Investigating radiative adjustments to anthropogenic aerosol perturbations using novel modelling techniques

Max Reeves Coleman

Department of Meteorology
School of Mathematical, Physical and Computational Sciences

This thesis is submitted for the degree of Doctor of Philosophy
PhD in Atmosphere, Oceans and Climate

March 2024

Abstract

Aerosols overall have a strong cooling effect on Earth's climate that masks warming from greenhouse gases, but likely reductions in future aerosol emissions will reduce this cooling. The magnitude of this reduction is uncertain however, partly because of the complex radiative effects of aerosol. To reduce the uncertainty of future climate change requires improved understanding and quantification of the mechanisms and magnitudes of aerosol radiative effects, particularly radiative adjustments. This thesis investigates variations on two existing methods of decomposing aerosol radiative effects, applying them to pre-industrial (1850) to present-day (2014) perturbations of anthropogenic sulphate and black carbon aerosol emissions using the UK Earth System Model (UKESM1). The first method applies nudging to model horizontal winds or horizontal winds and potential temperature to decompose aerosol effects mechanistically into circulation-mediated and atmospheric temperature-mediated adjustments. The second method, a variation of the partial radiative perturbation (PRP) method, decomposes aerosol radiative effects into contributions from individual atmospheric variables while conserving the total radiative effect. Circulation-mediated adjustments are found to be significant for sulphate (0.10$\pm$0.08 W m^{-2}; 7% of ERF) but not black carbon (-0.04$\pm$0.07 W m^{-2}; 8% of ERF) perturbation, while atmospheric temperature-mediated adjustments are significant for both sulphate (0.14$\pm$0.04 W m^{-2}; 10% of ERF) and black carbon (-0.25$\pm$0.08 W m^{-2}; 47% of ERF) perturbations, with both arising primarily from cloud adjustments. However, significant uncertainties are found with the nudging method that reduce confidence in quantification of temperature-mediated adjustments. The partial radiative perturbation method is difficult to implement but the variation applied here successfully conserves the total radiative effect, offering improvements over the existing method. Both methods offer insights into aerosol radiative effects and address gaps in existing techniques, but the nudging method suffers significant limitations. Isolating aerosol circulation-mediated adjustments does however help fill a gap in the literature by indicating that applying nudging to reduce simulation integration lengths omits significant adjustments due to suppressing circulation responses.

Declaration of Authorship

Declaration: I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Max Reeves Coleman

Acknowledgements

This work would not have been possible without the guidance, aid, and support of many people. Firstly, I would like to thank my supervisors, Bill Collins, Keith Shine, Nicolas Bellouin and Fiona O'Connor for their massive contributions. This included, amongst many other contributions, everything from guiding the conception and evolving aims of the project, to getting my first climate model simulation running, to providing encouragement and suggestions whenever obstacles arose. Next I would like to thank Mohit Dalvi for helping me to get the Unified Model nudging code working for an application it was never quite intended for and for being a constant source of knowledge on the inner workings of the Unified Model, answering the many questions I sent his way. Likewise, significant thanks go to James Manners for helping me setup and run the SOCRATES code and for helping with the seemingly endless debugging I required. My thanks also go to Tom Dunstan and Chris Smith for providing code that helped in developing the code for running PRP calculations with SOCRATES. Next I would like to thank my monitoring committee, Claire Ryder, Helen Dacre and Janet Barlow for useful guidance on the direction of this project and for always helping me to see my work in a more positive light.

On a more general level, my thanks go to all those who maintain and provide support for the Monsoon2 high performance computing system and the Reading Academic Computing Cluster, which the majority of this work was performed using. My thanks go to all the members of the Reading University Meteorology Department for the many useful training opportunities, seminars, lectures and conversations that have helped shape this work, and particularly all of the other PhD students who have helped make this journey more enjoyable and rewarding. My thanks also go to the SCENARIO DTP, and by extension NERC, and the Met Office who funded me and my work, making it possible.

Finally, I would like to thank my friends, my family, and my partner, Amy, all of whom have supported me with their everyday kindness and presence throughout both the challenges and successes of my time working on this thesis.

MC

Contents

List of Figures

Chapter 1

Introduction and literature review

1.1 Context and motivation

Anthropogenic climate change is a significant issue facing society. Climate change is expected to cause large shifts in temperature and precipitation patterns and these changes are expected to result in significant consequences for all of society globally, including changing frequency and intensity of natural hazards, sea level rise, biodiversity loss and economic and social disruption (IPCC, 2023a). To address climate change it is crucial to understand how anthropogenic activity affects Earth's climate presently, how that is likely to change in future, and how these effects can be mitigated.

Present day anthropogenic aerosol emissions overall cause a strong cooling effect, estimated to mask 0.5 °C (5 to 95 % confidence interval: 0.22 to 0.96 °C) of warming mostly due to greenhouse gas forcing relative to 1750 aerosol emissions (Forster et al., 2021). Global aerosol emissions are expected to reduce under increasingly stringent air quality policies in most future scenarios, with some aerosol emissions already declining (Szopa et al., 2021). As figure 1.1 illustrates, this is predicted to lead to an increase in warming in most of the core shared socio-economic pathways (SSPs). In particular, it can be seen that in the SSPs that avoid warming greater than 2 °C (SSP1-1.9 and SSP1-2.6) aerosol emission changes are expected to contribute between 0.1 to 0.7 °C of warming by 2100 relative to 2019 (Szopa et al., 2021). This is a potentially very large contribution towards the 1.5 or 2 °C limits on warming by the end of 21st century recommended by the Paris Agreement (IPCC, 2023a). Understanding the amount of warming and other climatic effects that reducing aerosol emissions will cause is therefore of great importance for mitigation of, and adaptation to, climate change.

To understand the effects of aerosol emission changes on climate it is useful to first understand their effects on the energy budget of Earth's atmosphere. This is because changes in the energy budget drive changes in aspects of the climate that have relevant impacts to society and the environment, such as temperature and precipitation, but energy budget changes can also be quantified without needing to include additional uncertainties and complexities associated with modelling the responses in these impact variables (Myhre et al., 2013; Sherwood et al., 2015; Forster et al., 2021). Likewise focusing on energy budgets is useful to compare the effects of aerosol changes with other natural and anthropogenic factors affecting climate (often referred

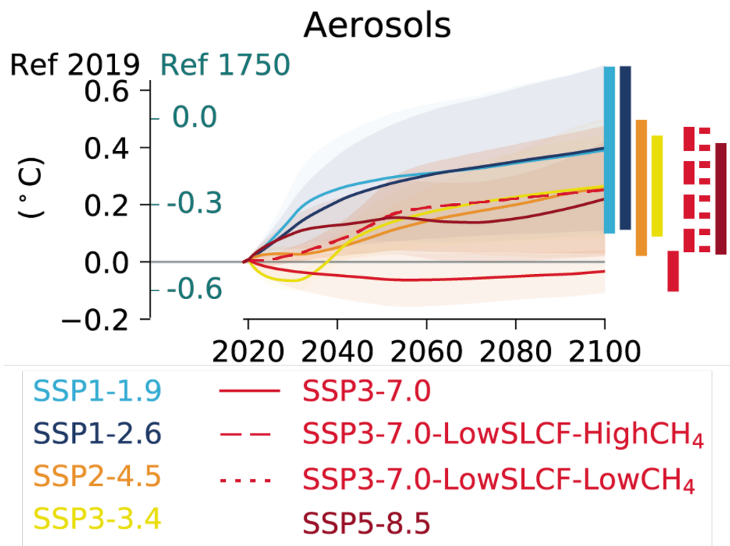

Figure 1.1: Predicted aerosol contribution to global surface air temperature change across a range of SSPs relative to 2019 and 1750. The predicted pathways are based on historic and future ERF assessed in Forster et al. (2021) with temperature responses determined from an impulse response function using an equilibrium climate sensitivity of 3.0 K for a doubling of CO_2. Uncertainty bars are 5-95% confidence interval. Modified from Szopa et al. (2021).

to as 'climate forcing agents'). To quantify and compare the effects of forcing agents on the Earth's energy budget, the concept of radiative forcing is therefore typically used: the change in net downward radiative flux in the atmosphere due to perturbation of a forcing agent.

However the effects of a forcing agent on radiative fluxes are complex. As illustrated in figure 1.2, these effects include (1) the forcing agent of interest causing an instantaneous radiative forcing by directly interacting with radiation; and (2) adjustments in other components of the Earth system in response to the forcing agent perturbation causing additional radiative forcing (Boucher et al., 2013; Forster et al., 2021). These adjustments (which exclude responses to surface temperature change, or 'radiative feedbacks') are termed 'radiative adjustments' and can significantly alter the total radiative forcing caused by the forcing agent. This is particularly true for aerosols, which include large radiative adjustments via numerous process, most significantly via their effects on clouds (Boucher et al., 2013; Sherwood et al., 2015; Forster et al., 2021). Understanding aerosol radiative adjustments and their contribution to aerosol radiative forcing is therefore critical to accurately modelling the climatic effects of future aerosol emissions.

Due to their indirect nature however, there is presently less certainty about the magnitudes and mechanisms of aerosol radiative adjustments than the direct effects of aerosols on radiation (Forster et al., 2021). Radiative adjustments therefore contribute a large amount of uncertainty to estimates of aerosol radiative forcing. This makes estimates of the radiative forcing of aerosol emissions reductions more uncertain than other major forcing agents, as illustrated in figure 1.3a, in which the 1750 to 2019 total aerosol radiative forcing uncertainty due to aerosols is at least twice that of other major forcing agents (Forster et al., 2021).

This uncertainty in aerosol radiative forcing due to aerosol radiative adjustments results in uncertainty in the temperature response due to aerosol emission changes. Indeed, not only does the 0.1 to 0.7 °C future warming due to aerosol emissions changes in SSP1-1.9 and SSP1-2.6

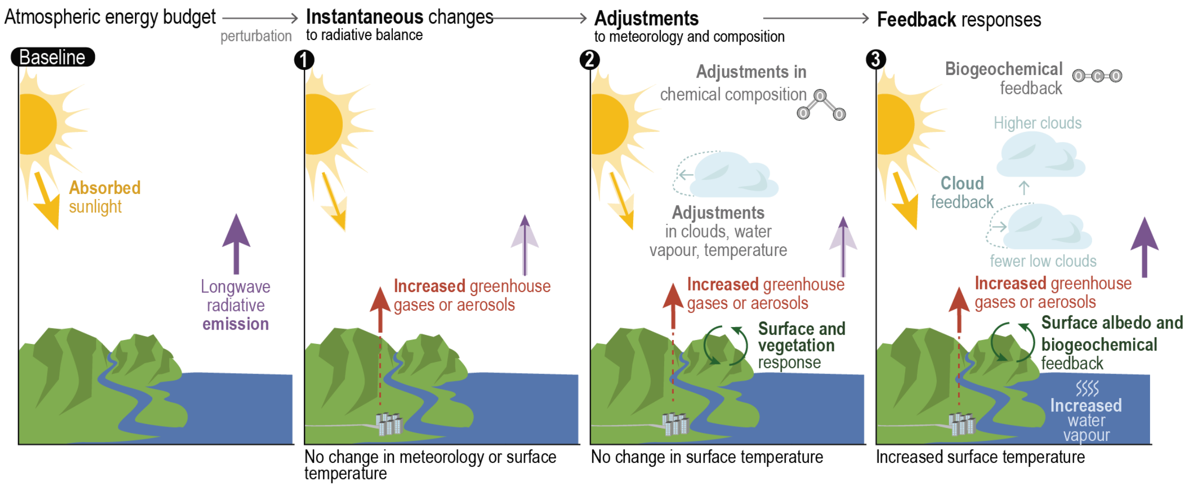

Figure 1.2: Illustration of the effects of forcing agents on the Earth's energy budget. The energy budget particularly depends upon absorption and scattering of incoming sunlight and emission and absorption of longwave radiation (baseline) which may be affected instantaneously by a perturbation of a forcing agent such as aerosol or greenhouse gases (1). Perturbation of the forcing agent prompts adjustments to meteorology and atmospheric composition which cause further changes in the energy budget (2). The changes to the Earth system as a result of the instantaneous radiative effects and adjustments cause a response in the surface temperature which then prompts further changes to the Earth system and energy budget (3). From IPCC (2023b).

mentioned earlier represent a potentially very large warming, it also represents a very large range of uncertainty. Some of this uncertainty will be due to uncertainty in surface temperature responses and radiative feedbacks rather than radiative adjustments. However, figure 1.3b illustrates that the majority of the uncertainty in aerosol contributions to surface temperature change between 1750 and 2019 is in the radiative forcing (approximately two thirds), and the relative contributions to uncertainty can be expected to be similar for future aerosol emission changes as past ones.

In summary, it has been established that future aerosol emission changes are expected to have significant but uncertain effects on future climate change, particularly surface temperature change. The magnitude and uncertainty of these climate effects due to aerosols are significantly contributed to by radiative adjustments, which significantly alter the radiative forcing due to aerosols. Therefore, because accurately and precisely estimating the climate effects expected to occur due to future aerosol emissions is important for informing climate mitigation and adaptation efforts, it is important to increase understanding of the magnitude and mechanisms of aerosol radiative adjustments.

1.2 Research aims and thesis outline

With the context above in mind, the overarching purpose for this thesis is to improve understanding of aerosol radiative adjustment processes and the magnitude of their radiative effects. To achieve this, this thesis implements variations on two methodologies for investigating radiative adjustments: firstly, the use of nudging as a mechanism denial tool; and secondly, a variation of the partial radiative perturbation (PRP) method. These methods are applied in model experiments perturbing either sulphate (SU) or black carbon (BC) aerosol emissions.

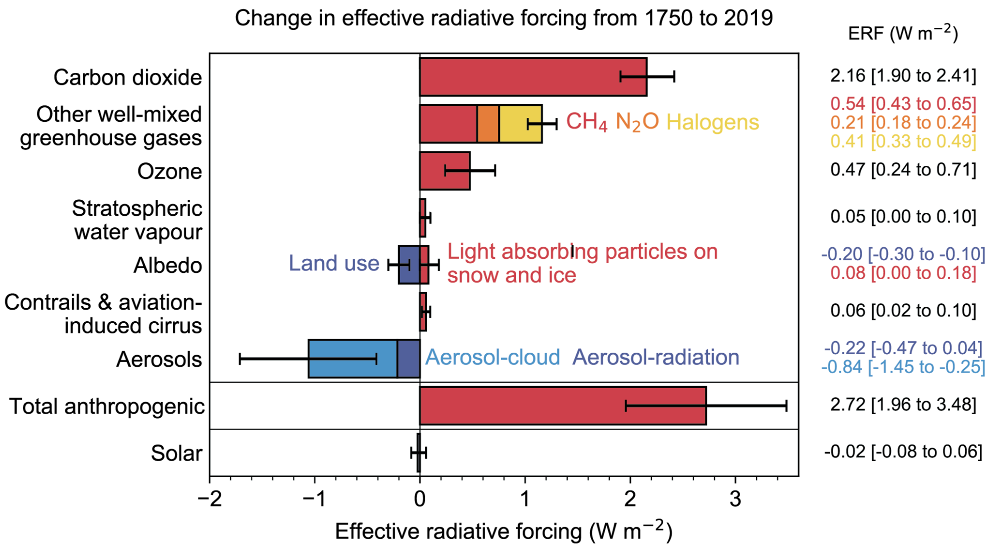

(a) SU

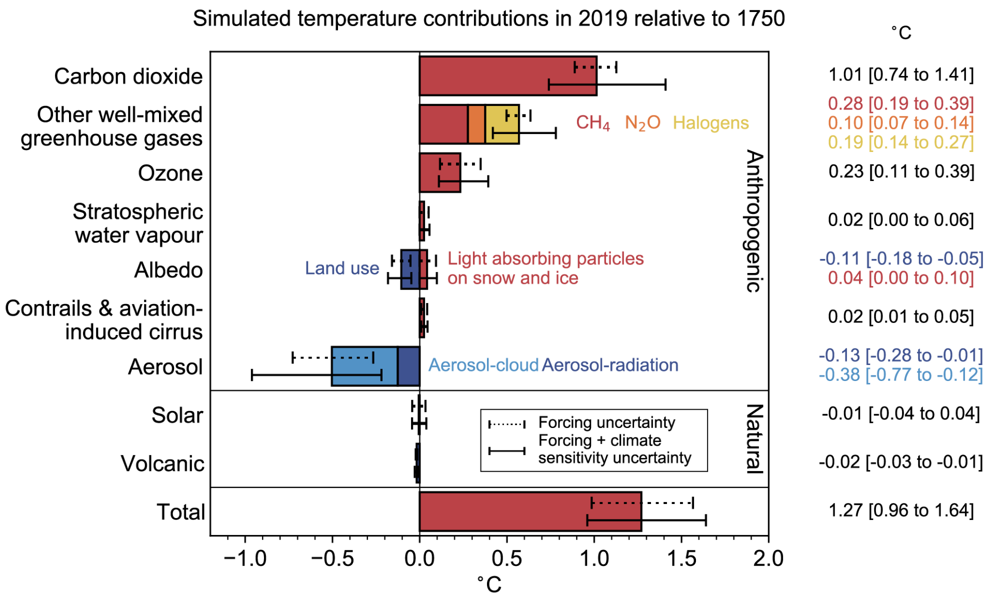

(b) BC

Figure 1.3: Assessed (a) radiative forcing and (b) contributions to surface temperature change due to perturbation of different categories of forcing agents from 1750 to 2019 conditions in the IPCC AR6 report. Solid coloured bars represent the best estimates of total (or 'effective') radiative forcing or surface temperature change due to each forcing agent and 5th/95th percentiles are illustrated by error bars. In (b) the dashed error bars represent 5th/95th percentiles due to uncertainty from radiative forcing alone, while solid error bars represent 5th/95th percentiles due to uncertainties from radiative forcing and the temperature responses of the climate. From Forster et al. (2021).

This allows the evaluation of the strengths and weaknesses of these methods for investigating aerosol radiative adjustments, and the gaining of new knowledge about radiative adjustments from SU and BC perturbation specifically. The specific research aims can thus be summarised as follows:

- Firstly, to implement and test variations of the nudging mechanism denial methodology and PRP methodology, assessing their usefulness for the purpose of understanding and quantifying radiative adjustments.
- Secondly, to apply both methodologies to anthropogenic SU and BC perturbations to gain

insights into radiative adjustments arising from these two aerosol climate forcing agents.

This thesis is presented as follows. For the remainder of this chapter a review of the literature relevant to this work is given, particularly concerning aerosol radiative adjustments and methods to investigate them. The rest of the thesis is then divided between the two methods used for investigating radiative adjustments here: two chapters are devoted to the nudging mechanism denial methodology, and one chapter to the PRP method. Chapter 2 evaluates the effects of using different settings for the nudging mechanism denial methodology and seeks to determine the optimal nudging setup to use for the application of this methodology. Chapter 3 then applies this methodology to anthropogenic SU and BC, evaluating the overall effectiveness of the method and drawing conclusions about SU and BC radiative adjustments. Chapter 4 applies the PRP method to anthropogenic BC perturbation, again evaluating the effectiveness of the method and drawing conclusions about BC radiative adjustments. Finally chapter 5 considers the conclusions that can be drawn from this work as a whole, their wider significance, and recommendations for future research to build on this work.

1.3 Background literature

This section offers a review of the literature of particular relevance to this work. Firstly a review of the radiative forcing-adjustment-feedback framework is given, highlighting what specifically is considered to be a radiative adjustment and the utility of the concept. Then a review of the key methods for determining radiative forcing and adjustments using general circulation models (GCMs) or Earth system models (ESMs) is given. This is followed by a review of methods applied in the literature to investigate specific radiative adjustments. Next is a deeper review of the application of nudging in GCMs and ESMs, particularly focusing on the choices concerning nudging settings and their consequences. Finally, a review of the current understanding of SU and BC radiative forcing and adjustments is given.

1.3.1 Forcing-adjustment-feedback framework for understanding climate responses

Assessments of the contribution of different climate forcing agents to climate change typically consider the effect of such agents on the atmosphere's energy budget (Forster et al., 2021). This is because perturbations to the atmosphere's energy budget can be linked to changes in temperature and precipitation, among other factors, which are key variables for understanding the effects of climate change. Forcing agents perturb this energy budget primarily via radiation, perturbing the transmission of shortwave (SW) and longwave (LW) radiation through the atmosphere via absorption and/or scattering. Therefore, a key metric used to compare different climate forcing agents is radiative forcing, which is the change in net downwards radiative fluxes due to a perturbation of a forcing agent (Sherwood et al., 2015; Forster et al., 2021). Often the perturbation is defined as the change between pre-industrial (PI) and present-day (PD). The net flux change is usually measured at the top of the Earth's atmosphere (TOA) (Forster et al., 2021). Estimates of radiative forcing can be determined for a range of climate forcing agents to compare the magnitude of their effects on climate.

Radiative forcing can be approximately related to global mean surface temperature change by a simple linear equation:

$$\Delta N = \Delta F + \alpha \Delta T \tag{1.1}$$

where ΔN is the change in net radiative flux balance at TOA measured in W m^{-2}, ΔF is the radiative forcing imparted by some forcing agent measured in W m^{-2}, ΔT is the surface temperature response to the forcing measured in K, and α is the feedback parameter (Sherwood et al., 2015; Forster et al., 2021) measured in W m^{-2} K^{-1}. This forcing-feedback framework can be understood as follows. When a radiative forcing ΔF is applied to the climate system, it causes an equal imbalance in the TOA flux balance ΔN. In response, the surface temperature changes, changing the outgoing longwave radiation (the Planck effect). However the change in ΔN in response to surface temperature increase depends on many processes, termed climate feedbacks. These feedbacks are described in α and can be seen as modulating the climate feedback to the radiative forcing. Typically α is negative and so the total climate feedbacks act to return the climate system to equilibrium (that is, return ΔN to zero). Therefore, if the radiative forcing and radiative feedbacks can be quantified, the surface temperature response to that radiative forcing can be determined.

For this equation to be useful, ideally the magnitude of radiative forcing between different forcing agents can be compared equally, with the same magnitude radiative forcing causing the same temperature response. However, early studies of radiative forcing tended to assess just the radiative forcing caused by a change in a forcing agent (e.g. aerosol concentration) without allowing the atmosphere to respond, or only allowing the stratosphere to respond. When using this 'instantaneous radiative forcing' (IRF) to compare forcing agents, their 'efficacy' in causing temperature change varies significantly (Hansen et al., 2005). This is primarily because, as well as interacting directly with radiation, many forcing agents affect other components of the climate system that interact with radiation. These responses are termed 'radiative adjustments'. Using only the IRF in equation 1.1 therefore omits some of the radiative forcing caused by a forcing agent.

To understand radiative adjustments it is useful to give an example. Figure 1.4 illustrates the pathways by which greenhouse gases and aerosols effect the TOA radiative energy balance (Boucher et al., 2013). In both cases, the forcing agent interacts with radiation: for example BC absorbs SW radiation, and so increased BC exerts a positive instantaneous radiative forcing. The positive radiative forcing from BC generally causes an increase in tropospheric temperature (Li et al., 2022). The increased tropospheric temperature will cause an increase in outgoing LW radiation; this is a radiative adjustment. In this case, this radiative adjustment is negative and so acts to reduce the total radiative forcing caused by BC. The tropospheric temperature increase will also then cause changes in circulation pattens and cloud fields which will cause further changes in TOA radiative balance, and so can also be described as radiative adjustments

(Sherwood et al., 2015). Furthermore, after ageing, BC can also act as cloud droplet condensation nuclei (CCN) in clouds (Bond et al., 2013). This mechanism causes an adjustment in cloud properties 'indirectly' (independent of the direct absorption of radiation by BC) that then cause their own radiative effect, and hence are also radiative adjustments. Thus it can be seen that, particularly for aerosols that interact with clouds in their role as CCN ('aerosol-cloud interactions' or aci), a forcing agent can cause a number of adjustments, which can lead to further adjustments, which can significantly alter the total radiative forcing compared to the IRF.

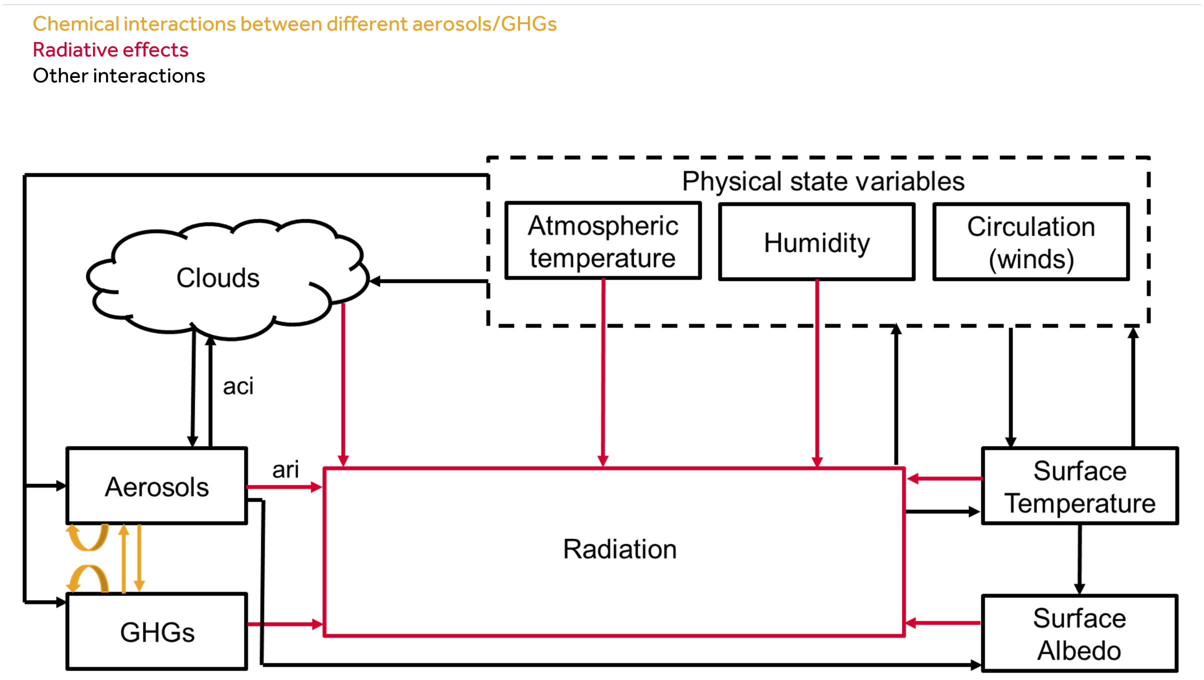

Figure 1.4: Illustration of interactions between aerosols and greenhouse gases (GHGs) and different components of the climate system that affect radiation. Changes in components of the system that arise due to initial perturbation of aerosol or GHG and their direct radiative effects that cause additional radiative forcing are radiative adjustments; while changes due to surface temperature change are radiative feedbacks. Radiative adjustments and feedbacks share some similar pathways. It is worth noting that interactions with the biosphere as well as feedbacks affecting emissions of aerosols and GHGs are not represented in this schematic.Furthermore, while aerosols do affect surface albedo directly via their deposition on surfaces, this effect is not represented in the models used in this work. After Boucher et al. (2013).

Including these radiative adjustments in the radiative forcing term in equation 1.1 brings the efficacy of different forcing agents closer to unity (Rotstayn and Penner, 2001; Shine et al., 2003; Hansen et al., 2005; Forster et al., 2021). This makes comparisons between forcing agents more useful, since they result in a similar global mean temperature response. This measure of radiative forcing that includes both the IRF and radiative forcing due to radiative adjustments is termed the 'effective radiative forcing' (ERF) (Myhre et al., 2013).

It is important to highlight that ERF does not include any responses to surface temperature which are considered instead to be radiative feedbacks (Boucher et al., 2013; Forster et al., 2021). Radiative feedbacks are considered separately from adjustments because they can be understood as responses directly to the surface temperature response, for the most part independently of the forcing agent causing the surface temperature response. Looking again at figure 1.4 it can be seen that radiative feedbacks include effects of surface temperature response on a number of Earth system components (e.g. clouds), many of which also respond to radiative adjustments.

Therefore it is important to be clear whether the response in the component arises from the effects of a forcing agent or the resulting surface temperature response.

This forcing-adjustment-feedback framework does however have its limitations. While using ERF does increase parity in the expected surface temperature response to forcing agents it does not do so perfectly. This is particularly true for forcing agents with spatially heterogeneous radiative forcing because when the same global mean radiative forcing is applied but with different spatial patterns, the spatial pattern of surface temperature response differs (Forster et al., 2000). Radiative feedbacks depend on the pattern of surface temperature response, so the same magnitude radiative forcing with different spatial pattern can result in a differing radiative feedback (Forster et al., 2000). Aerosols in particular have short atmospheric lifetimes and anthropogenic emissions of aerosols are not globally uniform, resulting in large spatial heterogeneity in their forcing (Li et al., 2022). Differences in the vertical distribution of how forcing agents affect radiative fluxes and atmospheric temperature can also result in deviations in efficacy. For example, contrail cirrus has been found to have a forcing efficacy much lower than one even when using ERF. This means that the surface temperature response to contrails estimated from their ERF is likely be substantially overestimated compared to directly modelling the temperature response (Bickel, 2023).

Another challenge posed by the use of ERF is the clarity and consistency with which radiative adjustments are described and calculated. This applies in particular to cloud responses to aerosol perturbation, and so it is worth considering the terminology used to describe cloud responses to aerosol perturbation. Aerosol effects on clouds via their role as CCN are often described as 'indirect effects' or microphysical effects (Boucher et al., 2013). These indirect effects include particularly the effect of aerosols increasing the cloud droplet number concentration (CDNC) and consequently increasing cloud reflectance (the 'Twomey effect' or first indirect effect) (Twomey, 1977); and the increase in CDNC causing a suppression of precipitation and thus increased cloud lifetime (the 'Albrecht effect' or second indirect effect) (Albrecht, 1989). Meanwhile cloud responses arising via atmospheric responses to the aerosol direct radiative effect have been described as 'semi-direct' effects. Semi-direct effects primarily include local heating or cooling of the atmosphere by aerosols causing changes in cloud liquid water paths, cloud fraction or other cloud properties, and consequently affecting scattering and absorption of radiation by clouds (Hansen et al., 1997; Sherwood et al., 2015).

Indirect cloud radiative adjustments contribute a significant portion of the ERF for total aerosol, and the cloud brightening effect in particular is the largest of these (Zelinka et al., 2014; Smith et al., 2020a). Boucher et al. (2013) set out new terminology in which the ERF due to aerosols is split into indirect cloud radiative forcing via aerosol cloud interactions (ERF$_{aci}$) and radiative forcing arising from aerosol radiation interactions (ERF$_{ari}$):

$$ERF = ERF_{\mathrm{aci}} + ERF_{\mathrm{ari}} \tag{1.2}$$

Using the form of Zelinka et al. (2023), the ERF$_{aci}$ can be further subdivided into the IRF$_{aci}$, which is the cloud brightening effect, and aci adjustments which include all other indirect cloud adjustments:

$$ERF_{\text{aci}} = IRF_{\text{aci}} + K^C \Delta C_{\text{indirect}} \tag{1.3}$$

where K^X is a factor describing the change in net downwards TOA radiative flux with infinitesimal changes in any variable, X:

$$K^X = \frac{\partial R}{\partial X} \tag{1.4}$$

where R is the net downwards TOA flux. Likewise the ERF$_{ari}$ can be split into the IRF$_{ari}$, which is the aerosol IRF, and ari adjustments, which include both semi-direct effects as well as all other non-cloud adjustments:

$$ERF_{\text{ari}} = IRF_{\text{ari}} + K^C \Delta C_{\text{semi-direct}} + K^T \Delta T + K^q \Delta q + K^\alpha \Delta \alpha \tag{1.5}$$

where T, q, and α are atmospheric temperature, specific humidity and surface albedo respectively. This terminology helpfully separates the large aci forcing (i.e. ERFaci) into a separate term from other adjustments and also defines the large cloud brightening effect (i.e. IRFaci) as part of the IRF, separating it from other aci adjustments.

In practice however, some techniques used for investigating aerosol ERFs are unable to distinguish between each of the ERF terms (Smith et al., 2020a). This has led to the sometimes confusing use of terminology in the literature. For example, all cloud responses are sometimes grouped into ERFaci despite having contributions from aci and ari adjustments (Zelinka et al., 2014; Gryspeerdt et al., 2020; Smith et al., 2020a). Consequently it is important that where cloud radiative effects are calculated, what quantity is being calculated or presented should be stated clearly. Since cloud radiative effects are not always separable between ari and aci responses (as is the case in the methods used in this work), it should be noted that in this work we refer to all cloud radiative responses as adjustments, including the IRFaci.

1.3.2 Methods for determining ERF

While using ERF enables greater parity when comparing forcing agents, it is more complicated to calculate than the IRF, since it requires accounting for the effects of many interactions between forcing agents and the Earth system. Two primary methods of determining ERF in GCMs and ESMs have emerged. The most commonly used method involves prescribing sea surface temperatures (SSTs) and sea ice (SI) in the model, and is often referred to as the Fixed SST or fSST method (Hansen et al., 2005; Forster et al., 2016). In this method, two simulations are run. The ‘control’ simulation has all boundary conditions set to a single state (e.g. pre-industrial). The ‘perturbed’ simulation is the same but with the forcing agent(s) of interest perturbed to a different state (e.g. using present day aerosol emissions). The simulations are both run for a number of model years using annually repeating boundary conditions (particularly SST and SI) to allow averaging over a number of model years representing the same climate state, allowing for a spin up time. While the model boundary conditions do not evolve over time, model variables are allowed to adjust, such that the perturbed simulation represents a climate which includes radiative adjustments to the perturbation of the forcing agent. The difference in net down TOA fluxes between the two simulations thus gives the ERF attributable to the forcing agent. Importantly prescribing the SSTs and SI prevents these changing in response to the perturbation, excluding most radiative feedbacks from the model response.

The other method for determining ERF requires a coupled model in which SST is allowed to change in a single simulation (Gregory et al., 2004). In this method the simulation is initialised from a control atmospheric state but has the forcing agent(s) of interest abruptly perturbed. The simulation is run for a number of model years with both adjustments occurring and surface temperature responding to the forcing, which results in radiative feedbacks which alter the TOA radiative balance. As illustrated in figure 1.5, fitting a linear regression of the mean TOA radiative imbalance against mean surface temperature of each model year acts as a solution to equation 1.1. Specifically, the regression gradient is the feedback parameter and the y-intercept (i.e. the radiative forcing when the surface temperature response is zero) is the ERF.

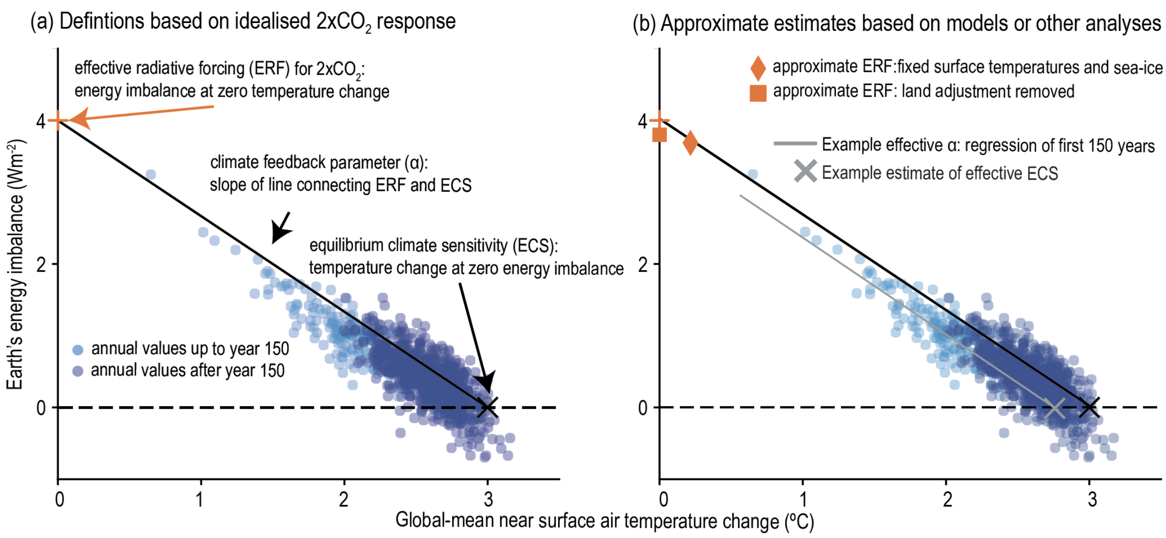

Figure 1.5: Illustration of fSST and regression based methods for calculating ERF and how they relate to each other, in (a) idealised calculations and in (b) approximate calculations which include the limitations of the two methods. From Forster et al. (2021).

Both methods have their limitations. One of the key limitations for the fSST method is that usually the surface temperature over land is not prescribed in models. Therefore most studies using this methodology report an ERF that includes a feedback term in response to land surface temperature change (Sherwood et al., 2015; Forster et al., 2016). This includes both the planck effect and responses in other variables to the surface temperature change. This can be seen in the approximate case in figure 1.5b. While this issue has long been acknowledged as a consequence of the method (Hansen et al., 2005), some studies have chosen to correct for this effect by subtracting the Planck response due to land surface temperature changes (Forster et al., 2021; Zelinka et al., 2023). Some studies have gone further and subtracted even estimates of the radiative feedbacks that would result from the land surface temperature response as well (Smith et al., 2020a). A few studies have addressed the issue directly by using the fSST method with land surface temperatures and humidity also fixed (Shine et al., 2003; Andrews et al., 2021). However this is difficult to implement in many ESMs with complex land models and requires prescribing not only land surface temperatures but soil temperature too (Ackerley and Dommenget, 2016; Forster et al., 2016). The linear regression method avoids this issue since the surface temperature does not need to be fixed at all.

One limitation with the linear regression method however is that it does not totally distinguish between radiative adjustments and feedbacks. This is because adjustments take time to occur in response to a forcing agent perturbation, allowing radiative feedbacks to begin to take place too. While radiative adjustments tend to take place more quickly than radiative feedbacks, there is overlap in their timescales. For example stratospheric temperature can take months to adjust (Stjern et al., 2023). Therefore, by the time the adjustments have occurred in the model the TOA forcing will already have been altered by feedbacks to the surface temperature response. Furthermore, this method assumes a linear fit between the surface temperature response and the TOA flux imbalance. This effectively assumes the feedback parameter remains constant. However, studies have demonstrated that the feedback parameter can change with climate state, with studies from a range of ESMs suggesting the feedback parameter increases with increasing surface temperature from pre-industrial to present day (Forster et al., 2021). This issue is illustrated in figure 1.5 where the ERF calculated will depend on the gradient of the regression, which in turn depends on the number of years used for the regression.

While both methods have limitations in terms of what exactly they calculate, the fSST method can achieve the same uncertainty level with much fewer model years (Forster et al., 2016; Pincus et al., 2016). Consequently most studies of ERF, including the Radiative Forcing and Aerosol Chemistry Model Intercomparison Projects (Pincus et al., 2016; Collins et al., 2017), tend to use the fSST method. All simulations in the present work use the fSST method. It should be noted that corrections for land surface temperature responses are not applied to ERFs reported in this work, and even though feedbacks may affect responses of other variables (notably surface albedo), these are not corrected for. However, the planck response associated with land surface temperature response is considered and evaluated.

1.3.3 Methods for isolating radiative adjustments

To understand and isolate the components of ERF, particularly radiative adjustments and the IRF, a number of methodologies have been employed in the literature using earth system models (ESM) and global climate models (GCM). These include, for example: using additional diagnostic radiation calls in models with and without aerosol and cloud included in radiation calculations (Ghan, 2013; Gryspeerdt et al., 2020); various partial radiative perturbation (PRP) methods (Smith et al., 2018, 2020a; Zelinka et al., 2023); processing model output with radiative kernels (Smith et al., 2018, 2020a,b); and mechanism denial experiments in which certain adjustment mechanisms are disabled (Johnson et al., 2019), variables are given prescribed values (Andrews et al., 2021; Huang et al., 2020), or changes in variables are suppressed (Johnson et al., 2019) to suppress and isolate adjustments. Here these methods are described and their merits and limitations considered in turn.

It should be noted that these methods can be used for quantifying radiative feedbacks as well as radiative adjustments. In most cases the methods described below do not specifically calculate radiative adjustments or feedbacks, but rather calculate the radiative effect of changes in specific variables in a pair of control-perturbed simulations. If these simulations fix surface temperature, only radiative adjustments are included; if they allow surface temperature to change, radiative feedbacks will be included also. In this work the focus in on radiative adjustments, so description of these methods is usually in the context of a pair of fSST simulations, such that they only return radiative adjustments.

Additional radiation calls

Additional radiation calls can be added to models to diagnose the radiative effects of certain components of a model in isolation. This typically works by having diagnostic radiation calls that calculate the radiative effect of changing certain variables, such as clouds and aerosols, without these changes affecting the model evolution (Ghan, 2013). This allows determination of the radiative effect of some change on the background of an unperturbed atmosphere i.e. without the atmosphere adjusting in response to the change.

A frequent application of additional radiation calls is for determining the direct radiative effects of clouds or aerosols. This works by having diagnostic radiation calls that exclude the presence of aerosol, clouds or both (Ghan, 2013). These quantities are often referred to as 'clean' (no aerosol), 'clear' (no cloud) or 'clear-clean' (no cloud or aerosol) sky fluxes. For example, the cloud radiative effect (CRE) can be determined as:

$$CRE = F - F_{\mathrm{clr}} \tag{1.6}$$

where F is the net downwards TOA radiative flux and F_{clr} is the clear sky net downwards TOA radiative flux (i.e. neglecting clouds). Similarly, the direct radiative effect (DRE) of aerosol can be determined as

$$DRE = F - F_{\mathrm{cln}} \tag{1.7}$$

where F_{cln} is the clean sky net downwards TOA radiative flux (i.e. neglecting aerosols).

These quantities can be differenced between control-perturbed simulation pairs to isolate the effects of the aerosol, cloud or other components of the model on radiation. For example, the difference in the all sky minus clear sky forcing between a control-perturbed fSST simulation pair gives an estimate of the cloud radiative adjustment:

$$\Delta CRE^{*} = \Delta(F - F_{\mathrm{clr}}) \tag{1.8}$$

where ΔCRE* is the change in cloud radiative effect between the control and perturbed simulations and is the radiative adjustment resulting from all cloud adjustments due to the forcing agent perturbation. However, Ghan (2013) suggested that the estimate of ΔCRE in equaiton 1.8 is biased, since it implies that ΔDRE (neglecting non-cloud or aerosol adjustments) can be determined as

$$\Delta DRE^{*} = \Delta F_{\mathrm{clr}} \tag{1.9}$$

They argue this would be biased in the SW since an increase in the scattering effects of aerosol would appear to be larger while an increase in absorption effects of aerosol appear smaller, if the scattering effects of clouds below the aerosol are neglected. Effectively the bias arises due to excluding the masking effect of cloud on aerosol perturbations as outlined by Zelinka et al. (2023). Ghan (2013) therefore suggested calculating ΔDRE as:

$$\Delta DRE = \Delta(F - F_{\mathrm{cln}}) \tag{1.10}$$

and calculating ΔCRE as:

$$\Delta CRE = \Delta(F_{\mathrm{cln}} - F_{\mathrm{clr,cln}}) \tag{1.11}$$

where $F_{\mathrm{clr,cln}}$ is the clear-clean sky net downwards TOA radiative flux. They also suggested a calculation for surface albedo effects as:

$$\Delta F_{\alpha} = \Delta F_{\mathrm{clr,cln}} \tag{1.12}$$

where F_{α} is the radiative forcing due to surface albedo adjustment. As noted by Zelinka et al. (2023) this term actually includes all other radiative adjustments aside from aerosols and clouds as well, particularly those from specific humidity and temperature adjustments.

Ghan (2013) demonstrated that the calculation of ΔCRE using equation 1.8 is positively biased relative to using 1.11. However, it is worth noting that in equation 1.11 one may consider that, by calculating cloud forcing in the absence of aerosol, cloud forcing may be biased. For example, neglecting either scattering or absorbing aerosol above cloud when calculating the radiative effect of a low cloud increase will exaggerate the negative SW forcing due to the low cloud increase since more SW radiation reaches the cloud to be affected by its scattering. Zelinka et al. (2023) outlined this effect algebraically, describing it as a masking effect of aerosol that describes how much the radiative impacts of changes in other variables (clouds, temperature, specific humidity and surface albedo particularly) are attenuated by aerosols. Given that a masking effect of cloud or aerosol applies whether ΔCRE is calculated under clean sky conditions or ΔDRE is calculated under clear sky conditions, there is arguably little distinction between equation 1.11 or calculating ΔCRE as:

$$\Delta CRE^{'} = \Delta(F - F_{\mathrm{clr}}) \tag{1.13}$$

and between equation 1.10 and calculating ΔDRE as:

$$\Delta DRE^{'} = \Delta(F_{\mathrm{clr}} - F_{\mathrm{clr,cln}}) \tag{1.14}$$

However, since the decomposition suggested by Ghan (2013) is already established, equations 1.11 and 1.10 are used for ΔCRE and ΔDRE unless otherwise stated going forward.

In the context of aerosol emission perturbation fSST experiments, these equations allow calculation of the radiative forcing due to changes in aerosol, clouds and all remaining factors separately. They are therefore a highly useful tool to separate out the major radiative adjustment terms. Consequently such diagnostics are now included in many GCMs and ESMs and used in many studies of radiative adjustments (Smith et al., 2018, 2020a; Thornhill et al., 2021).

Using additional radiation calls to separate radiative adjustments does have limitations. Firstly, these equations do not separate the ERF by adjustment mechanisms; rather they separate ERF into the radiative adjustment of three groups (aerosol, cloud, other) of variables. For clouds this means that ERFaci (microphysical) cloud adjustments and ERFari (semi-direct) cloud adjustments are grouped together into ΔCRE. For aerosol, the ΔDRE will include forcing due to all changes in aerosols. Therefore where only a subset of aerosol species are perturbed ΔDRE will still include any forcing due to adjustment of other aerosol species, and so does not represent the IRF of the subset of species perturbed. Likewise radiative forcing from all other adjustment components aside from clouds and aerosols (which will include adjustments in surface and atmospheric temperature, humidity, surface albedo, lapse rate, greenhouse gas via chemical adjustments, and biosphere responses) are not distinguished.

Secondly, as discussed earlier, and outlined in detail in Zelinka et al. (2023), ΔDRE and ΔCRE both include masking terms. Therefore these estimates of ΔDRE and ΔCRE are biased relative to their definitions in Forster et al. (2021). There are ways to quantify this error however. For example, (Shell et al., 2008) quantified the masking effects of temperature, surface albedo and water vapour included on $\Delta\mathrm{CRE}^*$ when investigating radiative feedbacks under doubling of CO_2, by differencing the radiative feedback calculated for each variable with radiative kernels under all sky minus under clear sky conditions.

Thirdly, even for perturbations of all aerosol and ignoring the masking effect outlined in Zelinka et al. (2023), ΔDRE is not exactly equivalent to the definition of IRF for aerosol perturbation outlined in Forster et al. (2021). This is because, while the radiative effects of changing other model components are excluded in this method, the ‘background’ atmosphere on which the radiative effect of the aerosol perturbation is calculated changes. This could be described as a change in how other variables attenuate the radiative effect of an aerosol perturbation due to those variables adjusting, effectively altering the masking effect described in Zelinka et al. (2023), and is explained in further detail in appendix A. This effect is not necessarily a flaw however, because in the real atmosphere adjustments happen simultaneously, so changing the background atmosphere between control and perturbed states, and consequently altering the masking term, does not bias ΔDRE in the way that the presence of a masking term does. It is also likely to have a smaller effect than the presence of a masking term itself. However, this effect does defy the strict definition of IRF and so is worth noting.

Partial Radiative Perturbation

PRP is a method of determining the radiative forcing due to changes in individual variables (Wetherald and Manabe, 1988). The method works very similarly to additional radiation calls, the main difference being that it is usually performed offline using output from GCM or ESM simulations rather than applied online during model timesteps. As can be done with additional radiation calls, PRP involves using a radiation code with a variable of interest set to values from a perturbed simulation while all others are set to values from a control simulation. This can be repeated for many individual variables, each time differencing the calculated TOA fluxes with TOA fluxes calculated when all variables are set to a control state, to determine each variable's radiative forcing. PRP has been used in many studies investigating radiative feedbacks (Colman et al., 1997, 2001; Colman, 2015; Bickel et al., 2020) and more recently for radiative adjustments (Smith et al., 2018; Mülmenstädt et al., 2019; Smith et al., 2020a). It can even be used to separate out radiative adjustments due to changes in cloud fraction, cloud liquid water content and droplet effective radius (Mülmenstädt et al., 2019).

Using an offline radiation code for PRP offers greater flexibility than adding radiation calls to an online model. Firstly, PRP can used with input from sources other than a GCM/ESM simulation, such as using reanalysis output (Pincus et al., 2020). Secondly, offline radiation codes can be more easily run with just a subsample of the grid points or timesteps from GCM/ESM simulations. However, additional online radiation calls may be more efficient than offline PRP calculations, particularly for frequently studied variables such as clouds and aerosols, for which clear and clean sky diagnostics are readily available in many models.

One challenge with using PRP is that it implicitly assumes decorrelation of variables in any state (e.g. control or perturbed), which is not a reasonable assumption (Colman and McAvaney, 1997; Soden et al., 2008). For example, specific humidity is expected to be correlated with cloud fraction since they follow the same spatial patterns. However the spatial pattern of the specific humidity field taken from one simulation (e.g. a control simulation) will not be as closely spatially correlated with the cloud fraction field from another simulation (e.g. a perturbed simulation). This is primarily due to internal variability affecting the spatial patterns of humidity and clouds in the different simulations, which are unlikely to match between the two simulations at any given time (though systematic changes in cloud processes due to the perturbation could also affect correlation). Consider the effect of this correlation when calculating ΔCRE with PRP, which uses the difference in TOA radiative flux calculated with cloud fields from a perturbed simulation but all other variable fields from a control simulation, minus TOA radiative flux calculated using all variable fields from a control simulation. In the first term the perturbed cloud fields and control specific humidity field will not be strongly co-located on any given timestep, and so both variables will have a stronger radiative effect than expected due to not overlapping. Thus ΔCRE will be positively biased. This issue could affect perturbations of other variables too and and the magnitude of the bias will depend on how strongly variables are correlated with others.

To counter this decorrelation assumption Colman and McAvaney (1997) suggested using a 'two-sided' calculation. This involves conducting both a 'forward' and 'backward' calculation. The forward calculation is the difference in TOA radiative fluxes calculated with one variable set to a perturbed state and all others set to a control state minus with all variables set to the control state. The backward calculation is the difference in TOA radiative flxues with all variables set to the perturbed state minus with the variable of interest set to a control state and all others set to the perturbed state. The radiative forcing associated with the variable of interest is then calculated as a mean of the forward and backward calculations. The forward and backward calculations therefore have a bias term of opposite sign. As demonstrated algebraically by Soden et al. (2008), when summing the forward and backward calculations the bias terms mostly cancel (leaving a smaller error term), and so the mean of the two calculations will be unbiased.

It is also worth considering that as well as (mostly) removing the decorrelation bias, using the two-sided calculation averages over calculation of the perturbation on a background of control or perturbed fields of the other variables. Strictly speaking this means that, if using output from a fSST experiment with a perturbation of total aerosol emissions and a control experiment, calculating the aerosol perturbation with two-sided PRP would not return the IRF as defined in Forster et al. (2021), but rather a very similar quantity. This is because the backward calculation considers the aerosol perturbation on the perturbed atmospheric state, i.e., with adjustments having occurred. This can be expected to slightly change the forcing of the aerosol perturbation compared to the forward case, even though the forcing arising directly from adjustments is cancelled in both forward and backward cases. By contrast using just the forward case would return the IRF plus a decorrelation bias. Since other methods such as using clear and clean sky diagnostics also calculate a pseudo IRF as described in section 1.3.3 the two-sided calculation is likely to be more useful since it excludes the decorrelation bias which can only be construed as an error.

Another challenge with the PRP method is that it generally requires sub-daily input to produce accurate results (Wetherald and Manabe, 1988; Colman et al., 2001; Bellouin et al., 2020b). For example, the radiative effect of clouds depends strongly on how they overlap on different model levels and on the state of other variables such as surface albedo. The variation of clouds with other variables and the variation of cloud fraction in different atmospheric layers will both have some degree of temporal correlation, over both sub-daily timescales and longer inter-day timescales. These relationships and correlations are lost in monthly means of clouds fields, in which most grid points across the globe would have some fraction of cloud due to averaging over cloudy and non-cloudy timesteps. Consequently, using monthly mean cloud fields will produce a different cloud radiative adjustment than if sub-daily instantaneous input were used instead, particularly in the SW (Bellouin et al., 2020b). While using input at the frequency of radiation timesteps will give greater accuracy, studies have suggested using less frequent input to sample the diurnal variation while reducing the computational burden of calculations relative to timestep frequency input (Colman et al., 2001). For example, Mülmenstädt et al. (2019) found similar results using PRP with 3 hourly averaged or timestep input, and Bellouin et al. (2020b) found similar results using PRP with 21 hourly or 3 hourly instantaneous input.

Approximate Partial Radiative Perturbation

The approximate partial radiative perturbation (APRP) method is a variation on the PRP method that, using a simplified model of radiation in the atmosphere, can be used to separate SW ERF into radiative effects of cloud, surface albedo and 'non-cloud factors' (Taylor et al., 2007). The non-cloud factors term includes any change in absorption or scattering in the atmosphere, which will be primarily due to adjustments/perturbations of aerosols, greenhouse gases and humidity. Moreover the APRP method allows separation of the cloud term into absorption, scattering and cloud amount terms and the non-cloud factors term into absorption and scattering terms (Zelinka et al., 2023). For aerosol perturbation experiments this is particularly useful, as it enables separation of radiative adjustments due to cloud brightness and cloud amount (Zelinka et al., 2023). The implementation of APRP is described in detail in Taylor et al. (2007) and Zelinka et al. (2023).

Taylor et al. (2007) tested the application of the APRP method for isolating radiative feedbacks from a doubling of (PD) CO_2, and compared it with calculations made using the PRP method. Using monthly mean climatology input to the APRP method, they found that estimates of the surface albedo, cloud and non-cloud radiative feedbacks all agreed with 10% with calculations from the PRP method. Spatial correlation between the the APRP and PRP calculations for surface albedo and cloud feedbacks were 0.997, but worse for the non-cloud term (0.794). Using annual mean rather than monthly mean climatology input, the APRP method produced a surface albedo estimate with a 13% error relative to the PRP estimate, and with close to double the root mean square (RMS) error. Zelinka et al. (2023) applied the APRP method to isolate radiative adjustments multi-model fSST experiments in which anthropogenic aerosols are perturbed from PI to PD conditions. They compared the results from the APRP method to using the double call method outlined in 1.3.3, finding that the two methods agreed closely on the total SW ERF, but the APRP method produced more negative SW ERF_{ari} and more positive SW ERF_{aci}. However they suggest this difference may be expected due to the masking terms included in the double-call method but not the APRP method. They also found that the global mean RMS difference between the two methods for the SW ERF was less than 10% of the standard deviation in SW ERF calculated from either method.

The main advantage of the APRP method over the PRP method is the ability to use coarser time resolutions of input to achieve similar results (Taylor et al., 2007). That said while APRP is expected to be less affected by using monthly mean input data compared with sub-daily input, it could also be used with finer scale input data. The main limitations of the APRP method are that it is difficult to apply to LW radiative adjustments, it is less flexible in what adjustments can be separated than the PRP method, and is less accurate than the PRP method.

Radiative kernels

Like the PRP method, radiative kernels can be used to determine radiative forcing due to perturbations of individual radiatively active variables (Held and Soden, 2000; Soden et al., 2008). Like the PRP method, this works by calculating the radiative response to changes in each variable individually. However, rather than calculating the effect of a perturbation in a given variable between a specific control and perturbed simulation directly with a radiation code, instead the differential between the radiative effect and the perturbation of a given variable is computed first (Held and Soden, 2000; Soden et al., 2008). This 'kernel' can thus be thought of as a 'weighting function' (Wu et al., 1993), specifically for the radiative response arising from a perturbation in the given variable. For the given variable, any perturbation in that variable's field can then be multiplied by the kernel to determine the radiative effect of that perturbation; hence, a kernel can be used multiple times for different experiments.

