

A Lagrangian vertical coordinate version of the ENDGame dynamical core. Part II: Evaluation of Lagrangian conservation properties

I. Kavčič^{a*} and J. Thuburn^a

College of Engineering, Mathematics and Physical Sciences, University of Exeter, Devon, EX4 4QF, UK

*Correspondence to: Iva Kavčič, Met Office, FitzRoy Road, Exeter EX1 3PB, UK. E-mail: iva.kavcic@metoffice.gov.uk

Accepted Article

A baroclinic instability test case is used to compare the Lagrangian conservation properties of three versions of a semi-implicit semi-Lagrangian dynamical core: one using a height based vertical coordinate and two using a Lagrangian vertical coordinate. The Lagrangian coordinate versions differ in the choice of target levels to which model levels are reset after each step—the first uses the initial model level heights while the second uses quasi-Lagrangian target levels. A range of diagnostics related to Lagrangian conservation are computed, including global entropy, unavailable energy, cross-isentrope mass flux, and consistency of potential temperature and potential vorticity with passive tracers and parcel trajectories. The global entropy, unavailable energy, and cross-isentrope fluxes do not suggest any clear advantage or disadvantage from the use of a Lagrangian vertical coordinate, though the cross-isentrope flux reveals a flaw in the formulation of the remapping of potential temperature in the Lagrangian coordinate model at the top boundary. The use of a Lagrangian vertical coordinate with quasi-Lagrangian target levels improves the consistency among potential temperature as a dynamical variable, potential temperature as a tracer and potential temperature on Lagrangian particle trajectories. It also improves consistency between a potential vorticity tracer and potential vorticity on Lagrangian particle trajectories. However, it degrades the consistency between model and tracer potential vorticity, as well as between model potential vorticity and potential vorticity on Lagrangian trajectories. This degradation appears to be related to the slopes of model levels, which are greater in the version with quasi-Lagrangian target levels.

This article has been accepted for publication and undergone full peer review but has not been through the copyediting, typesetting, pagination and proofreading process, which may lead to differences between this version and the Version of Record. Please cite this article as doi: 10.1002/qj.3375

1. Introduction

Throughout much of the atmosphere a number of important quantities are approximately conserved following the motion of air parcels. These quantities include specific entropy, potential temperature or equivalent potential temperature, composition (total specific humidity and other long-lived chemical constituents), and potential vorticity. Accurate numerical simulation of the atmosphere requires these Lagrangian conservation properties to be accurately captured by the numerical model. See Kavčić and Thuburn (2018) (hereafter Part I) and references therein and also section 3 below for further discussion.

It has often been argued (again see Part I and references therein) that the Lagrangian conservation properties of a numerical model can be improved by the use of a Lagrangian or quasi-Lagrangian vertical coordinate. However, when comparisons are made between a Lagrangian-coordinate model and a model with a more conventional (e.g. pressure based or height based) vertical coordinate, it is often the case that the formulations of the entire dynamical cores are very different, not just the type of vertical coordinate (e.g. Johnson *et al.* 2000; Rasch *et al.* 2006; Whitehead *et al.* 2015). This, then, makes any differences in the results difficult to attribute solely to the different vertical coordinates. It would be very desirable to have a ‘clean’ comparison, in which the only significant difference between models is the vertical coordinate, and other aspects of the formulation are as similar as possible.

Such a clean comparison was carried out by Mahowald *et al.* (2002) for a chemical transport model. They found that Lagrangian conservation properties such as age of air in the lower stratosphere were indeed better captured with an isentropic vertical coordinate than with a hybrid-pressure coordinate. Zhu and Schneider (1997) and Webster *et al.* (1999) made clean comparisons of general circulation models using a hybrid-isentropic coordinate and a hybrid-pressure coordinate. They found some notable improvements in the global circulation with the hybrid-isentropic coordinate, particularly a reduced cold bias in the winter polar lower stratosphere. More recently, Lauritzen *et al.* (2014) investigated the conservation of the global axial angular momentum (AAM) in NCAR’s Community Atmosphere Model Spectral Element (CAM-SE) dynamical core. They observed that using a Lagrangian vertical coordinate instead of the fixed Eulerian (hybrid-sigma) coordinate did not change the global AAM conservation properties of CAM-SE. However, those studies did not directly evaluate the Lagrangian conservation properties in the comparisons.

The present paper compares the Lagrangian conservation properties of three versions of a dynamical core, one using a height based coordinate (abbreviated to HB below) and the others using a Lagrangian vertical coordinate (abbreviated to LVC), but which are otherwise as similar as possible. All three versions use semi-implicit semi-Lagrangian integration schemes on a longitude-latitude spherical C-grid with a Charney-Phillips vertical staggering. The HB version uses numerical methods almost identical to those of the ENDGame dynamical core (Wood *et al.* 2014) used for operational forecasting at the Met Office. The LVC versions use the same numerical methods or analogous numerical methods where possible.

To prevent the folding over of model levels, the LVC model versions reset the model level heights at the end of each time step to certain ‘target levels’ and prognostic fields are remapped to these new heights. The two LVC versions differ in their choice of target levels: the first (LVC-R0) uses the initial level heights, which are the same as those used in the HB version; the second (LVC-QL) uses quasi-Lagrangian target levels, chosen to follow isentropes on small horizontal scales away from boundaries and to follow the initial height levels on large horizontal scales. These two versions are compared in order to test the hypothesis that frequent remapping to levels at fixed heights might introduce errors that are comparable to those associated with vertical transport in a height based coordinate and so reduce or remove any benefits of a Lagrangian vertical coordinate. The details of the numerical methods, including the specification of the target levels and the remapping to target levels, are presented in Part I.

The HB and LVC versions of the dynamical core are tested on a standard baroclinic instability test case; some details are presented in section 2 below. The flow field generated is sufficiently complex to exercise the numerics of the dynamical core, involving horizontal and vertical transport, the formation of coherent synoptic-scale vortices through a wrapping up process, and the formation of sharp fronts. A range of Lagrangian diagnostics are computed for each version. The diagnostics are motivated and described in section 3. The results are presented and discussed in section 4 and the conclusions are summarized in section 5.

2. Test case and model setup

The dynamical core is nonhydrostatic and does not make the shallow atmosphere approximation. Therefore we use the baroclinic instability test case proposed by Ullrich *et al.* (2014), which is suitable for such models.

Except for the addition of the Lagrangian diagnostics discussed below, the model setup is the same as described in Part I. The horizontal resolution is 192×96 grid points, corresponding to a grid size of 1.88° in longitude and latitude, or about 209 km at the equator. The upper boundary is a rigid lid at $D = 30$ km altitude, and 30 model levels are used, initially distributed in height as suggested by Ullrich *et al.* (2014). The time step is 1200 s. A slight off-centring $\alpha = 0.51$ is used in the semi-implicit scheme, and the nested iterative quasi-Newton solver uses 4 outer iterations and 1 inner iteration per time step. As described in Part I, the solution procedure follows the approach from Section 5 of Wood *et al.* (2014) with two main adaptations. The first is in the definition of Helmholtz coefficients (compare Wood *et al.* (2014) Appendix D with Part I Appendix B). The second is in the choice of solver for the Helmholtz problem (vertical line relaxation combined with a horizontal geometric multigrid method).

For the LVC-QL version the parameters defining the target levels are the same as in Part I. Both LVC versions use conservative parabolic spline remapping to remap mass to target levels, and simple cubic interpolation with a limiter to remap θ (option $\mathcal{M}\text{-}\theta_I$ in the notation of Part I) as this was the only option found to be stable for both LVC-R0 and LVC-QL configurations out to 15 days.

3. Lagrangian diagnostics

A range of diagnostics are used to evaluate the Lagrangian conservation properties of the three model versions.

3.1. Entropy

Conservation of specific entropy following fluid parcels implies that the global integral of entropy should be conserved. The global entropy budget places an important constraint on the climate system (e.g. Goody 2000). It has even been hypothesized that the climate system might adjust so as to maximize its entropy production (e.g. Paltridge 1975; Ozawa *et al.* 2003). Excessive production of entropy has been proposed as a possible cause for systematic cold pole problems in climate models (Johnson 1997). Woollings and Thuburn (2006) diagnosed the entropy sources associated with numerics and scale-selective dissipation in a spectral hydrostatic primitive equation dynamical core during a baroclinic wave life cycle. They found a global average entropy source of around $0.5 \text{ mWm}^{-2}\text{K}^{-1}$, which was a residual between a larger source due to temperature diffusion and an entropy sink associated with numerical dispersion and Gibbs errors. Locally near fronts sources and sinks were much larger, of the order of several $\text{Wm}^{-2}\text{K}^{-1}$.

Here the global entropy \mathcal{S} is diagnosed as

$$\mathcal{S} = c_p \int_V \rho \ln \theta \, dV, \quad (1)$$

where c_p is the specific heat capacity of air at constant pressure, ρ is density, and θ is potential temperature, with the integral approximated as a sum over all model grid cells. In the LVC model versions ρ and θ are stored at cell centres, so the calculation is straightforward. For the HB version θ is stored at the lower and upper faces of grid cells, so it is linearly interpolated to cell centres to calculate the entropy production. All rights reserved.

