ORIGINAL ARTICLE

Reduced-precision parametrization: lessons from an intermediate-complexity atmospheric model

Leo Saffin[1] | **Sam Hatfield**[2] | **Peter Düben**[2] | **Tim Palmer**[3]

[1]School of Earth and Environment, University of Leeds, Leeds, United Kingdom

[2]European Centre for Medium Range Weather Forecasts, Reading, United Kingdom

[3]Atmospheric, Oceanic and Planetary Physics, University of Oxford, Oxford, United Kingdom

Correspondence
Leo Saffin, School of Earth and Environment, University of Leeds, Leeds, United Kingdom, LS2 9JT
Email: l.saffin@leeds.ac.uk

Funding information
ITHACA, European Research Council, Grant Number: 741112; ESiWACE2, Horizon 2020 European Research Council, Grant Number: 823988

Reducing numerical precision can save computational costs which can then be reinvested for more useful purposes. This study considers the effects of reducing precision in the parametrizations of an intermediate complexity atmospheric model (SPEEDY).

We find that the difference between double precision and reduced precision parametrization tendencies is proportional to the expected machine rounding error if individual timesteps are considered. However, if reduced precision is used in simulations that are compared to double precision simulations, a range of precision is found where differences are approximately the same for all simulations. Here, rounding errors are small enough to not directly perturb the model dynamics but can perturb conditional statements in the parametrizations (such as convection active/inactive) leading to a similar error growth for all runs. For lower precision, simulations are perturbed significantly.

Precision cannot be constrained without some quantification of the uncertainty. The inherent uncertainty in numerical weather and climate models is often explicitly considered in simulations by stochastic schemes that will randomly perturb the parametrizations. A commonly used scheme is stochastic perturbation of parametrization tendencies (SPPT). A strong test on whether a precision is acceptable is whether a low-precision ensemble produces the same proba-

bility distribution as a double-precision ensemble where the only difference between ensemble members is the model uncertainty (i.e. the random seed in SPPT). Tests with SPEEDY suggest a precision as low as 3.5 decimal places (equivalent to half precision) could be acceptable which is surprisingly close to the lowest precision that produces similar error growth in the experiments without SPPT mentioned above. Minor changes to model code to express variables as anomalies rather than absolute values reduce rounding errors and low-precision biases allowing even lower precision to be used.

These results provide a pathway for implementing reduced-precision parametrizations in more complex weather and climate models.

KEYWORDS
reduced precision; parametrization; stochastic physics; model error

1 | INTRODUCTION

Increases in computational power have been, and will continue to be, a key contribution to the continuous improvement of numerical weather and climate models (Bauer et al., 2015). With additional computational power, resources can be invested to improve our representation of the Earth system. Resources can be invested in increased spatial resolution, shorter timestepping or better sampling of probabilities with more ensemble members. Resources can also be invested in extra model complexity such as improved representations of processes or inclusion of more processes not currently represented. The challenge is to choose the allocation that gives the best model for a specific purpose given limited computational resources.

We may be investing too much computational power in all areas of models by performing calculations at high numerical precision. Reducing the precision of computations will save computational power at the expense of introducing truncation errors. By reinvesting the computational power saved by using low-precision computing model improvements that outweigh the errors due a precision reduction can be introduced. For example, low-precision high-resolution models can outperform low-resolution high-precision models (Düben and Palmer, 2014).

In many areas of models we may be able to use low precision without any penalty to performance because of inherent model uncertainties (Palmer, 2014). A numerical model of the atmosphere is essentially a dynamical core, which approximates the governing equations of fluid dynamics, coupled to parametrizations, which represent physical processes not represented or poorly resolved by the dynamical core. Parametrizations are an essential component of models; however, due to the simplifications made in parametrizations they are also a large source of uncertainties. Because these uncertainties are at the grid scale and errors at small scales grow and saturate faster than at large scales (Lorenz, 1969), there is a scale dependence to how precisely we need to represent the atmosphere (Palmer, 2014). This scale dependence can be of practical use in spectral-transform dynamical cores where less precision is required to represent spectral modes closer to the truncation scale (Thornes et al., 2018; Chantry et al., 2018).

In this study we focus directly on reduced precision parametrizations and how well rounding errors can be masked by the inherent model uncertainty. Significant improvements have been made by representing uncertainty explicitly through stochastic schemes (Palmer, 2012; Leutbecher et al., 2017). Since stochastic schemes are tuned to give models the right spread and can be considered to be an explicit representation of model uncertainty, they can be used to determine the acceptable levels of rounding errors. Düben and Dolaptchiev (2015) showed that stochastic forcing could be used as an upper limit to rounding errors using the 1D Burgers equation with stochastic sub-grid-scale forcing. Here, we determine the acceptable precision levels of parametrizations when directly affected by SPPT (stochastic perturbation of parametrization tendencies), a commonly used stochastic scheme in state-of-the-art weather and climate models. Whether the noise generated by well-adjusted rounding errors can actually be used to quantify uncertainty within simulations and to develop a new stochastic parametrisation scheme, as discussed in Düben and Dolaptchiev (2015), is beyond the scope of this paper and will not be investigated here.

Reduced-precision computing is already proving of practical importance. Until recently, most models have used double precision (64-bit, see section 2.2 for definitions) by default. Forecasting centres are now experimenting with single precision (32-bit). For example, Váňa et al. (2016) found that reducing the precision of the majority of ECMWF's Integrated Forecasting System (IFS) from double precision to single precision reduced the runtime by 40% with no noticeable change in forecast skill. The single-precision IFS will allow higher resolution to be used for forecasts and is planned to be operational by 2021 (Düben et al., 2018). Similar results have been found for a single-precision version of MeteoSwiss' COSMO which is used operationally (Rüdisühli et al., 2013).

While single precision appears to be sufficient for use in most model components, a further reduction in precision will be more difficult. Half precision (16-bit) is, for example, likely insufficient for use in many model components. It is therefore useful to consider a stronger precision reduction for individual model components instead of a global reduction in precision. Using less precision for individual model components has been shown to be applicable for single precision, such as a single-precision dynamical core (Nakano et al., 2018) or a single-precision microphysics parametrization in a double-precision model (Gilham, 2018).

In this study, we focus on reducing the precision of individual parametrizations in SPEEDY. In section 2 SPEEDY (2.1) and the changes applied to introduce reduced precision (2.2) and stochastic physics (2.3) are described. Results are presented in section 3. In section 3.1, the direct impact of rounding errors on the initial parametrization tendencies are quantified. In section 3.2, the impacts of rounding errors and SPPT on error growth are quantified using forecast-type experiments with varying precision in the parametrizations. In section 3.3 some code improvements to reach lower precision are presented. Results are discussed in section 4.

2 | METHODS

2.1 | SPEEDY

SPEEDY (Simplified Parametrizations primitivE Equation DYnamics) is an intermediate complexity global atmospheric model (Molteni, 2003). SPEEDY solves the hydrostatic primitive equations using a spectral transform dynamical core and has a suite of simplified parametrizations. SPEEDY is specifically designed to work at low horizontal and vertical resolution. Here, a spectral resolution of T30 with a 96x48 Gaussian grid, eight sigma levels and a 40-minute timestep are used.

Figure 1 gives an outline of a single timestep in SPEEDY. Each timestep the spectral prognostic variables are transformed to the Gaussian grid where the parametrization tendencies and nonlinear terms in the dynamics are calculated. These tendencies are then transformed back to spectral space and added to the dynamics and diffusion

tendencies computed in spectral space. The spectral prognostic variables are than advanced forward in time using a leapfrog scheme with a Robert-Asselin-Williams filter (Amezcua et al., 2010).

SPEEDY has parametrizations for convection, large-scale precipitation, long-wave and short-wave radiation, surface fluxes, and vertical diffusion. The parametrizations in SPEEDY are explicitly designed to work with the bulk approximations associated with the low vertical resolution (Molteni, 2003). The eight sigma levels are taken as representative of one boundary-layer level, five troposphere levels and two stratosphere levels.

The boundary conditions for SPEEDY are generated from the ERA-Interim dataset (Dee et al., 2011) from 1981-2010. Monthly means are calculated for sea ice, snow depth, land-surface temperature and soil moisture. The boundary values are then generated each day by linear interpolation between these monthly means. For sea-surface temperature, an anomaly field for each month in the dataset is included as well as the climatological monthly means. A constant climatology is taken for surface albedo and vegetation. SPEEDY also uses realistic orography and a land-sea mask.

2.2 | Emulating reduced precision

A floating point number is made up of a sign bit s, a Significand S (also known as the mantissa) and an Exponent E. The Significand represents a number between 1 and 2 and the Exponent represents an unsigned integer such that the floating-point number (F) is given as,

$$F = (-1)^s S \times 2^{E-B}, \tag{1}$$

where B is a bias added to centre the exponent to 0. The Significand and Exponent are constructed as

$$S = 1 + \sum_{i=1}^{N} s_i 2^{-i}, \tag{2}$$

$$E = \sum_{j=1}^{M} e_j 2^{j-1}, \tag{3}$$

with s_i and e_j being the individual bits, and N and M being the number of bits of the Significand and the Exponent.

Three commonly used formats defined by the IEEE (Institute of Electrical and Electronics Engineers) floating-point standards are double precision (N=52, M=11, B=1023), single precision (N=23, M=8, B=127) and half precision (N=10, M=5, B=15). Although current conventional processors only implement double and single precision arithmetic, future processors will implement half precision. Half-precision computing is also available on current GPUs (Graphical Processing Units) which have recently been used in numerical weather and climate models (e.g. Leutwyler et al. (2016)). Alternatively, FPGAs (Field Programmable Gate Arrays) allow for flexible precision arithmetic; however, due to the current complexity of porting to FPGAs, actual speed-ups for reduced precision have only been demonstrated for toy models (e.g. Düben et al. (2015); Jeffress et al. (2017)). The focus of this study is not on the hardware. Here, we use a reduced-precision emulator to look at the full range of precision and identify issues that will arise when reduced-precision hardware is adopted in the future.

