Fundamental Aerosol-Cloud Interactions and their Influence on the Aerosol Indirect Effect on Climate

by

Daniel Alexander Rothenberg

B.S. Atmospheric Science, Cornell University (2011)

Submitted to the Department of Earth, Atmospheric, and Planetary Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Atmospheric Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2017

Author .....
Department of Earth, Atmospheric, and Planetary Science
November 14, 2016

Certified by .....
Chien Wang
Senior Research Scientist
Thesis Supervisor

Accepted by .....
Robert D. van der Hilst
Schlumberger Professor of Earth and Planetary Sciences
Head, Department of Earth, Atmospheric, and Planetary Sciences

Fundamental Aerosol-Cloud Interactions and their Influence on the Aerosol Indirect Effect on Climate

by

Daniel Alexander Rothenberg

Submitted to the Department of Earth, Atmospheric, and Planetary Science
on November 14, 2016, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Atmospheric Science

Abstract

The influence of anthropogenic aerosol emissions on the optical properties of clouds and the radiative forcing arising from these interactions, known as the aerosol indirect effect on climate, constitutes a fundamental uncertainty in our understanding of 20^{th} century climate change. In this dissertation, we investigate the role of a keystone physical process, droplet activation, in contributing to this uncertainty. The first half of the ensuing work focuses on the parameterization of this process in global model, assessing both existing schemes and developing a novel one. The second half then quantifies the influence of activation by using a suite of aerosol-climate models which include a complete description of the physics which give rise to the indirect effect.

Parameterizations of droplet activation perform well for idealized single-mode aerosol populations, but show systematic biases in high-pollution, weak-updraft regimes. These are exacerbated when the aerosol in question is a complex mixture. We show that estimates of droplet nucleation are highly sensitive to changes in the accumulation mode size and number concentration; this mode is itself sensitive to anthropogenic aerosol emissions, which potentially further biases modeled cloud droplet number. Using a model emulation technique, we develop a framework for building efficient metamodels of activation, which greatly reduce the mean error in droplet number predicted across regimes.

The biases in these parameterizations raise questions the influence of activation on the indirect effect. Using different schemes, we calculate a spread of $1\,\mathrm{W\,m^{-2}}$ in the indirect effect, which we show is equal to the spread computed from an independent suite of global models with different aerosol and physics modules. The estimated indirect effect scales more strongly with the baseline cloud droplet number concentration simulated by each model than by its change from pre-industrial to present day, indicating a strong saturation effect. While present-day estimates of aerosol-cloud interactions derived from satellite-based instruments are inadequate at constraining the pre-industrial cloud droplet burden, we show that process-based measurements could overcome this problem.

Thesis Supervisor: Chien Wang
Title: Senior Research Scientist

Acknowledgments

I can't imagine having spent my graduate career anywhere other than MIT and Cambridge.

First and foremost, I want to thank Chien Wang for recruiting me and offering me the opportunity to study with him. I've learned so much about science, research, and big ideas from our office meetings over in E19 or across twelve time zones on Skype than I could possibly record here. I can't thank you enough for supporting my many different academic and professional interests, even when they took time away from challenging the scientific status quo! Many other mentors helped me reach this day and to them I also owe my gratitude, especially Ron Prinn who guided me through my first two years at MIT as my academic advisor; Dan Czizco, who graciously included me in his group's experimental projects and offered critical perspectives on my work; and my other committee members, Paul O'Gorman and Steve Ghan, who provided valuable comments and feedback over the past few years.

I'm blessed to have been surrounded by many brilliant and inspiring individuals during my graduate tenure, including Sarvesh Garimella, Alex Avramov, Rotem Bar-Or, Karin Ardon-Dryer, Colin Pike-Thackray, Patrick Blonigan, Alex Andrews, Dan Gilford, Evan Howard, Vince Agard, Andy Miller, Alec Bogdanoff, and so many others. I had the privilege to work closely with Sarah Rosengard, Kelsey Mulder, Rosimar Rios Berrios, Kristy Carter, Joshua Alland, and Aryeh Drager to plan some amazing scientific conferences. Any time I ran into trouble navigating PAOC or EAPS, Melissa Fox and Christine Maglio were there to bail me out. From conference calls to coffee breaks, your all's support and friendship means the world to me.

I always had a refuge when I needed it in the MIT Music and Theater Arts community, and for that I thank in particular my amazing MITSO director, Adam Boyles, and my many other teachers including Marcus Thompson, Aleksandra Labinska and Jin Kyung-Joen. Music is so much more fun and rewarding when you get to make it with great friends and talents like Nathan Gibson, Dan Amrhein, Lawrence Cheuk, and Eva Golos, on top of the great communities of the MIT Symphony Orchestra and

the Harvard Dudley House Orchestra.

But most importantly, I wouldn't be here today if weren't for the support and love of my parents, Art and Carol, my brother, Matt, and my girlfriend, Shene. Thank you all so much for standing by me, even when I had my blinders on and was stuck deep in the quagmire of my work. When I doubted myself, your encouragement kept me going, and I owe all my success in this endeavor and in life to you.

My graduate career was made by possible by a generous outpouring of financial support from many organizations. A National Science Foundation Graduate Research Fellowship (1122374) provided freedom and flexibility to develop my independent research project, and further NSF support (AGS-1339264) greatly improved it. Hundreds of thousands of computing hours on the Yellowstone supercomputing facility were made possible with the support of allocations P35171155 and UMIT0020 from the NCAR Computational and Information Systems Laboratory which is also funded by the NSF. The PAOC Houghton Fund provided for valuable meetings with colleagues and peers across the country, and the American Meteorological Society graciously supported my travel while I served as a co-chair of the Student Conference Planning Committee. Support from the DoE Office of Science and the National University of Singapore/Singapore-MIT Alliance in Research and Technology literally allowed me to travel the world to work on this project and collaborate on others. Finally, the sponsors of the MIT Joint Program on the Science and Policy of Global Change enabled a world-class inter-disciplinary environment to learn about climate science and policy, and I am very thankful for their continued support of mine and the entire program's research portfolio.

To the coffee machine in the Charney Library: we've spent so many wonderful hours (like 9AM, 2PM, and 11PM... really, all the hours) together over the past five years, but it's time to say goodbye. Thanks for the memories and the calories!

Contents

List of Figures

List of Tables

Chapter 1

Overview

Within the atmospheric sciences, a wide of variety of tools and idealized models have been developed to aid us in tackling fundamental questions about the climate system: How do we account for the geological record of frequent, large swings in climate? What internal and external factors set the climate state? How sensitive is the climate to perturbations in those factors? What is the relationship between meteorology and climate across scales? At the apex of our available toolchain are elaborate, coupled models of the atmosphere, ocean, land, and ecological systems which attempt to encompass the complete body of human knowledge about the climate system.

But, as Twomey (1991) pointed out, eschewing these tools for the simplest "climate equation" can be surprisingly insightful:

$$\sigma T_e^4 = (1 - \alpha) S_0 \tag{1.1}$$

Stating a fundamental energy balance, this equation asserts that the amount of energy the Earth radiates to space (LHS) must balance the amount of energy it intercepts from the sun (RHS). But not all solar energy (S_0) is up for grabs; there is some factor α (the planetary albedo) which the Earth never even realizes as it is immediately reflected back to space, either by clouds or the surface of the Earth itself. Interpreting T_e is slightly more difficult, but can be taken as a proxy for the emission, transmission, absorption, and re-emission of energy absorbed at the Earth's surface,

and depends on that surface's temperature along with the absorptivity and structure of the atmosphere above it. In this way, it is similar to a planetary "temperature".

Depending on one's scientific proclivities, it is possible to torture this climate equation in many different ways. One of the most basic would be to observe that there are two ways methods to perturb T_e: either turn up or down the output of the Sun (S_0), or change how much solar energy is reflected back to space (α). If we hold the Sun's output constant, then changes in the albedo can be quite potent. For instance, Budyko (1969) and Sellers (1969) offered concise models of the interplay between glacial extent (which strongly influences α as glaciers are very reflective) and the global temperature. But changes in the properties of clouds could accomplish the same climactic feat, and clouds can respond to perturbations on far shorter time scales than glaciers.

How specifically, then, can the properties of cloud influence albedo? On the one hand, since α represents a globally-integrated effect, one could simply increase the average fractional cloud cover. Doing so would incur enhanced cloud interactions with longwave radiation, producing a feedback on T_e whose sign depends on the altitude at which a given cloud forms (such that increases in the global coverage of low-level clouds could produce a cooling effect on the climate, whereas increases in high altitude clouds could produce a warming effect). Alternatively, one could modify cloud reflectance, say by increasing the droplet number concentration N_d in the cloud (Twomey, 1977). This would not necessarily create a counteracting effect in the longwave; if anything, one might expect a positive feedback, since increases in N_d are often associated longer cloud lifetimes (Albrecht, 1989) and therefore larger α, in an integral sense. This observation is foundational to a corner of the academic literature on climate change aimed at developing novel engineering schemes to mitigate global warming by large-scale modification of cloud albedo (Latham, 1990; Latham et al., 2012; Kravitz et al., 2013; Crook et al., 2016) Furthermore, changes in N_d necessarily invoke changes in the microphysical structure of a cloud (the spatiotemporally varying droplet or hydrometeor size distribution) which play additional, secondary roles in modifying its albedo and dynamical properties (such as precipitation), although the

complexity of the processes involved in this may buffer the system and mute its response under many circumstances (Stevens and Feingold, 2009).

The question of how to perturb N_d is inextricably linked with the atmospheric aerosol: ultimately, most droplets in the atmosphere form via condensation on some particle. As a consequence, the simplicity of Equation (1.1) obfuscates a far more complicated, underlying, coupled aerosol-cloud-climate system. The sheer complexity of the processes involved in the dynamics of each system would respectively take up a complete textbook (to name one each, Seinfeld and Pandis (2006), Pruppacher and Klett (1997), and Pierrehumbert (2010)). The potential influence of perturbations to the aerosol-cloud-climate system's components is encoded in α.

Unlike Equation (1.1), the aerosol-cloud-climate system is not readily amenable to a simple, analytical treatment. But we can still make an effort to understand its behavior with a reductionist approach. To borrow from the ecology literature, we might describe a fundamental or "'keystone" component in the aerosol-cloud-climate system which, when taken away or modified, can grossly affect the character of the system as a whole (Paine, 1966). In the aerosol-cloud-climate system, that keystone is the process of droplet nucleation, or aerosol "activation". Activation acts as a machine which, for a given meteorological context, translates the complexities of the ambient aerosol size distribution into an assemblage of cloud droplets, our N_d, with a narrow size distribution. Many processes then act to further modify N_d and its underlying size distribution. For instance, mixing of dry air, collision and coalescence, and precipitation are some ways that influence how N_d evolves over time in a cloud. Clouds themselves can be coupled in a dynamical cycle, something which has also been described using frameworks co-opted from the ecology community, such as the Lotka-Volterra or predator-prey model (Koren and Feingold, 2011). That cycle can doubly intersect with aerosol via cloud processing, where evaporating cloud droplets release the aerosol material they have collected, in a chemically-altered form. But prior to any such cycle, activation dictates a basic relationship between the aerosol and N_d.

This fundamental aerosol-cloud linkage has long been recognized, and served to

motivate the development of prognostic cloud microphysics schemes in global climate models—the apex tools which try to resolve detail hidden by Equation (1.1). Using data from aircraft campaigns, works like Jones et al. (1994) employed simple empirical relationships to relate aerosol properties like sulfate mass to N_d, enabling a first, crude online estimate of the indirect effect (the influence of aerosols on climate via modifying cloud properties). Further developments increased the complexity of each component in the aerosol-cloud-climate system. Ghan et al. (1997) added a prognostic scheme for N_d into the mix, including an activation parameterization very closely related to one widely used in present-day models (Abdul-Razzak et al., 1998), but with prescribed aerosol number. Lohmann and Feichter (1997) included an interactive sulfate aerosol chemistry scheme in a different model to enable prognostic aerosol number which was fed to a different activation scheme (Lohmann et al., 1999). Over the ensuing two decades, these models (and other spiritual descendants) have continued to evolve, including ever-increasingly complex descriptions of the ambient aerosol and cloud hydrometeor size distributions and the processes which act on them. Underpinning this ballooning complexity has been the hope that any given innovation could provide insight or the opportunity for new observational constraints or bounds on the behavior of the aerosol-cloud-climate system.

Yet recently, despite its instigation in the development of global modeling tools, activation has played a back-seat role to other physical processes in this pursuit. We see this is in at least two ways. First, there has been a proliferation of physically-based aerosol activation schemes which work with arbitrary modal or sectional descriptions of the aerosol size distributions (Ghan et al., 2011). Modifying the aerosol scheme in a given climate model no longer requires an investment in the fundamental development of a new, suitable activation scheme. As a result, there is less incentive to rigorously study the nuances of how changes in that aerosol description influence droplet number, either using an activation scheme or a detailed process model. A related consequence is a disincentive to validate activation of said aerosol populations against real-world observations. Available observations cover but an increasingly small set of the aerosol parameter space in these models, and thus may only provide weak constraints on the

aerosol-cloud linkage in them.

Second, there are many microphysical processes in global models which may be far more poorly constrained or understood than activation, and these are attractive targets for sensitivity analysis and deriving observational constraints. For instance, autoconversion of cloud-water to rain-water involves bulk parameterizations tuned to large-eddy simulation output (Khairoutdinov and Kogan, 2000) and which require additional scaling or tuning factors for deployment in a global model. Changes in such a tuning systematically modify cloud macro-physical properties but still take N_d as an input, and are thus an important link in the aerosol-cloud-climate chain. Another example is the modeled representation of sub-grid phenomena. Global models necessarily are run at far coarser spatial and temporal resolution than those on which most cloud processes. Resolving relevant cloud-scale meteorology parameters in global simulations is thus a very difficult problem. Within the aerosol-cloud framework, one such parameter is the distribution of local vertical velocities driving the plumes which generate clouds. Potentially-convergent lines of work seek to develop either higher-order moist turbulence closures (Bogenschutz et al., 2013) which can be used to drive probabilistic approaches to sub-grid scale cloud physics, or to identify characteristic values from the resulting distributions of vertical velocity (Morales and Nenes, 2010) which can be used to drive traditional, deterministic parameterizations. Propagation of uncertainty in these types of parameter(ization)s through to the climate scale under the constraints of available theory and observation is a means to disentangle the web of processes in the aerosol-cloud-system into some semblance of order or importance.

Uncertainty in the aerosol and cloud components of the aerosol-cloud-climate system has been systematically and rigorously quantified. But how does the linkage between the two—activation—contribute to uncertainty in the climate system? The work presented in this dissertation targets this basic question, seeking to retain Equation (1.1) in its broader perspective through the assessment of aerosol indirect effects on climate (AIE) and aerosol-cloud interactions (ACI).

In pursuit of answers to this question, this dissertation is divided into two halves.

In its first half (Chapters 2 and 3), we develop a novel method for parameterizing aerosol-cloud interactions in global models (droplet activation). Unlike recent developments in the field, this parameterization is built from the statistical emulation of a detailed process model. Such a parent model would be amenable to the inclusion of many physical and chemical details, such as entrainment or the influence of atmospheric chemistry on condensation and activation. The statistical emulation process yields a form of parameterization akin to a look-up table, but extendable to a very large number of inputs without sacrificing accuracy, speed, or computational complexity. With this tool, we systematically assess the influence of aerosol physical and chemical properties on aerosol activation and cloud droplet number concentration, N_d, first for an idealized lognormal aerosol size distribution and then for a highly-complex, multi-modal aerosol formulation. As part of the process used to develop the emulator, we employ a greedy, iterative analysis to assess the constraints exerted by different modes on activation. We find that the number concentration of accumulation mode sulfate particles strongly predicts N_d in the vast majority of cases, including in the presence of giant CCN such as dust or sea salt particles. Even the presence of a small number of these particles is enough to offset large variations in the number of smaller, Aitken mode particles.

We then use our emulators and reference process model to evaluate two widely-used activation schemes. Compared to the reference model, the emulators we build in this work produce much smaller mean errors in N_d for aerosol populations sampled across the modeled globe. However, there are strong contrasts in the performance of all the parameterizations and emulators between oceanic and continental aerosol regimes. In particular, all of our schemes systematically over-predict N_d in polluted regimes over continents, and under-predict it clean regimes over oceans. Comparing the sampling distribution of aerosol populations from our global model versus a much larger parameter space suggests that these systematic biases are related to the underlying, modeled aerosol size distribution; we tend to produce sample aerosol populations which the schemes do not necessarily handle successfully. As a result, using different activation schemes can lead to large differences in the model-simulated

climatology of N_d. We hypothesize that these differences could dramatically influence estimates of the indirect effect.

In the latter half of the dissertation (Chapter 4 and 5) we test this hypothesis, aided by a large ensemble of modern aerosol-climate models. Using a modern aerosol-climate model of our own implementation, we couple each of the previously-studied activation schemes, along with a set of heuristic variants, and estimate the magnitude of the indirect effect. The spread among estimates using models with each activation scheme considered is $1.1\,\mathrm{W\,m^{-2}}$, which is as large as both the spread in our reference ensemble and that reported by the IPCC (Boucher et al., 2013). The spread is also larger than that reported by other works which perturb different components of the emissions, aerosol, and cloud components of similar global models. In agreement with previous work, we find that the magnitude of the indirect effect in our simulations scales poorly with the change in N_d between the pre-industrial and present day, and instead strongly with the base state climatology of N_d in the pre-industrial case under natural emissions. Because of this, differences in the relationship between N_d and aerosol encoded in each activation scheme, particularly in the natural or clean aerosol regimes, contributes significantly to uncertainty in the indirect effect. This is somewhat distressing, because unlike empirical parameterizations between N_d and aerosol or microphysical processes, do not have free parameters which could be tuned to correct biases compared to observations.

In an effort to resolve this source of distress, we perform analyses of aerosol-cloud interactions (ACI) and the impact of the choice of activation scheme on their magnitude and relationship to the indirect effect. We note that AOD-based metrics are problematic for this purpose, as the influence of anthropogenic aerosol emissions into our global model produces a step change in the distribution of regional AOD. Thus, measures of ACI derived using the present day climate are a poor proxy for the pre-industrial one. However, using CCN obviates this problem, as there is a smooth transition in the distribution of CCN which overlaps between both climate states. This motivates an assimilation of new data sources to quantify CCN-N_d relationships for constraining models and extending observational constraints on the indirect effect.

In addition to examining this fundamental source of uncertainty, we show that ACI differs greatly between the models using different activation schemes and strongly correlates with regional diversity in the strength of the indirect effect. This further motivates the need for *in situ* measurements of CCN and cloud microstructure as a means to providing constraints on the physics simulated in global models.

Finally, in Chapter 6, we reflect on the lessons learned in this dissertation, and contemplate the future of assessments of the aerosol indirect effect and aerosol-cloud interactions. We focus these reflections on three broad topics. First, we consider how field measurements and observation campaigns could provide data that would help constrain the indirect effect by concentrating on fundamental microphysical processes. Second, we consider the possibility of new, emergent constraints on aerosol-cloud interactions arising from our investigations here. Finally, we propose a renewed investment in a holistic analysis of the climate response to aerosol perturbation in lieu of continued efforts to constrain the climate forcing to those same factors. Such an effort could directly contribute to the investigation into emergent constraints on the indirect effect, and would be more relevant to a society faced with an imminent, rapid decline in aerosol emissions and associated climate penalties.

Chapter 2

Metamodeling of Droplet Activation for Global Climate Models

2.1 Introduction

Interactions between aerosol and clouds yield one of the largest sources of uncertainty in understanding climate and future climate change on regional and global scales (Boucher et al., 2013). Within the Earth's atmosphere, homogeneous liquid water droplet formation is not thermodynamically favorable (Pruppacher and Klett, 1997); instead, the pathway to nucleating cloud droplets is aided by the presence of ambient aerosols, a subset of which possess physical and chemical characteristics which allow them to serve as cloud condensation nuclei (CCN). These CCN provide a linkage between the physiochemical processes of atmospheric particles and cloud microphysics.

Changes in the background aerosol population can directly affect the properties of a nascent cloud droplet population. For instance, holding liquid water content constant, an increase in the number of CCN would tend to increase the total cloud droplet number con-

This chapter is adapted from Rothenberg, D., and C. Wang, 2015. Metamodeling of Droplet Activation for Global Climate Models. *Journal of the Atmospheric Sciences.* doi:10.1175/JAS-D-15-0223.1.

centration (the "Twomey" effect) while necessarily reducing the average size of the droplets (Twomey, 1974). Such a change could enhance a cloud's albedo, an effect which could be further amplified through microphysical feedbacks since smaller droplets impede the production of drizzle and thus lengthen cloud lifetime (Albrecht, 1989). Mechanisms whereby aerosol influence the properties of clouds (and ultimately climate) are generally known as "aerosol indirect effects" (Haywood and Boucher, 2000; Lohmann and Feichter, 2004) and provide a path for changes in the ambient aerosol to produce cascading effects up to progressively larger scales of atmospheric motion (e.g. Wang, 2005a; Ekman et al., 2011; Morrison et al., 2011; Tao et al., 2011; Fan et al., 2012; Altaratz et al., 2014).

Aerosol indirect effects can either warm or cool the climate, but they all fundamentally depend on a subset of the ambient aerosols which function as CCN. The theory describing the dependency of cloud droplet nucleation (also known as aerosol activation) on CCN availability and ambient aerosol has been rigorously developed using adiabatic and entraining parcel theory (Seinfeld and Pandis, 2006; Pruppacher and Klett, 1997), and depends on details of the heterogeneous chemical composition, number, size distribution(s) and mixing state of the background aerosol (Mcfiggans et al., 2006) as well as local meteorology (Morales and Nenes, 2010). Under polluted conditions, effects relating to chemical composition could produce a climatic effect as large as the basic "Twomey" effect (Nenes et al., 2002; Lance et al., 2004).

The development of activation parameterizations was pioneered by Twomey (1959) and Squires and Twomey (1961), who derived a relationship between the number of activated particles and the environmental supersaturation based on an aerosol size distribution approximated by a power law. Ghan et al. (2011) presented a thorough overview of subsequent developments over the past five decades and an inter-comparison of several modern parameterizations. However, there is still an active effort to improve these parameterizations, as they are increasingly called upon to mediate between ever more complex aerosol models and the climate models to which they are coupled. For instance, the parameterization initially developed by Nenes and Seinfeld (2003) has seen continuous development, including modifications to handle condensation onto insoluble but wettable particles using adsorption-activation theory (Kumar et al., 2009), environmental entrainment (Barahona and Nenes, 2007), and numerical improvement of the population-splitting technique (Barahona et al., 2010; Morales Betancourt and Nenes, 2014a). Similarly, Ghan et al. (2011)

extended the parameterization of Abdul-Razzak and Ghan (2000) to account for non-unity values of the accommodation coefficient a_c. Beyond idealized testing and droplet closure studies (Meskhidze, 2005; Fountoukis et al., 2007), these modern parameterizations have been implemented in coupled climate-aerosol models such as the Community Earth System Model (CESM) to predict online cloud droplet number concentrations, where they have been shown to correct biases in global average cloud droplet number concentrations and improve agreement with cloud properties measured from satellite-born instruments (Gantt et al., 2014). Furthermore, adjoints of these parameterizations have been derived and coupled to chemical transport and global models in order to study the sensitivities of cloud droplet number to aerosol, chemical and microphysical factors (Karydis et al., 2012; Moore et al., 2013).

Additionally, following the original integral/geometric approach by Twomey (1959), analytical representations of supersaturation evolution from adiabatic parcel theory have been progressively generalized to relate aerosol distributions to activation kinetics (Cohard et al., 1998; Khvorostyanov and Curry, 2006, 2008; Shipway and Abel, 2010; Shipway, 2015). Although fundamentally analytical parameterizations, schemes of this class typically must rely on expensive numerical operations, such as in the evaluation of hypergeometric functions and iterative loops.

While most of these recent efforts towards improving activation parameterizations have focused on building highly generalized, "physically-based" tools, there is still an application for other parameterization approaches. Saleeby and Cotton (2004) parameterized droplet nucleation for a cloud-resolving model, the Regional Atmospheric Modeling System (RAMS), by constructing a four-dimensional lookup table based on temperature, vertical velocity, aerosol number concentration, and the median radius of a lognormal aerosol mode with assumed chemical composition. Ward et al. (2010) added a fifth dimension representing chemical composition via aerosol hygroscopicity (following κ-Köhler theory (Petters and Kreidenweis, 2007) to the lookup table, and later generalized to aerosol soluble fraction by Saleeby and van den Heever (2013)). Constructing lookup tables of detailed parcel model results can be considered a form of model emulation combining a cache of known results and local polynomial (linear) approximation.

As the degrees of freedom and number of parameters describing a given aerosol population in a model increase, the burden of saving enough known points to interpolate through

the parameter space via lookup table to some reasonable accuracy increases algebraically. For instance, the CESM features a modal aerosol population with three predefined, internally mixed lognormal modes, each with a fixed geometric standard deviation (Liu et al., 2012). Each mode is uniquely described by two moments (total number and total mass concentration) and the chemical composition of the mode by a single prognostic hygroscopicity term. Thus, the entire aerosol population has $N = 9$ degrees of freedom - too many to build a lookup table of activation statistics. Physically-based parameterizations were designed to accommodate these sorts of arbitrary mixtures of aerosol, but have a tendency to systematically underestimate activated fractions and subsequently cloud droplet number (Simpson et al., 2014). This is due to the parameterizations' use of a set of assumptions which become increasingly likely to be violated as the aerosol population becomes more complex, specifically (1) that proto-droplets grow in equilibrium with environmental changes in relative humidity, and (2) that there are no kinetic or inertial limitations to droplet growth. The presence of giant CCN (Barahona et al., 2010) and weak updrafts or excessively polluted conditions (Nenes et al., 2001) exacerbates this problem.