3.2. Mass below isentropes

In adiabatic flow potential temperature θ is materially conserved. This implies that the mass contained below any given surface of constant θ remains constant. Conservation of the mass below every θ surface implies conservation of all moments of the θ distribution

$$\int \rho \theta^m dV \quad (2)$$

for any exponent m , as well as the entropy S , as noted above.

In realistic flows, fluxes of air and trace constituents across isentropic surfaces are strongly constrained by diabatic heating. Accurately modelling such fluxes is essential for capturing the global atmospheric circulation and transport (e.g. Holton *et al.* 1995), and spurious fluxes due to numerical errors can lead to systematic biases (e.g. Gregory and West 2002; Hardiman *et al.* 2015).

The ability of the three dynamical core versions to preserve the mass below isentropes is evaluated by computing the mass below a chosen set of isentropes at every step during the 15 day test case. The mass below a given θ surface, $\theta = \hat{\theta}$ say, is computed by summing the masses in those cells that have $\theta < \hat{\theta}$. To reduce ‘quantization’ error and ensure a smoother evolution of the diagnostics, θ is expressed as a vertical linear function of mass in each grid cell, giving θ as a piecewise linear function of mass in each column; in this way fractional cells are counted when the range of θ in the cell encompasses $\hat{\theta}$. For the HB ENDGame, constructing the piecewise linear fit is straightforward, since θ is stored at the lower and upper faces of grid cells. For the LVC ENDGame, θ is stored at cell centres. In this case the piecewise linear fit is constructed to be continuous at the lower and upper faces of cells and to minimize the changes in vertical gradient of θ through the column.

For the results shown below, 201 isentropes are used with a uniform 3 K spacing for $\hat{\theta}$ between 210 K and 810 K.

3.3. Unavailable energy

The unavailable energy is defined to be the internal plus potential energy* that the atmosphere would have if it were adiabatically rearranged so as to minimize the internal plus potential energy (Lorenz 1955). It can be shown that the minimizing state is in hydrostatic balance and stably stratified, with θ surfaces horizontal (e.g. Tailleux 2013). Since, by definition, that internal plus potential energy cannot be reduced further by adiabatic motions, none of it is available for conversion to kinetic energy.

Because the unavailable energy is defined in terms of an adiabatic rearrangement of the current atmospheric state, it is determined by the distribution of θ as a function of mass, or, equivalently the distribution of mass as a function of θ . Consequently, the unavailable energy, as well as the total energy, is conserved in adiabatic flow. The total energy is dominated by the unavailable energy, with the available energy (total minus unavailable) of the order of 500 times smaller (Peixoto and Oort 1992). Thus, a small relative error in the conservation of unavailable energy could represent a non-negligible fraction of the climate system energy budget.

If the atmosphere can be assumed to be hydrostatic and the global distribution of θ is known on pressure surfaces then, under some reasonable approximations, the available internal plus potential energy can be estimated in terms of the variance of θ on pressure surfaces (Lorenz 1955). The unavailable energy is then the total internal plus potential energy minus the available internal plus potential energy. Here, however, we are looking for small changes in the unavailable energy, so such simplifying approximations cannot be made. In particular, the current model state cannot be taken to be hydrostatic, even though the energy minimizing state is. Finding this energy minimizing state is the most difficult part of the calculation. We must take as given data the mass distribution of θ , expressed as the masses below a large set of θ surfaces, and directly compute the state that minimizes the internal plus potential energy while preserving

*When the hydrostatic and shallow atmosphere approximations are made the potential energy is proportional to the internal energy, and their sum is often referred to as the potential energy. This article is protected by copyright. All rights reserved.

that mass distribution of θ . An additional complication is the presence of a rigid lid. The calculation of the energy-minimizing state proceeds as follows.

Let $\hat{\theta}_{k+1/2}$, $k = 0, \dots, n$ be the set of θ surfaces used to compute this diagnostic, and let $\mathcal{M}_{k+1/2}$ be the mass below each of these surfaces, calculated as in section 3.2. The range of $\hat{\theta}$ must be sufficiently wide so that $\mathcal{M}_{1/2} = 0$ and $\mathcal{M}_{n+1/2}$ is the mass of the entire model atmosphere. The task is to determine the distance from Earth's centre $r_{k+1/2}$ of each theta surface such that the resulting state is in hydrostatic balance. The boundary conditions are $r_{1/2} = r_{\min}$, where r_{\min} is the minimum r on the bottom boundary, and $r_{n+1/2} = r_{\min} + D$.

A discrete approximation to hydrostatic balance is

$$I_{k+1/2} = c_p \hat{\theta}_{k+1/2} \frac{\Pi_{k+1} - \Pi_k}{r_{k+1} - r_k} + \frac{\Phi_{k+1} - \Phi_k}{r_{k+1} - r_k} = 0, \quad (3)$$

where $I_{k+1/2}$ is the hydrostatic imbalance at level $k + 1/2$. Here, the Exner pressure Π_k is given by

$$\Pi_k = \left(\frac{R \rho_k \theta_k}{p_{00}} \right)^{\kappa/(1-\kappa)}, \quad (4)$$

where R is the gas constant for dry air, $\kappa = R/c_p$ and $p_{00} = 10^5$ Pa is a constant reference pressure. The potential temperature at integer levels θ_k is obtained by linear interpolation in height between $\hat{\theta}_{k-1/2}$ and $\hat{\theta}_{k+1/2}$, and the density ρ_k is obtained from

$$\rho_k = \frac{\mathcal{M}_{k+1/2} - \mathcal{M}_{k-1/2}}{V_{k+1/2} - V_{k-1/2}}, \quad (5)$$

where $V_{k+1/2}$ is the volume of the model domain below $r = r_{k+1/2}$. When the domain bottom boundary is flat, as it is here, $V_{k+1/2}$ is given simply by

$$V_{k+1/2} = 4\pi(r_{k+1/2}^3 - r_{\min}^3)/3. \quad (6)$$

More generally, if there is orography, $V_{k+1/2}$ must be obtained by

$$V_{k+1/2} = \int dA \max \left((r_{k+1/2}^3 - r_{\text{surf}}^3)/3, 0 \right), \quad (7)$$

where dA is the element of area on the unit sphere and r_{surf} is the value of r at the Earth's surface at that horizontal location.

Finally r at integer levels is given by $r_k = (9(r_{k-1/2} + r_{k+1/2}) - (r_{k-3/2} + r_{k+3/2}))/16$ or, near the top and bottom boundaries, by $r_k = (r_{k-1/2} + r_{k+1/2})/2$, and Φ_k is the geopotential evaluated at $r = r_k$.

Thus, given an estimate for the values of $r_{k+1/2}$, the imbalance at each level can be calculated. We wish to determine the $r_{k+1/2}$ that solve the system of coupled nonlinear equations

$$I_{k+1/2} = 0, \quad k = 1, \dots, n-1. \quad (8)$$

The system is solved by an approximate Newton method. The Jacobian of the system is dominated by the dependence of Π_k on ρ_k and the dependence of ρ_k on $V_{k\pm 1/2}$. Retaining only these terms in the Jacobian leads to a tridiagonal problem for the Newton update. Convergence of the approximate Newton method is slow unless a good first guess is provided, so 25 iterations are used. The method may also fail to converge if only a very small but non-zero mass is contained below some θ surface, so if this condition is detected the layer in this article is protected by copyright. All rights reserved.

It was found that to compute accurate and noise-free estimates of the unavailable energy a large number of θ surfaces were needed. As the available energy is about 500 times smaller than the unavailable energy, the relative error in the available energy will be ~ 500 times greater. To observe the non-conservation error in available energy its relative error needs to be small enough, which requires a relative error of order 10^{-5} in unavailable energy. The calculation of unavailable energy converges at roughly first order in the number of θ layers used. For the results shown below the unavailable energy was computed using 1001 θ surfaces uniformly distributed between the minimum and maximum θ in the domain.

3.4. Potential vorticity

Potential vorticity is materially conserved in adiabatic and frictionless flow, and is approximately materially conserved throughout much of the real atmosphere on time scales of a few days. Potential vorticity is dynamically important because it controls the balanced, vortical, component of the atmospheric flow, which dominates the evolution on synoptic scales (Hoskins *et al.* 1985). Mixing of potential vorticity can lead to the formation of sharp gradients, which may then act as transport barriers (e.g. McIntyre and Palmer 1984; Holton *et al.* 1995). There is also a tendency for long-lived constituents to become correlated with each other and with potential vorticity, which can be useful for diagnosing transport and mixing from observations and models (e.g. Newman *et al.* 1988). See also the discussion in Whitehead *et al.* (2015) and Saffin *et al.* (2016).