A reduced-precision version of SPEEDY, adapted from Hatfield et al. (2018), is used here. Reduced precision is emulated using the Fortran module from Dawson and Düben (2017) with modifications from Hatfield et al. (2018) to include the Fortran complex type. The reduced precision emulator works by representing numbers as double precision and then truncating the number of bits in the Significand (N) of this double-precision number after every operation following floating-point standards ("round nearest ties to even"). The number N is controlled by an "sbits" parameter.

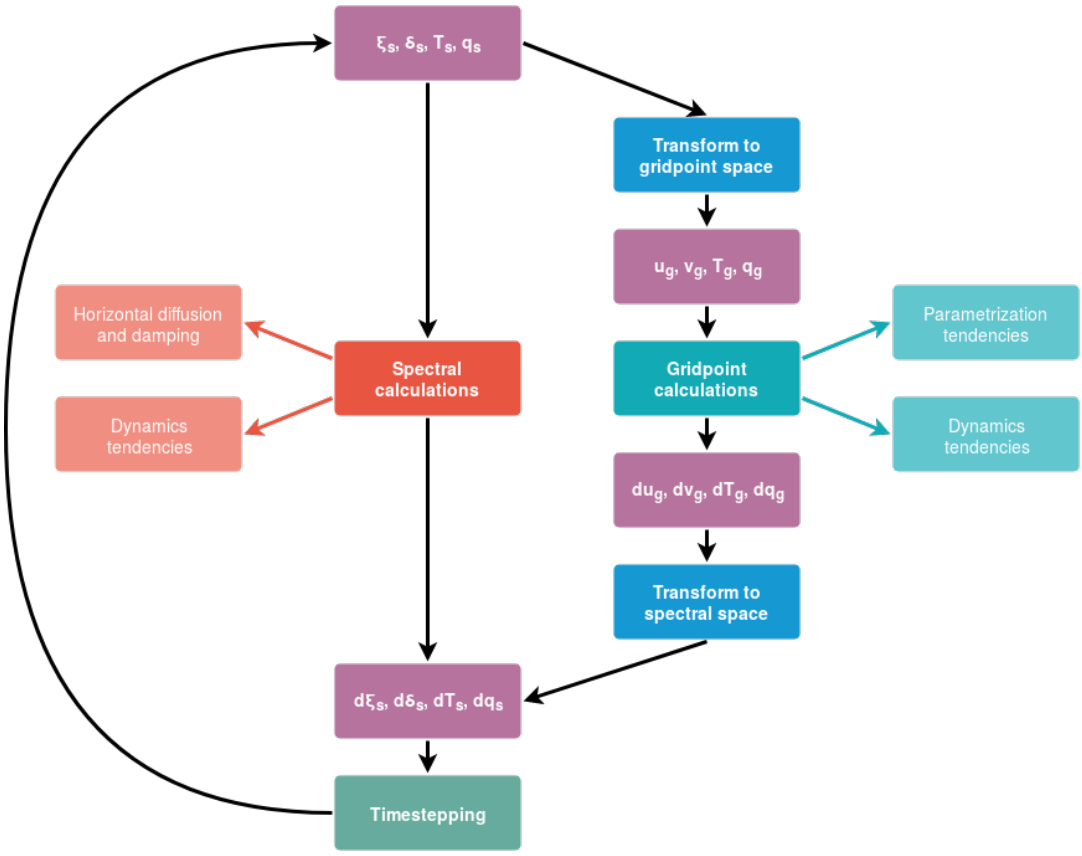

FIGURE 1 A diagram of a single timestep of SPEEDY. The model solution is based on a spectral-transform method. Prognostic variables vorticity (ξ_s), divergence (δ_s), temperature (T_s) and specific humidity (q_s) are stored as spherical harmonics, denoted by a subscript s, down to a truncation limit (T30). Linear terms in the dynamics, diffusion and timestepping are easily calculated in spectral space. Nonlinear terms and physical parametrizations must be calculated on a grid which requires the prognostic variables to be transformed to gridpoint space, denoted by subscript g, giving the gridpoint fields of zonal wind (u_g), meridional wind (v_g), temperature (T_g) and specific humidity (q_g). The derived tendencies for these variables in gridpoint space (du_g, dv_g, dT_g, dg_g) are then transformed back to spectral space and added to the tendencies computed in spectral space to give the total tendency for each prognostic variable in spectral space ($d\xi_g$, $d\delta_g$, dT_g, dg_g). These tendencies are then used to step the model forward in time. The focus of this study is on reducing precision in the calculation of parametrization tendencies in gridpoint space.

The reduced-precision emulator also has the option to emulate a half-precision exponent. SPEEDY works when using a single-precision exponent but crashes when using a half-precision exponent. Although there may be changes to the code that could be implemented to allow numbers to fit within the half-precision dynamic range such as rescaling variables, they are not attempted here. The focus of this study is on reducing precision in the Significand.

The reduced-precision emulator allows us to make meaningful studies of how a reduction in precision would impact on model results. However, by doing this, the emulator is reducing rather than increasing the speed of model simulations and it is therefore beyond the scope of this paper to measure performance improvements due to a precision reduction on real hardware.

There is a choice to be made as to how fine grained we test the reduced-precision because having too many different precisions could introduce overheads that outweighs the benefits of reduced precision. In this study we have followed the approach in the Met Office Unified Model where the microphysics parametrization can be run at single precision while the rest of the model uses double precision (Gilham, 2018). Gilham (2018) described this as a "bubble" of single precision where all data within the bubble is at single precision and all data passing in and out of the bubble is cast to the correct precision. Gilham (2018) showed good speedups within the microphysics parametrization and for the full model with single-precision microphysics proving that this "reduced-precision bubble" approach can be of practical use for individual parametrizations. In this study we implement a reduced-precision bubble for each parametrization within SPEEDY. To achieve this, constants are copied and stored at the reduced precision for each parametrization and all variables passed to a parametrization are copied and truncated before any calculations are performed.

2.3 | Stochastic physics

Stochastic perturbation of parametrization tendencies (SPPT) has been added to SPEEDY as a representation of model uncertainty. The scheme has been implemented following the SPPT scheme in ECMWF's IFS described by Palmer et al. (2009) and is summarised here. Each timestep, the total of the parametrized tendencies is perturbed by a randomly generated field giving the net tendency,

$$P = (1 + \mu r) \sum P_i, \tag{4}$$

where P_i is the tendency from an individual parametrization i, r is a 2d random field and μ is a vertical tapering function than can vary between zero and one. The total random field r is taken as the sum of three random gaussian patterns of different scales in gridpoint space. The parameters for the length and time scales are taken from table 1 in Leutbecher et al. (2017). The value of r is then limited to the range -1 to 1. Each random pattern is generated in spectral space to have a spatial autocorrelation in gridpoint space equivalent to a Gaussian on a sphere. At each timestep the spectral random field is advanced by first-order autoregression with a fixed decorrelation time scale. In the IFS the tapering function is smoothly reduced to zero in the boundary layer to avoid numerical instability and smoothly reduced to zero in the stratosphere to avoid strong perturbations of the radiative tendencies. In this study, the tapering is effectively switched off by setting μ to one at all levels. Setting μ to one at all levels was chosen for simplicity and to allow for an artificial model uncertainty at all levels. We would expect tapering to affect our results as it would effectively decrease this representation of model uncertainty. Since the parametrizations in SPEEDY are highly simplified and using SPPT in SPEEDY is an artificial representation of model uncertainty, we are not concerned with using a realistic tapering function.

3 | RESULTS

To assess the impact of reduced-precision parametrizations in SPEEDY, a set of different precision "forecast" experiments are compared to a "truth" run for a single initial state. The initial state was created by running SPEEDY from rest for 1 year (using boundary conditions starting from 1 January 1981), with SPPT switched on, to allow for spin up. From this initial state, SPEEDY is run with parametrizations in reduced precision and compared with double-precision (52 sbit) "truth" runs using the same initial state. Repeating these experiments using different initial states, generated by running the spin-up for further years, gave similar results so only the results from one initial state is presented here.

Similar to the suite of weather forecasts that are typically generated at operational weather forecast centres, we study both deterministic and ensemble "forecasts". We switch SPPT off for deterministic model runs and compare simulations with reduced precision against a single run at double precision (52 sbits). We switch SPPT on for ensemble model runs and generate 20 ensemble members for each ensemble at a given precision. All ensembles and ensemble members use the exact same initial state as the deterministic model runs. The only difference between each of the ensembles and ensemble members is the randomly generated seed used in the SPPT scheme.

The ultimate goal of these experiments is to find where differences from rounding errors are indistinguishable because of the inherent model uncertainty where the model uncertainty is represented by SPPT. All model runs use an identical initial state. In principle, the inclusion of initial condition uncertainty could make lower precisions more competitive. The initial state is taken from the spun-up model to minimise the effects of spin up in our experiments. There will still be some spin up in the model runs due to switching SPPT off in the deterministic experiments and due to the change in the SPPT pattern in the ensemble experiments. However, we expect the effect to be small.

3.1 | Differences in initial tendencies

In this section, the behaviour of rounding errors introduced in different parametrizations is investigated by quantifying the direct effect of rounding errors on the initial tendencies. We have output the tendency from each parametrization for the first timestep of the experiments with SPEEDY at different precision. Note that the inputs to each parametrization are identical for each precision, the only difference is due to the truncation of these inputs and constants as well as calculations performed at reduced precision within the parametrization. Since we are only looking at the initial tendencies, these are identical for the deterministic and ensemble experiments.