Typically radiative kernels are computed by calculating the radiative effect of applying an arbitrary increment to each variable in a radiation code. For example, to compute a temperature kernel, Smith et al. (2020b) took various output fields from simulations using the Hadley Centre General Environment Model 3 (HadGEM3-GA7.1) and used them in the model's offline radiative transfer code to calculate TOA and surface radiative fluxes. They then re-ran the offline radiative transfer code once for each model level, with temperature on each respective model level perturbed by +1 K, and repeated this for every 2 hours of one model year. For each level, taking the difference in TOA and surface radiative fluxes when using the control temperature field and the same field but with that level perturbed, gives the radiative kernel for temperature for that level. Repeating this over all levels gives the total radiative kernel for temperature. Thus, the level averaged difference in temperature fields between a control and perturbed simulation can be multiplied by this kernel to determine the total radiative effect of temperature change between the two simulations.

The ability to use a kernel for multiple experiments offers a major advantage over the PRP method. Running calculations in a radiation code is computationally expensive, and while the PRP method requires this to be done for every experiment, a kernel requires this only be done once (Soden et al., 2008). As such, radiative kernels tend to be used more frequently in recent studies of radiative adjustments than the PRP method (Smith et al., 2018, 2020a; Bickel, 2023).

Radiative kernels do suffer from some disadvantages compared to the PRP method however. Firstly, radiative kernels rely on the assumption that radiative responses are linear with the magnitude of perturbation of the radiative variable (Soden et al., 2008). How reasonable this approximation is depends on the size of the perturbation and the variable. For example, Martin and Quaas (2021) noted that surface albedo changes include strong non-linear processes, particularly including albedo over snow and ice, causing the kernel method to underestimate surface albedo radiative feedback.

Another limitation of kernels is that they are determined from a specific radiation code and for a specific baseline atmosphere (Soden et al., 2008). Consequently, if there are significant differences between model radiation codes, calculating radiative adjustments using outputs from one GCM, using a kernel generated from a different radiation code to that used by the GCM, may not by representative of or consistent with the radiative effects of the adjustments in the GCM. Likewise, if a kernel calculated with a given baseline climate (e.g. representing pre-industrial climate) is used to calculate radiative adjustments due to perturbations in a simulation representing a different baseline (e.g. representing present-day climate) the results may be less consistent with the radiative adjustments in the simulation.

A third challenge for radiative kernels is representation of cloud responses. Cloud radiative effects are particularly non-linear with changes in their properties (Shell et al., 2008), so few studies have sought to make radiative kernels for cloud responses. One example of a cloud radiative kernel was that developed by Zelinka et al. (2012) using pseudo-satellite observations developed for the International Satellite Cloud Climatology Project (ISCCP) (Klein and Jakob, 1999). The ISCCP kernel was calculated by determining the effect of increasing cloud fraction from 0 to 100% in different cloud-top pressure and optical depth bins, over different months, solar zenith angles, surface albedos and latitudes, and then dividing this by 100 to get the radiative feedback per percent of cloud fraction change in each bin. However, one of the primary limitations of the ISCCP kernel for investigating aerosol radiative adjustments is the inability to distinguish aci and ari cloud adjustments (Smith et al., 2020a). The ISCCP kernel can also only represent cloud changes in lit grid points, omitting the effects of cloud changes in unlit points on LW fluxes (Zelinka et al., 2012).

An alternative approach for quantifying cloud radiative effects using radiative kernels is to quantify all of the non-cloud adjustments using kernels and quantify the IRF (either using a kernel or other means, such as using ΔDRE for aerosol perturbation) and then assume the remainder of the ERF is due to cloud radiative adjustments (Taylor et al., 2007). The obvious limitation to this approach is again that it does not separate individual cloud radiative adjustment. This method also assumes one is interested in calculating all radiative adjustments and IRF, whereas the ISCCP kernel could be used to isolate the cloud radiative adjustment without needing to quantifying other adjustments also.

Finally, it is worth considering whether using two-sided calculations has a value in the context of radiative kernels. Bickel et al. (2020) noted that most kernels are calculated with only a 'forward' calculation (adding an increment to each variable), but since there are large differences in forward and backward calculations using the PRP method, this could result in differences between calculations using a radiative kernel versus using a two-sided PRP approach. As described earlier, the effects of using two-sided PRP are, firstly, the removal of the decorrelation bias due to intrinsic variability and resulting differences between control and perturbed simulations; and secondly, determining a mean of performing the perturbation with all other variables set to control or perturbed values. We can therefore consider whether two-sided kernel calculations would have value in the context of each of these effects. With regards to the first effect, Soden et al. (2008) demonstrated algebraically that using radiative kernels does not result in the spatial decorrelation bias due to intrinsic variability. Simply put, this is because rather than substituting a field from one simulation for another (with each simulation having differences due to natural variability at any given timestep), the kernel is calculated using arbitrary changes to each variables' field. However, while in principle the kernel method assumes that the kernel is calculated using infinitesimal perturbations, in practice they are calculated using finite increments to the variable of interest, usually at one model level at a time. While likely a much smaller effect than the spatial decorrelation bias due to intrinsic variability in a one-sided PRP calculation, this does result in a change to the spatial pattern of the variable without adjusting the spatial patterns of other variables, and so could be considered to cause a bias. With regards to the second effect, it could be argued that taking the average between calculating the radiative effect of perturbing a variable with control and perturbed values of other variables is more representative of actual radiative responses than just calculating the effect of a perturbation on the control state. For example, if performing a control and perturbed fSST simulation, one could calculate a forward kernel with the control simulation as the baseline and a backward kernel with the perturbed simulation as the baseline, apply both kernels to adjustments, and average the resulting two estimates of radiative adjustments.

Few studies have investigated the difference between using forward and backward radiative kernels. Martin and Quaas (2021) compared the application of using either forward and backward PRP or forward and backward radiative kernels for determining radiative feedback parameters under a quadrupling of CO_2. They computed the forward kernel using the control simulation as the baseline and the backward kernel using the perturbed simulation as the baseline. They found that the difference between forward and backward calculations was similar when using either kernels or PRP for determining the planck and albedo feedback parameters, but was considerably larger when using kernels for the lapse rate and water vapour feedback parameters. This would suggest that the seemingly arbitrary choice of many studies to use a forward kernel has a potentially significant effect on radiative feedbacks (or adjustments) calculated. While the justification for doing two-sided PRP calculations is clearer than for kernels because the decorrelation bias due to intrinsic variability affecting PRP calculations is a clearly detrimental bias while the difference in forward and backward kernel calculations could be primarily due to using either control or perturbed values for the other variables, the results of Martin and Quaas (2021) suggest further investigation of two-sided kernel calculations is warranted.

Mechanism denial

Mechanism denial methods effectively work by excluding processes from actually occurring in the model rather than calculating the effect of a change in a specific variable. Thus these methods tend to calculate the forcing associated with a mechanism of adjustment (e.g. circulation changes) that may affect one or a number of variables that interact with radiation, rather than the forcing due to a single variable adjusting (e.g. clouds) that may be affected by a number of mechanisms.

For example, Johnson et al. (2019) separated the effect of aci adjustments arising due to BC emissions perturbation by running two control-perturbed pairs of experiments, one prescribing CDNC, the other not. This effectively prevents the change in aerosol affecting CDNC, and hence, prevents aci adjustments. The difference in total radiative forcing between the two experiment pairs should therefore equal the radiative forcing due to aci adjustments. To be precise the difference will also include some forcing (likely of a relatively small magnitude) due to the effect of aci adjustments on the simulation's evolution, since the mechanism denial removes not just the aci forcing to the ERF but also its contribution to the model evolution. This is different to the aforementioned techniques which only include the direct radiative forcing associated with the variable(s) of interest not their effects on the model evolution.

Nudging is another example of a mechanism denial tool. In many radiative forcing studies nudging has been applied to reduce simulation lengths, by reducing noise due to internal variability differences between control and perturbed simulations (e.g. with different aerosol emissions) (Penner et al., 2018; Che et al., 2021; Mülmenstädt et al., 2019; Grosvenor and Carslaw, 2020; Zhu and Penner, 2020; Ghan et al., 2016). This is achieved by nudging the meteorology in both simulations towards the same reference state, thus reducing the simulation length that must be integrated over to determine robust signals (Kooperman et al., 2012; Forster et al., 2016). However, applying nudging can affect the ERF determined from these studies. For example, Kooperman et al. (2012) compared the aci radiative forcing resulting from pre-industrial (PI) to present-day (PD) aerosol perturbation, between a simulation pair without nudging ('free') and a simulation pair with nudged horizontal winds and dry static energy, using the Community Atmosphere Model 5 (CAM5). They noted a statistically significant positive bias in the magnitude of aci top of atmosphere (TOA) forcing in the nudged simulation pair, which they attributed to reductions in cloud fraction. While such a bias may prove problematic for studies using nudging to reduce simulation lengths, Kooperman et al. (2012) noted that this effect could be exploited as a mechanism denial tool to deliberately suppress model responses to understand their effects.

This principle was used by Johnson et al. (2019), who investigated the use of nudging horizontal winds to separate out the circulation adjustment arising from a PI to PD BC perturbation. Using the Hadley Centre Global Environment Model version 3 (HADGEM3), they found that nudging horizontal winds to reanalysis data did not affect the BC ERF compared to that determined in a free simulation pair. However, nudging did result in changes to ice water path and high cloud fraction, suppressing the reduction in both quantities compared to the free simulation pair. It is this application of nudging as a mechanism denial tool that is used in this thesis. A more detailed review of nudging application in the literature is given in the following section 1.3.4.

1.3.4 Nudging and considerations for its use

Nudging theory and applications

Nudging, or Newtonian relaxation, is a simple data assimilation tool that can be used to relax model variables towards a reference state (Jeuken et al., 1996). It is typically used to constrain the large scale model dynamics or 'weather' in a model to some reference state (Jeuken et al., 1996; Telford et al., 2008) without prescribing model variables directly (which could generate instabilities). As such, most studies using nudging nudge variables describing the large-scale dynamics, such as the horizontal winds or potential vorticity, and often temperature or related quantities such as dry static energy or potential temperature are nudged too (Jeuken et al., 1996; Li et al., 1998; Telford et al., 2008; Kooperman et al., 2012). Nudged simulations are sometimes also referred to as 'specified dynamics' simulations (Davis et al., 2020, 2022).

The basic formulation of the nudging equation may be written as:

$$\frac{\delta X}{\delta t} = F(X) + G(X_{\text{ref}} - X_{\text{mod}}) \tag{1.15}$$

where X is the variable to be nudged, $\frac{\delta X}{\delta t}$ is the rate of change of X with time, F(X) is the rate of change of X due to all other factors in the model, G is the relaxation parameter, X_{ref} is the reference value of X, and X_{mod} is the model value of X (Jeuken et al., 1996). $(X_{\text{ref}} - X_{\text{mod}})$ can be thought of as the difference between the reference and model value of X at a given time, and G can be thought of as the reciprocal of the e-folding timescale, the time for the difference to decrease by a factor e (van Garderen et al., 2021). Thus G may also be usefully thought of as the nudging 'strength', since the larger the value of G the more $(X_{\text{ref}} - X_{\text{mod}})$ is reduced by the nudging over time.

In a model with discrete timesteps the nudging equation can be implemented as:

$$\Delta X = F(X) + G\Delta t(X_{\text{ref}} - X_{\text{mod}}) \tag{1.16}$$

where ΔX is the change in X over the timestep Δt (Telford et al., 2008). Nudging is typically applied at every model timestep, using input data interpolated (usually linearly) to the time resolution of the timesteps (Jeuken et al., 1996), though some studies have experimented with nudging on only some timesteps finding limited success (Sun et al., 2019). Consequently the nudging term can be thought of as adding a fraction $G\Delta t$ of the difference between the (interpolated) model and reference values of X to the value of X at each timestep. Again G can be understood as the nudging strength, where the larger the value of G, the larger the fraction of the difference that is added each timestep, effectively increasing the degree to which X in the model is nudged towards the reference value.

The specific applications of nudging in the literature have been varied. Many studies apply nudging for the purpose of model validation, constraining meteorological conditions in a model to observations. This is useful when the process of interest is affected strongly by dynamical conditions and so it is important to ensure the model accurately simulates the meteorological conditions (Jeuken et al., 1996). Some studies do this for specific periods such as during field campaigns or particular meteorological events for investigating, for example, clouds (Dean et al., 2007; Gettelman et al., 2020) and trace gases (van Aalst et al., 2004). Other studies apply nudging over longer periods, to determine how well models represent different processes, for example, stratospheric chemistry (Davis et al., 2020).

As described in the previous section, for investigating radiative forcing, adjustments and feedbacks, nudging tends to be used instead as a way of constraining natural variability, or more rarely for mechanism denial experiments (Forster et al., 2016; Johnson et al., 2019). For these purposes nudging is generally not applied to ensure the model accurately simulates specific meteorological conditions, but rather to reduce the meteorological variability. As such the nudging reference used may be a climatology rather than a single time series, though reanalysis datasets are still often used as the source for the climatology rather than model output.

Considerations when using nudging

In all use cases of nudging there are a number of choices to be made about how nudging is implemented that can significantly affect results. These include which variables are nudged, the source of the reference dataset nudged towards, the value of G, the time resolution of the reference dataset, and the model levels over which nudging is applied. While the first two settings can usually be determined by the intended aim of the study, the other three are foten determined empirically. Some studies have investigated techniques to determine appropriate settings a priori (Omrani et al., 2012); however this is relatively unusual particularly for studies investigating radiative forcing.

An overview of these settings and the considerations required is given below for each. In all cases these choices depend upon the reason for applying nudging, and this overview focuses on the consequences for studies using nudging to investigate radiative forcing, particularly for mechanism denial experiments, since that is the application of nudging used in this thesis. That said, it is useful to draw on studies using nudging for a range of purposes, rather than just those investigating radiative forcing, for informing these choices.

The first choice to consider is which variables are nudged. As aforementioned, typically nudging includes potential vorticity, divergence and/or horizontal wind fields. Nudging these variables is a basic step to constraining the large scale circulation and related natural variability of the model simulation (Jeuken et al., 1996). Nudging just horizontal winds or potential vorticity has been observed to have little effect on radiative forcing estimates, but has been found to affect other model processes such as convective precipitation (Lin et al., 2016), high cloud fraction and ice water path (Johnson et al., 2019) and dust emissions (Sun et al., 2019).

Many studies also choose to nudge a temperature-related quantity such as potential temperature or dry static energy to constrain the model temperature field. A number of studies investigating aerosol radiative forcing that applied nudging to reduce the required integration time concluded it is better to avoid nudging temperature related quantities for such studies (Zhang et al., 2014, 2022). This is because nudging temperature related quantities was found to significantly affect calculated radiative forcing, for example by causing changes to high cloud formation. Likewise, in a mechanism denial experiment investigating circulation adjustments to BC perturbation Johnson et al. (2019) also avoided nudging temperature to allow thermodynamic adjustments to BC heating to take place, deciding to nudge only the horizontal winds.

Other variables are also nudged in some studies. For example, a number of studies have tried nudging specific humidity. Sun et al. (2019) found that nudging specific humidity in addition to horizontal winds in Energy Exascale Earth System Model Atmosphere Model Version 1 (EAMv1) towards a reanalysis caused significant biases in precipitation, cloud forcing, and aerosol burdens relative to the reanalysis. Studies have also found that nudging both temperature and specific humidity can degrade model performance (Sun et al., 2019). Other studies have nudged surface pressure, finding it improves representation of dynamical processes such as cyclogenesis (Brill et al., 1991) though may offer only small improvements (Jeuken et al., 1996).

The next choice is what reference dataset is used to nudge towards. For studies attempting to recreate specific observed atmospheric conditions typically a reanalysis dataset is used as the nudging source. However, nudging towards reanalysis data can cause a number of issues due to differences in the model producing the reanalysis and the model being nudged. For example, Zhang et al. (2014) investigated the effects of nudging horizontal winds and atmospheric temperature for a PI to PD anthropogenic and biomass burning aerosol perturbation using the Community Atmosphere Model 5 (CAM5) and ECHAM6-HAM2 models. Due to a cold bias in CAM5 compared to the reanalysis, they found that nudging the atmospheric temperature towards a reanalysis dataset caused a significant decrease in the LW cloud forcing, by suppressing the increase in formation of cloud ice that was seen in a parallel non-nudged simulation pair due to increased aerosol concentrations. However, for the ECHAM6-HAM2 model the sensitivity of LW cloud forcing and cloud ice formation to nudging the atmospheric temperature field was much smaller. This highlights the risk that systematic biases between the nudged model and reanalysis may induce unintended responses but that this risk varies by model. Likewise Davis et al. (2020) note that in the Whole Atmosphere Community Climate Model (WACCM) the tropical tropopause is almost 1km higher than in the Modern-Era Retrospective analysis for Research and Applications, Version 2 (MERRA2) reanalysis dataset. Nudging WACCM to MERRA2 they suggest could therefore result in changing the model's tropical tropopause height resulting in a number of artificial changes to the stratosphere such as upwelling.

Another issue with nudging to renalayses is differences in the orography fields used in the model producing the reanalysis and the nudged model. For example, between the Unified Model (UM) and models used to generate the ERA reanalysis, orography can differ by as much as hundreds of metres over the Andes and Antarctica (Telford et al., 2008). This can affect the interpolation of reanalysis data onto the nudged model grid, introducing errors. When nudging to a reanalysis nudging in the lowermost model levels can consequently generate instabilities (Forster et al., 2016).

An alternative to nudging to a reanalysis dataset is nudging to output from the model itself. This avoids the issues of orography differences and systematic biases, although is only useful where the aim is not to reproduce observed meteorological conditions, only to constrain the model, as in most studies of radiative forcing. For example Zhang et al. (2014) demonstrated a method of nudging using output from a control simulation of CAM5 with only the synoptic variability from a reanalysis to eliminate the systematic bias. With this nudging method they found the LW cloud forcing difference was reduced compared to nudging to a reanalysis. Using the EAMv1 model Sun et al. (2019) likewise found that the impact of nudging is smaller when nudging to the model's own meteorology and recommended using model climatology for nudging reference data where replicating observations or specific meteorological events is not required.

However even nudging to the model's own climatology can affect the simulation. The study of Lin et al. (2016) found that in simulations with CAM5 nudging the horizontal winds affected convective precipitation. They also found that nudging the horizontal winds and temperature significantly suppressed both the change in precipitation and cloud forcing. Sun et al. (2019) and Zhang et al. (2022) found that when nudging horizontal winds and temperature in EAMv1 to the model's own climatology there was a large bias in CRE over the marine stratocumulus regions and trade cumulus regions, particularly in the Pacific Ocean (figure 1.6b and d). They also concluded that when nudging horizontal winds or horizontal winds and temperature there were decreases in dust emissions due to decreased horizontal wind speeds near the surface. However, it is worth noting that their figure S18 demonstrates the decreases in dust emissions and near surface horizontal winds were greater for simulations nudged to reanalyses.

The next choice to consider is the value of G used. Ideally the value of G used should be large enough to sufficiently constrain the nudged variables but not too strong to introduce undue artificial forcing into the model or greatly interfere with model processes (Jeuken et al., 1996; Telford et al., 2008; Lin et al., 2016). A wide range of values have been used in the literature, though fewer studies have presented results comparing the performance of using different values of G. These studies typically compare the bias or spatio-temporal errors between the nudged model and reference dataset nudged towards across values of G. For example, Merryfield et al. (2013) tested varying G between 1/96 and 1/6 h^{-1} when nudging horizontal winds and temperature to the ERA40 reanalysis. They decided that the optimum value of G was that for which the root mean square difference (rmsd) in temperature between the nudged simulation and the nudging input was similar to the rmsd between various reanalysis datasets covering the same period. Over a range of model levels they determined that a value of 1/24 or 1/48 h^{-1} for G were had most similar rmsd to the rmsd between reanalyses. $G = 1/6\ h^{-1}$ constrained the model more strongly resulting in a lower rmsd while using $G = 1/96\ h^{-1}$ generated a higher rmsd. (It should be noted that Merryfield et al. (2013) only used a similar technique to nudging, which they refer to as constant increment nudging, though using a spatial filter they find this behaves very similarly to nudging). Davis et al. (2022) found that in simulations with a dynamical timestep of 30 minutes, when nudging horizontal winds and temperature to its own model output, the optimum value of G considering a number of metrics was between 1/24 and 1/12

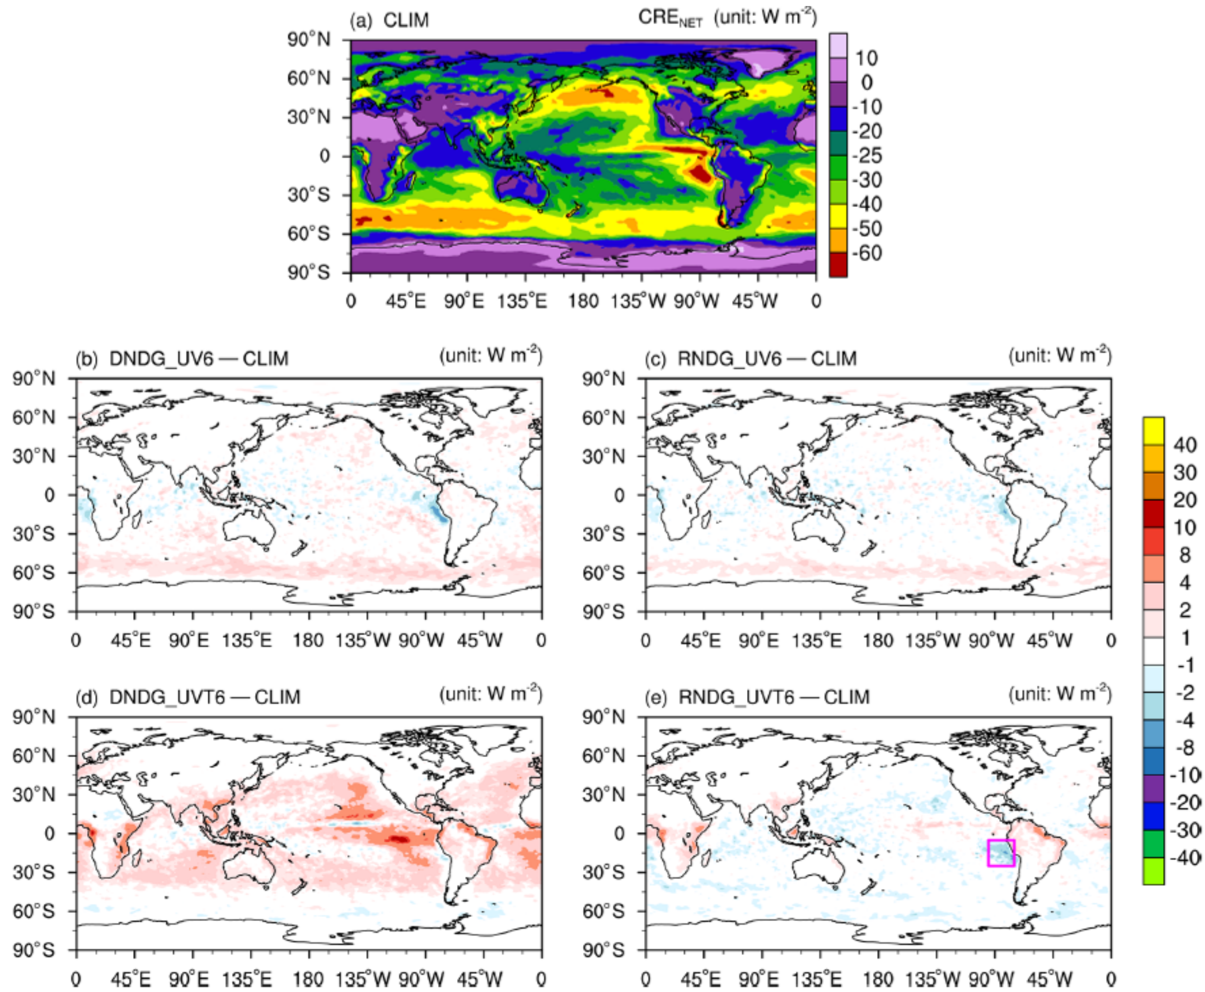

Figure 1.6: Cloud forcing errors arising from nudging to model climatology in the study of Zhang et al. (2022). (a) illustrates the annual mean CRE in a free simulation while (b-e) illustrate the difference in CRE between the free simulation and in simulations nudged to the free simulation. (b) and (c) have horizontal winds nudged only while (d) and (e) have horizontal winds and temperature nudged. The nudged simulations in (b) and (d) used a default nudging workflow (DNDG) in which the nudging tendency is calculated before convection and cloud physics schemes while the nudging reference data from the free simulation is calculated after these. For (c) and (e) a revised nudging workflow (RNDG) was used in which the nudging tendency is instead calculated at the same point in the timestep as when the reference data is generated. From Zhang et al. (2022).

h^{-1}. This varied slightly depending on the metric, with 1/12 or 1/24 h^{-1} most appropriate for trace gas concentrations, convective mass flux and horizontal wind spatial errors, but 1/24 and 1/48 h^{-1} being better for temperature spatial error and for biases across several fields.

Another way to assess the appropriate value of G is to consider the timescale of processes in the model. For example, Kooperman et al. (2012) considered that nudging may suppress only adjustments at longer timescales than the nudging timescale. They suggested this could be useful for isolating certain model processes (e.g. cloud feedbacks) in aerosol perturbation experiments. However, this could be detrimental where such processes are not intended to be suppressed.

These studies illustrate the importance that the aim of the study has for determining the appropriate value of G (like many of the other settings discussed here). Firstly, when tuning G, one should consider what variables and processes are of most importance (e.g. meteorology versus trace gas concentrations). If one variable is the focus of the study then errors introduced into other variables may be of less importance. Secondly, if the aim of the study is to approximately recreate the meteorology over a given period while nudging to a reanalysis, then the approach of Merryfield et al. (2013) is a good example of how to achieve sufficient, but not excessive, nudging that matches the real world uncertainty in observations. However, if using nudging for a mechanism denial experiment, achieving a very similar value to that of the nudging reference is important, so a higher value of G may be more appropriate. Also, while spatio-temporal correlations and rmsd are useful metrics, studies should not neglect to assess the actual bias too, especially where biases exist between the model and reference dataset, such as in mechanism denial experiments. Finally, one should also consider the resolution of the model and whether any processes are expected to be suppressed by the nudging. The effect of varying G may also depend on the model timestep, which will vary between models.

Another choice to be made is the nudging reference input time resolution to use. Many studies in the past have used 6 hourly input. However, Zhang et al. (2022) illustrated well how poorly this can represent a diurnal cycle in temperature and horizontal winds, particularly when many nudging schemes use a linear interpolation between input points. This is apparent from figure 1.7, where use of 6 hourly input causes large deviations in temperature and horizontal wind speed, particularly around peaks or troughs in either variable. Consequently they found that using 3 hourly nudging input significantly improved representation of a number of metrics. Increasing the time resolution of input to 1 hourly resulted in smaller gains. Davis et al. (2022) similarly tried a range of nudging input time resolutions, finding that generally the higher the resolution of input, the smaller the biases in nudged fields between the nudged model and the nudging input were.

As with G, the nudging input time resolution may also affect what processes are suppressed. In figure 1.7 one can see that nudging with coarser temporal resolution input data could fail to reproduce shorter timescale processes on the order of hours as the linear interpolation tends to relax the model towards a coarser representation of the diurnal cycle. By contrast this would be unlikely to affect processes that have timescales on the order of days. Jeuken et al. (1996) and Zhang et al. (2022) both noted that as spatial resolution in models increases, smaller scale (both spatially and temporally) processes will be resolved by models. It is therefore understandable that 6 hourly nudging input may have been acceptable when used in simulations with a spatial resolution of 500 km by Jeuken et al. (1996), but the simulations conducted by Zhang et al. (2022) having a resolution of $\sim$110 km benefit significantly from a higher resolution input.

Another choice to be made when nudging is what model levels to nudge. At the lower altitude limit of nudging, some studies suggest that nudging in the boundary layer can generate instabilities (Telford et al., 2008; Forster et al., 2016). However, this is mainly in the case of nudging to reanalyses that have differing orography to the nudged model, as discussed earlier. Sun et al. (2019) tested simulations with or without nudging applied at pressures above 850 hPa and found no significant difference between the two simulations. These simulations were nudged towards a reanalysis but they only nudged horizontal winds not temperature. At the upper end,

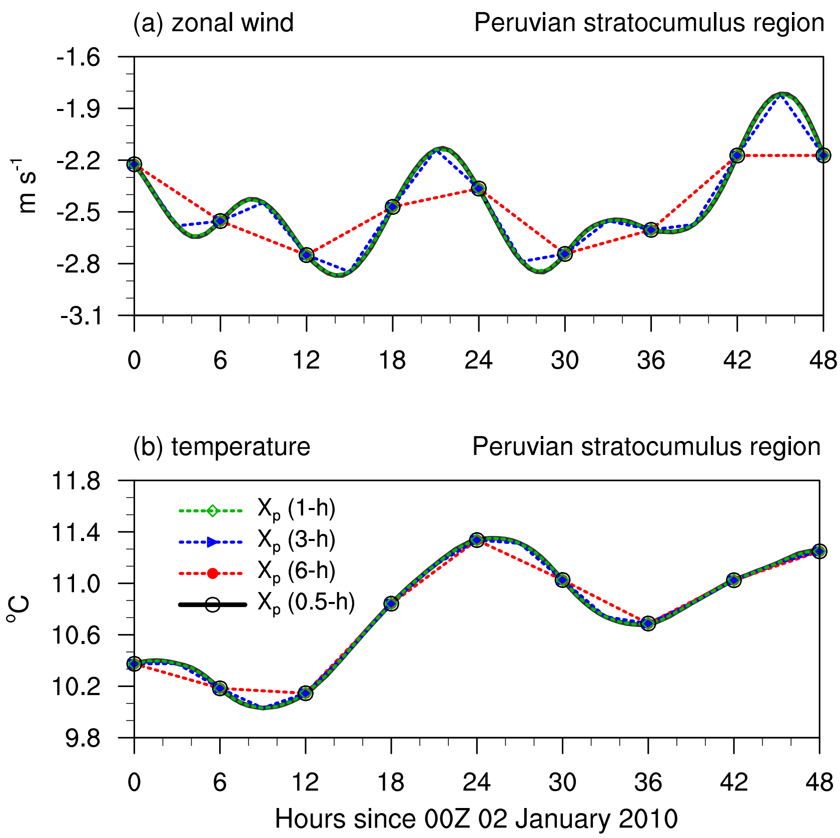

Figure 1.7: Plot illustrating the effect of varying the time resolution of nudging input data for (a) eastward wind speed, in m s^{-1}, and (b) atmospheric temperature, in degC. The thick black lines are area mean values output from a simulation over the area outlined by a pink box in figure 1.6e at 700 hPa during a 48 h period, witha timestep of 0.5 h. The dotted lines represent values interpolated to the same timesteps using 1 (green), 3 (blue) or 6 (red) hourly output from the same simulation. From Zhang et al. (2022).

Davis et al. (2022) linearly tapered G to 0 between 1 hPa and 0.1 hPa to avoid interfering with atmospheric tides and gravity waves.

Finally, while not always a choice, one should consider where in the model timestep sequence nudging is applied. Zhang et al. (2022) illustrated that, when nudging to input generated from a control run of the same model, if the nudging tendency term in equation 1.16 is calculated at a point in the timestep sequence of the nudged simulation that is different to the point the input is generated in the control run, an extra forcing term is introduced. This term will include the forcing from processes that occur in the model timestep between the point the nudging tendency is calculated and the point in the control simulation when the output is generated. In the study of Zhang et al. (2022) for example, this included deep convection, turbulence and stratiform cloud parametrizations. In their study this effect caused significant biases between nudged simulations and nudging input across a range of metrics. In particular this effect was responsible for the large bias in the CRE over marine stratocumulus regions noted by Sun et al. (2019) when nudging temperature, illustrated in figure 1.6.

While the point in the timestep that nudging tendency calculation and application are important, it is worth noting that it is not an easy aspect of the nudging to change in an existing model. It is also something that may pose an issue for nudging to reanalysis datasets, where the sequence of model processes cannot practically be altered in the reanalysis dataset to match the sequence in the nudged model. However, this term may be small compared to that of any biases between the nudged model and reanalysis dataset.

Summary of nudging considerations

Having considered all of the above choices it is worth highlighting that, since most of these settings are determined empirically, there is typically no ‘correct’ nudging setup, only a best case (Lin et al., 2016; Sun et al., 2019). Furthermore, even in an optimum setup, nudging may not achieve the desired effects. For example, Gettelman et al. (2020) found that in nudged simulations with CAM6 the smaller the bias achieved in atmospheric temperature and boundary layer structure, the worse the simulation of clouds became. While they noted this may be partly due to nudging to a reanalysis dataset, they also suggested this may result from nudging interfering with cloud processes.

In summary, nudging is a technique with a range of applications but is always an approximate one. What is clear is that the optimum settings for nudging depend on both the model and the study aim. Hence it is advisable to test a number of settings with the model to be used, or refer to studies that use the same model and/or that apply nudging for a similar purpose.

1.3.5 Effective radiative forcing and radiative adjustments due to anthropogenic black carbon and sulphate aerosols

As noted earlier, the ERF due to aerosols arises from a number of radiative adjustment processes in addition to their IRF. Figure 1.8 captures this large range of processes as a schematic (Bellouin et al., 2020a). As can be seen from figure 1.8, there are many radiative adjustment processes documented in the literature that involve cloud responses, with these responses varying by cloud type and height and aerosol species. These radiative responses are often divided into aci and ari effects as described in section 1.3.1. Non-cloud ari effects are also important and divided by many studies into adjustments associated with surface albedo, humidity, atmospheric temperature and/or lapse rate and surface temperature (Smith et al., 2018, 2020a). In this section a brief account of the radiative adjustments specific to SU and BC is given, particularly focusing on the differences between the effects of these two aerosols and what processes are important for each.

SU and BC can be categorised as scattering and absorbing aerosols respectively (Li et al., 2022). Scattering aerosols strongly scatter incident SW radiation, with a portion of that radiation 'backscattered'. This reduces the amount of solar radiation reaching the Earth's surface, increasing the albedo of the atmosphere. This generally causes a decrease in net downward radiation, resulting in a negative radiative forcing and cooling of the surface and atmosphere (Li et al., 2022). By contrast, absorbing aerosol strongly absorbs incident SW radiation, and re-radiates that energy as longwave radiation. This warms the local atmosphere and tends to decrease planetary albedo and cause a positive radiative forcing (Li et al., 2022). However, absorbing aerosol tends not to warm the surface much or even causes a cooling at the surface since the heating effect from absorption is higher in the atmosphere (Satheesh and Ramanathan, 2000; Li et al., 2022).

Aside from their direct interactions with radiation, aerosols have strong effects on radiation via clouds. SU aerosol has a large negative ERF_{aci} (Boucher et al., 2013; Smith et al., 2020a; Forster et al., 2021). This is due to SU being water soluble and acting as a CCN (Grosvenor and Carslaw, 2020; Li et al., 2022) and thus increasing the amount of radiation scattered by clouds due to the indirect effects described in section 1.3.1 (Twomey, 1977; Albrecht, 1989). The aci effects of SU are large and significantly increase the negative ERF due to SU (Boucher et al., 2013; Smith et al., 2018, 2020a; Forster et al., 2021).

ERF_{aci} due to BC perturbation is smaller than for SU (Johnson et al., 2019) since BC contributes a smaller fraction of CCN (Block et al., 2024). This is because freshly emitted BC is water insoluble and so does not act as cloud droplet condensation nuclei (Li et al., 2022). However, after emission BC undergoes ageing in the atmosphere, becoming coated with hydrophilic species such as organic carbon or sulphate, which increases its ability to act as a CCN (Bond et al., 2013), though still to a lesser degree than SU. BC is also thought to be able to act as an ice nucleating particle, but the mechanisms of this process and its significance relative to background concentrations of ice nucleating particles is not well understood (Li et al., 2022).

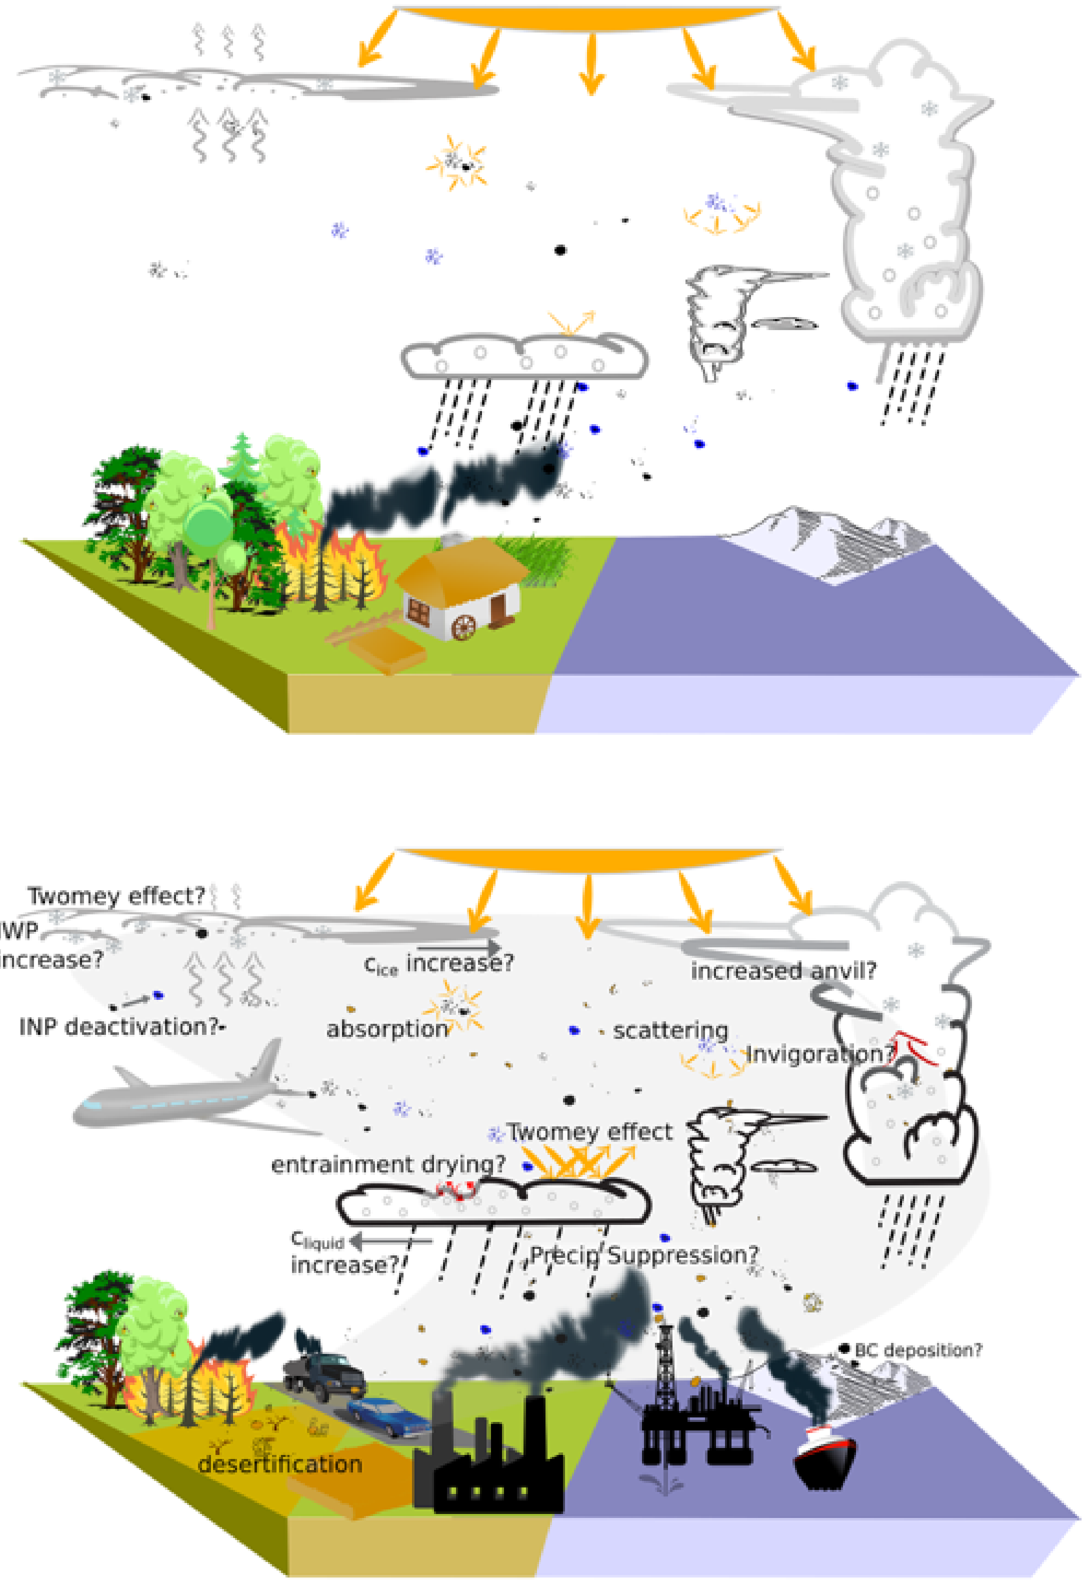

Figure 1.8: Schematic depiction of aerosol radiative effects in the (a) pre-industrial atmosphere and (b) with anthropogenic emissions of aerosols. Question marks indicate processes that are uncertain in their effects on ERF. C_{liquid} and C_{ice} refer to liquid and ice cloud fractions respectively; LWP and IWP refer to liquid and ice water paths respectively; INP refers to ice nucleating particles. From Bellouin et al. (2020a).

In terms of ari cloud adjustments, BC is generally thought to cause significant cloud radiative adjustments by affecting cloud properties via semi-direct effects (Johnson et al., 2004; Bond et al., 2013; Samset and Myhre, 2015; Sand et al., 2015; Smith et al., 2018, 2020a). Many of these effects occur due to changes in atmospheric stability and lapse rate and depend on the vertical distribution of BC aerosol (Samset et al., 2013; Samset and Myhre, 2015; Stjern et al., 2017; Herbert et al., 2020; Li et al., 2022). For example, where absorbing aerosol such as BC is at approximately the same altitude as clouds, it may cause cloud 'burn-off', where cloud reduces due to the local heating from absorbing aerosols decreasing vapour saturation (Koch and Del Genio, 2010). However, absorbing aerosol above clouds may increase cloud liquid water paths and fractions by increasing atmospheric stability (Koch and Del Genio, 2010; Samset and Myhre, 2015; Stjern et al., 2017; Herbert et al., 2020). This effect is particularly strong for marine stratocumulus clouds (Johnson et al., 2004; Koch and Del Genio, 2010; Herbert et al., 2020). Absorbing aerosol below cloud may also increase cloud by increasing convection (Koch and Del Genio, 2010). Overall these effects are thought to lead to increases in low cloud and decreases in high cloud, the former causing a negative radiative effect due to scattering incoming SW radiation, and the latter also causing a negative radiative effect due to decreasing absorption of outgoing LW radiation by cloud (Koch and Del Genio, 2010). Since these adjustments broadly depend on the change in atmospheric temperature and heating (and consequent changes in atmospheric stability and convection) resulting from the presence of BC or other absorbing aerosol, these cloud adjustments could be approximately described as atmospheric 'temperature-mediated' cloud adjustments.

As well as these atmospheric temperature-mediated adjustments, significant changes in ERF_{ari} can arise from changes in large-scale circulation and dynamical adjustments due to local and remote aerosol forcing (Sherwood et al., 2015; Myhre et al., 2017; Wilcox et al., 2023). For example, Williams et al. (2022) found significant differences in the cloud response to absorbing aerosol perturbations in different locations, with negative cloud radiative adjustments when perturbations were applied over south east Asia and the western Pacific, and generally positive cloud radiative adjustments when perturbations were applied over the eastern Pacific. They suggested this effect was caused by the local heating from absorbing aerosol interacting with the Walker circulation. They found that increasing absorbing aerosol over the western Pacific caused an enhancement of the Walker circulation , which increased liquid water content in the tropics. This resulted in a negative radiative forcing due to increased scattering of SW by clouds. Whereas when increasing absorbing aerosol over the eastern Pacific, the heating due to aerosol perturbation was confined locally by the opposing descent of the Walker circulation, resulting in decreased liquid water content due to decreased relative humidity. Other studies have likewise identified adjustments to clouds and precipitation due to interactions between aerosol forcing and dynamics, for example, due to interactions with the east Asian monsoon (Wilcox et al., 2019) and from aerosol emissions reductions over eastern China during the Covid-19 pandemic (Fahrenbach and Bollasina, 2023).

Considering all these ERF_{ari} cloud adjustments together, for BC perturbation ari cloud adjustments have been generally found to cause a negative radiative forcing that opposes the positive BC IRF (Stjern et al., 2017; Smith et al., 2018; Thornhill et al., 2021; Li et al., 2022). However this can vary significantly by model, with studies using some models having close to zero BC ERF_{ari} contributions from cloud adjustments (Johnson et al., 2019; O'Connor et al., 2021), and others reporting positive radiative adjustments (Allen et al., 2019).

Fewer studies have quantified ari cloud adjustments for SU, with most studies focusing on aci adjustments to SU perturbation, or grouping all cloud adjustments together (Smith et al., 2018; Thornhill et al., 2021). This is because, as a scattering aerosol, SU tends to have a stronger effect on surface temperature than atmospheric temperature (Li et al., 2022). However, circulation adjustments are known to arise from SU perturbations (Dong et al., 2016; Wilcox et al., 2019) and these will contribute to ERF_{ari}. Likewise, SU perturbation does affect atmospheric temperatures and so should be expected to cause some ari cloud radiative adjustments. Additionally, in fSST experiments where the land surface temperature can respond, cloud feedbacks will arise in response to this land surface temperature response and, while technically a radiative feedback, will be included in ari cloud adjustments calculations.

In terms of non-cloud adjustments, BC has generally been found to cause a negative atmospheric temperature radiative adjustment and a positive water vapour radiative adjustment (Smith et al., 2018; Thornhill et al., 2021). This is because BC heats the atmosphere, which both increases the amount of LW radiation emitted by the atmosphere and increases specific humidity, which increases the greenhouse effect of water vapour. SU has been found to cause the opposite effect (positive atmospheric temperature radiative adjustment and negative water vapour radiative adjustment), but with smaller magnitudes than for BC (Smith et al., 2018). Again this could be ascribed to SU having less effect on atmospheric temperature than surface temperature compared to BC.

Although the forcing from surface temperature response, the planck effect, is considered to be a radiative feedback not an adjustment, previous studies have quantified the radiative forcing from land surface temperature change under aerosol perturbations in fSST experiments. Smith et al. (2018) found this effect to be of the same sign as the atmospheric temperature radiative adjustment for both BC and SU perturbations, but of much smaller magnitude. By contrast Thornhill et al. (2021) found a close to zero surface temperature radiative response comparing results from 7 different models for BC perturbation.

Aside from temperature and water vapour adjustments, both aerosols can cause responses in albedo. BC deposition on snow or ice has been found to decrease surface albedo and enhance melting of snow and ice, causing a positive radiative adjustment (Clarke and Noone, 1985; Hansen and Nazarenko, 2004; Flanner et al., 2009; Bond et al., 2013; Forster et al., 2021). This mechanism has been assessed to cause only a small positive radiative adjustment, on the order of 0.01 W m^{-2} in more recent studies (Lin et al., 2014; Namazi et al., 2015). While SU does not affect surface albedo in this way, it has also been assessed to cause a small surface albedo response (Smith et al., 2018). This could occur because of changes in land surface temperatures in fSST experiments, which result in changes in albedo such as via changes in snow cover.

On a basic level one can summarise the key difference between SU and BC as perturbation of SU causes a negative ERF and tends to cause cooling of the atmosphere while perturbation of BC causes a positive ERF and tends to causing warming of the atmosphere. The differences in radiative effects between the two aerosols are more complex however and do not have the same symmetry as their ERFs, but rather involve different processes, as well as the same processes but to different degrees. Overall, much of these differences in radiative effects between the two aerosols arises from the fact that SU is a soluble and scattering aerosol, while BC is an insoluble and absorbing aerosol.

Chapter 2

Development and testing of model nudging process denial experiments for investigating radiative adjustments in UKESM1

2.1 Introduction

This work applies a similar method to that used by Johnson et al. (2019) described in section 1.3.3, in which nudging is used as a mechanism denial tool to isolate different radiative adjustments to aerosol radiative forcing. In this case, as well as using nudging of horizontal winds to suppress circulation adjustments, an additional set of experiments are conducted that nudge potential temperature as well to suppress atmospheric temperature mediated adjustments.

These experiments were conducted with perturbation of either sulphate (SU) or black carbon (BC) aerosol from pre-industrial (PI; 1850) to present day (PD; 2014) emissions to investigate their radiative adjustments. The motivation for conducting the experiments with perturbations of two different aerosols separately is firstly to investigate how their radiative adjustments differ, and also to demonstrate application of the methodology to two aerosols that behave differently.

Before running these simulations it is important to determine what are the appropriate nudging settings to use. This chapter presents an assessment of various nudging settings: the variables nudging is applied to; the value of G used; the model levels over which nudging is applied; and the time resolution of nudging input. This assessment primarily aims to identify an optimum nudging setup for the methodology used here and its application to SU and BC radiative adjustments in the following chapter. However, it also seeks to provide a wider assessment of optimal nudging settings, particularly for application of this same methodology with other forcing agents or models, and more broadly for application of nudging to simulations intended to investigate radiative forcing and adjustments. The earlier section 1.3.4 provides useful context for the assessment of nudging settings, explaining what they are and their significance.

This chapter is laid out as follows. First a detailed description of the theory behind the methodology is presented in section 2.2 as well as some of the higher level choices made in the experiment design. Then details of the model and simulation configurations used are given in section 2.3. Following this, a description of the nudging code used is given, and a description of the modifications to the code made as part of this thesis in section 2.4. Next are details of what nudging variables were investigated and the simulations conducted to achieve this in section 2.5. Then the results of the simulations testing different nudging variables are presented in section 2.6. And finally a discussion of the optimal nudging settings to be used in the following chapter, as well as any wider significance of the results for each variable, is presented in section 2.7.

2.2 Methodology for using nudging for process denial of radiative adjustments

The principle of the methodology is outlined as follows. By nudging variables to the same reference state, X_{ref}, across simulations with or without a forcing agent perturbed, the difference in those variables (and consequently, any variables dependent on the nudged variables) between the simulations should be suppressed. This principle can be applied to radiative adjustments. For example, in two fixed SST simulations, one with pre-industrial (PI) and the other present-day (PD) emissions of aerosols, the atmospheric temperature will differ between the simulations, primarily due to the difference in aerosol radiative forcing. By nudging the atmospheric temperature in the two simulations to the same reference state, the difference in atmospheric temperature due to the aerosol forcing is suppressed, as is the radiative effect of the temperature change. In addition, differences between the simulations in other fields that respond to the atmospheric temperature change (e.g. specific humidity, cloud fields, surface temperature) should also be suppressed. By comparing such a nudged simulation pair with a non-nudged simulation pair it should then be possible to determine the effect of the aerosol emission change on the atmospheric temperature, the consequent effects on dependent fields, and the radiative forcing associated with changes in these fields.

Here, nudging of the model horizontal wind fields (eastward wind speed, u, and northward wind speed, v) and atmospheric potential temperature (θ) are used to isolate the contributions of circulation and atmospheric temperature mediated adjustments that arise under perturbations of either anthropogenic SU or BC aerosol emissions. Figure 2.1 outlines the experiments required and how the adjustments are determined from these. Firstly, the difference in top of atmosphere (TOA) net downwards fluxes between a pair of free running simulations, one with a PD (control) state and the other with PI (perturbed) anthropogenic emissions of either SU or BC, determines the ERF. This uses the fSST method of determining ERF (Hansen et al., 2005; Forster et al., 2016). The difference in other fields that respond to the aerosol perturbation determines the free adjustment of those fields. Secondly, the difference in any field between a control and perturbed simulation both with u and v nudged (uv-nudged), instead returns the free adjustment of the field minus those resulting from circulation adjustments. The corresponding difference in TOA fluxes is the ERF minus the radiative effect of circulation adjustments. Thirdly, the difference in any field between a control and perturbed simulation with their horizontal wind fields and θ nudged (uvθ-nudged) instead returns the adjustments minus both circulation and atmospheric temperature mediated adjustments. The corresponding difference in TOA fluxes is the ERF minus the radiative effects of circulation and atmospheric temperature mediated adjustments.

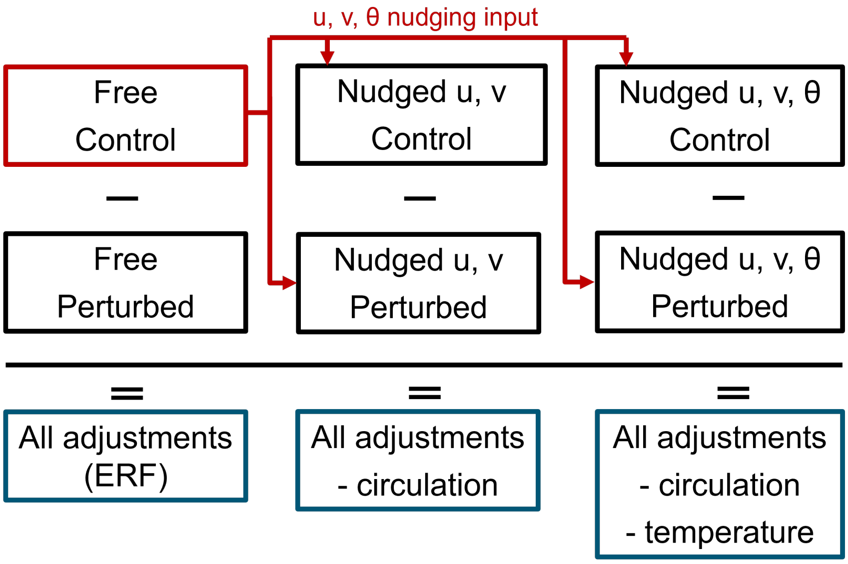

Figure 2.1: Schematic illustrating the experiments conducted and the determination of circulation and atmospheric temperature adjustments from the differences between simulations. Control simulations have all conditions set to 2014; perturbed simulations have emissions of the aerosol of interest set to 1850.

To look at this in reverse, in the uvθ-nudged simulation pair the key remaining radiative responses should be the DRE, aci adjustments (though temperature changes could plausibly affect the second indirect effect), cloud masking effects, part of the surface albedo adjustment (that is not caused by temperature), and any chemical adjustments. In the uv-nudged simulation pair, the radiative responses should include all of these plus those mediated by temperature, particularly including: cloud adjustments relating to adjustments in vertical heating rates and stability, tropospheric and stratospheric temperature adjustments (including lapse rate adjustment) and humidity adjustment. The free simulation pair radiative responses should include all those in the uv-nudged simulation pair as well as circulation-mediated adjustments. Circulation adjustments relate to changes arising from adjustment to large-scale dynamics, which can arise particularly due to changes in atmospheric absorption and land surface temperature change (Myhre et al., 2017), particularly including the effects of shifts in cloud patterns (Sherwood et al., 2015).

The effect of this nudging can also be better understood by reflecting back on figure 1.4 and considering how nudging would affect the interactions between aerosol and other components of the climate system. Figure 2.2 illustrates the same schematic but with the suppression due to nudging represented also. As can be seen, nudging itself only directly suppresses adjustment of atmospheric temperature and horizontal winds. In the case of temperature, this directly suppresses radiative adjustment due to atmospheric temperature response. However, the suppression of adjustment of temperature and winds also suppresses adjustments that respond to these variables, particularly adjustment of clouds (primarily semi-direct effects), as well as the distribution of aerosols and greenhouse gases, and surface temperature. The suppressed adjustment of these other components of the climate system then reduce radiative responses that would have otherwise arisen due to their adjustment.