Here the full form of the Ertel's potential vorticity is calculated:

$$Q = \frac{\boldsymbol{\zeta} \cdot \nabla \theta}{\rho}, \quad (9)$$

where $\boldsymbol{\zeta} = \nabla \times \mathbf{u} + 2\boldsymbol{\Omega}$ is the absolute vorticity vector, with $\boldsymbol{\Omega}$ the Earth's angular velocity vector. The hydrostatic and shallow atmosphere approximations, which are often made in diagnosing potential vorticity from models or analyses, are not made here.

On the staggered grid used by the model, the three components of $\nabla \times \mathbf{u}$ are most naturally calculated at different grid locations; the same is true for the components of $\nabla \theta$. Therefore, some averaging is unavoidable in order to calculate the full potential vorticity. The additional terms that arise from converting gradients along model levels to gradients at constant height ('bent terms', see Part I equation (8)) also introduce averaging. To minimize the effects of this averaging, the potential vorticity is calculated at the natural vertical vorticity points of the C-grid, that is, at points staggered in the north-south direction relative to u and in the east-west direction relative to v . In this way the contribution involving the vertical vorticity, which usually dominates the potential vorticity, experiences the least averaging.

It is worth noting that, because the bottom boundary is flat in the Ullrich *et al.* (2014) test case used here, the bent terms will vanish for the HB model version and will generally be very small for the LVC-R0 version. Therefore the LVC-QL version is expected to be most strongly affected by the averaging of the bent terms.

3.5. Tracers

Johnson *et al.* (2000, 2002) noted that the equivalent potential temperature θ_e is approximately materially conserved even in the presence of condensation and evaporation. In most atmospheric models θ_e is not directly predicted, but is a derived quantity obtained from other, predicted, quantities such as θ and specific humidity. Johnson *et al.* (2000, 2002) therefore proposed a test of numerical models in which the evolution of the derived quantity θ_e is compared with the evolution of a passive tracer initialized with the same distribution as θ_e . They argued that agreement between θ_e and the corresponding tracer would be a valuable check on the *consistency* between the tracer and the model. All rights reserved.

In a similar way, potential vorticity is a derived quantity in most atmospheric models. Whitehead *et al.* (2015) proposed an analogous test in which the evolution of potential vorticity is compared with the evolution of a tracer initialized to have the same distribution. See the discussion in Johnson *et al.* (2000, 2002); Whitehead *et al.* (2015) regarding the advantages of such consistency.

Here, two tracers were included in the test case for the three dynamical core versions. The first one was initialized to equal potential temperature: $T_\theta = \theta$; the second one was initialized to equal potential vorticity: $T_{PV} = Q$. After the initial time the tracer mixing ratio is materially conserved:

$$\frac{DT_\theta}{Dt} = 0, \quad \frac{DT_{PV}}{Dt} = 0 \quad (10)$$

The tracers are stored at the cell-centre pressure points in all three model versions. They are advected using a semi-Lagrangian advection scheme with cubic Lagrange interpolation. This is very similar to advection of the other model prognostic variables; however, in contrast to the treatment of the dynamical θ , the tracer advection does not use a limiter to prevent overshoots in the interpolation. In the LVC model versions the tracers are vertically remapped to target levels at the end of each time step using using cubic Lagrange interpolation (see Part I), again without any limiter.

3.6. Air parcel trajectories

Johnson *et al.* (2000, 2002); Whitehead *et al.* (2015) proposed comparing derived model fields that should be materially conserved with tracers stored on the model grid and advected by an Eulerian transport scheme. An alternative way to estimate the evolution of a materially conserved quantity is to compute purely Lagrangian air parcel trajectories. Each parcel may carry a number of labels corresponding to the materially conserved quantities of interest, for example the initial potential vorticity. At subsequent times the potential vorticity label may be compared with the potential vorticity derived from model prognostic fields and interpolated to the parcel's current location. Agreement between the two calculations gives another test of consistency between different calculations of Lagrangian conservation.

Here we combine this trajectory idea with the passive tracer idea of Johnson *et al.* (2000, 2002); Whitehead *et al.* (2015) by making a three-way comparison between (i) directly predicted or derived model dynamical fields interpolated to the air parcel location, (ii) passive tracers interpolated to the air parcel location, and (iii) air parcel trajectory labels.

For each model version a set of 44310 air parcels trajectories was computed during the 15 day integration. The trajectory starting points were distributed uniformly over the Northern Hemisphere on each model level at the heights of the pressure points. The trajectories were stepped forward using a Crank-Nicolson scheme using the same $\Delta t = 1200$ s as the main model integration, and using linear interpolation in space to determine the departure point and arrival point velocity. Thus, the calculation is similar to that used to compute the trajectories in the semi-Lagrangian advection scheme; the key differences are that here the departure point rather than the arrival point is known, and neither the arrival point nor the departure point can be assumed to coincide with a grid point. The work of McDonald (1986) and McDonald and Bates (1987) shows that a similar centred-in-time trajectory calculation using bilinear interpolation of velocities in space gives good accuracy.

Each trajectory carries two labels: its initial potential temperature and its initial potential vorticity, both obtained by interpolating to the trajectory starting point. At regular intervals, six values are output for each trajectory:

1. the initial value of θ ;
2. the current value of θ interpolated to the trajectory location;
3. the current value of T_θ interpolated to the trajectory location;
4. the initial value of Q ;
5. the current value of Q interpolated to the trajectory location;

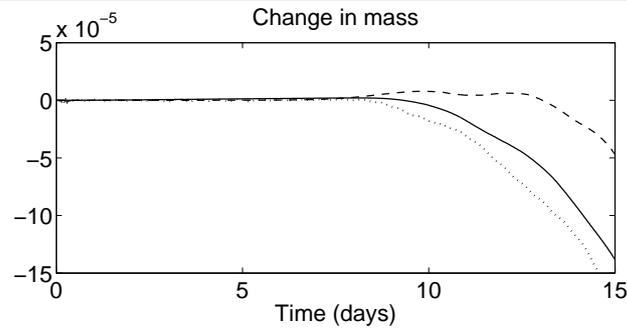


Figure 1. Fractional change in mass for the three model versions: HB (solid line); LVC-R0 (dashed line); LVC-QL (dotted line).

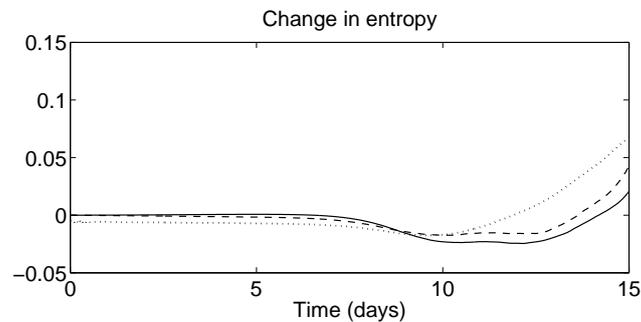


Figure 2. Change in entropy per unit mass ($\text{Jkg}^{-1}\text{K}^{-1}$) for the three model versions: HB (solid line); LVC-R0 (dashed line); LVC-QL (dotted line).

6. the current value of T_{PV} interpolated to the trajectory location.

Linear interpolation is used to interpolate fields to the trajectory location; this is sufficiently accurate since the errors do not accumulate over time.

4. Results and discussion

Baroclinic instability is triggered by a perturbation in the initial conditions. The disturbance remains linear and grows roughly exponentially for the first few days, becoming noticeable in surface plots around day 6. Subsequently the disturbance becomes strongly nonlinear and wraps up into a series of cyclones and anticyclones. Sharp surface fronts begin to form around day 8, and the flow continues to become more complex and turbulent up to the end of the run at day 15. See Ullrich *et al.* (2014) and Part I for more details.

4.1. Entropy

Figure 7 of Part I shows the fractional change in total entropy over the 15 days of the test for a number of configurations of the dynamical core. For the three versions tested here the total entropy in fact decreases, by about 1.5×10^{-4} for the HB and LVC-QL version and by about 0.5×10^{-4} for the LVC-R0 version. However, these entropy losses are dominated by the imperfect conservation of mass; this is clear from Figure 1, which shows the fractional change in mass for the three dynamical core versions. Interestingly the conservation of mass is noticeably better for the LVC-R0 version than for either the HB or LVC-QL versions, and this pattern also holds for the other remapping options tested in Part I (see their Figure 6). It is not clear why this is the case.

To compensate for the imperfect mass conservation, Figure 2 shows the change in the entropy *per unit mass* for the three model versions. To put these values in context, a typical global mean value for the entropy per unit mass is $5.8 \times 10^3 \text{ Jkg}^{-1}\text{K}^{-1}$. For the LVC-QL version there is a small but non-zero change at the end of the first time step. This coincides with a significant change in model level locations from near their initial heights to their QL target levels when remapping first occurs. Subsequently all three versions show a slight loss of entropy by day 8 and 10, followed by a systematic increase.