The goal of this study is to apply a surrogate modeling or emulation technique commonly used in the uncertainty quantification literature to a detailed parcel model capable of describing aerosol activation; this yields an efficient parameterization optimized for the high-dimensional parameter space affecting droplet nucleation in coupled aerosol-climate model. In essence, employing the derived emulator as an activation parameterization would be akin to directly coupling a detailed parcel model to a global model. Such a parameterization would be directly physically-based, but rely on fewer assumptions which affect the condensational growth of aerosol into CCN. However, it would also incorporate the efficiency of a lookup table, since the emulator would be designed to require a scarce amount of cached information and to be computationally cheap to evaluate. Additionally, it would improve upon the framework of lookup table and be extensible to a very high-dimensional parameter space and thus be compatible with aerosol-climate models of increasing complexity.

Section 2.2 describes the parcel model and the probabilistic collocation method (PCM) used to build its emulator. Section 4.3 presents results from applying the PCM to build a parcel model emulator designed to simulate the activation of a single lognormal aerosol mode under a wide variety of background environments and compares the new emulator to existing activation parameterizations. Section 2.4 motivates an extension of the technique

to an emulator suitable of mediating aerosol activation in a coupled aerosol-climate model.

2.2 Methodology

2.2.1 Parcel Model

Adiabatic parcel models are commonly used to study droplet activation and its sensitivity to factors such as environmental conditions and ambient aerosol properties. For this work, a novel parcel model based on previous studies (Leaitch et al., 1986; Nenes et al., 2001; Seinfeld and Pandis, 2006) was designed and implemented to accommodate diverse, chemically heterogeneous, polydisperse aerosol populations. The model simulates droplet growth on the initial aerosol population due to condensation within a constant-speed adiabatic updraft.

Although an arbitrary aerosol size distribution function can be supplied as an input to the model, for the purposes of this study the initial aerosol distribution is assumed to be lognormal and described by the equation

$$n_N(r) = \frac{dN}{d\ln r} = \frac{N_t}{\sqrt{2\pi}\ln\sigma_g}\exp\left(-\frac{\ln^2(r/\mu_g)}{2\ln^2\sigma_g}\right) \tag{2.1}$$

where the parameter set (N_t, μ_g, σ_g) correspond, respectively, to the total aerosol number concentration, the geometric mean radius, and the geometric standard deviation of the distribution. Within the model, this distribution is discretized into 200 size bins equally spaced over the logarithm of particle radius (r) and covers the size range ($\min\left[0.1\,\text{nm},\ \mu_g/10\sigma_g\right], \mu_g \times 10\sigma_g$). The mean radius in each bin grows due to condensation so that the activation of wetted aerosol into droplets is calculated in a Lagrangian sense. To relate size-dependent droplet growth to its embedded aerosol's chemical composition, each bin is prescribed a hygroscopicity following κ-Köhler theory (Petters and Kreidenweis, 2007).

To simulate droplet activation, the parcel model first computes an equilibrium wet-size distribution from the given initial aerosol population and initial environmental temperature, pressure, and relative humidity. Then, a set of conservation equations which describe the evolution of the parcel temperature, supersaturation, liquid/vapor water content, and pressure are integrated forward in time using a solver suitable for stiff systems (VODE; Brown et al., 1989). The complete system of equations and further details on the parcel model can be found in Appendix A of this chapter.

2.2.2 Polynomial Chaos Expansion

We construct an emulator of the parcel model in order to assess droplet activation by applying the probabilistic collocation method (PCM; Tatang et al., 1997). The PCM maps an output from the parcel model to a set of input parameters by building a response surface using a polynomial chaos expansion. The polynomial which results from this process is optimally a computationally-efficient, high-fidelity reproduction of the detailed parcel model simulation. Although often used for conducting global sensitivity analyses (Pan et al., 1997; Calbó et al., 1998; Mayer et al., 2000; Lucas and Prinn, 2005; Anttila and Kerminen, 2007) chaos expansion-based emulators have also been used to build deterministic parameterizations (Cohen and Prinn, 2011). To apply and build the chaos expansions discussed here, the open-source Design Analysis Kit for Optimization and Terascale Applications (DAKOTA; Adams et al., 2014) version 6.1 was used, which automates the sampling of the PCM collocation points and the computation of the coefficients of the polynomial chaos expansion given a user-generated interface to a numerical model (the parcel model described in Section 2.2.1) and a description of the inputs/outputs to and from that interface.

A review of the theoretical basis of polynomial chaos expansion and its potential applications is provided by Sudret (2008); here, we highlight the important details of the technique as applied via PCM for the benefit of the reader. PCM is a non-intrusive polynomial chaos expansion technique; rather than require complex, significant modifications to the model being emulated, PCM instead considers the model to be a black box and constructs a map from an input parameter space to the model output parameter space. To accomplish this, PCM re-casts the input parameters to a model as a set of M independent random variables, $\mathbf{X} = X_1, \ldots, X_M$, each with an associated probability density function. For each input in $\mathbf{X}$, the associated PDF is used as a weighting function to derive an orthogonal polynomial which adds to the bases for the polynomial chaos expansion, ϕ_j. Using a finite number of these bases, the chaos expansion for a given model response, R, is then

$$R \approx \sum_{j=0}^{P} \alpha_j \phi_j(\mathbf{X}) \tag{2.2}$$

The complete basis of polynomials up to a fixed total-order p is retained in the expansions computed here. For such a total-order expansion, equation (2.2) has $N_t = P + 1 = (M + p)!/(M!p!)$ terms as it contains each of the $p+1$ orthogonal basis polynomials for each input

parameter. PCM provides an experimental design for determining the coefficients of the expansion, α_j, by evaluating the model response for a set of N_s total input parameter sets, $\mathbf{X}^1, \ldots, \mathbf{X}^{N_s}$, corresponding to the roots of ϕ_j and solving a regression problem

$$\mathbf{\Phi}\boldsymbol{\alpha} = \mathbf{R} \tag{2.3}$$

where $\mathbf{R}$ is the vector of model responses, $\boldsymbol{\alpha}$ is the vector of expansion coefficiens, and the matrix $\mathbf{\Phi}$ contains rows for each of the polynomial terms ϕ_j evaluated for a given input parameter set $\mathbf{X}^j$.

A practical consideration in applying the PCM to a particular problem is what subset of the N_s potential points to use in solving for the coefficients. In general, there exists a full factorial design of size $N_s = (p+1)^M$ available for use (all the roots of the orthogonal basis polynomials for all inputs). However, for even moderately-sized p and M, the number of potential model evaluations grows very rapidly. In our application we choose a subset of N_s' parameter sets when applying the PCM by using two rules of thumb:

1. Choose parameter sets with roots closest to the origin (Sudret, 2008);
2. Cross-validate the regression result using $3N_t$ parameter sets chosen according to rule (1).

These rules will always produce an over-determined system for Equation (2.3). The accuracy of the resulting emulators derived in this study were not sensitive to increasing N_s', and $3N_t$ does not produce an excessive number of required parcel model simulations. Three different techniques were tested for solving this system: typical linear regression by ordinary least squares (OLS), Least Angle Regression (LARS; Tibshirani et al., 2004), and Least Absolute Shrinkage and Selection Operator (LASSO; Tibshirani, 2011). Both LARS and LASSO involve computing

$$\boldsymbol{\alpha} = \arg\min ||\mathbf{\Phi}\boldsymbol{\alpha} - \mathbf{R}||_{l_2}^2 \quad \text{such that } ||\boldsymbol{\alpha}||_{l_1} \leq \tau \tag{2.4}$$

in an iterative, greedy fashion with the potential to yield sparse solutions with some coefficients $\alpha_j = 0$. This would be desirable for high-order chaos expansions for many input parameters, as it would reduce the number of coefficients necessary to save for re-using the expansion as an emulator. Additionally, this greedy characteristic helps to avoid overfitting

the chaos expansions. Each time a potential term is added to the trial expansion, an error estimate is calculated based on leave-one-out sampling (Blatman and Sudret, 2011); if the error estimate increases, the potential term is rejected. Inspecting the resulting terms gives a metric to compare the OLS-derived chaos expansion. The same error calculation using leave-one-out sampling can be applied using the larger, independent sampling dataset used to evaluate the chaos expansions in Section 2.3.1 for each higher-order chaos expansion to help identify when overfitting is occurring.

2.2.3 Emulation of Parcel Model

The PCM was applied to emulate the activation of a single, lognormal aerosol mode embedded in a constant-speed adiabatic updraft as simulated by the parcel model described in Section 2.2.1. Specifically, the model was used to predict the base-10 logarithm of the maximum supersaturation $S_{\max}$ given aerosol of different lognormal size distributions and hygroscopicities, for different environmental and thermodynamic conditions. The mechanics of the PCM - and polynomial chaos expansion more generally - permit the use of arbitrary PDFs to describe the input parameters over their physically relevant values. In this application, uniform distributions were chosen to emphasize that the derived chaos expansion should perform equally well anywhere within the input parameter space. Each uniform distribution is defined by minimum and maximum permissible bounds for each input parameter, a and b, such that its probability distribution is just given as $f(x) = 1/(b-a)$ for $a < x < b$. The complete list of input parameters and their bounds is summarized in Table 2.1.

In order to utilize the PCM, the uniform distribution for each parameter must be rescaled to the range $[-1, 1]$. This produces a new set of random variables for each parameter X_i:

$$Z_i = \frac{2(X_i - a_i)}{b_i - a_i} - 1 \tag{2.5}$$

The orthogonal polynomials used in the basis of the chaos expansion which correspond to a uniform PDF over the interval $[-1, 1]$ are the canonical Legendre polynomials which follow the three-term recurrence relation:

Table 2.1: Input parameters and bounds used in computing chaos expansion for emulating droplet activation from a single, lognormal aerosol mode embedded in a constant-speed updraft.

Symbol	Name	Units	Bounds
$\log_{10} N$	Mode Number Concentration	$\log_{10}$ cm^{-3}	[1, 4]
$\log_{10} \mu_g$	Mode Geometric Mean Radius	$\log_{10}$ µm	[-3.0, 1.0]
σ_g	Mode Standard Deviation	-	[1.2, 3.0]
κ	Mode Hygroscopicity	-	[0.0, 1.2]
$\log_{10} V$	Updraft Velocity	$\log_{10}$ m s^{-1}	[-2.0, 1.0]
T	Air Temperature	K	[240, 310]
P	Air Pressure	Pa	[50000, 105000]
a_c	Accommodation Coefficient	-	[0.1, 1.0]

$$P_0(Z) = 1 \tag{2.6}$$

$$P_1(Z) = Z \tag{2.7}$$

$$P_{n+1}(Z) = \frac{(2n+1)ZP_n(Z) - nP_{n-1}(Z)}{n+1} \tag{2.8}$$

The roots of these Legendre polynomials can be inverted using Equation (2.5) to determine values in the original, physical parameter space to use in sampling the parcel model.

The bounds for the physical parameters supplied to the PCM were chosen in order to characterize activation near cloud base (Table 2.1). The logarithm of several variables (aerosol number concentration, aerosol geometric mean radius, and updraft velocity) is used because the supersaturation maximum computed by the parcel model is sensitive to changes in these parameters over several orders of magnitude. Updraft velocity is permitted to vary between 0.01 and 10.0 m s^{-1}; over this range (which covers a spectrum from weakly convecting, stratiform clouds to strong, deeply convecting ones) and the range of aerosol number concentration (which includes clean and very polluted regimes), activated fraction can range from virtually nothing to complete activation of the entire aerosol population. The lower bound of updraft speeds considered here is less than the minimal value allowed in many climate models (Golaz et al., 2011). The aerosol mode geometric mean radii μ_g span a variety of smaller Aitken-type modes to large, coarse aerosol modes and potentially giant CCN (Seinfeld and Pandis, 2006). The mode geometric standard deviation is fixed

in some global model aerosol schemes (e.g. Kim et al., 2008; Liu et al., 2012) and the range chosen here cover many potential values. Hygroscopicity values are based on Table 1 of Petters and Kreidenweis (2007) and span values for materials ranging from organic aerosol to highly-hygroscopic salts. The accommodation coefficient was limited to a globally representative range based on observations of CCN activation kinetics from many campaigns (Raatikainen et al., 2013). Temperature and pressure ranges were chosen to reflect typical lower troposphere values. Note that the bounds on the parameters considered here are expanded from those considered by Ghan et al. (2011) in their intercomparison of activation schemes.

Many parameterizations of droplet nucleation diagnose activation directly by applying equilibrium Köhler theory. To do this, the maximum supersaturation achieved by a cloudy parcel is used as a threshold; particles with a Köhler-predicted critical supersaturation lower than this maximum environmental supersaturation are considered to be activated. However, physically, for a droplet to activate it must grow beyond a critical size corresponding to this critical supersaturation. Due to kinetic limitations on droplet growth, this may not be realized for droplets growing on very large CCN (Nenes et al., 2001). Ghan et al. (2011) suggests particles with radius larger than 0.1 µm or those whose critical supersaturations are close to the environmental maximum supersaturation are likely to suffer from this effect. By directly considering a detailed parcel model, the emulators constructed here consider kinetic limitations on droplet growth and their feedback on the evolving parcel supersaturation. In existing, physically-based parameterizations in the literature, an instantaneous growth-rate assumption must be applied. This assumption causes parameterizations to under-predict supersaturation maximum because instantaneous growth will tend to condense water from the vapor phase too quickly and release surplus latent heat which suppresses the increase of the supersaturation.

Because the computed supersaturation maximum in a parcel model activation simulation can also vary over several orders of magnitude, we use $\log_{10} S_{\max}$ as the response function emulated by the PCM. However, in order to apply this transform to the response function, it must be assumed that the cloudy parcel always supersaturates with respect to water vapor, i.e. $S_{\max} > 0$. To ensure this, all simulations performed during sampling by the PCM start with an aerosol population equilibrated to 100% relative humidity and an initial environmental supersaturation of 0. Many existing parameterizations in the literature implicitly

make this same assumption by representing the aerosol population with respect to a coordinate derived from the critical supersaturation for a given size [Ghan et al. (2011); equations (12)-(17)]; in this case the integral over the size distribution spans $0 \leq S \leq S_{\max}$, and thus considers the same situation with respect to the growth of the nascent droplet population as equation (3) of Ghan et al. (2011).

For the 8-parameter input space governing single-mode activation considered here, the 3rd- and 4th-order chaos expansions produced by the PCM have 165 and 495 terms, respectively. The number of terms is equivalent to the number of coefficients one must store in order to re-use a given chaos expansion. This small memory footprint afford chaos expansions a huge advantage over similar parameterizations based on detailed look-up tables. An isotropic look-up table with M parameters and n sample points for each parameter would require n^M values to be stored - a value which for even small numbers of parameters can be several orders of magnitude larger than even a high-order chaos expansion. A more detailed description of how the chaos expansions are saved and later evaluated is given in Appendix 2.6.

The parcel model emulator yields $\log_{10}(S_{\max})$ as a function of an input parameter set drawing from the terms defined in Table 2.1,

$$\log_{10}(S_{\max}) = f(\log_{10} N, \log_{10} \mu_g, \sigma_g, \kappa, \log_{10} V, T, P, a_c) \tag{2.9}$$

From this value of $\log_{10}(S_{\max})$ the number concentration of cloud droplets activated, N_{act}, can be obtained by integrating over the original lognormal aerosol size distribution re-expressed as a function of critical supersaturation rather than droplet radius (Ghan et al., 2011), yielding the expression

$$N_{\text{act}} = \frac{N}{2}\left(1 - \text{erf}\left[2\ln\left(\frac{S_m}{S_{\max}}\right)/(3\sqrt{2}\ln\sigma_g)\right]\right) \tag{2.10}$$

where S_m is the critical superasturation for the geometric mean radius, μ_g.

2.2.4 Global Sensitivity Analysis

We supplement the assessment of our new droplet activation emulator by calculating a set of global sensitivity metrics not previously applied to this problem. The method deployed here

is a variance-based decomposition, which seeks to assign uncertainty in a model response to uncertainty in both individual model input parameters and their interactions with one another. Two different quantities, called Sobol' indices (Sobol', 2001), are produced by this method: main (S_i) and total (T_i) effect indices. Sobol' indices can be used to rank the relative importance of model inputs in influencing its response and for identifying potential inputs which are unimportant and also candidates to be held fixed without grossly biasing the accuracy of a model emulator (Sobol', 2001). Additionally, these indices can identify higher-order terms in expansions like Equation (2.2) (Sobol', 2001).

The main effect index, S_i, indicates what fraction of the uncertainty in a model response, R, is attributable to a single member of the model parameter set, X_i, by comparing the variance of the model response conditioned on X_i against the total variance in R. That is,

$$S_i = \frac{\mathrm{Var}_{X_i}(E[R|X_i])}{\mathrm{Var}(R)} \tag{2.11}$$

This is in contrast with the total effect index, T_i, which instead compares the variance of R conditioned on all the input parameters save for X_i (notated as $X_{\sim i}$). T_i quantifies the variance of R attributable to X_i and the sum of its interaction with other input terms. Similar to the main effect index,

$$T_i = \frac{\mathrm{Var}(R) - \mathrm{Var}(E[R|X_{\sim i}])}{\mathrm{Var}(R)} \tag{2.12}$$

Sudret (2008) derives alternative equations which further clarify the meanings of these terms in the context of studying model emulators. Critically, these terms can be expressed as multi-dimensional integrals over model input parameters which can be approximated via Monte Carlo or other sampling techniques, although it is very computationally expensive to do so. We adopt the column swap-out sampling method of Weirs et al. (2012) which combines Latin Hypercube Sampling with perturbed combinations of parameters to efficiently approximate S_i and T_i (see also Saltelli et al. (2010)). However, the sampling procedure adopted still requires a large number of model evaluations; for an initial n-size sample of the M parameters being studied, the method requires $(2 + M) \times n$ evaluations of the full-complexity model. We found that the computed S_i and T_i converged to stable values for the parcel model (and other parameterizations studied here) for $n \sim \mathcal{O}(10^3)$, and used $n = 1280$ to derive the values reported here.

Although the sampling procedure can be repeated for the emulators derived via polynomial chaos expansion, the orthogonality of the basis terms comprising each expansion lends itself to a more direct computation of S_i and T_i. Following Sudret (2008), we compute these indices directly from the coefficients of the derived chaos expansions. The sampling technique used to derive Sobol' indices for the parcel model, applied to the chaos expansions, produces similar estimates to those computed from the coefficients.

2.3 Results

2.3.1 Evaluation of Emulators

In order to assess the performance of the emulator, two sets of $n = 10,000$ samples were drawn using maximin Latin Hypercube Sampling from the parameter space defined in Table 2.1. This randomized design ensured that representative, equal numbers of samples were drawn from across the multi-dimensional parameter space. In the first set, variables whose logarithms were used to build the emulator were sampled in logarithmic space; in the second set, these variables were transformed back to their original values (e.g. from $\log_{10} N$ to N) before the sample was constructed. The two independent sets were blended together to assess the emulator. This helps ensure that both very high and very low values of the log-transformed are thoroughly represented within the sample. The set of sample parameter sets were run through all the derived chaos expansions of all orders, as well as the detailed parcel model as a reference benchmark for activation dynamics.

Figure 2-1 illustrates the performance of a 4th-order expansion whose coefficients were derived using ordinary least squares. The large range of initial temperatures, pressures, aerosol populations, and updraft speeds sampled here leads to a very large range of supersaturation maxima achieved by the ascending parcel. Weaker updraft speeds are generally associated with lower maximum supersaturations and corresponding to lower aerosol activated fraction; the opposite is true when strong updrafts are present, although there are some cases where a strong updraft activates a small fraction of aerosol. This typically occurs when initial aerosol size distribution is shifted towards larger radii and under polluted conditions with aerosol number concentrations greater than $3,000$ cm^{-3}. However, over the large parameter space sampled, the chaos expansion accurately reproduces the parcel

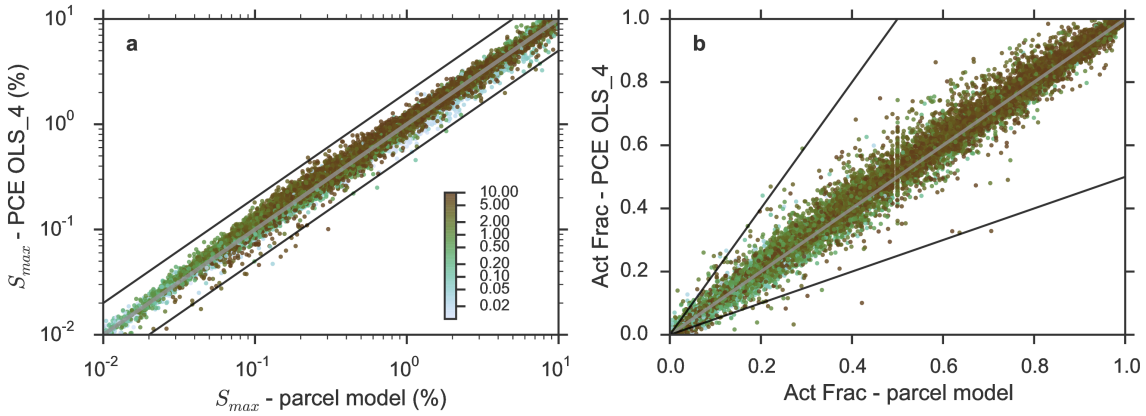

Figure 2-1: One-one plot comparing predicted supersaturation maximum (a) and diagnosed equilibrium droplet activated fraction (b) between parcel model and a polynomial chaos expansion of order $p = 4$, with coefficients computed using ordinary least squares. Black lines denote factor of 2 difference between predicted values using parcel model and those computed with the parameterization. Glyph shading denotes updraft velocity, V, with corresponding scale on panel (a).

model's determination of $S_{\max}$ and corresponding activated fraction. This is even true for predictions of small $S_{\max}$, which could potentially have a larger bias since the predicted error is expected to be uniform in $\log_{10}(S_{\max})$.

For parameter sets leading to a large activated fraction of 0.8-1.0, the relative error of the chaos expansion (compared to the parcel model) rarely exceeds 5%, and on average (for all activated fractions) is 5.7%. While the mean relative error in each activated fraction decile is close to 0, the standard deviation in the relative error tends to increase for the lower ones; the standard deviation in relative error decreases from 19.7% for activated fractions ranging from 0.1-0.2 to 3.5% for those ranging from 0.8-0.9. This suggests that there is a non-linear component in the mapping from the input parameter space to the emulated maximum supersaturation and diagnosed droplet number concentration which is prevalent in the weak droplet activation regime; the predicted activated fraction is more sensitive to small changes in the input parameters in this regime than in others.

Increasing the order of the chaos expansion tends to improve the accuracy of the predicted $S_{\max}$, as recorded in Table 2.2. However, there is not much difference between the methods used to compute the coefficients of the expansion beyond expansion order. For example, for the 4th- and 5th-order expansions, the expansions perform equally well regardless of what method (OLS, LARS, or LASSO) was used to compute the coefficients when con-

Table 2.2: Summary statistics for supersaturation maxima predicted by chaos expansions derived in this study compared to detailed parcel model calculations. The chaos expansions are organized by the method used to derive their coefficients and the expansion order in the two left-most columns. From left-to-right, the reported statistics are the normalized root mean square error (NRMSE), coefficient of determination (r^2, mean relative error (MRE), and standard deviation of the mean relative error (MRE std. dev.)

		NRMSE	r^2	MRE	MRE std. dev.
LASSO	2	0.292766	0.877600	-4.963330	31.098459
	3	0.193247	0.946670	-1.012509	19.797674
	4	0.112323	0.981983	3.941698	13.396632
	5	0.119231	0.979699	0.772525	9.904727
LARS	2	0.325442	0.848752	3.055232	39.776865
	3	0.250264	0.910559	-5.433902	17.853369
	4	0.104401	0.984435	-0.517486	13.092187
	5	0.135634	0.973729	-0.572383	9.962430
OLS	2	0.266774	0.898368	7.956622	40.762613
	3	0.220034	0.930861	-1.232172	20.441521
	4	0.201934	0.941768	4.243836	14.483854
	5	0.128687	0.976351	-0.259165	10.430909

sidering the mean and spread of the relative error to the parcel model reference simulations. In all cases, the chaos expansions produce very large r^2 values and normalized fractional root mean square errors (RMSE/ $\sum_{i=1}^{n}(X_i^2/n)$) which decrease as the order of the expansion increases.

These same statistics, computed for the diagnosed droplet number concentration given the predicted supersaturation maximum, are summarized in Table 2.3. Here, the trend is similar to before; increasing the order of the expansion tends to improve the accuracy of the diagnosed number concentration in terms of mean relative error, and also tends to decrease spread around that value. Third-order expansions tend to produce more accurate results with respect to the mean relative error, but this is overshadowed by the fact that there is far more variance in their predicted values as indicated by the standard deviation of their relative errors, which are almost twice as large as those of the higher-order expansions.

Table 2.3: Same as Table 2.2, but for predicted droplet number concentration.