Figure 2 shows the average absolute difference in the initial boundary-layer temperature tendency with respect to double precision for the first timestep. Each line is for a different parametrization in reduced precision. The initial boundary-layer temperature tendencies were chosen because it is the only tendency where each parametrization has a direct contribution (apart from cloud which does not directly produce tendencies). Figure 2a shows the difference in the initial tendency of the individual parametrization and Fig. 2b shows the difference in the total initial tendency.

One difference between Fig. 2a and Fig. 2b is that the latter accounts for any knock-on effects that rounding errors in the reduced-precision parametrization will have on subsequent calculations. This knock-on effect will depend on the order the parametrizations are called. In SPEEDY the convection scheme is called first which then affects the condensation calculation. The results of convection and condensation then determine the cloud distribution used in the radiation calculations. Short-wave radiation is called first then long-wave radiation is integrated downwards and upwards with a call to the surface fluxes scheme in the middle. Finally the vertical diffusion scheme is called. This means rounding errors in vertical diffusion will have no affects on other initial tendencies whereas rounding errors in convection will have a knock-on effect on all the other parametrizations.

Note that gridpoints in which the initial tendency is zero for both reduced and full precision are ignored in Fig. 2.

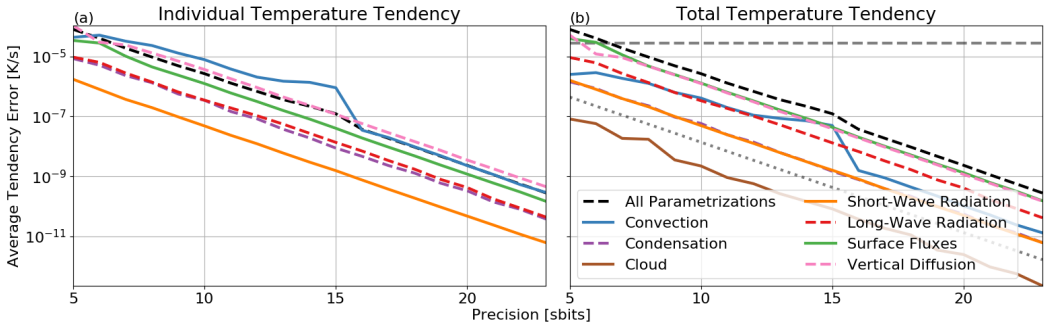

FIGURE 2 The mean of the absolute difference in each initial parametrization tendency with respect to double precision for the first model timestep. Shown is the initial boundary-layer temperature tendency because each parametrization directly affects this tendency. (a) The difference in the initial tendency from the individual parametrization. (b) The difference in the total initial tendency with the individual parametrization in reduced precision. The horizontal grey dashed line in (b) shows the mean absolute initial temperature tendency to show where the difference in the initial tendency becomes comparable to the actual initial tendency. The grey dotted line in (b) shows this mean initial tendency multiplied by the expected machine rounding error to show how the differences in the initial tendencies compare to a simple rounding error.

This means that parametrizations that do not affect all gridpoints (condensation, convection and vertical diffusion) will have differences averaged over all gridpoints in Fig. 2b but over a subset of gridpoints in Fig. 2a. This is most noticeable for convection which is second only to vertical-diffusion in terms of differences in the individual parametrization's initial temperature tendency with respect to double precision but appears less important in terms of differences in the total initial temperature tendency.

The difference in the initial temperature tendency with respect to double precision largely follows the expected machine rounding error ($2^{-(sbits+1)}$). This is shown by the dotted grey line in Fig. 2b which shows the machine rounding error multiplied by the mean absolute initial temperature tendency. The differences for each parametrization have the same gradient with a constant multiplier. The differences in the convection parametrization do not follow the machine rounding error as closely as the other parametrizations. This is because differences in the convection parametrization are dominated by the diagnosis of whether convection is triggered which will be discussed later.

Expressing the difference in the initial parametrization tendency as a multiple of the machine rounding error could be useful for model developers because it quantifies how sensitive the piece of code is to rounding errors and could be a useful benchmark when considering reduced precision. Table 1 shows a summary of the relative difference with respect to double precision for the initial tendencies as a multiple of the machine rounding error. The numbers shown are the minimum and maximum found by taking each line in Fig. 2b divided by the grey dotted line. Table 1 also summarises the differences in the initial tendency for other variables and vertical levels. The vertical levels have been grouped into boundary layer ($\sigma = 0.95$), lower troposphere ($\sigma = 0.835, 0.685, 0.51$), upper troposphere ($\sigma = 0.34, 0.2$) and stratosphere ($\sigma = 0.095, 0.025$).

The largest differences in the initial boundary-layer temperature tendency with respect to double precision come from convection, surface fluxes and vertical diffusion. Large differences in the initial boundary-layer specific-humidity tendencies can also be seen for these three parametrizations. Only the surface-fluxes parametrization perturbs the wind speeds and gives relatively small differences.

As expected, only rounding errors in the radiation parametrizations have any notable impact in the stratosphere.

TABLE 1 The average difference in the initial parametrized tendency with respect to double precision with different parametrizations in reduced precision expressed as a multiple of the machine rounding error. The vertical levels have been grouped into boundary layer ($\sigma = 0.95$), lower troposphere ($\sigma = 0.835, 0.685, 0.51$), upper troposphere ($\sigma = 0.34, 0.2$) and stratosphere ($\sigma = 0.095, 0.025$). Dashes show where the parametrizations have zero initial tendency.

Parametrization	Boundary Layer	Lower Troposphere	Upper Troposphere	Stratosphere
Temperature				
All Parametrizations	$167 - 285\varepsilon$	$66 - 545\varepsilon$	$42 - 762\varepsilon$	$78 - 111\varepsilon$
Convection	$6 - 115\varepsilon$	$17 - 492\varepsilon$	$25 - 744\varepsilon$	$0 - 2\varepsilon$
Condensation	$3 - 4\varepsilon$	$2 - 3\varepsilon$	$0 - 1\varepsilon$	$0 - 1\varepsilon$
Short-Wave Radiation	$4 - 4\varepsilon$	$3 - 4\varepsilon$	$7 - 12\varepsilon$	$71 - 107\varepsilon$
Long-Wave Radiation	$21 - 31\varepsilon$	$17 - 26\varepsilon$	$9 - 11\varepsilon$	$22 - 25\varepsilon$
Surface Fluxes	$86 - 133\varepsilon$	$2 - 3\varepsilon$	$0 - 1\varepsilon$	$0 - 1\varepsilon$
Vertical Diffusion	$54 - 117\varepsilon$	$24 - 38\varepsilon$	$1 - 7\varepsilon$	-
Specific Humidity				
All Parametrizations	$74 - 463\varepsilon$	$16 - 533\varepsilon$	$19 - 1078\varepsilon$	-
Convection	$9 - 370\varepsilon$	$11 - 528\varepsilon$	$18 - 1077\varepsilon$	-
Condensation	$2 - 2\varepsilon$	$2 - 2\varepsilon$	$1 - 1\varepsilon$	-
Short-Wave Radiation	$< 1\varepsilon$	-	-	-
Long-Wave Radiation	$1 - 1\varepsilon$	-	-	-
Surface Fluxes	$66 - 104\varepsilon$	-	-	-
Vertical Diffusion	$3 - 88\varepsilon$	$3 - 52\varepsilon$	-	-
Zonal Velocity				
Surface Fluxes	$13 - 18\varepsilon$	-	-	-
Meridional Velocity				
Surface Fluxes	$12 - 16\varepsilon$	-	-	-

The small but nonzero impact from other parametrizations is due to the knock-on effect of small rounding errors. The short-wave radiation parametrization results in the largest differences with respect to double precision in the stratosphere but the long-wave radiation parametrization gives similar differences for all model levels.

For the troposphere, the largest differences with respect to double precision come from the convection parametrization; however, there is a large spread because the large numbers are where the convection parametrization is switching on/off. Figure. 3a shows the numbers of gridpoints with nonzero initial tendencies as a function of precision for the convection parametrization. With lower precision, gridboxes are increasingly likely to be diagnosed as inactive in the convection parametrization. This introduces a bias rather than just random noise at low precision. This is also why the differences in the convection parametrization tendencies with respect to double precision do not follow the machine rounding error as well as other parametrizations. The differences become dominated by gridpoints activating/deactivating. The tendencies from the convection parametrization are comparatively large so any gridpoint that activates/deactivates will have a difference that is approximately equal to the tendency.