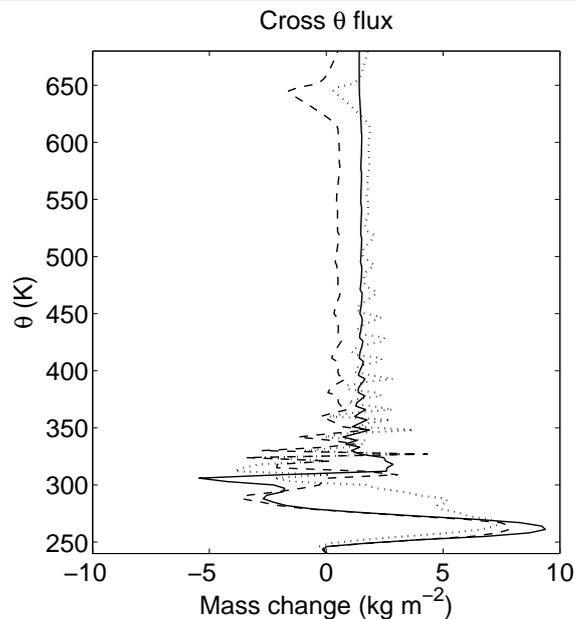


Figure 3. Mass flux across θ surfaces accumulated over the 15 day integration for HB (solid), LVC-R0 (dashed), and LVC-QL (dotted) model versions. A mass loss below a given θ surface is plotted as a positive value to indicate an upward cross-isentrope flux.

For comparison with Woollings and Thuburn (2006), the rates of change of entropy over the last seven days, expressed in appropriate units, are $0.35 \text{ mWm}^{-2}\text{K}^{-1}$ for HB, $0.72 \text{ mWm}^{-2}\text{K}^{-1}$ for LVC-R0, and $1.14 \text{ mWm}^{-2}\text{K}^{-1}$ for LVC-QL. These are certainly comparable to the values of around $0.5 \text{ mWm}^{-2}\text{K}^{-1}$ found by Woollings and Thuburn, despite the use of very different numerical methods and dissipation mechanisms, and a somewhat different test case.

4.2. Mass below isentropes

The time evolution of the mass below a set of 201 isentropes was diagnosed, as described in section 3.2. A convenient way to visualize the results is to interpret the change in mass below an isentrope as a mass flux across that isentrope, with a decrease in the mass below corresponding to an upward flux. The results for the three dynamical core versions are shown in Figure 3.

There is an important caveat to the interpretation of the mass change as a mass flux when the total model mass is not conserved. Thus, the non-zero mass changes for isentropes in the range $400 \text{ K} < \theta < 600 \text{ K}$ are indicative of imperfect mass conservation at lower altitudes rather than a lack of Lagrangian conservation in the stratosphere.

Interestingly, all three model versions show an upward mass flux of around 7 to 9 kg m^{-2} across isentropes around 260 K to 270 K . There is a smaller downward mass flux across higher isentropes, around 290 K for the HB and LVC-R0 versions and around 310 K for the LVC-QL version. The timing of these fluxes, which become significant from day 8 onwards, as well as their location in θ -space, is consistent with the idea that they are associated with numerical diffusion related to the formation of sharp fronts near the Earth's surface. These results may be compared with those of Woollings and Thuburn (2006), who found a significant downward flux, rather larger in amplitude, across θ surfaces around 290 K to 300 K .

Expressed as a mass flux per unit time, assumed to occur over the last seven days of the integration, the peak cross-isentrope mass flux found here is around $1.5 \times 10^{-5} \text{ kg m}^{-2}\text{s}^{-1}$. This is quite small compared to estimates of the real atmosphere cross-isentrope mass flux near the Earth's surface of around $10^{-4} \text{ kg m}^{-2}\text{s}^{-1}$ (e.g. Juckes 2001). However, when the HB version was initialised with an initial perturbation in both hemispheres and run for 30 days, giving a more realistic population of weather systems, we found the peak cross-isentrope mass flux per unit time to be comparable to the real atmosphere (around $1.6 \times 10^{-4} \text{ kg m}^{-2}\text{s}^{-1}$ over the last 15 days). This article is protected by copyright. All rights reserved.

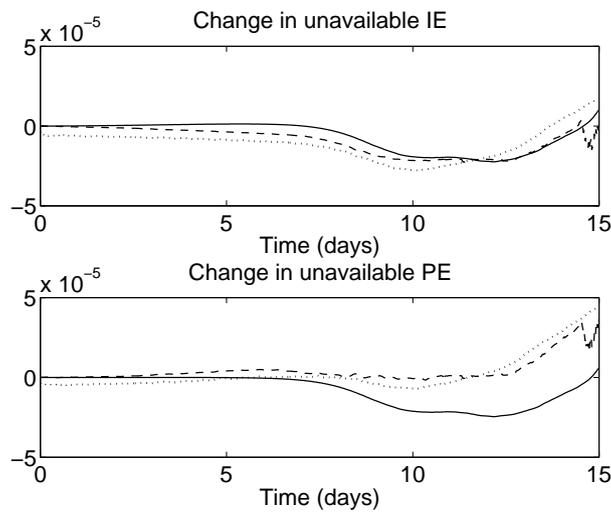


Figure 4. Fractional change in unavailable internal energy and unavailable potential energy *per unit mass* for HB (solid), LVC-R0 (dashed), LVC-QL (dotted) dynamical core versions.

Both LVC versions show a downward peak in the cross isentrope flux around 640 K. This corresponds to the uppermost level in the model; the highest value of θ in the domain is approximately 679 K. Calculating the mass below isentropes both before and after remapping shows that this downward flux is associated with the remapping of the θ field, which reduces to a linear interpolation or extrapolation in the uppermost model layer. Further tests (not shown) indicate that the feature is sensitive to the limiter used with the remapping of θ . The limiter is applied even when the remapped θ is extrapolated at the top of the model, effectively imposing a constant (rather than linear) extrapolation.

4.3. Unavailable energy

As for entropy, the changes in unavailable energy are dominated by the imperfect conservation of mass. Therefore in Figure 4 we show the fractional change over time of the unavailable energy per unit mass, decomposed into internal energy and potential energy contributions, for the three dynamical core versions. The relative changes are small, a few times 10^{-5} , and are of similar magnitude for the three versions. The small jump in the time series for the LVC-R0 version around day 14.5 is associated with the merging of isentropic layers needed to ensure convergence of the calculation of the minimum energy state, as discussed in section 3.3.

4.4. Consistent evolution of potential temperature

Figure 5 shows scatter plots of θ versus initial θ , T_θ versus initial θ , and T_θ versus θ at the trajectory locations at day 15 for the three model versions (as defined in Subsection 3.6, “initial θ ” denotes the label carried by a Lagrangian parcel interpolated to the starting point of the parcel trajectory). The plots are restricted to the range $\theta < 350$ K because at higher θ values the dynamics remains very simple and, with one exception noted below, the correlations are extremely good.

Imperfect correlations between θ and initial θ and between T_θ and initial θ are clearly visible. The scatter appears to be reduced slightly for the LVC-QL version compared with the HB and LVC-R0 versions, and this is borne out by the root mean square differences shown in Table 1. Note also that the magnitude of the root mean square differences is quite small, around 1–1.5 K; the vast majority of points on the scatter plots do lie very close to the diagonal.

The scatter is much smaller for the T_θ versus θ scatter plots than for the other scatter plots. This is to be expected, since the two fields are predicted using very similar algorithms. The agreement is not perfect because (i) the advection of θ uses a simple limiter to prevent overshoots whereas the advection of T_θ does not; (ii) θ is updated along with the density and wind fields as part of the iterative nonlinear dynamics procedure (see Part II) whereas T_θ is updated after the dynamical step once the winds are known; (iii) in the HB version