		NRMSE	r^2	MRE	MRE std. dev.
LASSO	2	0.165043	0.929439	5.697161	49.813356
	3	0.108790	0.969342	-0.930886	20.270639
	4	0.074406	0.985659	2.545998	19.834428
	5	0.058872	0.991022	1.441665	14.497129
LARS	2	0.167168	0.927611	3.524631	39.031227
	3	0.121896	0.961510	-0.030537	34.613813
	4	0.075984	0.985044	-0.280769	16.430850
	5	0.058961	0.990995	0.884227	17.817749
OLS	2	0.174550	0.921076	7.640377	43.978774
	3	0.125045	0.959496	0.380943	28.562715
	4	0.079971	0.983434	2.556756	20.929598
	5	0.061762	0.990119	1.295857	17.903202

2.3.2 Comparison with Other Parameterizations

We compare the performance of the chaos expansion-based emulators to two existing parameterizations from the literature. The scheme by Abdul-Razzak and Ghan (2000) (ARG) - which is widely used in global models - utilizes a psuedo-analytical solution to an integro-differential equation derived from the adiabatic parcel system with embedded aerosol growing via condensation. This is in contrast to the scheme by Morales Betancourt and Nenes (2014a) (MBN), which is based on an iterative scheme to separate the aerosol population into subsets whose growth is inertially limited or not, and uses this information to derive a maximum supersaturation for a given parcel system. The MBN scheme is generally more expensive to evaluate than the ARG scheme due to its iterative nature, but is often more accurate owing to its consideration of the potentially important effect of large albeit unactivated aerosol particles (Simpson et al., 2014). In contrast with these schemes, our emulators simulate the activation process based on the explicit numerical solution obtained from a detailed parcel model which is similar to the one used to build and evaluate the MBN scheme (e.g. Nenes and Seinfeld, 2003)

The parameter sets used in Section 2.3.1 were also used to compute droplet activation with the ARG and MBN parameterizations. The relative errors between the supersaturation maximum and droplet number predicted by these schemes and the chaos expansions compared to the parcel model are illustrated in Figure 2-2. Both the ARG and MBN pa-

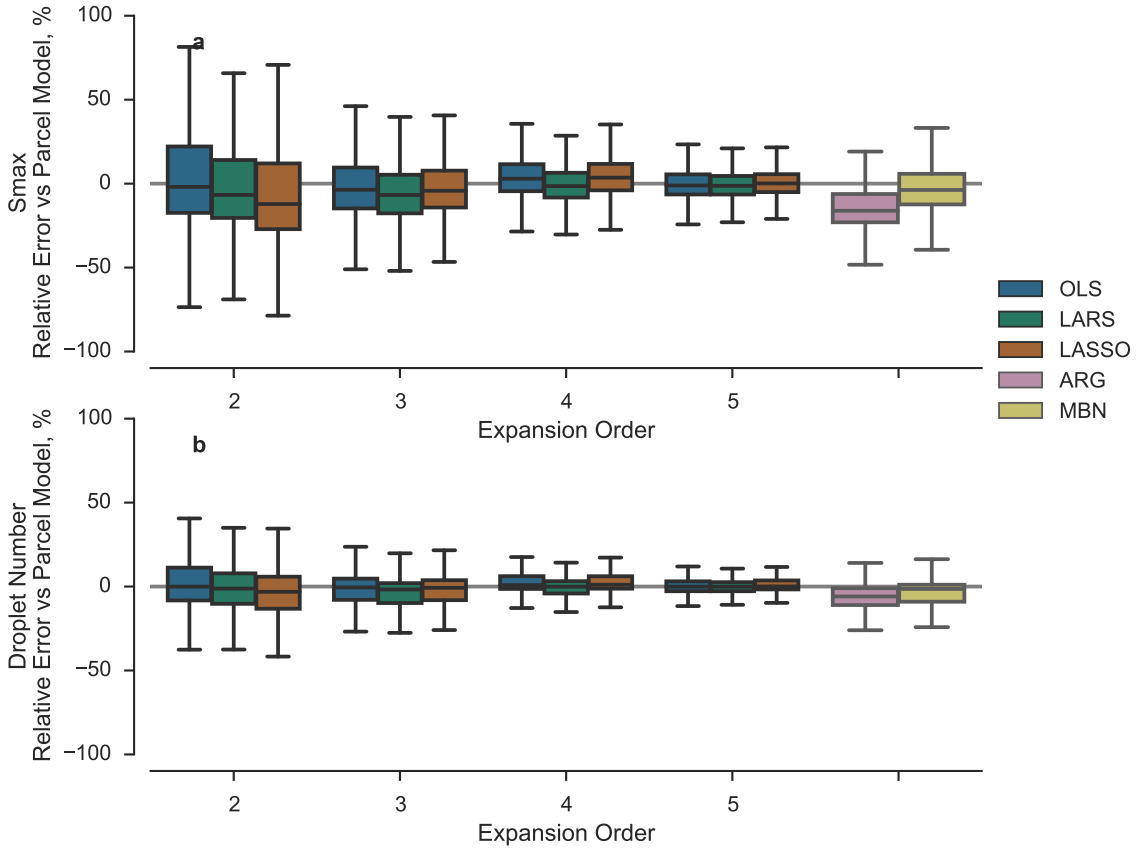

Figure 2-2: Boxplots illustrating mean relative error between supersaturation max (a) and droplet number concentration (b) predicted by chaos expansions and parameterizations versus detailed parcel model. The chaos expansions have been grouped by expansion order (x-axis) and method for computing their coefficients (OLS, LARS, and LASSO; hue). "ARG" refer to the scheme of Abdul-Razzak and Ghan (2000); "MBN" refers to the scheme of Morales Betancourt and Nenes (2014a).

rameterizations are more accurate than the 2nd and 3rd order chaos expansions. The ARG scheme tends to underpredict the maximum supersaturation, which is consistent with previous investigations into its performance (Abdul-Razzak and Ghan, 2000; Ghan et al., 2011; Simpson et al., 2014). This tends to produce a bias towards under-prediction of droplet number. The MBN scheme tends to yield more accurate predictions of both maximum supersaturation and droplet number. However, both schemes are out-performed by the 4th and 5th order chaos expansions, both on average and in terms of the variance of the predictions; for instance, the OLS-derived 4th and 5th order expansions yield relative error in predicted $S_{\max}$ with a mean and standard deviation of $4.2\% \pm 14.4\%$ and $-0.32\% \pm 10.4\%$, whereas the ARG and MBN schemes yield $-13.7\% \pm 18.2\%$ and $-2.4\% \pm 15.8\%$, respectively. This

is larger than other studies have reported, but we explore a much larger parameter space in our sampling for the purposes of deriving the chaos expansion.

As a consequence of tending to slightly under-predict $S_{\max}$, both the ARG and MBN schemes under-predict the number of activated droplets in the framework considered here. The mean relative error in droplet number predicted by the ARG and MBN schemes for the samples here are $-9.6\% \pm 23.4\%$ and $-4.9\% \pm 16.8\%$, respectively. All of the chaos expansions outperform the mean relative error of the ARG scheme, and those of order $p \geq 3$ do so with less variance.

In addition to producing low mean relative error in predicted $S_{\max}$ and droplet number activated, the chaos expansions also reproduce the dependence of activation dynamics on aerosol physical properties and updraft speed, as illustrated in Figure 2-3. In response to increasing aerosol number concentration, the $S_{\max}$ reached by an ascending parcel tends to decrease because there is a larger surface area available where condensation can occur, producing a larger source of latent heat release which limits the production of supersaturation. Overall, though, the droplet number concentration increases in spite of this effect as the aerosol activated fraction only decreases by a factor of 4 when the total number of initial aerosol increases by an order of magnitude (Figure 2-3a-b). Shifting the aerosol population to larger sizes (Figure 2-3c-d) produces a similar effect in inhibiting the increase in a parcel's supersaturation; however, Köhler theory predicts that these larger particles will more easily activate, which offsets the increase in $S_{\max}$ and yields larger droplet number concentrations. A similar effect occurs as aerosol hygroscopicity increases (Figure 2-3e-f).

The chaos expansions, as well as both the ARG and MBN schemes, capture these subtleties of activation dynamics as well as the detailed parcel model. More importantly, the expansions reproduce the sensitivity of activation to updraft speed (Figure 2-3g-h) which is the most important factor controlling $S_{\max}$ and setting the droplet number. At the largest updraft speeds of a few meters per second - indicative of deep, vigorous convection - the MBN scheme outperforms both the chaos expansions and the ARG scheme. However, for the aerosol population considered in Figure 2-3g-h (with $N = 1000\,\mathrm{cm}^{-3}$, $\mu = 0.05\,\mu\mathrm{m}$, and $\sigma = 2.0$), the relative error in predicted $S_{\max}$ by the chaos expansions at high updraft speeds does not substantially affect the diagnosed droplet number concentration, since in this case all but the smallest aerosol particles activate under equilibrium considerations. Note that this $S_{\max}$-overprediction coupled with an accurate assessment of activated fraction occurs

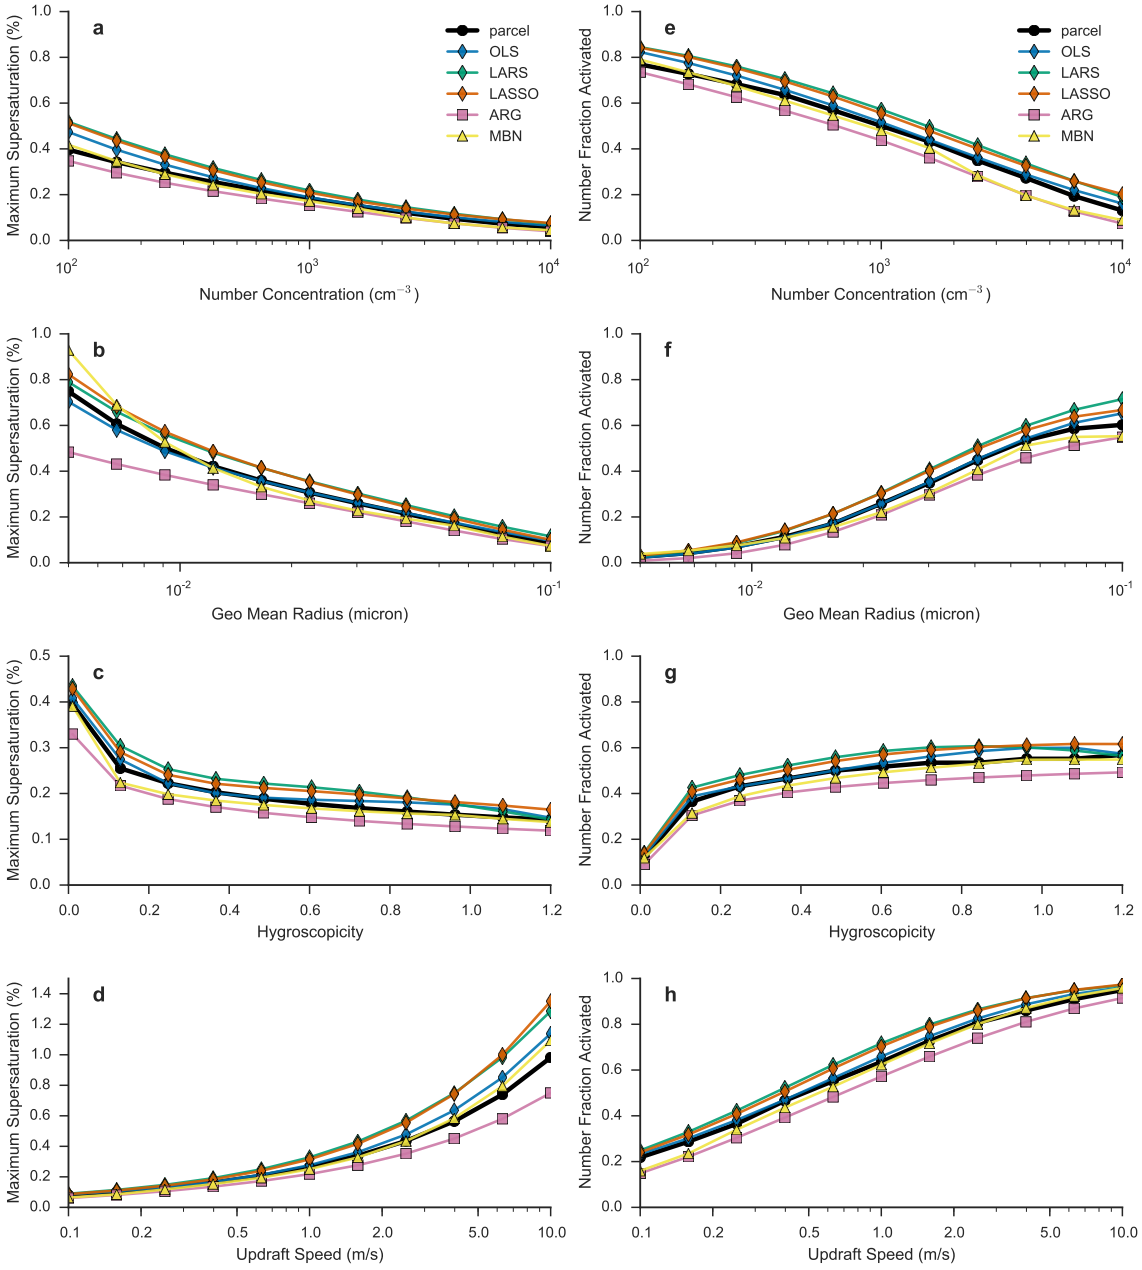

Figure 2-3: Sensitivity of parameterized and simulated maximum supersaturation (a-d) and activated number fraction (e-h) to changes in mode number concentration, mode geometric mean radius, mode hygroscopicity, and updraft speed with all other parameters held fixed at the values $T = 283\,\mathrm{K}$, $P = 850\,\mathrm{hPa}$, $V = 0.5\,\mathrm{m\,s^{-1}}$, $a_c = 1.0$, $\mu = 0.05\,\mu\mathrm{m}$, $\kappa = 0.54$, $N = 1000\,\mathrm{cm^{-3}}$, and $\sigma = 2.0$. "MBN" and "ARG" correspond to the schemes of Morales Betancourt and Nenes (2014a) and the update by Ghan et al. (2011) to Abdul-Razzak and Ghan (2000), respectively; the curves correspond to 4th order chaos expansions with coefficients derived using the named method.

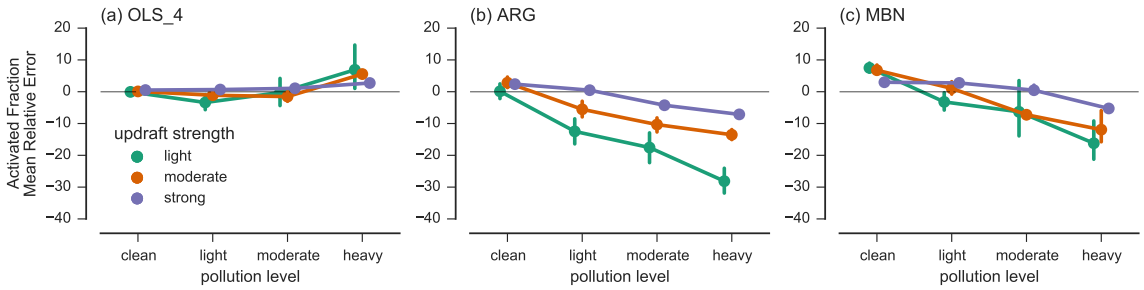

Figure 2-4: Mean relative error in activated fraction for 4th-order OLS-derived chaos expansion (a), ARG (b), and MBN (c) schemes relative to detailed parcel model. Updraft speeds range from light (10-50 $\mathrm{cm\,s^{-1}}$), moderate (0.5-2.0 $\mathrm{m\,s^{-1}}$), and strong (2.0-10.0 $\mathrm{m\,s^{-1}}$); pollution levels range from clean (10-250 $\mathrm{cm^{-3}}$), light (250-1000 $\mathrm{cm^{-3}}$), moderate (1000-2500 $\mathrm{cm^{-3}}$), and heavy (2500-10000 $\mathrm{cm^{-3}}$). Error bars denote 95% confidence interval on mean relative error from in-bin samples; samples with $\mu < 10\,\mathrm{nm}$ were omitted from these calculations.

for many different single-mode, lognormal aerosol populations.

Although all the chaos expansion results in Figure 2-3 appear biased high compared to the parcel model, this bias does not hold in general. Note that the ARG scheme is biased low in this particular analysis; this is generally just an artifact of fixing seven of the eight parameters and analyzing one-dimensional transects. Fixing the non-varying parameters at different values tends to shift the bias positive or negative in a non-systematic way. Critically, the choice of values for these parameters does not affect the chaos expansion's ability to reproduce the sensitivity to the varying parameter, which lends confidence that the expansions accurately reproduce the behavior of the parcel model.

Since Figure 2-3 highlights the fact that different schemes potentially perform better in different parts of the parameter space governing droplet activation, we stratified the sampling results based on level of pollution and updraft speed strength and computed activated fraction relative error statistics in each of these bins as shown in Figure 2-4. All of the schemes are accurate in clean and lightly polluted conditions (with aerosol number concentration, $N < 1000\,\mathrm{cm^{-3}}$). However, there is a tendency for both the ARG and MBN schemes to under-predict droplet number in heavily polluted conditions ($N > 2500\,\mathrm{cm^{-3}}$).

The 4th-order OLS-derived chaos expansion is plotted in Figure 2-4a as a representative example of the chaos expansions, and it retains its accuracy across the pollution level-updraft strength spectrum. The combination of light updrafts and heavy pollution tends to produce the largest under-prediction in activated droplet number, ranging from 10-30% for the ARG

and MBN schemes. Unsurprisingly, relative error in activated fraction tends to be least sensitive to increasing aerosol number concentration in the strong updraft regime. In this case, the vigorous updraft produces strong adiabatic cooling which overwhelms latent heat release from condensation as the droplets in the parcel grow, contributing to a rapid and large $S_{\max}$ (Equation (2.13)) and thereby activating a significant fraction of the aerosol.

It should be noted that parts of the parameter space we considered in applying the PCM, evaluating its output, and comparing to existing parameterizations may not be typical of real atmospheric cases. We chose a large parameter space in order to derive the most general emulator possible for this particular single mode aerosol case. In the real world, there should be some correlation between the ambient temperatures, pressures, and updraft speeds used when diagnosing aerosol activation, while we sample these factors as if they were independent from one another. The output from applying the parcel model to non-realistic activation scenarios could tend to inflate the computed relative errors and exacerbate differences between the parcel model and the existing, physically-based parameterizations.

2.3.3 Global Sensitivity Analysis

Total Sobol' indices, T_i, for each of the input parameters summarized in Table 2.1 corresponding to the prediction of $S_{\max}$ are plotted for each of the chaos expansions, parameterizations, and parcel model in Figure 2-5. The order of the chaos expansion does not impact the relative importance of each term between different schemes or for higher-order expansions of the same scheme. Both updraft speed and the aerosol distribution size parameter most strongly contribute to the variance in $S_{\max}$, followed by the total number of aerosol. This is in contrast with the chemical parameters in the model, a_c and κ, provide the weakest constraints, suggesting the importance of the total aerosol surface area in dominating the potential for droplet activation by controlling $S_{\max}$.

Indices derived using the chaos expansions very closely approximate those derived by sampling the full parcel model. This is expected since they simply provide an alternative framework for calculating the indices from the parcel model. However, the differences between the ARG and MBN schemes and the parcel model highlight the potential for biases in these schemes due to over-sensitivity to particular model parameters. For instance, the ARG scheme is less sensitive to variations in σ_g than the full parcel model; however, the

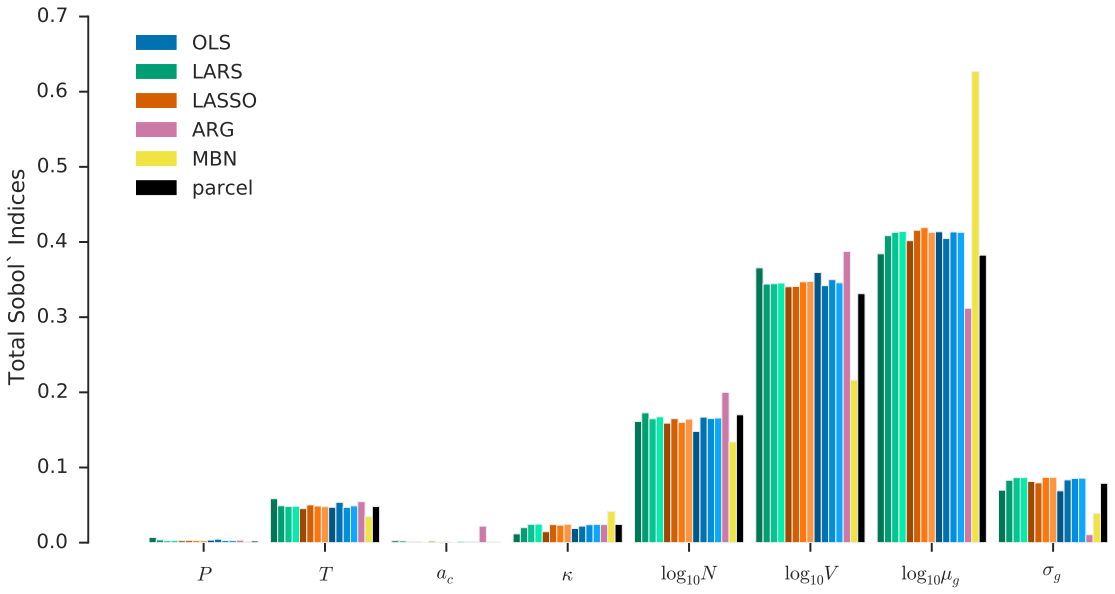

Figure 2-5: Total Sobol' indices corresponding to the prediction of $S_{\max}$ for each parameter in Table 2.1 for each chaos expansion (OLS, LARS, and LASSO), parameterization (ARG and MBN), and parcel model. For the chaos expansions, the ligher colors indicate successively higher-order expansions.

dependence of $S_{\max}$ in that scheme on σ_g is tuned to their own numerical calculations which may differ from ours (Abdul-Razzak et al., 1998). Furthermore, the ARG scheme is more sensitive to variations in a_c than our parcel model and emulators, which parameterizes the dependence of the condensational growth coefficient by a simple re-scaling against a reference value computed for $a_c = 1$ rather than explicitly account for it.

Maximum supersaturations produced by the MBN scheme are more sensitive to the geometric mean size of the aerosol than those from the parcel model or ARG scheme, and generally less sensitive to the strength of the updraft speed. The relative importance of each term for both the ARG and MBN scheme generally agrees with the estimates from the parcel model and chaos expansions, although the ARG scheme is most sensitive to updraft speed rather than geometric mean aerosol size. By means of those parameters' higher T_i, the aerosol size distribution parameters exert far more influence over $S_{\max}$ robustly for both parameterizations and the parcel model.

The chaos expansions contextually the relative importance of each input parameter with the potential for important interactions between terms. The leading eight terms ranked by main Sobol' index computed using the OLS-dervied chas expansions are summarized in

Table 2.4: Input parameters and combinations ranked by Main Sobol' index for the OLS-based chaos expansions.

Order	2		3		4		5	
Rank	term	Main (S_i)	term	Main (S_i)	term	Main (S_i)	term	Main (S_i)
1	$\log_{10}\mu_g$	0.379	$\log_{10}\mu_g$	0.358	$\log_{10}\mu_g$	0.360	$\log_{10}\mu_g$	0.359
2	$\log_{10} V$	0.331	$\log_{10} V$	0.311	$\log_{10} V$	0.315	$\log_{10} V$	0.312
3	$\log_{10} N$	0.122	$\log_{10} N$	0.129	$\log_{10} N$	0.127	$\log_{10} N$	0.129
4	σ_g	0.049	σ_g	0.058	σ_g	0.055	σ_g	0.056
5	T	0.042	T	0.052	T	0.045	T	0.047
6	$\log_{10} V \log_{10}\mu_g$	0.016	$\log_{10} N \log_{10}\mu_g$	0.022	$\log_{10} N \log_{10}\mu_g$	0.021	$\log_{10} N \log_{10}\mu_g$	0.021
7	$\log_{10} N \log_{10}\mu_g$	0.016	κ	0.017	κ	0.019	κ	0.019
8	κ	0.015	$\log_{10} V \log_{10}\mu_g$	0.017	$\log_{10} V \log_{10}\mu_g$	0.018	$\log_{10} V \log_{10}\mu_g$	0.017

Table 2.4. For all four (2nd-4th order) chaos expansions summarized, all the terms with rank greater than eight were an order of magnitude less important than those in the table. With the exception of the 2nd-order scheme, the rankings of the top eight terms do not change relative order, and the main terms dominate the higher-order ones. The only higher-order terms which contribute grossly to the variance in $S_{\max}$ are combinations of the updraft speed and aerosol size distribution parameters, reiterating their importance compared to the chemistry terms (only κ appears in the top eight).

2.3.4 Computational Efficiency of Chaos Expansions

As detailed in Appendix 2.6, evaluating the emulator produced by the chaos expansion requires two sets of straightforward floating-point operations. The first set requires the projection of the input parameters into the vector space spanned by the basis polynomials of the chaos expansion, which can then be used to evaluate the basis polynomials up to the required order. The remaining operations simply multiply these intermediate evaluations together and sum them to evaluate the full expansion. In general, this procedure should lie in between the ARG and MBN schemes in terms of computational complexity. The ARG scheme relies on straightforward floating point operations to derive an estimate for $S_{\max}$ which at worst involve evaluating a logarithm. However, the MBN scheme requires sets of iterations, each of which necessitates a costly evaluation of the error function and the complementary error function.