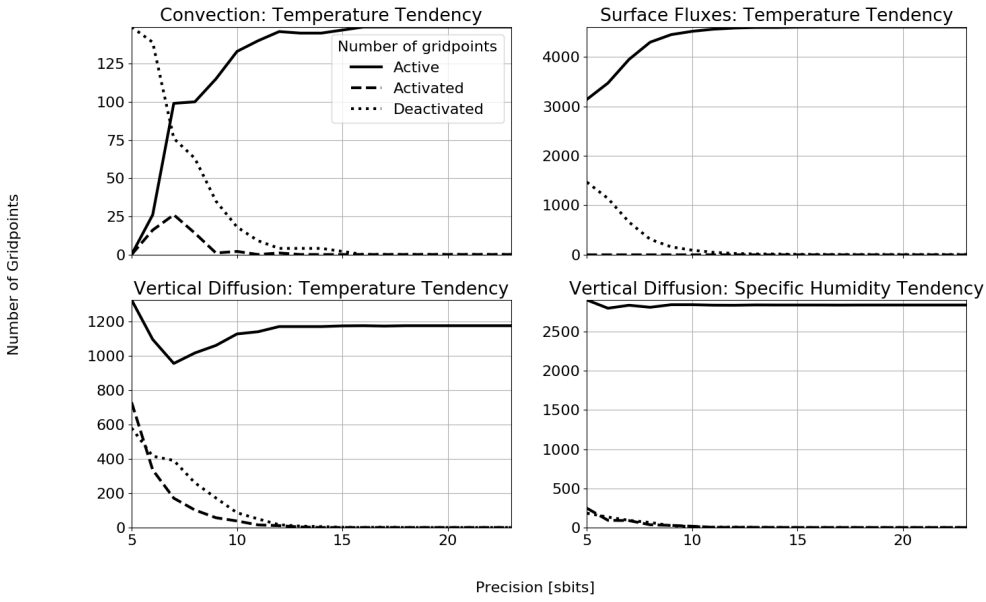

FIGURE 3 Number of gridpoints defined as active from the initial boundary-layer temperature tendency due to an individual parametrization as a function of precision that the parametrization is run at. A gridpoint is active if the initial tendency is nonzero. The gridpoint is defined as deactivated if the initial tendency is nonzero at double precision but zero at reduced precision and defined as activated if the initial tendency is zero at double precision and nonzero at reduced precision.

The surface-fluxes parametrization also shows issues with gridboxes becoming increasingly likely to be inactive (Fig. 2b). However, the relative differences for the initial surface-fluxes tendencies with respect to double precision in table 1 do not show as large a spread as for the initial convection tendencies. This is because the gridboxes where initial temperature tendencies are becoming zero at lower precision are those where the initial temperature tendencies were small at higher precision (not shown). Therefore, these differences do not dominate the average difference. This can still result in a bias because the temperature tendencies in the surface-fluxes parametrization are always positive in SPEEDY.

The vertical-diffusion parametrization also shows a large spread of differences in the initial tendencies with respect to double precision. It is less clear than the convection parametrization that the vertical diffusion parametrization is activating or deactivating because it is a combination of three different terms. The vertical diffusion parametrization includes shallow convection where there is conditional instability between the boundary layer and lowest tropospheric level, diffusion of moisture between the boundary layer and troposphere, and redistribution of dry-static energy in regions with a super-adiabatic lapse rate. This means that just diagnosing the parametrization as active or inactive as in Fig 3c and d does not tell the full story because one process could be deactivated but another process still gives nonzero tendencies. This is consistent with the large variation of differences in the initial specific-humidity tendencies with reduced precision (table 1) but no large changes in the number of active gridpoints (Fig. 3d).

3.2 | Differences in model evolution

In this section, differences between the evolution of experiments with reduced-precision parametrizations and double-precision experiments are quantified. Figure 4 shows differences with respect to the double precision experiment for the deterministic experiments with all parametrizations using the same reduced precision in the range 5-51 sbits. Shown is the root-mean-square (RMS) difference in geopotential height at 500 hPa in gridpoint space. The choice of variable and vertical level makes no practical difference to the results presented here (not shown). As expected, any change in precision results in differences in the model state. This can be seen by the fact there is error growth in the experiment with 51 sbits (Fig. 4a and c). Note that this difference is masked by the precision of the model output for the first couple of days because SPEEDY uses double precision but data is output in single precision.

In general, the lower the precision the faster the differences with respect to double precision will grow; however, error growth is not a simple, monotonic, function of precision. The differences for half precision (10 sbits) are always larger than single precision (23 sbits) which are in turn larger than differences with a single bit truncated (51 sbits). For intermediate precision however, this is less straightforward. For example, we can see that the 35-sbit experiment can have smaller differences than the higher-precision 51-sbit experiment. Also, the 11-sbit and 22-sbit experiments have very similar error growth despite the noticeable difference from the 10-sbit and 23-sbit experiments.

Looking at differences with respect to double precision as a function of precision (Fig 4b and d), we can categorise the error growth into three groups. At intermediate precision (11-22 sbits), the differences grow to moderate levels (> 1 m) within the first day and then follow a similar error growth. The differences in this range are almost independent of precision with slightly larger or smaller differences for some experiments.

At higher precision (23-51 sbits), the differences with respect to double precision are small initially ($< 10^{-4}$ m) but later rapidly grow to moderate levels (> 1 m) and then slowly grow from this point. Apart from the lower precision in this range, the timing of the rapid error growth is more random chance than a function of precision. This can be seen most clearly for the differences at 14 days on a logarithmic scale (orange dots in Fig. 4d): differences for experiments with 27-51 sbits are either ≈ 1 m or $\approx 10^{-4}$m with nothing in between and no obvious relation to the precision.

At lower precision (<11 sbits), the differences with respect to double precision are greater and increase the more

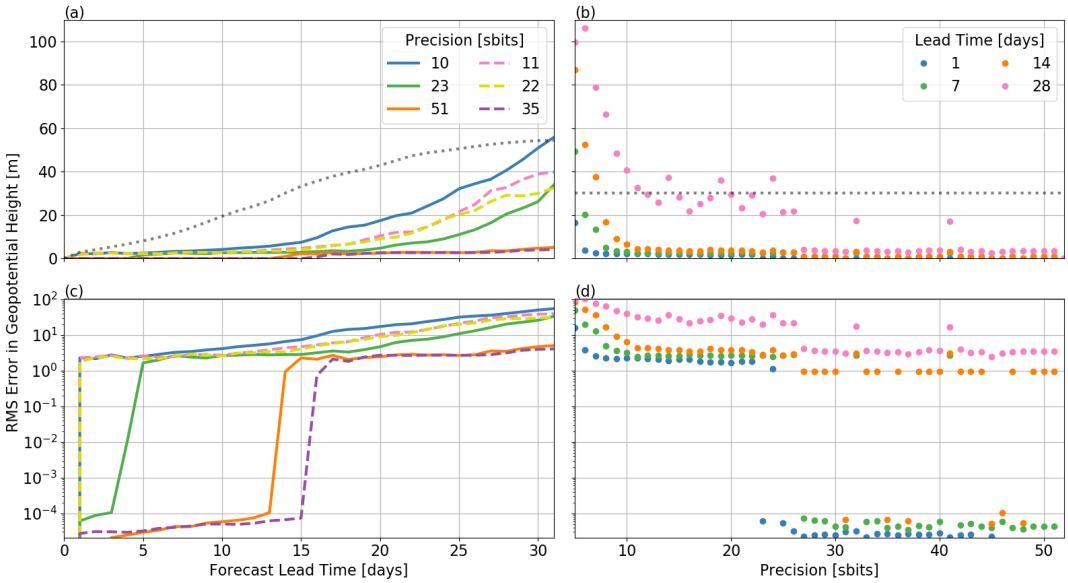

FIGURE 4 RMS difference in geopotential height at 500 hPa for deterministic experiments with reduced-precision parametrizations compared to a double-precision "truth" run. (a) Difference vs time for fixed precision. (b) Difference vs precision for fixed lead times. (c) and (d) show the same as (a) and (b) respectively but on a logarithmic scale. The dotted grey line in (a) shows the ensemble standard deviation of 500-hPa geopotential height for the 20-member, double-precision ensemble. The dotted grey line in (b) shows the same but only for 14-days lead time.

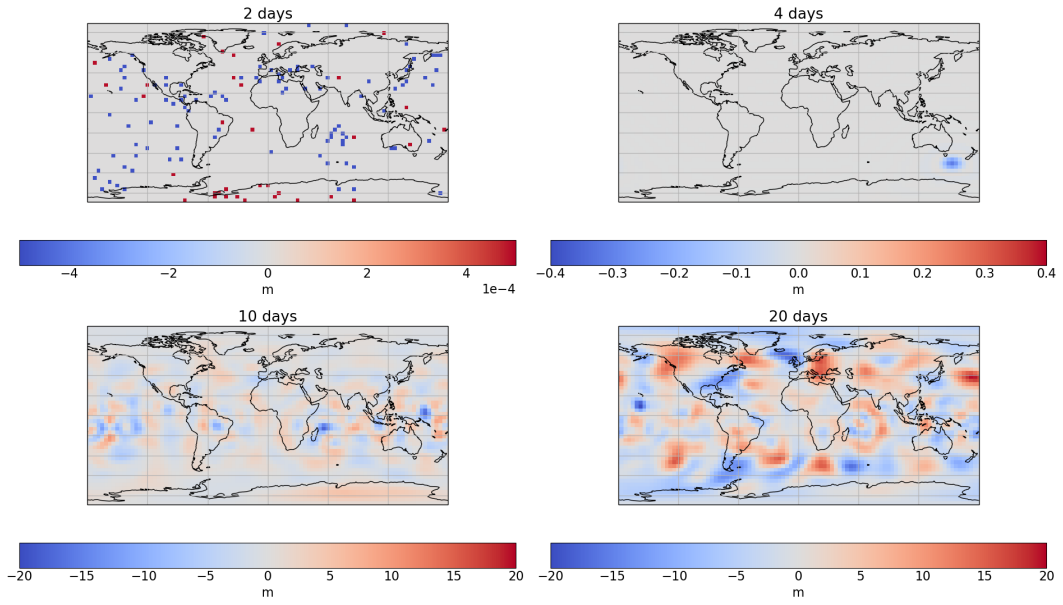

FIGURE 5 Difference in geopotential height at 500 hPa between the single-precision (23 sbit) and double-precision (52 sbit) deterministic model runs at various lead times. Please note the differences in the magnitude of the colour schemes between the plots.

precision is reduced. The increase in differences compared to intermediate precision is most distinct for 14-days lead time. At shorter lead times some of the lower precision experiments have similar differences as the intermediate precision experiments and at longer lead time the effects of chaos are more noticeable with more variation in the differences for intermediate precision.