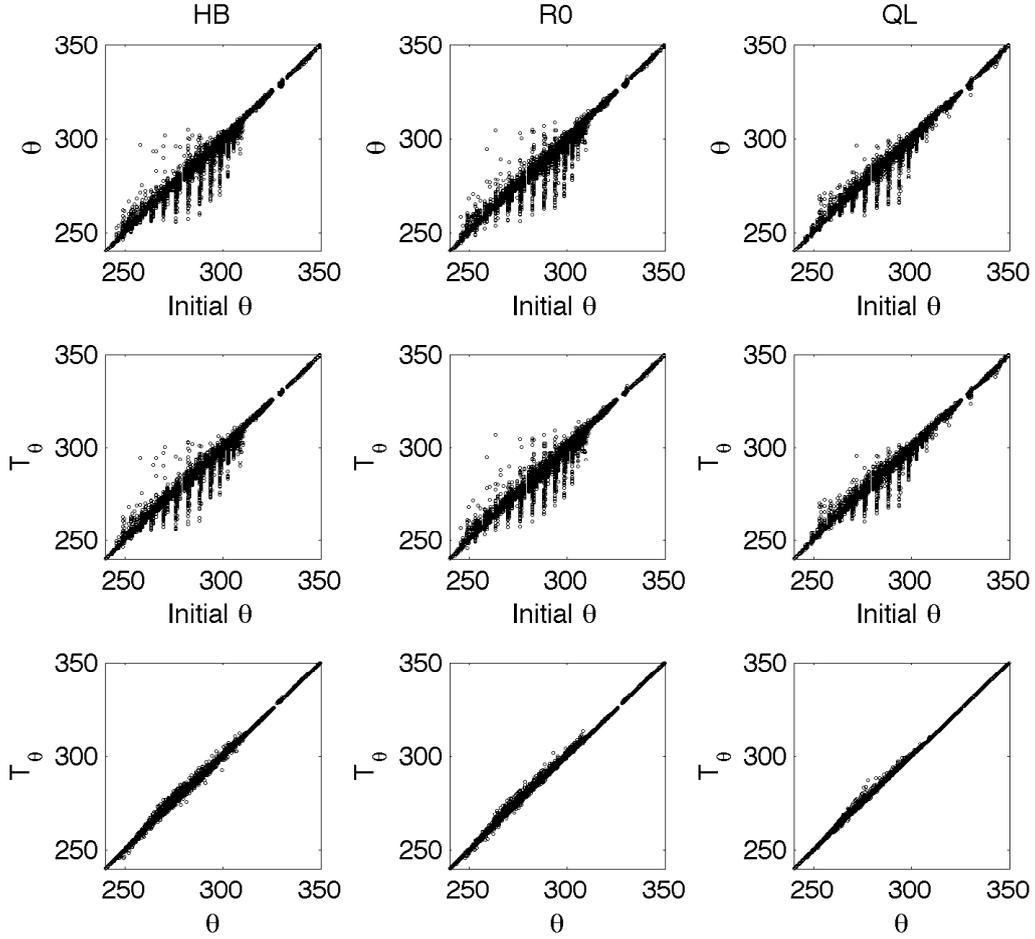


Figure 5. Scatter plots comparing θ , T_θ , and initial θ , (K) sampled at the trajectory locations at day 15. Left column: HB version; centre column: LVC-R0 version; right column LVC-QL version. Only trajectories with initial $\theta < 350$ K are shown.

θ is stored at the same locations as w , staggered vertically with respect to T_θ . Again, the scatter is reduced slightly for the LVC-QL version compared with the other versions (Table 1).

Table 1. Root mean square differences between different pairs of potential temperature variables at day 15, for parcels with initial $\theta < 350$ K. In the table, θ_0 indicates initial θ . The units are in K.

	Height-based	LVC (R0)	LVC (QL)
$\theta - \theta_0$	1.50	1.56	1.08
$T_\theta - \theta_0$	1.31	1.41	0.99
$T_\theta - \theta$	0.36	0.34	0.22

The three quantities θ , T_θ and initial θ are generally very well correlated for $\theta > 350$ K. The exception is at the very top of the model, where the LVC versions, particularly the LVC-R0 version, have some significant scatter. Figure 6 shows T_θ versus θ for the full model domain at day 15 for the LVC-R0 version. This and other diagnostics show that the errors at the top of the model are primarily in the θ field, which shows a drift to lower values. This error is consistent with the spurious cross-isentrope descent near the model top noted in section 4.2, associated with the application of the limiter in the remapping of θ .

4.5. Consistent evolution of potential vorticity

Scatter plots comparing Q , T_{PV} , and initial Q behave differently at different altitudes (as defined in Subsection 3.6, “initial Q ” denotes the label carried by a Lagrangian parcel interpolated to the starting point of the parcel trajectory). Figure 7 shows the scatter plots for those trajectories whose initial potential vorticity is 60 K^2 . There is clearly greater scatter for the LVC model versions, particularly LVC-R0,

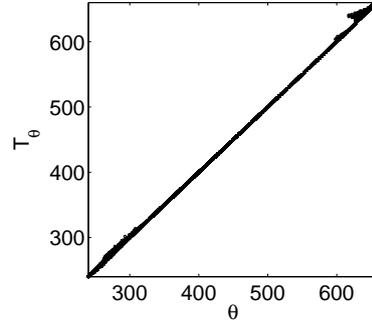


Figure 6. Scatter plot of T_θ versus θ (K) at day 15 for the LVC-R0 version.

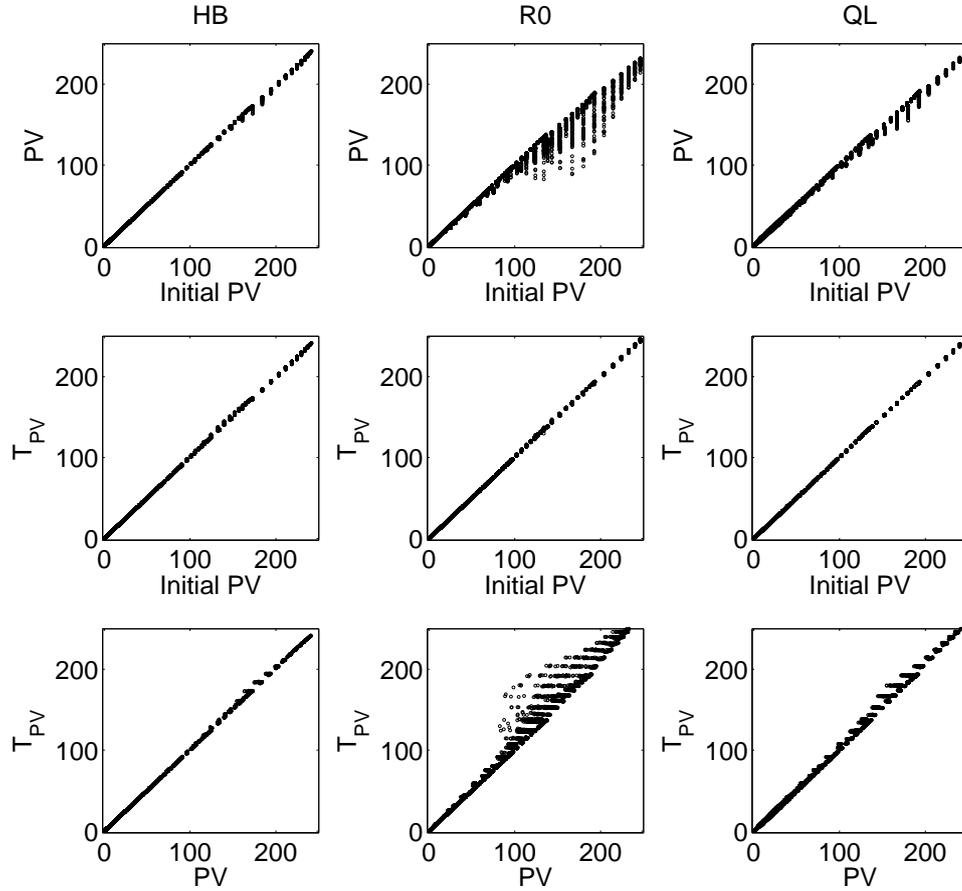


Figure 7. Scatter plots of potential vorticity Q , potential-vorticity-like tracer T_{PV} , and initial potential vorticity, sampled at the trajectory locations at day 15. Left column: HB version; centre column: LVC-R0 version; right column LVC-QL version. Only trajectories with initial $\theta > 360$ K are shown. The units are $10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

than for the HB version, with the greatest errors apparently in Q itself rather than T_{PV} or initial Q . Examining different ranges of θ shows that the large scatter occurs for $\theta > 600$ K. This is also where the LVC model versions have significant errors in their treatment of θ (sections 4.2, 4.4), suggesting that the potential vorticity errors might be a direct result of the θ errors. To test this hypothesis Q was recomputed using T_θ in place of θ in (9). The scatter was greatly reduced, confirming the hypothesis.

Figure 8 shows the potential vorticity scatter plots for those trajectories with initial θ between 315 K and 360 K. For the T_{PV} versus initial Q scatter plots the scatter appears to be somewhat reduced for the LVC-QL version compared with the HB and LVC-R0 versions. This impression is confirmed by the root mean square differences (Table 2). On the other hand, for Q versus initial Q and T_{PV} versus Q the LVC-QL version shows increased scatter compared with the other versions, particularly for large values of Q . A plausible explanation for this increased scatter is the contribution to Q from bent terms, which, as noted in section 3.4, are greatest in the LVC-QL version. Figure 9 shows a scatter plot of $|Q - T_{PV}|$ versus the slope of model levels at

