Geoscientific
Model Development

The Extrapolar SWIFT model (version 1.0): fast stratospheric ozone chemistry for global climate models

Daniel Kreyling, Ingo Wohltmann, Ralph Lehmann, and Markus Rex

Alfred Wegener Institute, Helmholtz Centre for Polar and Marine Research, Potsdam, Germany

Correspondence: Daniel Kreyling (daniel.kreyling@awi.de)

Abstract. The Extrapolar SWIFT model is a fast ozone chemistry scheme for interactive calculation of the extrapolar stratospheric ozone layer in coupled general circulation models (GCMs). In contrast to the widely used prescribed ozone, the SWIFT ozone layer interacts with the model dynamics and can respond to atmospheric variability or climatological trends.

The Extrapolar SWIFT model employs a repro-modelling approach, in which algebraic functions are used to approximate the numerical output of a full stratospheric chemistry and transport model (ATLAS). The full model solves a coupled chemical differential equation system with 55 initial and boundary conditions (mixing ratio of various chemical species and atmospheric parameters). Hence the rate of change of ozone over 24 h is a function of 55 variables. Using covariances between these variables, we can find linear combinations in order to reduce the parameter space to the following nine *basic* variables: latitude, pressure altitude, temperature, overhead ozone column and the mixing ratio of ozone and of the ozone-depleting families (Cl_y, Br_y, NO_y and HO_y). We will show that these nine variables are sufficient to characterize the rate of change of ozone. An automated procedure fits a polynomial function of fourth degree to the rate of change of ozone obtained from several simulations with the ATLAS model. One polynomial function is determined per month, which yields the rate of change of ozone over 24 h. A key aspect for the robustness of the Extrapolar SWIFT model is to include a wide range of stratospheric variability in the numerical output of the ATLAS model, also covering atmospheric states that will occur in a future climate (e.g. temperature and meridional circulation changes or reduction of stratospheric chlorine loading).

For validation purposes, the Extrapolar SWIFT model has been integrated into the ATLAS model, replacing the full stratospheric chemistry scheme. Simulations with SWIFT in ATLAS have proven that the systematic error is small and does not accumulate during the course of a simulation. In the context of a 10-year simulation, the ozone layer simulated by SWIFT shows a stable annual cycle, with inter-annual variations comparable to the ATLAS model. The application of Extrapolar SWIFT requires the evaluation of polynomial functions with 30–100 terms. Computers can currently calculate such polynomial functions at thousands of model grid points in seconds. SWIFT provides the desired numerical efficiency and computes the ozone layer 10^4 times faster than the chemistry scheme in the ATLAS CTM.

1 Introduction

Modern climate models include an increasing number of climate processes and run with ever higher model resolutions. Many processes that are relevant for the climate system are already well understood, but they remain computationally too demanding to be incorporated into climate models. One of these processes is the stratospheric ozone chemistry. The feedbacks between the ozone layer and the changing climate system have been investigated in various studies (e.g. Thompson and Solomon, 2002; Rex et al., 2006; Baldwin et al., 2007; Nowack et al., 2014; Calvo et al., 2015). All of them emphasize the importance of the interactions between climate change and the ozone layer. Climate simulations with a more accurate representation of the ozone layer lead to significant changes in tropospheric and surface variables. However, a frequently used approach to represent the

ozone layer in general circulation models (GCMs) is the use of prescribed zonal mean ozone climatologies, as in many of the Coupled Model Intercomparison Project 5 (CMIP5) simulations (IPCC, 2014). By using prescribed ozone, the atmospheric dynamics cannot interact with the ozone field, the ozone hole is a static, zonally symmetric feature that does not interact with atmospheric waves and the ozone layer does not respond to climate change and vice versa. But this approach is computationally cheap and does not impede the GCM capacity for ensemble simulations. The incorporation of an interactive ozone layer instead of climatologies allows the ozone field to actually match the model dynamics and enables two-directional feedbacks. Chemistry climate models (CCMs) with a highly resolved stratosphere usually provide such an interactive ozone layer, but the computational cost of CCMs still limits their usefulness for long-term ensemble simulations (Eyring et al., 2010). In recent years different approaches were taken to efficiently incorporate interactive ozone in climate simulations (Eyring et al., 2013). One of these approaches is the development of stratospheric ozone chemistry schemes with a very low computational burden in comparison to the computation time of the GCM, for example the Cariolle scheme (Cariolle and Teyssèdre, 2007) or the Linoz scheme (Hsu and Prather, 2009). In this paper we introduce the extrapolar part of the numerically efficient and interactive stratospheric ozone chemistry scheme SWIFT. Its goal is to provide sufficient accuracy and efficiency to enable ensemble simulations with atmosphere–ocean coupled GCMs, while maintaining the physical and chemical properties of the processes that govern ozone chemistry in the stratosphere so that the SWIFT approach is valid for a wide range of climatic conditions, including future climate scenarios.

SWIFT is subdivided into a polar and an extrapolar module. The two sub-modules follow separate approaches due to the differences in polar and extrapolar ozone chemistry. The lack of sunlight and very low temperatures during polar night extend the chemical lifetimes of various trace gases relevant for ozone depletion. Under these conditions the individual species within the chemical families Cl_y, Br_y, NO_y and HO_y are too far from chemical equilibrium so that their time evolution needs to be calculated with differential equations. The Polar SWIFT model simulates the time evolution of polar-vortex-averaged mixing ratios of ozone and four key species during Arctic and Antarctic winters. A small coupled differential equation system containing empirically determined fit parameters models the most relevant processes of polar ozone depletion. The first Polar SWIFT version was described by Rex et al. (2014) and the updated version was published by Wohltmann et al. (2017).

In extrapolar conditions the diurnal average concentrations of the individual species within the chemical families (partitioning) mentioned above are sufficiently close to photochemical steady state because the photochemical lifetimes of the involved species are sufficiently short compared to

the transport timescales. In a good approximation the chemically induced change in ozone over 24 h is a function of the concentrations of the chemical families, ozone itself and the physical boundary conditions. The Extrapolar SWIFT model is based on the substitution of a comprehensive differential equation system describing the ozone changes by algebraic functions. This approach is also referred to as repro-modelling and has been successfully applied to chemical models; see Spivakovsky et al. (1990), Turányi (1994) and Lowe and Tomlin (2000). As in the previous studies we obtain the algebraic functions by fitting the numerical solution of the chemical differential equation system with orthonormal polynomial functions. Following the approach of Turányi (1994) we use a wide range of input and output values of a full chemical model to create a data set that is then used for fitting the polynomial functions. However, a few modifications were introduced, most prominently in the selection of the most suitable polynomial terms. Moreover, we developed a termination criterion that does not require the selection of arbitrary thresholds. It is important to note that the repro-model is not a shortened subset of the full chemical system. By approximating the output of the full system, we ensure that all physical and chemical properties of the full chemical model are maintained in the repro-model. In this application the rate of change of ozone in the lower and middle stratosphere is parameterized by one polynomial function per month. Each of these polynomials is a function of nine *basic* variables, which are sufficient to parameterize the rate of change of ozone in the full chemical system. The *basic* variables are latitude, pressure altitude, temperature, the overhead ozone column, the volume mixing ratio (VMR) of the ozone-depleting substances (ODSs) combined into four chemical families and ozone itself. The calculation of the polynomial function values instead of solving the chemical differential equation system drastically reduces the computational cost and makes SWIFT a suitable candidate for coupling to a GCM.

Existing fast ozone schemes for climate models like the Cariolle scheme (Cariolle and Teyssèdre, 2007) or the Linoz scheme (McLinden et al., 2000; Hsu and Prather, 2009) use a first-order Taylor-series expansion of the rate of change of ozone around mean atmospheric states of ozone mixing ratio, temperature and the overhead ozone column. In comparison to SWIFT, these schemes do not explicitly include the abundance of ODS as a variable in the model. Handling changes in stratospheric ODS abundance requires the repeated determination of production and loss rates and their derivatives. Including the ODS as additional degrees of freedom in the Extrapolar SWIFT model increases its resilience to ODS variability. Moreover, the linear Taylor-series functions tend to produce larger deviations when the rate of change of ozone is not linear with respect to the variability of the three variables. The Extrapolar SWIFT polynomial functions are continuous throughout the stratosphere and can cope with the non-linear parts of the rate of change of ozone.

In Sect. 2 of this paper the application of repro-modelling to the rate of change of ozone is described. First we introduce the set-up of the repro-model, containing polynomial coefficients as free parameters. Further, the approximation algorithm determining these coefficients is described and its modifications in comparison to previous studies are explained. Section 3 focuses on the domain of definition of the polynomial functions and how outliers are handled in Extrapolar SWIFT. A validation and error estimation of the polynomial functions are presented in Sect. 4. Eventually, two different simulations with SWIFT are discussed in Sect. 5. A 2-year simulation focuses on the error in the ozone field caused by the monthly polynomial functions. A 10-year simulation mimics the set-up of SWIFT in a GCM and demonstrates the stability of the model over a longer simulation period. The development of Extrapolar SWIFT and the results of the simulations are also discussed in Kreyling (2016).

2 Application of repro-modelling to stratospheric ozone chemistry

2.1 Setting up the repro-model

The Extrapolar SWIFT repro-model calculates the rate of change of ozone over 24 h by evaluating polynomial functions of fourth degree. Each polynomial function is valid during 1 month of the year. To determine these polynomial functions we use multivariate fitting of a representative data set which comprises a wide range of stratospheric conditions, as suggested by Turányi (1994). As a source for the rate of change of ozone we use the comprehensive Lagrangian stratospheric chemistry and transport model ATLAS. The ATLAS model is described in detail in Wohltmann and Rex (2009) and Wohltmann et al. (2010). It contains 49 stratospheric trace gases interacting with each other in over 170 gas-phase and heterogeneous chemical reactions. Together with atmospheric and geographic initial and boundary conditions the differential equation system contains 55 variables and parameters. The rate of change of ozone may be represented as a function of 55 arguments:

$$\frac{\mathrm{dO_x}}{\mathrm{d}t} = \hat{F}(x_1, x_2, \ldots, x_{55}), \tag{1}$$

where O_x is the VMR of the odd oxygen family containing O_3, O and $O(^1D)$, and $\hat{F} : \mathbb{R}^{55} \to \mathbb{R}$. The O_x family has a longer chemical lifetime than ozone, which is beneficial to our approximation approach. Moreover, in the lower and middle stratosphere odd oxygen almost entirely consists of ozone. Thus in Extrapolar SWIFT O_x substitutes O_3.