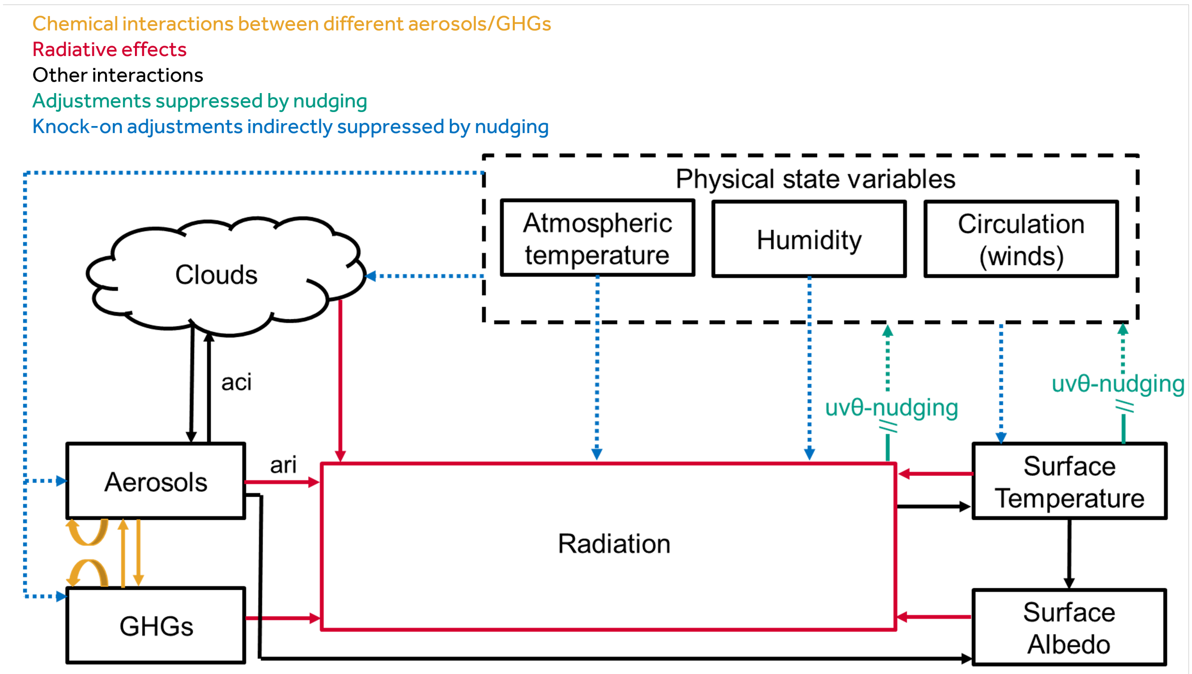

Figure 2.2: Illustration of interactions between aerosols and greenhouse gases (GHGs) and different components of the climate system that affect radiation, but with nudging applied to suppress adjustment of horizontal winds and atmospheric temperature. After Boucher et al. (2013).

For any given field, taking the difference between the control-perturbed difference of the field in the free simulation pair minus the control-perturbed difference of the field in the uv-nudged simulation pair, gives an estimate for the contribution of circulation adjustments to the overall adjustment in that field. Doing the same for the uv-nudged control-perturbed difference minus the uvθ-nudged control-perturbed difference gives an estimate for the contribution of atmospheric temperature mediated adjustments to the overall adjustment in that field. Similarly, taking the difference between radiative forcing in the free and uv-nudged pairs gives the total radiative adjustment due to circulation-mediated adjustments; and the difference in radiative forcing between uv-nudged and uvθ-nudged pairs gives the total radiative adjustment due to temperature-mediated adjustments. This can be illustrated algebraically:

$$\Delta X_{\mathrm{circ}} = \Delta X_{\mathrm{free}} - \Delta X_{\mathrm{uv\text{-}nudged}} \tag{2.1}$$

and:

$$\Delta X_{\mathrm{T}} = \Delta X_{\mathrm{uv\text{-}nudged}} - \Delta X_{\mathrm{uv\theta\text{-}nudged}} \tag{2.2}$$

where ΔX_{circ} is the adjustment in any variable (including TOA radiative adjustment forcing) due to circulation adjustments in response to the aerosol forcing, ΔX_{T} is the adjustment in any variable (including TOA radiative adjustment forcing) due to atmospheric temperature adjustments in response to the aerosol perturbation, and ΔX_{free}, $\Delta X_{\mathrm{uv\text{-}nudged}}$, and $\Delta X_{\mathrm{uv\theta\text{-}nudged}}$

are the adjustments in the variable in the free, uv-nudged and uvθ-nudged control-perturbed simulation pairs respectively. In this way, the adjustments to individual variables resulting from circulation and atmospheric temperature responses can be determined, as well as the radiative adjustment forcing caused by circulation- and temperature-mediated adjustments.

As illustrated in figure 2.1 the nudging inputs for u, v, and θ are derived from the free control experiment. Nudging towards output from the same model used for the experiments was chosen instead of nudging towards reanalysis data (as is often done with nudging) for two main reasons (see section 1.3.4 for further context). Firstly, it prevents differences in the underlying models used for the simulations and the reanalysis from affecting the response to nudging (Kooperman et al., 2012). This would be particularly detrimental if there were differences in orography (Telford et al., 2008), or systematic differences in temperature (Zhang et al., 2014), between the nudged model and reanalysis providing nudging input. Secondly, having the same orography between the nudging input and the model runs means nudging can be applied to the lowermost model levels, which could otherwise generate model instabilities (Forster et al., 2016). Since significant adjustments in winds and temperature resulting from SU and BC perturbations are expected in the lower levels, nudging these levels is expected to be important in suppressing adjustments.

As well as choosing from what source the nudging input is derived, another choice to be made was how to process the u, v, and θ fields from the free control to generate the nudging input fields. The simplest approach is to nudge u, v, and θ at each timestep in the nudged simulations to u, v, and θ values from the corresponding timestep in the free control experiment. To reduce computational requirements, u, v, and θ data can be taken from only every few timesteps (i.e. a coarser temporal resolution) of the free control, with their values on intermediate timesteps calculated from interpolation. This approach results in the nudged variables in each nudged simulation having a similar time series of natural variability to that of the free control, and by extension, to each other nudged simulation. Alternatively, one could determine a climatology for u, v, and θ in the free control, by averaging over multiple years of the free control for each timestep, and use this as the nudging input. Again u, v, and θ data can then be taken from only every few timesteps of the free control climatology using interpolation to fill the gaps. This alternative approach would cause the nudged variables in each year of the nudged simulations to follow roughly the same annual cycle, rather than matching the natural variability of each year of the free control. One potential downside to this second method is that if only a few years are averaged over, there may be a discontinuity in the nudging input between the first and last timesteps (i.e. the last timestep of December and first of January), since the first timestep of the free control does not follow its last timestep. However, the more years are averaged over the more this discontinuity would be smoothed out. Either approach should produce similar results for calculation of radiative adjustments and so in this work the first method - nudging at each timestep to values interpolated from the corresponding timestep in the free control - was chosen for simplicity.

It should also be noted that while the atmospheric temperature adjustments are determined as the difference between the uv-nudged and uvθ-nudged simulation pairs, they could alternatively be determined using a θ-nudged simulation pair differenced with the free pair. The primary reasoning for choosing to do the former here is that while it is not uncommon to nudge just the horizontal winds, it is more uncommon to nudge just θ in the literature. However, the latter method could be equally valid and should be explored in future work.

2.3 Model/simulation details

All simulations conducted use version 1.0 of the UK Earth System Model, UKESM1 (Sellar et al., 2019), based on the UK Met Office's Unified Model, UM (Brown et al., 2012). Simulations are based on the UKESM1 time-slice piClim experiments used in the Aerosol Chemistry Model Intercomparison Project (AerChemMIP) (O'Connor et al., 2021). Specifically, the model is used in an 'atmosphere-only' configuration, using the Global Atmosphere 7.1/Global Land 7.0 (Walters et al., 2019) version of HadGEM3, coupled with the interactive chemistry for stratosphere and troposphere scheme of the UK Chemistry and Aerosol model (UKCA; Morgenstern et al. (2009); O'Connor et al. (2014)). The model has 85 vertical levels going up to 85 km above sea level and uses N96 resolution grid spacing (grid cells of 1.875° longitude and 1.25° latitude). The model timestep length is 20 minutes.

Unlike the piClim experiments, the experiments conducted here are reversed to set 2014 (PD) as the control state and 1850 (PI) as the perturbed. Hence, in the control simulations all boundary conditions including SSTs, sea ice, greenhouse gases, chemistry and aerosol emissions, and vegetation properties (such as canopy height and leaf area index) are representative of 2014 conditions. This change was made to make the results more policy relevant by representing the effects of reducing aerosol emissions in a present atmosphere (whereas using a PI control atmosphere may be more useful for understanding the contributions of aerosols to climate change). In the perturbed SU experiments emissions of anthropogenic sulphur dioxide are set to an estimate of its emissions in 1850. In the perturbed BC experiments all emissions of BC (fossil fuel and biomass burning) are set to an estimate of its emissions in 1850. Biomass burning emissions were included in the perturbation to produce a clearer forcing and adjustment signal and because the anthropogenic component of biomass burning is difficult to separate out. The emissions data used are from estimates made by van Marle et al. (2017) and Hoesly et al. (2018). A further minor difference to the piClim configuration used in O'Connor et al. (2021) is that methane and nitrous oxide use prescribed concentrations using 2014 values in Meinshausen et al. (2017), rather than being determined interactively in the model.

For this work, the choice to use an ESM (over a GCM without interactive chemistry, such as HadGEM) was considered necessary to achieve the most accurate representation of aerosols and their interactions in the atmosphere. For example, including interactive chemistry enables representation of non-linear effects of BC emissions on CCN such as by BC reducing the uptake of aerosol precursor gases by soluble aerosol particles which more readily act as CCN (Koch et al., 2011; Johnson et al., 2019). While the nudging mechanism denial methodology here focuses on separating circulation- and atmospheric temperature-mediated adjustments, which are dynamical or physical rather than chemical adjustments, representing the atmospheric chemistry accurately is important to ensure the mechanism denial experiments start from an accurate representation of aerosol effect in the free case before nudging. Atmospheric temperature in particular could also have direct effects on chemistry processes, many of which are temperature dependent.

The simulations also employ additional radiation calls to separate out the DRE, CRE and residual or 'clear-clean sky' radiative forcing. The additional radiation calls calculate the TOA downwards radiative flux either without aerosol or without aerosol and clouds at each time step, and use the method of Ghan (2013) to calculate DRE, CRE and clear-clean sky radiative forcing. As noted by Zelinka et al. (2023) it is expected that the clear-clean sky forcing includes more adjustments than just the surface albedo effect, including particularly the surface, tropospheric and stratospheric temperature adjustments, the humidity adjustment and any chemical adjustments.

2.4 Nudging code and modifications

2.4.1 The UM nudging code

This work uses the existing nudging code in the UM, with modifications implemented as part of this thesis, which are outlined in the following section (2.4.2). An earlier version of the nudging code is documented in detail in Telford et al. (2008), and a more up to date description of the code can be found in Dalvi and Rodriguez (2020). Here a brief description of the nudging code is provided.

The existing UM nudging code allows nudging of three variables: u, v and θ. The nudging code requires 6-hourly time varying input for each nudged variable as the reference to nudge towards, which is linearly interpolated onto the model timesteps. 6-hourly input is required because the initial implementation of nudging in Telford et al. (2008) used the ERA-40 reanalysis dataset which was available at 6 h intervals. This input must be pre-processed into the appropriate form for the nudging code to read. The model then nudges to the input data using equation 1.16 on every model timestep. In this work the 6-hourly input is taken from the free control simulation, using

A number of settings can be varied in the nudging code. Aside from choosing which of the three variables are to be nudged, the settings of greatest relevance here are: G; the levels that are nudged; and ramps for increasing nudging linearly over a range of levels at the top and bottom of the range of levels nudged. The ramps scale the value of G applied across the specified number of levels, from 0 at the lowest level nudged to G at the specified number of levels above for the lower ramp; or 0 at the highest level nudged to G at the specified number of levels below for he upper ramp.

The nudging scheme is applied at the end of the physics and dynamics schemes. The model output fields for u, v, and θ used as the nudging input (from the free control) are also generated after the physics and dynamics schemes. The nudging input should therefore be generated (and applied) at the same point in the timestep cycle as when the nudging is applied, meeting the recommendation of other studies discussed earlier in section 1.3.4 (Zhang et al., 2014; Sun et al., 2019; Zhang et al., 2022).

2.4.2 Modifications to the nudging code

Initial testing of the nudging indicated that significant nudging tendencies were being applied to potential temperature even after only a few model timesteps. There was also a difference between the model potential temperature value output from the free control simulation and after that output was read into the model. Figure 2.3a illustrates this issue for an example timestep and grid point: the potential temperature read into the model ('input-read') is warmer than the raw input ('input-raw') at all levels, and there is a positive nudging increment ('nudge inc.'), which nudges the model value towards the warmer input values. There is also a persistent negative increment due to 'other' factors (i.e. simulated physical processes) that opposes the nudging increment ('other inc.'). This pattern was observed for many timesteps and over multiple grid points suggesting an error with the implementation of the nudging.

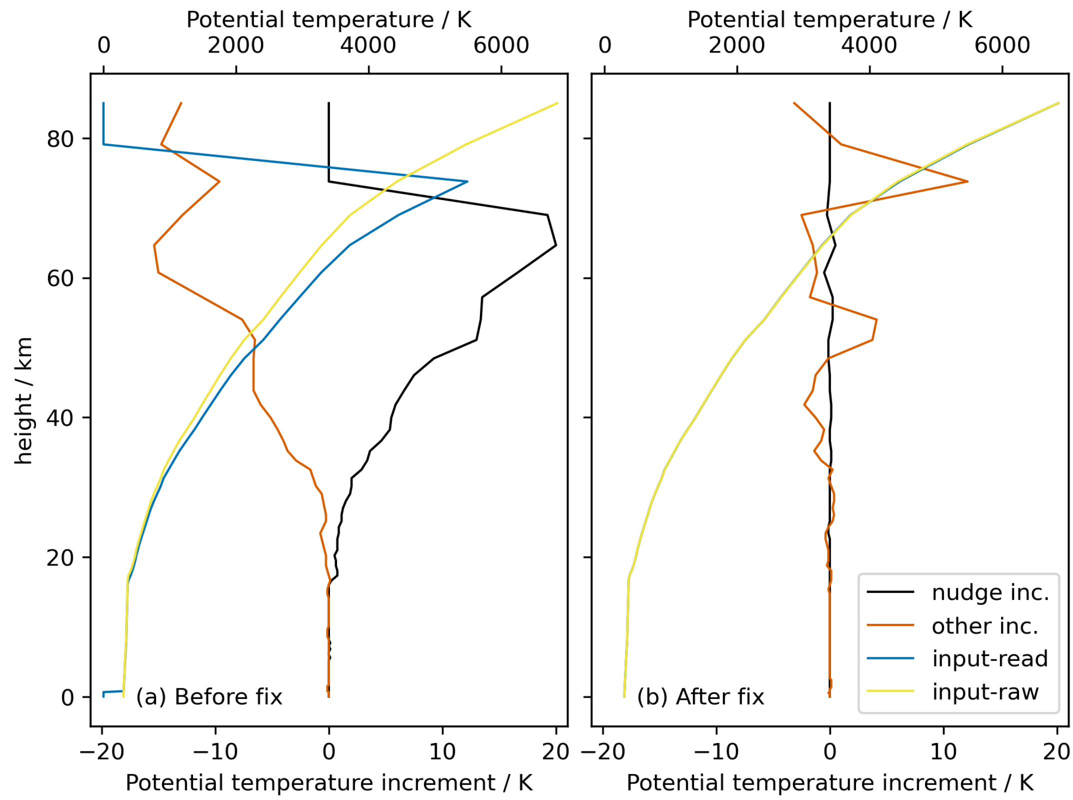

Figure 2.3: Nudging input for θ before ('input-raw') and after ('input-read') being read into the nudging code, and increments of θ due to nudging ('nudge inc.') or all other ('other inc.') processes. (a) before fixing nudging code indexing error (as described in text) and (b) after fix. All the profiles are for the same single timestep and lat/lon grid point. Increments are plotted against the bottom axis, input values against the top axis.

These tests identified the issue was a result of an indexing error in the nudging code when applied to climate simulations (as used in this work) rather than numerical weather prediction (NWP) simulations. The indexing error arises due to an extra theta level present in NWP simulations, that is not present in climate simulations. Consequently the nudging input was shifted one model level down when read into the nudging code, resulting in it being positively biased at each model level. This resulted in the positive nudging increments as the nudging attempts to force the model to warmer potential temperatures. I fixed this issue via a branch to the nudging code that corrects indexing, resulting in a better match between the raw input value of θ and the value after being read in by the model, as well as causing smaller nudging increments, as illustrated in figure 2.3b.

With the indexing error fix applied, there were still differences between the nudging input before and after being read into the nudging code. The ‘grid interpolated’ line in figure 2.4 illustrates this error, plotting the atmospheric profile of the difference between the input before and after being read into the nudging code, for an example timestep and grid point. It can be seen that below 15 km the difference could be up to ±0.5 K, while in the stratosphere and mesosphere the difference could reach several kelvin.

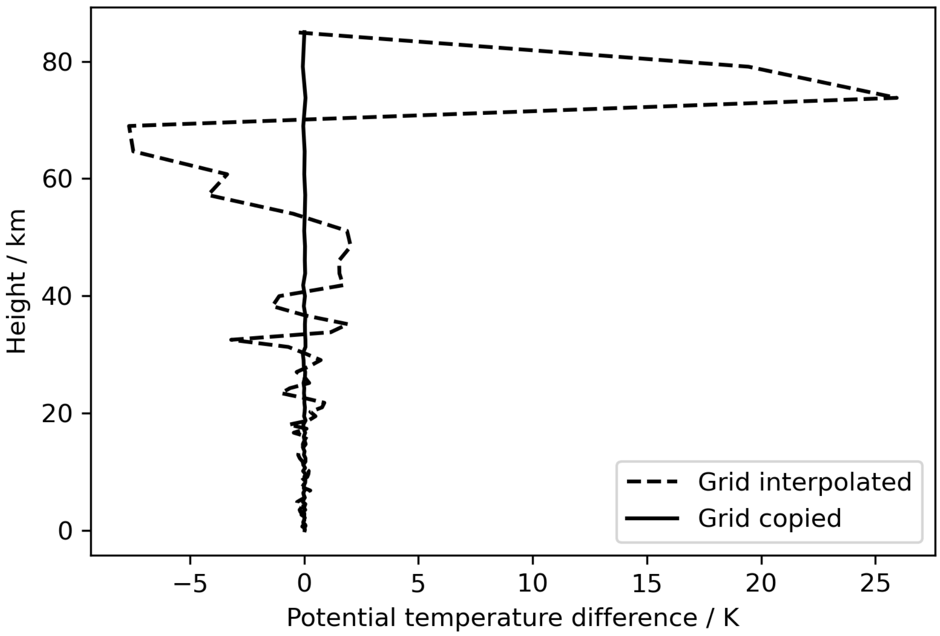

Figure 2.4: Atmospheric profile of the difference between potential temperature supplied as nudging input to the model before and after being read by the model nudging code, for two different cases: one in which the input is interpolated spatially and one in which the input is just copied. All the profiles are for the same single timestep and lat/lon grid point.

This issue was identified to be a result of regridding being applied to the nudging input when read into the nudging code. Since the nudging code is able to accept inputs from different resolutions and sources, nudging input is spatially interpolated onto the model grid. However, even though in this work the nudging input uses the same model grid, it still goes through a spatial interpolation step and this appears to introduce errors. To address this issue a further change was made to the nudging code branch, allowing the code to directly copy the input, skipping the spatial interpolation step. This reduced the difference between the nudging input before and after being read into the nudging code to less than 0.1 K, as illustrated by the 'grid copied' line in figure 2.4.

The small remaining bias between input before and after being read into the model may result from a similar effect being caused by the time interpolation applied to the input. However unlike the spatial grid, the input data used here is not the same time frequency as the model timesteps (20 min), and so temporal interpolation is required. While it would be possible to nudge the model using input data with the same frequency as the model timestep, this would be considerably more expensive computationally, and the remaining difference in figure 2.4 is small regardless.

Aside from these fixes, two further changes were made to the nudging code to enable extra aspects of the nudging to be tested beyond the existing scope of the nudging code. The first was a change to allow the lowermost model atmospheric level to be nudged. In the existing nudging code an exception is raised if the lower nudging ramp is set to less than 1, resulting in it not being possible to nudge the lowermost level. With a small change applied to the branch to adjust the exception, the first model level can be nudged.

The second change was to allow the nudging input time resolution to be varied. The existing nudging code could only take 6 hourly input. However, particularly with linear interpolation onto model timesteps being used here, this provides a relatively poor representation of the diurnal cycle in potential temperature, as discussed earlier in section 1.3.4. This could reduce the effectiveness of nudging in suppressing adjustments in the mechanism denial methodology used here, as well as generating spurious responses to the nudging. Therefore further edits were made via a branch to allow the nudging code to use varying time frequencies of nudging input.

The fixes to the indexing error and regridding were applied to all further simulations presented in this and the following chapter. The changes to allow nudging the first model level and varying the input time resolution were only used for their respective simulations, as detailed in the next section.

2.5 Experimental design for testing nudging setups

As discussed in section 1.3.4, determining appropriate nudging settings is typically done empirically, which is the approach taken in this work. Four sets of simulations were conducted, each set varying a different aspect of the nudging setup: the value of G used; the model levels nudged; whether the lowermost model level is included in nudging; and the time frequency of nudging input. These four sets of simulations and their differences are outlined in Table 2.1 and can be summarised as follows:

- One set of simulations investigated varying G, testing a range between 1/24 (weaker nudging) and 1/1 h^{-1} (stronger nudging).

- Another set of simulations investigated varying the model level range nudged to include the lowermost model levels, with a variation on the lower ramp up of nudging, using the central value of G from the previous set of simulations (1/6 h^{-1}).

- A further simulation set investigated extending nudging to the lowermost model level, using the central G value from the first set of simulations and level range from one of the second set of simulations.

- And a final set of simulations tested varying the nudging input time resolution, using two different values of G and the level range from one of the second set of simulations.

It is worth noting why each set of simulations was conducted and what specifically was being looked for in the context of the methodology presented in section 2.2. For the simulations varying G, the aim was to find a value of G that optimised the balance between better suppression of the u, v and θ adjustments, whilst avoiding introducing significant non-physical responses in the simulations. These simulations were done without nudging the lower levels on the basis that this could cause instabilities in the model (see section 1.3.4). The second set of simulations sought to test whether nudging the lower levels helped with suppressing adjustments in the lower levels, as well as the effect of including a lower nudging ramp. The second set of simulations did not extend nudging to the lowermost model however, since by default the nudging code did not originally permit this. The third test therefore investigated whether nudging the lowermost level too had a significant impact, using the modifications to the nudging code described in section 2.4. All of these three sets of simulations used 6-hourly nudging input. The final set of simulations was then conducted to test whether increasing the time resolution of the input to 3 or 1 hour intervals would improve either the adjustment suppression or reduce the control errors. One of these simulations used a stronger relaxation parameter (G = 1/3 h^{-1}) with 3-hourly input to test whether a higher value of G (and so stronger adjustment suppression) could be used without introducing non-physical responses if the input resolution was higher.

For each of these nudging settings, a control and perturbed simulation were run to be used as a pair for calculating adjustments and associated radiative forcing. SU emissions perturbation was used for the perturbed simulations for all tests because of the higher radiative forcing signal from a PI to PD SU emissions perturbation than for BC. For the first two sets of simulations (testing G and nudging lower model levels), these were conducted with both uv-nudged and uvθ-nudged simulation pairs; for the latter two sets of experiments (testing nudging the lowermost model level and varying the input time resolution), these were only conducted with uvθ-nudged simulation pairs. In all nudged simulations the upper bound of the nudging was model level 83 ($\sim$74 km amsl) with an upper ramp of 2 levels (meaning the full value of G was used up to level 81, $\sim$65 km amsl).

As well as the nudged simulations, a free control simulation and a free SU perturbed simulation were run. As aforementioned, the free control simulation was used to generate the nudging input. The free control and free SU perturbed simulations were also used as a pair for

Table 2.1: Nudging settings applied for each nudging setup tested. G is the relaxation parameter in h^{-1}. The nudging bottom refers to the level at which the nudging ramp begins. The bottom ramp is the number of levels upwards over which the value of G applied is increased linearly to its specified value, reaching the specified value at level = nudging bottom + bottom ramp. Model levels 1, 4, 10 and 14 correspond to heights of 20, 160, 800 and 1493 m above mean sea level. The input time interval is the time resolution of nudging input from the free control simulation used in each simulation. The bold row is the best case nudging setup identified in section 2.7

Setting tested	G / h^{-1}	Nudg. bottom	Bottom ramp	Input time interval / h
G	1/24	10	4	6
	1/12	10	4	6
	1/6	10	4	6
	1/3	10	4	6
	1/1	10	4	6
Lower nudging range	**1/6**	**1**	**1**	**6**
and ramp	1/6	1	4	6
Nudging level 1	1/6	1	0	6
Input time resolution	1/6	1	1	3
	1/6	1	1	1
	1/3	1	1	3

calculating adjustments and associated radiative forcing, to allow comparison with the nudged simulations. These free simulations were conducted for 35 model years each, with the first year removed as spinup and the remaining 34 years used for calculating radiative adjustments. Note that the nudging input generated from the free control was not averaged; the nudging input for the nudged simulations was comprised of unique 6 hourly instantaneous nudging input from the corresponding timestep in the free control simulation.

For the first three sets of nudging test simulations, the nudged simulations were conducted for only 5 model years, with the first year used as spinup and the remaining 4 used for averaging. The shorter integration time for nudged simulations than the free simulations was possible because the pattern of internal variability is similar in the control and perturbed nudged simulations, because both follow the natural variability of the free control simulation which they are nudged to; consequently the difference in the nudged control and perturbed TOA net down radiative fluxes (=radiative forcing) is less variable over time without nudging. To test the validity of using only 4 years averaging, a uvθ-nudged control and perturbed SU experiment were each extended by 5 years to test the effect of averaging over more years. The difference in net down TOA flux determined using either a 4 or 9 year long control simulation was 0.012 W m^{-2}. The difference in ERF determined using either 4 or 9 years of integration in the control and perturbed simulations was <0.001 W m^{-2}. Hence, using more than 4 model years from the nudged simulations makes little difference to the results. Furthermore, with 4 years of integration the standard error of nudged simulation ERFs is already comparable or smaller than that of the free simulations, similar to findings in Kooperman et al. (2012).

The final set of nudging test simulations (varying time resolution of input) were conducted for only one model year. This was due to the extra computational costs and time required (which includes generating the higher time resolution input from a second free control simulation). Thus they are intended to be only limited tests to indicate whether increasing time resolution of input is likely to improve upon the existing results.

2.6 Nudging performance across different nudging settings

Here are presented the results from all of the nudging test simulations described in the previous section. All results presented are determined using multi-annual means integrated over the run length (minus 1 year spinup) of simulations unless otherwise stated. Likewise, unless otherwise stated, all uncertainties given are calculated as two times the standard error of the multiannual mean. Thus the uncertainty only captures the uncertainty due to interannual variability not other forms of uncertainty. References to statistical significance for all results are based on difference with two standard errors, implying statistical significance at the 95 % confidence limit assuming a normal distribution.

The results are divided into different aspects of the nudging performance for the first three sets of experiments (investigating effects of varying G and applying nudging to the lower and lowermost model levels). These include the effectiveness with which u, v and θ adjustments are suppressed; the presence of biases in u, v and θ in the nudged control simulations; biases in the nudged control simulation TOA fluxes; biases in the nudged control simulation cloud fields; the effect of nudging lower model levels on stability and cloud adjustments; and finally the effect of varying the nudging setup on ERF. This is followed by results from the final set of experiments (varying the time resolution of nudging input) which only focused on the effectiveness of suppressing the θ adjustment.

2.6.1 Suppression of adjustments in u, v and θ

First we deal with the performance across the first three sets of simulations (testing varying G and the levels nudged, including whether to nudge the surface or not). To assess the performance of these different nudging setups, firstly their suppression of adjustment in u, v, and θ fields can be compared. Profiles of multi-annual global level mean adjustments for each nudged simulation pair are illustrated in fig. 2.5, for θ, u and v. Compared to θ, the u and v adjustments are harder to characterise with simple level-averaged means, involving positive and negative changes in different areas at each level. Therefore for the horizontal wind fields, instead of a mean adjustment, the RMS of the multi-annual mean adjustment determined over each level is illustrated. This better represents the magnitude of adjustment in horizontal winds at each level and makes it easier to see how well nudging suppresses this adjustment. The adjustments in the free simulation pair are also included so that the θ bias and horizontal wind RMS of the free simulation pair can be compared with that of the nudged simulations to assess how well the nudging suppresses the adjustment. In an ideal case, adjustment of u and v should be zero in all the nudged simulations, and adjustment of θ zero in the uvθ-nudged simulations.

In fig. 2.5c it can be seen that for all of the uvθ-nudged simulations the θ adjustment is negative both in the troposphere and stratosphere. The nudging strength can be seen to directly affect the adjustment suppression in the uvθ-nudged simulations, with the stronger nudged (higher G) simulations having a smaller adjustment (i.e. more suppressed). In the free troposphere, the adjustment varies from approximately 50 % of the free adjustment in the G = 1/24 h^{-1} simulation to less than 5 % of the free adjustment in the G = 1/1 h^{-1} simulation. However, in all of the uvθ-nudged simulations that do not nudge the lower levels, the adjustment suppression is much less effective in the lowermost 1 kilometre, with the adjustment being as large or larger than the free adjustment. By contrast, the two experiments with lower level nudging suppress the adjustment more strongly in the lowermost 1 km, although not entirely. However, in the rest of the troposphere, the lower level nudged simulation pairs have a weaker suppression of the adjustment, with an adjustment similar to the G = 1/24 h^{-1} simulation without lower level nudging. What causes this pattern is uncertain since the nudging in the lower levels simulations is identical to the other simulations above the lowermost kilometre. It is also worth noting that the lower level uvθ-nudged experiments have localised positive and negative θ adjustments in a few regions, particularly over central Africa, that are not present in the simulation pairs not nudging the lower levels (fig. B.1). The variation of the lower ramp between the two simulation pairs with lower level nudging makes little difference to the θ adjustment. The simulation pair with nudging extended to the lowermost level also has little effect, with a reduced adjustment in the lowest three model levels, but slightly increased adjustment at all other levels.

Whilst not apparent in figure 2.5c, in the uvθ-nudged cases the θ adjustment in the lowermost kilometre is stronger over land and the poles (fig. B.1). This likely occurs because while the SSTs are fixed, the land surface temperature and temperature over sea ice is allowed to adjust, which will influence the lowermost model levels in particular. This may explain why nudging, even when applied to the lower model levels, is ineffective at suppressing the adjustment in θ.

In the stratosphere in figure 2.5c, all of the uvθ-nudged simulations have a negative θ adjustment unlike the free adjustment. This is unexpected given that the nudging should suppress the adjustment, not reverse it, and it is unclear what causes this effect.

Looking at figures 2.5a and 2.5b, some similar observations can be made for the u and v adjustments in the uvθ-nudged simulation pairs as were made for θ: the stronger nudging cases have a lower adjustment RMS above 1 km, and the lower level nudged experiment pairs have the lowest RMS in the lowest 1 km (fig. 2.5a and 2.5b). However, unlike for θ adjustment, the lower level nudged simulation pairs have a similar or smaller adjustment RMS in the free troposphere and stratosphere to the G = 1/6 hr^{-1} simulation pair without lower level nudging. Of these two pairs, the pair using a lower ramp has a slightly reduced v wind adjustment RMS and similar u wind adjustment RMS compared to the pair without a lower ramp. Extending the nudging to the surface seems to have little effect. Again it is worth noting there are some localised stronger remaining adjustments over central Africa when applying lower level nudging,

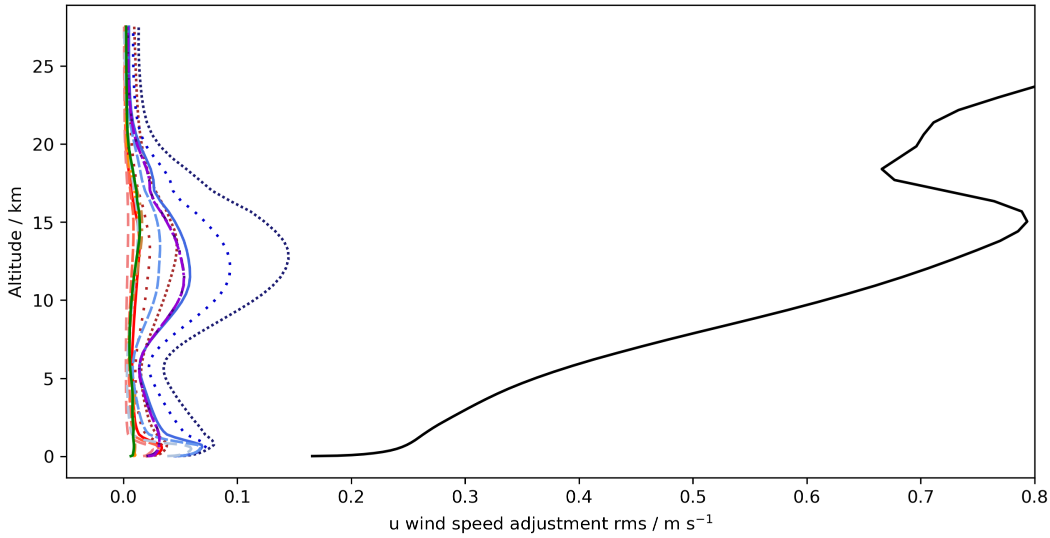

(a) u-wind

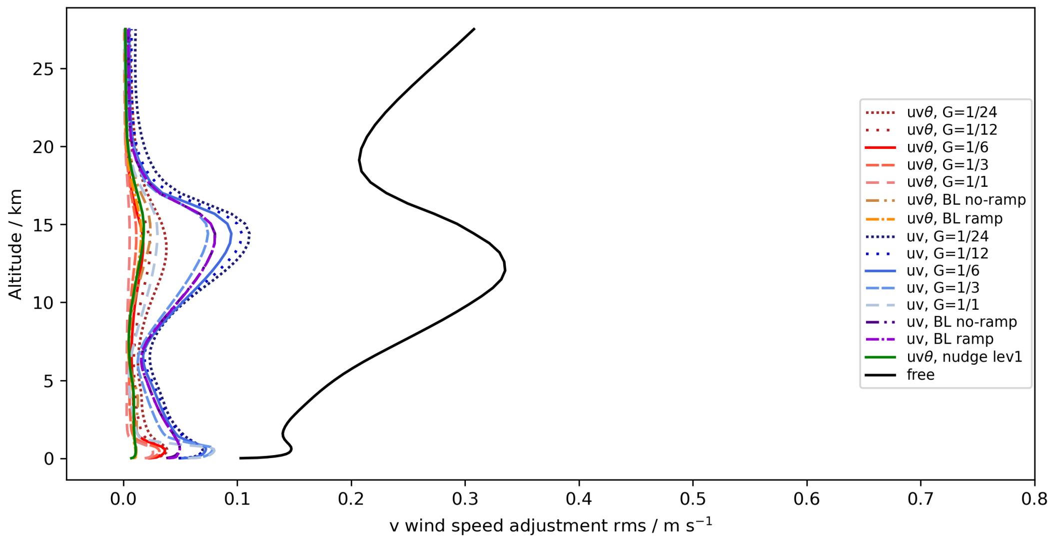

(b) v-wind

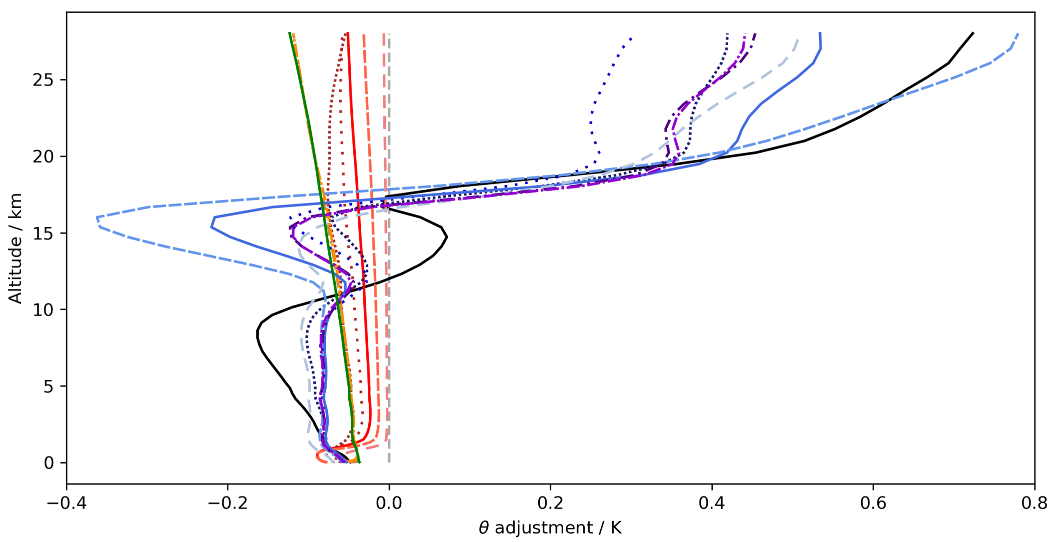

(c) θ

Figure 2.5: Adjustments in (a) eastward wind, (b) northward wind, (c) and potential temperature due to perturbation of SU from 1850 to 2014 for various nudging setups. For horizontal winds the RMS, calculated over each model level, of the global multiannual mean adjustment is given. For potential temperature, the model level mean of the global multiannual mean adjustment is given. The dashed grey line in (c) is the zero line.

but at higher altitudes than for θ (fig. B.2 and B.3).

While the aim is to achieve zero adjustment of θ in the uvθ-nudged simulation pairs, it could be considered ideal for the θ adjustment in the uv-nudged simulation pairs to not be suppressed at all instead (i.e, to equal the θ adjustment in the free simulation pair). If there is a large difference in θ adjustment between the uv-nudged and free simulation pairs it implies the uv-nudging may alter some of the atmospheric temperature response to the aerosol perturbation. This would make the difference between uv- and yvθ-nudged cases a less accurate representation of the atmospheric temperatures mediated adjustments. However, in practice the θ adjustment will be affected by uv-nudging, since the circulation adjustment is likely to affect atmospheric temperature in order to maintain thermal wind balance. In figure 2.5c, it is clear that this is true for all of the uv-nudged simulation pairs, which have considerable deviations in θ adjustment relative to the free adjustment. Specifically, the uv-nudged pairs have about half the adjustment in θ as the free simulation pair in the free troposphere; a negative adjustment in the lower stratosphere where the free adjustment is positive; and as low as half the adjustment in the mid stratosphere, with only the G = 1/3 hr^{-1} case approximating the free adjustment here. None of the setups could be described as particularly better, with those closer to the free adjustment in the mid stratosphere showing stronger negative adjustments in the lower stratosphere. There also is not a monotonic trend over the different values of G, with the G = 1/24 and 1/1 hr^{-1} pairs begin closer to each other than either G = 1/12 or 1/3 hr^{-1}. The lower level nudging has relatively little effect on the lower levels adjustment.

Very similar observations can be made from fig. 2.5a and 2.5b for the u and v adjustment in uv-nudged simulation pairs as for the uvθ-nudged simulation pairs. Firstly, the stronger nudged cases show smaller adjustment RMS; and secondly, lower level nudging reduces the adjustment RMS in the lower levels, and a small amount in the higher levels. Aside from these observations, the adjustment RMS can be seen to be greater for both u and v for each uv-nudged pair compared with their uvθ-nudged equivalents. Similar to the point above concerning θ suppression by uv-nudging, while the temperature adjustment is expected to affect circulation (and so nudging θ might cause suppression of some adjustment in u and v as seen here), ideally the uv-nudged pair would have the same adjustment in u and v as the uv θ-nudged pairs in order to separate out the circulation adjustment entirely. Since the uv-nudging does not suppress the entire horizontal wind adjustment relative to the uvθ-nudged pairs, some amount of the differences between uv and uvθ-nudged simulations could be ascribed to the circulation adjustment.

2.6.2 Biases in u, v, and θ in nudged control simulations

Next, the effectiveness of the nudged control simulations in reproducing the free control can be assessed to give an indication as to whether the nudging is introducing unintended non-physical responses in the model. In principle, since the nudged control experiments have identical conditions to the free control (except for the application of nudging of the u, v and θ fields to the free control), the nudged controls should be close to identical to the free control. Therefore the difference in any field between the free control minus each nudged control experiment - referred to as the 'control error' - should ideally be close to zero. This assumption is only valid for equivalent years in the free and nudged controls due to natural variability. As such, a time average over only years 2 to 5 from the free control (instead of years 2 to 35 as done for adjustments) is used for differencing against fields from the nudged control simulations for determining control errors.

To aid interpretation of the control errors, it should be noted that since the adjustments are calculated from differences in control and perturbed simulations for each nudging setup, the control error is not necessarily important for the separation of circulation and temperature adjustments; if the perturbed simulations included precisely the same errors as the control simulations, then when differenced, the error would have no effect on the adjustment calculated. However in practice, in a simulation with a slightly different atmosphere, one might expect the same perturbation in emissions to generate different adjustments because of the slightly different conditions, biasing the adjustments. While the control error is not a direct measure of this bias itself, a larger control error will indicate that this bias is likely to be greater. Furthermore, a larger control error may indicate that nudging is introducing spurious changes that are not real to the atmosphere, which may affect the representation of adjustment processes.

Figure 2.6 illustrates the control errors for each nudged control simulation for u, v, and θ, again using a multi-annual level mean bias for θ and multi annual mean level RMS for u and v. Considering u$v\theta$-nudged controls first (fig. 2.6c), for θ the stronger nudged controls have a negative error, which is particularly significant in the lower 1 km. Indeed the G = 1/1 hr^{-1} control has a control error over an order of magnitude larger than the free θ adjustment seen in fig. 2.5c for the G = 1/1 hr^{-1} simulation pair. Weaker nudged controls have a positive control error error, which is greater in the stratosphere (fig. 2.6c). The G = 1/6 hr^{-1} control has the lowest error (though it should be noted this arises somewhat from a cancelling of negative error over land and positive over the poles). The lower level nudged controls have a positive θ control error, which is greater than the G = 1/6 hr^{-1} control without lower level nudging. The lower level nudging has little effect on the control error at lower levels. Likewise applying a lower ramp or not to the nudging can be seen to make little difference to the θ control error, and extending nudging to the lowermost level only increases the control error slightly at most levels.

For the u and v control errors (fig. 2.6a and 2.6b), in the lower 1 km the u$v\theta$-nudged controls with G = 1/6, 1/3 and 1/1 hr^{-1} have a large error. The controls with lower level nudging or weaker nudging have a smaller error in the lower 1 km. In the free troposphere and stratosphere the controls with weakest nudging have the lowest error for v, whereas the strongest nudged control has the lowest error for u. Compared to the G = 1/6 hr^{-1} nudging control without lower nudging, the lower level nudged controls have smaller control error in v in the free troposphere and stratosphere and similar error in u. Of the lower level nudging controls, the one without a lower ramp up of nudging has a slightly lower u wind control error at most levels but similar v wind control error to the control with a lower ramp. Extending nudging to the lowest level appears to decrease both the u and v wind control errors relative to the lower level nudged controls by a small amount.

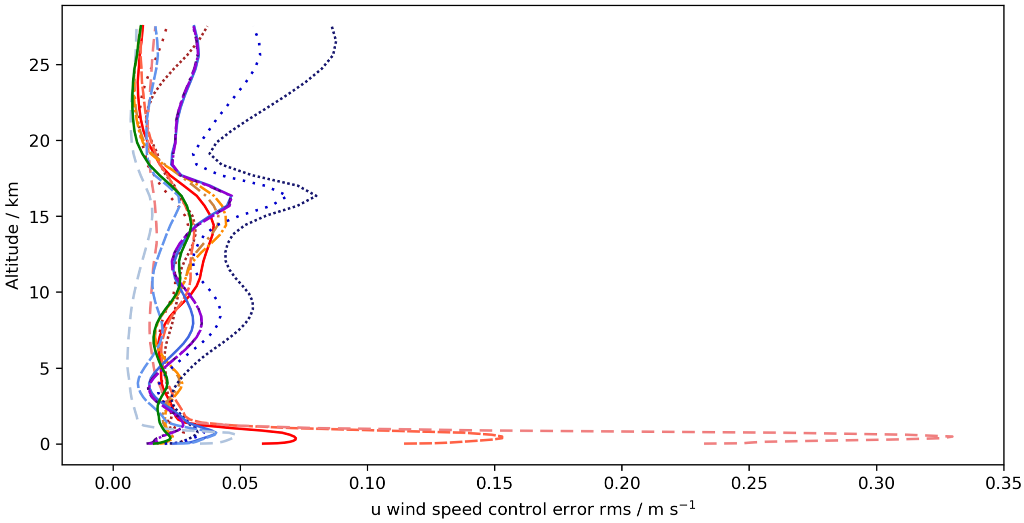

(a) u-wind

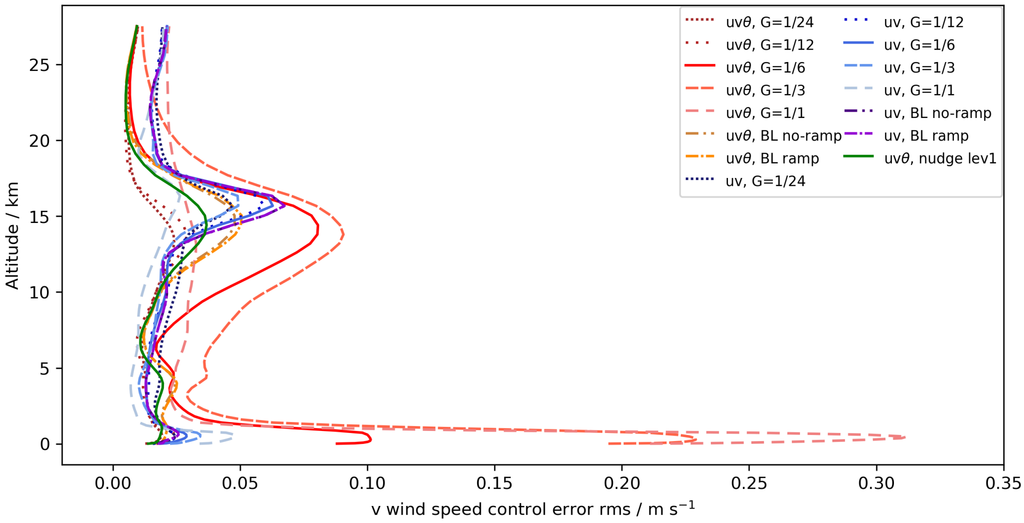

(b) v-wind

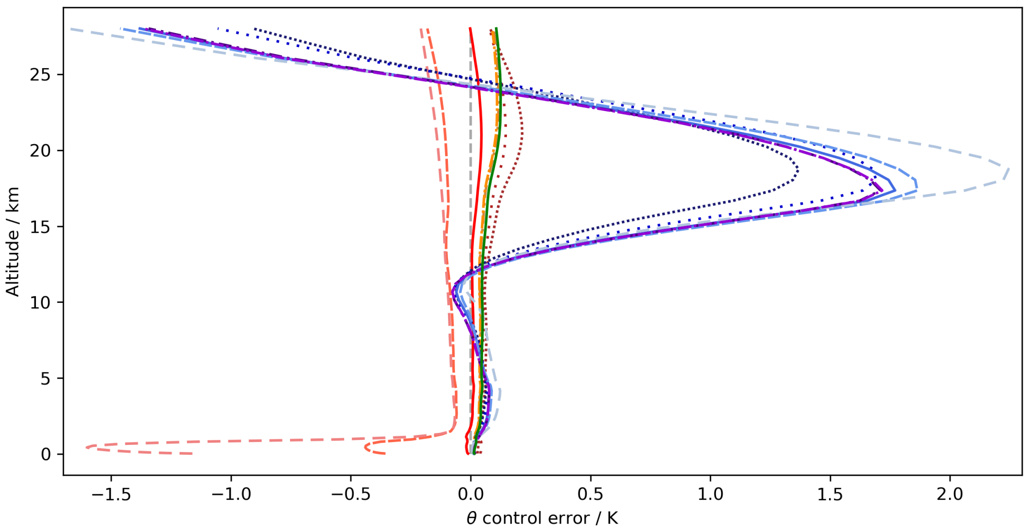

(c) θ

Figure 2.6: Difference in (a) eastward wind, (b) northward wind, (c) and potential temperature fields between the free control minus each nudged control simulation. The difference is determined using only years 2 to 5 from the nudged and free controls. For horizontal winds the RMS, calculated over each model level, of the multiannual mean bias is given. For potential temperature, the model level mean of the multiannual mean bias is given. The difference is determined using only years 2 to 5 of the free control simulation.

The uv-nudged control simulation in figure 2.6c have a strong positive θ control error in the lower stratosphere up to $\sim$24 km, where the errors become strongly negative. This means the uv-nudged controls all have a cooler lower stratosphere and warmer mid-stratosphere than the free control. Unlike for the uvθ-nudged controls, the error is roughly monotonic with G: the largest error is in the strongest nudging case, and smallest error in the weakest nudging case. However, for all the uv-nudged cases this large stratospheric control error indicates that the uv-nudging introduces non-physical responses in the stratosphere. The lower level nudging has little effect on the control error, decreasing the control error by only a small amount. In the troposphere all of the uv-nudged pairs have only a small error, <0.1 K in magnitude (except G = 1/1 hr^{-1}).

In the uv-nudged control simulations the spread in u-wind errors is greater than the uvθ-nudged controls, but there is less spread in the v-wind errors. Generally all of the uv-nudged control simulations have small errors, except for the u-wind control error in the weaker nudged controls. The stronger uv-nudged controls tend to have smaller control errors than the weaker nudged controls, with exceptions at some heights. The addition of lower level nudging does decrease the control error in the lower levels, but has little effect at higher model levels.

2.6.3 Biases in nudged control simulation TOA fluxes

As well as comparing the control error for u, v, and θ, the difference in multiannual mean TOA radiative fluxes between nudged control minus free control can be compared for each nudging control experiment, illustrated in fig. 2.7. There is a large variation of control flux error across the nudged controls. For uvθ-nudged controls, the stronger nudging cases have very large errors of $\sim$-2.5 and $\sim$-7 W m^{-2} for G = 1/3 and 1/1 hr^{-1} respectively. The weaker nudged cases have a smaller but non-negligible error. The G = 1/6 hr^{-1} control has a net TOA flux error within uncertainty of zero (and so within uncertainty of the free control), though with slightly negative SW and positive control LW errors. The two controls with lower level nudging perform best, with net, SW and LW errors all within uncertainty of zero. Since both these two controls are within error of zero, this suggests that the lower nudging ramp has no significant effect. The control with nudging applied at level 1 is also within error of these two controls but with a smaller uncertainty range. This suggests that nudging the lowest level does not significantly change the control flux error but does reduce noise.

Most of the uv-nudged control experiments have a net TOA flux control error not significantly different from zero, but with offsetting negative SW and positive LW errors. The LW and SW errors increase in magnitude with the nudging strength (except G= 1/1 SW error), whereas nudging the lower levels has little effect. Hence, the G = 1/24 hr^{-1} nudged control has the smallest magnitude SW and LW control errors, though is still significantly different from zero.

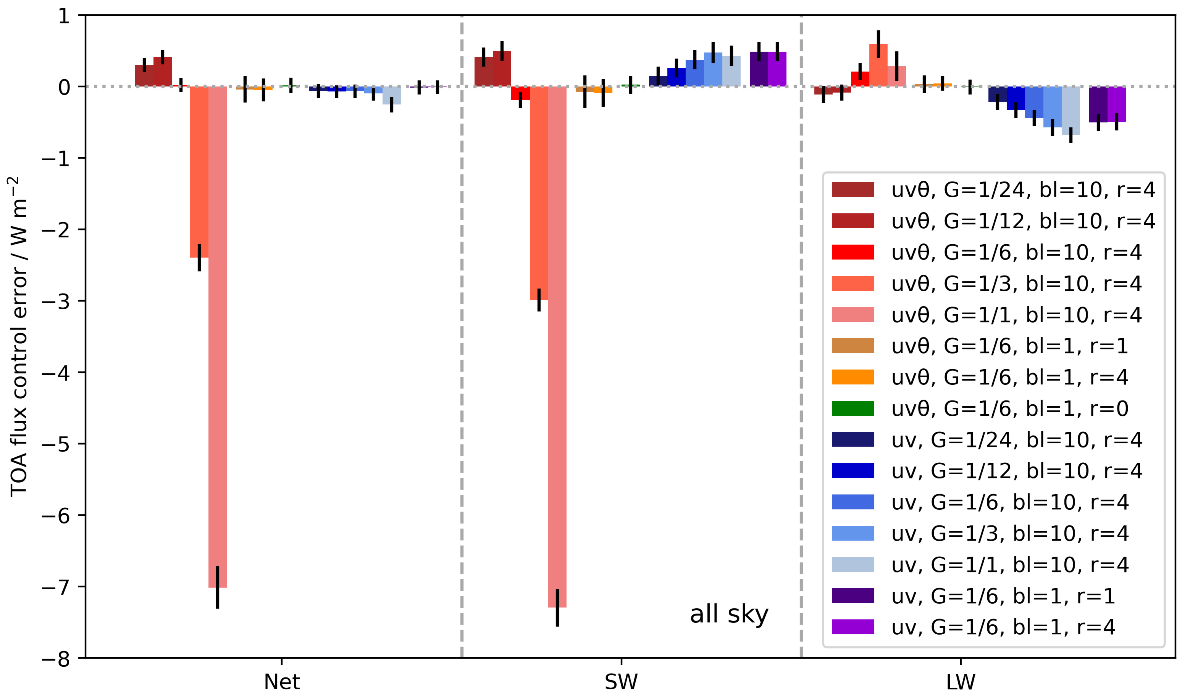

Figure 2.7: Comparison of differences in net, shortwave (SW) and longwave (LW) TOA flux between the free control minus each nudged control. The difference is determined using only years 2 to 5 of the free control simulation. Black lines illustrate uncertainty ranges, calculated as two standard errors over the annual mean flux difference. G is the relaxation parameter applied; bl is the bottom model level nudged; r is the lower nudging ramp which is the number of model levels over which the relaxation parameter is linearly increased from 0 to G. The first 5 bars are from the first set of nudging experiments (varying G) for uvθ-nudging; the next two are from the second set of nudging experiments (nudging the lower model levels) for uvθ-nudging; and the next bar is from the third nudging experiment set (nudging the lowermost model level) for uvθ-nudging; the remaining bars are from the first two sets of experiments but for uv-nudging (see table 2.1 for experiment details).

2.6.4 Biases in nudged control simulation cloud fractions

Since a significant portion of aerosol radiative adjustments are expected to be via clouds, the cloud control errors should also be compared between the different setups. Figure 2.8 illustrates the cloud volume fraction (the three dimensional fraction of a gridbox containing cloud) control error for three uvθ-nudged control simulations and two uv-nudged control simulations. The uvθ-nudged G = 1/1 h^{-1} case has a much stronger control error in all parts of the atmosphere than the other cases, exceeding 0.15 at some grid points (fig. 2.8c; note the different scale). The large error in cloud fraction likely explains the large control TOA flux error for this experiment seen in figure 2.7. Nudging the lower levels does not improve or worsen the control error with uvθ-nudging (fig. 2.8b compared with 2.8d) generally, though does change the pattern of the error in the upper troposphere between 20 ° north and south. The control errors for uv-nudged simulations are similar, though with larger errors in the high cloud but smaller in the low cloud (fig. 2.8e). For the uv-nudged simulations the lower level nudging has a negligible effect on the cloud fraction control error, with figure 2.8fe and f being almost identical.