These three groups of error growth are a result of how SPEEDY responds to the small errors in tendencies introduced by the reduced-precision parametrizations. Figure. 5 shows the difference between the 23 sbit (single precision) and 52 sbit (double precision) model runs at various lead times. Initially the differences are small (Fig. 5a). At this time, the small differences in the tendencies have led to differences at the smallest scales of the model but the differences are too small to affect the model dynamics. At a later stage, the small differences at the gridscale are enough that it changes the diagnosis of convection in a single gridbox (Fig. 5b) leading to a much larger difference between the two model runs that rapidly grows and propagates leading to more and more gridpoints with different branches in the parametrizations. The O(1) RMSE is then dominated by small-scale differences in the tropics (Fig. 5c) and grows slowly (see Fig 4a and c). At a later time the differences become more dominated by larger-scale patterns in the midlatitudes (Fig. 5d) and the total RMSE grows more rapidly.

At intermediate precision the rounding errors do not directly affect the large-scale dynamics but will affect the diagnosis of convection and indirectly result in rapid error growth within the first few timesteps such that the differences with respect to double precision (and each other) are similar after the first day. At higher precision it takes longer for the rounding errors to result in differences in the diagnosis of convection and therefore the rapid error growth is delayed. Since the rounding noise changing the diagnosis of convection is largely random, this explains why lower precision can have smaller differences with respect to double precision (i.e. 35 sbits vs 51 sbits in Fig 4). The reason that RMS

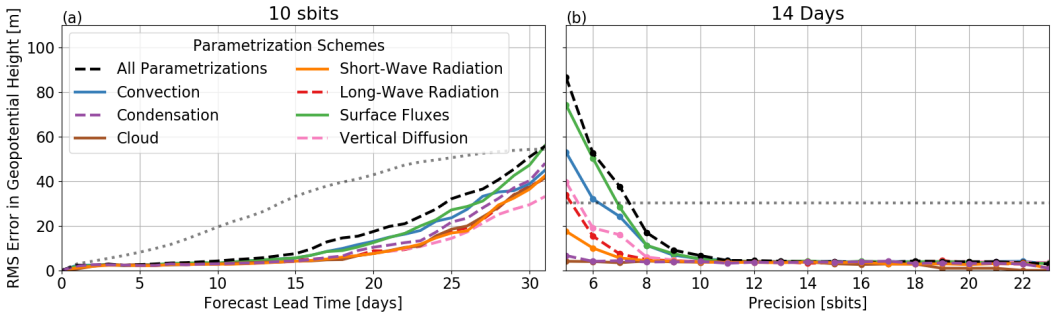

FIGURE 6 RMS difference geopotential height at 500 hPa for deterministic experiments with reduced-precision in individual parametrizations compared to a double-precision "truth" run. (a) Difference vs time for 10 sbit experiments. (b) Difference vs precision at 2-weeks lead time. The dotted grey line in (a) shows the ensemble standard deviation of 500-hPa geopotential height for the 20-member, double-precision ensemble. The dotted grey line in (b) shows the same but only for 14-days lead time.

differences with respect to double precision for experiments with 27-51 sbits are either $\approx$1 m or $\approx 10^{-4}$m with nothing in between is because the error growth is so rapid so no points here are sampling the middle of that error growth.

Figure 4 shows that differences with respect to the double precision experiment become large when reducing the precision of all parametrizations below 11 sbits; however, it should be possible to have individual parametrizations at lower precision without large errors. To determine the precision levels for each parametrization we have run experiments with a single parametrization in reduced precision in the range 5-23 sbits, for each parametrization. Figure 6 shows the RMS difference in geopotential height at 500 hPa with respect to the double precision experiment for these experiments with individual parametrizations in reduced precision. Figure 6a shows the difference as a function of time for each experiment using 10 sbits and Fig. 6b shows the difference as a function of precision for each experiment at two-weeks lead time. As expected, the differences with respect to double precision with individual parametrizations in reduced precision are less than or equal to the differences with all parametrizations in reduced precision.

The convection and surface-fluxes parametrizations result in larger differences with respect to double precision at 10 sbits compared to the other parametrizations. The error growth for the other parametrizations is similar to the intermediate-precision (11-22 sbits) experiments in Fig. 4. The difference with individual parametrizations in reduced precision has a similar behaviour to Fig. 4b with a flat intermediate region and rapidly increasing differences at lower precision. The only changes appear to be that this is shifted left (to lower precision) dependent on the parametrization. This shift to lower precision makes sense because the individual parametrizations will introduce less errors than all parametrizations combined. The surface-fluxes and convection parametrizations are the dominant source of errors at low precision: all other parametrizations could individually be reduced to 10 sbits or lower without a noticeable increase in the differences with respect to double precision compared to the differences at intermediate precision.

To truly determine the acceptable precision for parametrizations we need to consider whether the reduced-precision errors are within a given uncertainty. One way of doing this is to compare the differences with respect to double precision to the uncertainty from the ensemble spread. The dotted grey line in Fig. 4a shows the global mean standard deviation of 500-hPa geopotential height for the 20-member double-precision ensemble. The dotted grey line in Fig. 4b shows the same but for 14-days lead time. The ensemble spread grows faster than the differences shown in Fig. 4a and only experiments with precision lower than 8 sbits have differences larger than the ensemble spread at

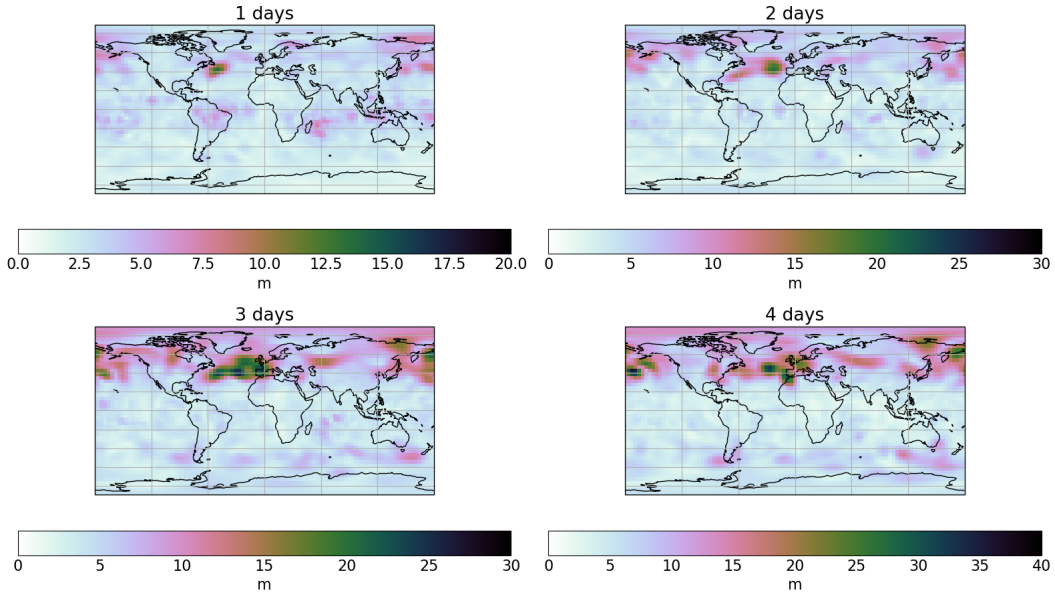

FIGURE 7 Standard deviation of geopotential height at 500 hPa for the 20-member double-precision ensemble at various lead times.

14-days lead time (Fig. 4b).

Comparing differences in deterministic experiments with ensemble spread is useful to see where the differences are small compared to the uncertainty; however, it is not a strong constraint on the precision. Comparing deterministic experiments to the stochastic ensemble is inconsistent because the model evolution and climatology will differ due to the introduction of SPPT. Figure 7 shows the spread of the double-precision ensemble at various lead times. The growth of ensemble spread is more rapid than the initial error growth in the intermediate precision deterministic experiments. Within the first two days the ensemble spread is dominated by large-scale differences in the midlatitudes (Fig. 7b) which grow and propagate (Fig. 7c and d) similar to the later stages of error growth in the single-precision deterministic experiment (Fig. 5d). This makes sense because the SPPT scheme is perturbing the parametrizations more than the reduced precision and globally. However, this does not rule out whether rounding errors in the parametrization lead to errors in the probability distribution predicted when they are included in ensembles and perturbed by SPPT. Instead we compare different precision ensembles so that we are comparing like-for-like. A precision can then be deemed acceptable if an ensemble run at that precision produces the same probability distribution as the double-precision ensemble. Each ensemble is run with 20 members. The double-precision ensemble is compared to 23 sbit (single precision), 10 sbit (half precision) and 8 sbit (low precision) ensembles as well as ensembles with individual parametrizations at 8 sbits.