In order to set up a repro-model, we need to determine a set of *basic* variables which are sufficient for the parameterization of all the physical and chemical processes in the full chemical system. The determination of *basic* variables is a crucial aspect since their number should be large enough so

that the function in Eq. (1) is approximated with sufficient accuracy. On the other hand, their number should be as small as possible so that the repro-model is numerically efficient. This is partly achieved by lumping the chemical species into chemical families. The following four chemical families are relevant for ozone depletion in the stratosphere and therefore constitute four of the *basic* variables:

$$\mathrm{Cl}_y = \underbrace{\mathrm{Cl} + \mathrm{ClO} + \mathrm{Cl_2O_2}}_{\text{short-lived}} + \underbrace{\mathrm{ClONO_2} + \mathrm{HCl}}_{\text{reservoir}}$$

$$\mathrm{Br}_y = \underbrace{\mathrm{Br} + \mathrm{BrO} + \mathrm{HBr} + \mathrm{HOBr}}_{\text{short-lived}} + \underbrace{\mathrm{BrONO_2} + \mathrm{BrCl}}_{\text{reservoir}}$$

$$\mathrm{NO}_y = \underbrace{\mathrm{N} + \mathrm{NO} + \mathrm{NO_2} + \mathrm{NO_3}}_{\text{short-lived}} + \underbrace{\mathrm{HNO_3}}_{\text{reservoir}}$$

$$\mathrm{HO}_y = \underbrace{\mathrm{H} + \mathrm{OH} + \mathrm{HO_2}}_{\text{short-lived}} + \underbrace{\mathrm{H_2O}}_{\text{reservoir}}$$

The stratospheric ozone depletion is driven by catalytic cycles involving the short-lived species of the above-listed chemical families. Consequently, the repro-model requires information on the concentration of the short-lived compounds. This may be derived from the concentrations of the chemical families. In the extrapolar regions the short-lived reactive species (e.g. ClO_x or BrO_x) are sufficiently close to chemical equilibrium determined by the local conditions (e.g. pressure, temperature, radiation and the abundance of reaction partners). Consequently, in the chemical families containing only one reservoir gas (NO_y and HO_y) the concentration of the short-lived species is uniquely determined by the abundance of the total family; i.e. we assume local chemical equilibrium between the short-lived and reservoir species. For Cl_y and Br_y the partitioning between the reservoir species needs to be considered. However, in most regions of the extrapolar stratosphere the lifetime of $ClONO_2$ is shorter than the timescales of vertical or meridional transport so that $ClONO_2$ also comes close to equilibrium state. The same can certainly be assumed for $BrONO_2$, which has an even shorter lifetime than $ClONO_2$.

Apart from the VMR of the chemical constituents, the reaction rates depend on temperature, air density and in the case of photolysis rates on the actinic flux, particularly on the ultraviolet attenuation (UV attenuation). These parameters must also be implicitly or explicitly included into the set of *basic* variables. Table 1 summarizes the nine *basic* variables we have identified. The column "Remarks" points out different properties and processes parameterized by the variable. A function of these nine variables (Eq. 2) sufficiently approximates the function in Eq. (1), but reduces the dimensionality from 55 to 9.

$$\Delta O_x = F\left(\phi, z, T, \mathrm{topO_3}, \mathrm{Cl}_y, \mathrm{Br}_y, \mathrm{NO}_y, \mathrm{HO}_y, O_x\right), \tag{2}$$

where ΔO_x is the rate of change of ozone over 24 h and $F : \mathbb{R}^9 \to \mathbb{R}$. After determining an approximation for $F \approx \widetilde{F}$ in

Table 1. Nine *basic* variables of the Extrapolar SWIFT model. The column "Remarks" lists properties and processes parameterized by the variable. Pressure altitude is defined as $z = -H \log(\frac{p}{p_0})$ and overhead ozone is the integrated ozone column above a specific location in the atmosphere.

Variable	Unit	Symbol	Remarks
Latitude	(°)	ϕ	Solar zenith angle and actinic flux
Pressure altitude	(m)	z	Air density and actinic flux
Temperature	(K)	T	Kinetics of reactions
Overhead ozone column	(DU)	$topO_3$	Attenuation of UV radiation
Chlorine family	(ppb)	Cl_y	Catalytic ClO_x chemistry
Bromine family	(ppt)	Br_y	Catalytic BrO_x chemistry
Nitrogen oxide family	(ppb)	NO_y	Catalytic NO_x chemistry
Hydrogen oxide family	(ppm)	HO_y	Catalytic HO_x chemistry
Odd oxygen	(ppm)	O_x	Chapman and catalytic chemistry

Eq. (2) by SWIFT, the chemical change of the ozone VMR at each grid point in a GCM simulation can be calculated by Eq. (3).

$$O_x(t + 24\,h) = O_x(t) + \Delta O_x(t) \cdot 24\,h$$
$$\approx O_x(t) + \widetilde{F}(\phi(t), z(t), T(t), topO_3(t), Cl_y(t),$$
$$Br_y(t), NO_y(t), HO_y(t), O_x(t)) \cdot 24\,h \quad (3)$$

2.2 Approximation algorithm

The algebraic equation of the repro-model is a polynomial function of fourth degree (i.e. the sum of the exponents of a term is ≤ 4). The polynomial uses the same nine *basic* variables as in Eq. (2) and yields the rate of change of ozone over 24 h. The ΔO_x function in Eq. (2) can be approximated by a polynomial function $\widetilde{F}$:

$$F(x_1, \ldots, x_9) \approx \widetilde{F}(x_1, \ldots, x_9) = \sum_{j=1}^{n} c_j f_j(x_1, \ldots, x_9), \quad (4)$$

where $x_1, \ldots, x_9$ represents the *basic* variables, f_j represents polynomial terms (e.g. $f_1 = x_1^2 x_2$) and c_j represents their coefficients for $j = 1 \ldots n$, where n is the number of polynomial terms. For a polynomial function of fourth degree with nine variables, the maximum number of terms is 715, including all mixed terms. The coefficients c_j in Eq. (4) are determined such that the rate of change of O_x is calculated by the ATLAS CTM for m different values of the *basic* variables $x_{i1}, \ldots, x_{i9}$, $i = 1, \ldots, m$, which are approximated with best accuracy:

$$F(x_{i1}, \ldots, x_{i9}) \approx \sum_{j=1}^{n} c_j f_j(x_{i1}, \ldots, x_{i9}). \quad (5)$$

The m different values of the *basic* variables will be referred to as training data points or the training data set. In order to write Eq. (5) in matrix notation we define an $m \times n$ matrix $\mathbf{A}$ with $\mathbf{A}_{ij} = f_j(x_{i1}, \ldots, x_{i9})$ and a vector $\boldsymbol{F}$ with

$F_i = F(x_{i1}, \ldots, x_{i9})$, $i = 1, \ldots, m$. The polynomial coefficients c_j are grouped into a vector $\boldsymbol{c}$. Then, the linear equation system in Eq. (5) can be expressed as

$$\mathbf{A}\boldsymbol{c} = \boldsymbol{F}. \quad (6)$$

To determine $\boldsymbol{c}$ we employ the least-squares method, which is to minimize the Euclidian norm ($\|\,\|$) of the deviation between the approximation and $\boldsymbol{F}$:

$$\|\mathbf{A}\boldsymbol{c} - \boldsymbol{F}\| \rightarrow \min. \quad (7)$$

The minimization in Eq. (7) can be made more efficient and numerically stable by first transforming the matrix $\mathbf{A}$ into an orthogonal matrix. Spivakovsky et al. (1990) achieve this with successive Householder transformations which finally yield the **QR** decomposition of matrix $\mathbf{A}$. Turányi (1994) and Lowe and Tomlin (2000) use the Gram–Schmidt process for orthogonalization. The literature suggests (e.g. Golub and Van Loan, 1996) that the unmodified Gram–Schmidt process has worse numerical properties which can impair the orthogonalization. In our approach we are using a **QR** decomposition based on the Householder transformation.

We start the fitting procedure with one polynomial term ($n = 1$) on the right-hand side of Eq. (5). During the following iterations the polynomial function is consecutively extended by one additional term. This corresponds to an extension of the matrix $\mathbf{A}$ by one column. Turányi (1994) started the approximation with the constant term and continued with linear terms, then quadratic terms and so on up to terms of maximum degree, also including all mixed terms. In each iteration the residuum $\|\mathbf{A}\boldsymbol{c} - \boldsymbol{F}\|$ was calculated. If the current residuum was reduced by a certain threshold relative to the previous residuum, then the current term was accepted to be added to the polynomial function. This method tested the terms in a given arbitrary order. If the order of testing had been a different one, other polynomial terms would be accepted and the overall quality of the polynomial function could potentially be better.

In our approach we are circumventing this problem by testing all polynomial terms individually as the next addi-

tional term. In other words, in each iteration each of the still available polynomial terms is temporarily added to the already selected terms and the fitting procedure is carried out. The term which reduces the residuum the most is permanently added to the polynomial function and removed from the pool of available terms. In the next iteration all remaining terms are fitted in combination with the previously accepted ones. By simply choosing the best fitting term we also avoid setting an arbitrary threshold for the minimum required reduction of the residuum. This polynomial term selection method makes the fitting procedure computationally much more extensive. However, the fitting procedure has to be carried out only once so that this additional computation time imposes no disadvantage during the application of SWIFT.

The more polynomial terms are added to the function, the better the approximation will be; i.e. the residuum can be reduced further and further. If as many polynomial terms (corresponding to columns of $\mathbf{A}$) are fitted as there are training data points (corresponding to rows of $\mathbf{A}$) then the linear equation system in Eq. (6) is no longer overdetermined. In this case any small-scale structure originating from the random distribution of training data points would have been fitted and the polynomial function would contain an impractically large number of terms. An overfitted polynomial function actually causes the residuum to be higher when it is applied to an independent data set (i.e. not a subset of the data the polynomial was fitted to). Consequently the fitting procedure should be terminated before the random fluctuations in the training data set are fitted. This termination criterion can be defined by applying the selected polynomial terms and their coefficients to an independent data set instead of the training data set. The independent data set is named the testing data set here. The quality of the approximation is expressed by

$$r = \|\mathbf{A}^{\text{Test}}c - F^{\text{Test}}\|, \qquad (8)$$

where $\mathbf{A}^{\text{Test}}$ is like the matrix $\mathbf{A}$, only the rows of $\mathbf{A}^{\text{Test}}$ correspond to the testing data points and the vector F^{Test} contains the rate of change of ozone at the testing data points. The polynomial coefficients c are the ones determined via Eq. (7), and r is the residuum corresponding to the polynomial function with one temporarily added term. At some iteration during the fitting procedure, the residuum r will not be reduced by any of the available additional terms. This defines the termination of the approximation algorithm.