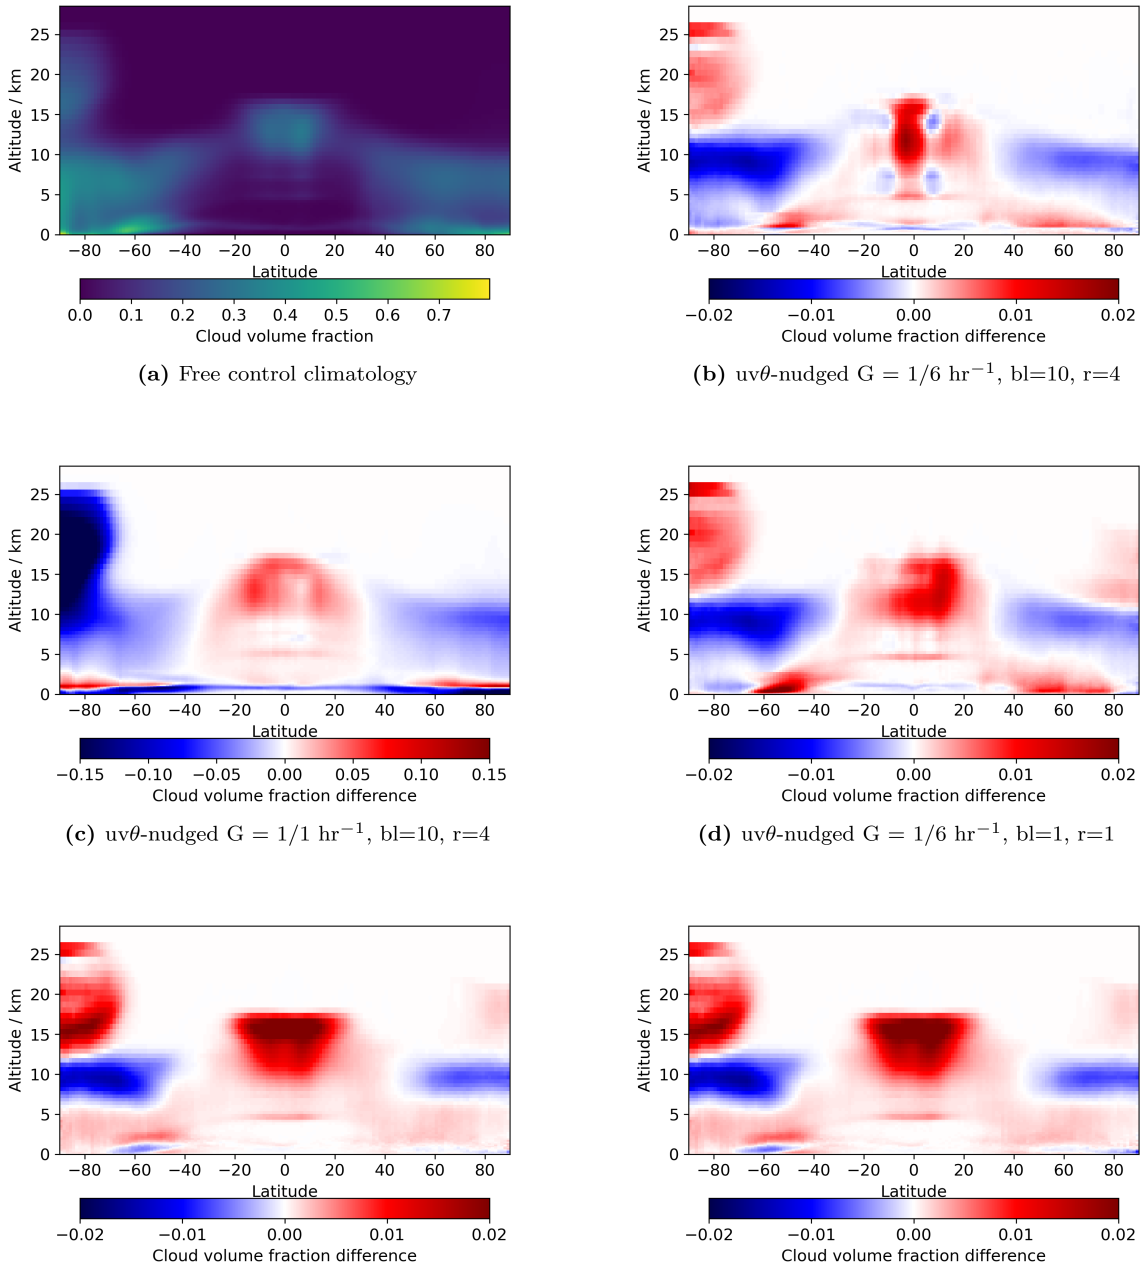

(a) Free control climatology

(b) uvθ-nudged G = 1/6 hr^{-1}, bl=10, r=4

(c) uvθ-nudged G = 1/1 hr^{-1}, bl=10, r=4

(d) uvθ-nudged G = 1/6 hr^{-1}, bl=1, r=1

(e) uv-nudged G = 1/6 hr^{-1}, bl=10, r=4

(f) uv-nudged G = 1/6 hr^{-1} , bl=1, r=1

Figure 2.8: (a) Climatology of zonal mean cloud volume fraction from free control, and (b-f) differences in multiannual zonal mean cloud volume fraction between selected nudged controls minus the free control. Note the larger scale for (c). The differences are determined using only years 2 to 5 of the free control simulation.

2.6.5 Nudging effects on lower tropospheric stability and cloud fraction adjustment

One more comparison found to be particularly relevant for the effect of applying lower level nudging is the adjustment in cloud fields and atmospheric stability. Figure 2.9 illustrates the cloud adjustments in the G = 1/6 h^{-1} uvθ-nudged simulation pair without lower level nudging and in one of the lower level nudged simulation pairs, and the adjustment in lower tropospheric stability (LTS) for both simulation pairs, due to perturbation of SU from 1850 to 2014. LTS is approximated as $\theta_{20} - \theta_1$, where the subscript denotes the model level, with level 20 at approximately 3 km above mean sea level (amsl) and level 1 the lowermost atmospheric level. While this is a crude metric for stability it has been shown to have use for understanding low cloud changes (Andrews et al., 2021). Without lower level nudging, there is a strong positive cloud fraction adjustment in the lowest 1 km in response to SU perturbation (fig. 2.9a). Since the θ adjustment is poorly suppressed in the lower 1 km when not nudging the lower levels (fig. 2.5c), the unsuppressed cooling of these levels relative to the more strongly suppressed cooling above increases the lower tropospheric stability below the height that nudging is applied (fig. 2.9c). Hence, this cloud fraction increase is a spurious effect arising from the nudging. With lower level nudging applied it can be seen that the stability adjustment is much smaller (fig. 2.9d). Accordingly, while there remains a cloud increase in the lowest 1 km in the lower level nudged simulation pair, it is weaker (fig. 2.9b), suggesting that nudging the lower model levels reduces this spurious increase in cloud fraction.

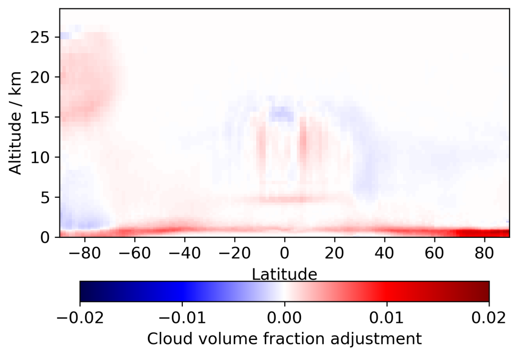

(a) Cloud adjustment without lower level nudging

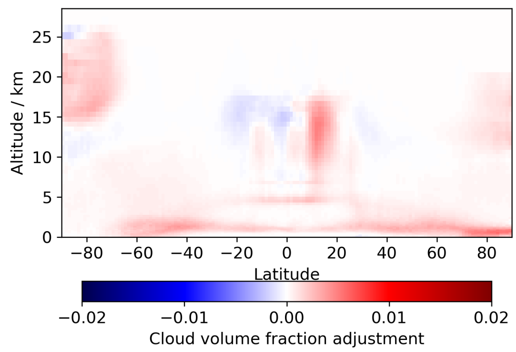

(b) Cloud adjustment with lower level nudging

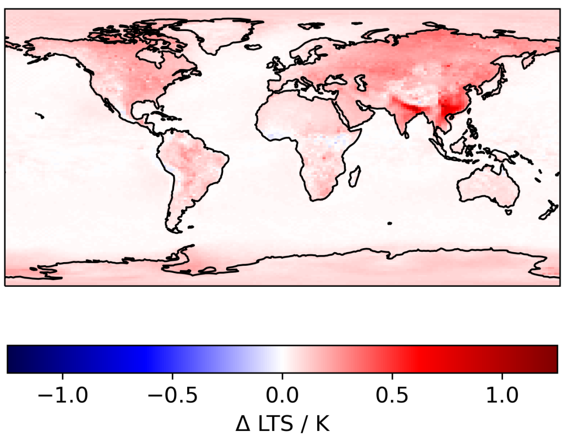

(c) LTS adjustment without lower level nudging

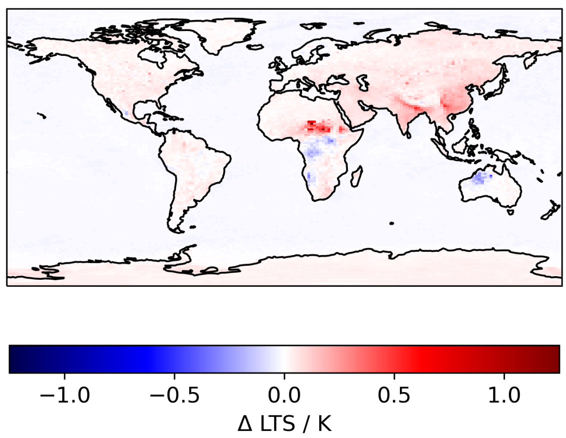

(d) LTS adjustment with lower level nudging

Figure 2.9: Effect of nudging height on cloud adjustment via stability adjustment. Plots illustrate the adjustment in multiannual zonal mean cloud volume fraction (a, b) and adjustment in multiannual mean lower tropospheric stability (c, d) between the uvθ-nudged control (2014 SU emissions) simulation minus the uvθ-nudged SU perturbed simulation (1850 SU emissions), both with G = 1/6 h^{-1}, when nudging is applied either from (a, c) model level 11 ($\sim$1 km above mean sea level) or (b, d) model level 2 ($\sim$50 m above mean sea level) upwards. Lower tropospheric stability is approximated as $\theta_{20} - \theta_1$, where the subscript denotes the model level, with level 20 at approximately 3 km above mean sea level and level 1 the lowermost atmospheric level.

2.6.6 ERF under varying nudging setups

Having compared the performance of the different setups, the effects on the overall ERFs determined can be compared. Figures 2.10 and 2.11 illustrate the ERF due to PI (1850) to PD (2014) SU emission perturbation calculated from the free simulation pair and each of the nudged simulation pairs, for all sky and the CRE. For the free simulation pair, the all sky ERF is negative as expected for SU aerosol forcing. The negative SW forcing results from a negative CRE and DRE (not shown) in approximately equal measure. The positive LW forcing results mainly from the CRE response, and to a smaller extent residual clear-clean sky responses (not shown).

Across the different uvθ-nudged simulation pairs the ERFs determined vary significantly, with a range of ~0.4 W m^{-2} between G = 1/1 and 1/6 hr^{-1}. Most have a more negative ERF than the free simulation pair (except for the strongest nudging case) with the ERF having a large dependence on G. The ERF does not vary monotonically with G, with the middle case (G=1/6 h^{-1}) having the most negative ERF. The application of lower level nudging causes a significant but smaller change in ERF than the variation of G, with the ERFs of the two lower level nudged simulation pairs being less negative than the ERF of the simulation pair without lower level nudging but the same G. The two lower level nudged simulations pairs are within uncertainty, suggesting no significant effect of the lower nudging ramp. However, in contrast to the TOA control flux errors seen in figure 2.7, nudging the lowermost level increases the uncertainty of the radiative forcing calculated. Figure 2.11 illustrates that the majority of the variation in ERF across the different nudging setups arises from a varying CRE response.

By contrast, for the uv-nudged experiment pairs the ERFs are fairly invariant with G and whether lower level nudging is applied or not. All of the uv-nudged experiments have a more negative net ERF than the free experiment pair, with a slightly more negative SW and less positive LW forcing in the uv-nudged experiments.

2.6.7 Effect of varying nudging input time resolution

One of the main deficiencies of all the nudging setups tested is the weak suppression of θ adjustment with uvθ-nudging. The simulations testing higher time resolution input were primarily conducted to see if this could be improved upon, and so the only aspects of the higher nudging input time resolution simulation pairs presented here are the θ adjustment suppression and control error. Since the simulations were run for one model year only, the θ adjustment and control error are plotted as monthly means for 4 model levels to give an approximate indication of the θ adjustment and control error (rather than using a time mean full vertical profile as before) with the limited dataset. These results are illustrated in figure 2.12.

From figure 2.12a it can be seen that for all of the simulation pairs the θ adjustment is similar for levels 1 (~50 m), 11 (~1 km) and 26 (~5 km), and larger for level 55 (~20 km). This is consistent with the pattern in figure 2.5c for all other simulation pairs. Increasing the time resolution alone appears to offer little improvement of the adjustment suppression, with

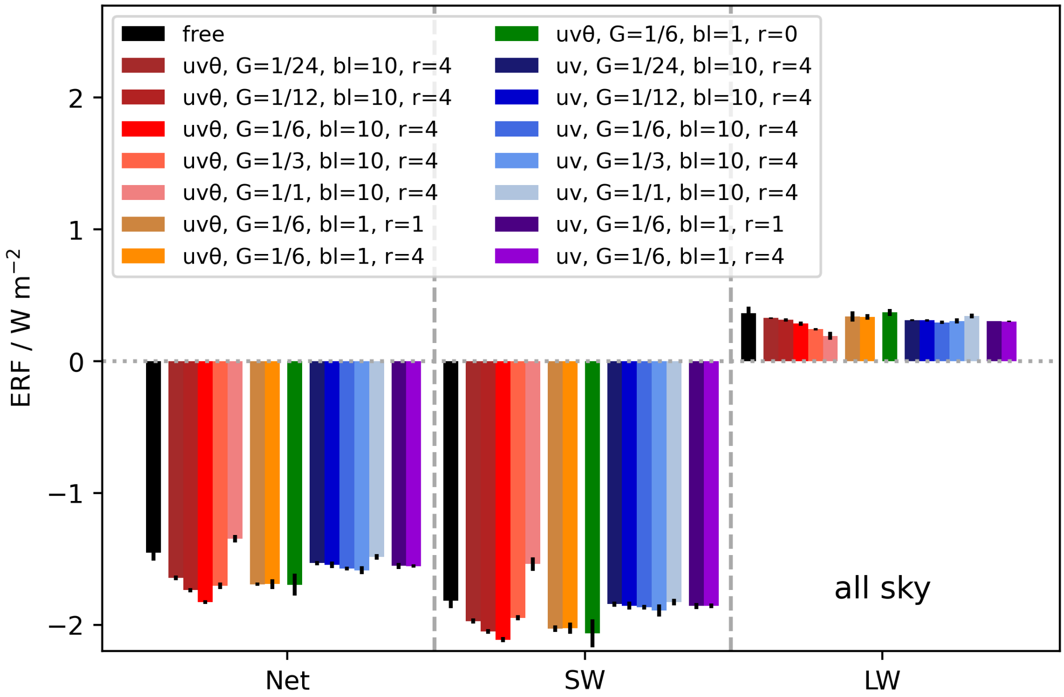

Figure 2.10: Comparison of all-sky ERF due to 1850 to 2014 SU perturbation determined from each pair of control and perturbed simulations. Black lines illustrate uncertainty ranges, calculated as two standard errors over the annual mean ERF. G is the relaxation parameter applied; bl is the bottom model level nudged; r is the lower nudging ramp which is the number of model levels over which the relaxation parameter is linearly increased from 0 to G. The first bar is from the free experiment pair; the next 5 bars are from the first set of nudging experiments (varying G) for uvθ-nudging; the next two are from the second set of nudging experiments (nudging the lower model levels) for uvθ-nudging; and the next bar is from the third nudging experiment set (nudging the lowermost model level) for uvθ-nudging; the remaining bars are from the first two sets of experiments but for uv-nudging (see table 2.1 for experiment details)

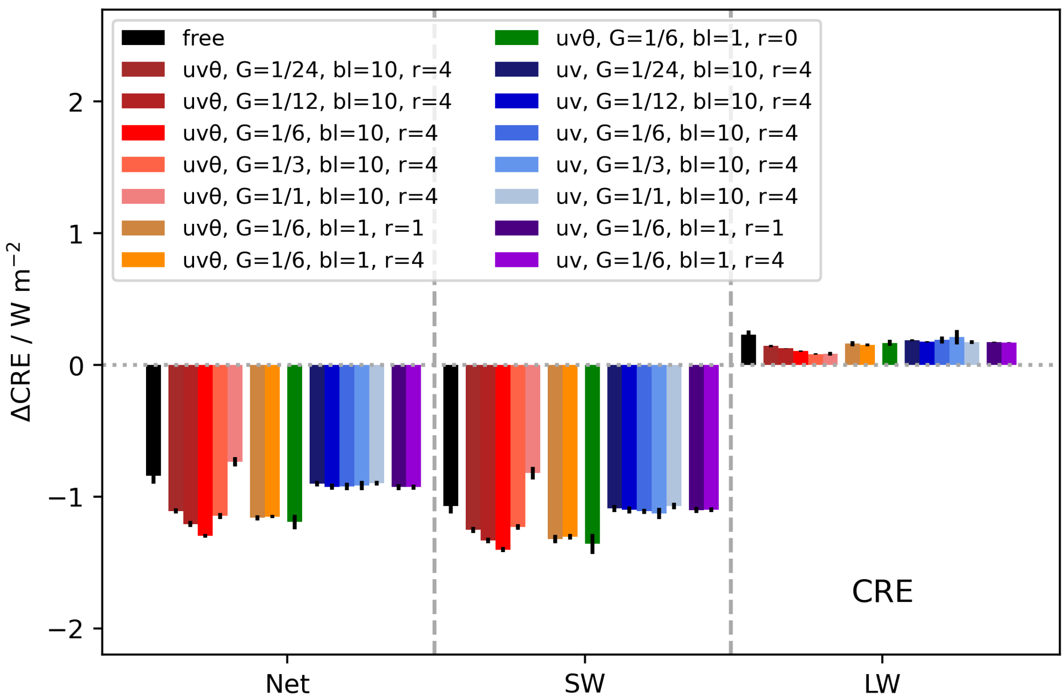

Figure 2.11: As for fig. 2.10 but for CRE.

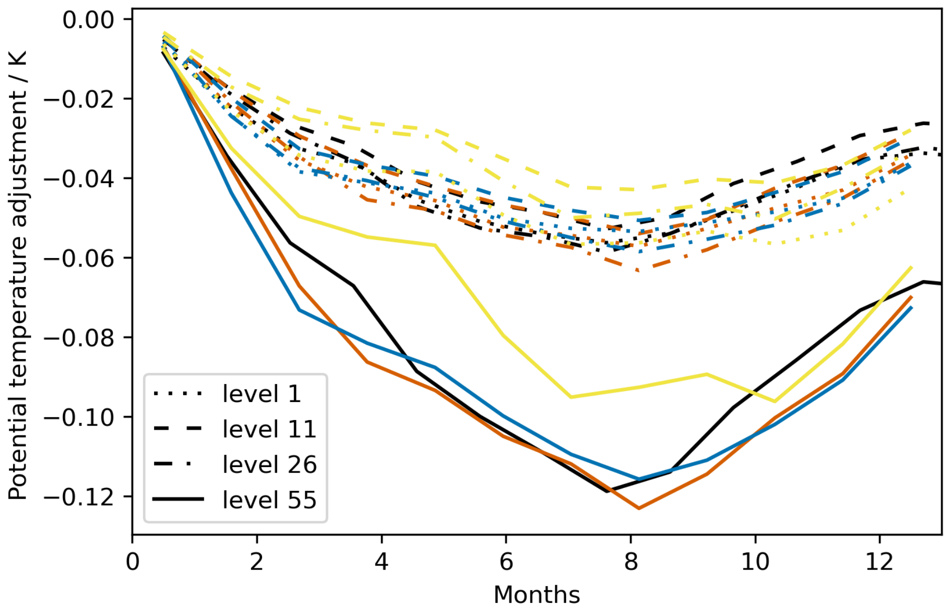

(a) Adjustment

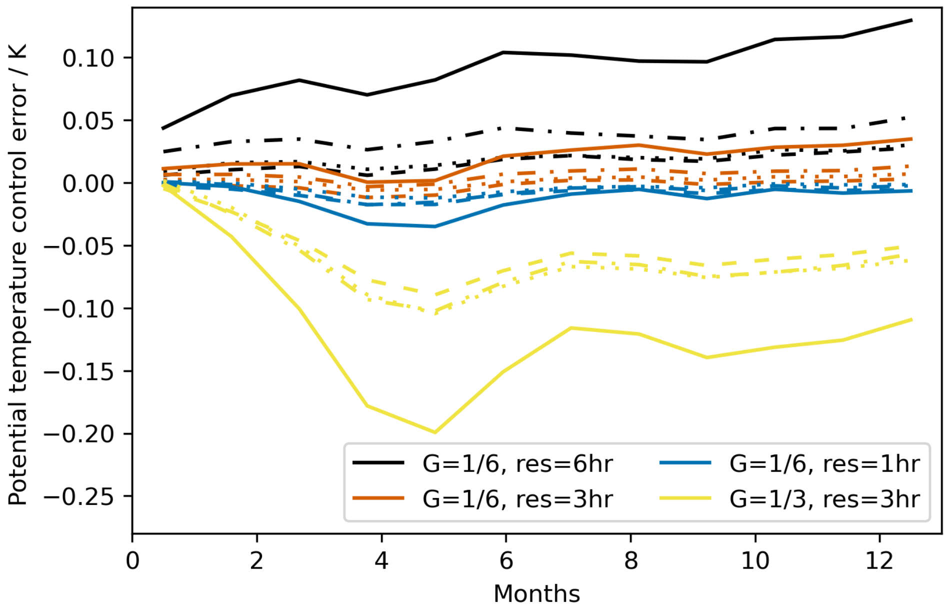

(b) Control error

Figure 2.12: Global monthly mean (a) adjustment and (b) control error in θ over a selection of model levels, using four different nudging setups. The adjustment is calculated as the difference between θ fields due to perturbation of SU from 1850 to 2014 emissions between the control and SU perturbed UKESM simulations. Control error is calculated as the difference in θ fields between the free control and nudging control UKESM suite. The four nudging setups used are the first experiment testing lower nudging range and ramp (G=1/6, resolution (res) = 6hr), and the three experiments testing input time resolution, as listed in table 2.1. Level heights are: 1 = ∼50 m; 11 = ∼1 km; 26 = ∼5 km; 55 = ∼20 km.

both the 3 hourly input and 1 hourly input simulation pairs with G = 1/6 h^{-1} having a similar adjustment magnitude to the 6 hourly input simulation pair across all levels. The simulation pair with 3 hourly input and G = 1/3 h^{-1} does however have a smaller θ adjustment throughout most of the 12 month period than the other simulation pairs, similar to the G = 1/3 h^{-1} simulation pair with 6 hourly input in the first set of nudging experiments (figure 2.5c).

Turning to the control error, from figure 2.12b it can be seen that there is considerably more variation in control error with nudging input time resolution. The control simulation with 3 hourly nudging input and G = 1/6 h^{-1} has a smaller magnitude control error than the 6 hourly nudging input control simulation, with a reduction of more than 60 % on levels 1, 26 and 55, and smaller reduction on level 11, over most of the 12 month period. Despite the clear decrease in control error going from 6 hourly to 3 hourly input, there is not a clear decrease in error going from 3 hourly to 1 hourly input, with the simulation using 1 hourly nudging input having similar magnitude control errors, although with control errors generally negative rather than positive. By contrast to the other simulations, the control simulation with 3 hourly nudging input and G = 1/3 h^{-1} has a larger magnitude control error. For levels 26 and 55 this control error is similar to that of the G = 1/3 h^{-1} simulation pair with 6 hourly input illustrated in 2.6c, but lower for levels 1 and 11 due to the lower level nudging employed in the higher time resolution simulations.

2.7 Discussion and identifying an optimum nudging setup

Across the different variables investigated it is clear that both the choice of G and whether nudging is applied to lower model levels have a significant influence on the uvθ-nudged simulations, particularly evidenced by the effect on the SU ERF calculated from the different simulation pairs. Clearly for the methodology outlined in section 2.2 to be effective, careful consideration must be given to the nudging settings for uvθ-nudged simulation pairs. By contrast, for uv-nudged simulation pairs the G value and levels nudged were less important, making some difference to the control error and adjustment suppression, but relatively little difference to the overall SU ERF calculated.

Because of the sensitivity of the uvθ-nudged simulations to the nudging settings used, a single value for each setting is identified as the optimum value of those tested. These settings are used for the application of the methodology to investigate SU and BC aerosol radiative adjustments in the following chapter. Due to the smaller variation across the uv-nudged simulations when using different settings, this comparison is focused more on the optimum nudging settings for the uvθ-nudged simulations. The same settings are however, used for both the uv- and uvθ-nudged simulations in the following chapter for consistency.

Firstly, the performance of different values of G across the first set of nudged simulation pairs can be summarised as follows:

- Weaker nudging cases are less effective at suppressing u, v, and θ adjustments; have low control errors for u and v; and have middling control errors for theta and TOA fluxes;
- Stronger nudging cases are most effective at suppressing adjustments; but generate larger control errors for u, v, and θ, especially in the levels below which nudging was not applied (lowermost 1 km), and large control cloud fraction and TOA flux errors;
- The middle nudging case is middling at adjustment suppression; performs best on control theta and TOA flux errors; and has middling control errors for u and v.

Because of the significant control errors, the stronger nudged cases (G = 1/3 and 1/1 h^{-1}) can be dismissed. With the cloud fields and TOA fluxes so different to the free control, the atmosphere in these control simulations is unlikely to be representative of the real PD atmosphere, and there will be differences in the representation of adjustments compared with the free simulation due to the differing background atmosphere. From the remaining three values of G tested, the 1/6 h^{-1} case performs better at suppressing adjustment of u, v, and θ and is comparable across the control errors to the weaker nudged cases.

Secondly, the impact of including nudging in the lower model levels can likewise be summarised as follows:

- Lower level nudging improves suppression of the u, v, and θ adjustments in the lower levels;
- In the free troposphere and stratosphere, lower level nudging nudging negatively impacts suppressing the θ adjustment but improves suppression of the u and v adjustment;
- Lower level nudging reduces control error for u, v and TOA flux, but increases the control error for θ;
- Lower level nudging also introduces some spurious control errors and spurious localised areas of adjustment in u, v, and θ, particularly over central Africa;
- The more consistent suppression of θ across the atmosphere when applying nudging to the lower levels reduces erroneous stability changes and corresponding cloud adjustment.
- The lower nudging ramp had little effect.

The fifth point is potentially a very significant one given the importance of clouds in aerosol radiative adjustments. It could be the reason for the less negative ERF when nudging the lower levels relative to the simulation pair with the same G value (1/6 h^{-1}) without lower levels nudged (fig. 2.10), since nudging the lower levels avoids the artificial increase in low cloud fraction. Based particularly on this point, as well as the improvement in adjustment suppression in the lower model levels, using lower level nudging appears to be better for process denial experiments involving θ nudging. This is despite the poorer suppression of θ adjustment in the free troposphere and stratosphere, and areas of spurious localised adjustment and control error, when nudging the lower model levels. Including the lower nudging ramp or not appears to make little difference.

These results demonstrate that, when nudging θ or other temperature-like quantities, nudging to lower levels is important to avoid introducing artificial stability changes at the level below which nudging is applied from. For the nudging methodology used here this is important: perturbed simulations are expected to have systematic differences in temperature to the control simulation that they are nudged towards due to the perturbation of aerosol, and so artificial stability changes should be expected if no nudging the lower levels. With cloud being of particular interest in the adjustments arising from aerosol forcing, avoiding artificial cloud changes is particularly important. However, this effect could also apply to studies nudging temperature-like quantities more generally: for example, if a model with a known bias in temperature is nudged towards a reanalysis one might expect similar artificial changes in stability to occur. While differences in orography prevent reliably nudging the lower models when nudging to reanalysis datasets to avoid this issue, other studies should consider this effect may affect results, particularly where stability or cloud responses are of particular importance to the study. However it should be noted that in the uv-nudged simulations the lower level nudging made less difference, with u and v adjustments and control errors in the uv-nudged simulations without lower level nudging being similar over the lower levels to the nudged levels, particularly for lower values of G, similar to the findings of Sun et al. (2019).

Thirdly, including nudging of the lowermost model level or not made little difference. While nudging the lowermost level does decrease u and v control errors by a small amount, the main differences are the improved suppression of θ adjustment at the lower few model levels, worse suppression of θ adjustment at all higher levels, and a larger uncertainty in the ERF when nudging the lowermost level. While the improved θ adjustment suppression at the surface could help ensure temperature mediated adjustments are suppressed, the increased ERF uncertainty would make it harder to draw significant conclusions from the results. It is also not obvious what causes the ERF uncertainty. While nudging this extra level seems less important than the value of G or whether to nudge the lower 10 levels, it is concluded here that not nudging this level is more effective.

Finally, increasing the time resolution of the nudging input does appear to decrease the θ control error but does not seem to decrease the θ adjustment. Effectively this suggests that with higher time resolution input the nudging is able to better reduce the bias when nudging towards the same atmospheric state, in agreement with the findings of Zhang et al. (2022) and Davis et al. (2022). However, when nudging towards an atmospheric state with different aerosol emissions the increased time resolution of nudging input does not reduce the bias. This could be because in the perturbed simulations the nudging is trying to overcome systematic differences in θ between the nudging input and the 'natural' state of the simulations, meaning small differences in the 'target' (nudging input) of the nudging make little difference. Conversely in the control simulations the nudging is only working to constrain differences in internal variability, meaning much of the control error will arise from smaller differences that the extra precision in the target can help reduce. Additionally, the results of the simulation pair with 3 hourly input and G = 1/3 h^{-1} indicate that increasing the time resolution of the input does not allow stronger nudging to be used without loss to the control error. Indeed the control error and adjustment appeared similar to the simulation pair using the same value of G but with 6 hourly input aside from the differences due to lack of lower level nudging in that simulation pair. This may suggest the input resolution is a less important setting to consider than the value of G and the levels nudged.

In terms of an optimum setting for input time resolution, the results indicate that control errors can be reduced by a large percentage just by increasing the resolution to 3 hourly from

6 hourly. Increasing resolution to 1 hourly makes less difference. This is similar to the findings of Zhang et al. (2022) who found only marginal reductions in the bias between nudging input and a nudged simulation when going from 3 to 1 hourly input. However, the reduced control error of the 3 hourly input simulation must be measured against the extra data storage required, which was doubled relative to 6 hourly nudging input; and the extra computation time, which only increased slightly for the actual simulation but did require longer pre-processing time also. While not a rigorous investigation, from the results presented here it was determined here that, without any decrease in the adjustment as well as control error, it was not worth increasing the time resolution to 3 hourly for the final experiments used in the following chapter. However, it should be stressed the investigation into varying time resolution here was limited to only 1 year simulations and investigating θ adjustments and control error only; a more detailed investigation of the effects of higher resolution input on this methodology could therefore be valuable.

Following from the analysis of each investigated aspect of the nudging setup the optimum settings were determined as G = 1/6 h^{-1}, with nudging applied to the lower levels, without including nudging of the lowermost level or a lower nudging ramp, and using 6 hourly time resolution of input. This corresponds to one of the simulation setups in the second set of test simulations. While the decision to not use a lower nudging ramp was somewhat arbitrary, because having a lower ramp does not appear to improve the simulation at all it could be considered better to not include it by default. This is the nudging setup used in chapter 3 to apply the nudging methodology outlined in this chapter to aerosol perturbation and its radiative adjustments.

2.8 Conclusions

This chapter presents an assessment of the effects of varying a number of nudging settings, specifically in the context of using nudging as a mechanism denial tool for investigating radiative adjustments, using SU as an example. The results suggest that the optimum value of G for this methodology using UKESM1 was found to be 1/6 h^{-1}. This is at the high end of values typically used in studies of radiative forcing that apply nudging in general circulation models. This stronger nudging may be necessary with this methodology because, unlike most nudging studies in which nudging is used to 'course correct' a simulation by nudging to a reanalysis, here nudging is used to suppress an intentional difference between the nudged simulation and nudging input, namely aerosol adjustments.

Nudging the lower model levels was found to be important due to the risk of causing artificial stability changes at the lowest height nudging is applied at, which can cause an artificial cloud response. This effect could occur in other studies if the nudging input used is expected to have a systematic difference in temperature to the simulation being nudged, such as if a model is known to have a large temperature bias relative to a reanalysis dataset being nudged to. Studies using nudging should therefore check whether this artificial stability change occurs and consider whether it affects cloud responses. This could be done by testing the application of nudging from different levels in the lower troposphere. However, this may only be possible when nudging to output from a model using the same orography rather than a reanalysis.

From these results, using G=1/6 h^{-1} with lower level nudging is identified as the optimal nudging setup for applying the mechanism denial methodology to SU and BC radiative adjustments with UKESM1. However, it is clear there are limitations even in this optimal case. In particular, across a number of different nudging setups large differences in cloud in the nudged control simulations versus the free control were identified, which could affect the representation of cloud adjustments in the nudged simulations pairs. Across all of the uvθ-nudged simulation pairs there was a weak suppression of the θ adjustment, achieving only ~60% reduction through much of troposphere but less over the lower levels, and a θ adjustment in the stratosphere of opposite sign to that in the free simulation pair. While using higher time resolution nudging input might be useful, a limited test suggested that using higher time resolution nudging input does not improve suppression of θ adjustment. However this was only a limited test and using higher time resolution did reduce the bias in θ between the nudged control simulations and nudging input, as found in other studies. Further work is therefore required to fully investigate the use of higher resolution nudging input.

2.9 Recommendations for nudging settings for different purposes

As well as the discussion given above, a list of approximate recommendations is given for different applications of nudging. These recommendations draw from both the results found in this chapter as well as the literature review given in section 1.3.4. As such these recommendations are most applicable to the use of nudging in UKESM1 (and the UM more broadly). Some recommendations, such as which variables to nudge and the nudging input source are likely to be fairly applicable to nudging in other models too, while others such as G and input time resolution may vary considerably between models. The recommendations are outlined in table 2.2, and explained in more detail below.

When thinking about the most appropriate settings to use, the aims in terms of both adjustments and control errors can be considered. When constraining a single simulation to specific meteorology, the aim should generally be to minimise the control error, this simply being the difference between the simulation and the input. When reducing the integration length for calculations of ERF or radiative adjustments, the aim is to constrain the natural variability. However, this is not intended to constrain the simulation to any particular values, only to increase consistency of variability between a pair of simulations. As such, the control error should be minimised but the adjustment ideally would be unaffected, since suppressing it can affect the ERF. Some suppression of adjustments will be unavoidable, however. By contrast, when seeking to deny specific mechanisms to isolate their effects the primary aim is to suppress the adjustments, whilst ideally minimising the control error too. It is also worth noting that, in the first case when constraining to specific meteorology, when constraining a pair of simulations

Table 2.2: A list of recommended settings for nudging for different applications of nudging. The applications considered are specifically: nudging to constrain a model to a specific (usually observed/reanalysis) meteorology; nudging to reduce necessary integration time by constraining natural variability to be the same in a pair of simulations; and nudging to suppress the circulation and/or atmospheric temperature response to isolate the effects of these responses. These recommendations are primarily made for using nudging in UKESM1, but may apply more broadly to other models in some cases. It should also be noted that some of these settings will depend on others: for example G may depend on the input time resolution, and whether θ is required to be nudged for the intended purpose. 'BL' = boundary layer

Nudging setting	Constraining to specific meteorology	Reducing integration time	Mechanism denial
Variables nudged	u,v; θ if necessary	u,v	u,v for circulation; θ for atmospheric temperature
Nudging input source	Source of specific meteorology (usually reanalysis)	Same model output (e.g. PiClim for CMIP)	Same model output
G	Depends on exact aim	Weaker: 1/24 h^{-1} (or possibly weaker)	Stronger: 1/6 (or 1/12 h^{-1} if just uv-nudging)
Nudged level range	Depends on exact aim	Above BL, unless nudging θ	Above BL, unless nudging θ
Input time resolution	3 hourly	3 hourly	3 hourly is better but 6 hourly may be sufficient
Is nudging suitable?	Yes	Depends on outputs of interest: just aci or WMGHGs - yes; overall ERF - maybe	Okay for circulation adjustment; less effective for temperature adjustment

with some perturbation between them the recommendations will be similar to constraining a single simulation, but may depend on the exact purpose. With this thought process in mind each setting can be considered in turn below:

Variables nudged

For constraining to a specific meteorology, the variables nudged will be largely determined by the specifics of the intended purpose. In most case just nudging the horizontal winds may be sufficient. However, if processes strongly dependent on temperature, such as atmospheric chemistry, need to be matched between the simulation and the input then it may be necessary to nudge θ as well (Davis et al., 2022). For reducing integration time it is best to only nudge the horizontal winds, as nudging θ can affect ERF too much. For mechanism denial the variables nudged obviously depend on what processes are intended to be suppressed.

Nudging input source

When constraining to specific meteorology the source will be whatever meteorology is intended to be replicated by the simulation, usually a reanalysis dataset. For reducing integration time where the aim is to reduce control error but leave adjustments as unaffected as possible, nudging to output from the same model is likely to be better than nudging to reanalysis since there will be fewer differences between the simulation and the nudging input. One obstacle here is that a suitable model simulation may not already be available to use for nudging input, whereas reanalysis data is readily available. However, in many cases a suitable simulation form the same model could be found. For Climate Model Intercomparison Project (CMIP) experiments for example, a PI fSST control (PiClim) experiment is likely to be a reasonable choice of nudging input for any timeslice fSST experiments. The same recommendation applies to mechanism denial experiments.

G

For constraining to specific meteorology the value of G to use depends on the exact purpose. If only approximate conditions are required a lower value of G (1/12 or 1/24 h^{-1}) may be sufficient; if replicating the input more exactly is required, stronger nudging (1/6 h^{-1}) will help reduce the control error more in most cases. For reducing integration length, the results found here suggest that, when nudging only horizontal winds, the lowest value of G used (1/24 h^{-1}) may be best. Using G = 1/24 h^{-1} with uv-nudging resulted in the least suppression of u, v and θ adjustments relative to the free simulation pair, generally low control errors across u, v, θ and TOA radiative fluxes, and only a small change in ERF. It is worth noting that an even lower value of G could be more effective for this purpose, but this is not pursued further here. For mechanism denial experiments, G must be higher because of the need to actively suppress processes. As described in the discussion of the optimum nudging setup, this was found to be G = 1/6 h^{-1} (though using G = 1/12 h^{-1} could be sufficient for uv-nudging only).

Nudged level range

For constraining to specific meteorology, again the exact purpose should determine whether nudging lower levels is important, for example, whether the focus of the study is on lower troposphere changes or upper troposphere and stratosphere. However, if θ is to be nudged, then the lower model levels should be nudged regardless. For reducing integration time the results suggest that nudging the lower levels had little effect on control error and adjustments in the uv-nudged case. Assuming the constraint of natural variability is approximately the same whether nudging the lower levels or not, one could suggest that avoiding nudging those levels may be better, simply to avoid over-constraining the model. However, this has not been tested thoroughly, and could be further investigated. For mechanism denial the recommendation is similar: if nudging just horizontal winds, nudging the lower levels makes little difference; for nudging θ, the lower levels should be nudged.

Input time resolution

While only a small test was carried out here, the results indicated that control error in θ can be reduced by nudging to 3 hourly input vs 6 hourly input, but with little further improvement using 1 hourly input. The studies by Zhang et al. (2022) and Davis et al. (2022) also noted improvements when nudging to higher resolution nudging input. As such, for constraining to specific meteorology or reducing nudging integration time it may be advisable to use 3 hourly nudging. However, the benefits of using 3 hourly nudging input should still be considered against available computing resources. For mechanism denial increasing the nudging input time resolution did not increase the suppression of adjustments much, which is more important for this use case than minimising the control error. As such, it may be less worthwhile to use 3 hourly input for mechanism denial experiments, though further work could investigate this.

Suitability of nudging

While somewhat subjective, it is worth commenting on how suitable nudging is for achieving each purpose. For constraining a simulation to specific meteorology, nudging could be considered suitable, since this is the most direct effect of nudging and nudging can achieve low control errors when optimal settings are used. For reducing integration time, the effect of nudging on radiative adjustments, and consequently ERF, brings into question how suitable nudging is as a tool. In some cases these effects may be insignificant, such as calculating aci effects from an aerosol perturbation, as suggested by (Johnson et al., 2019). For example, one could calculate aci radiative effects due to an aerosol perturbation, by taking the difference in ERF calculated from two pairs of nudged simulations, one with aci effects switched on and one without. While nudging will affect radiative adjustments other than aci it should do so similarly in each pair and so the effects will cancel, while the effect of aci should be relatively unaffected by nudging. However, if overall ERF is of interest and something is perturbed that causes significant radiative adjustments, such as aerosols, the results shown here suggest that any nudging can effect the overall ERF. If lower values of G than 1/24 h^{-1} are used, this difference may be reduced, and could be investigated in future work. Finally for mechanism denial experiments, as explained in more detail in the conclusions, nudging may be an effective technique for investigating circulation-mediated adjustments, but is less effective for atmospheric temperature-mediated adjustments.

Chapter 3

Results from applying model nudging process denial experiments to anthropogenic sulphate and black carbon aerosol radiative adjustments

3.1 Introduction and methods

This chapter presents results from using the nudging mechanism denial method outlined in section 2.2 to isolate radiative adjustments from PI to PD perturbation of anthropogenic SU and BC emissions in UKESM1. These results are produced using the same UKESM1 control and SU perturbed fSST free running simulations used in chapter 2 and control and SU perturbed uv and uvθ-nudged simulations that used the optimum nudging setup identified in section 2.7. Specifically this nudging setup uses G = 1/6 h^{-1}, nudging applied from model level 2 ($\sim$50 m) to 82 ($\sim$70 km), and 6 hourly nudging input. In addition, one free, one uv-nudged and one uvθ-nudged simulation with BC perturbed were conducted, using the same nudging setup for the nudged simulations. To clarify, these simulations use the model/simulation setup and nudging code described in sections 2.3 and 2.4 respectively, with the free simulations being 35 years with 1 year for spinup, and the nudged simulations 5 years with 1 year for spinup.

While the results presented in chapter 2 focused on evaluating the performance of different nudging setups for the mechanism denial methodology, this chapter focuses on the results from applying this methodology, particularly focusing on different adjustments and their associated radiative impacts. As with the results from chapter 2, all results presented in this chapter are determined using multi-annual means integrated over the run length (minus 1 year spinup) of simulations unless otherwise stated and all uncertainties given are calculated as two times the standard error of the multiannual mean. References to statistical significance for all results are based on difference with two standard errors, implying statistical significance at the 95 % confidence limit assuming a normal distribution.

This chapter is structured as follows. Firstly, results concerning the effectiveness of the nudging methodology in suppressing circulation and atmospheric temperature adjustments are presented. These results somewhat overlap with those presented for SU perturbation experiments in 2.6, but focus solely on the best case nudging setup, which is used in this chapter, and introduce results for the BC perturbed simulations as well as SU. Secondly, a breakdown of radiative forcing across the free, uv-nudged and uvθ-nudged and different sky types is presented. This is followed by results of the cloud adjustments and then other adjustments (temperature, humidity) across the different simulations. Finally a discussion of the results and what can be learned about aerosol radiative adjustments is given, as well as suggestions for improvements to the nudging methodology.

3.2 Effectiveness of nudging for suppressing adjustments

Figure 3.1 illustrates profiles of multi-annual level mean adjustments in θ and northward wind speed arising from SU and BC Pi to PD perturbation. The adjustment in the free simulation pair can be compared with that of the nudged simulations to assess how well the nudging suppresses adjustments. As explained in section 2.6, ideally nudging would cause the wind speed adjustments to be zero and the θ adjustment zero in the uvθ-nudged simulation pair. Also because adjustment of the horizontal wind fields is harder to characterise with simple level-averaged means the root mean square (RMS) of the multi-annual mean adjustment determined over each level is used instead.

The free adjustment in θ to the SU forcing is characterised by a cooling of the troposphere, with warming in the stratosphere (fig. 3.1c). In the uvθ-nudged simulation pair, the tropospheric cooling response is partially suppressed, with an adjustment approximately half that of the free θ adjustment across most of the troposphere. This however excludes the lowermost model levels (<100m) where the adjustment remains almost as big as in the free simulation pair. This response can be partly ascribed to the land surface temperature not being fixed and the nudging extending to model level 2 ($\sim$50 m) not the lowermost model level. In the stratosphere, the uvθ-nudged simulation pair also shows a cooling adjustment, increasing roughly linearly with height, despite the free adjustment having a warming response. This suggests the nudging is not just suppressing the adjustment, but introducing a new signal into the simulation.

Under BC forcing the θ adjustment in the free simulation pair is approximately opposite that of SU in the troposphere, causing a warming effect. The uvθ-nudging is less effective at suppressing the tropospheric temperature adjustment to BC than SU, with no reduction of the θ adjustment between the free and uvθ-nudged below $\sim$1 km and between 3 to 5 km, and a reduction of only <50 % between 1 and 3 km and in the upper troposphere. It is not clear what could cause this difference between the two aerosols, but would suggest nudging is less successful at suppressing temperature adjustments due to BC than SU in these experiments. In the stratosphere, the uvθ-nudged BC perturbed case mirrors the SU perturbed case, continuing to show a positive θ adjustment despite the free adjustment being negative.

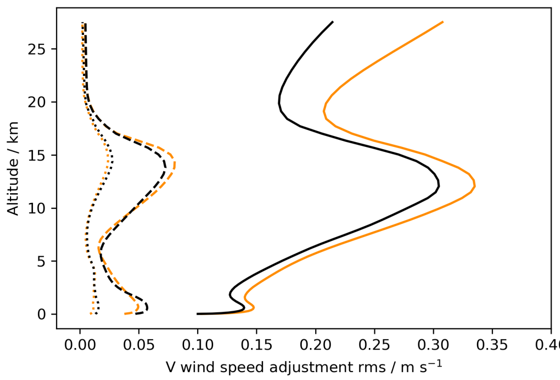

(a) v-wind

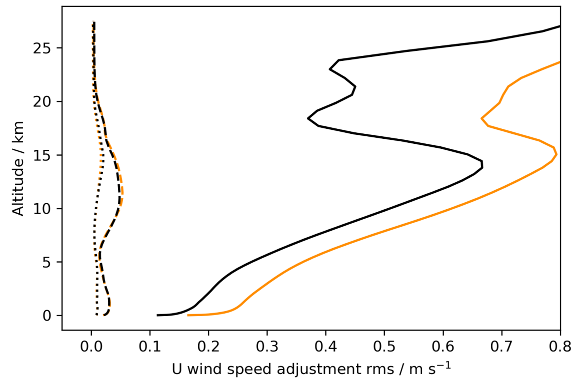

(b) u-wind

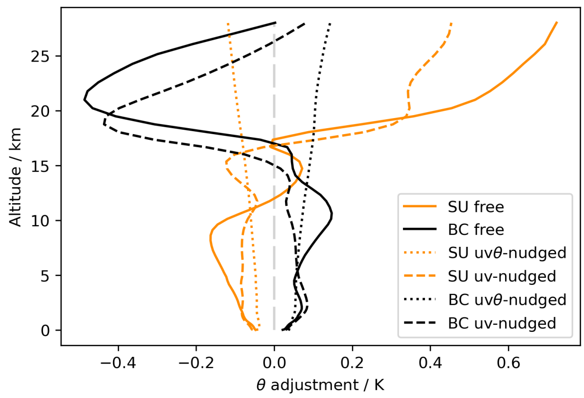

(c) θ

Figure 3.1: Profiles of global level-mean adjustments in (a) northward wind (v), and (b) eastward wind (u), both in ms^{-1}, and (c) potential temperature, in K, in response to perturbation of SU or BC from 1850 to 2014 emissions calculated from free, uv-nudged and uvθ-nudged control/perturbed simulation pairs using the optimum nudging setup identified in chapter 2. For potential temperature, the global model level mean of the multiannual mean adjustment is given. For horizontal winds the root mean square (RMS) of the multiannual mean adjustment calculated over each model level is given to better represent the mean adjustment.

For the SU perturbation, both the northward and eastward wind adjustments in both uv- and uvθ-nudged simulation pairs are substantially suppressed relative to the free simulation pair (fig. 3.1b and 3.1a). The uv-nudged simulation pair suppresses 60 to 70 % of the northward wind adjustment below 2 km, 70 - 90 % between 2 and 20 km and >90 % above 20 km. The uvθ-nudged simulation pair suppresses the northward wind adjustment even more effectively, with >80 % of the adjustment suppressed at all levels. A similar result is found for eastward wind adjustment, with uv-nudging suppressing >80 % of the adjustment below 2 km and >90 % above 2 km, and uvθ-nudging suppressing >90 % at all heights.

Under BC perturbation the wind adjustments are broadly similar. In the free case for both eastward and northward wind speeds the adjustment is smaller than under SU perturbation, but has very similar variation with altitude. Under uv- and uvθ-nudging the wind adjustments are very similar to the SU case. Because the adjustment in the free case is smaller for BC than SU, consequently the adjustment suppression could be considered slightly less effective in the BC perturbed nudged simulations.

Figure 3.2 illustrates profiles of control errors in θ and northward wind speed (Control error is defined in section 2.6). Like for the adjustments in figure 3.1, these use a multi-annual level mean for θ and multi-annual level RMS for northward wind. For the θ control error, the uvθ-nudged control has a small positive control error, being between 0.05 and 0.1 K warmer than the free control in the troposphere. This is smaller than the free θ adjustment in the troposphere seen in figure 3.1c, but of similar magnitude to the uvθ-nudged adjustment. In the uv-nudged case the θ control error is between -0.1 and 0.1 K in the troposphere, which is less than the free adjustment and similar to the uv-nudged adjustment. However in the stratosphere the uv-nudged θ control error is much larger, varying between 1.8 and -1.4 K in the low and mid stratosphere respectively. This raises some concern that only nudging the horizontal winds may introduce spurious temperature responses in the stratosphere, which could have knock on effects in the troposphere also.

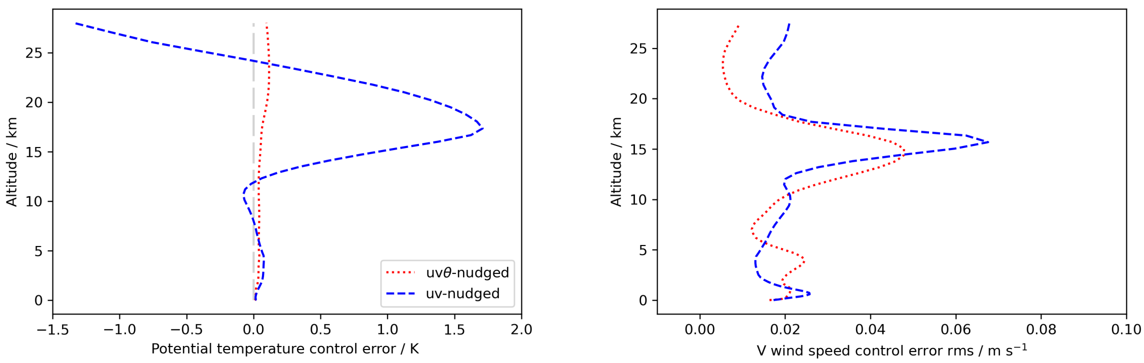

(a) θ **(b)** v-wind

Figure 3.2: Profiles of global level-mean control error in (a) potential temperature, in K, and (b) northerly wind, in ms^{-1}, calculated between free control minus uv-nudged or uvθ-nudged control simulations. For potential temperature, the global model level mean of the multiannual mean bias is given. For horizontal winds the root mean square (RMS) of the multiannual mean control error calculated over each model level is given to better represent the mean control error. The control error is determined using only years 2 to 5 from the free control to match the pattern of internal variability in the nudging simulations.

For the northward wind control error, the uv- and uvθ-nudged control simulations have similar control error magnitudes, with slight differences over height. These errors are of similar magnitude to the nudged northward wind adjustments in figure 3.1a, although are larger between 13 to 18 km, and are smaller than the free northward wind adjustment by 70 to 90 %. The control errors in the eastward wind have a similar profile. As such, the horizontal wind control errors can be considered small relative to the adjustments being suppressed, suggesting nudging is not introducing sizeable artefacts into the horizontal wind speed fields, increasing confidence in the results.

It is clear from these results that while the nudging does suppress some of the adjustment in θ and horizontal wind speeds in the troposphere, it does not totally suppress the adjustment, particularly for θ. Nudging also introduces a number of unintended effects, including the large cloud control errors noted in section 2.6. The incomplete suppression of adjustment and the artefacts introduced by nudging are likely to affect the results to some extent, particularly in the uvθ-nudged case. With these limitations in mind, the remaining results consider the radiative forcing of the circulation and atmospheric temperature adjustments and the mechanisms through which they affect radiative forcing.

3.3 ERF and radiative adjustments from temperature and circulation adjustments

The radiative forcing arising from circulation and temperature mediated adjustments are calculated for SU and BC from the method outlined in figure 2.1 and equations 2.1 and 2.2. These radiative forcings are illustrated in figure 3.3. For both aerosol types the circulation and temperature mediated adjustments have opposite forcing to their ERF, acting to offset their overall forcing. For SU the net temperature mediated (0.14 $\pm$ 0.04 W m^{-2}) and circulation adjustments (0.10 $\pm$ 0.08 W m^{-2}) are both positive and significantly different from zero, with the temperature mediated adjustment larger but within two standard errors of the circulation adjustment magnitude. For BC, both are negative though only the net temperature mediated adjustment (-0.25 $\pm$ 0.08 W m^{-2}) is significantly different from zero (circulation adjustment: -0.04 $\pm$ 0.07 W m^{-2}). The temperature mediated adjustment is also a much larger fraction of the ERF in the BC perturbed case than in the SU case, with a magnitude around 50 % of the ERF compared to around 10 % in the SU perturbed case. For both species the temperature adjustment arises from a larger shortwave (SW) component (positive for SU; negative for BC) with a smaller offsetting LW component (negative for SU; positive for BC). As reported later, this SW/LW pattern can be largely explained by cloud radiative adjustments, with increased cloud resulting in negative SW forcing due to enhanced scattering of incoming solar radiation, but positive LW forcing due to enhanced absorption of OLR (and vice versa for decreased cloud).

Table 3.1 further breaks down the radiative forcing due to 1850 to 2014 SU emissions perturbation into the all sky, CRE, DRE and clear-clean sky components for the free, uv-nudged and uvθ-nudged experiment pairs. In the free case, the negative ERF under SU perturbation can be seen to arise mostly ($\sim$60 %) from the CRE, but with a large ($\sim$40 %) contribution from the DRE, as expected from previous studies (Smith et al., 2018; O'Connor et al., 2021). The CRE is dominated by a large negative SW forcing, indicative of an increase in scattering by cloud in line with aci effects. There is also a smaller positive LW forcing, which is indicative of a decrease in outgoing longwave radiation (OLR) due to clouds. This could be via increased cloud fraction absorbing more OLR or decreased cloud top temperatures in response to aerosols as suggested in Quaas et al. (2009).

These values are similar to the study of O'Connor et al. (2021) who also used UKESM1 with perturbation of SU from PI to PD conditions. The all sky and SW CRE forcing from their study and this one are within 2 standard errors. However, one small difference is that the LW CRE found here is slightly larger than theirs (0.17 $\pm$ 0.02 Wm^{-2}).