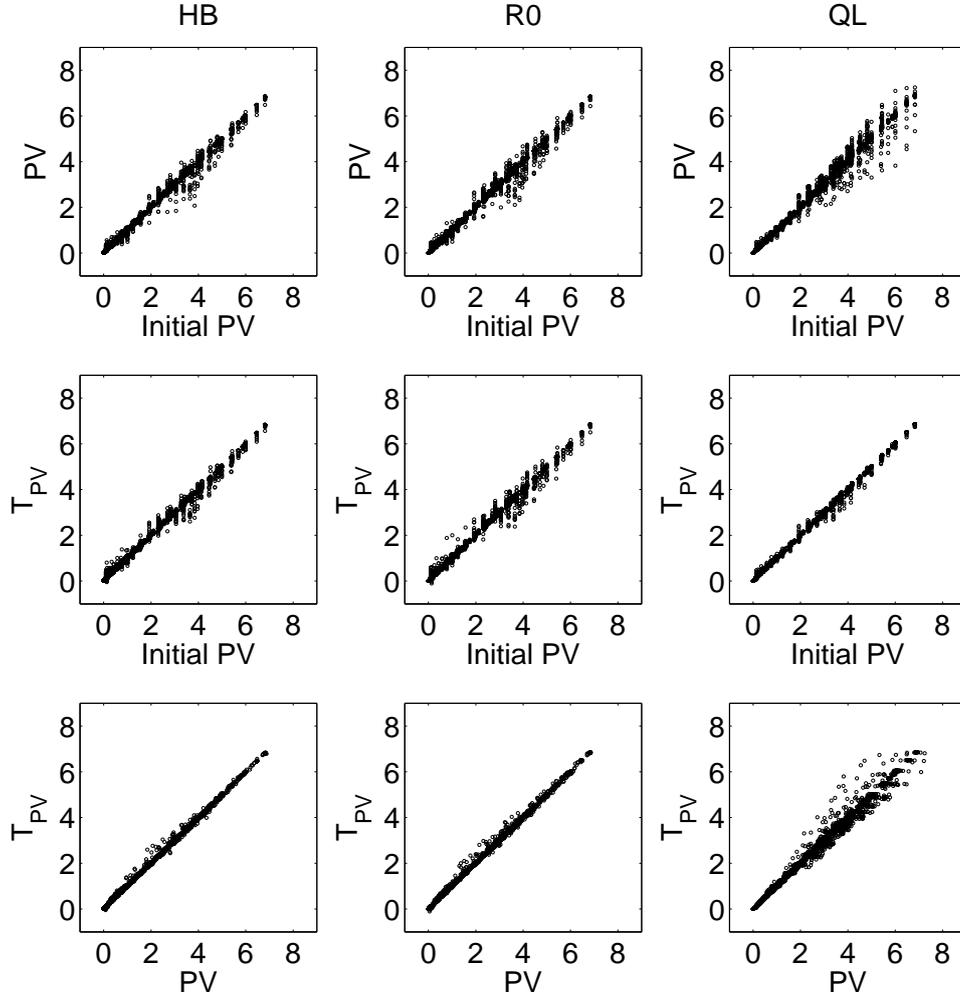


Figure 8. Scatter plots of potential vorticity Q , potential-vorticity-like tracer T_{PV} , and initial potential vorticity, sampled at the trajectory locations at day 15. Left column: HB version; centre column: LVC-R0 version; right column: LVC-QL version. Only trajectories with $315 \text{ K} < \theta < 360 \text{ K}$ initially are shown. The units are $10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

day 15 for the same set of trajectories for the LVC-QL version. The best fit line to the data has a gradient $0.225 \times 10^{-3} \text{ kg K m}^2 \text{ s}^{-1}$, and the correlation coefficient is 0.498, indicating that there is indeed a link between model level slopes and the degradation of Lagrangian conservation of potential vorticity.

Table 2. Root mean square differences between different pairs of potential vorticity variables for parcels with $315 \text{ K} < \theta < 360 \text{ K}$. The units are in $10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

	Height-based	LVC (R0)	LVC (QL)
$Q-Q_0$	0.036	0.041	0.055
$T_{PV}-Q_0$	0.029	0.035	0.019
$T_{PV}-Q$	0.015	0.017	0.049

To directly visualize the comparison between Q and T_{PV} , Figure 10 shows maps of both quantities on the 330 K isentrope at day 12 for the three dynamical core versions. The main point to note is that, at this model resolution and this time range, the agreement between Q and T_{PV} is very good. In fact, the differences between the three versions are noticeably larger than the difference between Q and T_{PV} for any given version. The differences between Q and T_{PV} do grow gradually at later times.

The most striking departures from perfect correlations in the potential vorticity variables occur for those trajectories with initial θ less than 315 K. Figure 11 shows scatter plots at day 15 for parcels with initial θ less than 315 K for the three

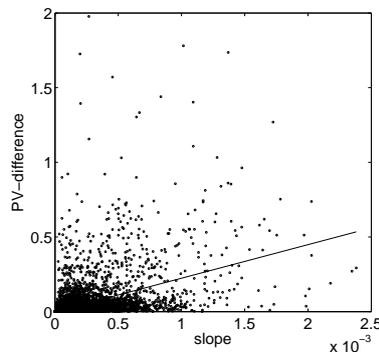


Figure 9. Scatter plot of $|Q - T_{PV}|$ versus model level slope, sampled at the trajectory locations at day 15 for the LVC-QL version. Only trajectories with $315 \text{ K} < \theta < 360 \text{ K}$ initially are shown. The units for $|Q - T_{PV}|$ are $10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

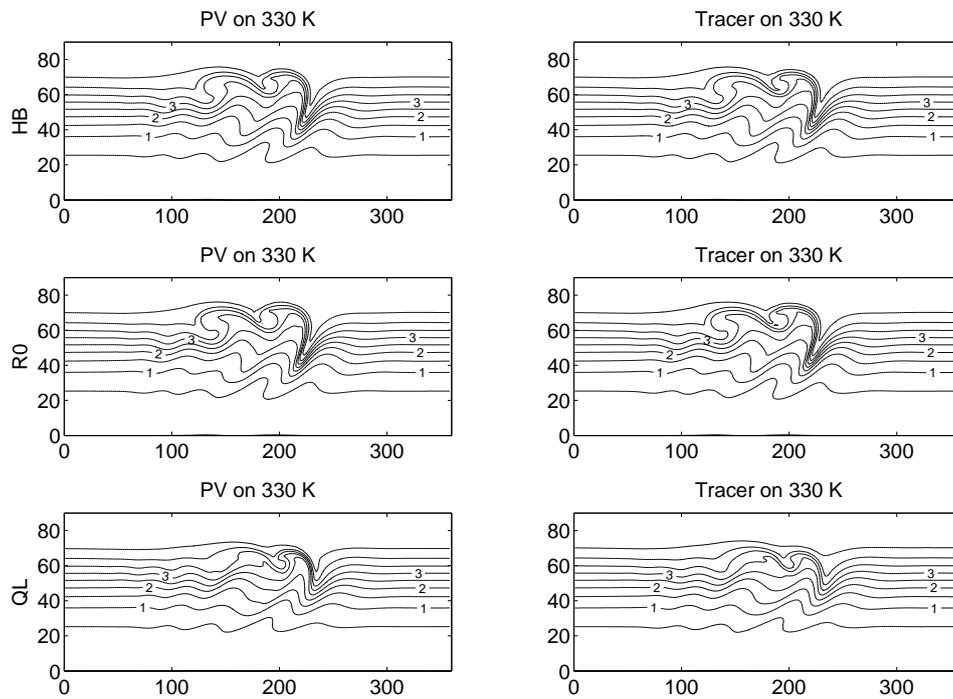


Figure 10. Longitude-latitude maps of Q (left) and T_{PV} (right) on the 330 K isentrope at day 12. Top: HB version; middle: LVC-R0 version; bottom: LVC-QL version. The contour interval is $0.5 \times 10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

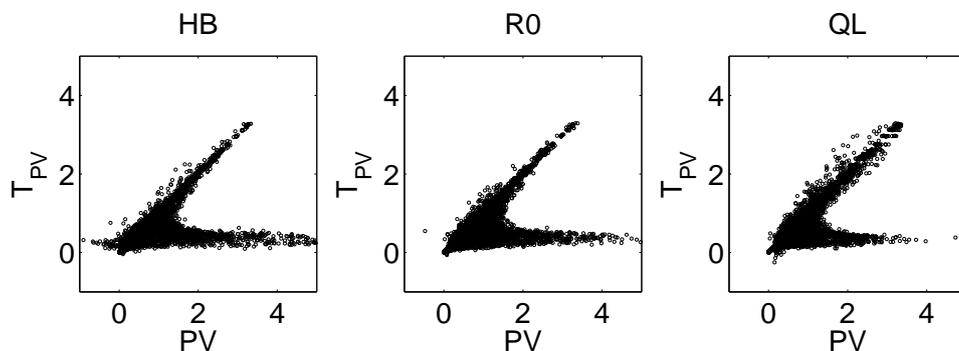


Figure 11. Scatter plots of potential-vorticity-like tracer T_{PV} versus potential vorticity Q , sampled at the trajectory locations at day 15. Left: HB version; centre: LVC-R0 version; right: LVC-QL version. Only trajectories with initial $\theta < 315 \text{ K}$ are shown.

dynamical core versions. For all three versions, some parcels with small values of T_{PV} have acquired very large values of Q , indicating large Lagrangian non-conservation of Q .