It is important that the testing data set has the same probability distribution of *basic* variables as the training data set. We achieve this by randomly separating the output of the ATLAS simulations into the training and the testing data set, containing 2 / 3 and 1 / 3 of the total output, respectively.

2.3 Latitude and altitude boundaries of the repro-model

In this section we discuss where in the stratosphere the Extrapolar SWIFT model can be used, i.e. for which latitudes
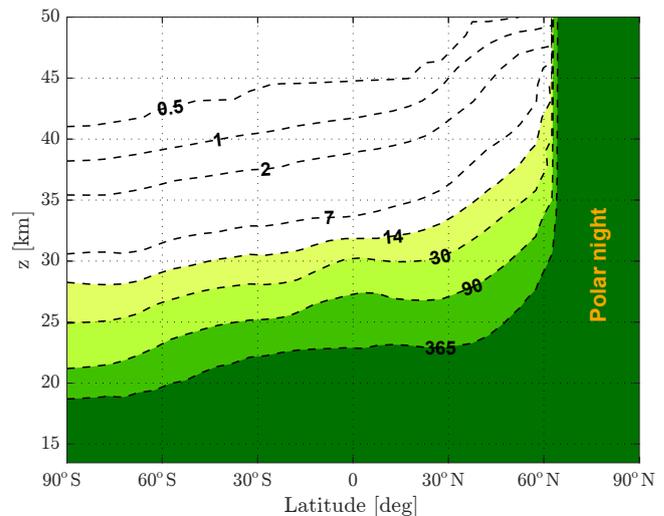

Figure 1. Zonal mean of O_x lifetime in January derived from ATLAS CTM data. Contour numbering shows the lifetime in days.

and altitudes the underlying assumptions are valid. A key aspect for the definition of this latitude–altitude region is the mean chemical lifetime τ of O_x.

$$\tau = \frac{[O_x]}{R}, \qquad (9)$$

where $[O_x]$ is the concentration instead of VMR and R is the sum of the rates of all O_x-depleting catalytic cycles. In Fig. 1 the mean chemical lifetime of O_x taken from ATLAS data for January is displayed. The contour labels specify the lifetime in days. In the lower stratosphere the O_x lifetimes exceed 365 days. The longer lifetimes in the lower stratosphere are a consequence of the slower reaction rates of the catalytic O_x-loss cycles mostly due to fewer O atoms. The O atom is produced via the photolysis of O_3, but its vertical distribution is mainly controlled by the three-body reaction $O + O_2 + M \rightarrow O_3 + M$. The rate constant of this reaction increases with increasing pressure. The latitudinal (and seasonal) variation of the O_x lifetime reflects the varying length of the day and the attenuation of solar radiation on its way through the atmosphere.

Above roughly 30 km of altitude the mean lifetime of O_x is shorter than vertical and meridional transport timescales. In this quasi-chemical equilibrium state the O_x concentration is determined by the local meteorological conditions and the abundance of ODS. Consequently O_x can be calculated as a function F_{eq} of the previously mentioned *basic* variables, but without the O_x VMR itself, so that $F_{\text{eq}} : \mathbb{R}^8 \rightarrow \mathbb{R}$.

$$O_x = F_{\text{eq}}(\phi, z, T, \text{top}O_3, Cl_y, Br_y, NO_y, HO_y) \qquad (10)$$

Accordingly, in the upper stratosphere the O_x VMR at a point in time $t + \Delta t$ is a function of eight *basic* variables at time $t + \Delta t$. The function F_{eq} can also be approximated by a poly-

nomial function $\widetilde{F}_{eq}$:

$$O_x(t + \Delta t) \approx \widetilde{F}_{eq}(x_1(t + \Delta t), \ldots, x_8(t + \Delta t)), \qquad (11)$$

where $x_1, \ldots, x_8$ corresponds to the eight variables in Eq. (10). In SWIFT the polynomial functions calculate ΔO_x and determine the O_x VMR of the next time step as in Eq. (3). The O_x VMR at time $t + \Delta t$ is a function of the nine *basic* variables $x_1(t), \ldots, x_9(t)$.

$$O_x(t + \Delta t) = O_x(t) + \widetilde{F}(x_1(t), \ldots, x_9(t)) \cdot 24\,h \qquad (12)$$

Both equations (Eqs. 11 and 12) yield $O_x(t + \Delta t)$. Setting them equal results in

$$
\begin{aligned}
O_x(t) &+ \widetilde{F}(x_1(t), \ldots, x_9(t)) \cdot 24\,h \\
&= \widetilde{F}_{eq}(x_1(t + \Delta t), \ldots, x_8(t + \Delta t)) \\
\widetilde{F}(x_1(t), &\ldots, x_9(t)) \cdot 24\,h \\
&= \widetilde{F}_{eq}(x_1(t + \Delta t), \ldots, x_8(t + \Delta t)) - O_x(t). \qquad (13)
\end{aligned}
$$

This means that in the quasi-equilibrium region of O_x, $\widetilde{F}$ is the result of the O_x polynomial function in equilibrium $\widetilde{F}_{eq}$ minus the linear O_x term. However, the polynomial function $\widetilde{F}$ contains various O_x terms of higher degree. These terms together with the higher O_x VMR in the upper stratosphere cause rather large errors. Consequently, the polynomial function $\widetilde{F}$ is not suited to be used in the region where O_x is in quasi-chemical equilibrium. The altitude of 30 km roughly marks the transition between the equilibrium and non-equilibrium state of O_x. The lifetime is roughly 14 days in this altitude (see Fig. 1). We defined the 14-day contour to be the upper boundary up to which the polynomial functions can be used or rather up to which the training and testing data sets reach.

Since the lifetime of O_x is a function of the incoming solar radiation, the altitude and tilt of the 14-day contour also depends on the season. In the course of the year the tilt of the lifetime contours will shift. For each monthly polynomial function we defined a separate upper boundary. In the quasi-equilibrium region (upper stratosphere) the SWIFT simulations currently require ozone values interpolated from stratospheric climatologies. For the future a similar repro-modelling approach will be applied to function F_{eq} in Eq. (10) by fitting the O_x VMR directly instead of the rate of change of O_x.

However, the upper stratosphere only contributes a few percent to the stratospheric ozone column. The bulk of ozone dominating the total column values is in the lower stratosphere below 30 km. This motivated our focus on this part of the stratosphere which we will refer to as the ΔO_x regime.

The lower boundary of the ΔO_x regime is set to 15 km of pressure altitude (roughly 120 hPa). In the tropics 15 km is approximately the altitude of the tropical tropopause layer (TTL) defining the boundary between tropospheric and stratospheric air. In the extratropical regions ozone-rich air

can also be found below 15 km, especially in the northern high latitudes. However, at theses altitudes and latitudes the rate of change of ozone is close to zero and the transport of ozone is much more relevant (see also Fig. 1). When running Extrapolar SWIFT in a GCM, treating ozone as a passive tracer below 15 km of pressure altitude is recommended.

The regime boundaries between Extrapolar SWIFT and Polar SWIFT are defined by the edge of the polar vortex. The horizontal extent of the polar vortex is defined by 36 mPV units, where mPV is the modified potential vorticity according to Lait (1994) (with $\theta_0 = 475$ K). In the vertical, the specified vertical extent of the Polar SWIFT domain goes from roughly 18 to 27 km of pressure altitude. Above and below Polar SWIFT the extrapolar module is used, although the rate of change of ozone is close to zero during polar night.

2.4 Training data

The monthly training and testing data set for Extrapolar SWIFT are generated with the stratospheric Lagrangian chemistry and transport model ATLAS (Wohltmann and Rex, 2009; Wohltmann et al., 2010). The data used in this work originated from two 2.5-year simulations, one from November 1998 to March 2001 and the second from November 2004 to March 2007. The chemistry module of ATLAS contains a comprehensive set of gas-phase chemical reactions and a heterogeneous chemistry scheme. Photolysis and reaction coefficients are taken from the recent Jet Propulsion Laboratory (JPL) catalog (Sander et al., 2011). All partial species of the 4 ozone-depleting chemical families (Cl_y, Br_y, NO_y and HO_y) are included in the 49 ATLAS species. The individual species are initialized from different sources. The VMR of H_2O, N_2O, HCl, O_3, CO and HNO_3 were initialized from the Aura Microwave Limb Sounder (MLS) climatologies for the 1998–2001 simulation. The 2004–2007 simulation used the measurements of Aura MLS directly (Waters et al., 2006). The VMR of CH_4 and NO_2 (substitute for NO_x) were taken from climatologies of the HALogen Occultation Experiment (HALOE) instrument (Grooß and Russell III, 2005). Initial values for Cl_y and Br_y were derived from tracer–tracer correlations to CH_4 and N_2O measured during an aircraft and ballooning campaign described in Grooß et al. (2002). The ATLAS trajectories are initialized in roughly 2 km thick pressure altitude layers with a horizontal resolution of 200 km. On each trajectory the chemistry is calculated like in a chemical box model. ATLAS solves a coupled system of differential equations to obtain the rate of change of the trace gases. The stiff numerical solver uses an automatic adaptive time step and is based on the numerical differentiation formulas (Shampine and Reichelt, 1997). After 24 h (mixing time step) the mixing algorithm merges or creates trajectories and interpolates the chemical species accordingly. The ATLAS trajectories are driven by ERA-Interim wind fields, temperatures and heating rates (Dee et al., 2011).

For each month of the year, daily snapshot values of the *basic* variables at the current location of the trajectories and the corresponding ΔO_x at a fixed time of day (00:00 UTC) are compiled into a data set which is later split into a training and testing data set. The number of trajectories computed in an average ATLAS run is roughly 10^5 throughout the lower and middle stratosphere. In order not to exceed the size of the computer's main memory, a random subsample of the 10^5 trajectories of each day of a month was taken. This was done so that all monthly data sets contain the same amount of data: 8 million and 4 million data points in the training and testing data set, respectively. The monthly data sets are chosen such that they also contain a fraction of data from the 10 days preceding and following the current month. We do this in order to ensure a smoother transition of polynomial functions from one month to the next.