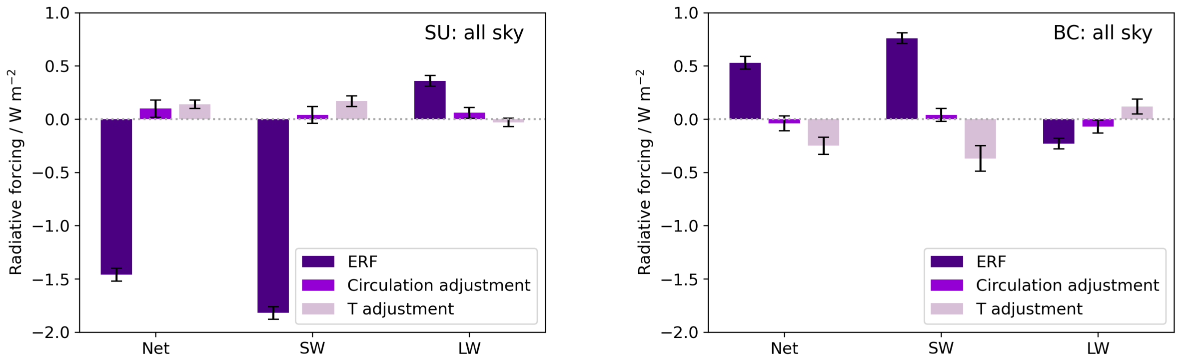

(a) SU **(b)** BC

Figure 3.3: Effective radiative forcing (ERF), circulation-mediated adjustment radiative forcing and atmospheric temperature-mediated adjustment radiative forcing, in Wm^{-2} due to perturbation of SU (a) and BC (b) from 1850 to 2014 emissions. The ERF is determined as the difference in TOA net down flux between the free control and free perturbed SU/BC simulations. The circulation adjustment radiative forcing is determined as the difference between the free ERF minus the radiative forcing from the uv-nudged simulation pair. The atmospheric temperature adjustment is determined as the difference in forcing from uv-nudged simulation pair minus the forcing from the uvθ-nudged simulation pair. All are global, multi-annual means in all-sky conditions. Numerical values are presented in tables 3.1 and 3.2.

The nudged simulation pairs have similar radiative forcing to the free pair in the DRE, but with significant differences in the CRE and clear-clean sky. It is likely that DRE is affected very little by the nudging because there should be little impact of circulation or temperature on the direct aerosol interaction with radiation. In contrast, in the uv-nudged case the CRE has a slightly stronger negative radiative forcing (-0.09 Wm^{-2}), while in the uvθ-nudged simulation pair the CRE has a considerably stronger negative forcing (-0.32 Wm^{-2}). Consequently, the circulation adjustment (free minus uv-nudged) and temperature adjustment (uv-nudged minus uvθ-nudged) radiative forcing both principally arise from a positive change in the cloud radiative forcing, with a stronger change in response to temperature adjustment.

While not an exact comparison, (Smith et al., 2018) reported ari cloud radiative adjustments (i.e. excluding aci and surface albedo) amounting to ~-0.4 and ~-0.8 Wm^{-2} due to a five times increase in SU from PD in HadGEM2 and HadGEM3 respectively (their figure S8). Dividing their values by 5 approximate the perturbation from PI to PD emissions performed here, equivalent to ~-0.08 and ~-0.16 Wm^{-2}. These values are of the opposite sign to the CRE circulation- (0.09 $\pm$0.08 Wm^{-2}) and atmospheric temperature-mediated (0.23 $\pm$0.04 Wm^{-2}) adjustments in table 3.1. While the CRE in table 3.1 will include some effect of aci, this is unlikely to be large, with most of the adjustment expected to be due to semi-direct effects. Therefore the results here suggest an ari (semi-direct effect) radiative adjustment of the opposite sign for SU to their study. It is worth noting however, 3 other models in their study reported smaller negative ari cloud radiative adjustments, and two reported positive ari cloud radiative adjustments.

It is also worth noting that the difference in CRE between the free and uvθ-nudged simulation pairs is primarily in the SW, with only a small change in forcing in the LW. This is consistent with the temperature adjustment causing a decrease in low cloud, which would decrease SW negative forcing but have a smaller effect on LW forcing, as described in the next section. By contrast the uv-nudged simulation pair CRE differs from the free pair by similar amounts in the LW and SW, suggesting the circulation adjustment affects high cloud too.

Aside from the CRE, the clear-clean sky component of the forcing has a net value close to zero in the free and uv-nudged cases. However, in the uvθ-nudged simulation pair both SW and LW components are positively shifted resulting in a positive clear-clean response. This partially offsets the stronger negative CRE in the uvθ-nudged case resulting in a smaller total (all sky) temperature adjustment forcing.

Table 3.1: Radiative forcing due to perturbation of SU from 1850 to 2014 emissions, determined from free, uv-nudged and uvθ-nudged control-perturbed simulation pairs. Circulation-mediated adjustment and temperature-mediated adjustment calculated using equations 2.1 and 2.2 are also given. For each experiment the radiative forcing is given for all sky, DRE, CRE, and clear-clean (residual) forcing components using the decomposition of Ghan (2013). Forcing components are further subdivided into net, SW and LW forcing. All values are given in Wm^{-2}.

Sky type	Band	Free	uv-nudged	uvθ-nudged	Circ. adjust.	Temp. adjust.
All	Net	-1.46±0.06	-1.55±0.02	-1.69±0.01	0.10±0.08	0.14±0.04
	SW	-1.82±0.06	-1.86±0.02	-2.03±0.02	0.04±0.08	0.17±0.05
	LW	0.36±0.05	0.30±0.00	0.34±0.04	0.06±0.05	-0.03±0.04
DRE	Net	-0.63±0.01	-0.65±0.00	-0.64±0.01	0.02±0.02	-0.01±0.01
	SW	-0.67±0.02	-0.69±0.00	-0.68±0.01	0.02±0.02	-0.01±0.01
	LW	0.04±0.01	0.04±0.00	0.04±0.00	0.00±0.01	0.00±0.00
CRE	Net	-0.84±0.06	-0.93±0.02	-1.16±0.02	0.09±0.08	0.23±0.04
	SW	-1.07±0.06	-1.10±0.02	-1.32±0.03	0.03±0.08	0.22±0.05
	LW	0.23±0.03	0.17±0.00	0.16±0.02	0.05±0.03	0.01±0.02
ClearClean	Net	0.02±0.04	0.02±0.00	0.11±0.02	-0.01±0.04	-0.09±0.03
	SW	-0.08±0.03	-0.07±0.00	-0.03±0.00	-0.01±0.03	-0.04±0.00
	LW	0.10±0.04	0.09±0.00	0.14±0.02	0.01±0.04	-0.05±0.03

Considering the same breakdown for BC (Table 3.2), the positive forcing in the free simulation pair arises primarily ($\sim$ 95%) from the DRE with only a small net positive contribution from the CRE. The small CRE results from a positive SW forcing mostly offset by a negative LW forcing. This is different to results from other models which tend to indicate a larger negative net CRE for BC resulting from increased low cloud and decreased high cloud (Stjern et al., 2017), a difference that has been noted in previous studies using UKESM1 (Johnson et al., 2019; O'Connor et al., 2021).

As with SU these values can be closely compared with those of O'Connor et al. (2021), who perofmred a similar PI to PD BC perturbation to that performed here. Unlike for SU, the total all sky radiative forcing here (0.53 ± 0.06 Wm^{-2}) is larger than in their study (0.37 ± 0.03 Wm^{-2}). This difference mostly arises from a larger SW forcing from the DRE here than their clear sky forcing (DRE + clearclean). The SW CRE found here is also larger (0.26 ± 0.04 Wm^{-2}) than theirs (0.15 ± 0.02 Wm^{-2}). However, it is worth noting that in both studies the LW and SW CRE are of similar magnitudes, unlike for SU where the SW effect dominated strongly, suggesting high cloud changes are important for BC radiative adjustments, not just low cloud.

As with the perturbed SU experiments, the DRE is unchanged between the free and nudged simulations. Unlike in the SU experiments, the uv-nudged simulation pair has a similar CRE forcing to the free case, within two standard errors. However the SW and LW components of the uv-nudged CRE are both significantly smaller in magnitude than in the free CRE. This could suggest the circulation adjustment suppresses some cloud response in the uv-nudged case, particularly high cloud that may affect the LW forcing more than the SW. The uvθ-nudged CRE by contrast is much more strongly positive than the CRE in the free simulation pair, with a SW forcing more than twice the amount but a similar LW forcing (within two standard errors). This may suggest a strong low cloud increase due to the temperature adjustment is present in the free case, which acts to reduce the net positive forcing compared to in the uvθ-nudged case where temperature adjustment is suppressed (see next section). As with the SU experiments, the clear-clean sky response is strongest in the uvθ-nudged case, acting to reduce the all sky forcing difference between the uvθ-nudged simulation pair and free and uv-nudged simulation pairs.

Comparing these results with other studies, Smith et al. (2018) reported a cloud radiative adjustment of $\sim$-0.7 $\pm$ 0.3 Wm^{-2} for a 10 times increase in BC emissions relative to PD. Dividing their value by 10 ($\sim$-0.07 $\pm$ 0.03 Wm^{-2}), their value is smaller than the total CRE circulation and atmospheric temperature-mediated radiative adjustment in table 3.2 (-0.36 $\pm$ 0.12Wm^{-2}). While the value reported here will include adjustment due to aci effects and theirs excludes the first indirect effect, one can assume that most of the CRE for BC will be via ari (semi-direct effects) rather than aci adjustments. As such the radiative adjustment due to ari cloud effects in this work for BC appear to be large compared to the average from their study. One model in their study did have a much larger cloud radiative adjustment, however even this when divided by 10 ($\sim$ 0.24 Wm^{-2}; their figure S3) would be smaller than the combined circulation- and atmospheric temperature-mediated adjustment found here.

Figure 3.4 illustrates the spatial pattern of the all sky net TOA forcing for the 1850 to 2014 SU and BC perturbations. The spatial pattern of the ERF is broadly similar for the two species (some exceptions being over the Congo basin, Arabia, the Amazon and the coast of Peru) but of opposite sign (fig. 3.4a and 3.4b). This could be explained by the two aerosols having a broadly similar spatial distribution of emissions, aside from some areas (e.g. tropical or boreal forest biomass burning emissions). Both species show a stronger radiative forcing over land than sea, with especially strong forcing near to significant emission regions in south and east Asia. Over sea, the forcing is generally negative for SU except in areas of the Southern Pacific and Southern Ocean, whereas for BC the ERF over sea is more varied.

The uv-nudged forcing has a similar spatial pattern to the ERF for both species (fig. 3.4c and 3.4d). The main difference to the free simulation pair is a suppression of noise, particu-

Table 3.2: Radiative forcing due to perturbation of BC from 1850 to 2014 emissions, determined from free, uv-nudged and uvθ-nudged control-perturbed simulation pairs. Circulation-mediated adjustment and temperature-mediated adjustment calculated using equations 2.1 and 2.2 are also given. For each experiment radiative forcing is given for all sky, DRE, CRE, and clear-clean (residual) forcing components using the decomposition of Ghan (2013). Forcing components are further subdivided into net, SW and LW forcing. All values are given in Wm^{-2}.

Sky type	Band	Free	uv-nudged	uvθ-nudged	Circ. adjust.	Temp. adjust.
All	Net	0.53±0.06	0.57±0.01	0.82±0.07	-0.04±0.07	-0.25±0.08
	SW	0.76±0.05	0.72±0.00	1.09±0.12	0.04±0.06	-0.37±0.12
	LW	-0.23±0.05	-0.15±0.01	-0.27±0.06	-0.07±0.06	0.12±0.07
DRE	Net	0.50±0.01	0.50±0.01	0.51±0.03	0.00±0.02	-0.01±0.04
	SW	0.50±0.02	0.50±0.01	0.51±0.03	0.00±0.03	-0.01±0.03
	LW	0.00±0.01	0.00±0.00	0.00±0.00	0.00±0.01	0.0±0.00
CRE	Net	0.07±0.05	0.11±0.00	0.44±0.05	-0.03±0.06	-0.33±0.06
	SW	0.26±0.04	0.21±0.00	0.57±0.09	0.05±0.05	-0.37±0.09
	LW	-0.19±0.03	-0.10±0.01	-0.13±0.04	-0.09±0.03	0.03±0.05
ClearClean	Net	-0.04±0.04	-0.04±0.00	-0.13±0.02	0.00±0.04	0.09±0.02
	SW	0.00±0.03	0.01±0.00	0.01±0.01	-0.01±0.04	0.01±0.01
	LW	-0.04±0.04	-0.06±0.01	-0.14±0.02	0.01±0.05	0.08±0.03

larly noticeable over oceans, illustrating the effect of nudging the horizontal winds bringing the nudged simulation intrinsic variability closer to the nudging reference data. However, there are also differences over areas with larger forcing, in particular over ocean areas along the western coast of North and South America and over eastern China under SU perturbation.

The uvθ-nudged forcing has a similar spatial pattern to the ERF over land for both species, but differs over the ocean particularly for BC (figures 3.4e and 3.4f). Specifically, there is a significant positive forcing along the western coasts of several landmasses, with a smaller positive forcing along the storm tracks, in the BC perturbed case, that is not seen in the ERF. For SU, although the ERF shows areas of negative forcing in similar regions along the western costs of landmasses already, there is a stronger negative forcing in these regions in the uvθ-nudged case, and a negative forcing along the southern hemisphere storm track. As discussed in section 1.3.5, previous studies have found that increased BC emissions can cause an increase in marine stratocumulus cloud in these regions as part of the semi-direct effect, resulting in a overall negative radiative forcing due to increased scattering in the SW. Therefore, it seems likely the stronger positive forcing in these regions in the uvθ-nudged BC case is a result of less marine stratocumulus cloud. For SU, given the pattern is very similar, it seems likely the stronger negative forcing in the uvθ-nudged case is related to increased marine stratocumulus, and thus more scattering in the SW occurring.

By contrast over land there is relatively little difference between the uv-nudged and uvθ-nudged forcings. There is a weaker negative forcing over northern mid- and high-latitude land in the uvθ-nudged SU perturbation, though this is smaller than the difference in marine stra-

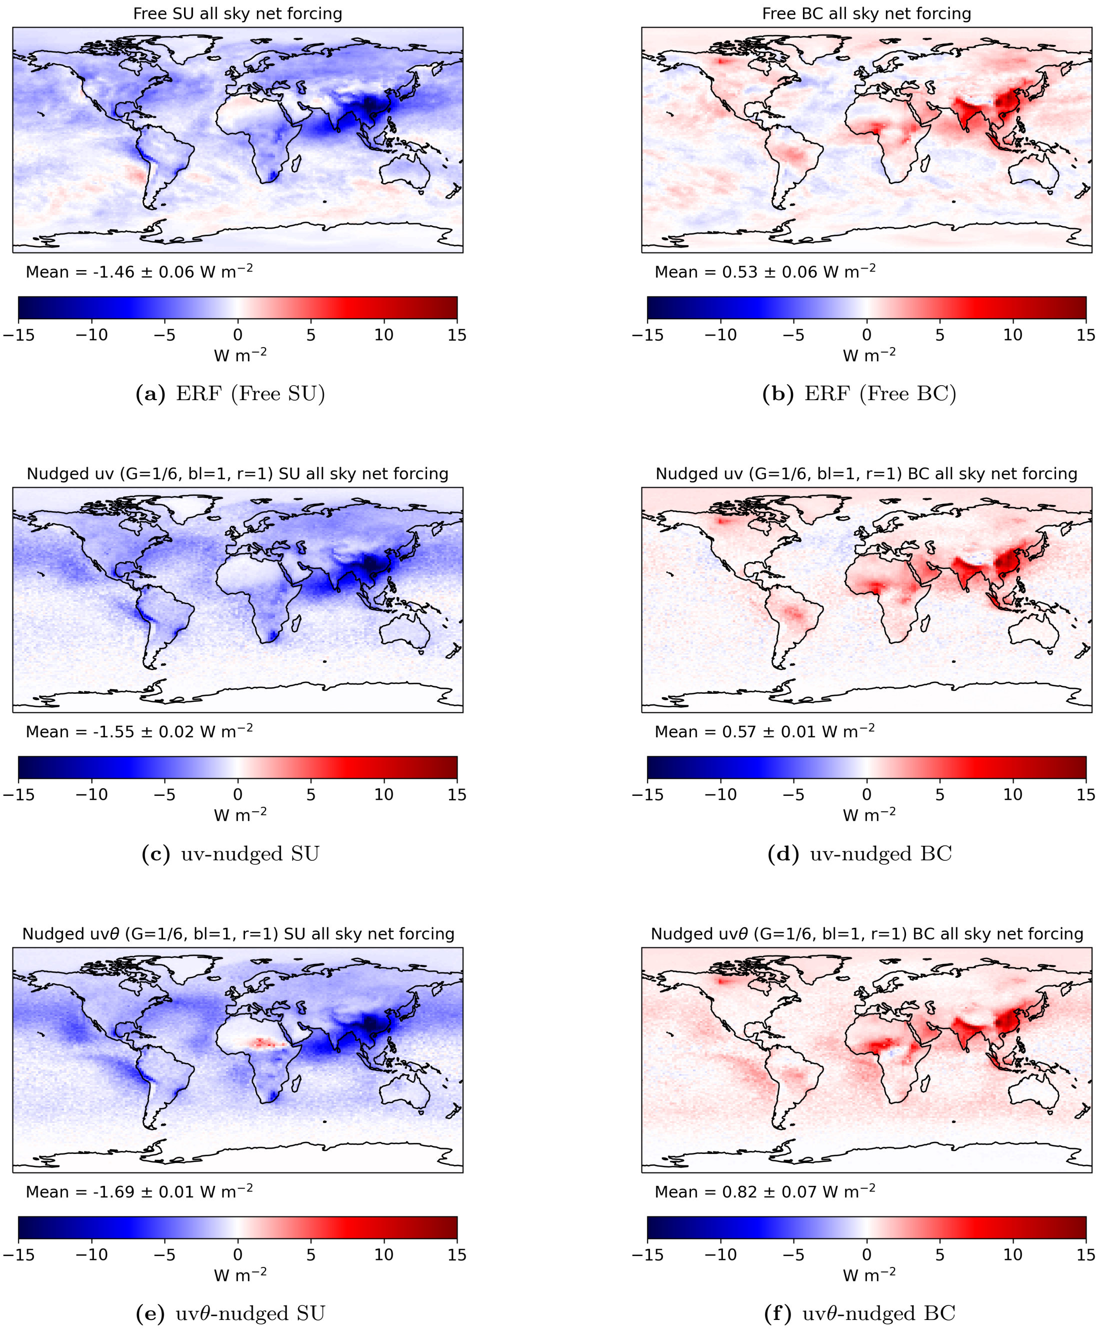

(a) ERF (Free SU)

(b) ERF (Free BC)

(c) uv-nudged SU

(d) uv-nudged BC

(e) uvθ-nudged SU

(f) uvθ-nudged BC

Figure 3.4: Multiannual mean TOA all sky radiative forcing due to 1850 to 2014 perturbation of SU (a, c, e) or BC (b, d, f), determined for (a,b) free (=ERF), (c,d) uv-nudged, and (e,f) uvθ-nudged simulation pairs.

tocumulus regions. In the BC perturbed case there is less change over land, with only a small decrease in the positive forcing over central South America and south east Asia. The SU perturbed uvθ-nudged case also exhibits a positive forcing over central Africa, and a slight negative forcing in a similar region for the BC perturbation, however this is likely due to spurious adjustments in the horizontal winds and potential temperature fields introduced by nudging rather than a real effect, as noted earlier in section 2.6.

3.4 Cloud response to aerosol perturbation

3.4.1 Low cloud adjustments

Given that most of the temperature- and circulation-mediated radiative adjustment forcing can be ascribed to the CRE, it should be possible to identify spatial patterns in the cloud adjustment that explain the radiative adjustment forcing. Figure 3.5 illustrates the adjustment in low cloud cover arising in each of the experiment pairs for SU and BC perturbations as well as the low cloud cover from the free control as a baseline for comparison. Low cloud is represented by taking the maximum cloud volume fraction across levels 3 to 16 ($\sim$75 to $\sim$2000 m above the surface) for each grid point. Low cloud adjustment is represented by taking the maximum cloud volume fraction across model levels 3 to 16 for the control and perturbed simulations for each grid point, and differencing these fields. This representation is chosen because of the variation in typical cloud heights over different latitudes, which would make a mean over a level range less useful. However, similar interpretations can be drawn from means across the same model level ranges (not shown).

In the free simulation pair with SU perturbation (fig. 3.5b) there is an increase in low cloud cover over some land areas, particularly in Eurasia and North America. Over oceans the signal is less clear with areas of both increased and decreased cloud cover. In the free BC perturbation (fig. 3.5c), the response over land is less clear, with the sign of the response varying by region. Over the oceans there is a similarly variable pattern in the sign of cloud response. The low cloud response is similar for both aerosols over the mid and high southern latitudes.

The main effect of the uv-nudging on low cloud adjustment appears to be the suppression of noise. In both SU perturbed (fig.3.5d) and BC perturbed (fig. 3.5e) simulation pairs the cloud adjustment is reduced considerably in all areas, except close to strong anthropogenic aerosol emission sources in south-east Asia. This is particularly noticeable over the Southern Hemisphere, where there are few anthropogenic aerosol sources.

Beyond suppressing noise, in the SU case, the uv-nudging reduces the low cloud increase over Eurasia and North America seen in the free simulation pair. Over the oceans there are also areas of increased cloud in the marine stratocumulus regions along the western edges of Africa and North and South America in particular, though the signal is weak. Following equation 2.1, these responses can be inferred as the result of suppressing circulation adjustments in the uv-nudged case, which appears to affect the low cloud response to SU forcing over these regions. In the BC case by contrast there is no distinct change in the adjustment beyond the noise suppression.

For the uvθ-nudged pair with SU perturbation (fig. 3.5f), two patterns are apparent in particular. Firstly, over land the positive cloud response is even weaker than in the uv-nudged case. This indicates the atmospheric temperature adjustment is responsible for some of the cloud increase over land seen in the free case, since suppressing the temperature adjustment here reduces the cloud response. Secondly, there is a pattern of strongly increased cloud over

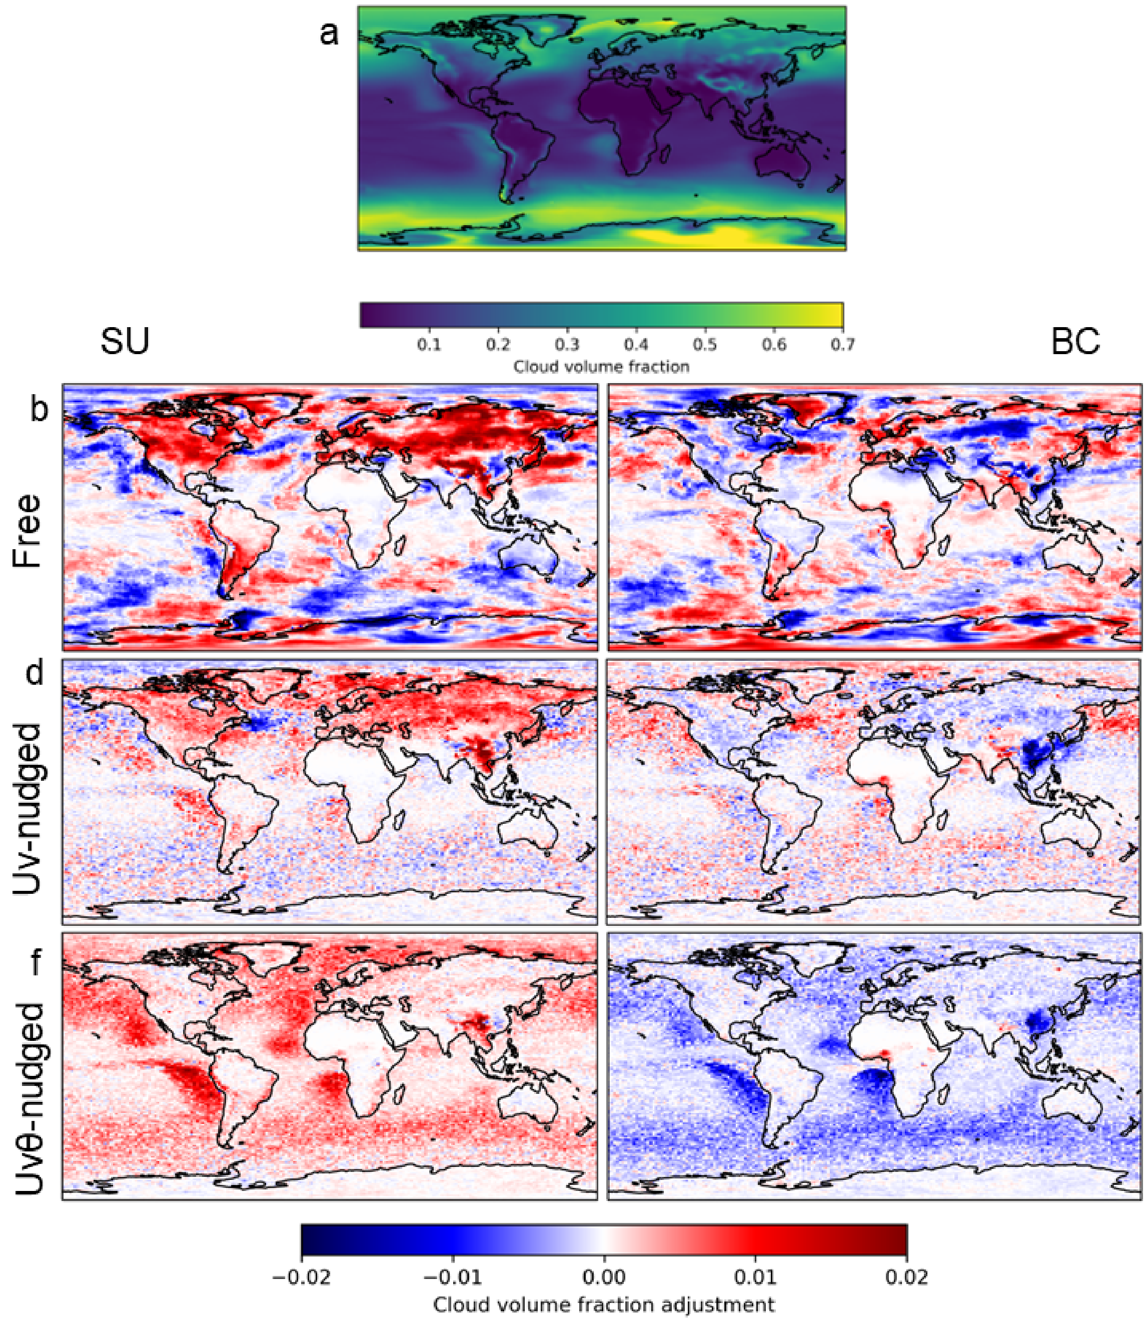

Figure 3.5: Multiannual mean low cloud volume fraction and its adjustment across simulations. (a) Low cloud in the free control, which is represented by the maximum cloud volume fraction over model levels 3 to 16 ($\sim$75 to $\sim$2000 m above mean sea level for points over sea) at each grid point for the control. Low cloud adjustment due to 1850 to 2014 perturbations of SU (b, d, f) or BC (c, e, g) determined for (b,c) free, (d,e) uv-nudged and (f,g) uvθ-nudged simulation pairs, represented by taking the level with maximum cloud volume fraction over model levels 3 to 16 for each grid point for the control and perturbed simulations, and differencing these fields.

the marine stratocumulus regions and along the jet stream paths. This suggests that the atmospheric temperature adjustment causes a decrease in cloud fraction over these regions, that would otherwise be seen in the free case, since the pattern only emerges when the temperature adjustment is suppressed. In the BC case (fig. 3.5g) the same spatial pattern of cloud response is also evident, though with a decreased cloud fraction, implying the atmospheric temperature adjustment causes an increase in cloud in these regions.

Comparing the low cloud adjustment with the findings from the radiative forcing maps in figure 3.4, it is clear that the negative forcing over marine stratocumulus regions in the uvθ-nudged SU perturbed simulation pair (fig. 3.4e) corresponds to an increase in low cloud in these regions (fig. 3.5f). A similar observation can be made for the BC case with a positive forcing in these regions (fig. 3.4f) and decreased low cloud (fig. 3.5g). What is less clear is what causes these opposite, but similar magnitude, marine stratocumulus responses. In the SU perturbed uvθ-nudged case, the marine stratocumulus increase could possibly be explained by aci adjustments, with the increase in SU acting to increase the (time averaged) cloud fraction by increasing cloud lifetime. However, the same mechanism could not apply in the BC case, since there is a decrease in marine stratocumulus cloud cover in the BC perturbed uvθ-nudged case. While BC is known to affect marine stratocumulus via semi-direct effects as described earlier, the decreased cloud fraction in these regions is greatest in the uvθ-nudged case, where θ adjustment is suppressed, and largely absent from the uv-nudged and free cases. Therefore it seems unlikely this response results from the semi-direct effect.

An alternative hypothesis for the marine stratocumulus decrease in the BC case could be that BC emissions cause a decrease in other CCN in these regions. This effect has been noted previously by other studies which have found that over polluted areas BC particles (particularly from fossil fuel emissions) can cause a decrease in viable CCN by acting as a sink for aerosol precursor gases, limiting growth of existing soluble particles that more readily act as CCN (Koch et al., 2011; Johnson et al., 2019). In the BC perturbed case we find a reduction in SU mass mixing ratio around south and east Asia, where BC burden is particularly high, in line with this mechanism, illustrated in figure 3.6. However, over marine stratocumulus regions, which are not so polluted, we found that SU mass mixing ratios decreased by only a few percent or not at all. The SU burden also typically decreased most in the BC perturbed free case (fig. 3.6a), less in the uv-nudged case (fig. 3.6b), and least in the uvθ-nudged case (fig. 3.6c) over marine stratocumulus regions and south and east Asia. Since no clear pattern of cloud decrease is observed in the stratocumulus regions in the BC perturbed free and uv-nudged cases, and over south and east Asia the cloud decrease is greater in the nudged cases than the free case, this suggests these changes in SU caused by BC are too small to result in the large low cloud increases seen over the marine stratocumulus regions.

As well as the cause of the low cloud response in the BC perturbed uvθ-nudged case being unclear, it is not clear why the cloud adjustment in the marine stratocumulus regions is absent in the uv-nudged case. This could possibly be explained by temperature mediated adjustments (specifically semi-direct effects) in the uv-nudged case causing a positive low cloud adjustment that offsets the decrease in the uvθ-nudged case (where temperature-mediated adjustments are suppressed). It is also unclear why the magnitude of low cloud adjustment in these regions is similar for BC and SU; it may suggest that both are a result of a similar mechanism rather than the low cloud response in the SU uvθ-nudged case being a result of aci adjustments.

To further understand the low cloud response, we can consider the adjustment in lower tropospheric stability (LTS). This is particularly useful for understanding atmospheric temperature adjustments, where cloud responses under the semi-direct effect are known to relate to changes

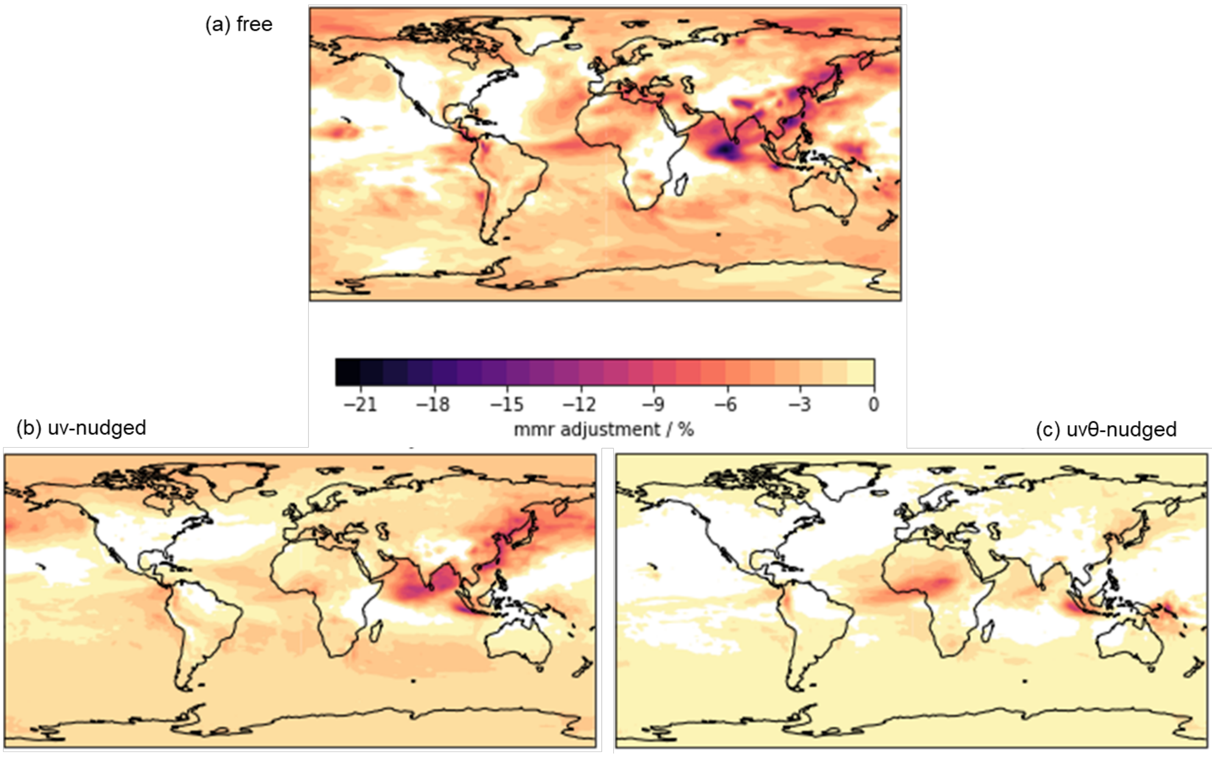

Figure 3.6: Multiannual mean adjustment in SU mass mixing ratio burden at model level 11 ($\sim$ 1 km) due to 1850 to 2014 BC emissions perturbation, for (a) free, (b) uv-nudged and (c) uvθ-nudged simulations pairs. Plotted as a percentage change of the SU burden in the free control, and only negative adjustments are shown.

in stability caused by aerosol forcing (Herbert et al., 2020). Figure 3.7 illustrates the lower tropospheric stability and its adjustment across the simulation pairs. LTS is approximated as $\theta_{20} - \theta_1$, where the subscript denotes the model level, with level 20 at approximately 3 km above mean sea level and level 1 the lowermost atmospheric level.

In the free control it can be seen that there is already higher LTS over the marine stratocumulus regions than elsewhere over oceans (fig. 3.7a). In the free SU perturbed case, the LTS generally decreases over oceans, especially over the northern storm track, and to a lesser extent over most of the marine stratocumulus regions, while it increases particularly over areas of land in Eurasia and North America. In the free BC perturbed case there is generally an increase in LTS over both land and oceans, though excluding the Southern Ocean and Antarctica.

In the uv-nudged SU perturbed case, compared with the free adjustment, much of the LTS adjustment in the southern hemisphere is absent, suggesting much of it was noise in the free case. The LTS increase over North America and Eurasia is still present but reduced, while there remains a decrease in LTS over mid and high northern latitude ocean which is more uniform than in the free case. In the uv-nudged BC perturbed case there is a very similar pattern LTS adjustment over ocean to the SU perturbed case, but of opposite sign, with an increase in LTS.

In the uvθ-nudged case the LTS adjustment is further suppressed for both SU and BC as would be expected since suppressing the atmospheric temperature response is the aim of nudging θ. Aside from over Central Africa, where there are known artefacts introduced by the nudging as

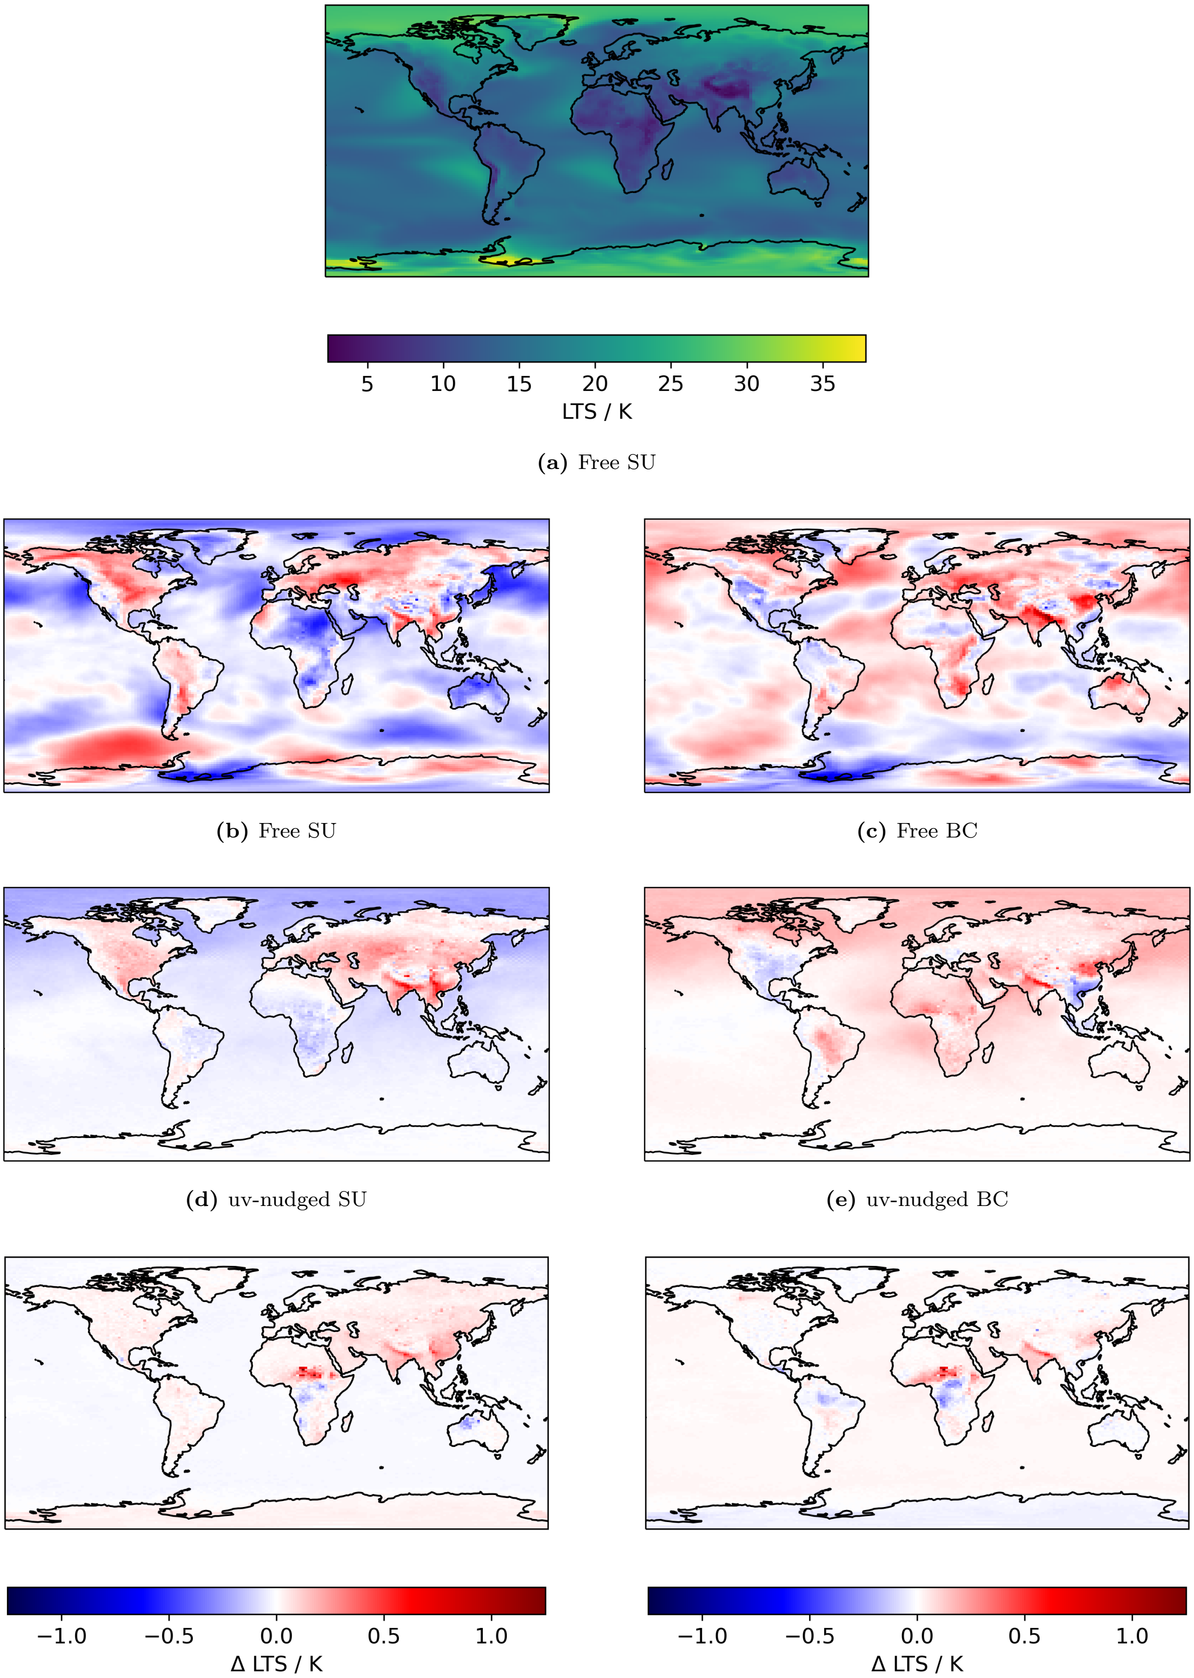

(a) Free SU

(b) Free SU

(c) Free BC

(d) uv-nudged SU

(e) uv-nudged BC

(f) uvθ-nudged SU

(g) uvθ-nudged BC

Figure 3.7: (a) Multiannual mean lower tropospheric stability (LTS) from the free control, and lower tropospheric stability adjustment due to 1850 to 2014 SU (b, d, f) or BC (c, e, g) perturbation determined for (b,c) free, (d,e) uv-nudged and (f,g) uvθ-nudged simulation pairs. LTS is calculated as the difference in θ at level 20 minus level 1.

noted previously, the only areas with much LTS adjustment are over south and east Asia in the SU case, and specifically near to high emission areas in India and Eastern China for both species. Even still these adjustments to LTS are smaller in magnitude than in the uv-nudged or free cases.

While some changes can be put down to removal of noise by the nudging, the progressively decreased magnitude of the positive LTS adjustment over North America and Eurasia between the free, uv-nudged and then uvθ-nudged SU perturbed cases correlates well with the progressively decreasing positive cloud fraction adjustment noted earlier from figure 3.5. This could point to a feedback arising from SU forcing whereby cooling of the land surface (see figure 3.9) more strongly than the troposphere from SU forcing causes an increase in LTS, that then causes an increase in low-cloud fraction. This is however in contrast to the BC perturbed cases, where cloud adjustment is also generally negative over land in the nudged cases, despite the LTS adjustment being generally positive.

What is conspicuously absent from the uvθ-nudged LTS adjustments for both aerosols is an adjustment over the marine stratocumulus regions or along the storm tracks. While it is expected that nudging θ should suppress adjustment in LTS, this confirms that a semi-direct effect cannot be responsible for the changing low cloud in these regions under uvθ-nudging seen in figure 3.5f and 3.5g. It can also be seen that while the BC perturbed uv-nudged case shows an LTS response in the marine stratocumulus regions west of southern and western Africa, there is no significant response west of northern South America or Central America. Consequently, it seems unlikely that the lack of low cloud adjustment in these areas in the BC perturbed uv-nudged case could be explained by semi-direct effects offsetting the negative adjustment seen in the uvθ-nudged case.

While there is no obvious physical explanation for the low cloud adjustments in the uvθ-nudged simulation, it is worth noting that the pattern is similar to the bias identified by Zhang et al. (2022) and illustrated in figure 1.6. In their study, the bias was caused by calculating nudging increments in nudged simulations at a different point in the model timestep to when the nudging input data was generated from a free control. A similar pattern of bias is observed in the cloud control error for the uvθ-nudged control simulation conducted here (not shown). While the control error is not the same as an adjustment, and so does not prove that the cause of the marine stratocumulus adjustments in the uvθ-nudged simulation pairs is due to the nudging timestep sequence, and Zhang et al. (2022) employed a different model that may exhibit different behaviour, this possibility definitely requires further investigation.

3.4.2 High cloud adjustments

Figure 3.8 illustrates high cloud volume fraction from the free control and its adjustment in each simulation pair due to SU and BC perturbations. This uses the same representation as for low cloud: the maximum cloud volume fraction over model levels 32 ($\sim$7 km) to 45 ($\sim$14 km) is used for the control volume fraction, while cloud adjustment is represented by taking the maximum cloud volume fraction over these levels for the control and perturbed simulations and differencing.

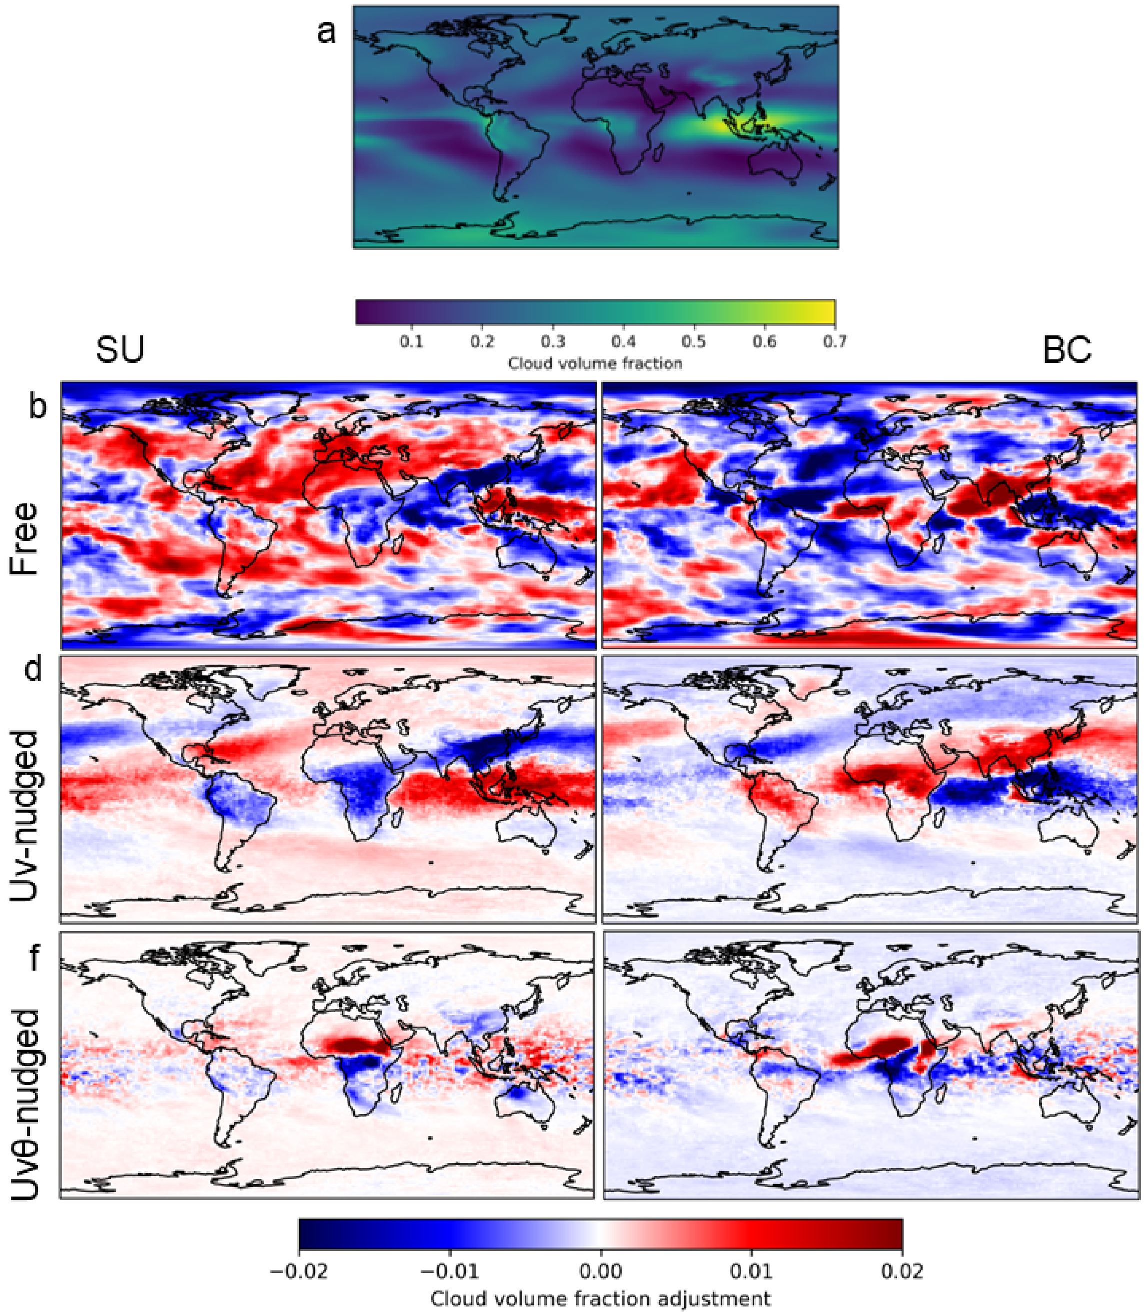

Figure 3.8: Multiannual mean high cloud volume fraction and its adjustment across simulations. (a) High cloud in the free control, which is represented by the maximum cloud volume fraction over model levels 32 to 45 ($\sim$7 to $\sim$14 km above mean sea level for points over sea) at each grid point for the control. High cloud adjustment due to 1850 to 2014 perturbations of SU (b, d, f) or BC (c, e, g) determined for (b,c) free, (d,e) uv-nudged and (f,g) uvθ-nudged simulation pairs, represented by taking the level with maximum cloud volume fraction over model levels 32 to 45 for each grid point for the control and perturbed simulations, and differencing these fields.

For the free case, the SU and BC high cloud adjustments are generally opposite in spatial pattern but with similar magnitudes, the main exceptions being close to the poles. In the uv-nudged case the spatial pattern is even more closely mirrored. For both species the high cloud adjustment is suppressed outside of approximately 30°S to 50°N, with a generally similar spatial pattern within this latitude range to in the free case, except for over the North Pacific, North Indian Ocean and Northern South America where the sign of cloud adjustment has changed. In the uvθ-nudged case the high cloud adjustment is limited to areas along the intertropical convergence zone (ITCZ), with particularly strong adjustment over central Africa. The response over central Africa is very likely related to the aforementioned control errors in theta and horizontal winds at least in part, and can explain the anomalous radiative forcing in this area in figure 3.4 for the uvθ-nudged case. It is not clear whether the response generally along the ITCZ in the uvθ-nudged cases is related to this anomaly or representing a real process.

Comparing the high cloud response in free, uv-nudged and uvθ-nudged cases, the results indicate that outside of 30°S to 50°N the adjustment is primarily a result of circulation adjustments or noise, since it is suppressed in the uv-nudged case; while the adjustment within this range is largely a result of temperature adjustment, since much of the adjustment in the uv-nudged case is suppressed in the uvθ-nudged case (excluding the anomalous pattern over central Africa). The similarity between the SU and BC perturbed high cloud adjustments is perhaps surprising given the different interactions by which SU (scattering) and BC (absorbing) affect climate. This could possibly be explained by the high cloud adjustment being generally above the altitude anthropogenic aerosol emissions typically reach and so more dependent on the large scale changes to atmospheric temperatures, and thus convection, than local responses to aerosol forcing.

While the high cloud response does not explain as much of the TOA forcing signal seen in figure 3.4 as the low cloud response, some areas of the high cloud adjustment pattern are clearly correlated with the radiative forcing. For example, in the SU perturbed uv-nudged case (fig. 3.4c) across south and south-east Asia, southern Africa and northern South America the negative forcing follows the spatial pattern of the decrease in high cloud (fig. 3.8d). Although part of the negative forcing in these areas is likely to result from the DRE of SU, there is a negative LW radiative forcing in these areas (not shown), likely related to a decrease in absorption of outgoing LW by high clouds due to decreased high cloud fraction. Furthermore, in table 3.2 it is apparent that the uv-nudged BC case has a less negative LW CRE forcing than the free BC case, which is likely explained by differences in the high cloud response rather than low cloud. While a difference in the net high cloud change in the uv-nudged and free BC perturbed cases is not obvious from figures 3.8c and 3.8e, zonal mean cloud adjustments show that high cloud reduces less in the uv-nudged case than the free case and is likely to be the explanation for the more negative LW CRE forcing in the free BC case (not shown). This is consistent with the findings of Johnson et al. (2019) who also found that uv-nudging suppresses the high cloud decrease due to BC perturbation, though here it can be seen that while the free and uv-nudged cases are not significantly different in net ERF, this high cloud change is associated with a significant change in the LW CRE forcing.

3.5 Non-cloud adjustments

Whilst the contribution of clear-clean sky forcing to the ERF is small relative to the CRE, it is significant, particularly in the uvθ-nudged simulations (Table 3.1 and 3.2). Two of the primary contributions to the clear-clean sky forcing are the surface temperature response and specific humidity response. It should be noted that since SSTs are fixed as part of the experiment design, ocean temperatures can only change where sea ice is present, but the land surface is allowed to respond.

Figure 3.9 illustrates the surface temperature response in each of the simulations pairs for both aerosols. The land surface temperature response for the two species is quite different with a negligible global mean change caused by the BC perturbation and a small global mean cooling caused by the SU perturbation (-0.06 K) in the free simulation pairs (fig. 3.9). However, locally the SU perturbation produces stronger changes, particularly a cooling of up to 1.5 K over most of North America and Eurasia (fig. 3.9a). The areas of cooling correspond with areas of greatest low cloud fraction increase (fig. 3.5f) and LTS increase (fig. 3.7b). The BC perturbation, while also showing some regional change in surface temperature, only causes large surface temperature changes over India and north eastern China where there is up to 1.5 K cooling (fig. 3.9b).

For the SU perturbation the uv-nudging decreases the surface temperature response over southern hemisphere land, suppressing noise, but also reduces the cooling over North America and Eurasia by approximately 0.5 K (fig. 3.9c). This suggests the circulation adjustment for SU is responsible for around half of the surface cooling over North America and Eurasia. In the BC case (fig. 3.9d), the uv-nudging likewise suppresses the surface temperature response, indicative of the noise suppression effect of nudging, and weakens the cooling over India and north eastern China.

In the uvθ-nudged case the cooling over land due to SU perturbation is reduced further, except over central Africa which is likely due to nudging artefacts (fig. 3.9e). With BC perturbation, the uvθ-nudging also decreases the magnitude of the surface temperature adjustment (fig. 3.9f).

It is to be expected that the atmospheric temperature adjustment will affect the land surface temperature (when not fixed), but it is noteworthy that the uv-nudging can also suppress the land surface temperature adjustment, both locally and overall, illustrated in the SU perturbed experiment pairs. Furthermore, the change in land surface temperature in the SU perturbed cases highlights that, while the uvθ-nudging is intended to suppress atmospheric temperature mediated adjustments, by not fixing the land surface temperature adjustment, some adjustments may still take place. For example, in many areas the cooling of the land surface in the SU perturbed experiments is greater than the LTS increase (fig. 3.7). If the land surface temperature were fixed, the LTS increase in these areas would likely be reduced or even become negative in these areas (though fixing the land surface temperature would in turn affect the tropospheric temperature). Changing the LTS adjustment could then significantly change low cloud adjustments. A similar effect (acting in reverse) was noted by Andrews et al. (2021) for CO_2 quadrupling, with the land temperature increase causing a decrease in LTS over land, associated with decreased low cloud over land compared to if land surface temperature is fixed.