On average, the 2nd and 5th-order OLS-derived chaos expansion was 10-17 times faster than the MBN scheme given the same single mode aerosol population. The exact speedup depended on the background velocity; for weak updraft speeds, the performance of the MBN

scheme fared better, although it became much worse for updrafts where $V < 2\,\mathrm{m\,s^{-1}}$. The ARG scheme was consistently 1-3 times faster than those same chaos expansions. Since the pathway for evaluating either the ARG or chaos expansion schemes do not change depending on the input parameters, their performance was the same regardless of what inputs were provided.

2.4 Summary and Conclusions

An efficient parameterization of droplet activation for a single aerosol model under a wide variety of different physico-chemical properties and thermodynamic conditions was developed via statistical emulation of a detailed parcel model using polynomial chaos expansion. The emulators predict the maximum supersaturation achieved by a parcel, which is then used to diagnose activated droplet number using Köhler theory in a similar framework to existing activation parameterizations. The 4th- and 5th-order chaos expansions derived from the detailed parcel model are more accurate on average than two commonly-used, physically-based parameterizations from the literature (Abdul-Razzak and Ghan, 2000; Morales Betancourt and Nenes, 2014a). Additionally, the chaos expansions are all at least 10 times faster to evaluate than the MBN scheme and only about twice as expensive as the ARG scheme. A simple algorithm was suggested for evaluating a chaos expansion which requires a minimal amount of data about the expansion (such as the basis polynomials and the coefficients of the expansion terms) to be saved; in this way, the chaos expansions offer a method for extending lookup tables to very high dimensionalities without suffering from exponentially-rising storage costs.

Based on the large set of aerosol properties and thermodynamic conditions we sampled in order to derive and evaluate the chaos expansions, we observed that our emulators particularly outperform the existing schemes in conditions where a light updraft and heavy aerosol pollution (with respect to number concentration) are present. Because the ultimate goal of an activation parameterization is to couple the aerosol physics and chemistry to the cloud microphysics of a global-scale model, this deficiency in the existing parameterizations could be particularly important. Few global models have aerosol-microphysics connections in their deep convection parameterizations, but many source potential cloud droplet formation based on a detailed aerosol activation calculation for their shallow convection and stratiform

microphysics schemes. These schemes sometimes artificially restrict the lowest possible updraft speed available for estimating droplet activation, but as a consequence they ensure that weak updrafts make up a large portion of the activation conditions considered during a model run. In regions of the world with heavy anthropogenic aerosol pollution - such as southern and eastern Asia - this provides a recipe for systematically under-predicting droplet number and potentially impacting either a global model's simulated aerosol indirect effect on climate or the modeled aerosol-cloud interaction's sensitivity to changes in anthropogenic aerosol emissions.

The global sensitivity analysis framed on the input parameter set used to derive the new chaos expansion emulators provides an additional, new check on the performance of existing activation schemes compared to the detailed parcel model. The breakdown of main Sobol' indices calculated using the chaos expansions provides insight into the importance of interactions between the dominant first-order terms (updraft speed, number concentration, and geometric mean size of the aerosol distribution), which are further summarized by the total Sobol' indices derived for the parcel model and both ARG and MBN schemes. The over-sensitivity of the MBN scheme to the geometric mean radius and the under-sensitivity of the ARG scheme to the updraft speed contribute to their disagreement with the parcel model across the range of pollution levels and updraft speeds studied here. Further sensitivity analyses could shed additional light on the potential biases of activation schemes, and could provide useful metrics for evaluating the improvement of parameterizations more generally than simple ones based on relative or absolute error alone.

Critically, the framework from which the chaos expansions reported here are derived is extendable to the case where a complex, multi-species/multi-modal aerosol population is tracked by a global model; in that case, the number of parameters describing the aerosol size distribution and chemical composition simply increases. Future work will derive chaos expansions emulating activation for a multi-modal aerosol distribution specific to a particular global aerosol-climate model. Additionally, physical processes not considred here can also be introduced into the chaos expansion framework. For instance, entrainment can be incorporated into the parcel model following Seinfeld and Pandis (2006) and Barahona and Nenes (2007). Subgrid-scale variability in updraft speeds due to the coarse resolution of global model grids and the distribution of these updrafts can be represented either by a characteristic value (Morales and Nenes, 2010) or by numerical integration over a dis-

tribution (Lohmann et al., 1999; Golaz et al., 2011). In the latter case, many activation calculations must be performed, incurring a large computational cost. However, the entire integration over a spectrum of droplet speeds could be parameterized in the chaos expansion framework, greatly reducing the cost of this calculation and potentially improving the accuracy of diagnosed cloud droplet number.

As the complexity of global aerosol-climate models increases with respect to the number of aerosol modes and species tracked by the model, there is a pressing need to understand how biases in activation calculations across the high-dimensional parameter spaces defining the aerosol-climate model affect cloud properties and ultimately impact modeled climate. This work highlights a novel way to build efficient, accurate activation schemes for this purpose akin to customized lookup tables which cannot themselves extend to cover the necessary parameters. Employing such schemes should help improve simulated cloud microphysical properties and constrain modeled aerosol indirect effects on climate.

2.5 Appendix A: Parcel Model Description

The adiabatic cloud parcel model implemented for this study follows the basic equations of Pruppacher and Klett (1997) and adopts the framework used by Nenes et al. (2001) to account for kinetic limitations on droplet growth. Fundamentally, the model integrates a system of coupled ordinary differential equations which describe the thermodynamic evolution of an adiabatically-lifted, non-entraining parcel. In all the simulations described here, we use the Variable-coefficient Ordinary Differential Equation solver (VODE; Brown et al., 1989) to integrate the system forward in time.

The model tracks the evolution of supersaturation, S, with respect to water as

$$\frac{dS}{dt} = \alpha(T,P)V - \gamma(T,P)\frac{dw_c}{dt} \tag{2.13}$$

where $\alpha(T,P) = \frac{gM_wL}{c_pRT^2} - \frac{gM_a}{RT}$ and $\gamma(T,P) = \frac{PM_a}{e_sM_w} + \frac{M_wL^2}{c_pRT^2}$ are functions which are weakly dependent on temperature and pressure (Leaitch et al., 1986), M_w and M_a are the molecular weights of water and air, L is the latent heat of vaporation of water, c_p is the specific heat of dry air at constant pressure, R is the universal gas constant, g is the acceleration due to gravity, e_s is the saturation vapor pressure, and w_c is the liquid cloudwater mass mixing

ratio. Equation (2.13) expresses the supersaturation as a balance between production due to adiabatic cooling and loss due to latent heat release. This same framework describes the parcel's change in temperature over time,

$$\frac{dT}{dt} = -\frac{gV}{c_p} - \frac{L}{c_p}\frac{dw_v}{dt} \tag{2.14}$$

where V is the updraft velocity and w_v is water vapor mass mixing ratio. Water mass is conserved as vapor condenses into cloud water,

$$\frac{dw_v}{dt} + \frac{dw_c}{dt} = 0 \tag{2.15}$$

Equations (2.13)-(2.15) are linked through the growth of the cloud droplet population from the initial aerosol. Given n bin sizes, each associated with a number concentration N and a radius r, the change in cloud water can be written as

$$\frac{dw_c}{dt} = \frac{4\pi\rho_w}{\rho_a}\sum_{i=1}^{n} N_i r_i^2 \frac{dr_i}{dt} \tag{2.16}$$

where ρ_w and ρ_a denote the density of water and air, respectively.

The diffusional growth rate for droplets in the ith bin is calculated by

$$\frac{dr_i}{dt} = \frac{G}{r_i}\left(S - S_{\text{eq}}\right) \tag{2.17}$$

where S is the environmental supersaturation, S_{eq} is the Köhler-predicted equilibrium supersaturation of the droplet, and G is a growth coefficient which is a function of both the physical and chemical properties of the particle receiving condensate,

$$G = \left(\frac{\rho_w RT}{e_s D_v' M_w} + \frac{L\rho_w[(LM_w/RT) - 1]}{k_a' T}\right)^{-1} \tag{2.18}$$

Non-continuum effects on the diffusivity (D_v') and thermal conductivity (k_a') factors are accounted for with the corrections

$$D_v' = D_v \Big/ \left(1 + \frac{D_v}{a_c r}\sqrt{\frac{2\pi M_w}{RT}}\right) \tag{2.19}$$

and

$$k'_a = k_a \Big/ \left(1 + \frac{k_a}{a_T r \rho_a c_p}\sqrt{\frac{2\pi M_a}{RT}}\right) \tag{2.20}$$

In these correction terms, the thermal accommodation coefficient, a_T, is assumed to be 0.96; the condensation coefficient, a_C, is allowed to vary as observations suggest it could take values between $0.1 - 1.0$ (Raatikainen et al., 2013). The instantaneous droplet growth rate is further modulated by the difference between the environmental supersaturation, S, and the saturation ratio over the surface of the aqueous droplet, S_{eq}. We treat the droplet-dependent S_{eq} following Petters and Kreidenweis (2007), who employ a single-term κ to parameterize particle hygroscopicity; values of κ can be derived from laboratory experiments. Under the framework of κ-Köhler theory the curvature effect term remains the same, while the solute effect term is re-written such that

$$S_{\text{eq}} = \frac{r^3 - r_d^3}{r^3 - r_d^3(1-\kappa)} \exp\left(\frac{2M_w\sigma_w}{RT\rho_w r}\right) \tag{2.21}$$

where r and r_d are the droplet radius and the dry radius of its embedded aerosol particle (which is tracked for each initial aerosol size in the model), and σ_w is the droplet surface tension, which we take to be independent of the droplet solution composition and described following the recommendation of Pruppacher and Klett (1997), $\sigma_w = 0.0761 - 1.55\times10^{-4}(T - 273)$. A limitation of the this approach for computing S_{eq} is that it is not convenient to derive analytical expressions for the critical supersaturation and radius; they must be computed numerically by finding the value r_{crit} such that

$$\left.\frac{\partial S_{\text{eq}}}{\partial r}\right|_{r_{\text{crit}}} = 0 \tag{2.22}$$

and then computing $S_{\text{crit}} = S_{\text{eq}}(r_{\text{crit}})$ for a given κ and r_d. This is accomplished using Brent's method (Brent, 1973) and by bounding r_{crit} from below with the observation that $r_{\text{crit}} > r_d$.

Finally, the parcel thermodynamic description is closed by predicting the pressure change within the ascending parcel following the hydrostatic relationship, which can be written using the ideal gas law as

$$\frac{dP}{dt} = -\frac{gPV}{R_d T_v} \tag{2.23}$$

where T_v is the virtual temperature, which is employed to account for changes in air density due to loss of water vapor to condensate. Equations (2.13)-(2.16), (2.23), and (2.17) applied to each n droplet size bins form a closed system which conserves total water mass.

2.6 Appendix B: Chaos Expansion Emulator Evaluation

The PCM as applied here produces two outputs: a P-length vector of coefficients α comprised of real values, and a $P \times M$ matrix of orthogonal polynomial orders Φ comprised of integers. Each term in matrix Φ contains a multi-index component for each term in the chaos expansion and indicates the order of the orthogonal polynomial corresponding to term $1 \leq j \leq M$ for expansion term $0 \leq i \leq P$. For any expansion, $\max[\Phi] = p$, the desired order of the chaos expansion. Algorithm 1 describes the evaluation of a chaos expansion.

Algorithm 1 Psuedo-code for evaluating a polynomial chaos expansion of the form given in Equation (2.2), applied to the computation of $S_{\max}$

for all X_j **do**
    $Z_j \leftarrow$ project X_j
end for
$\hat{S} \leftarrow 0$
for row $i = 0$; $i \leq P$ **do**
    $\hat{S}_i \leftarrow 1$
    for column $j = 0$; $j \leq M$ **do**
        $k \leftarrow \Phi[i, j]$
        $P_{j,k} \leftarrow \phi_j^{\mathrm{k}}(Z[j])$
        $\hat{S}_i = \hat{S}_i \times P_{j,k}$
    end for
    $\hat{S} = \hat{S} + \alpha[i] \times \hat{S}_i$
end for
$S_{\max} \leftarrow 10^{\hat{S}}$

Evaluating the chaos expansion involves two parts. First, the input parameters must be projected to conform to the space supported by the PDFs associated with each basis polynomial type, producing a set of parameters Z_j. In general, a set of mixed orthogonal polynomials could be used to derive a chaos expansion, but here only Legendre polynomials were used, and each parameter can be projected using Equation (2.5). Second, the

polynomial can be evaluated by treating Φ as a look-up table for the orders of each basis orthogonal polynomial. In practice, the evaluation of these orthogonal polynomials at Z_j for orders up to k can be efficiently precomputed (before the polynomial evaluation loop) by existing orthogonal polynomial libraries (Gautschi, 1994).

In general, the only computationally complex part of the chaos expansion evaluation algorithm is the projection from X_j to Z_j; given certain basis orthogonal polynomials and their associated PDFs, this procedure could involve numerical integration or otherwise complicated function evaluations. In the case of simple uniform PDFs, though, the process is achieved entirely by re-scaling the parameters, with little computational overhead. Furthermore, although evaluating a chaos expansion requires looping over each of its terms, each term can be computed independently from one another and efficiently optimized. This is in contrast with an iterative scheme which could involve numerical integration or other costly operations and must be performed in sequence.

Chapter 3

Development and Assessment of an Aerosol Activation Metamodel for the CESM/MARC Aerosol-Climate Model

3.1 Introduction

Ambient aerosol play a critical role in the climate system by interacting with radiation through several different mechanisms. Depending on its composition, aerosol can directly scatter or absorb incoming solar radiation, leading to a direct radiative effect and rapid changes in the energy budgets of the surface and atmosphere. Additionally, aerosol mediate the production of clouds by providing surface area on which water vapor may condense to form droplets. Through this second pathway, changes in the aerosol population perturb the radiative properties of clouds by altering their microstructure and lifecycle, thereby impacting the planetary radiative budget. Despite decades of focused research by the scientific

This chapter is adapted from Rothenberg, D., and C. Wang, 2016. An aerosol activation metamodel of v1.2.0 of the pyrcel cloud parcel model: Development and offline assessment for use in an aerosol-climate model. *Geoscientific Model Development Discusssions.* doi:10.5194/gmd-2016-228.

community, the radiative forcing produced through this second pathway, known as aerosol-cloud interactions, remains one of the largest uncertainties in understanding contemporary and future climate change on both regional and global scales (Boucher et al., 2013).

Contemporary Earth System Models are a valuable tool for assessing this uncertainty. Compared to their predecessors, these models incorporate aerosol microphysics schemes which represent the global variation in particle size distributions and aerosol composition. These aerosol size distributions can then be used to predict cloud droplet number concentrations. As a result, changes in the aerosol size distribution due to anthropogenic emissions can impact cloud optical properties and produce an aerosol-climate indirect effect. The interactions between aerosol particles, water vapor, and cloud droplets are often described using the conceptual model of a possibly-entraining, adiabatic cloud parcel (e.g. Feingold and Heymsfield, 1992; Nenes et al., 2001; Ervens et al., 2005; Topping et al., 2013). This parcel model framework provides a convenient means for both direct numerical and laboratory simulation of aerosol-cloud interactions, and has been used to establish the physical basis through which change sin the available aerosol perturb cloud radiative processes (Feingold et al., 2001).

However, it is not practical to directly include parcel model calculations in global models due to their coarse grid scales. Adiabatic parcel theory describes a process which occurs on a spatial scale of tens of meters, over the course of a few seconds—scales much finer than those resolved in many global model simulations. Some efforts have sought to incorporate fine-scale information about aerosol-cloud interactions by embedding higher-resolution models within each grid cell of a global model (for example, the "multi-scale modeling framework" of Wang et al., 2011). However, while this approach has improved the representation of cloud processes in global models, it still does not resolve the scales of motion associated with parcel theory. To bridge this gap in spatial scales, global models employ activation parameterizations which predict cloud droplet number concentrations using information about the subgrid-scale variability in meteorology and aerosol size distribution.

Twomey (1959) first developed a scheme for placing an upper bound on this number of droplets derived from geometrical arguments applied to parcel theory. In subsequent decades, these arguments were used in conjunction with aircraft measurements to broadly characterize typical cloud droplet numbers in different regimes. For instance, Boucher and Lohmann (1995) developed power-law relationships between droplet number and sulfate

mass concentration for separate continental and maritime regimes; similarly, Jones et al. (2001) developed a global parameterization for droplet number based on total aerosol number concentration using a survival function. While appropriate for earlier climate models which did not seek to resolve the complexity of the ambient aerosol, contemporary coupled aerosol-climate and aerosol-cloud-resolving models include more details about aerosol properties which can be used to predict cloud droplet number through explicit activation calculations. Broadly speaking, efforts to improve these activation calculations have proceeded down two avenues: caching of pre-computed, detailed parcel model calculations; and extensions to Twomey's geometrical arguments using physically-based approximations.

The simplest technique applied to parameterizing aerosol activation is to pre-compute, using a detailed parcel model, a set of aerosol and meteorological conditions sampled over some pre-determined parameter space (Saleeby and Cotton, 2004; Ward et al., 2010). These samples can then be used to construct a look-up table, using some form of interpolation to compute activation results for points within the domain of the initial parameter space. However, changing or increasing the number of parameters—such as including additional aerosol modes or a composition dimension—in the look-up table rapidly increases its computational cost, since it is costly to both store and access the required information in memory during run-time.

Alternatively, one can extend Twomey's analytical approach to bounding the maximum supersaturation achieved in the constant-speed, adiabatic ascent of a parcel (e.g. Cohard et al., 1998; Cohard J-M. and Pinty, 2000). Several techniques have been employed to account for the sensitivity of the maximum supersaturation achieved by such a parcel to variations in the initial dry aerosol particle size distribution (Khvorostyanov and Curry, 2006; Shipway and Abel, 2010; Shipway, 2015). These techniques yield closed-form expressions which relate the parameters describing the initial aerosol size distribution and updraft speed to the supersaturation maximum, and can be used to predict droplet number concentrations for a given scenario.

Such studies have led to a number of explicit activation schemes, each applying a different computational and analytical approach but fundamentally relying on the same set of physical approximations (see Ghan et al., 2011, for a more complete review). Some of these schemes have been extended to account for additional physical or chemical processes which can influence activation, such as the presence of organic surfactants on the surface of

droplets which tends to reduce surface tension (Abdul-Razzak and Ghan, 2004), or to account for entrainment in the ascending parcel (Barahona and Nenes, 2007). These schemes generally rely on iterative calculations to settle on an estimate of how many aerosol particles activate into cloud droplets (e.g. Shipway and Abel, 2010; Nenes and Seinfeld, 2003; Ming et al., 2006), although some employ pseudo-analytical solutions to avoid this process (Abdul-Razzak et al., 1998).

Since they offer generalized formulations independent of the representation of the aerosol particle size distribution provided to them, physically-based schemes have been the preferred tool for coupling with global aerosol-climate models. But, they are not without their shortcomings. Simplifying assumptions used to construct physically-based schemes, such as particles' equilibrium growth in response to changes in the ambient relative humidity, neglect kinetic limitations on growth and also lead to an underestimate in cloud droplet number in both polluted conditions and ones with weak updraft speeds (Nenes et al., 2001). Some schemes tend to under-estimate the number of droplets nucleated in the presence of multiple, competing aerosol modes, owing to their representation of water vapor uptake by particles (Simpson et al., 2014). Although Gantt et al. (2014) showed that using a variety of activation schemes in a modern aerosol-climate model with a complex aerosol particle size and composition distribution can lead to a global increase of estimated cloud droplet number by 155%, many evaluations of activation scheme performance have focused on the same set of relatively simple aerosol particle size distributions. For instance, Abdul-Razzak (2002) used the Whitby (1978) aerosol particle size distributions but with only two aerosol composition scenarios, varying the insoluble mass fraction in the coarse aerosol mode between either 0% or 90%. Nenes and Seinfeld (2003) and Fountoukis and Nenes (2005) used the same aerosol particle size distributions, but assumed just one additional aerosol composition scenario, exploring the impact of aerosol which are composed of 20% (by mass) of an organic which displays surface-active behavior; Ming et al. (2006) employed an identical set of evaluation simulations. Shipway and Abel (2010) restricted the analysis of their scheme's performance with multi-modal aerosol particle size distributions to an idealized bi-modal distribution with non-varying, homogeneous particle composition. Ghan et al. (2011) summarized the performance of all of these simulations with the same set of tri-modal aerosol particle size distributions, but extended their analysis to evaluate droplet number simulated by two schemes (Abdul-Razzak and Ghan, 2000; Fountoukis and Nenes, 2005) when employed in a

global model. Meskhidze (2005) used in situ data collected from two different aircraft observation campaigns to evaluate the sectional formulations of the Nenes and Seinfeld (2003) and Fountoukis and Nenes (2005) schemes; these evaluations were extended by Fountoukis et al. (2007), which analyzed the modal formulations of the same schemes, and Morales et al. (2011) which showed that accounting for entrainment reduced the over-prediction of droplet number in stratiform clouds by the activation schemes.

However, look-up table approaches do not necessarily solve these shortcomings. Detailed parcel models can more easily accommodate additional physics and chemistry which influences aerosol activation, but it would not be practical to incorporate this information into a look-up table. However, in previous work, Rothenberg and Wang (2016) developed a framework for producing an emulator of a detailed parcel model, which could be used to extend the benefits of look-up tables to high-dimensional parameter spaces. In their example, which focused on the activation of a single lognormal aerosol mode, the meta-models produced by such a framework had lower error statistics on average than existing activation schemes when compared to a benchmark parcel model.

In this work, we extend the methodology developed in Rothenberg and Wang (2016) to develop an emulator suitable for inclusion in a modern, coupled aerosol-climate model. Furthermore, we assess the performance of the emulator against existing activation schemes across a large input parameter space and a subset reproducing the tendencies of the aerosol-climate model. By reproducing the original, detailed parcel model on which it is trained, such an emulator could reduce biases in estimates of cloud droplet number concentration in cloud regimes characterized by either high pollution or relatively weak forcing and ascent. This could in turn lead to improved estimates of the aerosol indirect effect from global models.

This chapter is organized to reflect the exploratory process used to develop the activation emulator presented here, with the hope that clearly delineating the steps involved will encourage other groups to pursue similar lines of work. Section 3.2 provides background on the parent aerosol-climate model for which our activation emulator was derived. Section 3.3 details the construction of the emulators and the tools used to produce them. Two different evaluations of the emulators are presented in Section 3.4. Finally, in Sect. 3.5 we motivate future projects using these emulators and, more broadly, this approach to building parameterizations.

3.2 Activation Parameter Space

The parent aerosol-climate model for which we seek to derive a droplet activation emulator is the MultiMode, two-Moment, Mixing-state resolving Model of Aerosols for Research of Climate (MARC, version 1.0.1 here Kim et al., 2008, 2014). MARC extends the NCAR Community Earth System Model (CESM; version 1.2.2 here), which is a fully-coupled global climate model with sub-components for simulating climate processes in the land, ocean, atmosphere, and sea ice domains. In particular, MARC replaces the default modal aerosol treatment (Liu et al., 2012) in the atmosphere component of the CESM, the Community Atmosphere Model (CAM5; version 5.3 here), with a scheme which simultaneously resolves both an external mixture of different aerosol species and internal mixtures between others (Wilson et al., 2001). In this sense, MARC refers to a configuration of the CESM with the CAM5 atmosphere component and the alternative aerosol formulation.

The aerosol population within MARC is comprised of a tri-modal sulfate distribution (nucleation [NUC], aitken [AIT], and accumulation [ACC] modes), discrete modes for pure black carbon (C) and a generic organic carbon species (OC), and two internally-mixed modes, one of each for mixed sulfate-black-carbon (MBS) and sulfate-organic-carbon particles (MOS. The ratio of the masses within each mixed species evolves over time, changing the optical and chemical properties of those particles. For each of these seven modes, MARC predicts total particle mass (M) and number concentrations (N) for a corresponding log-normal size distribution with a prescribed width (geometric standard deviation; σ_g).

Additionally, both sea salt ("SSLTn") and dust ("DSTn") particle size distributions are described within MARC by a 4-bin, single-moment scheme with fixed particle sizes. For these single moment modes, MARC predicts a total number concentration for each bin and then diagnoses a total mass in order to simulate a lognormal mode with prescribed geometric mean radius (which is narrower for super-micron dust and sea-salt). Each mode has a prescribed particle hygroscopicity follow κ-Köhler theory (Petters and Kreidenweis, 2007) except for the MOS mode, which has a composition-dependent κ computed assuming that the carbon and sulfate in the particle forms a simple internal mixture. Unlike the MOS mode, the MBS mode assumes a core-shell structure with sulfate coating a black carbon nucleus, and has a fixed hygroscopicity corresponding to that of sulfate, $\kappa = 0.507$. The sea salt modes are assumed to be comprised of NaCl with $\kappa = 1.16$, and the dust modes are assumed to be a

mixture of minerals with a hygroscopicity of $\kappa = 0.14$ (Scanza et al., 2014). The organic and black carbon modes are assumed to be non-hygroscopic and not significant players in aerosol activation. These assumptions about the aerosol size distribution simulated by MARC are summarized in Table 4.1.