Individual chemical species in ATLAS are grouped into their respective families and summed up to generate the mixing ratios of Cl_y, Br_y and NO_y. HO_y is simply substituted by water vapour, since the H_2O VMR is a factor of 1000 larger than the sum of all other HO_y constituents. The ΔO_x value is defined by the difference of the O_x VMR between two snapshots along the Lagrangian trajectory. A ΔO_x value is associated with the beginning of a 24 h period:

$$\Delta O_x(t) \cdot 24\,\text{h} = O_x(t + 24\,\text{h}) - O_x(t). \tag{14}$$

Before the fitting procedure, the *basic* variables are normalized to a range from 0 to 1. Otherwise the order of magnitude of the polynomial coefficients would vary extremely due to the strongly varying magnitude of the *basic* variables (e.g. pressure altitude $\approx 10^4$ m vs. Br_y VMR $\approx 10^{-11}$).

The Lagrangian trajectories in ATLAS are not distributed homogeneously. In general, higher trajectory densities can be found where there is strong horizontal and vertical wind shear, e.g. at the edge of the polar vortex. This is caused by the trajectory mixing algorithm in ATLAS, which initializes new or deletes existing trajectories based on their rate of divergence or convergence in a region of the model atmosphere. The regions of increased trajectory densities coincide with strong gradients of chemical constituents and meteorological parameters. Thus these gradients are well resolved in ATLAS, which is beneficial to Extrapolar SWIFT. The training and testing data sets simply contain the same unmodified sampling as in ATLAS and therefore also resolve the gradients well.

3   Validity of repro-model in a changing climate

3.1   Domain of definition of polynomial functions

The extensive training data set derived from the ATLAS CTM fills a portion of the nine-dimensional hyperspace, which defines the domain of definition of the fitted poly-
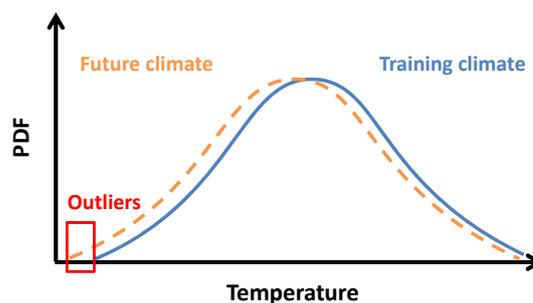

Figure 2. Schematic of a shift in the probability density function of stratospheric temperature in a future climate.

nomial functions. SWIFT is intended to be used in long-term climate simulations and it will certainly encounter inter-annual and decadal variability. Therefore we used data from ATLAS simulations covering a wide range of stratospheric variability. By taking the training and testing data from different decades we include maximum and minimum conditions of the solar cycle. The data also represent different QBO phases and the varying strengths and lifetimes of the Arctic and Antarctic polar vortices.

Climatological changes impacting the probability distribution of the *basic* variables can also be expected, e.g. changes in temperature and meridional circulation. The resilience of SWIFT to such trends is outlined in Fig. 2. Future climate scenarios will shift the current probability density function (PDF) of the *basic* variables. The schematic in Fig. 2 shows a shift of the temperature PDF, assuming a normal distribution of the temperature, with the eight other *basic* variables fixed. Most of the PDF in the training climate (blue) and the future climate (orange) overlaps. The slightly colder conditions of the future climate are thus mostly covered by the present domain. Only at low temperatures when the probability is small do outliers (red) occur. These outliers will force the polynomial functions to extrapolate and likely produce erroneous ΔO_x values. Therefore outliers need to be identified and the extrapolation must be prevented.

Apart from a PDF shift like the one illustrated in Fig. 2, there can be scenarios in which the shift of the PDF is too severe and the repro-model cannot be applied. An example would be the reduction of stratospheric chlorine by 50%. The majority of the Cl_y PDF would be outside the original PDF. In such a case the repro-model needs to be refitted to an adjusted training climate, which can easily be done by running the full ATLAS model for a few years driven by output from a climate model or with modified levels of the ODS.

3.2   Handling outliers

When running a SWIFT simulation, the polynomial function should not be evaluated outside the domain defined by the training data set. Polynomial functions of higher degree tend to rapidly increase or decrease when extrapolated. In order

to determine if a data point lies outside or inside the nine-dimensional domain of definition we need to be able to define its boundaries. This could be achieved by enveloping the nine-dimensional cloud of data points by a conjunction of nine-dimensional cells (cuboids) corresponding to a nine-dimensional regular grid (look-up table). These grid cells are either sampled by the training data set or not. A sampled grid cell is defined as being inside the domain, and all the non-sampled grid cells are outside. Dealing with a nine-dimensional grid with only a few nodes per dimension readily creates a grid with millions of cells. However, the majority of these grid cells represent combinations of *basic* variables which do not occur in the stratosphere (e.g. warm temperatures in the lowermost stratosphere). Consequently less than 0.1 % of the grid cells are actually sampled by the training data set. Using efficient ways to store and search this sparse data set would be a feasible option for identifying outliers. However, in our approach we make use of the regular grid but go one step further. Again we employ a fitting procedure to determine a polynomial function that yields positive values inside the sampled domain and negative values outside. This polynomial function is hereafter called the domain polynomial. The regular grid sampled by the training data set will be referred to as the training grid and is used for fitting the domain polynomial. The domain polynomial is obtained in the following way. First the cells of the training grid are assigned either positive or negative values. The positive values (inside the domain) are derived from the number of neighbouring cells also sampled as being inside the domain. Outside the domain the cells are assigned negative values derived from the cell's distance to the closest cell inside the domain. In order to improve the quality of the fit at the domain boundary, some smoothing operations were applied. By removing individual cells being isolated in the opposing region the transition from positive to negative values becomes more smooth. Additionally we removed outside cells which are adjacent to one cell inside the domain, but not to any other. These cells are assigned values of only -1 but are actually surrounded by outside cells with much lower values. Finally the grid cells which were assigned values close to zero are copied multiple times in the training grid to increase the weight of this region during the fit.

During the application of SWIFT within a GCM the following operations are carried out at each spatial grid point. The domain polynomial is computed for the values of the nine *basic* variables in order to determine whether these values reside inside or outside the domain of definition of the original polynomial function. If inside, the ΔO_x is calculated as usual. If the values of the nine *basic* variables prove to be outside, we need to determine a close location inside the domain of definition, where a ΔO_x can be calculated safely. Newton's method is applied to find a nearby null of the domain polynomial, which defines the boundary of the domain of definition. Within a certain margin of the null (± 0.5) the iteration of Newton's method is stopped and the ΔO_x value is

calculated at the current coordinates in the nine-dimensional space. An advantage of using the domain polynomial is that its derivatives can be computed easily and used in Newton's method.

4   Validation of polynomial functions

4.1   Comparison of the rate of change of ozone

As an initial validation step the rate of change of ozone in the testing data set is compared to the rate of change of ozone calculated by the polynomial functions. In Fig. 3 the ΔO_x in ATLAS and Extrapolar SWIFT is displayed as zonal averages. The ATLAS ΔO_x is taken from the testing data sets and the SWIFT ΔO_x from the polynomial functions evaluated on the testing data set. The four months shown (January, April, July and October) are selected as representative of each season. The data are binned into equivalent latitude (5°) vs. pressure altitude (1 km) bins and averaged. Grey shaded bins either mark areas outside the ΔO_x regime (e.g. polar vortex, upper or lower regime boundary) or indicate too few trajectories to yield a meaningful average. Since the effective area of the zonal bands decreases towards the poles, the bins with too few trajectories are found in high latitudes.

In general all four months show good agreement between ATLAS and SWIFT. Especially in the tropics and mid-latitudes the amplitude of ΔO_x and the extent of regions of production or loss compare very well. Even detailed structures like the two local maxima in the tropical ozone production region in January are visible in SWIFT. Steep gradients of ΔO_x, e.g. around 25 km at mid-latitudes, are well reproduced by the polynomial functions. Deviations between ATLAS and SWIFT occur at the upper boundary of the summer hemisphere and in high latitudes at the beginning of the winter season (e.g. Southern Hemisphere in April, Northern Hemisphere in October). In the nine-dimensional hyperspace some boundary regions of the training and testing data set are less densely populated with trajectories than more central regions. This can have different causes, but the most obvious one is the spatial difference of the trajectory density caused by the mixing algorithm in ATLAS (see Sect. 2.4). Moreover the extreme values of some of the nine *basic* variables occur less frequently if they are approximately Gaussian distributed. Finally the selection criteria for the trajectories described in Sect. 2.3 can cause sparsely sampled regions; for example, due to the variability of the polar vortex in 1 year, the January data set will include certain polar latitudes which will not be included in the January of the next year. In regions with lower trajectory density fewer squared errors need to be minimized during the least-squares minimization. Consequently these regions have less weight in the approximation than more densely sampled regions and the deviations will be larger. However, we decided not to manipulate the trajectory density in the training and testing data sets because we

wanted to maintain the frequency with which meteorological and chemical conditions occur in ATLAS. A sparsely populated region in the nine-dimensional space implies infrequent and therefore less relevant stratospheric conditions.

4.2 Error estimation

To estimate the error of Extrapolar SWIFT, we examine the difference of ATLAS ΔO_x minus SWIFT ΔO_x divided by the O_x VMR.

$$Q = \frac{\Delta O_{x_{\text{SWIFT}}} - \Delta O_{x_{\text{ATLAS}}}}{O_x} \cdot 100 \qquad (15)$$

Q is given in units of percent per day [$\% \, \text{day}^{-1}$]. Q describes the positive or negative percentage drift of O_x VMR per day due to the error in SWIFT. The division by O_x makes the differences at small and high O_x VMR more comparable, instead of just interpreting the absolute deviation. Similar to the relative error, Q tends to have larger values for very small O_x VMR. These properties of Q need to be taken into account when considering different regions with high or low ozone VMR. In the lower tropical stratosphere where very small O_x VMR can be found, the absolute errors of SWIFT are small in contrast to the Q values, which can exceed $\pm 50 \, \% \, \text{day}^{-1}$. However, for the calculation of the total ozone column the deviations at small O_x VMR are irrelevant. Also for the computation of the atmospheric heating rates based on the SWIFT ozone field, the absolute errors originating from other greenhouse gases (e.g. CO_2) with a much higher concentration are much more important than the deviations at small O_x VMR. In Fig. 4 we discuss the distribution of Q. In Sect. 5 we use the absolute deviations between the SWIFT and the ATLAS simulation to discuss the error quantitatively.