Considering specific humidity, the responses are largely similar but opposite for the two

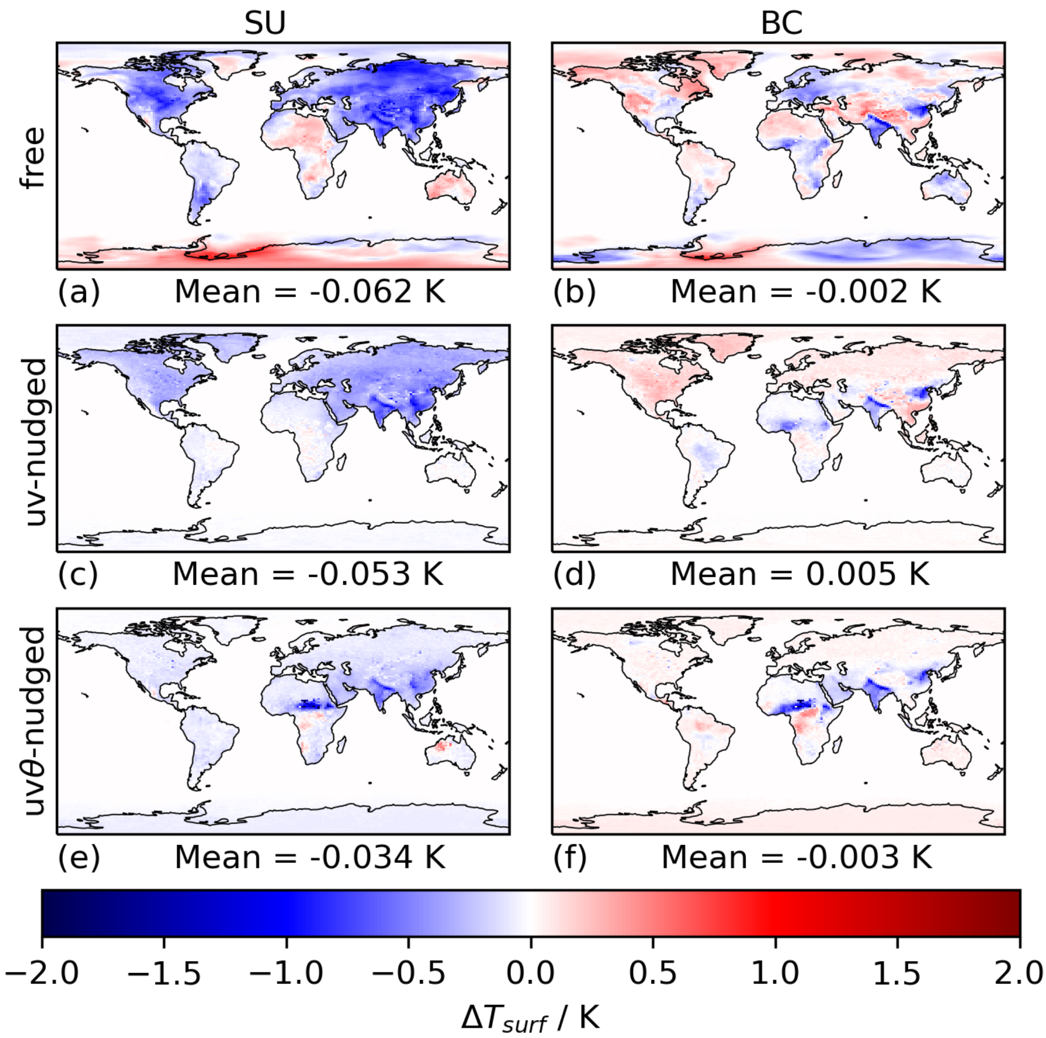

Figure 3.9: Multiannual mean surface temperature adjustment due to 1850 to 2014 SU (a, c, e) or BC (b, d, f) perturbations determined for free (a, b), uv-nudged (c, d), and uvθ-nudged (e, f) simulation pairs. Since SSTs are fixed, only the land surface and sea ice temperature varies.

species in free, uv-nudged and uvθ-nudged cases (fig. 3.10). The free perturbation of SU causes a decrease in specific humidity in the troposphere, as might be expected from the cooling of the atmospheric temperature. The BC perturbation has the opposite effect, increasing specific humidity in response to its warming effect. When the atmospheric temperature response is suppressed in the uvθ-nudged simulation pair the specific humidity response is reduced in the troposphere, being close to zero between 2 to 7 km. However, below $\sim$1 km the adjustment is up to $\sim$30% in the SU case and $\sim$50% in the BC case of the free adjustment, possibly relating to the less effective suppression of θ adjustment close to the surface noted in figure 2.5c.

The uv-nudging also suppresses the specific humidity response significantly. Below $\sim$ 2 km the uv-nudged cases have a similar adjustment in specific humidity to the free case, but above this height, are closer to the uvθ-nudged cases. The uv-nudging therefore contributes more to suppression of the specific humidity adjustment in the free troposphere than the additional nudging of θ, particularly for BC where the uv-nudged and uvθ-nudged cases have very similar positive adjustments. This could be explained in part by the uv-nudging suppressing around half of the free θ adjustment in the free troposphere for SU (fig. 2.5c).

In the stratosphere the specific humidity response is more different for the two aerosols. For SU the weighted adjustment in the free case is bigger than in the troposphere increasing to -6%. The uvθ-nudged SU perturbed case tends to have an adjustment close to zero, while the uv-nudged case suppresses most of the adjustment, except close to the tropopause. For BC, the the free adjustment is also bigger in the stratosphere but less so, only reaching 3% around the tropopause. The uv-nudged case is similar to for SU, suppressing most of the adjustment in the stratosphere, also less so near to the tropopause. The uvθ-nudged BC perturbed case is very different causing a peak in adjustment at ~22 km rather than near the tropopause, and a stronger adjustment than the free case, reaching 5%.

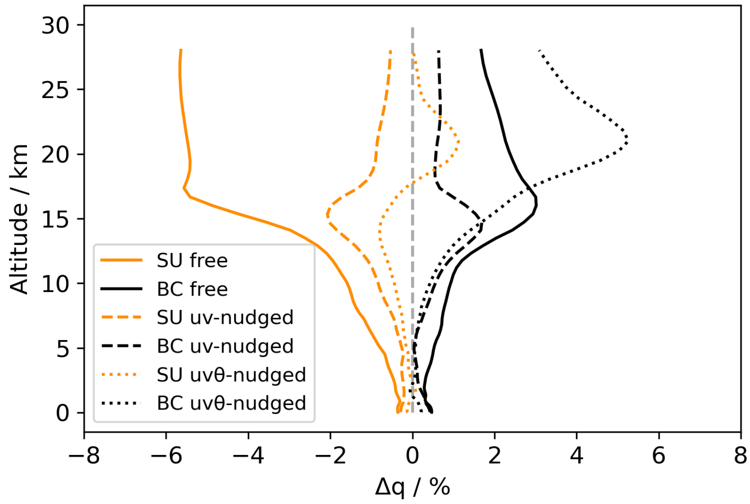

Figure 3.10: Multiannual, global model level mean adjustment in specific humidity, q, due to 1850 to 2014 SU or BC perturbations determined for free, uv-nudged and uvθ-nudged simulation pairs. Adjustment is given as a percentage of the multiannual global model level mean specific humidity in the free control.

3.6 Discussion

From the results it is clear that the primary component of circulation- and temperature-mediated radiative adjustments is via cloud adjustment. The CRE is the largest forcing component of both the circulation adjustment and temperature adjustment for both aerosols. However, there are significant contributions to the temperature adjustment forcing from residual clear-clean sky adjustments. It is perhaps unsurprising that the CRE is the most significant component given the significant contribution cloud adjustments have been found to make to aerosol ERFs in other studies (Smith et al., 2018; Bellouin et al., 2020a; Smith et al., 2020a).

Of the two nudging types, uvθ-nudging clearly has a larger effect on ERF than uv-nudging. This is associated with strong changes in low cloud, particularly over marine stratocumulus regions and the storm tracks, which are largely absent from the uv-nudged and free simulation pair adjustments. For both aerosols, the resulting temperature mediated adjustment forcing (negative for BC and positive for SU) opposes the overall ERF. This result agrees with previous studies that find that BC causes a negative forcing via semi-direct effects (Samset and Myhre, 2015; Sand et al., 2015; Thornhill et al., 2021). However, it is not clear that the semi-direct effect causes this cloud response in the uvθ-nudged BC perturbed case, and therefore whether the response is representative of the real effect of temperature adjustments arising from aerosol forcing, or an artefact introduced by nudging. Given the similarity (albeit of opposite sign) of the low cloud response in the SU perturbed case, the low cloud response in the SU case may also be an anomaly due to nudging. Consequently the temperature mediated radiative adjustment should be interpreted with caution for both aerosols.

While smaller, the effect of uv-nudging is found to be significant for SU, and is more readily understood than the effects of uvθ-nudging. The uv-nudging effect on radiative forcing occurs primarily via cloud adjustments. In the SU perturbed case this includes a decrease in low cloud over northern hemisphere land, and for BC, an increase in high cloud, relative to the free simulation pair adjustments. It should be highlighted that even though the forcing response is smaller than the uvθ nudged case, it is clear there are significant effects of nudging the horizontal winds. Aside from the cloud changes, for example, at most altitudes above the free troposphere the majority of specific humidity adjustment is suppressed by the uv-nudging.

In terms of the effectiveness of nudging as a mechanism denial methodology, there are a number of limitations that can be identified from this and the previous chapter. One significant issue is that the suppression of adjustment was imperfect: while the horizontal wind adjustment could be largely suppressed by nudging, nudging the potential temperature as well was less effective. Consequently the adjustment identified cannot be confidently said to include the total adjustment response for the temperature adjustment.

The other most significant issue is the possibility of unintended effects of nudging that are difficult to identify a physical explanation for. This particularly includes some of the cloud responses. For example, the control error in cloud fields illustrated in figure 2.8 demonstrates that even where nudging input is generated from an identical simulation to that being nudged (other than not being nudged), it can introduce responses in the cloud fields that are as large as the signal of the low cloud adjustments under uvθ-nudging. Also, the marine stratocumulus adjustments in the uvθ-nudged simulation pairs are similar to patterns of control error observed by Zhang et al. (2022) that they found were resolved by altering the sequence of nudging in the model timestep.

Unexpected or counter-intuitive responses from nudging have been identified in other studies as well. For example, Gettelman et al. (2020) found that nudging could reduce the bias of temperature and boundary layer structure against observations whilst simultaneously increasing bias in model clouds. They suggested this response may be due to the nudging disrupting links between microphysical and turbulent processes. While that study was concerned with the Southern Ocean specifically and used reanalysis data for the nudging input, it seems possible such links could be disrupted elsewhere even when nudging to data from the same model. This can cast doubt on whether even significant differences between the different nudged simulations have physical interpretations, or are caused by spurious effects of nudging.

To make the nudging mechanism denial methodology used here more useful requires addressing some of these limitations. One potential improvement could include fixing the surface temperature. This is because while nudging the lower model levels helped to increase suppression of the adjustment of horizontal wind and potential temperature fields in lowermost 1 km, the potential temperature adjustment was still suppressed by less than 50 % over these levels with SU perturbation, and less than 30 % with BC perturbation. This could be partly due to the land surface temperature being able to adjust freely, affecting the atmospheric levels close above. Future studies using the nudging mechanism denial method to investigate atmospheric temperature-mediated adjustments as done here could therefore try fixing the surface temperature, such as by using the method of Andrews et al. (2021), or possibly through implementing nudging of the surface temperature field. This would hopefully fix some undesirable effects of fixing SSTs but not land surface temperatures.

Aside from fixing the land surface temperature, future work with this methodology could include adjusting the values of G and the time frequency of the nudging input. While the 1 year tests using higher time resolutions described in 2.6.7 found no overall improvements to θ adjustment suppression and control error, a more detailed study by Zhang et al. (2022) found significant reductions in bias between nudging input and nudged simulations when using the EAMv1 model with higher time resolution nudging input. While the benefits may be model dependent, further studies using nudging as a mechanism denial tool could further explore increasing the time resolution of nudging input used. Furthermore, while nudging settings are often tuned by trial and error as done here, Omrani et al. (2012) demonstrate the possibility of determining the optimum values of G through understanding physical processes and the time over which internal variation will cause a variable to diverge from its target value, rather than trial and error. Future studies could seek to apply this approach to the use of nudging as a mechanism denial tool.

Finally, a more general point that can be taken from these results is that nudging the θ field will affect ERFs determined from UKESM1, particularly due to cloud responses but also clear-clean sky responses too. This agrees with the findings of several previous studies using other models that found nudging temperature or similar quantities could affect radiative forcing and cloud representation in particular (Kooperman et al., 2012; Zhang et al., 2014; Sun et al., 2019). Furthermore, this confirms the recommendations of Forster et al. (2016) and Zhang et al. (2014) that temperature should not be nudged when nudging is only being used to reduce natural variability so that integrations can be shorter. However, while both this study and that of Johnson et al. (2019) found the circulation-mediated radiative adjustment forcing for uv-nudging to be insignificant for BC perturbation, the SU circulation-mediated radiative adjustment forcing was found to be significant here. Furthermore, for both aerosols we found considerable differences in cloud, humidity and temperature fields under just uv-nudging compared to the free simulations. Therefore, where determining the ERF of some forcing agent is of interest, even nudging just the horizontal winds to suppress natural variability should be done with caution. This may require even greater care if regional, not just global, adjustments are being investigated.

3.7 Conclusions

This chapter investigated the effects of nudging on BC and SU aerosol radiative adjustments and its potential as a tool for isolating circulation and atmospheric temperature mediated adjustments and their associated radiative forcing. It was demonstrated that nudging horizontal winds resulted in a small radiative forcing difference that was found to be significant in the SU case. For both aerosols nudging the horizontal winds caused considerable changes in cloud, specific humidity, and temperature fields. This suggests that while the effects of nudging horizontal winds may be small, it can have considerable effects, particularly locally, suggesting circulation adjustments due to aerosol perturbations can be significant. The differences due to additional temperature nudging were even greater, strongly affecting marine stratocumulus cloud and radiative forcing, suggesting a large effect of atmospheric temperature mediated adjustments. However, uncertainty around the mechanisms by which temperature nudging affected cloud fields casts some doubt on whether these differences are associated with physical atmospheric temperature mediated adjustments or are solely unintended effects of nudging.

Indeed the study demonstrated that the methodology used has a number of limitations that affect its utility. These particularly include: achieving the effective suppression of adjustments by nudging, the introduction of unintended model responses not representative of physical processes by nudging, and the sensitivity of the method to the nudging settings chosen. These issues are a significant problem when nudging both the horizontal winds and potential temperature, and without addressing the first two issues, this method cannot be considered to give reliable information on temperature mediated adjustments. When nudging just horizontal winds, the horizontal wind adjustment suppression was considerably better and fewer artefacts were introduced by the nudging, so at least for global means and broad patterns this method may give useful insights into aerosol circulation adjustments. However for investigating finer scale adjustments in circulation, this method is unlikely to be reliable without further refinements. Such refinements should include testing the use of higher time resolution nudging input and implementing the methodology in different models.

Finally, the study demonstrated that nudging can affect ERF determined from timeslice experiments, even when nudging to model output to avoid biases with reanalysis models and only nudging horizontal winds. This suggests that caution must be used when using nudging to constrain internal variability to reduce the number of model years that must be integrated over for determination of radiative forcing estimates from perturbation studies.

Chapter 4

Investigating aerosol radiative adjustments using a variation of the partial radiative perturbation technique with the SOCRATES radiation code

4.1 Introduction

In this chapter the application of PRP calculations to aerosol radiative adjustments is investigated. This work uses outputs from the free control and BC perturbed UKESM simulations used in chapters 2 and 3 to run PRP calculations using the offline version of the UKESM radiation code to determine radiative forcing due to 1850 to 2014 BC perturbation and resulting radiative adjustments. This chapter presents details of the application of the PRP method, including a number of issues encountered, as well as estimates of radiative adjustments determined for a number of variables under BC perturbation from a one year PRP integration. See section 1.3.3 for an overview of the application of the PRP method.

Using PRP compliments the use of nudging in that PRP allows the separation of radiative forcing due to adjustment of specific variables, while nudging separates adjustment mechanisms. The decision to use PRP over a radiative kernel approach was primarily because this work sought to use a radiation code for calculation of radiative adjustments that is close to identical to the online radiation code, which is not possible using an existing radiative kernels. This replication of the online radiation code included using 1 hourly input to the offline radiation code to test assumptions about linearity of radiative adjustments and to directly represent the cloud adjustment with the PRP method. The intention is also to investigate a variation of the PRP method that may be more efficient and produces a sum of adjustment terms that is more consistent with the total radiative adjustment forcing.

This chapter begins with a breakdown of the PRP experiments conducted and the methodology used for them, including a description of the alternative PRP methodology employed here. Details of the offline radiation code settings follow. After this, a description of the code developed and platform specific methodology to perform the PRP experiments is given, primarily in the interest of making this work more easily replicable. Next is given results from testing differences between output from the offline and online radiation codes. This includes an account of bug fixes and other modifications that were made to the offline radiation code or the output from the UKESM simulations to overcome any differences between the offline and online codes. Next the results of the one year PRP integration for BC perturbation and resulting radiative adjustments are given. Then a discussion of the key results is given, and finally, a conclusion of the key findings from this work.

4.2 Methodology of PRP calculations

PRP calculations are performed for a number of radiatively active variables that adjust in response to a PI (1850) to PD (2014) anthropogenic BC emissions perturbation. The input for these PRP calculations is derived from the free fSST UKESM control and BC perturbed simulations described in 2.3, with added output diagnostics. The PRP calculations use the Suite Of Community RAdiative Transfer codes based on Edwards and Slingo (SOCRATES) offline radiation code (Manners et al. 2015). SOCRATES is also used for the online radiation code in the UM, though there are a number of differences between the offline and online versions. SOCRATES was chosen for the PRP simulations to ensure that the PRP calculations represented the online model radiative responses as closely as possible. As such, the settings used in SOCRATES were made identical to those in the online radiation code in the UKESM simulations that generated the input where possible, though in some cases this could not be achieved with the existing offline SOCRATES code.

PRP calculations are performed for a range of variables. Specifically, PRP calculations were used to determine the radiative effect of changes in aerosol, surface albedo, atmospheric temperature, surface temperature, specific humidity, and 'other' variables. Since BC emissions are perturbed, the radiative effect of aerosol change can be considered to be approximately the IRF, though strictly speaking this PRP term will include radiative effects of adjustments to any other aerosols, not just the IRF from BC perturbation. The radiative effect due to surface temperature response is technically a radiative feedback, though it should be re-iterated that only land surface temperature is allowed to respond. As with other studies (e.g. Smith et al. (2018)) the radiative response from surface temperature is included in the total ERF given. The remaining 4 PRP terms are all radiative adjustments. The 'other' PRP term includes the effect of adjustments to all other input variables to SOCRATES that were not fixed between control and perturbed UKESM simulations. These are specifically ozone (which used an interactive 3D field in the UKESM simulations) and pressure.

As with nudging, one can think of PRP in the context of the radiative processes illustrated in figure 1.4. Figure 4.1 outlines these interactions, noting where PRP is used to calculate the radiative effects of a change in a component of the climate system. As can be seen, unlike nudging, PRP was not used to suppress the adjustment of variables themselves but rather to calculate the direct radiative effects of each variable.

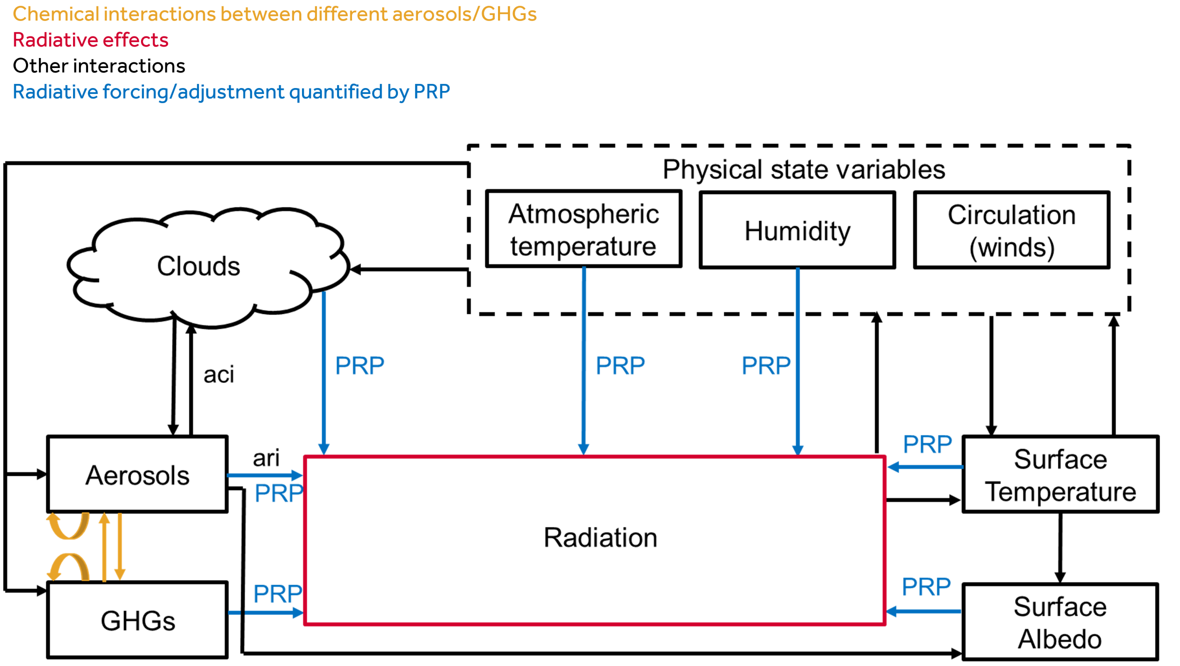

Figure 4.1: Illustration of interactions between aerosols and greenhouse gases (GHGs) and different components of the climate system that affect radiation, noting where the PRP method calculates radiative effects of changes in different components. After Boucher et al. (2013).

The initial intention was to perform all sky PRP calculations and also include a cloud PRP term. However, before conducting the PRP calculations, tests of the representation of TOA radiative fluxes from the control (i.e. PD) UKESM simulation in SOCRATES were conducted for all sky, clear sky, clean sky and clear-clean sky cases. These tests are described in section 4.5 and highlighted difficulties in including the online model cloud fields in offline SOCRATES accurately. These issues were unable to be resolved despite a number of fixes being implemented and following extensive consultation with Met Office colleagues. Consequently the PRP calculations presented here were performed under clear sky conditions and no cloud PRP term is calculated.

The input to SOCRATES from the UKESM simulations is 1 hourly for one model year. The input frequency matches the time resolution of the online radiation code. While this input resolution is higher than that used in other studies, the high resolution was chosen to best represent cloud radiative effect (see section 1.3.3). Only a one integration was used for the PRP calculations due to the computational and time limitations when using 1 hourly input data (see section 4.4).

4.2.1 Cumulative PRP method

When the IRF and adjustments are calculated individually using PRP (i.e. one variable is perturbed, with all other variables held at their control values) as is done in the conventional PRP method, the sum of the terms is not necessarily equal to the whole ERF. Therefore, a different implementation of the PRP method is used here. In this method the perturbations are conducted 'cumulatively'. More specifically, the first PRP term is calculated as the difference between the net TOA fluxes when all variables are set to their values in the control simulation minus the net TOA fluxes with one variable (e.g. aerosol) set to its value in the perturbed simulation. The second PRP term is calculated as the difference between net TOA fluxes with the latter state (i.e. aerosol set to the perturbed simulation values and all other variables set to control simulation values) minus net TOA fluxes with an additional variable set to its values in the perturbed simulation (e.g. aerosol and water vapour are set to perturbed values and all others set to control values). Subsequent PRP terms follow this pattern, with the sum of each adding up to the total clear sky radiative forcing.

To overcome the decorrelation bias described in section 1.3.3, both a forward and backward calculation is performed for each variable and the mean taken to give a two-sided calculation for each PRP term. This effectively results in two sets of cumulative PRP terms, one set beginning with all variables set to control simulation values and changing them to perturbed values cumulatively; the other set beginning with all variables set to perturbed simulation values and changing them (in the same order, not reverse order) to control simulation values.

Algebraically this method can be described as follows. Firstly, the calculation for the forward aerosol perturbation can be written as:

$$\Delta F_{A,f,\mathrm{clr}} = R(A_p, q_c, \alpha_c, T_c, T^*_c, o_c) - R(A_c, q_c, \alpha_c, T_c, T^*_c, o_c) \tag{4.1}$$

where A is aerosol, q is specific humidity, α is surface albedo, T is atmospheric temperature, T^* is surface temperature, o is the 'other' variables; $\Delta F_{A,f,\mathrm{clr}}$ is the clear sky (subscript 'clr') forcing due to aerosol perturbation (subscript 'A') in the forward case (subscript 'f'); and R is the TOA net downwards flux and is a function of A, q, α, T, T^*, and o. This is effectively the usual form of a forward PRP calculation. However the forward calculation for the next PRP term, specific humidity, is:

$$\Delta F_{q,f,\mathrm{clr}} = R(A_p, q_p, \alpha_c, T_c, T^*_c, o_c) - R(A_p, q_c, \alpha_c, T_c, T^*_c, o_c) \tag{4.2}$$

Thus the perturbation of specific humidity is performed with aerosol set to the perturbed simulation values and control simulation values of all other variables, unlike the typical forward calculation where all variables would be set to their control values except for specific humidity. It should be noted that choosing water vapour as the next PRP term is mostly arbitrary, as is the order of the remaining PRP terms (a point that is explained in more detail later in this section). The forward calculations for the remaining variables (α, T, T^* and other) are respectively:

$$\Delta F_{\alpha,f,\mathrm{clr}} = R(A_p, q_p, \alpha_p, T_c, T^*_c, o_c) - R(A_p, q_p, \alpha_c, T_c, T^*_c, o_c) \tag{4.3}$$

$$\Delta F_{T,f,\mathrm{clr}} = R(A_p, q_p, \alpha_p, T_p, T_c^*, o_c) - R(A_p, q_p, \alpha_p, T_c, T_c^*, o_c) \tag{4.4}$$

$$\Delta F_{T^*,f,\mathrm{clr}} = R(A_p, q_p, \alpha_p, T_p, T_p^*, o_c) - R(A_p, q_p, \alpha_p, T_p, T_c^*, o_c) \tag{4.5}$$

$$\Delta F_{o,f,\mathrm{clr}} = R(A_p, q_p, \alpha_p, T_p, T_p^*, o_p) - R(A_p, q_p, \alpha_p, T_p, T_p^*, o_c) \tag{4.6}$$

Backward calculations for each variable are then performed similarly. Again the backwards calculation for the aerosol perturbation is like the usual form, but for the remaining variables this differs to the conventional backwards calculation. The backward calculations are as follows:

$$\Delta F_{A,b,\mathrm{clr}} = R(A_p, q_p, \alpha_p, T_p, T_p^*, o_p) - R(A_c, q_p, \alpha_p, T_p, T_p^*, o_p) \tag{4.7}$$

$$\Delta F_{q,b,\mathrm{clr}} = R(A_c, q_p, \alpha_p, T_p, T_p^*, o_p) - R(A_c, q_c, \alpha_p, T_p, T_p^*, o_p) \tag{4.8}$$

$$\Delta F_{\alpha,b,\mathrm{clr}} = R(A_c, q_c, \alpha_p, T_p, T_p^*, o_p) - R(A_c, q_c, \alpha_c, T_p, T_p^*, o_p) \tag{4.9}$$

$$\Delta F_{T,b,\mathrm{clr}} = R(A_c, q_c, \alpha_c, T_p, T_p^*, o_p) - R(A_c, q_c, \alpha_c, T_c, T_p^*, o_p) \tag{4.10}$$

$$\Delta F_{T^*,b,\mathrm{clr}} = R(A_c, q_c, \alpha_c, T_c, T_p^*, o_p) - R(A_c, q_c, \alpha_c, T_c, T_c^*, o_p) \tag{4.11}$$

$$\Delta F_{o,b,\mathrm{clr}} = R(A_c, q_c, \alpha_c, T_c, T_c^*, o_p) - R(A_c, q_c, \alpha_c, T_c, T_c^*, o_c) \tag{4.12}$$

As noted earlier, using this cumulative PRP method ensures that the sum of the PRP terms adds up to the total clear sky radiative forcing, that is to say that:

$$\Delta F_{\mathrm{clr}} = F_{A,\mathrm{clr}} + F_{q,\mathrm{clr}} + F_{\alpha,\mathrm{clr}} + F_{T,\mathrm{clr}} + F_{T*,\mathrm{clr}} + F_{o,\mathrm{clr}} \tag{4.13}$$

where ΔF_{clr} is the total clear sky forcing due to BC perturbation. Equation 4.13 is true for both the sum of forward and sum of backward terms, and consequently, the sum of two-sided PRP terms also. To verify this, the total clear sky radiative forcing can be calculated independently with SOCRATES as the difference in net TOA fluxes with all variables set to control or perturbed values:

$$\Delta F_{\mathrm{clr}} = R(A_p, q_p, \alpha_p, T_p, T_p^*, r_p) - R(A_c, q_c, \alpha_c, T_c, T_c^*, r_c) \tag{4.14}$$

Equation 4.14 uses calculations already made for the PRP terms, and can also be used to check whether the SOCRATES code reproduces the total radiative forcing calculated from UKESM diagnostics.

As well as its conservative nature, there is also a small computational advantage of using this cumulative PRP method: one PRP term does not require any additional calculations of the radiation code to compute it. For example, for calculation of the forward other PRP term in equation 4.6 the first term is also required in the aerosol backward PRP term and the second term is also required in the surface temperature forward PRP term. By contrast if the forward other PRP term was to be determined with the conventional PRP method, an extra run of the radiation code would be required in which only the other variables are perturbed (and similar for the backward other PRP term). Consequently this method saves two runs of the radiation code (one forward, one backward) compared to the conventional PRP method (irrespective of the number of PRP terms calculated). A caveat worth noting is that the computational saving only applies where PRP terms account for all radiatively active variables that are able to adjust between the control and perturbed simulations. This will be true for many studies using PRP, but for example, if PRP was used to calculate only the aerosol and water vapour adjustments, then it would require the same number of calculations using either conventional or cumulative PRP methods.

It should be clarified that, as alluded to earlier, the forward and backward calculations for each term using this cumulative PRP method are expected to differ from the respective forward and backward calculations using the conventional PRP method. For example, in the forward water vapour PRP term (and the same logic can be applied for the backward term), the water vapour radiative adjustment is calculated on a background of all variables set to control states, except aerosol which is set to its perturbed state; whereas in the conventional method aerosol would also be set to its control state. This difference in background variables is likely to result in small differences in the masking effect of aerosol on water vapour adjustment. As such there should be a small difference between the forward water vapour PRP term calculated using the cumulative and conventional PRP methods. However, there is no reason for either method's forward water vapour term to be more or less 'correct' than the other; using either control or perturbed states for the other variables is arguably arbitrary because in the atmosphere all adjustments happen simultaneously anyway. Crucially though, while the forward term of the cumulative method perturbs water vapour on a background of perturbed aerosol, its backward term perturbs water vapour on a background of control aerosol (and the conventional PRP method does the opposite). As such, although it is not tested here, while the individual forward or backward terms are likely to differ between the two methods, it is expected that the two sided (mean) PRP for each variable should be approximately the same between the two methods. As an extension of this reasoning, the order in which the variables are perturbed in the cumulative PRP method is likewise arbitrary: a different order will result in differences to the forward and backward terms, but the two-sided PRP should change very little.

Finally, it must be noted that the equations above are written following the general convention in the literature that one is calculating the radiative effect of going from the control to the perturbed state. However in this work, by defining the perturbed simulation as having 1850 emissions and control simulation as having 2014 emissions of BC, these equations would return PRP terms corresponding to changing the emissions of BC from 2014 to 1850. This would be the reverse of the rest of the thesis and most comparable literature, where results are given as the radiative effect of present day minus pre-industrial emissions of aerosols. Therefore for consistency, the results presented from the PRP calculations using the above equations are multiplied by -1 (i.e. simply representing the radiative effect of going from 1850 to 2014 BC emissions, or from perturbed to control states).

4.3 Details of the SOCRATES radiative transfer codes

SOCRATES was run with a two-stream approximation, in which radiation is represented by upward and downward diffuse fluxes as well as a direct downward solar flux in the SW bands. Its broadband configuration was used, in which there are 6 SW and 9 LW bands. The spectral file used was that from the Met Office's Global Atmosphere Configuration 7 (Walters et al., 2019).

For cloud representation, SOCRATES (as in the online UKESM simulations) was set to use a Monte Carlo Independent Column Approximation scheme for representation of cloudy sky. Ice and water clouds were treated treated separately rather than homogeneously.

To run SOCRATES requires a number of general input variables. These include: temperature on model levels and layers; surface temperature; pressure on levels and layers; SW and LW surface albedo (the latter being 1 minus the surface LW emissivity); incoming TOA SW flux; the solar zenith angle; specific humidity; and greenhouse gas concentrations. Of the greenhouse gases only the ozone values come from an interactively determined field in the UKESM simulations. CO_2, O_2, CH_4, N_2O, CFC-12-eq and HCFC134a-eq (the latter two being equivalent concentrations for a number of halocarbons; see Meinshausen et al. (2017)) used fixed uniform values. For aerosol (specifically for using GLOMAP mode aerosol rather than the older CLASSIC scheme in the UM), additional inputs required were: model layer depth; aerosol optical depth (AOD) in SW and LW bands; aerosol scattering in SW and LW bands; and aerosol asymmetry factor in SW and LW bands. Finally for clouds, additional inputs were: total, liquid and ice cloud fractions; liquid and ice water mass mixing ratios; and liquid effective radius. Ice crystal effective radius was not required since the UKESM simulations used a parametrisation based on temperature and ice mass mixing ratio.

4.4 Code development

The SOCRATES code can be installed and run on a number of platforms. For this work SOCRATES was run on the Monsoon2 computing platform. SOCRATES has inbuilt tests that can be run to ensure installation is performed correctly. In this case these were satisfactory, generating slight discrepancies all at less than six significant figures.

To generate input for the SOCRATES PRP calculations the free control and BC perturbed simulations described in chapter 2 were modified to add the additional diagnostics required for SOCRATES, at 1 hourly frequency. SOCRATES requires the input variables to be pre-processed into a NetCDF format. A python module allowing conversion of the variables into the appropriate format is contained in the SOCRATES code distribution but requires additional steps.

Additional python code was written to achieve this pre-processing step working from code supplied by Tom Dunstan (Met Office).

To efficiently run SOCRATES across multiple timesteps, the preprocessing and SOCRATES calculations were performed using Rose suites (https://github.com/metomi/rose/). These allow for automated cycling of preprocessing and SOCRATES run tasks. The overall process was split up into a preprocessing suite and a 'run' suite: the preprocessing suite loaded, preprocessed, and archived the input data; while the run suite executed SOCRATES with the input. Two variations on the run suite were made: a 'test' run suite that used input from only the control UKESM simulation to compare the TOA fluxes output from SOCRATES against online model diagnostics (see section 4.5); and a 'PRP' run suite that used inputs from both the control and BC perturbed UKESM simulations to perform PRP calculations (see section 4.6). These steps are illustrated in figure 4.2.

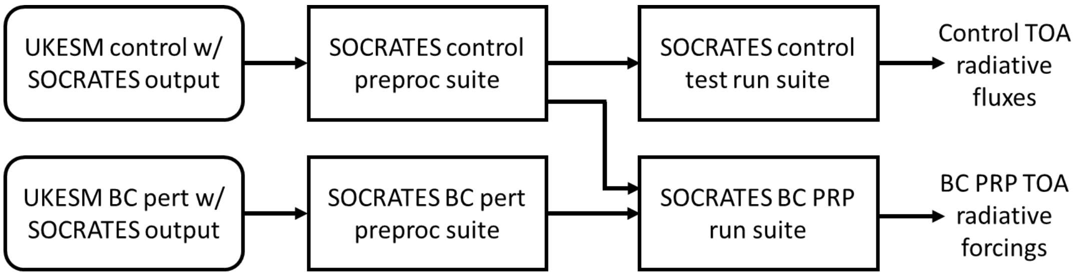

Figure 4.2: Schematic representation of the experimental steps in this chapter.

The preprocessing and SOCRATES calculations were split into separate Rose suites because the input data for each timestep was large ($\sim$ 600 MB), causing the preprocessing step to require a significant computing time itself. Specifically, using hourly input the preprocessing suites takes $\sim$2 days to preprocess 1 model year of input data and the run suite takes 10 to 15 days to do a 1 year integration. Furthermore, by running the preprocessing as a separate suite, the PRP calculations for different aerosols can be conducted without repeating the preprocessing for the control simulation. Likewise it is possible to re-run just the run suite for any repeat SOCRATES calculations, without needing to re-preprocess the input data.

4.5 SOCRATES testing and code fixes

Here are presented the results from test experiments, in which radiative fluxes calculated from SOCRATES using just input from the control UKESM simulation are compared with radiative flux diagnostics output from UKESM. These tests were conducted to test whether that the online model radiation code could be accurately reproduced offline with SOCRATES. As such, ideally the difference between SOCRATES and online radiative fluxes is zero. However, due to a number of issues encountered and differences between SOCRATES and the online radiation code, this was not always the case. These differences, and the fixes developed in this work to address them, are detailed here.

The issues encountered and corresponding fixes can be broken down by input variables, with the primary issues affecting: TOA incoming SW flux and solar zenith angle; aerosols; clouds; and surface emissivity. These are dealt with in turn in the following subsections, as well as one subsection outlining remaining differences that could not be resolved in the time available. It should be noted that for the results in section 4.6 all of the fixes described below (except the cloud fixes) were applied in the calculations and the unresolved differences remained.

4.5.1 Incoming solar flux and solar zenith angle

An early issue identified was that the incoming TOA flux was significantly biased in SOCRATES calculations against the online model. This bias followed the pattern of solar insolation at a given time and was due to the incoming TOA flux diagnostic being used as input to SOCRATES already including the correction for solar zenith angle. The correction for solar zenith angle is applied again in the SOCRATES code. To address this a reverse correction (inversely weighting by the cosine of solar zenith angle) is applied to the incoming solar flux diagnostic during preprocessing.

A related issue identified was that the incoming solar flux diagnostic was out of phase with the solar zenith angle. In particular, the area of Earth's surface illuminated according to the solar zenith angle (i.e. angle less than 90°) was larger than the illuminated area according to the incoming solar flux (i.e. incoming flux greater than 0 W m^{-2}) and appeared to be shifted ∼5° eastwards. This difference resulted from different timestamps for different diagnostics when output from UKESM beyond the specified time profile of output (i.e. two diagnostics with the same time profile do not necessarily represent the exact same time profile). Specifically, the incoming solar flux diagnostic represents a mean value over a single dynamical timestep, whereas the solar zenith angle diagnostic (as well as other radiation code specific diagnostics) are averaged over a (1 hour) radiation timestep. As such, the incoming solar flux was instead taken from the top level of a 3D flux diagnostic that is output on only radiation timesteps to address this issue.

4.5.2 Aerosols

Despite the fixes to incoming solar flux and solar zenith angle diagnostics, large biases and 'not a number' errors (NaNs) were observed in both all and clear sky fluxes calculated from offline SOCRATES relative to the online diagnostics. The majority of the bias and the NaNs did not persist in the clear-clean sky calculations, suggesting there were errors in the offline SOCRATES calculations involving aerosols.

Two code errors were thus identified with the input and representation of aerosol in SOCRATES. The first issue was a bug identified in the python module for preprocessing inputs into NetCDF. This bug effectively copied the values for aerosol absorption extinction into the scattering extinction and aerosol asymmetry parameter. As such I updated the python preprocessing module to include a fix for this bug and submitted a branch to be included in the next SOCRATES code release.

The second error was that the 3D aerosol input fields were interpolated in SOCRATES to pressure levels supplied as an input file. However, the pressure levels used were only for a single arbitrary column rather than a 3D field. With guidance from James Manners (Met Office), I added a change to the SOCRATES code to omit the interpolation and copy the aerosol input values directly without interpolation (since the input is from a UM run, there is no need to adjust inputs between different vertical levels). This removed the bias present in the clear sky case that was not present in the clear-clean sky case and the NaNs, suggesting that identical treatment of aerosol in offline SOCRATES and the online radiation code was achieved.

4.5.3 Clouds

While fixing the aerosol issues reduced the bias and number of NaNs, in the all sky fluxes a large bias and some NaNs still remained. This pointed to clouds also having differing representation between the offline SOCRATES and online radiation codes.

To explain the NaNs issue, the liquid cloud droplet single scattering albedo is parametrised in the radiation code using the liquid droplet effective radius. However the droplet effective radius on some grid points falls outside of the range of the parametrisation (sometimes being zero) even though cloud was present at those points. In the online model the droplet effective radius is limited to within the window of the parametrisation; however these limits were not applied in the offline SOCRATES code. Adding these limits to the offline SOCRATES code was done with guidance from James Manners, who has added the changes to the SOCRATES code base. This fixed all remaining NaNs.

Unlike the NaNs it was not possible in the time available to resolve the cloud bias. While some issues with cloud representation were resolved, a large bias remained and the cause(s) of the remaining bias are uncertain. The remaining bias in the all sky net TOA flux is illustrated in figure 4.3 for a one day mean. In the SW there is no clear spatial pattern to the bias, aside from being smaller at higher latitudes (presumably due to less insolation at higher latitudes) and over areas with typically little cloud cover (e.g. the Sahara). There is also a mixed signal of positive and negative biases, with bias magnitudes of up to 30 W m^{-2}, suggesting there is no systematic under- or overestimation of cloud scattering effects. Similar observations can be made for absorption in the LW, though there is less dependence on latitude and the magnitude of the bias is smaller than in the SW. The fact that both SW and LW fluxes are biased suggests the cause is not just due to a difference in representation of either scattering or absorption alone, but more likely in the representation of the cloud cover, liquid water path, or droplet effective radius. Given the magnitude of the bias due to clouds and the lack of a solution, it was decided that the actual PRP calculations should be performed under clear sky conditions only, since the bias in cloud was likely to bias the effects of radiative adjustments.

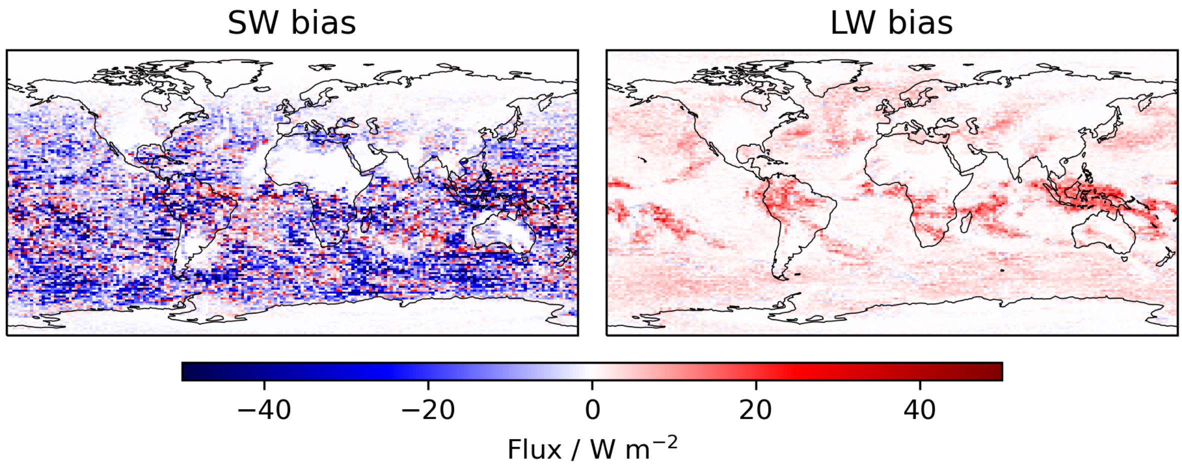

Figure 4.3: Difference in net downwards TOA fluxes under all sky conditions between SOCRATES and online radiation code for a mean over a one day mean for SW bands and LW bands.

4.5.4 Surface emissivity

Across calculations under all, clear, clean and clear-clean sky conditions there existed a persistent bias in the LW net downwards TOA fluxes in SOCRATES. As illustrated in figure 4.4a, there were particularly strong negative biases relative to the online fluxes over areas of desert. This arose from use of a single surface LW emissivity value (0.98) across all grid points in the initial SOCRATES calculations, whereas the online UKESM simulations vary LW emissivity by surface type. Bare soil (including deserts) in particular has a lower LW emissivity of 0.9, resulting in the strong negative bias over these surfaces which were emitting too much LW radiation.

The single uniform value of LW emissivity was initially used due to no LW emissivity diagnostic being present by default in the UKESM simulations. Adding a diagnostic to output the surface LW emissivity under direction from James Manners provides a mostly satisfactory fix: as can be seen in figure 4.4b, the LW net downwards TOA fluxes have a much reduced bias.

4.5.5 Remaining biases

Some differences between fluxes calculated with SOCRATES and the online radiation code remain. Aside from the unresolved cloud bias, a number of differences still remain. As noted before, the biases are identical between clear and clear-clean conditions, and so are not related to aerosols. The remaining biases are illustrated in figure 4.4b and 4.4c.

The main unresolved clear-clean bias in the LW is around coastlines, being a generally positive bias up to 0.5 Wm^{-2}, though along some coastlines the bias is negative, as illustrated in figure 4.4b. The exact cause of this bias is uncertain, but is possibly a result of different treatment of surface temperatures on grid points containing land and sea. Specifically, it may be the absence of a switch allowing for quadratic averaging of surface temperature at coastal points,

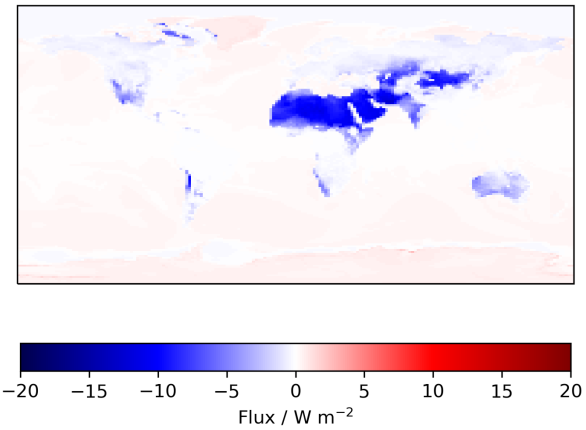

(a) LW: uniform emissivity

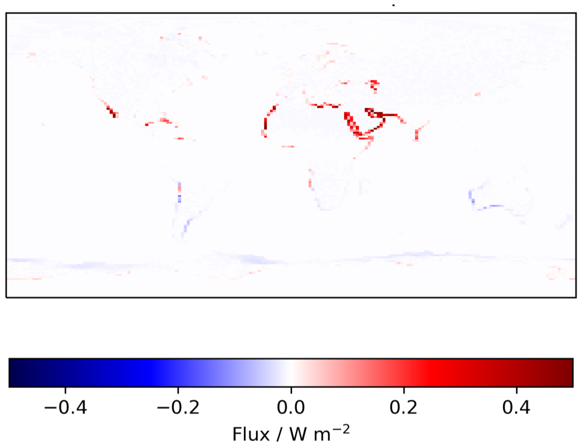

(b) LW: surface specific emissivity

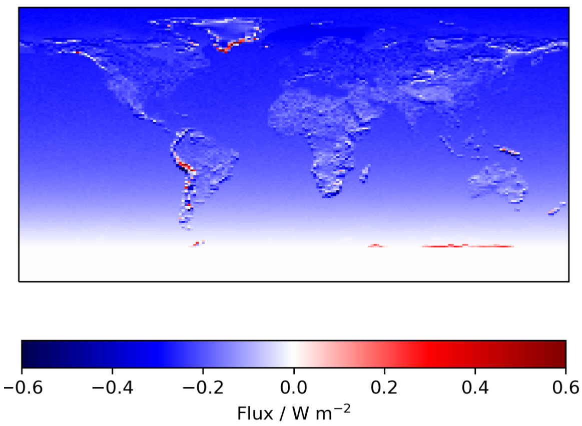

(c) SW

Figure 4.4: Difference in net downwards TOA fluxes under clear-clean sky conditions between SOCRATES and online radiation code for a mean over one day (specifically from late May). (a,b) illustrate the bias in the LW and (c) the bias in the SW. (a) illustrates the LW flux bias when using a single uniform value for LW emissivity in SOCRATES while in (b) the same emissivity field from the online simulations is used. Note the different scale bars.

which is used in the online simulations. Whatever the cause, the bias is more strongly positive over sub-tropical northern latitudes, with negative values at mid and high southern latitudes suggesting a link to insolation strength given the time of year in figure 4.4b (late May).

In the SW there are two clear patterns of clear-clean bias remaining, illustrated in figure 4.4c. The first is a negative bias that varies with latitude, being stronger at northern latitudes, up to -0.6 W m^{-2}. The variation with latitude is likely to be related to the solar zenith angle and associated varying insolation at different latitudes. The mechanism of this bias is uncertain. The other pattern is of stronger positive or negative biases near mountain ranges. Generally the bias is more negative on the shaded side of mountains (i.e north side at mid and high northern latitudes and south at southern or low northern latitudes) and positive or close to zero on the other side. The bias is also generally stronger at higher latitudes and weaker over sub-tropical northern latitudes. This bias is likely due to the absence of a correction for dependence of insolation on surface gradient in the offline SOCRATES code, but is included in the online model.

From figures 4.4b and 4.4c it can be seen that these remaining biases in the TOA fluxes calculated from SOCRATES can be up to approximately 0.5 W m^{-2} locally. While this is a potentially large bias, it should be noted that they should have no significant effect on radiative forcing calculations. This is because the causes of the bias are systematic and unlikely to vary with aerosol perturbation, and so should mostly cancel when differencing two TOA flux calculations from SOCRATES.

While it was not possible to resolve all of the issues identified in this section, this work has helped to identify some which were previously unknown (particularly those with aerosols and clouds). These have either been fixed or are at least now known problems which hopefully shall benefit future users of the offline SOCRATES code.

4.6 Results of 1 year clear sky PRP calculations

As described earlier, clear sky PRP calculations were performed for a limited 1 year case, determining radiative forcing from aerosol, humidity, surface albedo, temperature, surface temperature and residual terms due to perturbation of BC emissions from 1850 to 2014 levels. The results of these calculations, as well as the total forcing calculated with offline SOCRATES, are presented in table 4.1. The forward and backward forcings for SW and LW are also illustrated in figure 4.5.

Table 4.1: 1 year global mean TOA radiative forcing or radiative adjustment calculated for a range of adjusting variables using the partial radiative perturbation method due to 1850 to 2014 BC emissions perturbation. The forward (fwd), backward (bwd) and two-sided (mean) PRP calculations are given for each variable, for SW, LW and net forcing. Total is the total clear sky forcing calculated as the difference between TOA fluxes calculated in SOCRATES with all variables set to control or perturbed using equation 4.14 (i.e. not just the sum of the table column). All values are given in W m^{-2}.

Variable	**SW**			**LW**			**Net**		
	fwd	bwd	mean	fwd	bwd	mean	fwd	bwd	mean
Total			0.33			0.18			0.51
Aerosol	0.41	0.39	0.40	0.03	0.03	0.03	0.45	0.42	0.43
Specific humidity	-0.02	0.01	-0.01	0.52	-0.34	0.09	0.50	-0.34	0.08
Surface albedo	-0.04	-0.05	-0.05	0.00	0.00	0.00	-0.04	-0.05	-0.05
Temperature	-0.00	-0.00	-0.00	-0.29	0.25	-0.02	-0.29	0.25	-0.02
Surface temperature	0.00	0.00	0.00	-0.06	0.17	0.06	-0.06	0.17	0.06
Other variables	-0.02	-0.02	-0.02	-0.02	0.08	0.03	-0.05	0.06	0.01

Firstly we consider the total clear sky radiative forcing calculated from SOCRATES. The overall forcing includes a large positive SW forcing and a smaller positive LW forcing. The SW forcing arises primarily from the aerosol term, while the LW forcing arises primarily from a combination of water vapour and surface temperature adjustments. As a test of the validity of the implementation of these PRP experiments using SOCRATES, the total clear sky radiative forcing calculated from SOCRATES (eqn. 4.14) can be compared with the total calculated from UKESM diagnostics. The two estimates agree within 2 d.p. for SW, LW and net clear sky

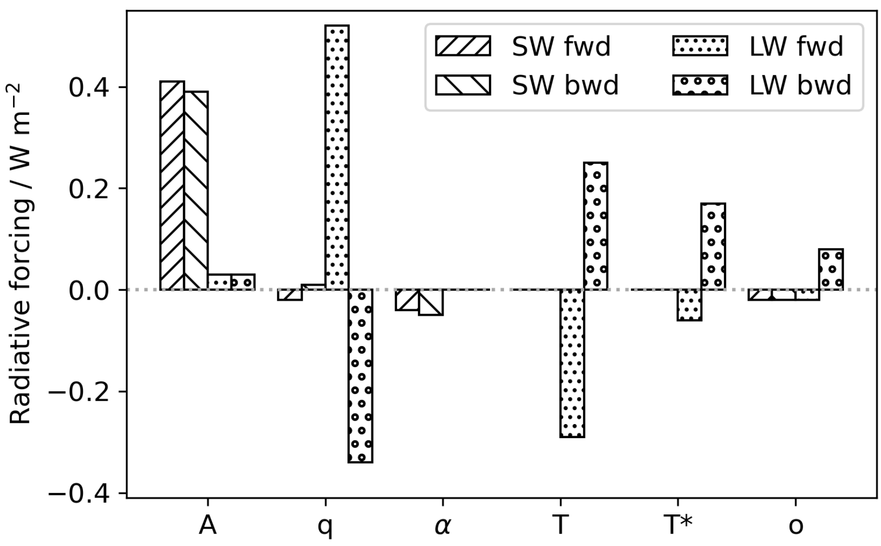

Figure 4.5: Plot illustrating radiative forcings calculated in table 4.1, using partial radiative perturbation, for SW and LW forward (fwd) and backward (bwd) calculations. A is all aerosol, q is the water vapour, α is the surface albedo, T is the atmospheric temperature, T* is the surface temperature, and o is other variables.

radiative forcing. This confirms that the remaining unfixed biases in TOA fluxes described at the end of section 4.5.5 have no significant effect on radiative forcing calculated from the PRP experiments.

As well as the total clear sky radiative forcing calculated from SOCRATES agreeing with that calculated from UKESM, the sums of the forward, backward or two-sided PRP calculations are all equal to the total clear sky radiative forcing too. That is to say that equations 4.14 and 4.13 produce equal estimates of ΔF_{clr}, as expected. This confirms that the cumulative PRP method does produce individual PRP terms with a sum consistent with the total radiative forcing.

Moving onto the specific variables, table 4.1 and figure 4.5 show that the aerosol PRP term is positive in the SW with a small positive forcing in the LW. The positive SW forcing is expected for aerosol given that BC has a positive IRF from absorbing incoming solar radiation (Li et al., 2022). Indeed the SW forcing, illustrated in figure 4.6, shares a similar spatial pattern to the all sky 34-year mean forcing illustrated in figure 3.4b, reflecting larger forcing in areas with stronger emissions of BC. The aerosol PRP term makes up 84 % of the net clear sky radiative forcing.

As with the total radiative forcing, the aerosol radiative forcing calculated from the PRP method can be compared with radiative flux diagnostics output from UKESM, specifically the ΔDRE, over the same one year period. Since the PRP calculations are performed under clear sky, the PRP aerosol forcing is compared with ΔDRE' from equation 1.13 (i.e. calculated as the difference between clear sky and clear-clean sky radiative forcing, rather than between all sky and clean sky as done in Ghan (2013)). The two measures of aerosol forcing are very similar having the same net value of 0.43 Wm^{-2}, but with ΔDRE' having a SW and LW forcing of 0.41

and 0.02 Wm^{-2} respectively, being slightly different from the SW and LW aerosol PRP terms (0.40 and 0.03 Wm^{-2} respectively). The spatial patterns of both ways of calculating aerosol forcing are also very similar.

