3.1 Algebraic structure of the model and the hypothesis of "damped oscillatory mode forced by the atmosphere", connections with linear inverse models, importance of the nonlinear feedbacks involving precipitation, etc. These are the issues the papers prompts (me) to think about. The key difference between the present approach and the previous data-driven approaches, I think, is the focus on a deterministic nonlinear model, rather than stochastic linear model. In my experience, the residual tendencies unexplained by the dynamical operators, say, on the right-hand side of (8–10) are typically large (maybe boxcar running-mean annual smoothing reduces them by a lot, maybe not). The apparent success of the deterministic model indicates that these residual dynamics are unimportant for the phenomenon of interest; at the same time, the "stochastic driving" is internally generated within the model. It is surprising to me that such low-order dynamics generate realistic spectra.

In any case, having the actual model (8-10) provides one with a tool to study the sources of the decadal variability and predictability in this model (by looking, for example, at the importance of various terms in the "dynamical" experiments with, say, suppressed, or one-way coupling between equations, just like it is done with more complex dynamical climate models, or by parameterizing further the model dynamics with linear operators and stochastic driving etc. etc.)

Once again, a very interesting paper!

> We thank you for this assessment.
>
> We have included new text which addresses all these points made above:
>
> *"The advance presented here lies in identifying low-order deterministic models whose nonlinear structure can reproduce key aspects of decadal variability without the need for explicitly prescribed stochastic forcing. While residual dynamics are not explicitly modelled as stochastic noise in our framework, we acknowledge that such processes are present in the real system. Some of these unresolved processes may be implicitly captured through the robust nonlinear terms identified by the low-order models. As mentioned in the discussion above, the observed coupled 20-year mode of variability is partly excited by stochastic NAO forcing and maintained by slow ocean dynamics. The "stochastic driving" of this mode of variability in the low-order models is internally generated through deterministic nonlinear interactions. A comparison of our low-order models to a Linear Inverse Model (LIM) would provide insight into whether the identified nonlinearities offer a better representation of the dynamics, allowing for improved predictive*

> *capability. We leave this as an important direction for future work. Furthermore, having these low-order dynamical models provides a tool to study the sources of decadal variability and predictability. For example, various experiments could be performed, such as suppressing terms to probe one-way couplings between equations and further parameterising the model dynamics with stochastic forcing. Further research is currently underway to study the state-dependent predictability of the low-order models and to diagnose the models' underlying physical processes by comparing them to coupled climate model simulations."*

Stéphane Vannitsem Reviewer #2

The manuscript presents a very interesting work on the development of a low-order model to describe the coupled dynamics at monthly time scale between the Atlantic Multidecadal Variability (AMV), the North Atlantic Oscillation and the precipitation over the North Atlantic, the latter being used as a proxy for latent heat exchange and fresh water fluxes between the ocean and the atmosphere. The authors apply a discovery method to isolate the terms appearing on the right-hand side of the dynamical equations, up to quadratic terms. They found several interesting models that do present relevant variability for the North Atlantic. The most interesting model composed of 26 terms displays an erratic behaviour reminiscent of chaotic dynamics. A strong low-frequency variability is identified on time scales between 20 to 35 years. The latter range is reminiscent of the coupled ocean-atmosphere variability identified over this region using a reanalysis dataset (Vannitsem and Ghil, 2017). The manuscript is well written, easy to follow and provide important information on the coupled dynamics between the ocean and the atmosphere over the Atlantic region.