Figure 4 shows the probability distribution of Q for the four representative months January, April, July and October. As in Fig. 3 the Q values of the roughly 4 million data points of each monthly testing data set are discussed. The bin width of 1 bar in Fig. 4 is $0.2 \, \% \, \text{day}^{-1}$. Thus over $20 \, \%$ of Q values reside within the interval of $\pm 0.1 \, \% \, \text{day}^{-1}$ in all four months. The majority of Q values lie within the $\pm 1 \, \% \, \text{day}^{-1}$ interval. The mean (pink dashed line) and the median (cyan dotted line) are close to zero. The strongest systematic biases (mean) are $-0.3 \, \% \, \text{day}^{-1}$ in July and $+0.25 \, \% \, \text{day}^{-1}$ in October; the median, however, is centred very close to $0.0 \, \% \, \text{day}^{-1}$ in both months. The grey shaded area shows the standard deviation (SD) around the mean. The variability of the SD indicates that the quality of the approximation actually varies significantly between the months. The errors of the October polynomial function (SD of $3.5 \, \% \, \text{day}^{-1}$) are spread more strongly than the errors of the April polynomial function (SD is roughly $0.6 \, \% \, \text{day}^{-1}$).

As mentioned before, individual Q values can surpass $\pm 50 \, \% \, \text{day}^{-1}$ when the O_x VMR is small, i.e. below 100 ppb. But these extreme deviations are rare, which is demonstrated

by the 5 and $95 \, \%$ quantiles (black dotted lines); $90 \, \%$ of the total Q values are located between the two quantile lines.

5 Simulations with SWIFT

5.1 SWIFT coupled to the ATLAS CTM

The Extrapolar SWIFT module was coupled to the ATLAS CTM in order to perform validation simulations. In this setup the SWIFT scheme replaces the detailed stratospheric chemistry model of ATLAS. Apart from the geographical and meteorological variables provided by ATLAS, Extrapolar SWIFT requires the VMR of the four ozone-depleting chemical families Cl_y, Br_y, NO_y and HO_y. We compiled monthly zonal climatologies to be distributed with the model if required. The H_2O climatology (substituting the HO_y family) is based on extensive observational data from Aura MLS. The Cl_y, Br_y and NO_y climatologies are composed of the two ATLAS simulations used in the training and testing data sets. All species in ATLAS contributing to one of the chemical families are summed up and weighted according to their yield of active chlorine, bromine or NO_x. The initialization of chemical species for the two ATLAS simulations was described in Sect. 2.4. For initialization and regions outside the ΔO_x regime, an additional O_x climatology is required. This climatology is also compiled from an extensive set of ozone measurements by Aura MLS.

The SWIFT in ATLAS simulations are driven by ERA-Interim data (Dee et al., 2011). Every 24 h the SWIFT module is called and the rate of change of ozone is calculated based on the current conditions at the beginning of each trajectory. The VMRs of the four ozone-depleting chemical families are interpolated from the trace gas climatologies. Latitude, pressure altitude and atmospheric temperature are defined by the trajectory and the overhead ozone column is integrated from the ozone values of the overhead trajectories. In combination with these eight parameters, the ozone VMR of the last time step (24 h before) is used to calculate the rate of change of ozone (ΔO_x) by evaluating the polynomial function. Eventually the ΔO_x is added to the O_x VMR from the last time step, according to Eq. (3). In order to smooth the transition between two polynomial functions corresponding to consecutive months, we linearly interpolate between the ΔO_x results of the two polynomial functions. All other components of the ATLAS CTM, like the trajectory transport or the mixing algorithm, remain unchanged. The SWIFT in ATLAS simulations apply outlier handling as described in Sect. 3.2.

Above the seasonally dependent upper boundary of the ΔO_x regime, as introduced in Sect. 2.3, climatology values of O_x are used in the simulation. In a layer that extends over 2 km below this upper boundary the O_x VMRs are determined by computing an altitude-weighted average between values from the climatological O_x values and the
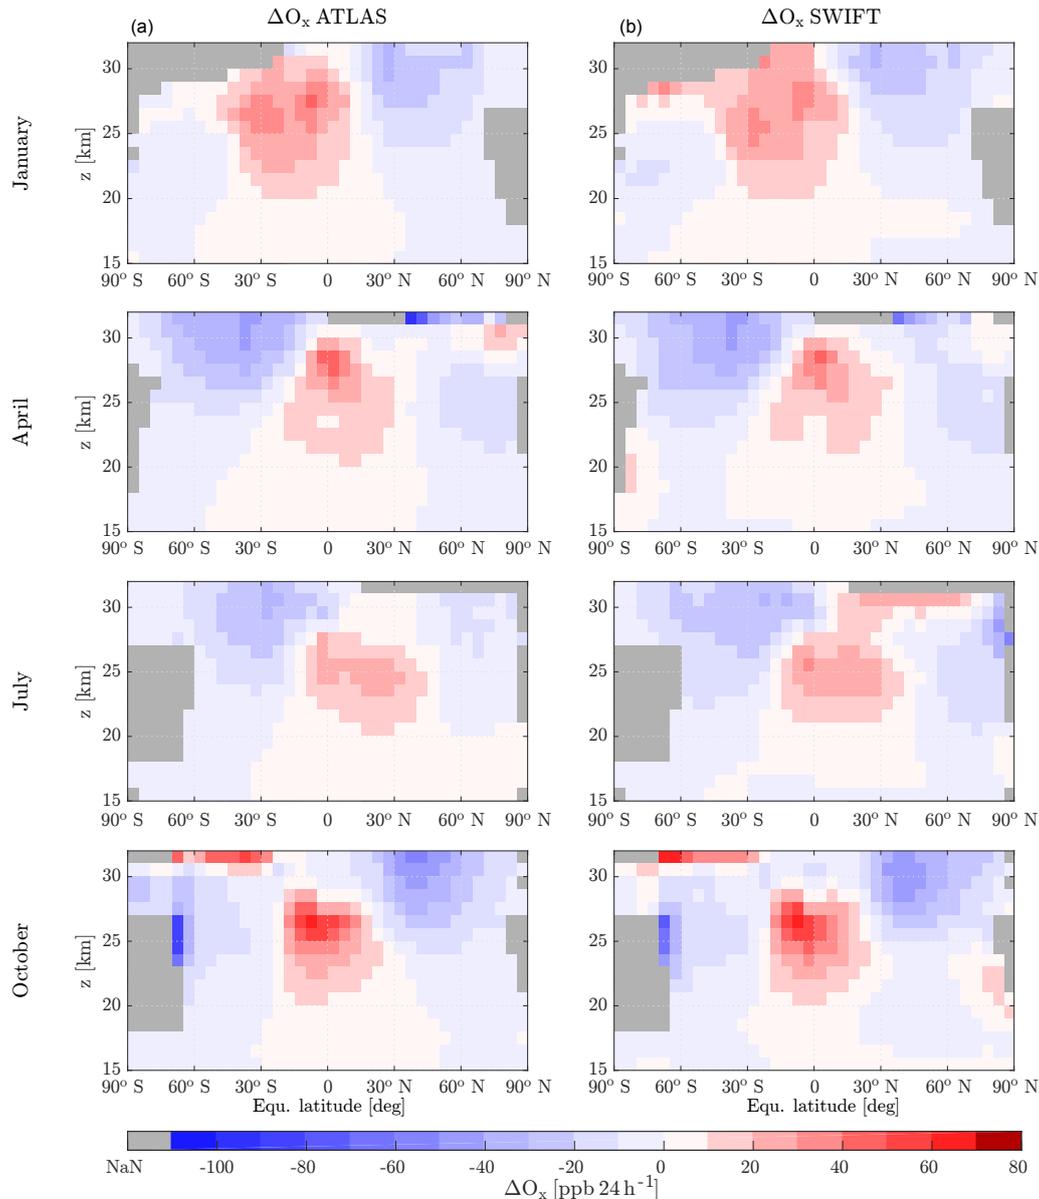
Figure 3. Zonal and monthly mean of ΔO_x from the testing data sets for four representative months. **(a)** ATLAS ΔO_x and **(b)** the result of the SWIFT polynomial functions evaluated at the testing data. Grey bins contain no or too little data.

ΔO_x regime. Inside the polar vortex O_x climatology values are used. The Polar SWIFT module is intentionally switched off to investigate only the performance of the extrapolar module.

5.2 2-year simulation

Initially, the Extrapolar SWIFT module coupled to ATLAS was used in a simulation over a period of 2 years. With this short simulation we want to compare the development of the ozone layer in SWIFT to a reference simulation with AT-LAS. The goal of the comparison is to investigate the error or drift caused solely by the SWIFT polynomial functions.

Therefore the simulation conditions of both runs should be as similar as possible. To achieve this, the SWIFT simulation does not use trace gas climatologies for Cl_y, Br_y, NO_y, H_2O and O_x, but uses zonally and daily averaged trace gas VMRs instead. These daily values are compiled from the reference ATLAS simulation. Thus, apart from the averaging, the background trace gas fields are identical in both simulations. Further, the simulation covers a 2-year time period which co-incides with the period from which half of the training data originated (years 2005 and 2006). By selecting this simulation period we ensure that the SWIFT polynomial functions were trained with the stratospheric conditions of those years.
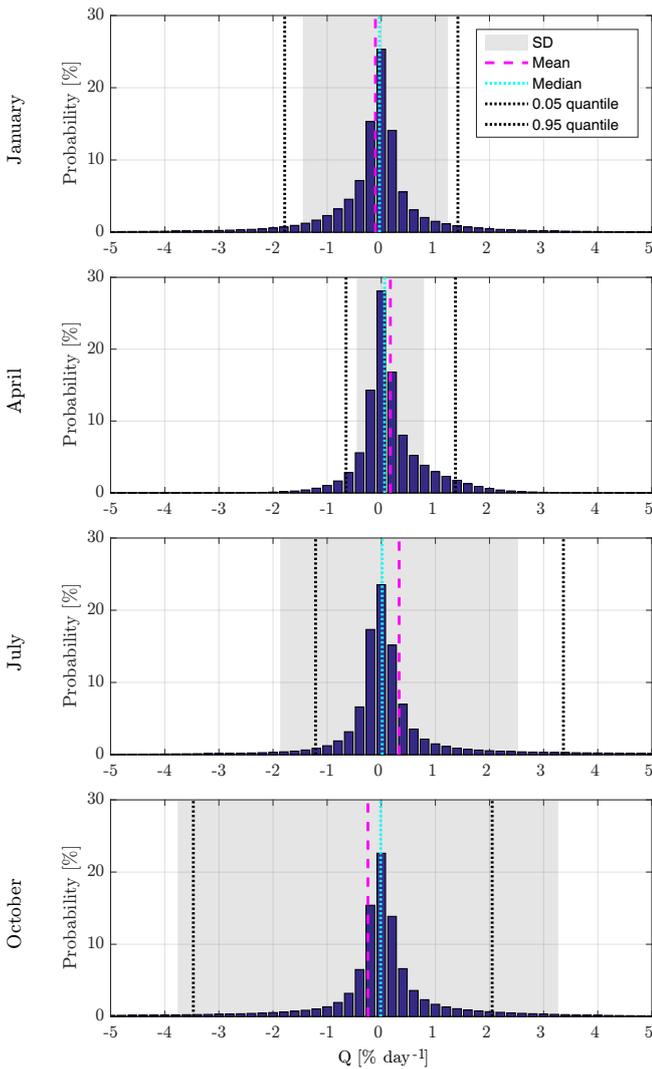
Figure 4. Probability distribution of error quantity Q of monthly testing data sets January, April, July and October. Grey shading indicates the standard deviation around the mean (pink dashed line). The dotted lines shows the median (cyan) and 5 and 95 % quantiles (black).