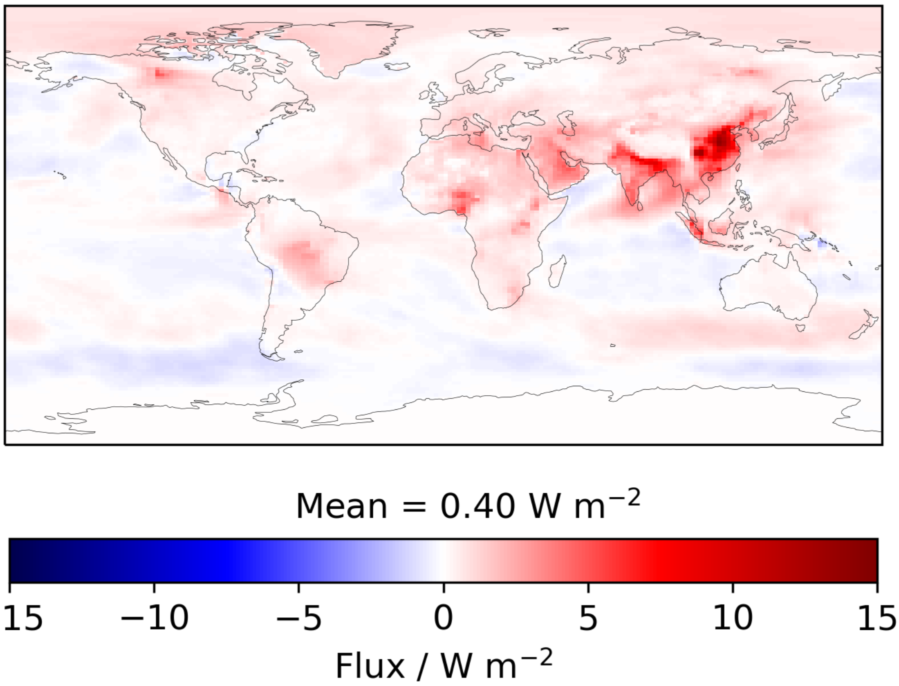

Figure 4.6: Map of 1 year mean clear sky aerosol radiative forcing under 1850 to 2014 BC emission perturbation, calculated using two-sided partial radiative perturbation.

Moving on to T, T* and q, it can be seen that, as expected, most of their radiative effect is in the LW, but with opposing sign in the forward and backward calculations. In figure 4.5 it is particularly evident that while the magnitude of radiative adjustment is large in the forward and backward cases, the opposing signals result in a relatively small two-sided radiative adjustment. This is in contrast to the aerosol and albedo PRP terms for which the forward and backward calculations are similar. Looking at individual monthly means in figure 4.7, while there is a large difference between forward and backward radiative adjustments for T, T* and q, they are not always of opposing sign. The cause of the difference in forward and backward PRP calculations is likely due to the decorrelation bias described in section 1.3.3.

For the albedo PRP term, Table 4.1 shows that there is a small negative radiative adjustment in the SW, which as noted earlier is similar in the forward and backward cases. Spatially, this adjustment is mainly confined to the mid and high northern latitudes, as illustrated in figure 4.8a. Comparing this pattern with the one month mean adjustment of the surface albedo field for July (fig. 4.8b) and January (fig. 4.8c) of the one year run, it is clear that most of this albedo forcing arises from albedo adjustment during northern hemisphere winter. The albedo adjustment is therefore most likely primarily due to changes to albedo arising from snow cover changes. With a one year integration it is not possible to suggest any albedo change due to changing snow cover is a statistically significant adjustment response, given that snow cover will depend strongly on meteorology. It should also be noted that while BC forcing could affect snow

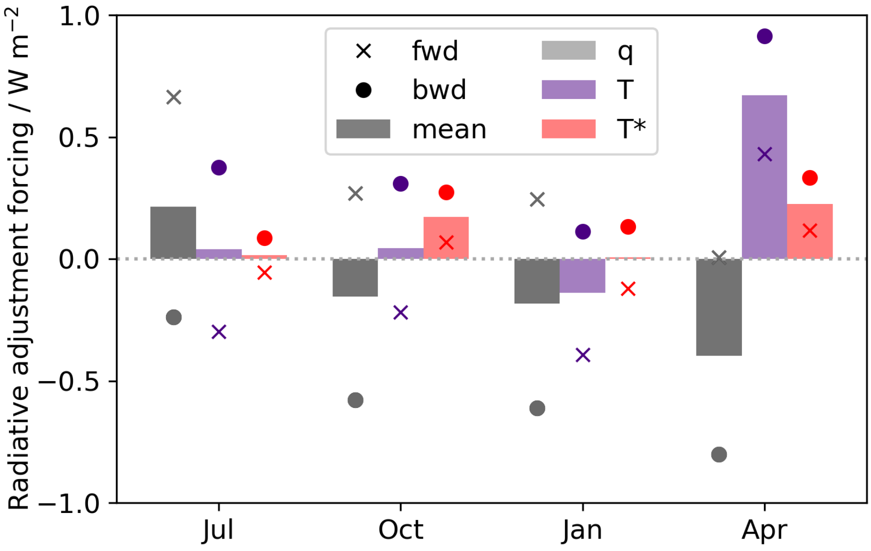

Figure 4.7: Clear sky 1 month mean LW radiative adjustments calculated from SOCRATES partial radiative perturbation calculations for temperature (T), surface temperature (T*), and water vapour (q) adjustment due to 1850 to 2014 BC emission perturbation, for mid season months. The two-sided (mean) calculation as well as forward (fwd) and backward (bwd) PRP calculations are given for each variable.

cover indirectly by affecting surface temperature, precipitation and dynamics, in UKESM1 there is no mechanism for BC deposition on snow to affect snow albedo.

The other PRP term, which includes ozone and pressure adjustments, causes a small net radiative adjustment of 0.01 W m^{-2}. This results from a small negative two-sided forcing in the SW (-0.02 W m^{-2}) and slightly larger positive two-sided forcing in the LW (0.03 W m^{-1}). This likely points to changes in O_3, which can both absorb incoming SW radiation in the stratosphere and outgoing LW radiation in the troposphere. However, this radiative adjustment is unlikely to represent a significant effect, so is not considered further.

4.7 Discussion of the PRP method and results

A major limitation for interpretation of the PRP results is the 1 year integration period, which is insufficient for determining uncertainties. This prevents identifying any statistical significance for the specific values calculated. However, even with this short integration time we can draw some tentative conclusions about the distribution of BC PI to PD ERF between the IRF and individual radiative adjustments.

Firstly, since the aerosol PRP term will mostly correspond to BC perturbation (with adjustments to other aerosol species possible but likely small), we can state that the IRF was determined to be 84% of the clear sky ERF, with the remaining 16% made up of radiative adjustments. The estimate for BC IRF here (0.43 W m^{-1}) is similar to that determined by (O'Connor et al., 2021) using the double call method (0.38 ± 0.01 W m^{-1}) also using UKESM1. Since the double call method applied to the same 1 year period here returns the same value

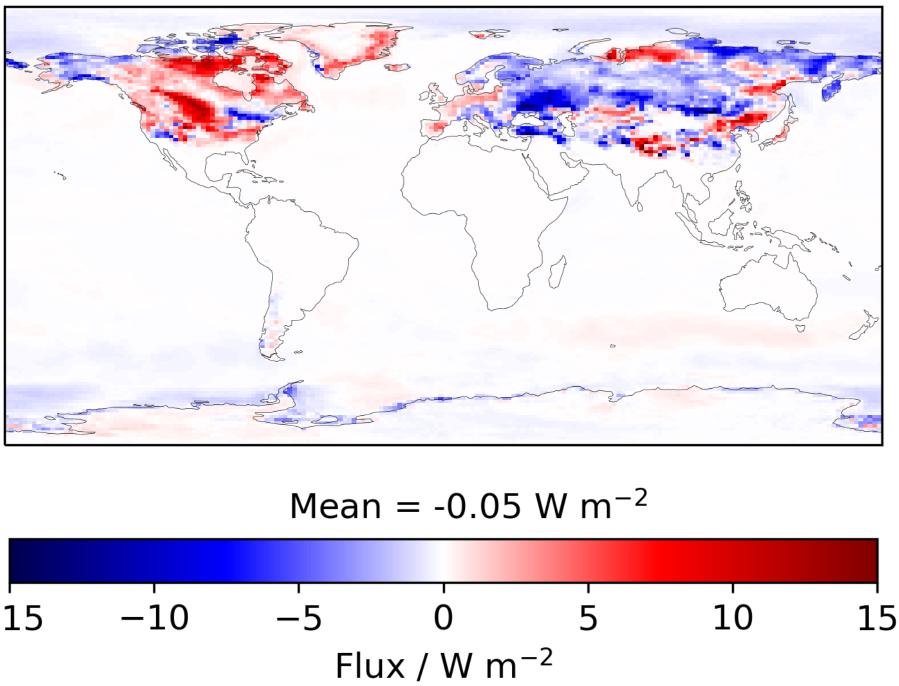

(a) SW albedo PRP radiative adjustment

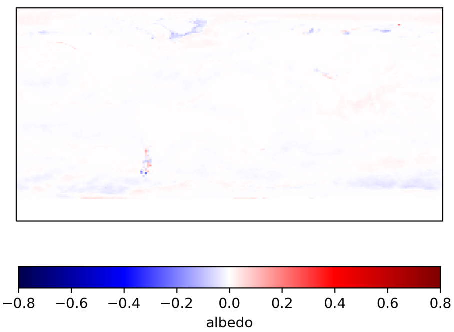

(b) Surface direct albedo July

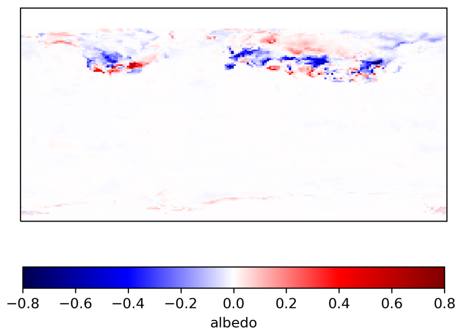

(c) Surface direct albedo January

Figure 4.8: Maps of (a) one year mean clear sky SW surface albedo radiative adjustment calculated from partial radiative perturbation and (b, c) month mean direct surface albedo adjustment for SW bands 1-3 (200 - 690 nm) for July 2014 and January 2015, due to 1850 to 2014 BC emissions perturbation.

as the PRP term, the difference between the estimate of BC IRF from this work and that of O'Connor et al. (2021) is likely to relate to the choice of PD as the baseline atmosphere in this work, while PI was used as the baseline atmosphere in theirs.

Turning to the remaining PRP terms, these can be compared with the results of Smith et al. (2018), who used the kernel method to calculate radiative adjustments from a 10 times increase in present day BC across multiple models and kernels. They determined radiative adjustments of approximately -0.1, -0.9, -0.3, 0.7, and 0.1 W m^{-1} due to surface temperature, tropospheric temperature, stratospheric temperature, water vapour, and surface albedo adjustments. Assuming linearity, dividing these values by 10 approximates the response expected from a PI to PD perturbation as modelled in this work. This results in similar radiative adjustments for water vapour between the two studies (0.08 in this work and 0.07 [=0.7/10]) W m^{-1} in Smith et al. (2018)). However, the surface temperature and surface albedo responses found in this work are much bigger, being half of the values report by Smith et al. (2018) before dividing by 10. Meanwhile the temperature adjustment is much smaller in this work when compared to their combined tropospheric and stratospheric adjustments (-0.02 compared to -0.12 [=(-0.9 + -0.3)/10] W m^{-1}). Some differences are expected since Smith et al. (2018) performed the radiative kernel calculations in all sky rather than clear sky conditions. However, there will be significant differences due to the short 1 year integration used in this work compared to their 10 year integrations.

Aside from the exact values, much can be learned from comparing the PRP terms relative to each other. One key finding from the 1 year clear sky BC PRP experiment is that the cumulative PRP method did return individual PRPs whose sum equals the total forcing. By comparison, using conventional PRP, Bickel et al. (2020) calculated the difference between the sum of individual radiative adjustments and the ERF - IRF for two experiments, one perturbing air traffic (specifically contrails) by 12 times, the other doubling CO_2. For the air traffic and doubling CO_2 experiments they determined differences of 0.03 and 0.18 Wm^{-2}, compared to ERFs of 0.26 and 3.55 Wm^{-2} respectively. While these differences are much larger percentages of the ERF (and individual adjustments), this may be because they determined the IRF term from a separate online experiment to that which generated input for the PRP calculations (and determined the ERF). Rieger et al. (2017), investigating radiative feedbacks under a doubling of CO_2, determined a much smaller difference between the sum of individual variable radiative feedback parameters determined with PRP and the total radiative feedback parameter, determining the difference to be less than 0.01 $Wm^{-2}K^{-1}$. These different studies with the conventional PRP method suggest that the sum of individual PRP terms could be very close to the total radiative forcing, but not identical, though with only two studies it is hard to suggest what the typical difference might be. By contrast the cumulative method guarantees that the sum and the total are equal, in addition to having a small computational advantage in saving two runs of the radiation code.

Another interesting result is the large difference between forward and backward calculations in the temperature, water vapour, and surface temperature LW PRP terms. These differences between forward and backward calculations are expected due to spatial decorrelation of perturbed variables as described in section 1.3.3. However, unlike the example of cloud decorrelation with humidity given in section 1.3.3, it is not immediately obvious what the decorrelation effect between variables in the absence of clouds would be. That said, one would expect some degree of correlation between humidity, atmospheric temperature and surface temperature, and so some decorrelation bias effects on radiation could be hypothesised.

For example, one could suppose that the surface temperature is positively correlated with the absorptivity of the atmosphere due to water vapour. Then, consider the outgoing LW radiation due to surface temperature in a multi layer grey gas model, where the emissivity is uniform between levels:

$$OLR_s = (1 - \epsilon)^n \sigma T*^4 \tag{4.15}$$

where OLR_s is the outgoing LW radiation from the surface, ϵ is the emissivity of each layer, n is the number of model layers, σ is the planck constant, and T* is the surface temperature. Assuming that ϵ is uniform across layers is generally a poor assumption for water vapour, but should be sufficient in this case. Also, assume that Kirchoff's law applies, such that absorptivity is equal to ϵ, and that while epsilon is uniform between levels, it varies spatially between different model atmospheric columns. Then, if surface temperature and absorptivity due to water vapour are positively correlated, in an atmospheric column where T* is high, ϵ is also high (and vice versa). If however, surface temperature and humidity were decorrelated, where T* is high ϵ would, on average, be lower (and vice versa). Looking at equation 4.15 one can see that if columns with higher T* are scaled by (1 - ϵ) with a lower value of ϵ, then OLR_s is higher. Although the reverse is true for columns with lower T but now higher ϵ, the mean OLR_s across columns would be higher. Consequently, when T* and ϵ are decorrelated there will be a negative net TOA downwards flux than when correlated.

This hypothesis, if true, could then partially explain the results presented in figure 4.5. To elaborate, the surface temperature forward PRP term includes the decorrelation bias as a positive (the second term in 4.5 is decorrelated, but as noted earlier, the equations are multiplied by -1). Consequently, since the bias is negative, the surface temperature forward PRP term should be negatively biased. This is seen in figure 4.5, where the surface temperature forward PRP LW term is indeed negatively biased (and the backward term positively biased).

This hypothesis could also partially explain the positively biased humidity forward PRP LW term, which should include the same magnitude bias, although positive (since the decorrelation bias is included as a negative in the humidity forward PRP term). However, since the difference between the forward and backward humidity PRP terms is larger than the difference in surface temperature forward and backward PRP terms, there is likely an additional decorrelation bias between humidity and atmospheric temperature also. Indeed, considering table 4.1, the sum of the differences between the forward and backward LW surface temperature (0.23 W m^{-2}) and atmospheric temperature (0.54 W m^{-2}) PRP terms (= 0.77 W m^{-2}) is similar to that of the humidity LW forward and backward PRP terms (0.86 W m^{-2}). A similar thought process to that applied to surface temperature and humidity could be used to suggest a mechanism by which decorrelation of humidity and atmospheric temperature causes a bias in TOA flux. However, this would be more complex given the temperature of multiple layers are being perturbed rather than just the surface. By contrast, there may be little decorrelation between atmospheric temperature and surface temperature, since they should have little effect on each other when considering equation 4.15.

To verify whether this (or any other) hypothesis is a realistic explanation for the difference in forward and backward PRP terms found here would primarily require investigating how surface temperature and humidity are correlated. Additionally, one could investigate whether the magnitude of OLR change from decorrelating surface temperature and humidity is approximately consistent with the difference in forward and backward terms observed here. This is not investigated here, and is left to future work.

While the differences between forward and backward terms seem large for humidity, temperature and surface temperature, and the mechanisms behind the decorrelation effects are not further investigated, similar findings have been made in other studies. The closest comparison is against Bickel (2023), who reported differences in forward and backward radiative adjustments of ∼0.5 W m^{-2} for both lapse rate and humidity PRP terms, and of ∼0.2 W m^{-2} for surface temperature, under clear sky conditions resulting from CO_2 perturbation (their figure C.1). These magnitudes are similar to those presented here (table 4.1), where the magnitude of differences between forward and backward PRP terms are 0.86, 0.54 and 0.23 W m^{-2} for specific humidity, temperature and surface temperature PRP terms respectively.

It is worth noting that although Bickel (2023) find that their forward humidity PRP term is negatively biased and the backward term is positively biased, this is consistent with the results here where the reverse is true. This is because in their calculations the forward PRP terms for all variables include the decorrelation biases as a positive, and as a negative in the backward PRP terms, whereas as outlined above, the specific humidity PRP term here includes the decorrelation bias as a negative. Conversely, they find that their forward surface temperature PRP term is negatively biased, which is the same as in this work where the decorrelation bias due to surface temperature changing will be included as a positive in the forward PRP term. Comparing their lapse rate term with the atmospheric temperature term in this work is more complicated since they are not equivalent terms. This does point to one small disadvantage of the cumulative PRP method, this being the slightly more challenging interpretation of decorrelation biases in the forward and backward terms, though as noted before this should not effect the two-sided mean value which the value of interest.

It is also worth considering that Bickel (2023) found a significant difference in their forward and backward calculated CO_2 PRP term. Since CO_2 is well-mixed, it is unlikely there is much spatial decorrelation bias present. Instead, this difference may result from the difference in masking effect from water vapour and temperature of the CO_2 perturbation, between performing the perturbation on a control or perturbed background. This effect could contribute to the spatial decorrelation seen in the PRP calculation presented here for temperature, surface temperature and humidity.

Now comparing the PRP results to the use of other methods for separating ERF components, it is worth highlighting the similarity of the aerosol PRP term and the ΔDRE' calculated from UKESM double radiation call diagnostics. These metrics both measure the radiative effect of the aerosol perturbation, so are expected to be similar, but there may be differences due to how they are calculated. While a longer integration period is required to confirm the value of the aerosol PRP term, after just one year the two metrics give a very similar result. This suggests that the PRP method here is giving sensible results. The two metrics might not give such good agreement if performed under all-sky conditions however, since cloud masking effects are excluded and may differ between the two methods.

A final point this work highlights is that implementing PRP with an offline radiation code that matches the settings of the online radiation code can be challenging, particularly with hourly input resolution. Firstly, while PRP in principle allows the online and offline calculations to have matching settings unlike radiative kernels (unless creating a kernel for a specific set of simulations), in practice it is still difficult to achieve. In particular, despite considerable efforts, it proved too challenging to get the representation of cloud in the online radiation code and offline SOCRATES to match. Secondly, the large input data size and number of timesteps SOCRATES needed to be run for when using 1 hourly input resolution resulted in long run times for both preprocessing and run suites. Minor issues encountered, such as the preprocessing suite occasionally failing to archive input for model days (requiring manual re-running), could also have been a consequence of the excessive burden placed upon the computing system. While the limits on computing and runtime resources will vary between different computing systems and by model, these do pose a real limitation on the utility of this methodology.

4.8 Conclusions

In this chapter an alternative implementation of the PRP method was devised and used to investigate radiative adjustments to BC perturbation. This cumulative PRP method is found to be a more efficient technique for PRP experiments with no obvious drawbacks compared to the conventional PRP method. The cumulative PRP method saves two calculations of the radiation code compared to the conventional method, whilst ensuring the sum of individual PRP terms equals the total radiative forcing, something that does not appear to be guaranteed with the conventional method based on the studies that have considered this matter. However, while it is expected that the two-sided PRP terms should be similar between the two PRP methods, future work is needed to verify this, such as by using both methods applied to the same dataset.

Due to unforeseen challenges representing cloud fields in SOCRATES it was not possible to conduct all-sky PRP calculations, and only a 1 year integration in clear sky conditions was performed here. This limits the interpretations that can be made about BC radiative adjustments. However, the radiative forcing calculated from the aerosol PRP term is found to be similar to ΔDRE' calculated using additional radiation call diagnostics, a result likely to be fairly robust even with the short integration time. To understand the magnitude of the remaining PRP terms these calculations should be performed under all-sky conditions and for a longer integration period than what was possible here.

Finally the challenges encountered highlight the difficulty in implementing the PRP method. In particular, ensuring the offline SOCRATES and online UKESM radiation codes used the same settings proved very challenging, and using hourly input required significant computing resources and time. While these issues will vary between models, this difficulty should be noted for future studies and requires further work to address in the model used here.

Chapter 5

Thesis conclusions and Recommendations

5.1 Key Conclusions

It is worth reflecting back on the motivation for, and aims of, this thesis. It was noted that there are large uncertainties concerning the climate impacts of declining aerosol emissions, and that this uncertainty is partly driven by uncertainty in radiative adjustments. The aims were firstly to implement variations on the nudging mechanism denial and PRP methodologies, assessing how useful they are for investigating aerosol radiative adjustments; and secondly, to apply these methodologies to anthropogenic SU and BC emissions and learn something about their radiative adjustments.

5.1.1 How useful are the methods used here for investigating aerosol radiative adjustments?

So, what can be learnt from this work about how useful these two methods are for quantifying aerosol radiative adjustments and understanding there mechanisms? Beginning with the nudging mechanism denial methodology, this method was able to produce estimates of the radiative effects of circulation- and temperature-mediated adjustments. However, confidence in these estimates is limited by a number of issues that were encountered with nudging. In particular, one can conclude that the radiative effect of temperature-mediated adjustments is unlikely to represent a reliable estimate of actual physical processes, while the estimate of radiative effects of circulation-mediated adjustments may be somewhat reliable but still uncertain.

The major issue contributing to this lack of certainty in the results is that nudging was found to introduce a number of unexplained effects and biases. This particularly included large adjustments in low cloud over marine stratocumulus regions and storm tracks in uvθ-nudged simulations, the signal from which dominates the estimate of the radiative effects of temperature-mediated adjustments. While it is plausible there could be a physical reason for the strong low cloud adjustment in uvθ-nudged simulations, the observation of close to opposite low cloud adjustments between SU and BC perturbed uvθ-nudged simulations despite the different interactions the two aerosol species have with clouds, would suggest this is more likely an error introduced by the nudging. Nudging also resulted in biases in cloud fraction between the free and (uv- and uvθ-) nudged controls. While the effect of these bias on calculated radiative adjustments is uncertain, the biases may indicate nudging is distorting the representation of cloud processes even in the absence of aerosol perturbations. While not as severe an issue as the anomalous low cloud adjustments in the uvθ-nudged simulations, this does reduce confidence in quantification of both circulation- and temperature-mediated adjustments here. It is worth noting also that these errors could even affect other, more conventional applications of nudging, particularly where temperature related quantities are nudged.

Another significant issue affecting conclusions from the uvθ-nudged simulations is the limited suppression of adjustments of θ. While this issue is to a degree inherent in the method due to the nature of nudging, the uvθ-nudging achieved typically less than a 50 % reduction in θ adjustment in the troposphere, which is certainly less than aimed for. This reduces confidence that any difference between the free and uvθ-nudged simulations do correspond to suppressed temperature-mediated adjustments, or that the estimate of radiative effects of temperature-mediated adjustments actually includes the effect of all temperature-mediated adjustment, since only part of the adjustment in θ is suppressed. Suppression of the horizontal wind adjustment however was much more effective, so this is less of a concern for the calculations concerning circulation-mediated adjustments.

Beyond these uncertainties, this work also highlighted that because of the strong interdependence of temperature and winds in the model, nudging either will affect the other. Consequently it is difficult to suggest no temperature-mediated adjustments are suppressed by uv-nudging or no additional circulation-mediated adjustment is suppressed by uvθ-nudging. While different implementations of the nudging methodology may be able to achieve better separation, it is unlikely the two could be separated totally, and this is a limitation inherent to the method.

Despite this rather negative assessment of the utility of the nudging denial methodology, it should be highlighted that these conclusions are derived from using only one ESM, and, while a number of nudging settings were tested, still others could be investigated. Similar studies with different models may not encounter the same errors or may achieve better adjustment suppression, as discussed later in section 5.3.2.

Furthermore, the method potentially could have more value for looking at the broad mechanisms of adjustments, particularly at smaller scales. In this thesis the focus was primarily on quantifying the magnitude of global multiannual mean radiative effects of circulation- and temperature-mediated adjustments. As noted, the uncertainties and temperature-wind interdependence make it difficult to trust the resulting values. But while less focus was directed towards it here, the patterns of responses, if not their magnitudes, may provide useful insights into circulation- and temperature-mediated adjustment mechanisms. For example, one could investigate why uv-nudging alone suppressed most of the adjustment in specific humidity in the upper troposphere seen in figure 3.10, how this response varied spatially, and whether it was due primarily to changes in temperature resulting from uv-nudging or could be attributed to suppression of specific dynamical responses. This could help to understand the mechanisms by which anthropogenic aerosol emissions have affected upper tropospheric humidity, even if the quantified radiative adjustment associated with humidity adjustments is relatively uncertain due to the aforementioned limitations. Focusing on such mechanisms more useful when focused on specific regions, or if perturbations to aerosol were made only over specific regions, since there are fewer effects to consider from global atmospheric responses. This point is discussed further in section 5.3.2.

As for the PRP methodology, the most challenging aspect was implementing the method. Specifically, ensuring that the TOA radiative fluxes matched between offline and online calculations in all-sky calculations proved impossible in the time available due to the complexity of cloud representation in UKESM and differences between the online and offline radiation codes. It was also considerably time consuming to implement the method generally and troubleshoot even the clear sky implementation, particularly with the use of hourly input to SOCRATES. These difficulties are in themselves relevant to answering the question of how useful the PRP method is, since it is the difficulty implementing the method alongside the computational constraints that is the main disadvantage of the PRP method noted in the literature. This difficulty is of course model dependent however, evidenced by the fact that the PRP method has been used in other studies, and clearly the difficulty of implementation will depend on the model used. This point also highlights that studies using PRP should be sure to confirm that offline radiation code calculations match, within a reasonable degree, online calculations from the GCM/ESM before conducting PRP calculations.

The use of hourly input to SOCRATES contributed to the difficulties implementing the method, since it made the simulation times considerably longer. Given hourly input is a higher resolution than that used in other studies using PRP (Colman et al., 2001; Mülmenstädt et al., 2019; Bellouin et al., 2020b) this may suggest that using hourly input is inadvisable. However, the effects of using this high input resolution were not truly assessed, particularly since the use of such high resolution input was primarily intended to help with representing cloud radiative adjustments which were not tested in the clear sky integration performed. Therefore it is not possible to conclude what the advantages could be of using hourly input from this work, only that it considerably adds to the difficulty implementing the method.

Despite the difficulties implementing the PRP methodology, the 1 year the clear-sky integration was successful in that it successfully separates out the different radiative adjustments from the IRF, as intended. The cumulative version implemented here also successfully ensures that the sum of the individual PRP terms and the ERF matched and does so with a slight computational advantage over the conventional PRP implementation.

5.1.2 What can be learned about anthropogenic SU and BC radiative adjustments?

Moving onto the second research aim, what can be learned from this thesis about SU and BC radiative adjustments? One could conclude from the nudging experiments that circulation adjustments do appear to be a significant contributor to ERF for SU but not BC (though both are close to the margin of significance used here), based on the difference between the uv-nudged and free simulations. This corroborates the findings of Johnson et al. (2019) that uv-nudging does not significantly alter the ERF from BC perturbation, whilst also showing that this may not be true for other forcing agents such as SU. It also shows how not just the ERF, but adjustments, particularly in clouds and specific humidity, may be affected by uv-nudging. To what extent these changes can be attributed to real circulation adjustments rather than nudging artefacts is somewhat uncertain however.

As for the differences between uv-nudged and uvθ-nudged experiments, while one might conclude that there are significant atmospheric temperature-mediated radiative adjustments that oppose the overall ERF for both aerosols, significant uncertainties make this conclusion highly uncertain, and suggest much of the difference in ERF between the experiments may instead be due to artefacts from nudging. For BC, a significant negative atmospheric-temperature-mediated radiative adjustment seems a plausible result since ERF$_{ari}$ cloud adjustments due to BC perturbation (i.e. semi-direct effects) have generally been estimated to be negative, with many of these adjustments occurring due to changes in atmospheric temperature profiles due to BC, as discussed in 1.3.5. For SU, there has been little quantification of ERF$_{ari}$ cloud adjustments so there is little to compare the SU atmospheric temperature-mediated radiative adjustment estimate with. While adjustments to cloud due to changes in atmospheric heating are not generally thought to be as significant for SU as BC, the ability for land surface temperature to respond and produce cloud feedbacks could, at least in principle, be a plausible mechanism for this estimate. However, the large differences in marine stratocumulus adjustment between uvθ-nudged and uv-nudged experiment pairs for both species do not have an obvious physical explanation consistent with the changes in lower tropospheric stability or known mechanisms of cloud adjustment. Likewise the differences in marine stratocumulus for the two species are strongly mirrored, which is not expected due to the different ways the aerosols affect radiation and clouds. As such, there can be limited confidence in any conclusions drawn from the uvθ-nudged simulations about atmospheric temperature-mediated adjustments.

From the PRP method there is less to be learned due to only conducting a 1 year integration in the clear sky only. The BC clear sky IRF was estimated to make up 84% the clear sky forcing due to BC, with clear sky adjustments comprising the remainder. Temperature, surface temperature and specific humidity responses, despite having strong magnitudes in individual forward and backward PRP calculations, are determined to be small terms overall. Surface albedo also contributes a small clear sky radiative adjustment to BC ERF, primarily dependent on snow and ice cover changes at high latitudes, that is consistent between forward and backward PRP calculations.

5.2 Wider Significance

Having considered the key conclusions from this thesis through the lens of the research aims, we can next consider what wider significance there might be in these conclusions. For nudging, one clear significance from this work is it helps to address a point raised in the latest AR6 IPCC report. Specifically Forster et al. (2021) noted that while nudging horizontal winds could reduce integration time for fSST experiments, there are not enough studies to conclude whether the exclusion of circulation adjustments by this method is significant. The work here found that, while borderline, there was a difference in ERF between free and uv-nudged SU perturbed experiments at the 95% confidence limit, but not in the BC case. There were also large differences in adjustments of cloud fraction and specific humidity between free and uv-nudged experiments. While it is not certain this difference could be attributed to actual physical responses comprising circulation adjustments rather than nudging artefacts, this at least suggests that uv-nudging could significantly alter the determined ERF, and should be investigated further if it is to be used for the purposes of reducing simulation lengths. This is particularly true where accurate quantification of adjustments of specific fields, such as of cloud fractions and specific humidity, are of interest.

Another point of significance is that this thesis illustrates the large potential for unintended effects of nudging. When nudging temperature related quantities, one should be wary of artificial stability changes affecting cloud cover. This could occur wherever large differences exist between temperature fields in the nudged model and nudging input, not just in the mechanism denial methodology used here. This thesis would also suggest there are potentially large effects on a number of variables from nudging only the horizontal winds, such as the large bias between uv-nudged and free cloud fields, even in the absence of perturbed aerosol emissions. While this bias might be overlooked when focusing on ERFs, it could have significance for studies using nudging to investigate cloud or precipitation responses more specifically.

As for the PRP experiments, the cumulative PRP method presented here offers some improvements on the conventionally applied PRP method. This method should reduce computation time relative to the conventional PRP method and ensures the sum of individual radiative adjustments adds up to the total. While this cumulative PRP method would still take more resources to employ than use of a pre-existing radiative kernel, radiative kernels (even if applied with forward and backward calculations) cannot calculate individual radiative adjustments that are conservative of the total radiative forcing. This is likely to be a small benefit that does not outweigh the computational cost of using PRP over an existing kernel, but could be of value in some studies.

5.3 Recommendations for future study

5.3.1 Direct extensions to this thesis

A number of directions for future study can be suggested from this work. Firstly, there are a number of very specific recommendations that mostly seek to address or clarify some of the challenges encountered with the use of the nudging mechanism denial and PRP methodologies. These recommendations could be achieved just with UKESM1 and the offline SOCRATES code, and could be seen as direct extensions to this thesis. Some of these recommendations were noted in the conclusions to each chapter, but are worth restating here.

For nudging these specific recommendations include:

- Investigating the impact of varying the nudging sequence in UKESM
- Testing higher resolution nudging input over a longer simulation
- Nudging with the land surface temperature fixed
- Nudging just potential temperature
- Investigating stratospheric responses to nudging

The first four of these recommendations are intended to address the issues with nudging θ that were noted in section 5.1: namely the anomalous responses when nudging θ and the weak suppression of θ adjustment. The first recommendation is based on the finding of Zhang et al. (2022) that strong biases in marine stratocumulus clouds may occur when nudging input is generated at a different timestep to when nudging is applied. Investigating whether a change to the nudging sequence used in this work might avoid the anomalous marine stratocumulus adjustments is a logical first step. This would definitely be worth doing, given that it may be an issue that affects representation of cloud in other studies using nudging in UKESM, not just for the mechanism denial methodology employed here.

The second recommendation is made because while higher time resolution nudging input was tested in this thesis, it was only a limited test. An extended study could investigate using 3 and 1 hourly nudging input resolution with 5 year model runs to verify whether the higher resolution would improve theta adjustment suppression or not.

The third recommendation is made firstly on the basis that if the surface temperature could be fixed when nudging theta, the theta adjustment suppression should be improved in the lower troposphere. Secondly, removing the land temperature response is necessary for determining ERF without including feedbacks to land surface temperature change. In particular, conducting the same experiments with fixed land surface temperature would help determine whether the cloud fraction adjustments observed in the SU perturbed cases are actually primarily responses to surface temperature (i.e feedbacks) or are indeed adjustments. However, fixing land surface temperature, and soil temperature and moisture as performed in Ackerley and Dommenget

(2016) and Andrews et al. (2021) would be difficult to implement in UKESM, as this was briefly investigated it as a possibility during this work. A possible alternative could be nudging the land surface and soil fields instead.

The fourth recommendation is suggested because, while nudging a temperature related quantity without nudging horizontal winds is seldom done in the literature, it would be useful to explore it as a way to achieve better suppression of the θ adjustment or reduce the unexplained low cloud adjustments. This could be the case since nudging both uv and θ could be over-constraining the model, preventing greater suppression of the θ. It would also be a relatively easy recommendation to implement with the existing UM nudging code.

The final recommendation is made since the effects of nudging in the stratosphere were largely ignored in this work. This was because anthropogenic aerosol adjustment processes are expected to occur primarily in the troposphere (Smith et al., 2018). However, understanding some of the stratospheric responses to nudging noted in this thesis, and possibly adjusting the nudging strength in the stratosphere, could help understand and address some of the issues with nudging. It may also be more important if the nudging mechanism denial methodology were applied to emissions of volcanic aerosol, for which stratospheric responses may be more important than here.

For the use of PRP, three specific recommendations are made:

- Performing a longer, all-sky, PRP integration, which first requires resolving the difference in TOA radiative fluxes between the offline SOCRATES and online UKESM calculations conducting in this work
- Comparing the cumulative PRP method with calculations using the conventional PRP method

With respect to the first recommendation, the primary reason for using 1 hourly input to SOCRATES was to provide accurate representation of cloud adjustments in SOCRATES, given that cloud adjustments were expected to be the largest adjustment for aerosols. As such it would be useful to represent these adjustments in SOCRATES, and also enable the other PRP terms to be conducted under all sky rather than clear sky. This could then be run for a longer integration period to determine PRP terms with meaningful statistical uncertainties.

For the second recommendation, whilst in principle the cumulative method used here is hypothesised to produce similar PRP terms to the conventional PRP method, this hypothesis should be verified. This could be done effectively with only a single year of integration using the conventional PRP method with the same input data used here. Such an experiment would also allow testing to what degree the conventional method produces a sum of PRP terms that equals the ERF. These tests would help verify whether the cumulative PRP method used here produces results consistent with the conventional PRP method and whether the cumulative method ensures better conservation of the total radiative effect across radiative adjustments than the conventional method.

5.3.2 Broader paths for extension

As well as the very specific recommendations made above, the discussion of significance earlier points to some broader ways this work could be extended if considerably more resources were available. These can be summarised as follows, and are expanded upon below:

- Testing the nudging mechanism denial methodology with different models, nudging setups and nudged variables
- Testing the nudging mechanism denial methodology with different forcing agents
- Applying nudging with a focus on regional scales
- Considering effects of nudging on precipitation
- Combining the use of nudging and PRP for separating both adjustment mechanisms and specific variable contributions to radiative adjustments
- Investigating whether radiative kernels calculated from different baseline climates states significantly affects radiative adjustment calculations using kernels
- A review of methods for decomposing ERF into radiative adjustments, particularly for aerosols

Testing the nudging mechanism denial methodology with different models would help establish whether different ways of implementing nudging affect the adjustment suppression and introduction of artefacts, and hence, the usefulness of the method. While a natural step for any work is to compare results from different models to see how responses vary, for this method in particular the only other study using this same methodology used the very similar HadGEM3 model (Johnson et al., 2019). Testing the methodology with different models would also implicitly test different implementations of nudging, since nudging schemes for different models will be implemented differently (e.g. the interpolation of nudging input onto model timesteps). Likewise, to suppress the same adjustments different variables can be nudged (e.g. temperature, potential temperature, or dry static energy) and the exact variables nudged will differ between models. Since there may be advantages to one nudging scheme, or choosing to nudge one variable over another, this would be a useful exercise. Testing this methodology with a range of models would also help to more definitively determine whether nudging horizontal winds suppresses significant circulation adjustments contributing to ERF, addressing the aforementioned question noted in Forster et al. (2021).

Testing the nudging mechanism denial methodology with different forcing agents would be another natural extension to this work. Radiative adjustments are typically a smaller source of uncertainty and better understood for other forcing agents than aerosols (Forster et al., 2021). However, some forcing agents, such as ozone, have been found to cause adjustments in circulation or via changing atmospheric heating profiles (MacIntosh et al., 2016) with few studies quantifying such effects and these adjustments could potentially be quantified with nudging. As with using different models, perturbing different forcing agents could also be useful to more definitively answer the question of whether uv-nudging suppressing circulation-mediated radiative adjustments significantly affects ERF, since the magnitude for circulation adjustments will vary by forcing agent, and this thesis has only addressed this question for aerosol perturbations.

Applying nudging to investigate aerosol radiative adjustments at regional scales would be an interesting extension to this work since, as discussed earlier, it may be a better use case for nudging as a mechanism denial tool. This is because while this thesis concludes that the nudging methodology used here is not likely to help understand the global aerosol response magnitude due to uncertainties introduced by spurious changes and inadequate θ adjustment suppression, nudging may be more useful at regional scales where there is generally greater uncertainty in aerosol effects already (Wilcox et al., 2023) and where studies may be interested in mechanisms more than global mean ERFs. This could include, for example, experiments perturbing aerosol emissions over a particular region (e.g. south Asia) which is expected to affect a particular climate process (e.g. the south Asian monsoon), with nudging applied to attempt to suppress these effects, isolating the effect of the aerosol perturbation. Or nudging could be applied over specific regions or latitudes only, with regional or global emission perturbations, to isolate the radiative adjustments unique to the nudged regions or latitudes. This potential avenue of extension is particularly relevant given the increasing interest in regional scale responses to aerosol perturbations (Wilcox et al., 2023). However, it is worth stressing the uncertainty associated with using nudging for mechanism denial, and it may be best used as an additional indicative tool rather than a precise method.

Investigating the effects of nudging on precipitation would be useful to understand whether nudging has application in understanding climate effects other than radiative effects. While this thesis has focused on radiative effects, precipitation is a key factor in climate change impacts (IPCC, 2023a). While the nudging mechanism denial methodology may not be effective for determining exact precipitation responses due to adjustments, it could help to see how precipitation responds broadly to suppressing adjustments. This would also help answer the question of whether uv-nudging suppresses significant circulation adjustments or not. Even if uv-nudging did not significantly alter ERF, studies of forcing agent perturbation interested in precipitation (or other responses) as well should avoid using nudging to decrease simulation integration length if precipitation is affected significantly.

Combining the use of nudging and PRP (or radiative kernels) in simulations would enable separating radiative adjustments by both adjustment mechanisms and by specific radiatively active variables. This would be useful to, for example, separate the circulation- and temperature-mediated adjustments in cloud with nudging and calculate the radiative effect of component of cloud adjustment with PRP. Likewise, using nudging in combination with PRP as implemented by Mülmenstädt et al. (2019) to separate out adjustments due to cloud amount, effective radius and liquid water paths would allow a proper separation of IRF_{aci}, the radiative adjustments comprising ERF_{aci}, and ERF_{ari}, unlike most methods that fail to separate the latter from aci adjustments. This would help decompose aerosol ERF more closely to the way it is decomposed in definitions and help quantify the relative magnitudes of aci and ari adjustments to aerosol forcing.

While not investigated in this work, the study of Martin and Quaas (2021) suggests there may be significant differences between use of radiative kernels calculated at different baseline conditions. For example, calculating the kernel from PI, PD or 2xCO_2 could produce different radiative adjustments when applied to outputs from control-perturbed simulation pairs. This is worth noting here since the need to conduct forward and backward calculations has been considered a weakness of PRP that does not apply to kernels, but calculating kernels at different climate states is a similar issue that has not been properly investigated in the literature. A simple study to investigate this could be using a pair of fSST simulations, one with PD climate, and another with 2xCO_2 as a perturbation. Two radiative kernels could then be calculated, by adding 1 K to temperatures, 1 percentage point to albedo, 1 percentage point to relative humidity etc, one using the PD simulation as a baseline the other using the perturbed simulation as a baseline. These kernels could then be applied to the adjustments between the control and perturbed simulation and compared to see if the use of different baselines for the kernels causes a significant difference in calculated radiative adjustments.

Finally, there could be great value in a review article that summarises the different methods of decomposing radiative adjustments, evaluates their relative merits for different applications, and compiles a list of resources to use those different methods. Currently in the literature there are a lot of studies using one, two or sometimes three different methods, sometimes in conjunction with each other and other times to evaluate one against the other. However, a review of all the major methods for decomposing radiative adjustments is lacking. Many of the methods are also non-trivial to implement and in some cases are implemented incorrectly (such as the implementation of APRP by Smith et al. (2020a) as rectified in Zelinka et al. (2023)); or have nuances and revisions noted or suggested in later studies than the original study outlining the method (such as the surface albedo adjustment term of Ghan (2013) which has later been noted to include adjustments other than just surface albedo (Zelinka et al., 2023)). These factors make it difficult to properly use or interpret these methods and could be particularly challenging for those uninitiated in the fields of aerosol climate effects and radiative adjustments from other forcing agents more generally.

A review article could address these issues by, firstly, comparing the application of the different methods to the same simulations, to ensure a standardised comparison of the different methods and their differences. Secondly, it could outline clearly what each method actually computes and how the quantities computed relate to formal definitions such as ERF_{aci} and ERF_{ari}. It could also include a decision tree to indicate what method is most appropriate given the diagnostics available, what adjustments are the focus of the study, and what forcing agent is being investigated. Thirdly, such a review could include links to existing tools and code to implement the different methods, as Zelinka et al. (2023) have recently done for the APRP method. Overall, this could make it easier for studies to select, implement and understand the appropriate method for their work.

5.3.3 Recommendations for quantifying radiative adjustments with different methods

Finally, reflecting on both the research conducted and the review of literature given in this thesis, a number of suggestions can be made about what methods to use for investigating radiative adjustments. These are listed as bullet points as follows:

- Studies should use consistent and clear definitions of radiative adjustments, particularly concerning cloud aci and ari responses (ideally using abbreviations such as ERFaci and ERFari exactly, rather than redefining them to include other adjustments).
- The use of diagnostics recommended in Ghan (2013) are highly useful, but studies should be clear to consider the masking effects outlined in Zelinka et al. (2023) and discussed in section 1.3.3.
- In most cases, for calculating radiative adjustments from individual variables, particularly non-cloud variables, radiative kernels are more effective than PRP due to the ease of use. However, studies should consider the differences between the radiation code in their particular model and that used to create the radiative kernel. This is particularly true for multi-model studies, which should consider if the aim is to sample wholly different models, and whether using radiative kernels calculated from another model obstructs the sampling of different models' responses.
- If PRP is desirable, the cumulative method presented here should produce comparable results but more efficiently than the conventional method presently used in most studies.
- For separating radiative adjustments by mechanisms rather than variables, use of model switches, such as those present in UKESM1 for the first and second indirect effects, is likely to be more reliable than using nudging, and less sensitive to the setup used.
- For some use cases nudging may be acceptable for reducing simulation integration time for aerosol ERF studies (e.g. if only aci adjustments are considered). However, if other cloud adjustments are of interest, based on the results in this thesis it is not recommendable to use nudging to reduce simulation integration time. If nudging is still intended to be used for this purpose, studies should compare the nudged simulation to the input used (which should be output from the same model) to see what effect the nudging is having beyond reducing natural variability.

Appendix A

Explanation of the deviation of $\triangle$DRE calculated with the double call method from IRF

Here is given a further explanation of the final point made in section 1.3.3, specifically that ΔDRE calculated using the method outlined in Ghan (2013) is slightly different from IRF as defined in Forster et al. (2021). Using the terminology of Zelinka et al. (2023), one can first define the radiative effect of a perturbation of a variable, X, between a control and perturbed as:

$$K^X \Delta X = K^X X_p - K^X X_c \tag{A.1}$$

where ΔX is the change in X between the control (X_c) and perturbed state (X_p) and K^X is defined as in equation 1.4 as a factor describing the change in net downwards TOA radiative flux with infinitesimal change in variable X.

Since the radiative effect of a change in X will depend on the atmospheric state, K^X can be considered a function of clouds, aerosols, surface albedo, atmospheric temperature and specific humidity and can be more precisely written as $K^X(C, A, \alpha, T, q)$. Calculation of aerosol IRF assumes that no adjustments takes place, and so can be defined as:

$$IRF = K^A(C_c, A_c, \alpha_c, T_c, q_c)A_p - K^A(C_c, A_c, \alpha_c, T_c, q_c)A_c \tag{A.2}$$

where subscript c indicates the control state of a variable. However, equation 1.10 can be rearranged to:

$$\Delta DRE = (F_p - F_c) - (F_{p,cln} - F_{c,cln}) \tag{A.3}$$

where subscript p indicates the control state of a variable. The first term in equation A.3 can be described as the ERF and the latter term as the forcing due to changes in all variables excluding aerosols (and excluding the masking effect of aerosol on these changes). It can therefore be seen that the aerosol forcing effect is derived from the first term, in which it is calculated as the difference in radiative effect of the perturbed concentration of aerosol with all other variables set to values from the perturbed simulation minus the radiative effect of the control concentration of aerosol with all other variables to values from the control simulation. Thus in the method of Ghan (2013) the aerosol radiative effect, ΔDRE, is calculated as:

$$\Delta DRE = K^A(C_p, A_p, \alpha_p, T_p, q_p)A_p - K^A(C_c, A_c, \alpha_c, T_c, q_c)A_c \tag{A.4}$$

The difference between equations A.2 and A.4 is subtle and likely to be smaller than the aerosol masking term outlined in Zelinka et al. (2023) that is excluded from ΔDRE. It also does not make it incorrect to use ΔDRE. As noted earlier, one could interchange the control and perturbed state in fSST experiments and validly argue either arrangement was the most useful calculation of ERF. Here one could argue that including the change in the masking effect of other variables due to them being perturbed when calculating the direct radiative forcing due to aerosol is more representative of the actual forcing due to the aerosol change than calculating it on a fixed control atmospheric state.

Appendix B

Cross sections of u, v and θ adjustments and control errors in nudged simulations

Here are presented plots of the adjustments and control errors in u, vm and θ from a subset of the different nudging setups explored in 2 and outlined in table 2.1. While the level mean profiles presented in figures 2.5 and 2.6 are generally sufficient, there are some features obscured by averaging over levels and presenting horizontal wind adjustments and control errors as an RMS. As such, the figures here are included to help illustrate the spatial variation of these adjustments and control errors in different dimensions. For the optimum nudging setup (G=1/6 h^{-1} with lower level nudging), maps over 5 models levels and a zonal mean of adjustments of θ (fig B.1), u (fig. B.2), and v (fig. B.3) and control errors of θ (fig. B.4), u (fig. B.5), and v (fig. B.6) are illustrated.

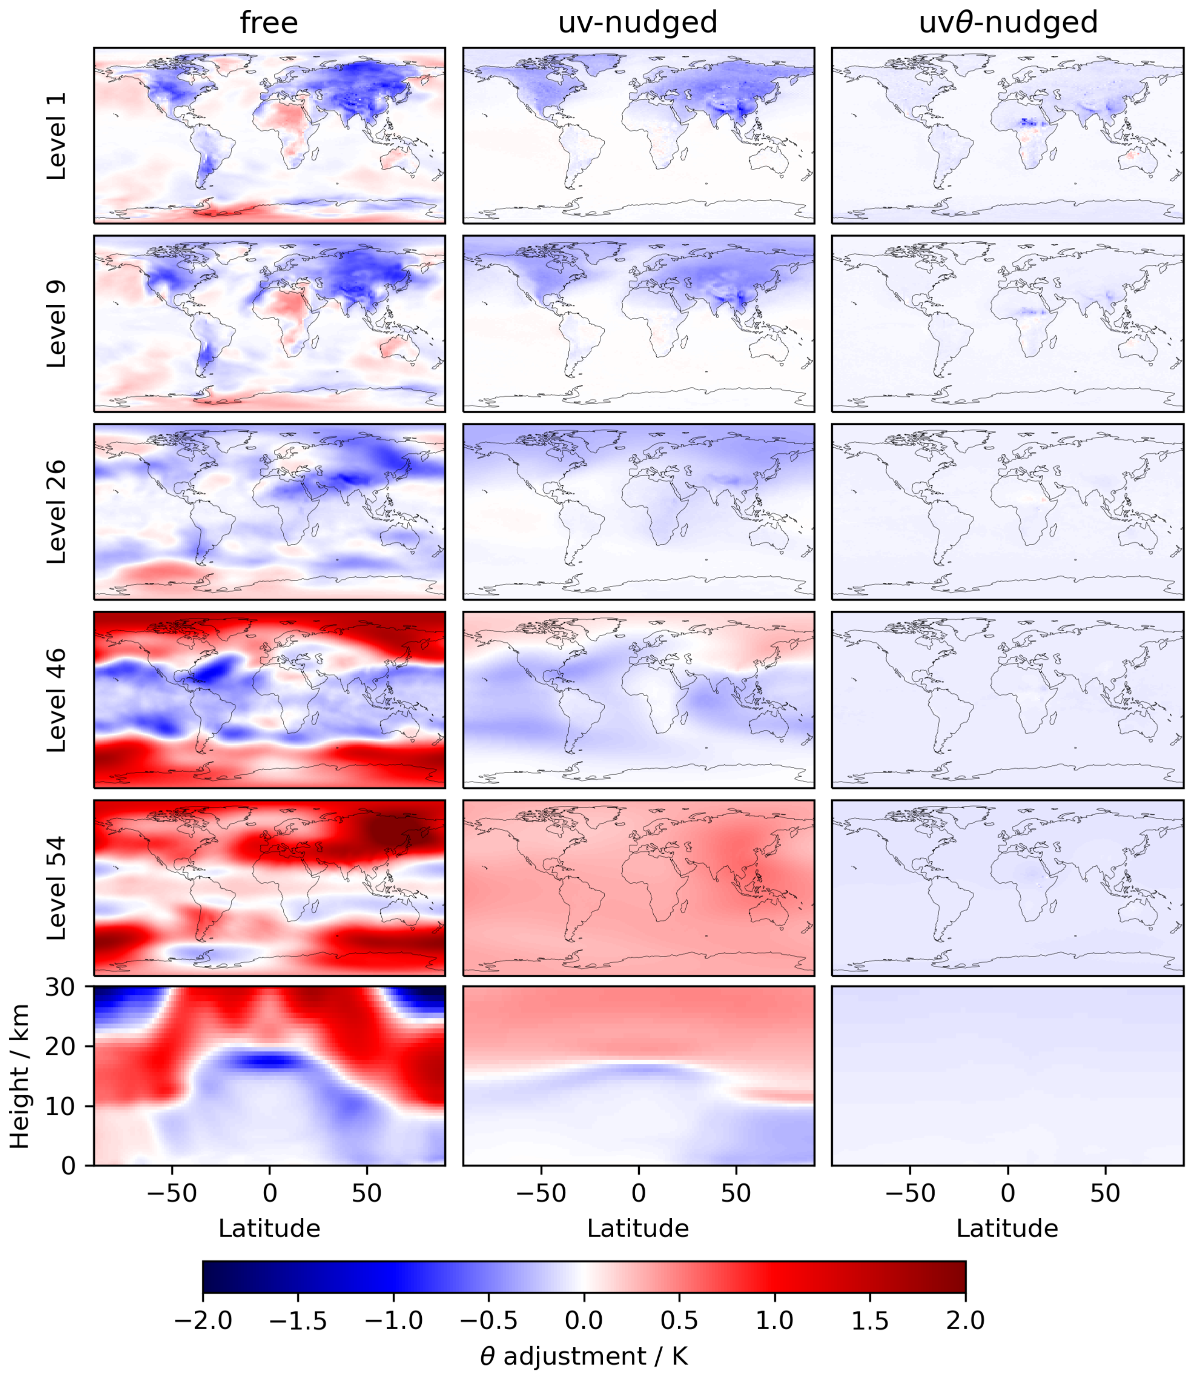

Figure B.1: Multiannual mean adjustment in θ over a range of model levels and as a zonal mean for the free, uv-nudged and uvθ-nudged simulation pairs. The nudged simulations use the optimal nudging case identified in table 2.1. Model levels are 1, 9, 26, 46, and 54 corresponding to heights amsl of 20 m, 660 m, 4.85 km, 14.72 km and 20.23 km respectively.