> We are grateful that the manuscript is regarded as an interesting read, easy to follow and provides important information on coupled ocean-atmosphere dynamics in the North Atlantic.
>
> We are pleased to cite (Vannitsem and Ghil, 2017) and have now included this in the manuscript:
>
> *"We note that the variability of the coupled ocean–atmopshere of the North Atlantic on a similar time scale was also identified using reanalysis data in Vannitsem and Ghil (2017)"*

A first aspect that is probably useful to mention/address is the role of processes/time scales that are not represented in the exploration, often described by noise. An example is of course the Linear Inverse Modelling (LIM) for which a stochastic noise is used to represent processes that are not described by the deterministic part of the model. In your forecasting setup, it would be nice to have a comparison with the LIM approach to see whether the identified nonlinearities provide a better representation of the dynamics, allowing for a better skill. This question is probably beyond the scope of the current manuscript but is at least worth to mention.

Thank you for this suggestion.

We have now included these new comments made by the reviewer above in the new text below:

"The advance presented here lies in identifying low-order deterministic models whose nonlinear structure can reproduce key aspects of decadal variability without the need for explicitly prescribed stochastic forcing. While residual dynamics are not explicitly modelled as stochastic noise in our framework, we acknowledge that such processes are present in the real system. Some of these unresolved processes may be implicitly captured through the robust nonlinear terms identified by the low-order models. As mentioned in the discussion above, the observed coupled 20-year mode of variability is partly excited by stochastic NAO forcing and maintained by slow ocean dynamics. The "stochastic driving" of this mode of variability in the low-order models is internally generated through deterministic nonlinear interactions. A comparison of our low-order models to a Linear Inverse Model (LIM) would provide insight into whether the identified nonlinearities offer a better representation of the dynamics, allowing for improved predictive capability. We leave this as an important direction for future work."

A second related aspect is to clarify what is the proportion of variability indeed represented by the dynamical model as compared with the observations. That could allow for making an estimation of the noise that should be added to the description to effectively represent the full variability. This characterization could also allow helping in selecting the best models.

We thank you for this point. We have now calculated explained variances of the model for all variables in the train and test datasets for LYR[1-1]. We have also noted its relevance as a potential model selection criterion.

The new text is included below:

"We also calculate the explained variance of the model in LYR[1-1] for both training and test. For the training data, the explained variances at LYR[1-1] for the AMV, NAO and precipitation are 0.65, 0.10, and -6.62 respectively. For the test data, the explained variances are 0.81, -0.48, and 0.40, respectively. We note that AMV variability is well represented in both training and test data in this lead year range. However, the low-order model struggles to capture the variability of the NAO and precipitation. Similarly to the skill score mentioned above, this difference is likely due in part to the enhanced persistence of the AMV initial state. We note that the explained variability of the low-order models could allow for an estimation of the noise that should be added to the models to effectively represent the full variability. In addition, explained variance in a given lead year range could also support the model selection process."

Finally, precipitation is used as an observable of your system. This is indeed an interesting way to see the coupling between the ocean and the atmosphere, but this choice probably would not be my first choice of variables. One could also use both sensible and latent heat fluxes and also the freshwater flux that would allow for discriminating the role of the different variables. It would be nice to have a comment on this and the ability of the approach to handle additional variables and additional nonlinearities.

> Thank you for this point. We agree that precipitation is not the only possible choice of variable to represent air–sea coupling, and that using other quantities (e.g. latent and sensible heat fluxes, freshwater flux) could help discriminate the underlying physical processes.
>
> We have now included all the points made by the reviewer in the new text below:
>
> *"We make a final remark on the use of precipitation to represent ocean–atmosphere coupling, where we have used it as an integrated proxy for both atmospheric latent heating and freshwater fluxes. Moreover, precipitation is a useful variable from a climate impacts perspective, particularly in relation to European precipitation variability associated with the NAO and AMV. To further develop this work, one could also incorporate both sensible and latent heat fluxes, as well as freshwater fluxes, which would allow for a clearer discrimination of the role of these different processes. The SINDy algorithm can easily accommodate additional variables and higher-order nonlinearities. However, increasing the dimensionality of the models in our case introduces problems, namely the risk of overfitting on such limited observational data. Accurately disambiguating the processes linking many co-varying parameters and identifying both stable and robust models would be more challenging when working with short datasets."*