In other words, the errors cannot be caused by stratospheric variability unknown to SWIFT.

The panels in Figs. 5 and 6 show monthly averaged ozone concentrations for the 2-year SWIFT simulation (middle column). The reference ATLAS simulation is shown in the left column and the difference between the two in the right column. Since it is the ozone concentrations and total ozone columns that are crucial for the feedback of ozone to the model radiation, we have transformed the mixing ratios produced by SWIFT into ozone concentrations here. In the regions outside the ΔO_x regime, e.g. inside the polar vortex (white contour) or above the upper boundary (black dashed line), O_x values from the daily averaged O_x fields are used.

Figure 5 shows the entire annual cycle of 2005 (first simulation year) in bimonthly intervals. Figure 6 repeats the sequence for the second simulation year 2006. Throughout both years SWIFT shows excellent agreement with the ozone layer of the ATLAS simulation. The seasonal cycle of the ozone layer is very well reproduced. The average deviation oscillates at $\pm 0.2 \times 10^{12}$ per cm^{-3}. Over the course of the year 2005 the positive differences in the lower stratosphere of the Northern Hemisphere change sign to negative differences in the second half of the year. This pattern can also be observed in the second simulation year 2006. If the polynomial functions produce similar deviations in the same month of different years, we can attribute the deviations to a suboptimal approximation. However, the discussed deviations are in a region of strong meridional transport where the residence time of air parcels is sufficiently short so that no significant accumulation of errors occurs.

Further, it is unlikely for the monthly polynomial functions to produce the same deviations in exactly the same regions. If we compare the magnitude of the positive differences in January and March 2005 vs. January and March 2006 we see that the more positive deviations have switched from one month to the other. The variability of the magnitude can probably be attributed to the inter-annual stratospheric variability of the Northern Hemisphere, in particular the extent and lifetime of the polar vortex. In general the deviations of the year 2006 are not larger or more extensive than in 2005. Apparently no significant error is propagated from the preceding year to the following year.

5.3 10-year simulation

A SWIFT simulation over a period of 10 years demonstrates the stability of the model. The set-up for this simulation mimics the coupling of SWIFT to a GCM, although SWIFT is actually running in the ATLAS CTM. The trace gas climatologies for Cl_y, Br_y, NO_y and H_2O are the monthly climatologies described in Sect. 5.1. The simulation starts in November 1998 and continues until December 2008. This period encompasses both training data periods, the time between the two and a period after the last training data period. The bright blue curve in Fig. 7 shows the seasonal and interannual variation of the stratospheric ozone layer simulated by SWIFT. The depicted value is the integrated stratospheric ozone column in Dobson units from 15 to 32 km of pressure altitude. In order to observe a strong seasonal signal, we choose to display a location in the Northern Hemispheric mid-latitudes (Potsdam at 52.4° N, 13.0° E). The orange and green shaded years in Fig. 7 are the simulation periods of the training data set. The red curve in both periods shows the values of the reference ATLAS simulation. In both periods SWIFT reproduces the seasonal signal seen in ATLAS quite well. Especially in the green shaded patch the agreement between SWIFT and ATLAS seems to be as good as in the orange patch, although SWIFT was running continuously
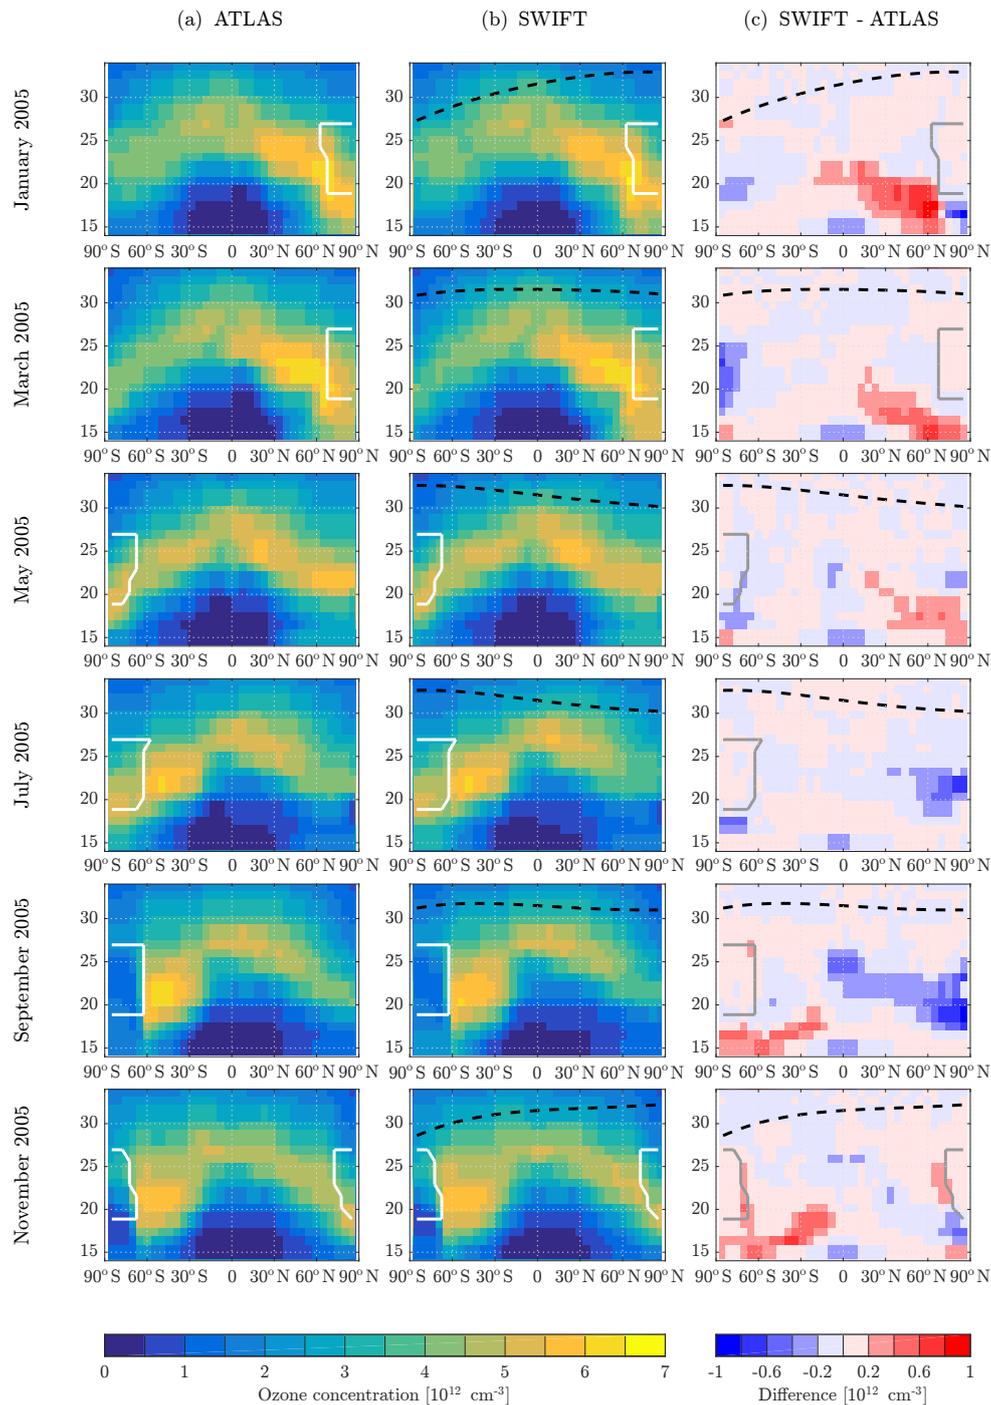
**Figure 5.** The 2005 zonal and monthly mean stratospheric ozone concentrations plotted in equivalent latitude vs. pressure altitude. **(a)** Reference simulation with ATLAS, **(b)** the SWIFT simulation and **(c)** the difference. The dashed black contour shows the upper boundary of the ΔO_x regime. The white contour indicates the location of the polar vortices.

for 4 years in between. To demonstrate this more clearly, the scatter plot in Fig. 8 shows daily averaged ozone columns of SWIFT on the x axis vs. the ones from ATLAS on the y axis. The colouring of the dots corresponds to the two time periods in Fig. 7. The scatter of data points from both periods over-laps entirely and the magnitude and distribution of deviations from the diagonal is identical. Clearly the errors of SWIFT did not accumulate over the course of the previous 6 years.