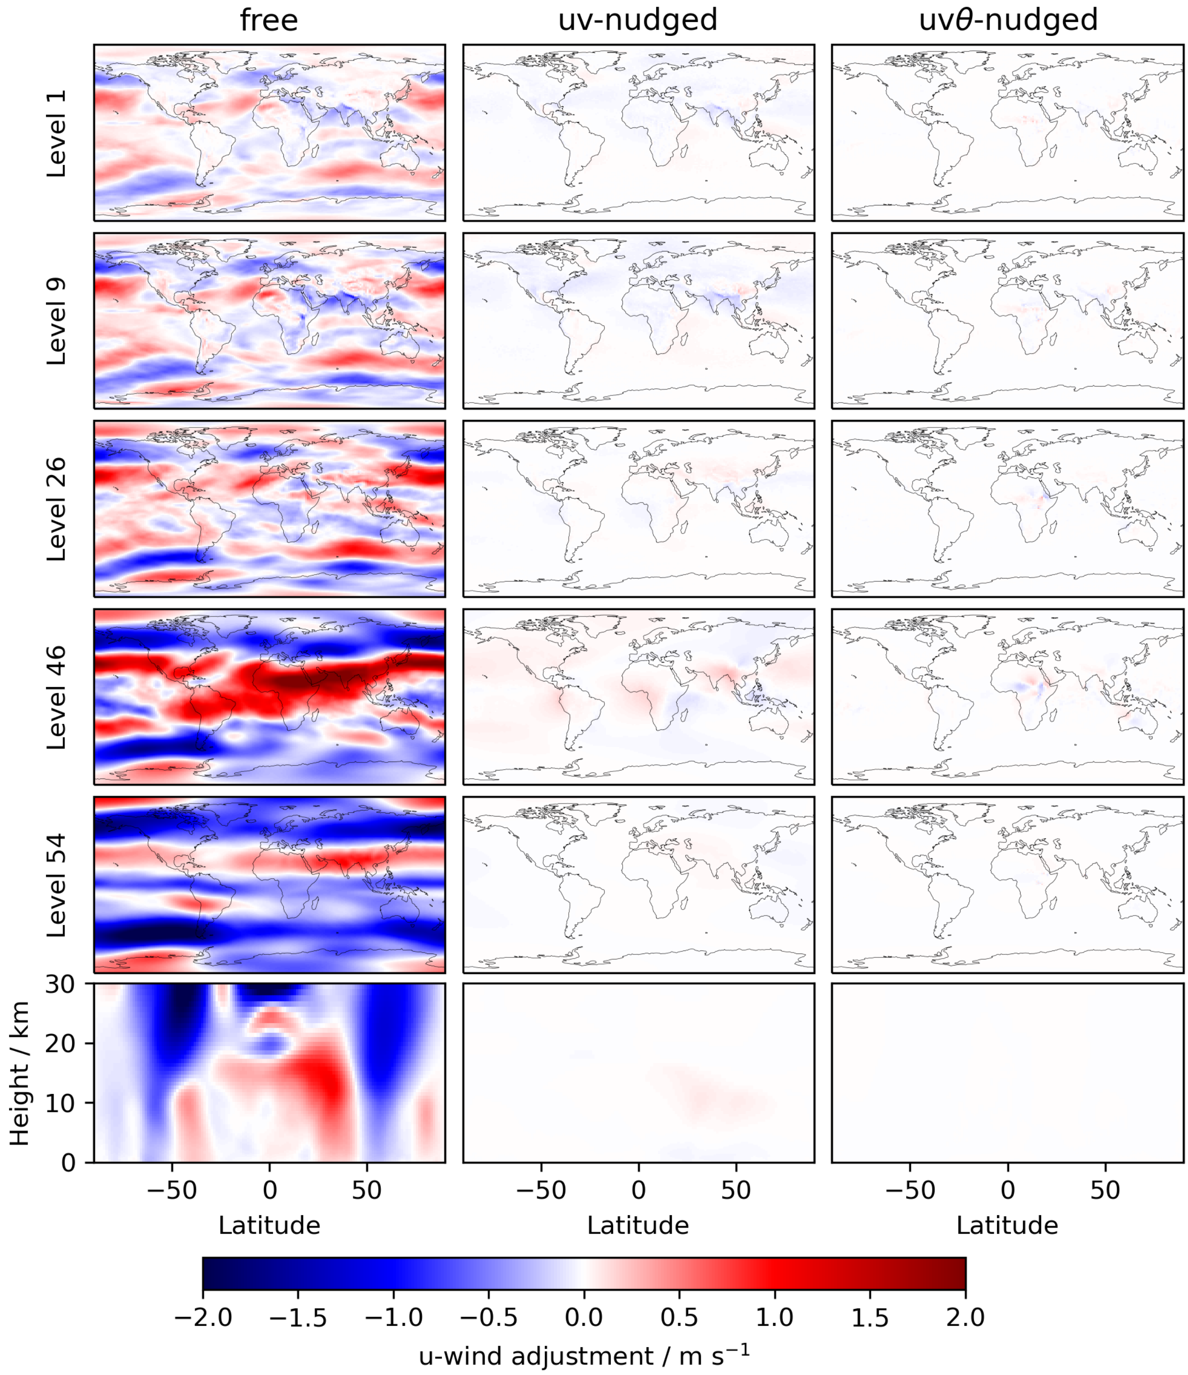

Figure B.2: As for figure B.1 but for u-wind adjustment.

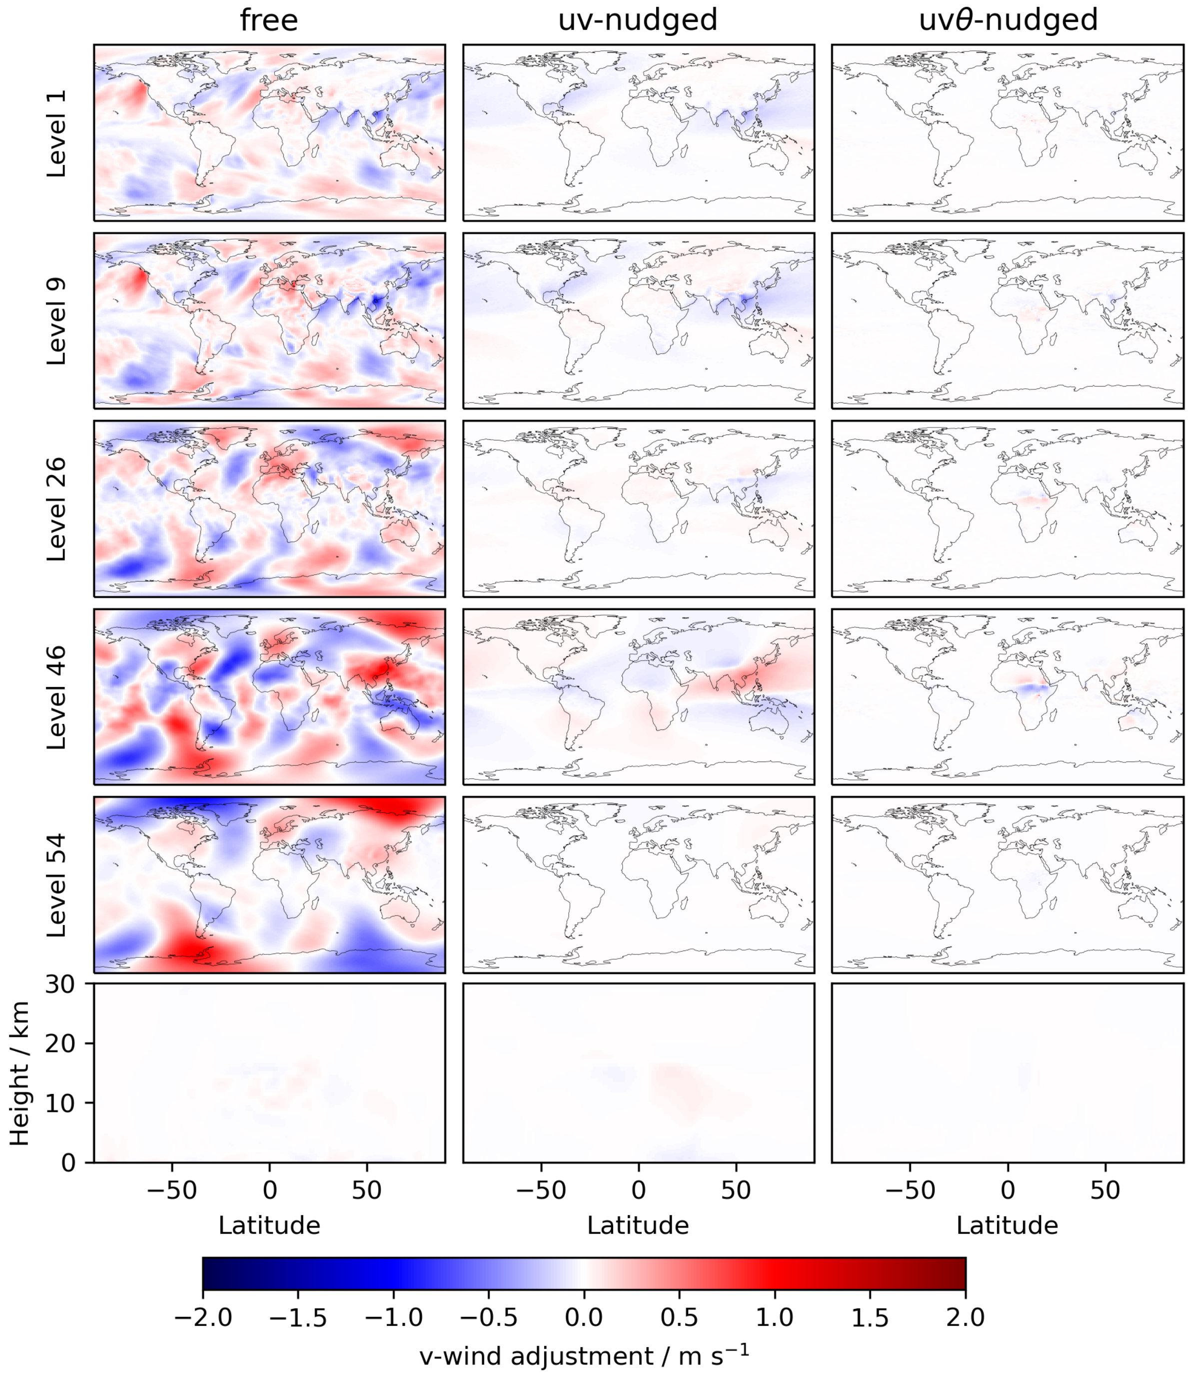

Figure B.3: As for figure B.1 but for v-wind adjustment.

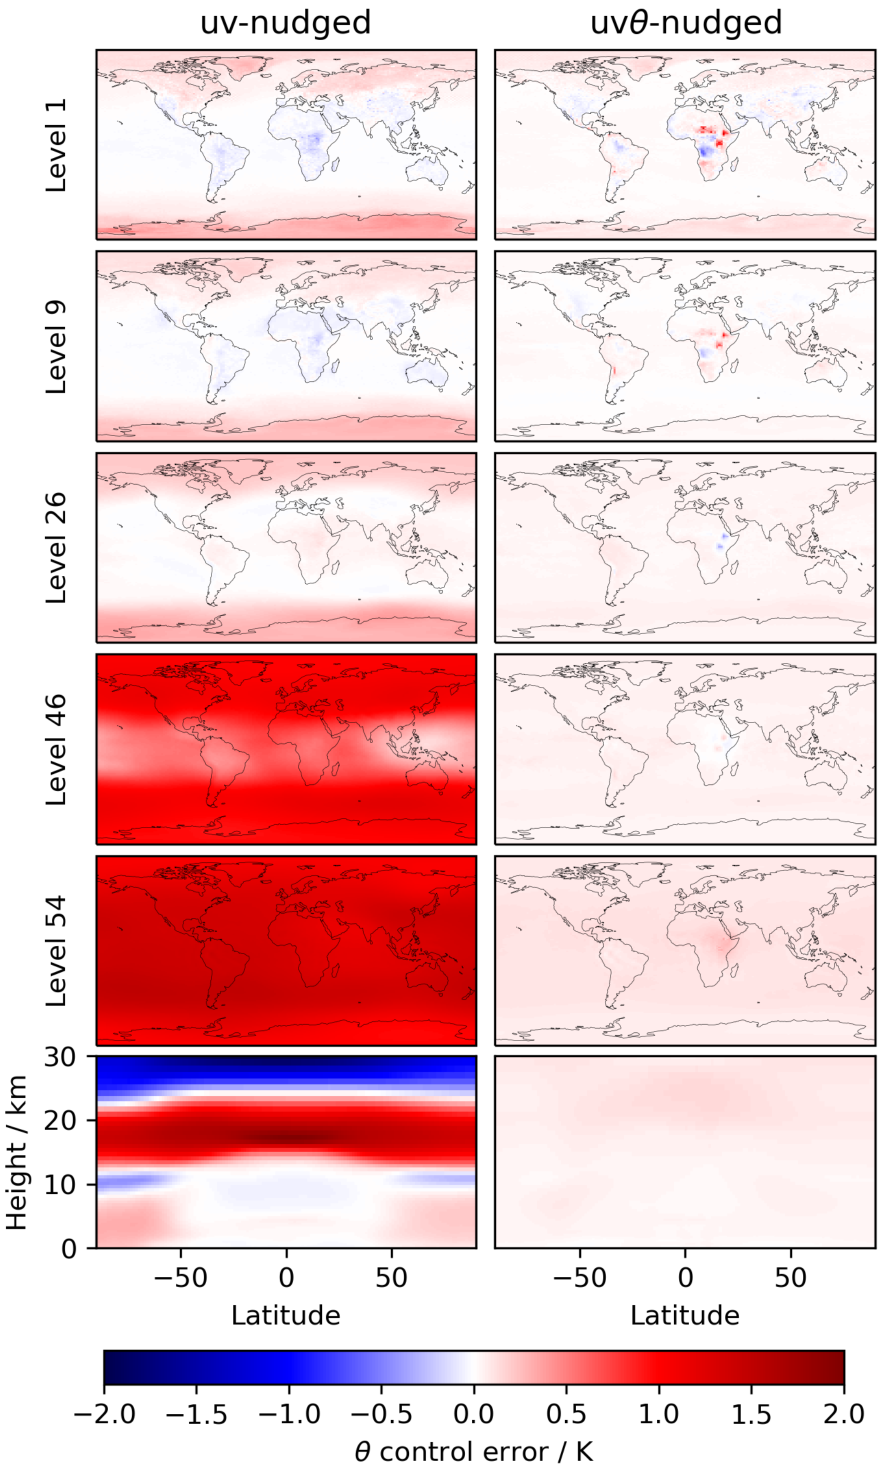

Figure B.4: Multiannual mean control error in θ over a range of model levels and as a zonal mean for the uv-nudged and uvθ-nudged control simulations. Control error is calculated as the difference in the field between the free control minus in the nudged control. The nudged simulations use the optimal nudging case identified in table 2.1. Model levels are 1, 9, 26, 46, and 54 corresponding to heights amsl of 20 m, 660 m, 4.85 km, 14.72 km and 20.23 km respectively.

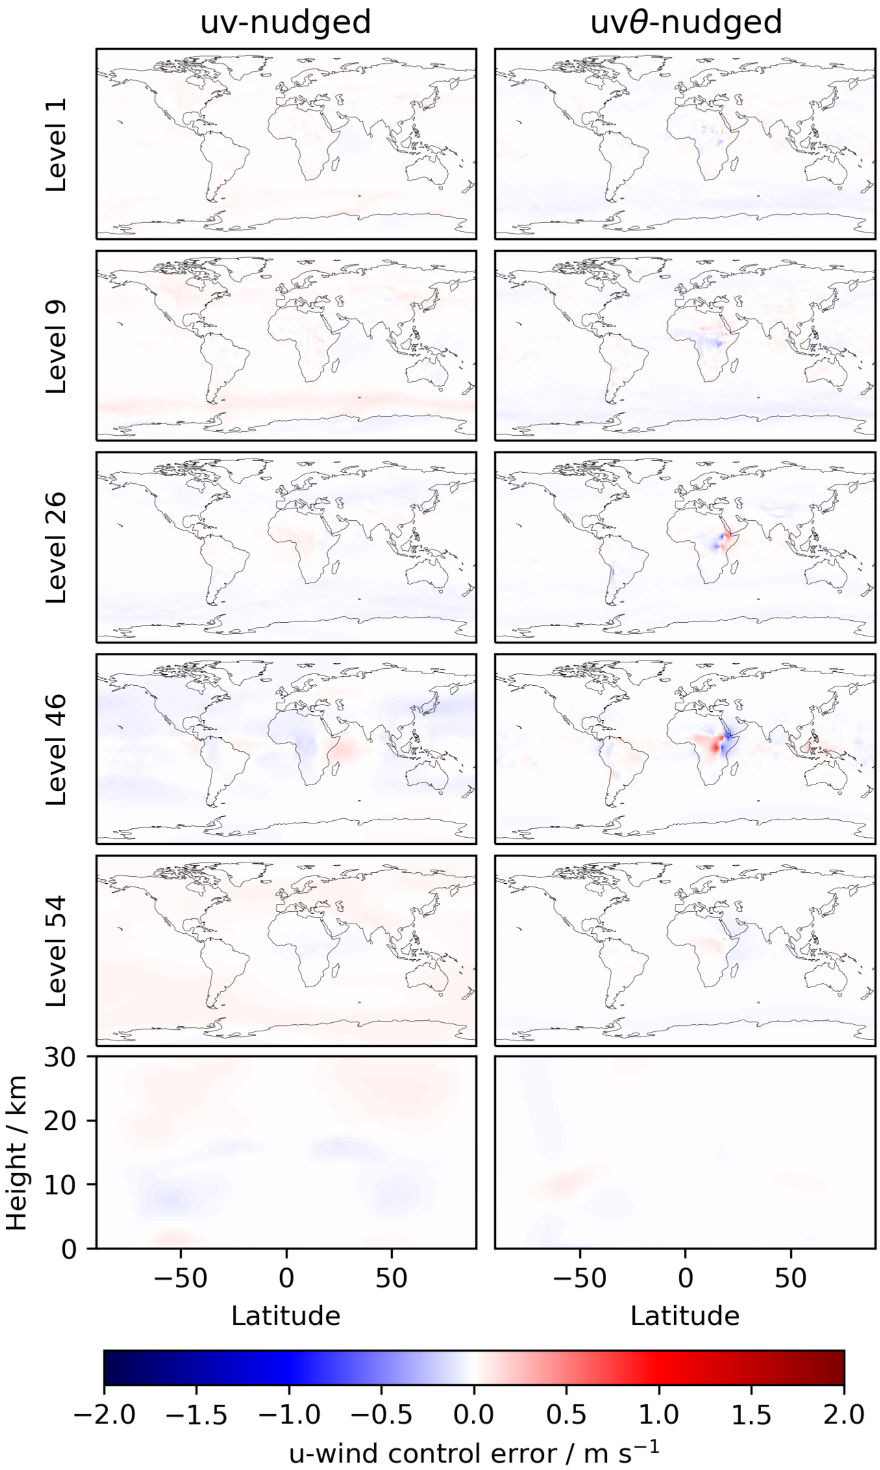

Figure B.5: As for figure B.4 but for u-wind control error.

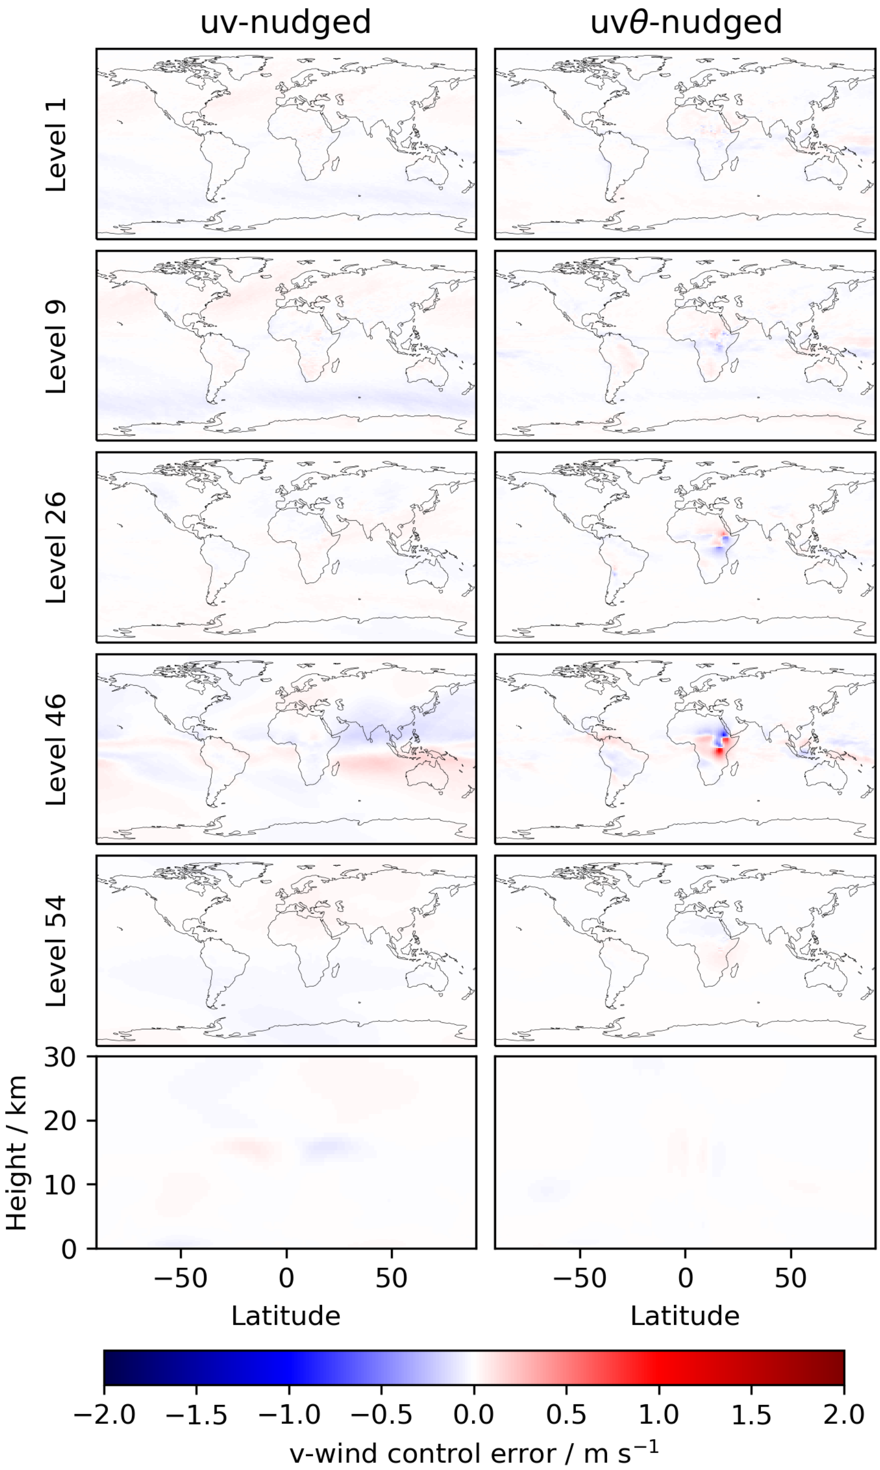

Figure B.6: As for figure B.4 but for v-wind control error.

Bibliography

Ackerley, D., and D. Dommenget, 2016: Atmosphere-only GCM (ACCESS1.0) simulations with prescribed land surface temperatures. *Geoscientific Model Development*, **9 (6)**, 2077–2098, https://doi.org/10.5194/gmd-9-2077-2016, URL 10.5194/gmd-9-2077-2016.

Albrecht, B. A., 1989: Aerosols, Cloud Microphysics, and Fractional Cloudiness. *Science*, **245 (4923)**, 1227–1230, https://doi.org/10.1126/science.245.4923.1227, URL 10.1126/science.245.4923.1227.

Allen, R. J., A. Amiri-Farahani, J.-F. Lamarque, C. Smith, D. Shindell, T. Hassan, and C. E. Chung, 2019: Observationally constrained aerosolcloud semi-direct effects. *Climate and Atmospheric Science*, **2 (1)**, 16, https://doi.org/10.1038/s41612-019-0073-9, URL 10.1038/s41612-019-0073-9.

Andrews, T., C. J. Smith, G. Myhre, P. M. Forster, R. Chadwick, and D. Ackerley, 2021: Effective Radiative Forcing in a GCM With Fixed Surface Temperatures. *Journal of Geophysical Research: Atmospheres*, **126 (4)**, https://doi.org/10.1029/2020JD033880, URL 10.1029/2020JD033880.

Bellouin, N., and Coauthors, 2020a: Bounding Global Aerosol Radiative Forcing of Climate Change. *Reviews of Geophysics*, **58 (1)**, https://doi.org/10.1029/2019RG000660, URL 10.1029/2019RG000660.

Bellouin, N., and Coauthors, 2020b: Radiative forcing of climate change from the Copernicus reanalysis of atmospheric composition. *Earth System Science Data*, **12 (3)**, 1649–1677, https://doi.org/10.5194/essd-12-1649-2020, URL 10.5194/essd-12-1649-2020.

Bickel, M., 2023: Climate Impact of Contrails Cirrus. Ph.D. thesis, Ludwig-Maximilians-Universität München, 1–154 pp., https://doi.org/10.57676/mzmg-r403, URL 10.57676/mzmg-r403.

Bickel, M., M. Ponater, L. Bock, U. Burkhardt, and S. Reineke, 2020: Estimating the Effective Radiative Forcing of Contrail Cirrus. *Journal of Climate*, **33 (5)**, 1991–2005, https://doi.org/10.1175/JCLI-D-19-0467.1, URL 10.1175/JCLI-D-19-0467.1.

Block, K., M. Haghighatnasab, D. G. Partridge, P. Stier, and J. Quaas, 2024: Cloud condensation nuclei concentrations derived from the CAMS reanalysis. *Earth System Science Data*, **16 (1)**, 443–470, https://doi.org/10.5194/essd-16-443-2024, URL 10.5194/essd-16-443-2024.

Bond, T. C., and Coauthors, 2013: Bounding the role of black carbon in the climate system: A scientific assessment. *Journal of Geophysical Research: Atmospheres*, **118 (11)**, 5380–5552, https://doi.org/10.1002/jgrd.50171, URL 10.1002/jgrd.50171.

Boucher, O., and Coauthors, 2013: Clouds and Aerosols. *Climate Change 2013 The Physical Science Basis*, Cambridge University Press, Cambridge, 571–658, https://doi.org/10.1017/CBO9781107415324.016, URL 10.1017/CBO9781107415324.016.

Brill, K. F., L. W. Uccellini, J. Manobianco, P. J. Kocin, and J. H. Homan, 1991: The use of successive dynamic initialization by nudging to simulate cyclogenesis during GALE IOP 1. *Meteorology and Atmospheric Physics*, **45 (1-2)**, 15–40, https://doi.org/10.1007/BF01027473, URL 10.1007/BF01027473.

Brown, A., S. Milton, M. Cullen, B. Golding, J. Mitchell, and A. Shelly, 2012: Unified Modeling and Prediction of Weather and Climate: A 25-Year Journey. *Bulletin of the American Meteorological Society*, **93 (12)**, 1865–1877, https://doi.org/10.1175/BAMS-D-12-00018.1, URL 10.1175/BAMS-D-12-00018.1.

Che, H., P. Stier, H. Gordon, D. Watson-Parris, and L. Deaconu, 2021: Cloud adjustments dominate the overall negative aerosol radiative effects of biomass burning aerosols in UKESM1 climate model simulations over the south-eastern Atlantic. *Atmospheric Chemistry and Physics*, **21 (1)**, 17–33, https://doi.org/10.5194/acp-21-17-2021, URL https://doi.org/10.5194/acp-21-17-2021.

Clarke, A. D., and K. J. Noone, 1985: Soot in the Arctic snowpack: a cause for perturbations in radiative transfer. *Atmospheric Environment (1967)*, **19 (12)**, 2045–2053, https://doi.org/10.1016/0004-6981(85)90113-1, URL 10.1016/0004-6981(85)90113-1.

Collins, W. J., and Coauthors, 2017: AerChemMIP: quantifying the effects of chemistry and aerosols in CMIP6. *Geosci. Model Dev*, **10**, 585–607, https://doi.org/10.5194/gmd-10-585-2017, URL 10.5194/gmd-10-585-2017.

Colman, R. A., 2015: Climate radiative feedbacks and adjustments at the Earth's surface. *Journal of Geophysical Research: Atmospheres*, **120 (8)**, 3173–3182, https://doi.org/10.1002/2014JD022896, URL 10.1002/2014JD022896.

Colman, R. A., J. Fraser, and L. Rotstayn, 2001: Climate feedbacks in a general circulation model incorporating prognostic clouds. *Climate Dynamics*, **18 (1-2)**, 103–122, https://doi.org/10.1007/s003820100162.

Colman, R. A., and B. J. McAvaney, 1997: A study of general circulation model climate feedbacks determined from perturbed sea surface temperature experiments. *Journal of Geophysical Research: Atmospheres*, **102 (D16)**, 19 383–19 402, https://doi.org/10.1029/97JD00206, URL 10.1029/97JD00206.

Colman, R. A., S. B. Power, and B. J. McAvaney, 1997: Non-linear climate feedback analysis in an atmospheric general circulation model. *Climate Dynamics*, **13 (10)**, 717–731, https://doi.org/10.1007/s003820050193, URL 10.1007/s003820050193.

Dalvi, M., and J. Rodriguez, 2020: Unified Model Documentation Paper 083: Nudging the Unified Model. Tech. rep., 1–4 pp. URL https://code.metoffice.gov.uk/doc/um/latest/papers/umdp{_}083.pdf.

Davis, N. A., P. Callaghan, I. R. Simpson, and S. Tilmes, 2022: Specified dynamics scheme impacts on wave-mean flow dynamics, convection, and tracer transport in CESM2 (WACCM6). *Atmospheric Chemistry and Physics*, **22 (1)**, 197–214, https://doi.org/10.5194/acp-22-197-2022, URL 10.5194/acp-22-197-2022.

Davis, N. A., S. M. Davis, R. W. Portmann, E. Ray, K. H. Rosenlof, and P. Yu, 2020: A comprehensive assessment of tropical stratospheric upwelling in the specified dynamics Community Earth System Model 1.2.2 Whole Atmosphere Community Climate Model (CESM (WACCM)). *Geoscientific Model Development*, **13 (2)**, 717–734, https://doi.org/10.5194/gmd-13-717-2020, URL 10.5194/gmd-13-717-2020.

Dean, S. M., J. Flowerdew, B. N. Lawrence, and S. D. Eckermann, 2007: Parameterisation of orographic cloud dynamics in a GCM. *Climate Dynamics*, **28 (6)**, 581–597, https://doi.org/10.1007/s00382-006-0202-0, URL 10.1007/s00382-006-0202-0.

Dong, B., R. T. Sutton, E. J. Highwood, and L. J. Wilcox, 2016: Preferred response of the East Asian summer monsoon to local and non-local anthropogenic sulphur dioxide emissions. *Climate Dynamics*, **46 (5-6)**, 1733–1751, https://doi.org/10.1007/s00382-015-2671-5, URL 10.1007/s00382-015-2671-5.

Fahrenbach, N. L. S., and M. A. Bollasina, 2023: Hemispheric-wide climate response to regional COVID-19-related aerosol emission reductions: the prominent role of atmospheric circulation adjustments. *Atmospheric Chemistry and Physics*, **23 (2)**, 877–894, https://doi.org/10.5194/acp-23-877-2023.

Flanner, M. G., C. S. Zender, P. G. Hess, N. M. Mahowald, T. H. Painter, V. Ramanathan, and P. J. Rasch, 2009: Springtime warming and reduced snow cover from carbonaceous particles. *Atmospheric Chemistry and Physics*, **9 (7)**, 2481–2497, https://doi.org/10.5194/acp-9-2481-2009, URL 10.5194/acp-9-2481-2009.

Forster, P., and Coauthors, 2021: The Earth's Energy Budget, Climate Feedbacks and Climate Sensitivity. *Climate Change 2021 The Physical Science Basis*, Cambridge University Press, 923–1054, https://doi.org/10.1017/9781009157896.009.

Forster, P. M., M. Blackburn, R. Glover, and K. P. Shine, 2000: An examination of climate sensitivity for idealised climate change experiments in an intermediate general circulation model. *Climate Dynamics*, **16 (10-11)**, 833–849, https://doi.org/10.1007/s003820000083, URL 10.1007/s003820000083.

Forster, P. M., and Coauthors, 2016: Recommendations for diagnosing effective radiative forcing from climate models for CMIP6. *Journal of Geophysical Research: Atmospheres*, **121 (20)**, 12,412–460,475, https://doi.org/10.1002/2016JD025320, URL 10.1002/2016JD025320.

Gettelman, A., and Coauthors, 2020: Simulating Observations of Southern Ocean Clouds and Implications for Climate. *Journal of Geophysical Research: Atmospheres*, **125 (21)**, e2020JD032 619, https://doi.org/10.1029/2020JD032619, URL 10.1029/2020JD032619.

Ghan, S., and Coauthors, 2016: Challenges in constraining anthropogenic aerosol effects on cloud radiative forcing using present-day spatiotemporal variability. *Proceedings of the National Academy of Sciences*, **113 (21)**, 5804–5811, https://doi.org/10.1073/pnas.1514036113, URL 10.1073/pnas.1514036113.

Ghan, S. J., 2013: Technical Note: Estimating aerosol effects on cloud radiative forcing. *Atmospheric Chemistry and Physics*, **13 (19)**, 9971–9974, https://doi.org/10.5194/acp-13-9971-2013, URL 10.5194/acp-13-9971-2013.

Gregory, J. M., and Coauthors, 2004: A new method for diagnosing radiative forcing and climate sensitivity. *Geophysical Research Letters*, **31 (3)**, 3205, https://doi.org/10.1029/2003GL018747, URL 10.1029/2003GL018747.

Grosvenor, D. P., and K. S. Carslaw, 2020: The decomposition of cloudaerosol forcing in the UK Earth System Model (UKESM1). *Atmospheric Chemistry and Physics*, **20 (24)**, 15 681–15 724, https://doi.org/10.5194/acp-20-15681-2020, URL 10.5194/acp-20-15681-2020.

Gryspeerdt, E., and Coauthors, 2020: Surprising similarities in model and observational aerosol radiative forcing estimates. *Atmospheric Chemistry and Physics*, **20 (1)**, 613–623, https://doi.org/10.5194/acp-20-613-2020, URL 10.5194/acp-20-613-2020.

Hansen, J., and L. Nazarenko, 2004: Soot climate forcing via snow and ice albedos. *Proceedings of the National Academy of Sciences*, **101 (2)**, 423–428, https://doi.org/10.1073/pnas.2237157100, URL 10.1073/pnas.2237157100.

Hansen, J., M. Sato, and R. Ruedy, 1997: Radiative forcing and climate response. *Journal of Geophysical Research: Atmospheres*, **102 (D6)**, 6831–6864, https://doi.org/10.1029/96JD03436, URL 10.1029/96JD03436.

Hansen, J., and Coauthors, 2005: Efficacy of climate forcings. *Journal of Geophysical Research: Atmospheres*, **110 (D18)**, 1–45, https://doi.org/10.1029/2005JD005776, URL 10.1029/2005JD005776.

Held, I. M., and B. J. Soden, 2000: Water Vapor Feedback and Global Warming. *Annual Review of Energy and the Environment*, **25 (1)**, 441–475, https://doi.org/10.1146/annurev.energy.25.1.441, URL 10.1146/annurev.energy.25.1.441.

Herbert, R. J., N. Bellouin, E. J. Highwood, and A. A. Hill, 2020: Diurnal cycle of the semi-direct effect from a persistent absorbing aerosol layer over marine stratocumulus in large-eddy simulations. *Atmospheric Chemistry and Physics*, **20 (3)**, 1317–1340, https://doi.org/10.5194/acp-20-1317-2020, URL 10.5194/acp-20-1317-2020.

Hoesly, R. M., and Coauthors, 2018: Historical (1750-2014) anthropogenic emissions of reactive gases and aerosols from the Community Emissions Data System (CEDS). *Geoscientific Model Development*, **11 (1)**, 369–408, https://doi.org/10.5194/gmd-11-369-2018, URL 10.5194/gmd-11-369-2018.

Huang, Y., Y. Wang, and H. Huang, 2020: Stratospheric Water Vapor Feedback Disclosed by a Locking Experiment. *Geophysical Research Letters*, **47 (12)**, https://doi.org/10.1029/2020GL087987, URL 10.1029/2020GL087987.

IPCC, 2023a: IPCC, 2023: Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, H. Lee and J. Romero (eds.)]. IPCC, Geneva, Switzerland. Tech. rep., Intergovernmental Panel on Climate Change. https://doi.org/10.59327/IPCC/AR6-9789291691647, URL 10.59327/IPCC/AR6-9789291691647.

IPCC, 2023b: Technical Summary. *Climate Change 2021 The Physical Science Basis*, Cambridge University Press, 35–144, https://doi.org/10.1017/9781009157896.002, URL 10.1017/9781009157896.002.

Jeuken, A. B. M., P. C. Siegmund, L. C. Heijboer, J. Feichter, and L. Bengtsson, 1996: On the potential of assimilating meteorological analyses in a global climate model for the purpose of model validation. *Journal of Geophysical Research: Atmospheres*, **101 (D12)**, 16 939–16 950, https://doi.org/10.1029/96JD01218, URL 10.1029/96JD01218.

Johnson, B. T., J. M. Haywood, and M. K. Hawcroft, 2019: Are Changes in Atmospheric Circulation Important for Black Carbon Aerosol Impacts on Clouds, Precipitation, and Radiation? *Journal of Geophysical Research: Atmospheres*, **124 (14)**, 7930–7950, https://doi.org/10.1029/2019JD030568, URL 10.1029/2019JD030568.

Johnson, B. T., K. P. Shine, and P. M. Forster, 2004: The semi-direct aerosol effect: Impact of absorbing aerosols on marine stratocumulus. *Quarterly Journal of the Royal Meteorological Society*, **130 (599)**, 1407–1422, https://doi.org/10.1256/qj.03.61, URL 10.1256/qj.03.61.

Klein, S. A., and C. Jakob, 1999: Validation and Sensitivities of Frontal Clouds Simulated by the ECMWF Model. *Monthly Weather Review*, **127 (10)**, 2514–2531, https://doi.org/10.1175/1520-0493(1999)127⟨2514:VASOFC⟩2.0.CO;2, URL 10.1175/1520-0493(1999)127{%}3C2514:VASOFC{%}3E2.0.CO;2.

Koch, D., and A. D. Del Genio, 2010: Black carbon semi-direct effects on cloud cover: review and synthesis. *Atmospheric Chemistry and Physics*, **10 (16)**, 7685–7696, https://doi.org/10.5194/acp-10-7685-2010, URL 10.5194/acp-10-7685-2010.

Koch, D., and Coauthors, 2011: Soot microphysical effects on liquid clouds, a multi-model investigation. *Atmospheric Chemistry and Physics*, **11 (3)**, 1051–1064, https://doi.org/10.5194/acp-11-1051-2011, URL https://acp.copernicus.org/articles/11/1051/2011/.

Kooperman, G. J., M. S. Pritchard, S. J. Ghan, M. Wang, R. C. J. Somerville, and L. M. Russell, 2012: Constraining the influence of natural variability to improve estimates of global aerosol indirect effects in a nudged version of the Community Atmosphere Model 5. *Journal of Geophysical Research: Atmospheres*, **117 (D23)**, https://doi.org/10.1029/2012JD018588, URL 10.1029/2012JD018588.

Li, J., and Coauthors, 2022: Scattering and absorbing aerosols in the climate system. *Nature Reviews Earth & Environment*, **3 (6)**, 363–379, https://doi.org/10.1038/s43017-022-00296-7, URL 10.1038/s43017-022-00296-7.

Li, Y., R. Ménard, L. P. Riishøjgaard, S. E. Cohn, and R. B. Rood, 1998: A study on assimilating potential vorticity data. *Tellus A: Dynamic Meteorology and Oceanography*, **50 (4)**, 490, https://doi.org/10.3402/tellusa.v50i4.14535, URL 10.3402/tellusa.v50i4.14535.

Lin, G., J. E. Penner, M. G. Flanner, S. Sillman, L. Xu, and C. Zhou, 2014: Radiative forcing of organic aerosol in the atmosphere and on snow: Effects of SOA and brown carbon. *Journal of Geophysical Research: Atmospheres*, **119 (12)**, 7453–7476, https://doi.org/10.1002/2013JD021186, URL 10.1002/2013JD021186.

Lin, G., H. Wan, K. Zhang, Y. Qian, and S. J. Ghan, 2016: Can nudging be used to quantify model sensitivities in precipitation and cloud forcing? *Journal of Advances in Modeling Earth Systems*, **8 (3)**, 1073–1091, https://doi.org/10.1002/2016MS000659, URL 10.1002/2016MS000659.

MacIntosh, C. R., R. P. Allan, L. H. Baker, N. Bellouin, W. Collins, Z. Mousavi, and K. P. Shine, 2016: Contrasting fast precipitation responses to tropospheric and stratospheric ozone forcing. *Geophysical Research Letters*, **43 (3)**, 1263–1271, https://doi.org/10.1002/2015GL067231, URL 10.1002/2015GL067231.

Martin, A., and J. Quaas, 2021: Determination of radiation couplings in climate change simulations: Analysis with two different linearization methods. Ph.D. thesis, University of Leipzig, 1–14 pp., URL https://ul.qucosa.de/api/qucosa{%}3A78017/attachment/ATT-0/.

Meinshausen, M., and Coauthors, 2017: Historical greenhouse gas concentrations for climate modelling (CMIP6). *Geoscientific Model Development*, **10 (5)**, 2057–2116, https://doi.org/10.5194/gmd-10-2057-2017, URL 10.5194/gmd-10-2057-2017.

Merryfield, W. J., and Coauthors, 2013: The Canadian Seasonal to Interannual Prediction System. Part I: Models and Initialization. *Monthly Weather Review*, **141 (8)**, 2910–2945, https://doi.org/10.1175/MWR-D-12-00216.1, URL 10.1175/MWR-D-12-00216.1.

Morgenstern, O., P. Braesicke, F. M. O'Connor, A. C. Bushell, C. E. Johnson, S. M. Osprey, and J. A. Pyle, 2009: Evaluation of the new UKCA climate-composition model Part 1: The stratosphere. *Geoscientific Model Development*, **2 (1)**, 43–57, https://doi.org/10.5194/gmd-2-43-2009, URL 10.5194/gmd-2-43-2009.

Mülmenstädt, J., E. Gryspeerdt, M. Salzmann, P.-L. Ma, S. Dipu, and J. Quaas, 2019: Separating radiative forcing by aerosolcloud interactions and rapid cloud adjustments in the ECHAMHAMMOZ aerosol-climate model using the method of partial radiative perturbations. *Atmospheric Chemistry and Physics*, **19 (24)**, 15 415–15 429, https://doi.org/10.5194/acp-19-15415-2019, URL 10.5194/acp-19-15415-2019.

Myhre, G., and Coauthors, 2013: Anthropogenic and Natural Radiative Forcing: In Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change. *Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA*, 659–740, URL https://www.ipcc.ch/site/assets/uploads/2018/02/WG1AR5{_}Chapter08{_}FINAL.pdf.

Myhre, G., and Coauthors, 2017: PDRMIP: A Precipitation Driver and Response Model Intercomparison ProjectProtocol and Preliminary Results. *Bulletin of the American Meteorological Society*, **98 (6)**, 1185–1198, https://doi.org/10.1175/BAMS-D-16-0019.1, URL 10.1175/BAMS-D-16-0019.1.

Namazi, M., K. von Salzen, and J. N. S. Cole, 2015: Simulation of black carbon in snow and its climate impact in the Canadian Global Climate Model. *Atmospheric Chemistry and Physics*, **15 (18)**, 10 887–10 904, https://doi.org/10.5194/acp-15-10887-2015, URL 10.5194/acp-15-10887-2015.

O'Connor, F. M., and Coauthors, 2014: Evaluation of the new UKCA climate-composition model Part 2: The Troposphere. *Geoscientific Model Development*, **7 (1)**, 41–91, https://doi.org/10.5194/gmd-7-41-2014, URL 10.5194/gmd-7-41-2014.

O'Connor, F. M., and Coauthors, 2021: Assessment of pre-industrial to present-day anthropogenic climate forcing in UKESM1. *Atmospheric Chemistry and Physics*, **21 (2)**, 1211–1243, https://doi.org/10.5194/acp-21-1211-2021, URL 10.5194/acp-21-1211-2021.

Omrani, H., P. Drobinski, and T. Dubos, 2012: Investigation of indiscriminate nudging and predictability in a nested quasigeostrophic model. *Quarterly Journal of the Royal Meteorological Society*, **138 (662)**, 158–169, https://doi.org/10.1002/qj.907, URL 10.1002/qj.907.

Penner, J. E., C. Zhou, A. Garnier, and D. L. Mitchell, 2018: Anthropogenic Aerosol Indirect Effects in Cirrus Clouds. *Journal of Geophysical Research: Atmospheres*, **123 (20)**, 11,652–11,677, https://doi.org/10.1029/2018JD029204, URL 10.1029/2018JD029204.

Pincus, R., P. M. Forster, and B. Stevens, 2016: The Radiative Forcing Model Intercomparison Project (RFMIP): experimental protocol for CMIP6. *Geoscientific Model Development*, **9 (9)**, 3447–3460, https://doi.org/10.5194/gmd-9-3447-2016, URL 10.5194/gmd-9-3447-2016.

Pincus, R., and Coauthors, 2020: Benchmark Calculations of Radiative Forcing by Greenhouse Gases. *Journal of Geophysical Research: Atmospheres*, **125 (23)**, e2020JD033 483, https://doi.org/10.1029/2020JD033483, URL 10.1029/2020JD033483.

Quaas, J., and Coauthors, 2009: Aerosol indirect effects general circulation model intercomparison and evaluation with satellite data. *Atmospheric Chemistry and Physics*, **9 (22)**, 8697–8717, https://doi.org/10.5194/acp-9-8697-2009, URL 10.5194/acp-9-8697-2009.

Rieger, V. S., S. Dietmüller, and M. Ponater, 2017: Can feedback analysis be used to uncover the physical origin of climate sensitivity and efficacy differences? *Climate Dynamics*, **49 (7-8)**, 2831–2844, https://doi.org/10.1007/s00382-016-3476-x, URL 10.1007/s00382-016-3476-x.

Rotstayn, L. D., and J. E. Penner, 2001: Indirect Aerosol Forcing, Quasi Forcing, and Climate Response. *Journal of Climate*, **14 (13)**, 2960–2975, https://doi.org/10.1175/1520-0442(2001)014⟨2960:IAFQFA⟩2.0.CO;2, URL 10.1175/1520-0442(2001)014{%}3C2960:IAFQFA{%}3E2.0.CO;2.

Samset, B. H., and G. Myhre, 2015: Climate response to externally mixed black carbon as a function of altitude. *Journal of Geophysical Research: Atmospheres*, **120 (7)**, 2913–2927, https://doi.org/10.1002/2014JD022849, URL 10.1002/2014JD022849.

Samset, B. H., and Coauthors, 2013: Black carbon vertical profiles strongly affect its radiative forcing uncertainty. *Atmospheric Chemistry and Physics*, **13 (5)**, 2423–2434, https://doi.org/10.5194/acp-13-2423-2013, URL 10.5194/acp-13-2423-2013.

Sand, M., T. Iversen, P. Bohlinger, A. Kirkevåg, I. Seierstad, Ø. Seland, and A. Sorteberg, 2015: A Standardized Global Climate Model Study Showing Unique Properties for the Climate Response to Black

Carbon Aerosols. *Journal of Climate*, **28 (6)**, 2512–2526, https://doi.org/10.1175/JCLI-D-14-00050.1, URL 10.1175/JCLI-D-14-00050.1.

Satheesh, S. K., and V. Ramanathan, 2000: Large differences in tropical aerosol forcing at the top of the atmosphere and Earth's surface. *Nature*, **405 (6782)**, 60–63, https://doi.org/10.1038/35011039, URL 10.1038/35011039.

Sellar, A. A., and Coauthors, 2019: UKESM1: Description and Evaluation of the U.K. Earth System Model. *Journal of Advances in Modeling Earth Systems*, **11 (12)**, 4513–4558, https://doi.org/10.1029/2019MS001739, URL 10.1029/2019MS001739.

Shell, K. M., J. T. Kiehl, and C. A. Shields, 2008: Using the Radiative Kernel Technique to Calculate Climate Feedbacks in NCAR's Community Atmospheric Model. *Journal of Climate*, **21 (10)**, 2269–2282, https://doi.org/10.1175/2007JCLI2044.1, URL 10.1175/2007JCLI2044.1.

Sherwood, S. C., S. Bony, O. Boucher, C. Bretherton, P. M. Forster, J. M. Gregory, and B. Stevens, 2015: Adjustments in the Forcing-Feedback Framework for Understanding Climate Change. *Bulletin of the American Meteorological Society*, **96 (2)**, 217–228, https://doi.org/10.1175/BAMS-D-13-00167.1, URL 10.1175/BAMS-D-13-00167.1.

Shine, K. P., J. Cook, E. J. Highwood, and M. M. Joshi, 2003: An alternative to radiative forcing for estimating the relative importance of climate change mechanisms. *Geophysical Research Letters*, **30 (20)**, 2047, https://doi.org/10.1029/2003GL018141, URL 10.1029/2003GL018141.

Smith, C., and Coauthors, 2020a: Effective radiative forcing and adjustments in CMIP6 models. *Atmospheric Chemistry and Physics*, **20**, 9591–9618, https://doi.org/10.5194/acp-20-9591-2020, URL 10.5194/acp-20-9591-2020.

Smith, C. J., R. J. Kramer, and A. Sima, 2020b: The HadGEM3-GA7.1 radiative kernel: the importance of a well-resolved stratosphere. *Earth System Science Data*, **12 (3)**, 2157–2168, https://doi.org/10.5194/essd-12-2157-2020, URL 10.5194/essd-12-2157-2020.

Smith, C. J., and Coauthors, 2018: Understanding Rapid Adjustments to Diverse Forcing Agents. *Geophysical Research Letters*, **45 (21)**, https://doi.org/10.1029/2018GL079826, URL 10.1029/2018GL079826.

Soden, B. J., I. M. Held, R. Colman, K. M. Shell, J. T. Kiehl, and C. A. Shields, 2008: Quantifying Climate Feedbacks Using Radiative Kernels. *Journal of Climate*, **21 (14)**, 3504–3520, https://doi.org/10.1175/2007JCLI2110.1, URL 10.1175/2007JCLI2110.1.

Stjern, C. W., and Coauthors, 2017: Rapid Adjustments Cause Weak Surface Temperature Response to Increased Black Carbon Concentrations. *Journal of Geophysical Research: Atmospheres*, **122 (21)**, 11,462–11,481, https://doi.org/10.1002/2017JD027326, URL 10.1002/2017JD027326.

Stjern, C. W., and Coauthors, 2023: The Time Scales of Climate Responses to Carbon Dioxide and Aerosols. *Journal of Climate*, **36 (11)**, 3537–3551, https://doi.org/10.1175/JCLI-D-22-0513.1, URL 10.1175/JCLI-D-22-0513.1.

Sun, J., K. Zhang, H. Wan, P. Ma, Q. Tang, and S. Zhang, 2019: Impact of Nudging Strategy on the Climate Representativeness and Hindcast Skill of Constrained EAMv1 Simulations. *Journal of Advances in Modeling Earth Systems*, **11 (12)**, 3911–3933, https://doi.org/10.1029/2019MS001831, URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2019MS001831.

Szopa, S., and Coauthors, 2021: Short-lived Climate Forcers. *Climate Change 2021 The Physical Science Basis*, Cambridge University Press, 817–922, https://doi.org/10.1017/9781009157896.008, URL 10.1017/9781009157896.008.

Taylor, K. E., M. Crucifix, P. Braconnot, C. D. Hewitt, C. Doutriaux, A. J. Broccoli, J. F. B. Mitchell, and M. J. Webb, 2007: Estimating Shortwave Radiative Forcing and Response in Climate Models. *Journal of Climate*, **20 (11)**, 2530–2543, https://doi.org/10.1175/JCLI4143.1, URL 10.1175/JCLI4143.1.

Telford, P. J., P. Braesicke, O. Morgenstern, and J. A. Pyle, 2008: Technical Note: Description and assessment of a nudged version of the new dynamics Unified Model. *Atmospheric Chemistry and Physics*, **8 (6)**, 1701–1712, https://doi.org/10.5194/acp-8-1701-2008, URL 10.5194/acp-8-1701-2008.

Thornhill, G. D., and Coauthors, 2021: Effective radiative forcing from emissions of reactive gases and aerosols a multi-model comparison. *Atmospheric Chemistry and Physics*, **21 (2)**, 853–874, https://doi.org/10.5194/acp-21-853-2021, URL 10.5194/acp-21-853-2021.

Twomey, S., 1977: The Influence of Pollution on the Shortwave Albedo of Clouds. *Journal of the Atmospheric Sciences*, **34 (7)**, 1149–1152, https://doi.org/10.1175/1520-0469(1977)034⟨1149:TIOPOT⟩2.0.CO;2, URL 10.1175/1520-0469(1977)034{%}3C1149:TIOPOT{%}3E2.0.CO;2.

van Aalst, M. K., and Coauthors, 2004: Trace gas transport in the 1999/2000 Arctic winter: comparison of nudged GCM runs with observations. *Atmospheric Chemistry and Physics*, **4 (1)**, 81–93, https://doi.org/10.5194/acp-4-81-2004, URL 10.5194/acp-4-81-2004.

van Garderen, L., F. Feser, and T. G. Shepherd, 2021: A methodology for attributing the role of climate change in extreme events: a global spectrally nudged storyline. *Natural Hazards and Earth System Sciences*, **21 (1)**, 171–186, https://doi.org/10.5194/nhess-21-171-2021, URL 10.5194/nhess-21-171-2021.

van Marle, M. J. E., and Coauthors, 2017: Historic global biomass burning emissions for CMIP6 (BB4CMIP) based on merging satellite observations with proxies and fire models (1750-2015). *Geoscientific Model Development*, **10 (9)**, 3329–3357, https://doi.org/10.5194/gmd-10-3329-2017, URL 10.5194/gmd-10-3329-2017.

Walters, D., and Coauthors, 2019: The Met Office Unified Model Global Atmosphere 7.0/7.1 and JULES Global Land 7.0 configurations. *Geoscientific Model Development*, **12 (5)**, 1909–1963, https://doi.org/10.5194/gmd-12-1909-2019, URL 10.5194/gmd-12-1909-2019.

Wetherald, R. T., and S. Manabe, 1988: Cloud Feedback Processes in a General Circulation Model. *Journal of the Atmospheric Sciences*, **45 (8)**, 1397–1416, https://doi.org/10.1175/1520-0469(1988)045⟨1397:CFPIAG⟩2.0.CO;2, URL 10.1175/1520-0469(1988)045{%}3C1397:CFPIAG{%}3E2.0.CO;2.

Wilcox, L. J., N. Dunstone, A. Lewinschal, M. Bollasina, A. M. L. Ekman, and E. J. Highwood, 2019: Mechanisms for a remote response to Asian anthropogenic aerosol in boreal winter. *Atmospheric Chemistry and Physics*, **19 (14)**, 9081–9095, https://doi.org/10.5194/acp-19-9081-2019, URL 10.5194/acp-19-9081-2019.

Wilcox, L. J., and Coauthors, 2023: The Regional Aerosol Model Intercomparison Project (RAMIP). *Geoscientific Model Development*, **16 (15)**, 4451–4479, https://doi.org/10.5194/gmd-16-4451-2023, URL 10.5194/gmd-16-4451-2023.

Williams, A. I. L., P. Stier, G. Dagan, and D. Watson-Parris, 2022: Strong control of effective radiative forcing by the spatial pattern of absorbing aerosol. *Nature Climate Change*, **12 (8)**, 735–742, https://doi.org/10.1038/s41558-022-01415-4, URL 10.1038/s41558-022-01415-4.

Wu, X., J. J. Bates, and S. Singh khalsa, 1993: A Climatology of the Water Vapor Band Brightness Temperatures from NOAA Operational Satellites. *Journal of Climate*, **6 (7)**, 1282–1300, https://doi.org/10.1175/1520-0442(1993)006⟨1282:ACOTWV⟩2.0.CO;2, URL 10.1175/1520-0442(1993)006{%}3C1282:ACOTWV{%}3E2.0.CO;2.

Zelinka, M. D., T. Andrews, P. M. Forster, and K. E. Taylor, 2014: Quantifying components of aerosol-cloud-radiation interactions in climate models. *Journal of Geophysical Research: Atmospheres*, **119 (12)**, 7599–7615, https://doi.org/10.1002/2014JD021710, URL 10.1002/2014JD021710.

Zelinka, M. D., S. A. Klein, and D. L. Hartmann, 2012: Computing and Partitioning Cloud Feedbacks Using Cloud Property Histograms. Part I: Cloud Radiative Kernels. *Journal of Climate*, **25 (11)**, 3715–3735, https://doi.org/10.1175/JCLI-D-11-00248.1, URL 10.1175/JCLI-D-11-00248.1.

Zelinka, M. D., C. J. Smith, Y. Qin, and K. E. Taylor, 2023: Comparison of methods to estimate aerosol effective radiative forcings in climate models. *Atmospheric Chemistry and Physics*, **23 (15)**, 8879–8898, https://doi.org/10.5194/acp-23-8879-2023, URL 10.5194/acp-23-8879-2023.

Zhang, K., and Coauthors, 2014: Technical Note: On the use of nudging for aerosol-climate model intercomparison studies. *Atmospheric Chemistry and Physics*, **14 (16)**, 8631–8645, https://doi.org/10.5194/acp-14-8631-2014.

Zhang, S., K. Zhang, H. Wan, and J. Sun, 2022: Further improvement and evaluation of nudging in the E3SM Atmosphere Model version 1 (EAMv1): simulations of the mean climate, weather events, and anthropogenic aerosol effects. *Geoscientific Model Development*, **15 (17)**, 6787–6816, https://doi.org/10.5194/gmd-15-6787-2022, URL https://gmd.copernicus.org/articles/15/6787/2022/.

Zhu, J., and J. E. Penner, 2020: Indirect Effects of Secondary Organic Aerosol on Cirrus Clouds. *Journal of Geophysical Research: Atmospheres*, **125 (7)**, https://doi.org/10.1029/2019JD032233, URL https://onlinelibrary.wiley.com/doi/abs/10.1029/2019JD032233.