Alternative frameworks exist for parameterizing aerosol composition and hygroscopicity for use in global or process models. Sorjamaa and Laaksonen (2007) proposed using the Frenkel, Halsey and Hill (FHH) two-parameter adsorption theory which modifies the treatment of the Raoult effect in traditional Köhler theory and can accommodate condensation of monolayers on particle surfaces Although each of the parameters in this particular theory must be fit to empirical data for a specific aerosol composition, it has been used to parameterize the activation of insoluble aerosol particles (Kumar et al., 2009). In laboratory studies, Garimella et al. (2014) showed that κ-Köhler theory agreed with FHH adsorption theory for predicting critical supersaturations for large dust particles of atmospherically-relevant sizes. However, for application in global models like MARC, κ is advantageous because the hygroscopicity of particles with multiple constituents (such as in the MOS mode) can be estimated using the volume-weighted average of each component.

The aerosol size distributions predicted by MARC interact with both radiation and cloud microphysics. With respect to the latter, MARC adopts the two-moment, stratiform cloud microphysics scheme found by default in CAM5 (Morrison and Gettelman, 2008) which provides an explicit source of cloud droplet formation via aerosol activation from aerosol. This interaction is facilitated by means of a physical parameterization which takes as input both the physical and chemical properties of the ambient aerosol as well as limited information about meteorology—in particular, the distribution of subgrid-scale vertical velocities. In MARC, a single subgrid-scale characteristic updraft velocity (V) diagnosed from the grid cell turbulent kinetic energy (TKE) provided by the moist turbulence scheme (Park and Bretherton, 2009) and assumed to be isotropic is used to estimate droplet nucleation following Ghan et al. (1997) and Lohmann et al. (1999), such that

$$V = \overline{V} + \sqrt{\frac{2}{3}\mathrm{TKE}}$$

where $\overline{V}$ is the large-scale resolved updraft velocity. Furthermore, we limit $V > 0.2\,\mathrm{m\,s^{-1}}$ because the processes driving turbulence are not represented well in MARC, particularly

Aerosol Mode	Geometric Mean Particle Diameter (µm)	Geometric Std Deviation (σ_g)	Density ($\mathrm{g\,cm^{-3}}$)	Hygroscopicity (κ)
NUC	0 to 0.005 84	1.59	1.8	0.507
AIT	0.005 84 to 0.031	1.59	1.8	0.507
ACC	>0.031	1.59	1.8	0.507
OC	-	2.0	2.0	10^{-10}
MOS	-	2.0	†	†
BC	-	2.0	2.0	10^{-10}
MBS	-	2.0	2.0	0.507
DST01	0.16	1.4	-	0.14
DST02	0.406	1.4	-	0.14
DST03	0.867	1.4	-	0.14
DST04	1.656	1.4	-	0.14
SSLT01	0.5	1.59	-	1.16
SSLT02	2.0	1.37	-	1.16
SSLT03	5.0	1.41	-	1.16
SSLT04	15.0	1.22	-	1.16

Table 3.1: MARC aerosol mode size distribution and chemistry parameters. The MOS mode (†) has a composition-dependent density and hygroscopicity which is computed using the internal mixing state of organic carbon and sulfate present at a given grid-cell and timestep.

those driven by cloud-top radiative cooling above the planetary boundary layer (Ghan et al., 1997). Morales and Nenes (2010) and West et al. (2014) have explored how using different characteristic updraft velocities to represent subgrid-scale variability can influence simulated aerosol-cloud interactions; in particular, West et al. (2014) showed that using a similar TKE-based parameterization produced more realistic spatial and temporal variability in V, but tends to produce an unrealistically high frequency of its minimum-permissible value. We further assume that activation only occurs in non-entraining, adiabatic updrafts which carry air up and into the base of stratiform clouds.

3.2.1 MARC Aerosol and Meteorology Parameters

The set of size distribution parameters describing each aerosol mode in MARC and the number of meteorological factors influencing droplet nucleation in MARC is large, and each parameter can vary over several orders of magnitude across the globe in even a single timestep. To assess this parameter space, we sample instantaneous snapshots of the 3D aerosol and meteorology fields from a MARC simulation run with present-day aerosol and precursor gas

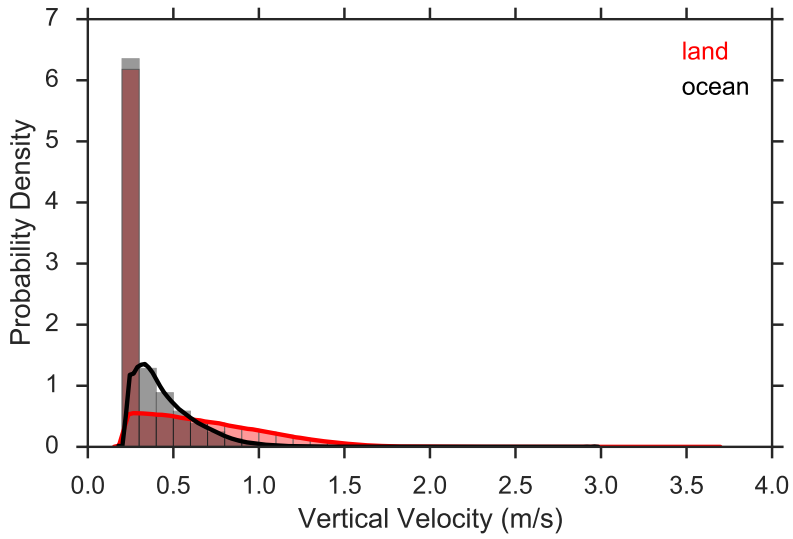

Figure 3-1: Distributions of model-predicted instantaneous sub-grid scale vertical velocity for near-surface (below 700 mb) grid-cells broken down by land (red) and ocean (black) regimes. Bars show a histogram of the distribution; the lines show kernel density estimate.

emissions. In total, 70 timesteps were sampled covering the complete seasonal and diurnal cycle at each model grid cell.

The variability in sub-grid scale vertical velocity as a function of continental versus maritime grid cells across all time samples in this output is summarized in Fig. 3-1. In both regimes, the mode updraft speed falls at the lower bound of $0.2\,\mathrm{m\,s^{-1}}$, and occurs about 50% of the time. These velocities rarely exceed $1\,\mathrm{m\,s^{-1}}$ - 10% and 1% of the time over land and ocean, respectively. On average, land velocities are slightly larger ($0.41\,\mathrm{m\,s^{-1}}$ vs $0.32\,\mathrm{m\,s^{-1}}$), but have higher variance. The distribution of vertical velocities in both regimes has a long positive tail, maxing out between $3\,\mathrm{m\,s^{-1}}$ to $4\,\mathrm{m\,s^{-1}}$ and never approaching the artificial cap imposed by MARC.

The different particle size distributions in MARC are influenced by both different emissions sources and acted upon by different physical processes. This leads to a great deal of spatial heterogeneity in the size distribution parameters. One aspect of this heterogeneity is depicted in Fig. 3-2, which shows distributions of the total number concentration of four of the modes aggregated into latitude and height bins. In general, number concentration for each mode decreases with height, as expected since (with the exception of the pure sulfate modes) all the modes have strong sources near the surface in the model. Natural aerosols

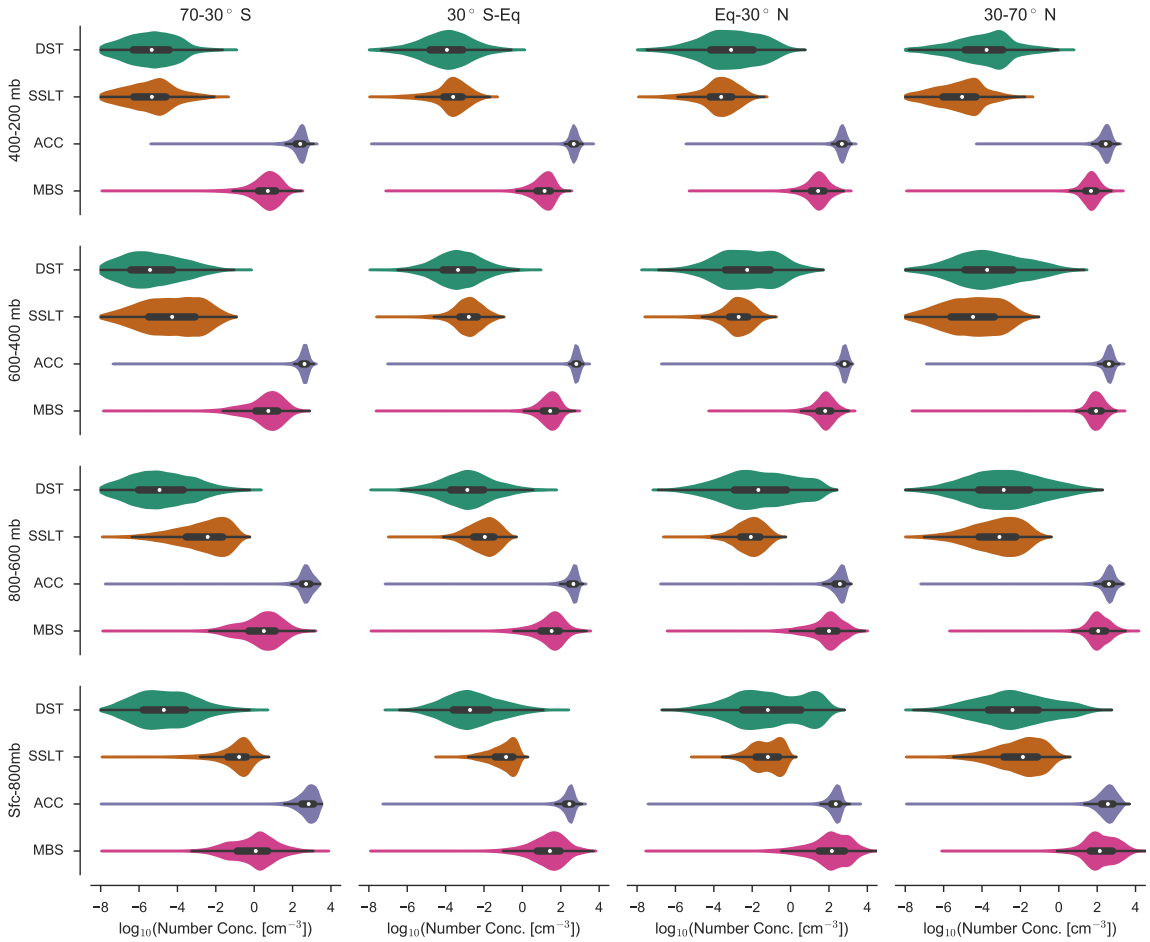

Figure 3-2: Violinplots showing the distribution of the logarithm of the number concentration for total dust (DST), total sea salt (SSLT), accumulation mode (ACC) and mixed sulfate-black-carbon mode (MBS) aggregated by latitude (columns) and vertical level (rows). The width of each violinplot is scaled by the number of observations for each mode in a given aggregation. The inner boxplot on each figure shows the median and interquartile range for reference.

(dust and sea salt) are generally much less abundant than anthropogenic ones. Furthermore, there are generally more aerosol by number in the Northern Hemisphere than in the Southern Hemisphere to the preponderance of anthropogenic emissions sources.

Many of the distributions featured in Fig. 3-2 have long tails extending towards very low number concentrations. These very low values can be problematic for activation parameterizations, especially those constructed from statistical methods or sampling, such as a lookup table, as they necessitate many saved interpolation points. However, aerosol activation produces at most one droplet per aerosol, so sensible lower bounds can be imposed to create a minimum threshold below which little activation is assumed to occur. Furthermore,

particle size is a critical factor in assessing aerosol activation; larger particles have a much lower barrier to activation following Köhler theory (Seinfeld and Pandis, 2006). To simplify the assessment of how activation is influenced by the simulated aerosol size distributions in MARC, we diagnose μ_g from the prognostic moments output by MARC, and study it in lieu of the total mass concentration (M).

Although there is a great deal of variability in the number concentrations simulated by MARC, for all but the nucleation mode sulfate and coarse dust and sea salt modes, those values are most often drawn from a range of just a few orders of magnitude. In general, the number burden in each mode approximately scales with the total aerosol burden. Put another way, the number concentrations in the main sulfate and carbonaceous modes tend to correlate strongly with one another, as does each mode's corresponding mean size. The exception to this rule is typically the sea salt mode, which is abundant in remote areas distant from the primary sources of anthropogenic sulfate and carbonaceous aerosol. The range of geometric mean mode particle sizes tends to be smaller than the range in number concentrations for each corresponding mode.

3.2.2 Reducing the Parameter Space

In total, MARC simulates 15 modes - seven double-moment and eight single-moment. As a result, we require 22 parameters to completely describe the aerosol size distribution.

Two parameters are needed to close the description of its composition; the hygroscopicity and density of the MOS mode evolves in response to its relative mixture of carbon and sulfate, and the MBS mode accrues sulfate mass through aging, which also impacts its particle density. Finally, the ambient temperature and pressure, as well as the vertical velocity of the updraft in which activation is occurring are meteorological parameters which can influence the droplet number nucleated. This yields a total activation parameter space with 27 independent dimensions.

The emulation method used by Rothenberg and Wang (2016) is designed to work with an arbitrary number of input parameters. However, in order to focus our analysis on those parameters most influential on predictions of droplet nucleation, we restrict the input parameter space by eliminating parameters which exert little or no influence on the activation process. For instance, the pure carbonaceous modes (OC and BC) are hydrophobic and

not sources of CCN, so we neglect them in the activation calculation. The nucleation mode sulfate typically has much fewer and much smaller particles than the Aitken mode and accumulation modes present in a grid cell. Furthermore, where nucleation mode particles are abundant in number, the other sulfate modes generally are, too. Thus, we also assume that the nucleation mode is not a source of CCN, and exclude it from our activation calculations. Additionally, the mixed black-carbon-sulfate mode (MBS) is assumed to be composed of particles with a carbon core and sulfate shell; we further assume that the entire surface of these particles are coated, effectively rendering the MBS particle hygroscopicity to be equal to that of sulfate. These assumptions effectively reduce the number of parameters we must consider by seven, bringing the initial parameters to a set of 17 aerosol ones and 3 meteorological ones.

3.2.3 Iterative Activation Calculations

To further reduce this number, we assess the relative importance of each individual aerosol mode and its influence on activation dynamics. This helps to identify a subset of aerosol modes to use as predictors in our emulator, avoiding the need to include all 17 potential aerosol parameters.

We accomplish this with an ensemble of iterative activation calculations using a detailed reference parcel model (Rothenberg and Wang, 2016), drawn from a sample of aerosol size distributions simulated by MARC. This strategy effectively employs a "greedy" algorithm to sort the set of available aerosol modes, ranking their influence on activation by their cumulative depression on the supersaturation maximum achieved for a given parcel ascent.

Given a set of n aerosol modes, the iterative calculations provide a sorted order for the modes, indicating their relative contribution to activation dynamics. Specifically, for each n modes, we pick a test mode and run a parcel model simulation to compute the minimum supersaturation maximum ($S_{\max}$) achieved in an updraft in which that mode is embedded. The mode which produces the minimum $S_{\max}$ is said to be the "dominant" mode, and we record its size distribution. We then remove that aerosol from the original set of n modes. At this point, we re-visit each of the $n-1$ remaining modes, and run parcel simulations in which the first "dominant" mode is present along with one additional mode. Again, we record the minimum $S_{\max}$ and remove the contributing mode from the original set, adding

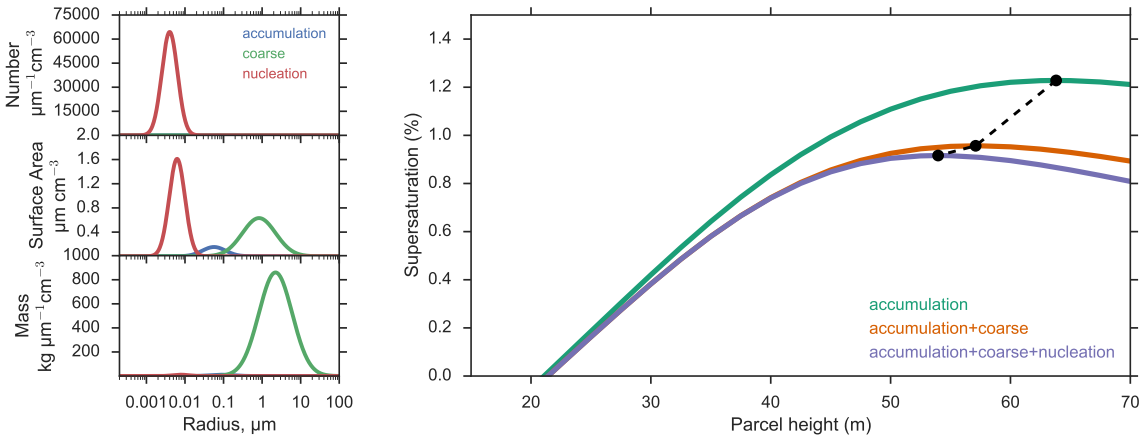

Figure 3-3: Illustration of particle number, surface area, and mass distributions (left) for typical marine size distribution from Whitby (1978), along with iterative activation calculations (right). In all calculations, we assume the aerosol are all sulfate particles with a hygroscopicity of $\kappa = 0.56$, and are activated in constant-speed updraft of $V = 0.5\,\mathrm{m\,s^{-1}}$. Note that the nucleation mode in the Whitby distribution is closer in size to the Aitken mode included in many moder aerosol-climate models

it to the set of dominant modes. The end result of $n-1$ iterations is a complete sorting of the modes, based on their contribution to reducing $S_{\max}$ in the parcel model simulations.

Figure 3-3 illustrates this iterative process for an example marine aerosol distribution following Whitby (1978). In this example, all particles are assumed to be pure sulfate. This size distribution is particularly interesting because by number, the vast majority of particles exist in the smaller, nucleation mode, but the coarse mode dominates the mass distribution. However, the most important factor for influencing the activation process is a mode's number-weighted surface area distribution, since this controls the rate of condensation and latent heat release, and thus the evolution of supersaturation in a cooling parcel. On the surface area distribution of Figure 3-3 we begin to visualize the accumulation mode, but both the nucleation and coarse modes have a larger presence. While the surface area distribution physically controls the activation process, its outcome—a diagnosed cloud droplet number concentration—is closely related to the number distribution, because we assume that droplets must form from individual particles. The interplay here is complex, though; the supersaturation maxima influenced by the surface area distribution sets the threshold which droplets must overcome in order to activate. Thus, it would be reasonable to assume that the nucleation mode particles (which feature prominently in both the number and surface area distributions) would "dominate" activation in this case, or that small changes in

the burden of these particles could have large consequences on how many cloud droplets will form.

However, that doesn't happen. On the right-hand panel of Fig. 3-3 are plotted traces of the supersaturation achieved in a parcel with the indicated aerosol population, as a function of height above the parcel's initial altitude. In the first iteration, the parcel achieved minimum supersaturation maxima of 1.2%, 1.4%, and 1.3% when just the accumulation, coarse, and nucleation modes are present, respectively. These are much higher supersaturations than the Köhler-theory critical supersaturation for the geometric mean particle size in both the accumulation and coarse modes (0.3% and 0.1%, respectively), and a large fraction of the particles in those two modes activate. However, the number concentration of particles in those two modes is very small (60 and 3.1 particles cm^{-3}), and it takes time before the latent heat release due to condensation balances the cooling in the ascending parcels. Although the nucleation mode particles are much smaller and fewer of them activate, their number concentration is much higher ($340\,\mathrm{cm}^{-3}$) and the total liquid water condensed on them is similar through the ascending parcel's trajectory, hence the similar values for $S_{\max}$ in all three cases.

In the second iteration—with the accumulation mode omnipresent—the coarse and nucleation modes only reduce $S_{\max}$ to 0.96% and 1.08%, respectively. The final iteration, which includes all three modes, produces an $S_{\max}$ of 0.92% with an ordering of (accumulation, coarse, nucleation). These reductions in $S_{\max}$ are due to the higher total particle number concentration and surface area available for condensation which leads to a more rapid balancing between the parcel's adiabatic cooling and warming due to latent heat release.

We apply this algorithm to a sample of 50,000 aerosol size distributions and meteorological parameters taken from our reference, present day MARC simulation. For each parameter set, we calculate the first four dominant modes, and record the supersaturation maxima produced by each successive combination. Using these supersaturation maxima, we diagnose the number of activated aerosol across the total aerosol population, including modes which are not present in the parcel model simulation at a given iteration. Additionally, for each parameter set we perform a reference parcel model calculation where all aerosol modes are included, for comparison with the iterative calculations.

Overall, the accumulation mode sulfate (ACC) is the dominant mode in 96.5% of the

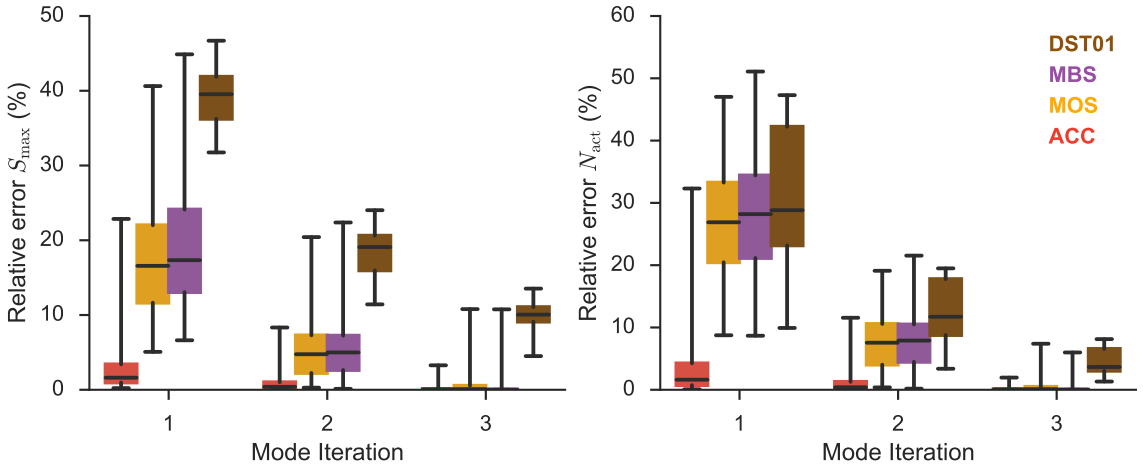

Figure 3-4: Relative errors in S_{max} (left) and N_{act} (right) in subsequent iterations of the greedy activation calculations. The coloring of each box indicates which mode was the first or dominant one. In each boxplot, box encompasses the interquartile range and the whiskers extend to the 1st and 99th percentiles in the corresponding sub-sample. Outliers beyond these percentiles are not plotted.

sample cases. Infrequently, the mixed-sulfate-carbon modes and smallest dust mode are the dominant mode, accounting for all of the remaining cases. When ACC dominates the activation dynamics, either the mixed-modes (MBS and MOS) or smallest sea salt mode is the second-most dominant one (in 10.3%, 36.2%, and 52.8% of cases, respectively). In fact, this ordering is so common that in 85% of cases, three of ACC, MOS, MBS, or SSLT01 comprise the top three dominant modes.

Figure 3-4 illustrates the potential error in calculating both S_{max} and N_{act} for each iteration of the activation calculations relative to using the full aerosol population, aggregated by which mode was the first-most dominant one. In all cases, using a subset of the modes tends to over-predict the droplet number activation as a consequence of predicting a high value for S_{max}. This is consistent with the physics of the activation problem; the presence of more aerosol surface area on which condensation can occur tends to produce a greater source of latent heat release to counter-balance adiabatic cooling in the ascending parcel, suppressing the development of a higher S_{max}. But as Fig. 3-4 shows, this over-prediction decreases rapidly as modes are included in the calculation. Part of this decrease is related to the fact that adding modes in each iteration captures a higher fraction of the total aerosol number; on average, the first dominant mode contains $70\% \pm 27\%$ of the total aerosol number, which increases to $80\% \pm 17\%$ and $89\% \pm 13\%$ after adding the second and third modes.

Following this increase in fraction of the total aerosol number included by the dominant mode set in each iteration is a decrease in the absolute error in N_{act} relative to the full aerosol population, with an average of less than $1\,\text{cm}^{-3}$ and max of $57\,\text{cm}^{-3}$ by the third iteration.

With respect to the goal of reducing the aerosol parameter space necessary for assessing aerosol activation, what's more important in these calculations than the frequently dominant modes, though, is the the absence of several modes altogether. In particular, the Aitken mode sulfate is never one of the first three dominant modes; beyond the modes depicted in Fig. 3-4, the only other modes in that set are the larger dust and sea salt modes. But in nearly all the cases sampled here, those modes have number concentrations that decrease as the bin's particle size increases. Also, the smallest dust and sea salt bins tend not to feature very high number concentrations in the first place. Even though the larger bins contribute particles which will almost always activate, owing to their size and composition, the particles in them simply aren't numerous enough to impact the activation calculations.

3.3 Emulator Development

We now seek to construct an emulator of a detailed adiabatic cloud parcel model capable of resolving droplet activation within an ascending, constant-speed updraft. The following sections briefly re-introduce the chosen cloud parcel model and emulation technique, polynomial chaos expansion; more information can be found in Chapter 2 of this dissertation.