Beginning in autumn 2004 observational data from the microwave limb sounder Aura MLS are available and we addi-
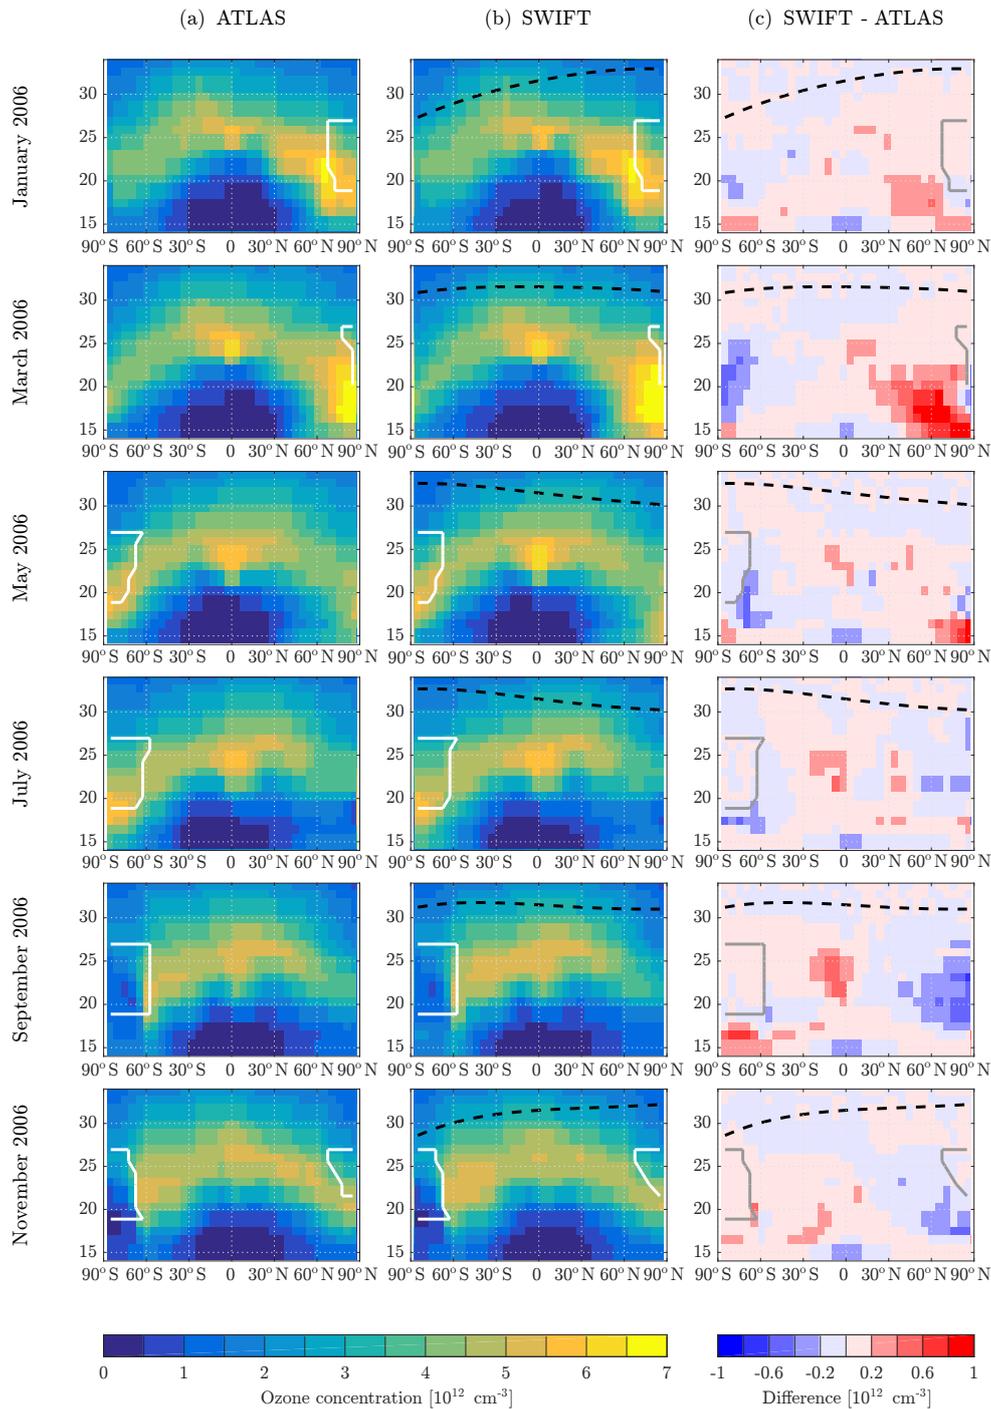
Figure 6. The 2006 zonal and monthly mean stratospheric ozone concentrations plotted in equivalent latitude vs. pressure altitude. See also Fig. 5.

tionally compare the SWIFT results with the Aura MLS observations (black line in Fig. 7). In autumn 2005 and 2006 ATLAS underestimates the ozone columns in comparison to the Aura MLS observations. Since SWIFT is trained with ATLAS data, SWIFT also reproduces this underestimation of about 30 DU and continues underestimating the autumn

stratospheric ozone columns in the years 2007 and 2008 (pink shaded patch). During the first half of each year, however, SWIFT matches the Aura MLS columns quite well and even captures the inter-annual variability shown by the observations (compare spring maximum 2007 vs. 2008). The scatter plot in Fig. 9 shows daily averaged ozone columns
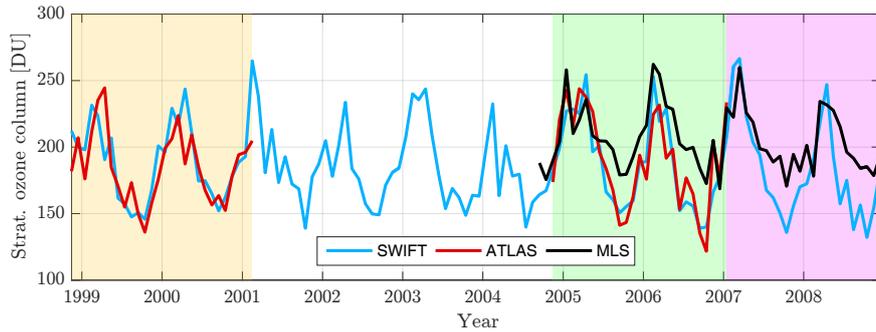
Figure 7. Monthly mean values of the stratospheric ozone column (15–32 km) over Potsdam (52.4° N, 13.0° E). The bright blue line shows the continuous 10-year simulation with SWIFT. The orange and green shaded patches are the periods from which the training data originate, and hence ATLAS data are available (red line). Beginning in fall 2004, Aura MLS data are available (black line) and during the pink period SWIFT and MLS are compared outside the training data period.
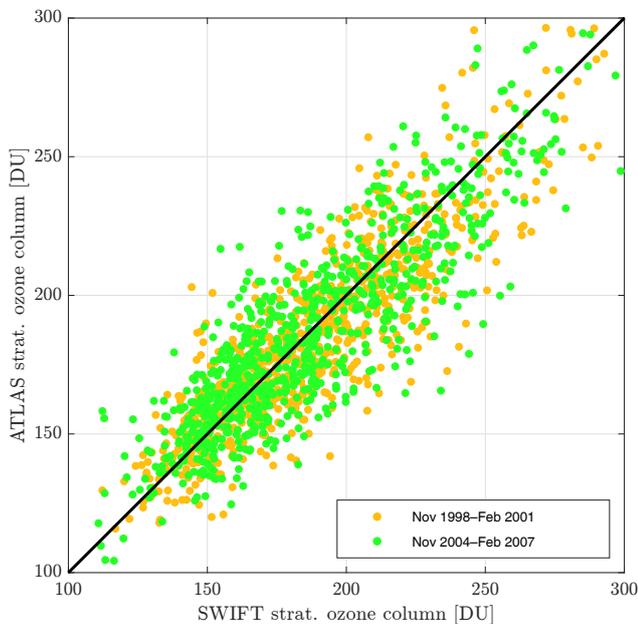
Figure 8. SWIFT vs. ATLAS scatter plot of daily averaged stratospheric ozone columns (15–32 km) over Potsdam. The orange and green dots correspond to the two training data periods in Fig. 7.

Figure 9. SWIFT vs. Aura MLS scatter plot of daily averaged stratospheric ozone columns (15–32 km) over Potsdam. The green dots correspond to the second training data period, and the pink dots correspond to the pink period in Fig. 7.

from the green and pink shaded years. Some amount of deviation in this figure is also caused by the difference in geolocation between the MLS profile and the selected location in the SWIFT simulation (Potsdam). Days on which no MLS measurement was taken in a 200 km radius of Potsdam are excluded, which reduces the total amount of days by about 50 %. Again the colouring of the dots corresponds to the periods in the time series (Fig. 7). As already seen in the monthly means in Fig. 7, SWIFT underestimates the smaller ozone columns (autumn values below 200 DU). Otherwise, the spread of the dots agrees well in both periods, proving that the SWIFT simulation is not less accurate outside the

training data period (pink) than under conditions which are part of the training data set (green).

5.4 Computational cost of Extrapolar SWIFT

The design of Extrapolar SWIFT enables full parallelization, since individual model nodes can independently evaluate the polynomial functions. A function consists of 30 to 100 polynomial terms, varying from month to month. Per model node and time step, three polynomial functions have to be evaluated, one domain polynomial and two ΔO_x polynomial functions for the interpolation between two months. During

the preparation of this paper Extrapolar SWIFT was coupled to the climate model ECHAM6.3. The Fortran SWIFT code is not fully optimized yet and the current estimates on the computation time are preliminary. An initial estimate of the increase in computation time caused by Extrapolar SWIFT is roughly 10 %. In comparison to an ECHAM version employing full stratospheric chemistry (ECHAM MESSy Atmospheric Chemistry model, or EMAC), the ECHAM + Extrapolar SWIFT requires 6–8 times less computation time (only estimated).

The version of Extrapolar SWIFT coupled to the ATLAS CTM is implemented in MATLAB because the ATLAS model was written in MATLAB. SWIFT in ATLAS is not optimized for speed and the evaluation of the polynomials is computed on a single core. However, when comparing the full stratospheric chemistry scheme of ATLAS vs. the evaluation of the SWIFT polynomial functions, the ozone layer can be computed 10^4 times faster than in the CTM.

6 Conclusions

The Extrapolar SWIFT model is a numerically efficient ozone chemistry scheme for global climate models. Its primary goal is to enable the interactions between the ozone layer, radiation and climate, while imposing a low computational burden to the GCM it is coupled to. We accomplished this by approximating the rate of change of ozone of the detailed chemistry model ATLAS by using algebraic equations. Orthogonal polynomial functions of fourth degree are used to approximate the rate of change of ozone over 24 h. An automated and optimized procedure approximates one globally valid polynomial function to a monthly training data set. In our repro-modelling approach we reduce the dimensionality of the model through exploitation of the covariance between variables. The polynomial functions are a function of only nine *basic* variables (latitude, pressure altitude, temperature, overhead ozone column, total chlorine, total bromine, nitrogen oxide family, water vapour and the ozone field). At the same time, all physical and chemical processes contained in the full model output are parameterized in the repro-model.