To compare ensemble runs, we calculated the overlap of each ensemble probability distribution with the double-precision ensemble probability distribution. The overlapping coefficient is given by

$$OVL = \sum_{\mathbf{X}} min(f_1(\mathbf{X}), f_2(\mathbf{X}))d\mathbf{X}, \tag{5}$$

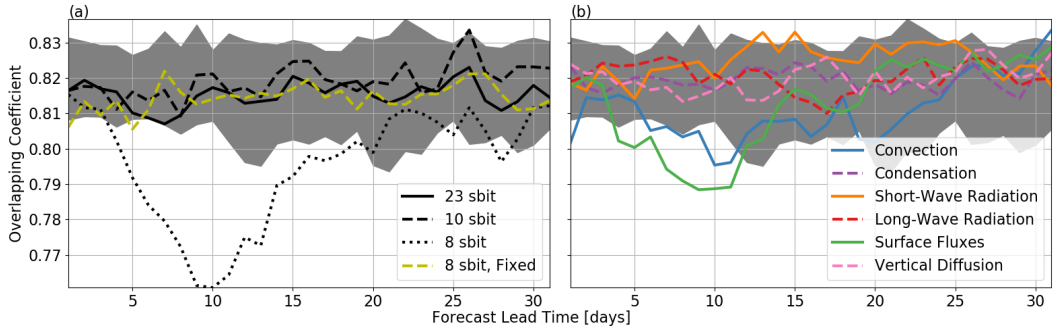

FIGURE 8 The overlapping coefficient for 500 hPa geopotential height between a double-precision reference ensemble and other ensembles with varying precision setups. (a) All parametrizations reduced to a single precision. (b) Individual parametrizations reduced to 8 sbits. All ensembles use 20 members. The grey shaded area shows the range in overlap calculated by randomly selecting two 20-member ensembles from 40 double-precision ensemble members 100 times.

where $f_1(\mathbf{X})$ and $f_2(\mathbf{X})$) are the two probability distributions being compared and the sum over $\mathbf{X}$ represents a discrete binning. The overlapping coefficient gives the fraction of probability mass common to both distributions (Inman and Bradley, 1989) and is used here as a measure of agreement between two ensembles.

The ensemble probability distributions ($f_1(\mathbf{X})$ and $f_2(\mathbf{X})$ in equation 5) are calculated at individual gridpoints by applying a Gaussian kernel filter to values of each ensemble member at that gridpoint. The standard deviation of the Gaussian filter is taken as $\sigma/2$ where σ is the standard deviation of the ensemble at that gridpoint. This was chosen to give a smooth probability distribution without oversmoothing and to adapt the smoothing to the spatial and temporal changes in spread.

Figure 8 shows the global area-weighted average of the overlapping coefficient for geopotential height at 500 hPa for each ensemble with respect to the reference 52-sbit ensemble as a function of time. To give a measure of uncertainty a second 20-member double-precision ensemble was run (with different randomly generated seeds in the SPPT scheme). The overlapping coefficient was then recalculated 100 times for two randomly selected 20-member ensembles taken from the combined 40 members from the two double-precision ensembles. The grey shading in Fig 8 shows the range in the overlapping coefficient calculated this way.

Both the 23 sbit and the 10 sbit ensembles remain in the same range suggesting that precision as low as 10 sbits is acceptable for all parametrizations. This is slightly lower than the region of interchangeable differences in the deterministic experiments showing that SPPT can mask these rounding errors. Although not exactly the same, the point at which differences rapidly increase with reducing precision in the deterministic experiments (see Fig. 6) does provide a good initial estimate of the acceptable precision when including SPPT.

Reducing precision below 10 sbits can degrade the ensemble: the 8-sbit ensemble is clearly worse than higher precision ensembles (Fig. 8a). After the first few days, the overlap of the 8-sbit ensemble starts to drop and reaches a minimum of $\approx$ 76% around 10 days before increasing again but remaining below the reference overlap. Of the ensembles with an individual parametrization at 8 sbits, the convection and surface-fluxes ensembles show smaller, but significant, decreases in overlap over a similar timescale (Fig. 8b). These are the two parametrizations that also gave the largest errors at low precision for the deterministic experiments.

The decrease in overlap for the 8 sbit ensemble with respect to the double-precision ensemble highlights the limitation of using ensemble spread as a constraint on differences in deterministic model runs. The difference between

the 8-sbit deterministic experiment and the double-precision (52 sbit) deterministic experiment is lower than the ensemble spread at 2-weeks lead time (Fig. 6b) but the 8-sbit ensemble is shown to be inconsistent with a double-precision ensemble. At longer lead times (>4 weeks, see Fig. 6a) the difference between the 10-sbit deterministic experiment and the double-precision (52 sbit) deterministic experiment becomes comparable to the ensemble spread; however, it is shorter lead times (1-2 weeks, see Fig. 8) that the differences in low-precision ensembles become apparent and the 10-sbit ensemble is shown to be consistent with the double-precision ensemble at all lead times anyway.

3.3 | Fixes to optimise precision for problematic parametrizations

The dominant issue for the convection parametrization is that low precision results in a reduction in the number of gridboxes where the convection parametrization is triggered. This is caused by the diagnosis of convectively unstable gridboxes in SPEEDY which is based on conditional instability. To calculate this, the static energy (SE) is input to the convection parametrization.

$$SE = c_p T + \phi, \tag{6}$$

where c_p is the specific heat capacity of dry air at constant pressure, T is temperature and ϕ is the geopotential. From SE the moist static energy (MSE) and saturated moist static energy (MSE_{sat}) are calculated.

$$\mathrm{MSE} = \mathrm{SE} + Lq, \tag{7}$$

$$\mathrm{MSE}_{sat} = \mathrm{SE} + Lq_{sat} \tag{8}$$

where L is the latent heat of evaporation, q is the specific humidity and q_{sat} is the saturated specific humidity.

Convection is diagnosed if MSE_{sat} on any tropospheric half level above the first is lower than MSE_{sat} in the boundary layer. Convection is then activated if either of the following two criteria are met

1. MSE_{sat} at the tropospheric half level is lower than MSE in the boundary layer or the lowest tropospheric level
2. Specific humidity in the boundary layer and the lowest tropospheric level exceed set thresholds

The problem with this diagnosis of convection at low precision is checking the differences between two values that are almost equal. With the coarser representation of numbers at lower precision two close numbers at double precision can often become the same number at low precision. For the convection parametrization, this means that checking if one number is greater than another number will preferentially go from "True" at high precision to "False" at low precision. Changing the check from "greater than" to "greater than or equal to" can remove this problem but introduces a similar problem in the other direction: checks that were previously "False" at high precision will become equal and therefore "True" at low precision leading to a bias where the convection parametrization is activating too much at low precision.

This issue with equality in logical checks is not easy to fix but the problem can be mitigated at low precision. This is because we are only interested in vertical differences in MSE but these are small compared to the absolute value of MSE. Therefore, because rounding errors are relative, the differences in MSE are excessively truncated. This means that this issue with MSE emerges at higher precision than other errors. By re-expressing MSE as an anomaly, the rounding errors in the differences in MSE can be put more in line with other variables allowing us to use lower precision. To achieve this, the surface value of SE is subtracted from each column before it is input to the convection parametrization.

The change to SE being stored as an anomaly is also propagated into the vertical diffusion parametrization. Although vertical diffusion was acceptable at 8 sbits, there were issues with the parametrization activating or deactivating at low precision with some checks based on vertical differences in SE.

The dominant issue for the surface-fluxes parametrization is that low precision results in small, positive, tendencies being rounded to zero. The temperature tendency from surface fluxes is the sensible heat flux (SHF) computed as,

$$SHF = \rho_s C |V_s| c_p (T_s - T_0), \tag{9}$$

where ρ_s is the density at the surface, C is an empirical coefficient which is different for land and sea, $|V_s|$ is the wind-speed magnitude including a gust factor, c_p is the specific heat capacity of dry air at constant pressure, T_s is the extrapolated surface temperature and T_0 is the temperature in the boundary layer.

The problem with this equation at low precision is diagnosing the difference between the boundary-layer temperature and the surface temperature. The difference between these two temperatures is small (typically less than one Kelvin). At 8 sbits, temperature in Kelvin close to 0° C gives a precision of ≈ 1 K. Therefore, at low-precision the difference between these numbers is increasingly likely to become zero. However, if the temperature is expressed in Celsius the precision is improved because the rounding error is relative to the absolute value of the number being stored. Unlike the static energy, temperature is also used multiplicatively and so this change is only applied within the surface-fluxes parametrization as it requires more careful code changes.

With these changes to SE and temperature implemented we have re-run the deterministic experiments as well as the 8-sbit ensemble. Table 2 shows the relative difference in the initial tendencies with respect to double precision as a multiple of the machine rounding error for the parametrizations that have been modified. Comparing with table 1 we can see that the relative differences in the initial tendencies are much improved. Even the initial velocity tendencies due to surface fluxes have improved which can only be due to converting temperature to Celsius.

The spread in differences in the initial convection and vertical-diffusion tendencies with respect to double precision has also reduced. This reduced spread corresponds to an improvement of the bias in the number of gridpoints with nonzero initial tendencies at low precision. Figure 9 shows the numbers of gridpoints with nonzero initial tendencies as a function of precision, with the code changes applied, for the same initial tendencies as Fig. 3. All the parametrizations show substantial improvement. The initial tendencies for surface fluxes and vertical diffusion don't have any flipping of active/inactive until very low precision and even then it is a small number of gridpoints. Convection is much better but still shows the largest differences which suggests further improvements to the convection diagnosis could be made.

These improvements have resulted in a reduction in the differences with respect to double precision from the deterministic experiments with all experiments at 10 sbits showing similar error growth (Fig. 10a). The convection parametrization now gives the largest differences at low precision, consistent with the differences in the initial tendencies. These differences are now much closer to the next worst parametrization, long-wave radiation, which hasn't been modified here (Fig. 10b). These improvements are also reflected in a much improved 8-sbit ensemble (Fig. 8a). The fixed 8-sbit ensemble is much better than the original 8-sbit ensemble and is almost entirely within the uncertainty of the double-precision ensemble apart from some small degradation near the beginning.