3.3.1 Parcel Model

Adiabatic cloud parcel models are a standard modeling tool for detailed assessments of aerosol activation and other studies focused on the composition of atmospheric particulates (Seinfeld and Pandis, 2006). In such a model, a constant-speed updraft drives adiabatic cooling in a closed, zero-dimensional air parcel within which are any number or configuration of aerosol particles. Initially prescribed a temperature, pressure, and water vapor content, the cooling parcel eventually develops a supersaturation with respect to water vapor. In a sufficiently supersaturated environment, water vapor condenses on particulate surfaces. However, condensation releases latent heat, which counter-balances the parcel's adiabatic cooling. This balance can be expressed

$$\frac{dS}{dt} = \alpha(T,P)V - \gamma(T,P)\frac{dw_c}{dt} \tag{3.1}$$

where V is the updraft speed, $\alpha(T,P) = (gM_wL/c_pRT^2) - (gM_a/RT)$ and $\gamma(T,P) = (PM_a/e_sM_w)+(M_wL^2/c_pRT^2)$ are functions weakly dependent on temperature and pressure (Leaitch et al., 1986), M_w and M_a are the molecular weights of water and air, L is the latent heat of vaporization of water, c_p is the specific heat of dry air at constant pressure, R is the universal gas constant, g is the acceleration due to gravity, e_s is the saturation vapor pressure, and w_c is the liquid cloud water mass mixing ratio (please refer to Appendix A of Chapter 2 for more details).

At some time t, the balance between heating due to latent heat release and cooling due to the parcel's adiabatic ascent will approximately balance such that $\frac{dS}{dt} = 0$ and a supersaturation maximum, $S_{\max}$, will occur. Thereafter, S generally decreases, relaxing to some value close to unity as condensation drives droplet growth, quenching the ambient water vapor surplus. Beyond this point, the aerosol bifurcates into two populations: protocloud droplets which will continue to grow due to condensation and eventually collision and coalescence, and interstitial haze particles which may have become hydrated, but upon which further condensation is not thermodynamically favorable.

In order to compute these effects, we simulate an arbitrary number of initial dry particle size distributions following a lognormal assumption:

$$n_N(r) = \frac{dN}{d\ln r} = \frac{N_t}{\sqrt{2\pi}\ln\sigma_g}\exp\left[-\frac{ln^2r/\mu_g}{2\ln^2\sigma_g}\right] \tag{3.2}$$

Each aerosol mode is thus defined uniquely by three parameters (N_t, μ_g, σ_g) corresponding to the total aerosol number concentration, the geometric mean particle radius, and the geometric standard deviation. The modes are each further discretized into a Lagrangian grid of 200 size bins, equally spaced over the logarithm of the particle radius such that each bin represents a different number of particles. The particles are then hydrated to an equilibrium size with respect to the initial relative humidity in the model; condensation acts to grow this size in response to the thermodynamic evolution of the parcel.

Furthermore, each bin is assigned a fixed hygroscopicity following κ-Köhler theory (Petters and Kreidenweis, 2007). By tracking the dry particle size, a hygroscopicity, and the wet radius of each particle, aerosol activation can be directly assessed within the model.

3.3.2 Polynomial Chaos Expansion

We emulate the behavior of the detailed parcel model by applying the probabilistic collocation method (PCM; Tatang et al., 1997). PCM is a method of polynomial chaos expansion which seeks to construct a model response surface by mapping input parameters related to the initial conditions and behavior of a model to some response measured from the model. This process yields a computationally efficient yet accurate reproduction of the model.

The PCM is a non-intrusive technique which does not require modifications to an existing model in order to be applied. Instead, the PCM treats the original, full-complexity model as a black box and the chosen set of M input parameters as independent, random variables, $\mathbf{X} = X_1, \ldots, X_M$, each with an associated probability density function. This PDF is used as a weighting function to derive a family of orthogonal polynomials which are used as the bases for the polynomial chaos expansion to be constructed, ϕ. Using a finite number of these bases and choosing some model response, R, we write the polynomial chaos expansion as

$$R \approx \sum_{j=0}^{P} \alpha_j \phi_j(\mathbf{X}) \tag{3.3}$$

Such an expression has $N_t = P + 1 = (M + p)!/(M!p!)$ total terms, since a given chaos expansion of order p will contain $p + 1$ basis terms for each input parameter and combinations thereof. The coefficients α_j are computed by evaluating the original model at a set of particular set of sample points, recording the response of the model, and solving a regression problem. Those sample points are generated by taking the roots of the orthogonal polynomials associated with each of the input parameters and their random variables.

In order to compute the polynomial chaos expansions, we use the Design Analysis Kit for Optimization and Terascale Applications (DAKOTA; Adams et al., 2014), version 6.1. This software automates the process of generating input parameter sets, sampling the full-complexity model to be emulated, and constructing the polynomial chaos expansion. Furthermore, it provides many useful statistical properties of the sample dataset and the chaos expansions themselves.

Symbol	Description	Lower Bound	Upper Bound
logN_ACC	Log of accumulation mode sulfate number concentration (cm^{-3})	-3 (1.2)	4 (100)
logN_MOS	Log of mixed sulfate-organic carbon number concentration (cm^{-3})	-5 (1.5)	4 (99.9)
logN_MBS	Log of mixed sulfate-black carbon number concentration (cm^{-3})	-5 (1.6)	4 (99.8)
logN_DST01*	Log of 0.16 micron dust particle number concentration (cm^{-3})	-5 (18.2)	2 (99.8)
logN_DST02*	Log of 0.4 micron dust particle number concentration (cm^{-3})	-5 (38.9)	1 (99.9)
logN_SSLT01*	Log of 0.5 micron sea salt particle number concentration (cm^{-3})	-5 (3.6)	1 (100)
logmu_ACC	Geometric mean size of accumulation mode (micron)	-3 (0.1)	0 (98.9)
logmu_MOS	Geometric mean size of mixed sulfate-organic carbon mode (micron)	-3 (0.06)	-1 (98.3)
logmu_MBS	Geometric mean size of mixed sulfate-black carbon mode (micron)	-3 (0.1)	-1 (98.5)
kappa_MOS	Hygroscopicity of mixed sulfate-organic carbon mode	0.1	0.6
log_V	Log of updraft velocity (m/s)	-2	1
T	Temperature (K)	240	310
P	Pressure (Pa)	50000	105000

Table 3.2: Input parameter space and bounds on associated uniform probability density functions used to derive polynomial chaos expansions for MARC activation. For the lower and upper bounds on the aerosol size distribution parameters, the parenthetical values denote the percentile of the distribution for that parameter at which the bound occurs. All terms are present for the main expansion; terms affixed with an (*) are added for the gCCN expansion.

3.3.3 Emulation of aerosol activation for MARC

We now apply the parcel model and chaos expansion technique in the previous sections to construct emulators of aerosol activation suitable for use in MARC. Following the analysis in Sect. 3.2, we identify a reduced-dimensionality input parameter space which covers the diverse set of aerosol and meteorology scenarios in which activation occurs in MARC, summarized by Table 3.2. Following the iterative calculations, we restrict the aerosol modes included in the activation calculation to just the accumulation mode sulfate and both mixed sulfate-organic carbon and sulfate-black carbon particles, as well as giant CCN particles from the dust and sea salt modes. To assess the importance of these less-abundant, coarse particles, we derive two emulators: a "main" scheme which includes just the ACC, MBS, and MOS modes, and a "gCCN" scheme which adds in the dust and sea salt modes. All the aerosol size distribution parameters are transformed using a logarithm, since they can take on values that span several orders of magnitude. Additionally, we consider the hygroscopicity of the mixed sulfate-organic carbon mode, as well as the updraft speed and ambient temperature and pressure in our input parameter space.

For each of these parameters we construct uniform size distributions, which are uniquely defined by a set of low and high bounds. These bounds are also noted in Table 3.2, along

with the percentile of the data they correspond to from our sampling study of the parameter space for the aerosol parameters. For most of those parameters, the bounds cover upwards of 99% of the sampled parameter space. Notably, the lower cut-off boundaries for dust and sea salt number concentration occur at much higher percentiles; the number concentration of droplets nucleated in our calculations is relatively insensitive to changes in the coarse mode number concentration at very low values, though, so we opt to constrain the input parameter space to these modes to a more physically relevant range. The MOS hygroscopicity, temperature, pressure, and updraft velocity ranges cover all plausible values that could be used in an online activation calculation within MARC.

We cast all of the input parameters as random variables with uniform probability density functions to emphasize that we care equally about computing activation with a parameter set drawn with any parameter values. However, in MARC, some combinations of parameters are extremely unlikely. For instance, sea salt has sources far remote, spatially, from the sources of black carbon; ergo, it is uncommon to see a high number concentration for sea salt as well as the mixed sulfate-black carbon mode. The trade-off here potentially lies in emulator performance, because the PCM will attempt to train the emulator to perform well for input parameters that we've assigned equal likelihood to, but may actually be far less likely to occur. A major benefit from this trade-off is simplicity in matriculating the emulator for use in a GCM. Using multiple probability density functions for different input variables, while offering a tuning knob to increase the accuracy of the emulators, would require a larger set of orthogonal polynomial bases than just the Legendre polynomials used for the uniform distributions here.

These parameters are used to drive parcel model simulations where we record the logarithm of $S_{\max}$ as the response variable. This value can then be used to diagnose the number concentration of droplets nucleated by assuming that any particles which experience their Köhler theory-predicted critical supersaturation. We note that although this does not resolve the issue of kinetic limitations on droplet growth and its potential to cause an under-prediction in droplet number (Nenes et al., 2001), unlike existing activation schemes, our emulator accounts for the feedback of these effects on $S_{\max}$, so one avenue of its impact is lessened. Furthermore, emulators predicting $S_{\max}$ were much more accurate at reproducing parcel model behavior than those which directly predicted estimates of N_{act} accounting for kinetic limitations on growth.

The end result of constructing the emulators is a function which maps $\log_{10}(S_{\max})$ to a set of values from our input parameter space,

$$\begin{aligned}\log_{10}(S_{\max}) = f(&\log_{10} N_{\mathrm{ACC}}, \log_{10} N_{\mathrm{MOS}}, \log_{10} N_{\mathrm{MBS}}, \\ &\log_{10} \mu_{\mathrm{ACC}}, \log_{10} \mu_{\mathrm{MOS}}, \log_{10} \mu_{\mathrm{MBS}}, \kappa_{\mathrm{MOS}}, \log_{10} V, P, T \\ &[, \log_{10} N_{\mathrm{DST01}}, \log_{10} N_{\mathrm{DST02}}, \log_{10} N_{\mathrm{SSLT01}}])\end{aligned} \tag{3.4}$$

From a prediction of the $S_{\max}$ achieved in an ascending parcel with the conditions passed to the emulator, we can then diagnose aerosol activation by re-writing the lognormal size distribution for each mode as a function of critical supersaturation (Ghan et al., 2011) to yield an expression

$$N_{\mathrm{act}} = \sum_{i=1}^{n} \frac{N_{t,i}}{2} \left(1 - \mathrm{erf}\left[2 \ln\left(\frac{S_{m,i}}{S_{\max}}\right) / (3\sqrt{2} \ln \sigma_{g,i})\right]\right) \tag{3.5}$$

where $S_{m,i}$ is the critical supersaturation for a particle of radius $\mu_{g,i}$ from mode i.

3.4 Evaluation of Emulators

We evaluate our emulators by applying them to both a synthetic sample of potential input parameters as well as real samples taken from a MARC simulation. In all of our comparisons, we study third and fourth order chaos expansions both excluding ("main") and including ("gCCN") the coarse dust and sea salt modes.

As a reference, we compute activation statistics for each sample from several different sources. First, we run the detailed parcel which the emulator aims to simulate. Second, as a further benchmark and comparison, we run two widely-used activation parameterizations from the literature. The first scheme, by Abdul-Razzak and Ghan (2000) (ARG) uses a pseudo-analytical solution to an integro-differential equation derived from the original adiabatic parcel model system. However, one part of the pseudo-analytical calculation involves a fit to parcel model calculations. The second parameterization, by Morales Betancourt and Nenes (2014a) (MBN), applies an iterative scheme to partition the aerosol population into two subsets, and uses different limits on the underlying analytical formulas to derive a maximum supersaturation. Because it requires a sequence of iterations to run, the MBN

scheme is more computationally expensive than the ARG scheme, but has the potential to include more detailed links between particle composition and condensation (Kumar et al., 2009) or entrainment into the parcel (Barahona and Nenes, 2007). Like the ARG scheme, though, one limiting case in the MBN scheme relies on a fit to parcel model simulations. In both cases, those simulations involved models conceptually similar to the one emulated here.

3.4.1 Input Parameter Space Sampling

Using the parameter space defined in Table 3.2, $n = 10000$ sample parameter sets were drawn using maximin Latin Hypercube Sampling (LHS). This randomized sampling method helps to ensure that the full aerosol and meteorology parameter space is studied while assessing its performance.

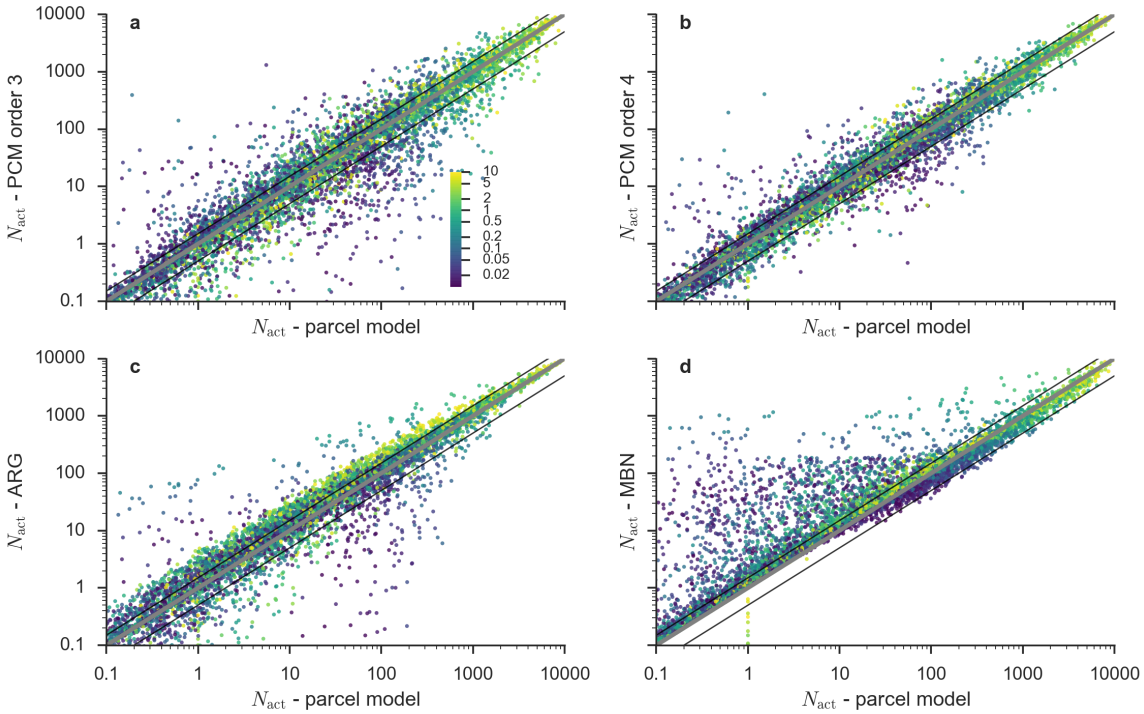

Figure 3-6: Like Figure 3-5, but plotting the predicted droplet number concentration activated for the aerosol "main" parameter set

Figure 3-5 compares the performance of each emulator and the two reference activation scheme against parcel model simulations using all of the LHS samples for the "main" aerosol parameter sets. In the simulations, higher updraft speeds (shaded) are nearly always asso-

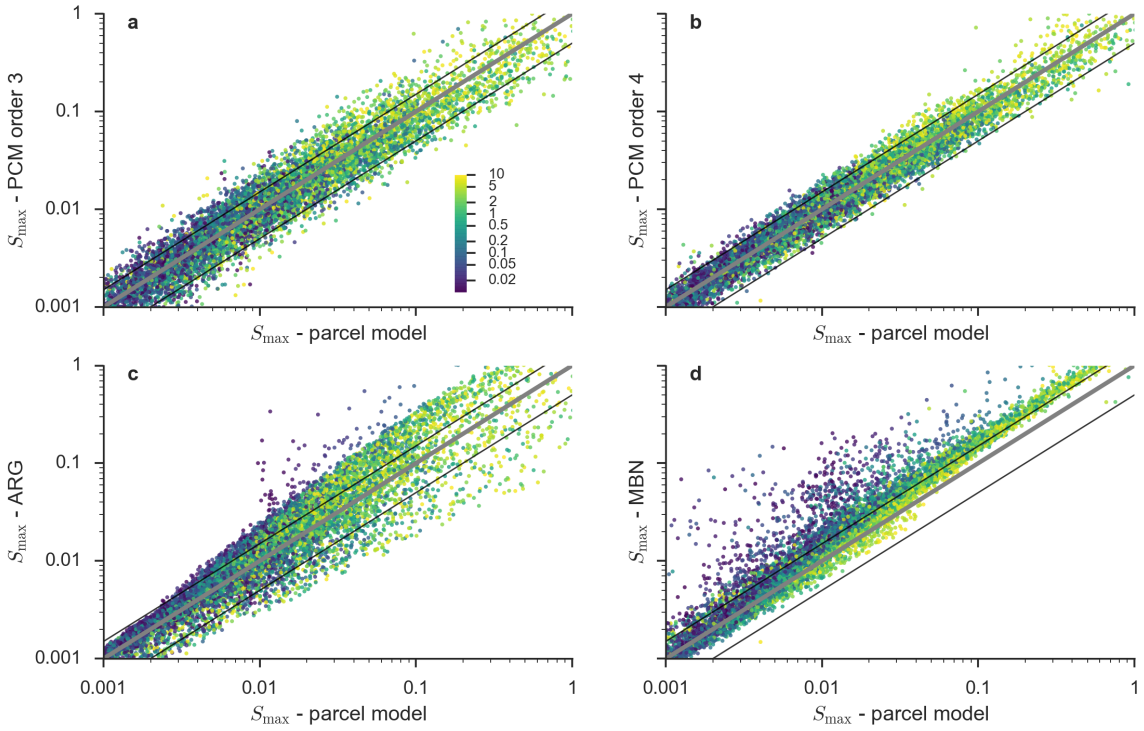

Figure 3-5: One-one plot comparing predicted supersaturation maxima between parcel model and activation parameterizations - 3rd-order emulator (a), 4th-order emulator (b), ARG (c) and MBN (d). The "main" aerosol parameter set (excluding the dust and sea salt as predictor modes) were utilized here. Glyphs are shaded to denote updraft velocity corresponding to each sample draw (in m/s), and are consistent for each panel. Solid black lines denote a factor-of-2 difference between predicted values from parcel model and corresponding parameterization evaluations

ciated with a much higher supersaturation maximum. For the emulators, accuracy tends to increase on average going from the 3rd order (Fig. 3-5-a) to the 4th order (Fig. 3-5-b) scheme, although there is slightly higher variance in the relative error compared to the parcel model at higher updraft speeds. With respect to the driving updraft speed, though, there isn't a consistent mode of bias - on average, the relative error is very low. The same does not hold true for the two reference schemes. The ARG scheme (Fig. 3-5-c) tends to predict both too-high and too-low supersaturation maxima at higher updraft speeds but is relatively well-calibrated at lower updraft speeds yielding a lower supersaturation maximum. On the other hand, the MBN scheme (Fig. 3-5-d) is generally more accurate and better-calibrated than either of the emulators or the ARG scheme, especially at higher updraft speeds - but tends to spuriously over-estimate S_{max} for weak updraft speeds.

Figure 3-7 takes the results depicted in Fig. 3-5 one step further by diagnosing droplet

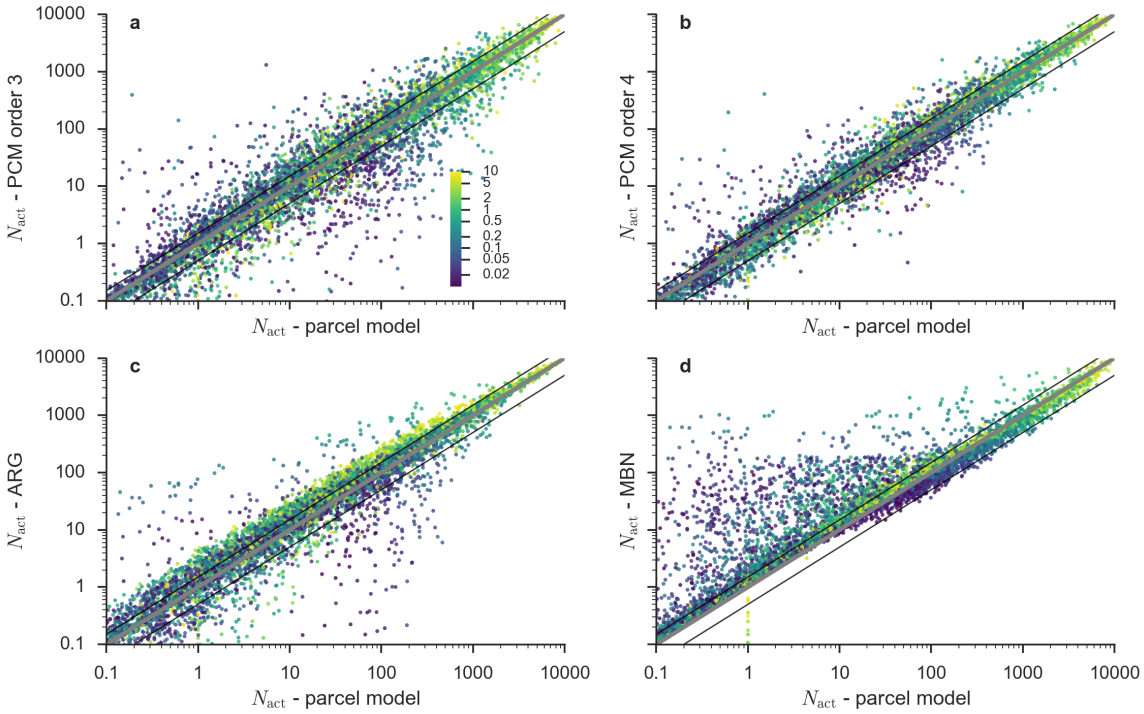

Figure 3-7: Like Figure 3-5, but plotting the predicted droplet number concentration activated for the aerosol "main" parameter set

number concentrations nucleated from each $S_{\max}$. For all the schemes, there can be substantial differences between the parcel model and each parameterization. This is particularly the case in regimes which give rise to fewer overall droplet number concentrations, either due to a lower driving updraft velocity, or a lower total aerosol number available to activate. Surprisingly, the MBN scheme tends to consistently activate a higher number of droplets with respect to the parcel model, especially in situations which should have very few droplets - below $10\,\mathrm{cm}^{-3}$. The ARG scheme does not have as consistent of a bias, but can both egregiously over-predict and under-predict droplet number, with these biases exaggerated at lower updraft speeds. By comparison, the emulators show much less overall bias. The mean error for the emulators follows that of $S_{\max}$ and is small, but there is variance which tends to impart a small low or high bias on its estimates.