Longitude-latitude maps of Q and T_{PV} on the 300 K isentrope at day 15 (Figure 12) show that these large Q values appear along a narrow band between 10° and 20° latitude and around 15° to 30° longitude. Comparison with maps of surface θ (Figure 13)

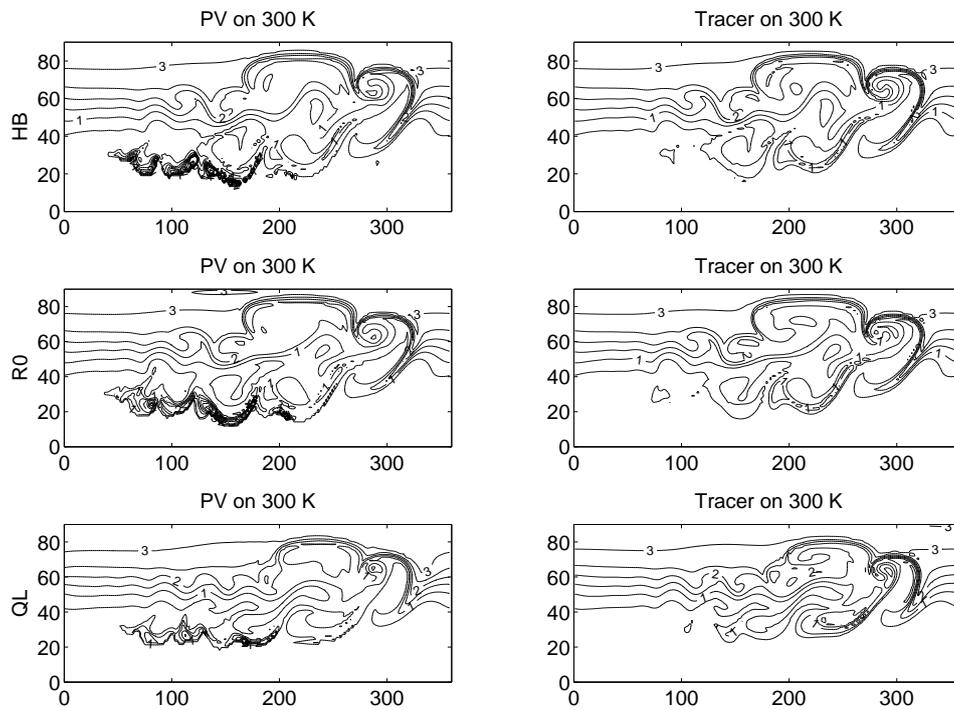


Figure 12. Longitude-latitude maps of potential vorticity Q and corresponding tracer T_{PV} on the 300 K isentrope at day 15. Top: HB version; middle: LVC-R0 version; bottom: LVC-QL version. The contour interval is $0.5 \times 10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

shows that the large Q values occur very close to the grounding line of the 300 K isentrope and mainly in the region of strong surface fronts.

Figure 14 shows Q and T_{PV} on an equator-to-pole vertical slice at 150° at day 15 for the HB version. A downward intrusion of high potential vorticity values from the stratosphere into the troposphere is captured in both fields, indicating a tropopause fold (e.g. Holton *et al.* 1995). The large Q values on the 300 K isentrope are clearly visible at 20° latitude (compare Figure 12). There are also large values between 40° and 60° latitude, which occur at lower θ values, around 260 K. At this time these large values of Q are confined to the lowest 2 km of the atmosphere.

Previous studies of frontal formation have shown that the appearance of such large values of Q near surface fronts is to be expected (e.g. Whitehead *et al.* 2015). In quasi-geostrophic theory the effect of surface variations in θ can be interpreted in terms of a surface δ -function contribution to the potential vorticity (Bretherton 1966). The idea generalizes to semi-geostrophic theory (e.g. Cullen and Purser 1984) and to the primitive equations (Schneider *et al.* 2003). In the semi-geostrophic theory of frontal formation a discontinuity in the surface θ can form in finite time (Hoskins and Bretherton 1972); when the Lagrangian form of the equations is solved for later times the surface front extends into the fluid interior as the surface potential vorticity δ -function is lifted (Cullen and Purser 1984). Large potential vorticity values also appear as fronts form in Eulerian primitive equation models (e.g. Nakamura and Held 1989), though here numerical diffusion is thought to play a role in lifting the potential vorticity δ -function and spreading it to the grid scale. A similar phenomenon occurs in numerical simulations of wake formation for flow past mountains (Schneider *et al.* 2003).

With these previous studies in mind, the appearance of large values of Q near the Earth's surface, as seen in Figures 11, 12, and 14, should not be interpreted as a failure of Lagrangian conservation in any of the dynamical core versions. Rather, it appears to be a legitimate finite-resolution approximate representation of adiabatic frictionless front formation. Note, by the way, that the tracer T_{PV} cannot be thought of as having a surface δ -function contribution, so we should not expect to see the appearance of large values of T_{PV}

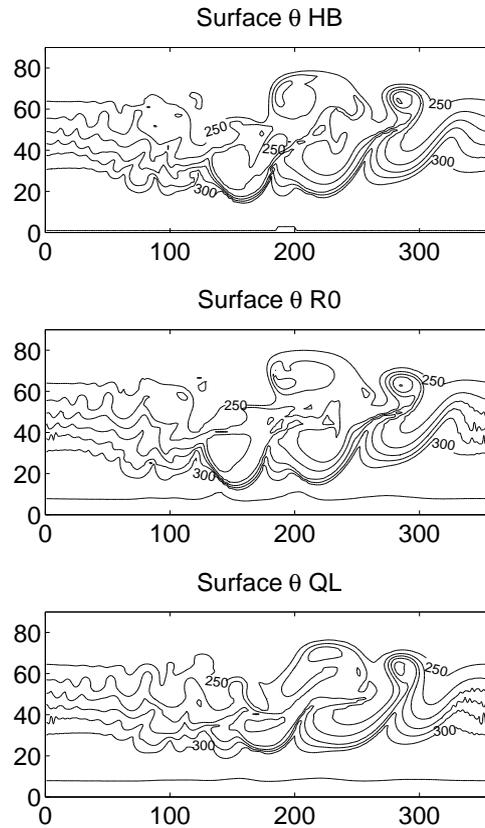


Figure 13. Surface θ at day 15 for HB version (top), LVC-R0 version (middle) and LVC-QL version (bottom). The contour values range from 250 K to 310 K in steps of 10 K.

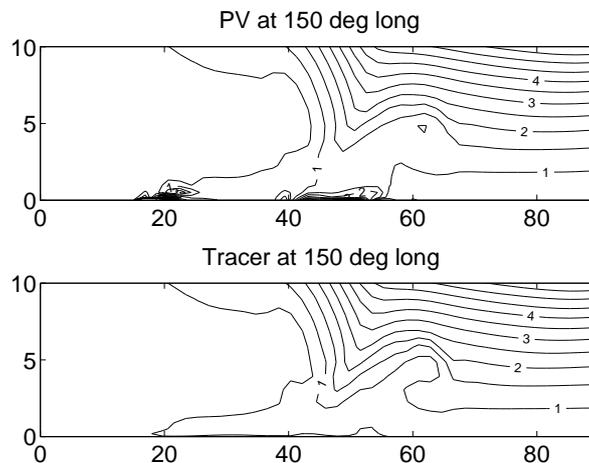


Figure 14. North-south vertical slice at 150° longitude day 15 showing Q and T_{PV} for the HB version. The contour interval is $0.5 \times 10^{-6} \text{ kg K m}^2 \text{ s}^{-1}$.

5. Conclusions

Using a standard baroclinic instability test case, we have compared the Lagrangian conservation properties of three versions of a nonhydrostatic global dynamical core: one using a height based vertical coordinate (HB), one using a Lagrangian vertical coordinate with resetting of levels after every step to their initial heights (LVC-R0), and one using a Lagrangian vertical coordinate with resetting of levels after every step to quasi-Lagrangian target levels (LVC-QL). The three versions use very similar semi-implicit semi-Lagrangian numerical methods based on those used in ENDGame (Wood *et al.* 2014), so that the effects of the Lagrangian vertical coordinate can be isolated as much as possible.

None of the dynamical core versions is inherently mass conserving, and all three show a fractional mass loss of order 10^{-4} over the 15 days of the test case. The mass loss for the LVC-R0 version is about one third of that for the HB and LVC-QL versions; the reasons for this difference are not clear. Changes in global integrals of entropy and unavailable energy are dominated by the change in mass. The fractional changes in the entropy per unit mass and in the unavailable energy per unit mass are of the order of a few times 10^{-5} , and none of the dynamical core versions is clearly better or worse than the others.

All three dynamical core versions show an upward mass flux of around 7 to 9 kg m^{-2} across isentropes around 260 K to 270 K , and a smaller downward mass flux across slightly higher isentropes. It is likely that these fluxes are associated with numerical diffusion near sharp θ gradients at fronts. The magnitude of these low altitude fluxes is similar in the three model versions, and no model version is clearly better or worse than the others. In the LVC model versions there are errors associated with the application of a limiter at the top model level in the remapping of θ ; these errors show up as a spurious downward mass flux across isentropes around 600 to 650 K . Finally, the diagnosis of cross-isentrope mass fluxes is somewhat obscured by the lack of exact mass conservation in the dynamical cores.

Scatter plots comparing θ , initial θ , and T_θ show some small but clear improvements in the consistency between these variables for the LVC-QL version compared with the HB and LVC-R0 versions. An exception to this occurs at the top model level, where the errors in remapping θ noted above degrade the scatter plots for the LVC versions.