4 | CONCLUSIONS

We have investigated reducing the precision of parametrizations in an intermediate-complexity atmospheric model (SPEEDY). Reducing precision in parametrizations introduces errors to the gridpoint tendencies; however, the errors

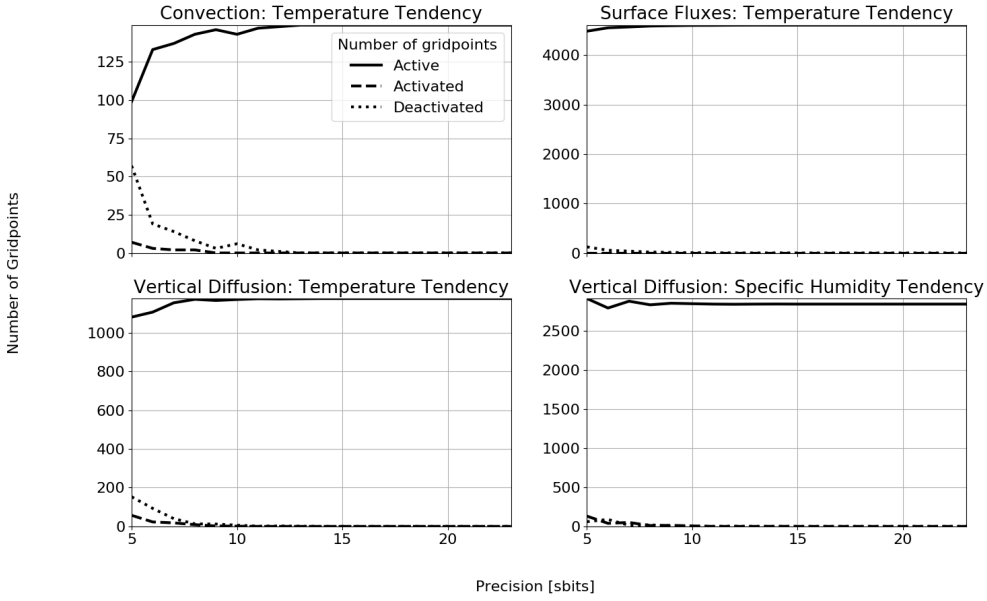

FIGURE 9 Same as Fig. 3 but with the changes to the model code described in section 3.3

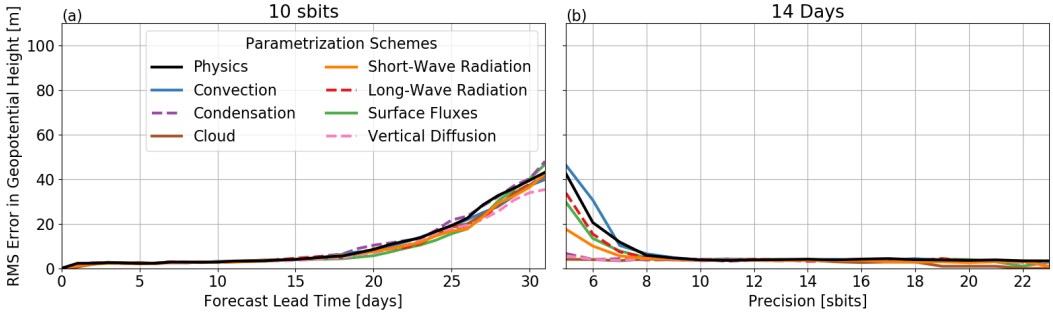

FIGURE 10 Same as Fig. 6 but with the changes to the model code described in section 3.3

TABLE 2 The same as table 1 but with the changes to the model code described in section 3.3. Only the parametrizations that have been modified are shown.

Parametrization	Boundary Layer	Lower Troposphere	Upper Troposphere	Stratosphere
Temperature				
All Parametrizations	$33-43\varepsilon$	$22-55\varepsilon$	$16-84\varepsilon$	$77-110\varepsilon$
Convection	$0-9\varepsilon$	$2-33\varepsilon$	$4-70\varepsilon$	$<1\varepsilon$
Surface Fluxes	$15-23\varepsilon$	$2-4\varepsilon$	$0-1\varepsilon$	$0-1\varepsilon$
Vertical Diffusion	$6-7\varepsilon$	$2-2\varepsilon$	$0-1\varepsilon$	-
Specific Humidity				
All Parametrizations	$13-39\varepsilon$	$7-40\varepsilon$	$7-112\varepsilon$	-
Convection	$1-26\varepsilon$	$2-35\varepsilon$	$6-111\varepsilon$	-
Surface Fluxes	$9-14\varepsilon$	-	-	-
Vertical Diffusion	$3-11\varepsilon$	$3-8\varepsilon$	-	-
Zonal Velocity				
Surface Fluxes	$2-2\varepsilon$	-	-	-
Meridional Velocity				
Surface Fluxes	$2-2\varepsilon$	-	-	-

can be masked by the inherent uncertainty in these tendencies. Including stochastic perturbation of parametrization tendencies (SPPT) as a representation of model uncertainty, we can reduce the precision of all parametrizations to 10 sbits (equivalent to a half precision mantissa) without seeing any degradation in our "forecast" experiment: a 10-sbit ensemble is indistinguishable from a double-precision (52 sbit) ensemble. Note that these ensemble "forecasts" use an identical initial state and the only difference between ensemble members is the randomly generated seed used in the SPPT scheme. This was done to make sure it was only model uncertainty, as opposed to any initial condition uncertainty, that could be masking rounding errors.

We have investigated the effects of rounding errors on the initial tendencies output from the parametrizations. The average difference in the tendency from each parametrization with respect to double precision was found to scale with the machine rounding error and can be expressed as a multiple of the machine rounding error (see table 1). The differences in the convection parametrization with respect to double precision do not follow the machine rounding error as closely as other parametrizations. This is because at low precision, differences in the convection parametrization tendencies are dominated by the diagnosis of convection as active or inactive over a subset of gridpoints. Looking at the number of active gridpoints showed that the convection parametrization is strongly biased towards deactivating when reducing precision. Other parametrizations also show tendencies becoming zero with reduced precision. For the surface-fluxes parametrization this is due to small tendencies becoming zero as opposed to a diagnosis of active or inactive.

We also compared deterministic experiments (no SPPT) where the only difference between each model run was the precision used in the parametrizations. Without SPPT, any reduction in precision will eventually result in large differences between experiments due to the chaotic nature of the atmosphere. This means that determining the acceptable precision can be difficult without some way of quantifying the acceptable level of differences. However,

with deterministic experiments using an identical initial state, there is a wide range of precision (11-22 sbits) where the error growth remains similar even though the differences in the initial tendency scales with the machine rounding error. The limit of this range of precision with similar error growth provides a surprisingly good estimate (< 11 sbits) of where differences in ensembles with SPPT will start to become apparent (< 10 sbits).

A similar error growth for a large difference in errors is indicative of the butterfly effect (Lorenz, 1969). Durran and Gingrich (2014) showed that large-scale errors in initial conditions dominate butterfly-effect type error growth in operational forecasts. The error growth is different here because we have excluded initial condition uncertainty to focus on model uncertainty. In ensemble experiments with differences generated only by a stochastic convection parametrization there is upscale error growth (Selz, 2018; Baumgart et al., 2019). Baumgart et al. (2019) analysed the dynamical mechanisms of uncertainty growth in these experiments and showed that short lead time differences were dominated by differences in convection which are then projected on to small differences in the upper-level midlatitude wave pattern within the first two days which then continue to grow over the next two weeks. The uncertainty growth in our ensemble experiments, where the only difference is the randomly generated seed used in the SPPT scheme, are qualitatively similar.

The error growth in the deterministic experiments initially behaves differently to the ensemble spread. For higher precision experiments, the noise from the parametrizations results in very small differences that do not affect the model dynamics. These differences eventually trigger a difference in a parametrization (i.e. convection active/inactive) which rapidly grows such that the error is dominated by small-scale differences in the tropics. The errors in the tropics grow slowly and it is then only at longer lead times (≈ 2 weeks later) that the differences become dominated by the midlatitude wave pattern.

Finding the acceptable precision by comparing deterministic experiments to each other gives a similar answer to comparing the ensembles with each other; however, comparing the deterministic experiments with the ensembles directly can make a lower precision appear acceptable. This is because we are comparing inconsistent models: the deterministic experiments without SPPT and the ensembles with SPPT. The error growth for deterministic experiments with rounding error is also different to the growth in ensemble spread which means the conclusion about acceptable precision is strongly dependent on lead time when only using the ensemble spread to constrain deterministic experiments.

We have also run experiments with individual parametrizations in reduced precision. Reducing precision of individual parametrizations has been shown to be useful in an operational model context: Gilham (2018) improved run times in the Met Office Unified Model by reducing the microphysics parametrization to single precision. We find that for the deterministic experiments and stochastic ensembles two parametrizations, convection and surface fluxes, are the dominant source of differences. This gives the option of saving more computational resources by setting precision on a parametrization-by-parametrization basis.

Some changes to the model code can be made to mitigate issues with the convection parametrization deactivating and the surface-fluxes parametrization rounding small tendencies to zero allowing for extra bits of precision reduction. In SPEEDY the convection parametrization is activated in regions of conditional instability diagnosed based on differences in moist static energy. The problem with these threshold-type checks at low precision is that a comparison between two similar numbers becomes a comparison between two identical numbers. This means that for comparisons which don't include equality ("<", ">") the comparison becomes more likely to return "False" at lower precision whereas comparisons including equality ("<=", ">=", "==") become more likely to return "True". While this comparison issue cannot be easily fixed, the likelihood of the two slightly-different numbers being rounded to the same number can be reduced by expressing them as an anomaly field. This will, in effect, increase the number of digits that are available to represent these fields when using reduced precision since rounding errors are relative. To achieve this, we replaced the

static energy field input to the convection parametrization with an anomaly of static energy relative to the surface.