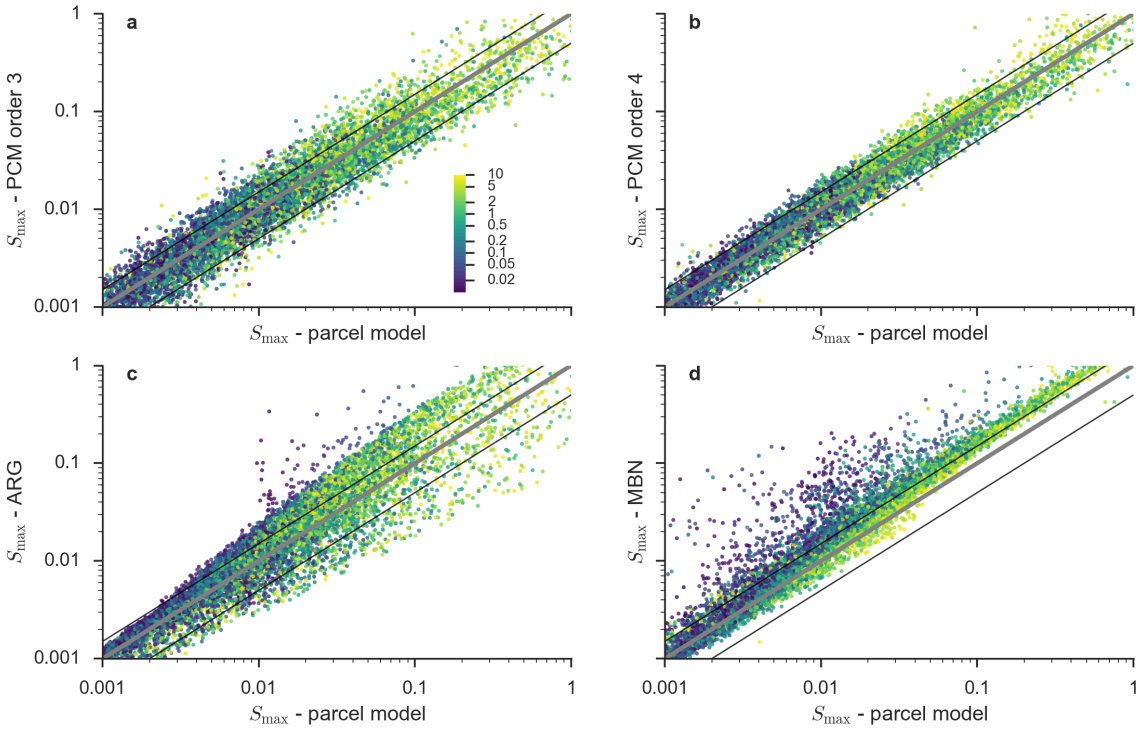

Figure 3-8: Like Figure 3-5, but for the "gCCN" parameter set

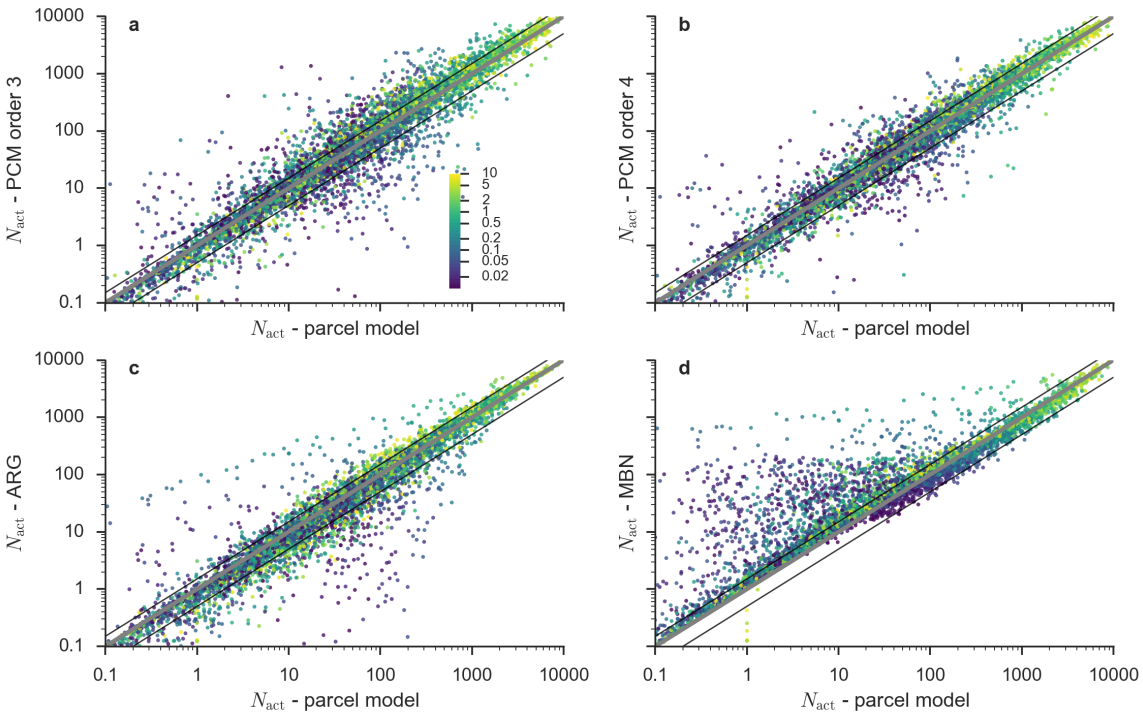

Figure 3-9: Like Figure 3-7, but for the "gCCN" parameter set

Both of these sets of plots are repeated in Figs. 3-8 and 3-9, but for the "gCCN" exper-

iment. Qualitatively, the results for all parameterizations are very similar, with the same overall biases - especially for the ARG and MBN parameterizations. The emulators tend not to perform as well overall in the "gCCN" cases, although they are still the best-calibrated scheme and do not have the velocity-regime errors that the MBN scheme has. In both the "main" and "gCCN" parameter sets, the MBN scheme tends to more regularly predict too many cloud droplets, save for polluted regimes where $N_{\rm act} < 100\,{\rm cm}^{-3}$, as the bias reverses and the scheme has a tendency to under-predict droplet number. Neither the emulators nor the ARG show this same tendency in bias.

These differences in bias are most likely related to the choice of parcel model used in testing and building the ARG and MBN schemes; because each scheme relies on some empirical tuning to parcel model calculations, details in the implementation of each parcel model which influence its sensitivity should show up in ensemble evaluations of each activation scheme. The "gCCN" case is more taxing to simulate with parcel models using a Lagrangian description of the particle size distribution, because condensational growth is computed for each particle bin simultaneously. The stiffness ratio of the underlying differential equations describing condensation in this case will be extremely large, as the small particles in the main aerosol modes will grow much more slowly than those in the giant CCN modes. Although modern ODE solvers can automatically handle these scenarios, the subjective choice of which particular solver and how to discretize the giant CCN population (how many bins per mode) could influence the sensitivity of $S_{\rm max}$ to changes in the model inputs and account for the differences observed here.

To better summarize the results in Figs. 3-5 and 3-7 to 3-9, summary statistics on the error of each scheme versus their corresponding parcel model calculations are shown in Table 3.3. In both sampling cases, all of the parameterizations show a strong linear correlation (r^2) between their predictions and the result of the parcel model. The emulators (PCM Order p) predict $S_{\rm max}$ with lower overall absolute and relative error, but with a much higher variance (not shown here). However, that lower error does not always translate into the emulators being the most accurate absolute predictors of $N_{\rm act}$. For the "gCCN" parameters, the ARG scheme predicts $N_{\rm act}$ with a lower average mean relative error. In both parameter sets, the MBN scheme is the least accurate compared to the parcel model used in these sampling calculations.

		$\log_{10} S_{\max}$				N_{act}			
exp	scheme	MAE	MRE	NRMSE	r^2	MAE	MRE	NRMSE	r^2
main	ARG	0.18	-3.26	0.10	0.94	40.14	25.39	0.15	0.98
	MBN	0.20	-11.79	0.18	0.81	59.05	44.95	0.30	0.90
	PCM Order 3	0.16	0.59	0.09	0.95	72.54	9.20	0.31	0.90
	PCM Order 4	0.10	-0.60	0.06	0.98	45.47	8.89	0.19	0.96
gCCN	ARG	0.17	8.54	0.09	0.93	37.41	-3.92	0.15	0.98
	MBN	0.20	-9.58	0.17	0.78	56.03	33.30	0.31	0.89
	PCM Order 3	0.16	0.59	0.08	0.95	81.19	15.14	0.34	0.87
	PCM Order 4	0.10	0.36	0.06	0.98	50.99	6.90	0.23	0.94

Table 3.3: Summary statistics for error in supersaturation maxima and droplet number activated predicted by emulators and activation parameterization relative to corresponding simulations with a detailed parcel model. From left-to-right, each column represents the coefficient of determination (r^2), mean absolute error (MAE), mean relative error (MRE), and the normalized root-mean-square error (NRMSE)

3.4.2 MARC Aerosol Sampling

Although the sampling in the previous section fully exhausts the input parameter space over which aerosol activation may need to be assessed, it undoubtedly samples from aerosol and meteorological conditions which may not be likely to occur in the real world. To better understand the performance and potential bias of the emulators developed here and the existing activation schemes, then, we also studied a sample of $n = 10000$ aerosol and meteorology parameter sets drawn directly from a MARC simulation. All of the schemes were evaluated again using these parameter sets and the detailed parcel model. This includes the "main" and "gCCN" emulators, which allows us to identify the importance of including the dust and sea salt modes as predictors in the chaos expansions. The parameters in these sets occasionally include values outside the ranges defined in Table 3.2 and studied in the previous section. These cases are more frequently associated with very low total aerosol number concentration, especially over the ocean where anthropogenic aerosols are limited and natural aerosols - which have a lower overall number burden - dominate. Because the aerosol samples from oceanic and continental grid cells differ in this fundamental way, we break down the following analysis to reflect those differences. As in the previous sampling experiment, summary statistics on the performance of each emulator, alongside the ARG and MBN schemes, are detailed in Table 3.4.

Qualitatively, all of the activation schemes perform similarly when evaluated against

	$\log_{10} S_{\max}$				N_{act}			
scheme	MAE	MRE	NRMSE	r^2	MAE	MRE	NRMSE	r^2
ARG	0.05	-0.16	0.03	0.92	25.5	2.87	0.16	0.94
MBN	0.06	0.05	0.05	0.71	26.7	-6.68	0.19	0.93
main Order 3	0.12	0.42	0.08	0.33	64.7	-1.81	0.44	0.59
main Order 4	0.04	-0.31	0.02	0.96	24	4.59	0.19	0.93
gCCN Order 3	0.14	-1.84	0.09	0.18	75.6	20.9	0.47	0.52
gCCN Order 4	0.12	4.61	0.10	-0.19	44.3	-19.9	0.33	0.76

Table 3.4: Same as Table 3.3, but for the sampling study using MARC aerosol and meteorology parameter sets.

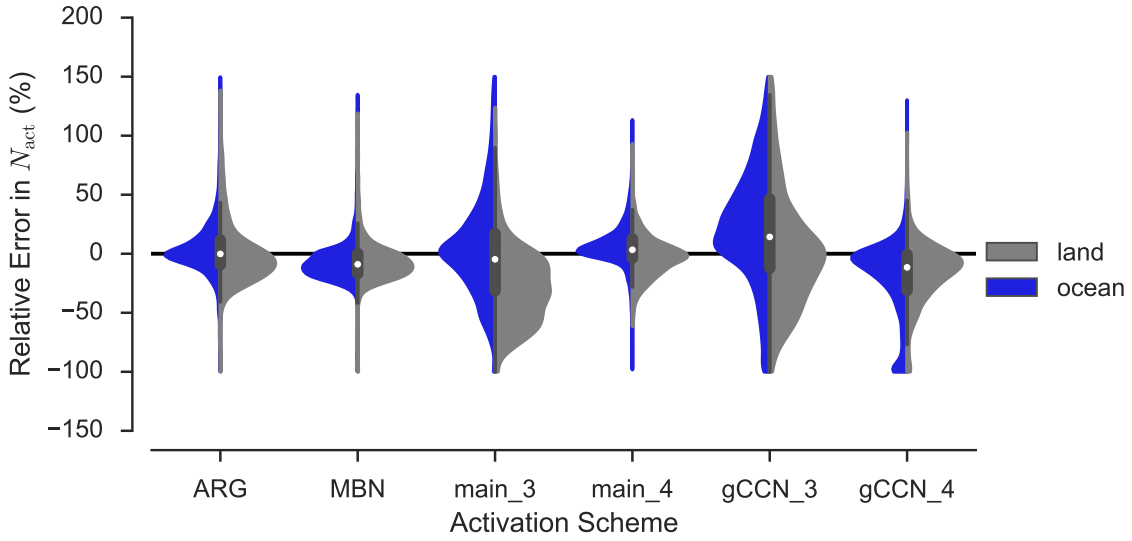

Figure 3-10: Distributions of relative error in scheme prediction of N_{act} versus detailed parcel, evaluated using samples taken from instantaneous MARC aerosol size distribution and meteorology and colored by geographical regime. The long tail of each distribution is clipped at the extrema for each scheme. The box plot in the center of each distribution shows the median and inter-quartile range of the total distribution of both land and ocean samples for each scheme.

the MARC parameters as compared to the more generic sampling in the previous section. Figure 3-10 summarizes distributions of relative error in N_{act} over land and ocean for each scheme. Neither the ARG nor the MBN scheme show much difference in error for the two regimes, although on average, the MBN tends to under-predict N_{act}. This under-prediction usually occurs in regimes with higher updraft speeds and thus higher overall droplet number concentrations. In conditions with weaker updraft speeds, the MBN scheme instead tends to slightly over-predict N_{act}. The ARG scheme is particularly well-calibrated in both regimes.

The emulators derived here do not perform as well as the physically-based parameterizations. Both 3rd order schemes tend to over-predict droplet number over oceans, and under-predict it over land, but with an extremely large variance extending to ±100%. However, including the effects of giant CCN measurably improves the performance of the 3rd order emulators in oceanic regimes. Increasing the order of the emulator also has a significant impact on their accuracy; the 4th order scheme which neglects giant CCN actually out-performs the ARG and MBN scheme on average, and shows little bias between land and ocean regimes, indicating good convergence with its parent parcel model. On the other hand, the gCCN scheme has not yet converged by including 4th order terms, even while its mean error statistics improve. Particularly troublesome is a secondary mode of extreme (over 100%) under-prediction of droplet number of oceanic regimes, but this metric is deceptive. Really, what is occurring is that for very low total aerosol number concentrations - with particle number in the single-digits per cubic centimeter - the 4th order "gCCN" scheme tends to predict half as many droplets as parcel model calculations indicate should form. This typically occurs when one or more of the input size distribution parameters (in particular, the number concentration) for the natural aerosol dips below the minimum threshold where the emulator was trained. When the emulators encounter inputs greater (lower) than these thresholds, they hold them to the maximum (minimum) value in its training range. This follows the assumption that the bounds for each parameter cover the entire range over which activation is sensitive to changes in that input. Put another way, activation should be relatively insensitive to changes near the maximum or minimum values in the range for each parameter. With respect to number concentration, this must be the case; populations with fewer than 10^{-3} particles cm^{-3} offer very little surface area for condensation, and simply cannot exert a strong influence over the developing supersaturation in the parcel. That the 4th order "gCCN" emulator produces too high of sensitivity in this regime suggests that statistical over-fitting is occurring near the extremes of the input parameter space.

To contextualize these differences in N_{act} bias over different geographical regimes, Fig. 3-11 re-maps the testing samples back to the original MARC grid. Here, the difference in regional biases becomes much clearer. Virtually everywhere, the MBN scheme is biased a little low, but there is no systematic difference in this bias between land or ocean, or by geographical areas. The ARG scheme and the 4th order "main" scheme show a different pattern; the ocean-land contrast in bias is clearly visible in the northern hemisphere. Furthermore, the

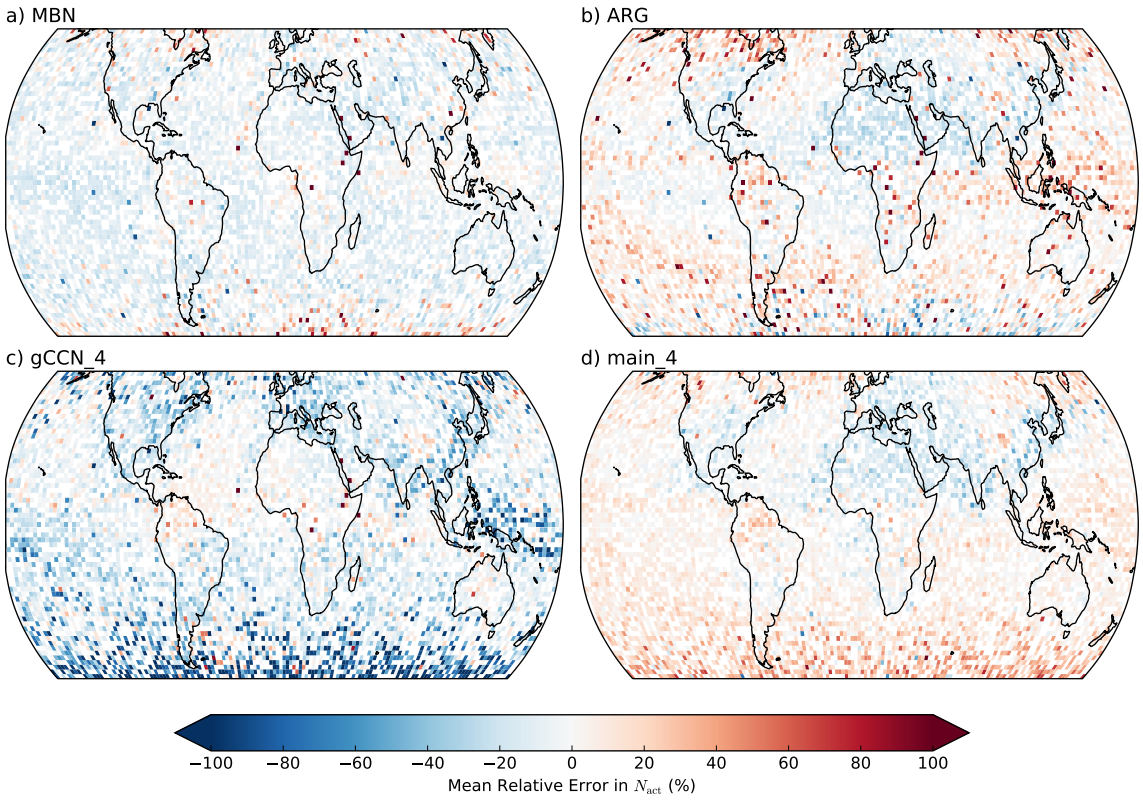

Figure 3-11: Mean relative error in scheme prediction of N_{act} versus detailed parcel model plotted against location on globe where those samples originated. At each grid location, all samples across timesteps and vertical levels (below 700 mb) are averaged together to compute the mean.

bias is typically positive over maritime regions, but negative over regions with anthropogenic aerosol influence. In particular, these regions include Europe and Southeastern Asia, where aerosol distributions are dominated by anthropogenic sulfate and black carbon, and over north central Africa, where the aerosol is a mixture of both dust and organic carbon emissions from biomass burning. In the zonal average, the main_4 scheme is virtually identical to the ARG scheme. However, both cases as well as with MBN, there are larger biases over the southern parts of the oceans, where the aerosol is predominantly comprised of sea salt and smaller sulfate particles produced indirectly through the emission of DMS.

Figure 3-11 also illustrates the poor performance of the gCCN_4 scheme, which under-predicts N_{act} nearly everywhere, but especially so in the southern portions of the ocean basins. The consistent under-prediction in this region explains the bimodal distribution over the ocean hinted at in Fig. 3-10. The gCCN_4 scheme does not perform too dissimilarly than the other schemes over regions with anthropogenic pollution or with mostly dust aerosol.

3.5 Discussion and Conclusions

In the present work, we extended the meta-modeling technique of Rothenberg and Wang (2016, Chapter 2 of this dissertation) in order to apply it to assess droplet nucleation from a complex, multi-modal aerosol mixture simulated by a modern aerosol-climate model. Simultaneously, we characterize the performance of both our new emulators for aerosol activation and two widely-used schemes from the literature, focusing on that same high-dimensional, complex aerosol parameter space. To identify the most important factors impacting activation in that complex parameter space, we apply a physically-based approach to assess the sensitivity of activation statistics to the composition of the aerosol size distribution. Finally, we explore contrasts between aerosol and meteorology regimes over land and ocean, noting the potential for different biases in assessed cloud droplet number depending on the choice of activation scheme used in a particular global modeling application.

In ensembles of iterative calculations using a large sample of aerosol size distributions from a coupled aerosol-climate model, we note that typically, a single mode tends to dominate activation or otherwise strongly predict the total number of droplets nucleated. This approach to understanding the sensitivities of activation dynamics on the underlying aerosol population is distinct from previously-published approaches in the literature. For instance, Karydis et al. (2012) and Morales Betancourt and Nenes (2014b) apply an adjoint approach to derive the sensitivity of aerosol activation to perturbations in input parameters supplied to activation schemes. Detailed calculations using this approach yield a map of local sensitivities or gradients in the relationship between, for example, N_{act} and one input parameter while holding all others constant, and are thus difficult to interpret. The iterative calculations performed here aim instead to address the global sensitivity of activation to configurations of an aerosol population.

It is somewhat surprising that the accumulation mode sulfate (ACC) successfully serves as such a strong proxy for the full aerosol population, even in the presence of giant CCN and a wide swath myriad updraft regimes. This result is likely model-dependent; the ACC mode in MARC is not only ubiquitous, but may be inadvertently (and subjectively) in a range of mean particle sizes for which aerosol activation is especially sensitive. At the same time, the coarse dust and sea salt modes in MARC, on average, hold too small a number concentration to dramatically impact activation calculations save for maritime regions far

remote from anthropogenic sources. However, the presence of sea salt as one of the modes most frequently ranked in the top three influencers of activation points to previous results which indicate the presence of giant CCN influence activation dynamics (e.g. Barahona et al., 2010).

The fact that a single mode can place such a strong constraint on aerosol activation is useful for attempts seeking to extend look-up table methods for building parameterizations. If two modes—an accumulation-size and a coarse-size—accurately predict aerosol activation, then one can constrain the look-up table to just a few key aerosol size distribution parameters. The inclusion of variable aerosol composition would still likely make employing a look-up table in a global model unwieldy, though, necessitating more sophisticated approaches. The emulation technique applied here is one such approach to tackling this problem which appears to work very well.

When sampling against the full training parameter space employed here, our emulators perform capably. Neglecting the influence of the giant CCN modes, the emulators built average a mean relative error of less than 1% in predicting $\log_{10} S_{\max}$, which translates to an mean relative error of 9.2% and 8.9% in predicting N_{act}. Including the giant CCN mode appears, at first, to dramatically increase the performance of the emulator, bringing those same metrics down to 0.3% and 6.9% for the 4th-order scheme. Relative to the ARG and MBN schemes, the emulators are much more accurate on average when compared to our reference parcel model. However, we note that both the ARG and MBN schemes contain components which themselves are tuned to parcel models employed by their developers. Thus, we should not expect those schemes to perfectly match the parcel model results calculated here. Instead, we emphasize that the comparison of our emulators with the ARG and MBN scheme is motivated as part of a broader attempt to understand how the fundamental activation process initiates a chain of physics which ultimately lead to the aerosol indirect effect on climate. Assessing the relative performance of activation schemes which, for all intents and purposes, perform extremely well at reproducing their own reference parcel models, is a critical step in establishing the parametric uncertainty in translating aerosol to droplet numbers and which underlies uncertainty in global model estimates of the indirect effect.

For this reason, we supplemented the evaluation of our emulators by using a second set of input parameter samples drawn from aerosol fields simulated by an aerosol-climate model.

In contrast with previous studies, we use instantaneous fields in lieu of monthly or annual averages for our samples. Activation is inherently a fast process; because the microphysics schemes in aerosol-cloud models directly account for a tendency of new droplets formed via nucleation, the activation parameterization in any model will be called every time-step and for every grid-cell where clouds are occurring. Assessing activation schemes using temporally-averaged aerosol fields risks missing some combinations of input parameters and limiting the range of values for which the scheme will need to accurately perform.

Most of the emulators and schemes tested here perform somewhat differently over land and ocean, owing to the presence (or lack thereof) of natural and anthropogenic aerosols in these different regimes. Unfortunately, when focusing on the narrower range of aerosol parameters present in MARC (in comparison with the larger parameter space on which the emulators were trained), the emulators which explicitly account for giant CCN perform poorly, especially over ocean regimes dominated by sea salt. However, their counterpart performs nearly identical to the ARG scheme, showing a slight over-prediction of N_{act} in maritime regimes and a slight under-prediction over continents. In the global average, the emulator agrees better with the detailed parcel model than the ARG scheme. By comparison, the MBN scheme, while prone to under-predicting N_{act} in both regimes in these calculations, shows far less variance in its mis-prediction. This would suggest the MBN scheme actually performs extremely well - it is simply calibrated against a different baseline (in this case, a different parcel model). In particular, the MBN scheme does not show a difference in relative error between ocean and land regions, suggesting it is appropriately sensitive to a large range of different aerosol populations.

The results presented here have important implications for global modeling studies seeking to quantify uncertain in the aerosol indirect effect on climate. While different activation schemes generally perform equally well when faced with idealized sets of input parameters (Ghan et al., 2011), their application in coupled aerosol-climate models may not be straightforward. Relative to parcel model calculations, activation schemes can likely show biases in predicting cloud droplet number in different regions of the world owing to spatial heterogeneity in the underlying aerosol and meteorology parameter distributions. This, in turn, will lead to biases in cloud radiative forcing and diagnosis of the indirect effect.

Some literature has already implicated the role of activation schemes in divergent model estimates of the indirect effect. Ghan et al. (2011) performed a pair of GCM experiments

using two different schemes and noted that between their simulations there is a 10% difference in the global average droplet number concentration, which produces a $0.2\,\mathrm{W\,m^{-2}}$ difference in the indirect effect. Using a sequence of increasingly-complex activation schemes, Gantt et al. (2014) performed similar simulations with just a present-day emissions scenario, showing large regional differences in average cloud droplet number concentration and, as a result, up to a difference of $0.9\,\mathrm{W\,m^{-2}}$ in global average shortwave cloud forcing. This perturbation in forcing naturally follows from results such as those highlighted in this work; changes in the base cloud droplet number concentration simulated in an aerosol-climate model have important consequences for the chain of cloud microphysical processes which ultimately give rise to the indirect effect. Those biases will necessarily be model-dependent, since the formulation of the basic activation diagnostic in each model is intertwined with regional and global variability in their simulated aerosol size distributions.

Future work should seek to systematically assess the differences in cloud microphysical processes and aerosol-cloud interactions arising from choice of activation schemes in aerosol-climate models. As this work illustrates, employing emulators of detailed parcel model calculations including various chemical and physical effects on the activation process will aid with this task, providing a way to quickly account for myriad facts which may be difficult or impossible to include in existing activation schemes or frameworks. However, this work must proceed in tandem with efforts to place strong constraints on the climatology and variability of cloud droplet number concentration across regions and meteorological regimes. The synthesis of these two lines of work may provide the necessary constraints to diagnose systematic biases in the representation of fundamental aerosol-cloud interactions in global aerosol-climate models and thus reconcile the disagreement between model- and satellite-derived estimates of the indirect effect.

Chapter 4

The Influence of Aerosol Activation on the Aerosol Indirect Effect on Climate

4.1 Introduction

The interaction between aerosols and water vapor underpins a gross uncertainty in our understanding of climate change and the way anthropogenic activity impacts it. By acting as the nuclei which seed the formation of clouds, aerosols can indirectly influence the climate system. However, constraining just how large this influence may be has proven elusive, despite nearly a half-century of investigation (Boucher et al., 2013). Twomey (1977) proposed the idea that for an idealized cloud system with a fixed liquid water content, a change in the number of aerosol particles or cloud condensation nuclei (CCN) could yield important changes in the cloud's microphysical structure, particularly with respect to droplet number and size. This "first indirect effect" has implications for a cloud's reflectivity and thus interaction with shortwave radiation. A "second indirect effect" is thought to arise from these initial changes in cloud microstructure, feeding back on the evolution of the cloud (its life-

This chapter is being prepared (along with the following chapter) for submission to *Atmospheric Chemistry and Physics*

time and precipitation) and ultimately its macrophysical properties (cloud fraction, liquid water content, etc.) (Albrecht, 1989; Haywood and Boucher, 2000). These indirect effects occur in liquid clouds, mixed-phase clouds (both in the Arctic and deeply-convecting clouds), and potentially in certain ice-phase clouds. The sum total of these influence of aerosol on cloud micro/macro-physical properties and radiation comprises the aerosol indirect effect on climate (AIE).