These errors in remapping θ at the top level also degrade the potential vorticity scatter plots for parcels near the model top for the LVC model versions, particularly the LVC-R0 version. For parcels in the range $315 \text{ K} < \theta < 350 \text{ K}$ all three model versions show very good consistency between Q and T_{PV} . The consistency between T_{PV} and initial Q is improved for the LVC-QL version compared with the HB and LVC-R0 versions. On the other hand, the consistency between Q and initial Q and between T_{PV} and Q is degraded for the LVC-QL version. A likely cause of this degradation is the contribution to the diagnosed potential vorticity Q from bent terms, which involves significant averaging and therefore reduced accuracy. This contribution is much greater in the LVC-QL model version, which has significant model level slopes, than in the HB and LVC-R0 model versions.

For fluid parcels close to the surface ($\theta < 315 \text{ K}$) there is a strong Lagrangian source of potential vorticity in the vicinity of surface fronts in all three model versions. This source is consistent with previous numerical simulations of frontal formation and with theoretical ideas involving the lifting of a surface potential vorticity δ -function into the fluid interior. It should not be interpreted as a failure of Lagrangian conservation in the models.

In summary, in Part I and the present paper (Part II) we have used a standard baroclinic instability test case to make a clean comparison between versions of a dynamical core using a height based vertical coordinate and using a Lagrangian vertical coordinate but with otherwise almost identical numerical methods. Part I shows that the Lagrangian vertical coordinate versions are considerably cheaper computationally than the height coordinate version, and have comparable or even slightly better global conservation properties. However, the Lagrangian vertical coordinate versions are less robust, even when a range of mitigating measures are taken. Part II shows, for some diagnostics, a small but clear improvement in the consistency between dynamical quantities, tracers, and trajectories for the Lagrangian vertical coordinate version with quasi-Lagrangian target levels. An exception to this, however, is that potential vorticity appears to be calculated less accurately because of its greater model level slope.

Acknowledgement

We thank two anonymous reviewers for their constructive comments on an earlier version of this paper. This work was funded by the Natural Environment Research Council Grant NE/H006834/1.

References

- Bretherton FP. 1966. Critical layer instability in baroclinic flows. *Q. J. Roy. Meteorol. Soc.* **92**: 325–334.
- Cullen MJP, Purser RJ. 1984. An extended Lagrangian theory of semi-geostrophic frontogenesis. *J. Atmos. Sci.* **41**: 1477–1497.
- Goody R. 2000. Sources and sinks of climate entropy. *Q. J. Roy. Meteorol. Soc.* **126**: 1953–1970.
- Gregory AR, West V. 2002. The sensitivity of a model's stratospheric tape recorder to the choice of advection scheme. *Q. J. Roy. Meteorol. Soc.* **128**: 1827–1846.
- Hardiman SC, Boutle IA, Bushell AC, Butchart N, Cullen MJP, Field PR, Furtado K, Manners JC, Milton SF, Morcrette C, O'Connor FM, Shipway BJ, Smith C, Walters DN, Willett MR, Williams KD, Wood N, Abraham NL, Keeble J, Maycock AC, Thuburn J. 2015. Processes controlling tropical tropopause temperature and stratospheric water vapor in climate models. *J. Clim* **28**: 6516–6535.
- Holton JR, Haynes PH, McIntyre ME, Douglass AR, Rood RB, Pfister L. 1995. Stratosphere–troposphere exchange. *Rev. Geophys.* **33**: 403–439.
- Hoskins BJ, Bretherton FP. 1972. Atmospheric frontogenesis models: Mathematical formulation and solution. *J. Atmos. Sci.* **29**: 11–37.
- Hoskins BJ, McIntyre ME, Robertson AW. 1985. On the use and significance of isentropic potential vorticity maps. *Q. J. Roy. Meteorol. Soc.* **111**: 877–946.
- Johnson DR. 1997. “General coldness of climate models” and the second law: Implications for modeling the Earth system. *J. Clim* **10**: 2826–2846.
- Johnson DR, Lenzen AJ, Zapotocny TH, Schaack TK. 2000. Numerical uncertainties in the simulation of reversible isentropic processes and entropy conservation. *J. Clim* **13**: 3860–3884.
- Johnson DR, Lenzen AJ, Zapotocny TH, Schaack TK. 2002. Numerical uncertainties in the simulation of reversible isentropic processes and entropy conservation: Part II. *J. Clim* **15**: 1777–1804.
- Judkes M. 2001. A generalization of the transformed Eulerian–mean meridional circulation. *Q. J. Roy. Meteorol. Soc.* **127**: 147–160.
- Kavčić I, Thuburn J. 2018. A Lagrangian vertical coordinate version of the ENDGame dynamical core. Part I: Formulation, remapping strategies, and robustness. *Q. J. Roy. Meteorol. Soc.* Submitted.
- Lauritzen PH, Bacmeister JT, Dubos T, Lebonnois S, Taylor MA. 2014. Held–Suarez simulations with the Community Atmosphere Model Spectral Element (CAMSE) dynamical core: A global axial angular momentum analysis using Eulerian and floating Lagrangian vertical coordinates. *J. Adv. Model. Earth Syst.* **6**: 129–140.
- Lorenz EN. 1955. Available potential energy and the maintenance of the global circulation. *Tellus* **7**: 157–167.
- Mahowald NM, Plumb RA, Rasch PJ, del Corral J, Sassi F, Heres W. 2002. Stratospheric transport in a three-dimensional isentropic coordinate model. *J. Geophys. Res.* **107**: ACH 3, 1–14.
- McDonald A. 1986. A semi-Lagrangian and semi-implicit two time level integration scheme. *Mon. Weather Rev.* **114**: 824–830.
- McDonald A, Bates JR. 1987. Improving the estimate of the departure point position in a two-time level semi-Lagrangian and semi-implicit scheme. *Mon. Weather Rev.* **115**: 737–739.
- McIntyre ME, Palmer TN. 1984. The “surf zone” in the stratosphere. *J. Atmos. Terr. Phys.* **46**: 825–849.
- Nakamura N, Held IM. 1989. Nonlinear equilibration of two-dimensional Eady waves. *J. Atmos. Sci.* **46**: 3055–3064.
- Newman PA, Schoeberl MR, Plumb RA, Rosenfield JE. 1988. Mixing rates calculated from potential vorticity. *J. Geophys. Res.* **93**: 5221–5240.
- Ozawa H, Ohmura A, Lorenz RD, Pujol T. 2003. The second law of thermodynamics and the global climate system: A review of the maximum entropy production principle. *Rev. Geophys.* **41**: 1–24.
- Paltridge GW. 1975. Global dynamics and climate – a system of minimum entropy exchange. *Q. J. Roy. Meteorol. Soc.* **101**: 475–484.
- Peixoto JP, Oort AH. 1992. *Physics of climate*. American Institute of Physics.
- Rasch PJ, Coleman DB, Mahowald N, Williamson DL, Lin S, Boville BA, Hess P. 2006. Characteristics of atmospheric transport using three numerical formulations for atmospheric dynamics in a single GCM framework. *J. Clim* **19**: 2243–2266.
- Saffin L, Methven J, Gray SL. 2016. The non-conservation of potential vorticity by a dynamical core compared with the effects of parametrized physical processes. *Q. J. Roy. Meteorol. Soc.* **142**: 1265–1275.
- Schneider T, Held IM, Garner ST. 2003. Boundary effects in potential vorticity dynamics. *J. Atmos. Sci.* **60**: 1024–1040.
- Tailleux R. 2013. Available potential energy and exergy in stratified fluids. *Annu. Rev. Fluid Mech.* **45**: 35–58.
- Ullrich PA, Melvin T, Jablonowski C, Staniforth A. 2014. A proposed baroclinic wave test case for deep- and shallow-atmosphere dynamical cores. *Q. J. Roy. Meteorol. Soc.* **140**: 1590–1602.
- Webster S, Thuburn J, Hoskins BJ, Rodwell M. 1999. Further development of a hybrid-isentropic GCM. *Q. J. Roy. Meteorol. Soc.* **125**: 2305–2331.
- Whitehead JP, Jablonowski C, Kent J, Rood RB. 2015. Potential vorticity: Measuring consistency between GCM dynamical cores and tracer advection schemes. *Q. J. Roy. Meteorol. Soc.* **141**: 1001–1017.

- Wood N, Staniforth A, White A, Allen T, Diamantakis M, Gross M, Melvin T, Smith C, Vosper S, Zerroukat M, Thuburn J. 2014. An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the deep-atmosphere global non-hydrostatic equations. *Q. J. Roy. Meteorol. Soc.* **140**: 1505–1520.
- Woollings TJ, Thuburn J. 2006. Entropy sources in a dynamical core atmospheric model. *Q. J. Roy. Meteorol. Soc.* **132**: 43–59.
- Zhu Z, Schneider EK. 1997. Improvement in stratosphere simulation with a hybrid $\sigma - \theta$ coordinate GCM. *Q. J. Roy. Meteorol. Soc.* **123**: 2095–2113.

Accepted Article