The surface-fluxes parametrization rounds small temperature tendencies to zero at low precision. The calculated tendency is proportional to differences between the surface temperature and the boundary-layer temperature. The differences between these two temperatures are small compared to the absolute values of temperature when expressed in Kelvin; therefore, the two temperatures are increasingly likely to be expressed as the same number as precision is reduced resulting in zero tendency. This can be mitigated by expressing temperature in Celsius (or as an anomaly) such that the range of numbers is centred close to zero with the result that more decimal digits are available to express the actual physical signal. However, care needs to be taken if model fields are re-scaled to make sure that all model fields are working with the correct physical units.

Rerunning experiments with these improvements included shows a substantial improvement, allowing us to use lower precision than initially expected. This means that the errors at low precision were dominated by a few operations. Both improvements were changing an absolute field to an anomaly field which is good practice when trying to use low precision. Beyond these improvements certain sensitive operations can be performed at higher precision (Dawson et al., 2017) or a more fine grained approach can be used where each variable uses the lowest precision necessary (Düben et al., 2017).

The code changes to improve the model's resilience to low precision were implemented based on differences found in the initial parametrization tendencies with respect to double precision. For more complex models this means that these improvements could be made using a single-column model rather than having to run a full model which would save computational resources. However, the initial tendencies alone did not tell us where the rounding errors would lead to degraded forecasts so the single-column model could only be used for the first step of improving the model code.

The results here are not exhaustive: we only considered a single "forecast" and did not use a bespoke stochastic scheme for SPEEDY. Instead, these results help us to design future experiments for implementing reduced-precision parametrizations in fully complex numerical weather and climate models which can be summarised in a step-by-step guide.

1. Using a single column model, find the best version of a parametrization for low-precision computing. By modifying the code, minimise the amplification of the rounding error found in the initial tendencies and reduce biases that emerge at low precision
2. With the improved parametrizations implemented in the full model, run deterministic forecasts over a range of precision to find the point at which differences rapidly increase relative to double-precision reference forecasts
3. Using the precision setup determined from the deterministic forecasts, run ensembles of forecasts, using an internal representation of model uncertainty as the only difference between forecasts, to find a low-precision setup that is still indistinguishable from double precision

Each step is increasingly computationally expensive so this step-by-step guide allows computational resources to be saved while still allowing a thorough implementation of reduced precision.

ACKNOWLEDGEMENTS

We would like to thank the two anonymous reviewers whose comments helped improve the content and clarity of this paper. This work was funded by the European Research Council project ITHACA (grant No. 741112). PD gratefully acknowledges funding from a Royal Society University Research Fellowship and the ESiWACE2 project. The ESiWACE2 project has received funding from the European Union's Horizon 2020 research and innovation programme under grant

agreement No. 823988.

ENDNOTES

REFERENCES

Amezcua, J., Kalnay, E. and Williams, P. D. (2010) The Effects of the RAW Filter on the Climatology and Forecast Skill of the SPEEDY Model. *Mon. Weather Rev.*, **139**, 608–619.

Bauer, P., Thorpe, A. and Brunet, G. (2015) The quiet revolution of numerical weather prediction. *Nature*, **525**.

Baumgart, M., Ghinassi, P., Wirth, V., Selz, T., Craig, G. C. and Riemer, M. (2019) Quantitative View on the Processes Governing the Upscale Error Growth up to the Planetary Scale Using a Stochastic Convection Scheme. *Mon. Weather Rev.*, **147**, 1713–1731.

Chantry, M., Thornes, T., Palmer, T. and Düben, P. (2018) Scale-Selective Precision for Weather and Climate Forecasting. *Mon. Weather Rev.*, **147**, 645–655.

Dawson, A. and Düben, P. D. (2017) rpe v5: an emulator for reduced floating-point precision in large numerical simulations. *Geosci. Model Dev.*, **10**, 2221–2230.

Dawson, A., Düben, P. D., MacLeod, D. A. and Palmer, T. N. (2017) Reliable low precision simulations in land surface models. *Clim. Dyn.*

Dee, D. P., Uppala, S. M., Simmons, A. J., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M. A., Balsamo, G., Bauer, P., Bechtold, P., Beljaars, A. C. M., van de Berg, L., Bidlot, J., Bormann, N., Delsol, C., Dragani, R., Fuentes, M., Geer, A. J., Haimberger, L., Healy, S. B., Hersbach, H., Hólm, E. V., Isaksen, L., Kållberg, P., Köhler, M., Matricardi, M., McNally, A. P., Monge-Sanz, B. M., Morcrette, J.-J., Park, B.-K., Peubey, C., de Rosnay, P., Tavolato, C., Thépaut, J.-N. and Vitart, F. (2011) The ERA-Interim reanalysis: configuration and performance of the data assimilation system. *Q. J. R. Meteorol. Soc.*, **137**, 553–597.

Düben, P., Diamantakis, M., Lang, S., Saarinen, S., Sandu, I., Wedi, N. and Wilhelmsson, T. (2018) Progress in using single precision in the IFS. *ECMWF Newsl.*, **157**, 26–31.

Düben, P. D. and Dolaptchiev, S. I. (2015) Rounding errors may be beneficial for simulations of atmospheric flow: results from the forced 1D Burgers equation. *Theor. Comput. Fluid Dyn.*, **29**, 311–328.

Düben, P. D. and Palmer, T. N. (2014) Benchmark Tests for Numerical Weather Forecasts on Inexact Hardware. *Mon. Weather Rev.*, **142**, 3809–3829.

Düben, P. D., Russell, F. P., Niu, X., Luk, W. and Palmer, T. N. (2015) On the use of programmable hardware and reduced numerical precision in earth-system modeling. *J. Adv. Model. Earth Syst.*, **7**, 1393–1408.

Düben, P. D., Subramanian, A., Dawson, A. and Palmer, T. N. (2017) A study of reduced numerical precision to make superparameterization more competitive using a hardware emulator in the OpenIFS model. *J. Adv. Model. Earth Syst.*, **9**, 566–584.

Durran, D. R. and Gingrich, M. (2014) Atmospheric Predictability: Why Butterflies Are Not of Practical Importance. *J. Atmos. Sci.*, **71**, 2476–2488.

Gilham, R. (2018) 32-bit Physics in the Unified Model. *Tech. Rep. 626*, UK Met Office.

Hatfield, S., Düben, P., Chantry, M., Kondo, K., Miyoshi, T. and Palmer, T. (2018) Choosing the Optimal Numerical Precision for Data Assimilation in the Presence of Model Error. *J. Adv. Model. Earth Syst.*, **0**.

Inman, H. F. and Bradley, E. L. (1989) The overlapping coefficient as a measure of agreement between probability distributions and point estimation of the overlap of two normal densities. *Commun. Stat. - Theory Methods*, **18**, 3851–3874.

Jeffress, S., Düben, P. and Palmer, T. (2017) Bitwise efficiency in chaotic models. *Proc. R. Soc. A Math. Phys. Eng. Sci.*, **473**.

Leutbecher, M., Lock, S.-J. J., Ollinaho, P., Lang, S. T. K., Balsamo, G., Bechtold, P., Bonavita, M., Christensen, H. M., Diamantakis, M., Dutra, E., English, S., Fisher, M., Forbes, R. M., Goddard, J., Haiden, T., Hogan, R. J., Juricke, S., Lawrence, H., MacLeod, D., Magnusson, L., Malardel, S., Massart, S., Sandu, I., Smolarkiewicz, P. K., Subramanian, A., Vitart, F., Wedi, N. and Weisheimer, A. (2017) Stochastic representations of model uncertainties at ECMWF: state of the art and future vision. *Q. J. R. Meteorol. Soc.*, **143**, 2315–2339.

Leutwyler, D., Fuhrer, O., Lapillonne, X., Lüthi, D. and Schär, C. (2016) Towards European-scale convection-resolving climate simulations with GPUs: a study with COSMO 4.19. *Geosci. Model Dev.*, **9**, 3393–3412.

Lorenz, E. N. (1969) The predictability of a flow which possesses many scales of motion. *Tellus*, **21**, 289–307.

Molteni, F. (2003) Atmospheric simulations using a GCM with simplified physical parametrizations. I: model climatology and variability in multi-decadal experiments. *Clim. Dyn.*, **20**, 175–191.

Nakano, M., Yashiro, H., Kodama, C. and Tomita, H. (2018) Single Precision in the Dynamical Core of a Nonhydrostatic Global Atmospheric Model: Evaluation Using a Baroclinic Wave Test Case. *Mon. Weather Rev.*, **146**, 409–416.

Palmer, T. N. (2012) Towards the probabilistic Earth-system simulator: a vision for the future of climate and weather prediction. *Q. J. R. Meteorol. Soc.*, **138**, 841–861.

— (2014) More reliable forecasts with less precise computations: a fast-track route to cloud-resolved weather and climate simulators? *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.*, **372**.

Palmer, T. N., Buizza, R., Doblas-Reyes, F., Jung, T., Leutbecher, M., Shutts, G. J., Steinheimer, M. and Weisheimer, A. (2009) Stochastic Parametrization and Model Uncertainty. *ECMWF Tech. Memo.*, **598**, 1–42.

Rüdisühli, S., Walser, A. and Fuhrer, O. (2013) COSMO in single precision. *COSMO Newsl.*

Selz, T. (2018) Estimating the intrinsic limit of predictability using a stochastic convection scheme. *J. Atmos. Sci.*

Thornes, T., Düben, P. and Palmer, T. (2018) A power law for reduced precision at small spatial scales: Experiments with an SQG model. *Q. J. R. Meteorol. Soc.*, **0**.

Váňa, F., Düben, P., Lang, S., Palmer, T., Leutbecher, M., Salmond, D. and Carver, G. (2016) Single Precision in Weather Forecasting Models: An Evaluation with the IFS. *Mon. Weather Rev.*, **145**, 495–502.