Running the Extrapolar SWIFT model requires only the 12 monthly polynomial functions and information about the nine *basic* variables. The domain of the polynomial function is defined by the nine-dimensional training data set. A wide range of stratospheric variability needs to be included in the training data set to increase the robustness of the polynomial functions. We have shown that the SWIFT model can cope with a certain degree of unknown variability induced, for example, by climate change. We estimate that the polynomial functions can handle changes of up to a 10 % increase or decrease in stratospheric chlorine loading without adjusting the current training data set. More extreme changes, e.g. a 50 % reduction of chlorine, requires an extension of the training data with values of disturbed chemistry simulations. For handling occasional outliers, i.e. combinations of the nine *basic* variables outside the domain of definition, Extrapolar SWIFT includes a procedure to prevent extrapolation of the polynomial functions.

Simulations with the Extrapolar SWIFT model coupled to the ATLAS CTM have shown good agreement to the reference model ATLAS. The stability of SWIFT has been proven with a simulation over a 10-year period in which SWIFT was validated against model and observational references. Errors did not accumulate over the extended simulation period. Average deviations of the integrated stratospheric ozone column (15–32 km) are ±15 DU between ATLAS and SWIFT. The comparison to Aura MLS measurements showed an equally good agreement with Extrapolar SWIFT, except for the periods of underestimation of the stratospheric ozone column in autumn. This underestimation, however, is a bias that originates from the source model ATLAS. The computation of the solution of a polynomial function with up to 100 terms is significantly faster than solving a chemical differential equation system. Extrapolar SWIFT requires 10^4 times less computation time than the chemistry scheme of the ATLAS CTM.

References

Baldwin, M. P., Dameris, M., and Shepherd, T. G.: How Will the Stratosphere Affect Climate Change?, Science, 316, 1576–1577, https://doi.org/10.1126/science.1144303, 2007.

Calvo, N., Polvani, L., and Solomon, S.: On the surface impact of Arctic stratospheric ozone extremes, Environ. Res. Lett., 10, 094003, https://doi.org/10.1088/1748-9326/10/9/094003, 2015.

Cariolle, D. and Teyssèdre, H.: A revised linear ozone photochemistry parameterization for use in transport and general circulation models: multi-annual simulations, Atmos. Chem. Phys., 7, 2183–2196, https://doi.org/10.5194/acp-7-2183-2007, 2007.

Dee, D., Uppala, S., Simmons, A., Berrisford, P., Poli, P., Kobayashi, S., Andrae, U., Balmaseda, M., 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.: The ERA-Interim reanalysis: Configuration and performance of the data assimilation system, Q. J. Roy. Meteor. Soc., 137, 553–597, 2011.

Eyring, V., Shepherd, T. G., and Waugh, D. W.: SPARC CCMVal Report on the Evaluation of Chemistry-Climate Models, Tech. rep., SPARC Office, available at: http://www.sparc-climate.org/publications/sparc-reports/ (last access: 25 February 2018), 2010.

Eyring, V., Arblaster, J. M., Cionni, I., Sedláček, J., Perlwitz, J., Young, P. J., Bekki, S., Bergmann, D., Cameron-Smith, P., Collins, W. J., Faluvegi, G., Gottschaldt, K.-D., Horowitz, L. W., Kinnison, D. E., Lamarque, J.-F., Marsh, D. R., Saint-Martin, D., Shindell, D. T., Sudo, K., Szopa, S., and Watanabe, S.: Long-term ozone changes and associated climate impacts in CMIP5 simulations, J. Geophys. Res.-Atmos., 118, 5029–5060, https://doi.org/10.1002/jgrd.50316, 2013.

Golub, G. H. and Van Loan, C. F.: Matrix computations, Johns Hopkins studies in the mathematical sciences, 3rd Edn., Johns Hopkins University Press, Baltimore, 1996.

Grooß, J.-U. and Russell III, J. M.: Technical note: A stratospheric climatology for O_3, H_2O, CH_4, NO_x, HCl and HF derived from HALOE measurements, Atmos. Chem. Phys., 5, 2797–2807, https://doi.org/10.5194/acp-5-2797-2005, 2005.

Grooß, J.-U., Günther, G., Konopka, P., Müller, R., McKenna, D., Stroh, F., Vogel, B., Engel, A., Müller, M., Hoppel, K., Bevilacqua, R., Richard, E., Webster, C. R., Elkins, J. W., Hurst, D. F., Romashkin, P. A., and Baumgardner, D. G.: Simulation of ozone depletion in spring 2000 with the Chemical Lagrangian Model of the Stratosphere (CLaMS), J. Geophys. Res., 107, 8295, https://doi.org/10.1029/2001JD000456, 2002.

Hsu, J. and Prather, M. J.: Stratospheric variability and tropospheric ozone, J. Geophys. Res., 114, D06102, https://doi.org/10.1029/2008JD010942, 2009.

IPCC: Climate change 2013: the physical science basis: Working Group I contribution to the Fifth assessment report of the Intergovernmental Panel on Climate Change, Cambridge University Press, New York, 2014.

Kreyling, D.: Das extrapolare SWIFT-Modell: schnelle stratosphärische Ozonchemie für globale Klimamodelle, PhD thesis, Freie Universität Berlin, available at: http://www.diss.fu-berlin.de/diss/receive/FUDISS%5fthesis%5f000000102777 (last access: 25 February 2018), 2016.

Lait, L. R.: An alternative form for potential vorticity, J. Atmos. Sci., 51, 1754–1759, 1994.

Lowe, R. M. and Tomlin, A. S.: The application of repro-modelling to a tropospheric chemical model, Environ. Modell. Softw., 15, 611–618, https://doi.org/10.1016/S1364-8152(00)00056-6, 2000.

McLinden, C. A., Olsen, S. C., Hannegan, B., Wild, O., Prather, M. J., and Sundet, J.: Stratospheric ozone in 3-D models: A simple chemistry and the cross-tropopause flux, J. Geophys. Res., 105, 14653–14665, https://doi.org/10.1029/2000JD900124, 2000.

Nowack, P. J., Luke Abraham, N., Maycock, A. C., Braesicke, P., Gregory, J. M., Joshi, M. M., Osprey, A., and Pyle, J. A.: A large ozone-circulation feedback and its implications for global warming assessments, Nat. Clim. Change, 5, 41–45, https://doi.org/10.1038/nclimate2451, 2014.

Rex, M., Salawitch, R. J., Deckelmann, H., von der Gathen, P., Harris, N. R. P., Chipperfield, M. P., Naujokat, B., Reimer, E., Allaart, M., Andersen, S. B., Bevilacqua, R., Braathen, G. O., Claude, H., Davies, J., De Backer, H., Dier, H., Dorokhov, V., Fast, H., Gerding, M., Godin-Beekmann, S., Hoppel, K., Johnson, B., Kyrö, E., Litynska, Z., Moore, D., Nakane, H., Parrondo, M. C., Risley, A. D., Skrivankova, P., Stübi, R., Viatte, P., Yushkov, V., and Zerefos, C.: Arctic winter 2005: Implications for stratospheric ozone loss and climate change, Geophys. Res. Lett., 33, L23808, https://doi.org/10.1029/2006GL026731, 2006.

Rex, M., Kremser, S., Huck, P., Bodeker, G., Wohltmann, I., Santee, M. L., and Bernath, P.: Technical Note: SWIFT – a fast semi-empirical model for polar stratospheric ozone loss, Atmos. Chem. Phys., 14, 6545–6555, https://doi.org/10.5194/acp-14-6545-2014, 2014.

Sander, S. P., Abbatt, J., Barker, J. R., Burkholder, J. B., Friedl, R. R., Golden, D. M., Huie, R. E., Kolb, C. E., Kurylo, M. J., Moortgat, G., Orkin, V. L., and Wine, P. H.: Chemical kinetics and photochemical data for use in atmospheric studies, Evaluation No. 17, JPL Publication 10-6, available at: http://jpldataeval.jpl.nasa.gov (last access: 25 February 2018), 2011.

Shampine, L. F. and Reichelt, M. W.: The MATLAB ODE Suite, SIAM J. Sci. Comput., 18, 1–22, 1997.

Spivakovsky, C. M., Wofsy, S. C., and Prather, M. J.: A numerical method for parameterization of atmospheric chemistry: Computation of tropospheric OH, J. Geophys. Res., 95, 18433–18439, https://doi.org/10.1029/JD095iD11p18433, 1990.

Thompson, D. W. J. and Solomon, S.: Interpretation of Recent Southern Hemisphere Climate Change, Science, 296, 895–899, https://doi.org/10.1126/science.1069270, 2002.

Turányi, T.: Parameterization of reaction mechanisms using orthonormal polynomials, Comput. Chem., 18, 45–54, https://doi.org/10.1016/0097-8485(94)80022-7, 1994.

Waters, J., Froidevaux, L., Harwood, R., Jarnot, R., Pickett, H., Read, W., Siegel, P., Cofield, R., Filipiak, M., Flower, D., Holden, J., Lau, G., Livesey, N., Manney, G., Pumphrey, H., Santee, M., Wu, D., Cuddy, D., Lay, R., Loo, M., Perun, V., Schwartz, M., Stek, P., Thurstans, R., Boyles, M., Chandra, K., Chavez, M., Gun-Shing Chen, Chudasama, B., Dodge, R., Fuller, R., Girard, M., Jiang, J., Yibo Jiang, Knosp, B., LaBelle, R., Lam,

J., Lee, K., Miller, D., Oswald, J., Patel, N., Pukala, D., Quintero, O., Scaff, D., Van Snyder, W., Tope, M., Wagner, P., and Walch, M.: The Earth observing system microwave limb sounder (EOS MLS) on the aura Satellite, IEEE T. Geosci. Remote, 44, 1075–1092, https://doi.org/10.1109/TGRS.2006.873771, 2006.

Wohltmann, I. and Rex, M.: The Lagrangian chemistry and transport model ATLAS: validation of advective transport and mixing, Geosci. Model Dev., 2, 153–173, https://doi.org/10.5194/gmd-2-153-2009, 2009.

Wohltmann, I., Lehmann, R., and Rex, M.: The Lagrangian chemistry and transport model ATLAS: simulation and validation of stratospheric chemistry and ozone loss in the winter 1999/2000, Geosci. Model Dev., 3, 585–601, https://doi.org/10.5194/gmd-3-585-2010, 2010.

Wohltmann, I., Lehmann, R., and Rex, M.: Update of the Polar SWIFT model for polar stratospheric ozone loss (Polar SWIFT version 2), Geosci. Model Dev., 10, 2671–2689, https://doi.org/10.5194/gmd-10-2671-2017, 2017.