Difficulty in constraining the role of the indirect effect on contemporary climate change arises from two different but closely-related sources. Ghan et al. (2013) and Carslaw et al. (2013) illustrated, using two very different approaches, that lack of constraints on the emissions and physico-chemical size distribution of the pre-industrial aerosol provide a source of epistemic uncertainty surrounding AIE. However, interactions between aerosols and clouds are also state-dependent; the sensitivity of cloud processes and properties such as precipitation or albedo to aerosol perturbations can also vary widely across aerosol and cloud regimes. Furthermore, the array of aerosol-cloud interactions (ACI) which give rise to these sensitivities is diverse and complex. Sufficiently representing this complexity in global models which can only afford spatial and temporal resolution far too coarse to explicitly–and sometimes consistently–resolve them has proven an enduring scientific challenge, and is a second source of uncertainty towards AIE.

Here, we consider a fundamental component of ACI—the nucleation of cloud droplets from the ambient aerosol population, a process known as droplet activation. Droplet activation plays a key role in setting the climatology of cloud droplet number in a global model by providing the initial inputs to cloud microphysical processes, and provides a direct link between otherwise-independent aerosol and microphysics modules. Hoose et al. (2009) showed that in one global model, the indirect effect scaled nearly linearly with an artificial constraint placed on the minimum permissible cloud droplet number, a constraint or tuning parameter first developed to mitigate unrealistically-low droplet numbers in remote maritime clouds. This effect has also been observed when assessing ensembles of different aerosol-climate models (Quaas et al., 2009). Similarly, Storelvmo et al. (2009) showed that different empirical relationships between aerosol burden and droplet number produced grossly different estimates of the indirect effect.

The increasing complexity and scope of aerosol modules and activation schemes used in tandem with global models has obviated some of the need for this type of tuning parameter.

However, the choice of activation scheme can directly influence simulated droplet number climatologies. Ghan et al. (2011) found a 10% difference in AIE when using two different activation schemes in the same global model, despite a 20-50% difference in simulated droplet number, which is much smaller than the typical inter-model spread in AIE (e.g. Lohmann and Ferrachat, 2010; Boucher et al., 2013) By using a variant on that same global model with a more-complex chemistry and aerosol module, Gantt et al. (2014) dramatically increased the difference in simulated droplet number by 55% for the same two schemes; further, incremental improvements to those schemes yielded successively higher average droplet burdens. Although they did not estimate the indirect effect, they increased the global average shortwave cloud forcing by 13% for a present-day emissions scenario, illustrating the influence of activation on clouds' radiative impacts. In another study using the same global model, Gettelman (2015) observed a 28% decrease in the indirect effect when altering the numerics of activation such that droplets are nucleated before other microphysical processes act on the droplet number.

Of course, activation is just the first linkage in the chain of ACI, and other microphysics processes also contribute greatly to variance in the indirect effect. For instance, the parameterization of autoconversion, which relies on empirical relationships which are extremely difficult to constrain using observational data, can also grossly influence the indirect effect (Gettelman et al., 2013; Rotstayn and Liu, 2005; Golaz et al., 2011); implementing prognostic precipitation reduced the indirect effect by 14% in one model by modulating the sensitivity of cloud liquid water path to aerosol perturbations (Gettelman, 2015). As previously mentioned, though, activation stands apart from these processes because of its role in serving as the interface between aerosol and cloud microphysics modules in global models. While many global models use similar microphysical parameterizations and schemes, there is a wide diversity in approaches towards representing the physics and chemistry of the atmospheric aerosol. In particular, while many of these aerosol models do a good job at reproducing the regional and vertical distribution of chemistry-resolved aerosol mass and particle number, they can simulate vastly different particle size distributions (Mann et al., 2014). Activation schemes directly translate these particle size distributions into cloud droplet number, setting the baseline which other microphysical processes modify.

In this study, we quantify the influence of the representation of activation on estimates of the indirect effect by using a suite of state-of-the-science parameterizations, including

an emulator of droplet activation based on an adiabatic parcel model (Rothenberg and Wang, 2016). The first set of analyses highlights the fundamental changes in the simulated cloud and radiation climatologies attendant with the model configurations using each of the different activation schemes. We include, as a point of reference, a similar analysis on a suite of aerosol-climate model simulations explicitly targeted at assessing inter-model variance in the indirect effect. We then use the results of these analyses to motivate a deeper investigation into aerosol-cloud interactions simulated by our model configurations and how they influence the indirect effect on both regional and global scales.

This chapter is organized as follows. Section 4.2 introduces the MARC aerosol-climate model used as a host model for investigating the influence of droplet activation on the indirect effect, as well as the observational and model intercomparison datasets we use to evaluate this influence and our simulation design. In Sections 4.3.1 to 4.3.2 we consider how the different activation schemes influence the simulation of clouds and radiation in a present-day emissions scenario, and their sensitivity to aerosol perturbations. We follow this with analysis in Sections 4.3.3 to 4.3.5 on the influence of droplet activation on the indirect effect. Section 4.4 summarizes our findings and discusses implications for future studies.

4.2 Methods

4.2.1 MARC Global Aerosol-Climate Model

In order to better assess aerosol impacts on climate, we utilize an improved aerosol-climate model based on previous work (Kim et al., 2008; Ekman et al., 2012; Kim et al., 2014). Originally known as the CAM-MIT, the aerosol-climate model directly couples a multimodal, mixing-state-resolving aerosol scheme with a three-dimensional, global climate system model. In this latest iteration, the parent climate system model is the National Center for Atmospheric Research's (NCAR) Community Earth System Model (CESM; version 1.2) with Community Atmosphere Model (CAM; version 5.3) (Neale et al., 2012), while the aerosol model remains the multimodal model of Wilson et al. (2001). The aerosol model has previously been used to resolve aerosol physics in cloud-resolving simulations (Ekman et al., 2004, 2006; Wang, 2005a,b; Ekman et al., 2007; Engström et al., 2008). This new implementation of the CAM-MIT based on CESM/CAM has been renamed the two-Moment,

Multi-Modal, Mixing-state-resolving Aerosol model for Research of Climate (MARC).

MARC explicitly simulates a complex mixture of aerosol species, each with an associated lognormal size distribution function, using a two-moment scheme. Within MARC, the aerosol species is divided into a set of externally-mixed modes, including three distinct sulfate modes (nucleation or "NUC", Aitken or "AIT", and accumulation "ACC"), pure black carbon (BC), and organic carbon (OC). Additionally, two internally-mixed modes are separately resolved, including a mixture of sulfate-black-carbon (MBS) and a mixture of sulfate-organic-carbon (MOS). With the MBS mixture, particles are assumed to consist of a black carbon core coated with a sulfate shell; within the MOS mixture, particles are totally internally mixed according to the volumetric ratio of sulfate and organic carbon present. In addition to the number and mass moments predicted for each mode, the MOS and MBS modes have an associated prognostic variable indicating how much carbon is present in the mixture, allowing the mixing state to freely evolve. Additionally, MARC includes a sectional, single-moment (fixed-size) representation of dust (0.16 µm, 0.406 µm, 0.867 µm, 1.656 µm bins), and sea salt aerosols (0.2 µm, 2 µm, 5 µm, 15 µm bins). All of the MARC, dust, and sea salt modes are fully coupled to and interact with both the radiation model as well as with the cloud microphysics scheme (through droplet activation, discussed below) in the CESM, enabling the model to fully resolve aerosol direct radiative effect and the full gamut of aerosol indirect effects. Sea salt is assumed to be composed of NaCl, while dust is assumed to be a mixture of minerals (Albani et al., 2014; Scanza et al., 2014). Prescribed size distribution and hygroscopicity parameters for each mode are summarized in Table 4.1.

Sulfate particles are formed in MARC via binary nucleation of H_2SO_4−H_2O (Vehkamäki, 2002), with prognostic gaseous sulfuric acid predicted by the default CAM interactive sulfur chemistry module (Barth et al., 2000). Both gas-phase oxidation of SO_2 and dimethyl sulfide (DMS) provide sources for H_2SO_4, as well as aqueous reactions of S(IV) with both H_2O_2 and O_3. Coagulation between modes produces both pure (externally-mixed sulfate) and mixed (MOS and MBS) particles; the pure carbon (BC and OC) modes age into their mixed counterparts through a prescribed constant-time aging scheme (40 and 20 h for each) which is limited by the availability of H_2SO_4 for condensation. Both primary and secondary organic carbon aerosol are emitted into the OC mode in the model; biogenic volatile organic vapors (specifically isoprene and monoterpenes) are converted upon emission into OC using a simple yield coefficient suggested by Griffin et al. (1999). We assume that both pure

Aerosol Mode	Geometric Mean Particle Diameter (µm)	Geometric Std Deviation (σ_g)	Density ($g\,cm^{-3}$)	Hygroscopicity (κ)
NUC	0 to 0.005 84	1.59	1.8	0.507
AIT	0.005 84 to 0.031	1.59	1.8	0.507
ACC	>0.031	1.59	1.8	0.507
OC	-	2.0	2.0	10^{-10}
MOS	-	2.0	†	†
BC	-	2.0	2.0	10^{-10}
MBS	-	2.0	2.0	0.507
DST01	0.16	1.4	-	0.14
DST02	0.406	1.4	-	0.14
DST03	0.867	1.4	-	0.14
DST04	1.656	1.4	-	0.14
SSLT01	0.5	1.59	-	1.16
SSLT02	2.0	1.37	-	1.16
SSLT03	5.0	1.41	-	1.16
SSLT04	15.0	1.22	-	1.16

Table 4.1: MARC aerosol mode size distribution and chemistry parameters. The MOS mode (†) has a composition-dependent density and hygroscopicity which is computed using the internal mixing state of organic carbon and sulfate present at a given grid-cell and timestep.

carbonaceous modes are hydrophobic.

The sulfate particles grow and their size distribution evolves in response to the condensation of sulfuric acid and coagulation between particles. Particles from all modes can be lost through dry deposition, gravitational settling and impaction scavenging via precipitation, and each mode undergoes these processes with different efficiencies related to their size (Ekman et al., 2006) and hygroscopicity (Petters and Kreidenweis, 2007). Additionally, nucleation scavenging occurs in both deep convective and stratiform clouds. In stratiform cloud systems, nucleation scavenging is calculated through a prognostic droplet activation scheme, taking into account both local meteorology (sub-grid scale updraft speeds) and the total availability of ambient aerosol. In contrast, nucleation scavenging in deep convective clouds is computed using a prescribed supersaturation max of 0.1%. Although several aerosol species (sulfate and dust) play a role in heterogeneous ice formation in MARC (following Liu et al., 2007), this process does not remove ambient aerosol.

MARC adopts the stratiform cloud microphysics scheme from CAM5.3 (Morrison and Gettelman, 2008) and includes the updates to code structure and droplet activation tenden-

cies referred to as MG1.5 by Gettelman and Morrison (2015). The contribution of droplet nucleation to the cloud droplet number tendency, $\frac{\partial N_d}{\partial t}$, is computed following Ovtchinnikov and Ghan (2005), and can be non-zero in both newly formed and pre-existing clouds. It is restricted to cloud-base in pre-existing clouds but can occur at all levels of newly-formed clouds where cloud water mass is predicted to develop. Additionally, $\frac{\partial N_d}{\partial t}$ includes sink terms such as accretion of cloud water, self-collection of hydrometeors, evaporation, autoconversion, advection, and inter-type scavenging of hydrometeors. Autoconversion is parameterized as a function of cloud water content and droplet number, N_d (Khairoutdinov and Kogan, 2000). Ice and mixed-phase cloud microphysics are based on Liu et al. (2007) and Gettelman et al. (2010).

4.2.2 Simulation Design and Setup

We perform a set of simulations with MARC using different activation schemes (see Appendix A for more details). For both the ARG (**arg**) and MBN (**nenes**) schemes, we run a **_comp** and a **_min_smax** scheme, where the latter incorporates the minimum-$S_{\max}$ heuristic. For the PCM schemes, we use both the **main4** and **gCCN3** schemes derived in Chapter 3, and the 4th-order OLS scheme (**ols4**) of Chapter 2 for the supplemental heuristic. We note that for the MBN scheme, we apply the κ-Köhler theory formulation to handle dust instead of the adsorption mechanics implemented by Kumar et al. (2009). Other than the change in activation schemes, the simulations are identical and use the same emissions scenarios.

For each scheme, we performed a pair of 6-year simulations using a horizontal grid resolution of $1.9^\circ \times 2.5^\circ$ and 30 vertical levels. Each simulation is run with prescribed sea surface temperatures and ice cover running an annual cycle for the year 2000. To focus this work on the indirect effect, we diagnose the aerosol direct radiative effect during runtime by running repeated radiative transfer calls with each aerosol in isolation from one another, but do not include these effects in the heating rates used to forward-integrate the model. The pairs of simulations differ only in their use of emissions scenarios; here, we use a present day (PD and pre-industrial (PI) corresponding scenario to the years 2005 and 1850, respectively. Following Kim et al. (2008), we use constant monthly emissions derived using an offline modeling process (Mayer et al., 2000; Wang, 2004) for BC, primary OC, and SO_2; emissions

of DMS and VOC (isoprene and monoterpene) also vary on a monthly basis.

For all simulations, we output monthly mean fields and analyze the final 3 years of output for both the PD and PI cases. Additionally, we output a suite of instantaneous cloud micro- and macrophysical variables sampled at either cloud-top or one kilometer above the surface and saved every three hours over the duration of the simulations. For consistency with the radiative transfer calculations in the model, the maximum-random overlap hypothesis is used to derive cloud-top quantities. To focus our study on warm, liquid-water clouds (which should be the most susceptible to changes in activation dynamics), we restrict our analysis of the instantaneous output to grid-points where the cloud-top temperature is greater than $-10\,^{\circ}\mathrm{C}$.

4.2.3 Other Data

Satellite Data

To assess MARC's performance in simulating present-day cloud and radiation fields, we use a climatology of observations derived from satellite-based sensors. Cloud micro- and macrophysical fields were derived using six years of monthly-average data from the MODerate Resolution Imaging Spectroradiometer (MODIS; Collection 5.1), an instrument on board the Terra satellite. Cloud droplet number is derived from Level 1 data from the same instrument using a technique employing an adiabatic cloud assumption (Bennartz, 2007, with updates provided via personal communication); for this reason, it is only suitable for maritime cloud regimes equator-ward of 60°. The global radiative budget at the top of the atmosphere is estimated using data from the Clouds and Earth's Radiant Energy System (CERES) sensor, which also flies aboard Terra. All data is re-gridded to the MARC simulation grid before time-averaging for analysis.

Cloud Droplet Number Observations

Because the MODIS-derived cloud droplet number concentration (CDNC) retrievals have a high degree of uncertainty, we also evaluate simulated droplet numbers against a large collection of *in situ* observations previously compiled by Karydis et al. (2011). We compare these observations to instantaneous output of in-cloud droplet number from our present-day simulations, first interpolated to 850 mbar, and then averaged over the seasons and

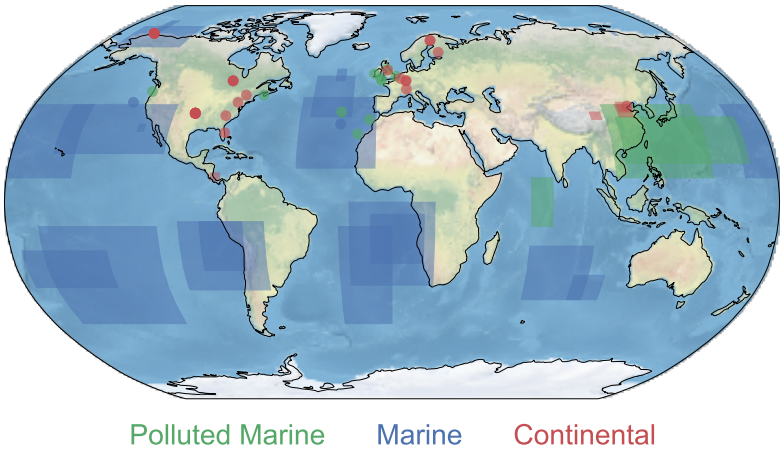

Figure 4-1: Locations of *in situ* observational data reported by Karydis et al. (2011). Different colors correspond to classifications of different aerosol regimes.

observation areas for the final three years in each model run (Table 2 Karydis et al., 2011). For observations from a specific location, we locate the model grid-cell containing that location for analysis. As a reference, we include the modeled CDNC corresponding for each observation produced by the chemical transport model simulations performed by Karydis et al. (2011). Figure 4-1 plots the global distribution from where the observations are sourced.

AeroCom Model Experiments

We supplement our investigation into the influence of droplet activation on the aerosol indirect effect by analyzing an additional set of climate models which participated in the Aerosol Comparisons between Observations and Models (AeroCom) Indirect Effects Experiment (as used by Ghan et al., 2016; Zhang et al., 2016). Similar to the experiment conducted here, pairs of integrations (using present-day and pre-industrial emissions scenarios) were performed with each model, using the same IPCC emissions scenarios for primary aerosol and precursors (Lamarque et al., 2011). Each simulation uses prescribed sea surface temperatures, sea-ice extent, and atmospheric greenhouse gas concentrations, and was run for a length of five model years. Unlike the integrations performed here, though, these simulations differ in two key ways.

First, all of these simulations were run using a Newtonian relaxation method in order to constrain the evolution of the atmosphere in each model (Kooperman et al., 2012). To

accomplish this, the prognostic model equations were nudged using re-analysis wind fields corresponding the period 2006-2010. It has been previously shown that orchestrating ensembles of models in such a way tends to provide more robust estimates of the indirect effect in shorter simulations, increasing the statistical significance of regional differences due to the aerosol-induced perturbation (Kooperman et al., 2012). Furthermore, Kooperman et al. (2012) demonstrate that compared to free-running simulations, nudged simulations yield a slightly smaller estimate of the indirect effect. Some AeroCom models supplemented nudging of the wind fields by simultaneously nudging temperature fields towards re-analysis. The simulations we perform with MARC are conducted with a free-running model configuration. We therefore stress that the AeroCom simulations provide a robust multi-model reference to evaluate the potential contribution of activation to uncertainty in the indirect effect; we do not seek to identify a model configuration which provides a "best" estimate.

Second, the present-day carbonaceous aerosol emissions used in the AeroCom simulations differ from those employed here. The offline-procedure used in Wang (2004) produces a larger estimate of globally- and yearly-integrated BC and OC mass emissions by factors of 1.8 and 1.6, respectively. For both fields, the difference is approximately evenly distributed across latitude bands from -30S to 50N, although OC emissions used in MARC are about 10% lower than in the AeroCom simulations northward of 50N. MARC derives primary emissions in terms of number concentration using these mass emissions and an assumed size distribution. Thus, our simulations will tend to simulate a larger burden of OC and BC in terms of particle number, particularly in regions where there is strong anthropogenic emission. The simulations we perform here do not include atmospheric heating due to absorption by BC and OC, which further are incapable of nucleating cloud droplets. However, both BC and OC can exert an indirect effect through aging and mixing with sulfate.

Where possible, we have attempted to include output from as many participating AeroCom models as possible in our analysis. We discriminate inclusions based on a quality-control process. Not every model reported the full set of requested output, which limits our ability to analyze some climatological fields. Also, some models reported inconsistent units, particularly for cloud-top and in-cloud microphysical quantities. We have corrected these where possible, or attempted to re-derive from hourly instantaneous output missing data. For our analysis, we re-grid each model to the simulation grid used in our MARC runs. A detailed summary of each model and its suite of parameterizations relevant for modeling the

indirect effect can be found in Appendix A of (Zhang et al., 2016).

4.3 Results

4.3.1 Evaluating Cloud Droplet Number

Predicted cloud droplet number concentration (CDNC) from each variant of MARC are compared against observations sourced from around the globe (Figure 4-1) in Figure 4-2. MARC generally under-predicts CDNC in regimes where the highest concentrations were observed, particularly over continents. This is partly a function of the activation scheme; both of the nenes schemes and each of the min_smax schemes (relative to their full-competition reference) captures some of the variability towards higher CDNC. The same pattern of under-prediction in CDNC occurs in polluted marine regimes for each activation scheme. In clean marine regimes, all of the MARC simulations produce little variance in simulated CDNC, although most of the comparisons are within ±50% of the observations. The detailed aerosol, chemistry, and activation treatment in the NASA GMI model utilized by Karydis et al. (2011) produces much better agreement with observations, although their model also tends to consistently predict much too-high CDNC over continental regions; over half of their reported values are greater than their corresponding observed CDNC values by 50%.

Compared to the observations, MARC much more frequently simulates very low CDNC less than $50\,\mathrm{cm}^{-3}$ than shown in observations (Karydis et al., 2011). In fact, observations rarely probe cloud regimes with CDNC this low; the exceptions in the dataset compiled by Karydis et al. (2011) are both extremely far removed from anthropogenic emissions, either north of Alaska or near the Canary Islands. This does not necessarily mean that droplet number is infrequently below this threshold, or that MARC is accurately simulating aerosol-cloud interactions in these regimes. For instance, TKE-based diagnostics used to assess subgrid-scale updraft speed variability tend to simulate extremely low updraft speeds much too frequently, which can bias CDNC estimates low (West et al., 2014). Additionally, techniques for estimating droplet number using satellite measurements (e.g. Bennartz, 2007) rely on an adiabatic cloud approximation which can not account for entrainment and dilution of droplet number in some marine dynamical regimes, further lending to over-estimates.

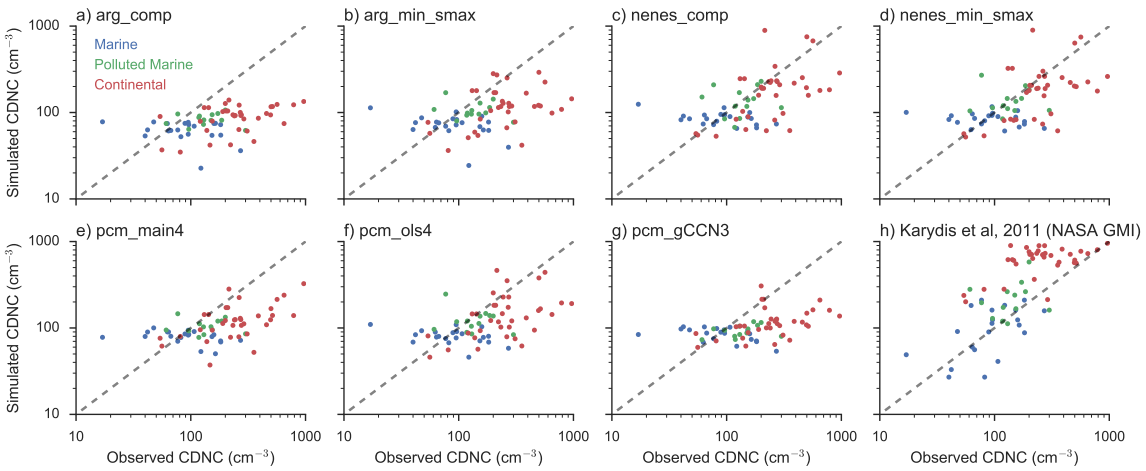

Figure 4-2: One-to-one comparisons between observed and simulated cloud droplet number concentrations from regions across the globe. Panels (a)-(g) show results from the MARC simulations using the indicated droplet activation schemes; panel (h) shows results from CDNC modeled by a chemical transport model with detailed aerosol and activation treatments.

The potential for regionally-dependent biases in predicted CDNC are more apparent when aggregating relative error, as in Figure 4-3. On average, MARC fares the worst in continental regimes, regardless of activation scheme. In contrast, the average simulated CDNC in clean marine regimes is well-calibrated, but has much higher variance. Polluted regimes tend to have the least variance, and the model performs better in these regimes than in purely continental ones. These biases obfuscate the fact that in the global average, CDNC burden is relatively well-predicted in comparison with recent modeling estimates. Estimates of global-average CDNC in the simulations performed here range from $60\,\mathrm{cm}^{-3}$ to $91\,\mathrm{cm}^{-3}$ (for the arg_comp and nenes_min_smax cases, respectively). This is mostly in agreement with recent studies, albeit on the lower side of estimates ($75\,\mathrm{cm}^{-3}$ to $135\,\mathrm{cm}^{-3}$ by Penner et al. (2006) using a suite of models employing the Abdul-Razzak and Ghan (2000) activation scheme; $83\,\mathrm{cm}^{-3}$ by Leibensperger et al. (2011) using an empirical relationship between aerosol and droplet number; $96\,\mathrm{cm}^{-3}$ by Barahona et al. (2011) using an earlier variant of the Morales Betancourt and Nenes (2014a) activation scheme).

For a more extensive assessment of simulated cloud droplet number, we turn to estimates derived from MODIS observations (a recent update to Bennartz (2007) provided by the author). These estimates Figure 4-4 are based on a simple model of adiabatically-stratified liquid clouds which is most appropriate for thin, stratiform clouds capping the maritime

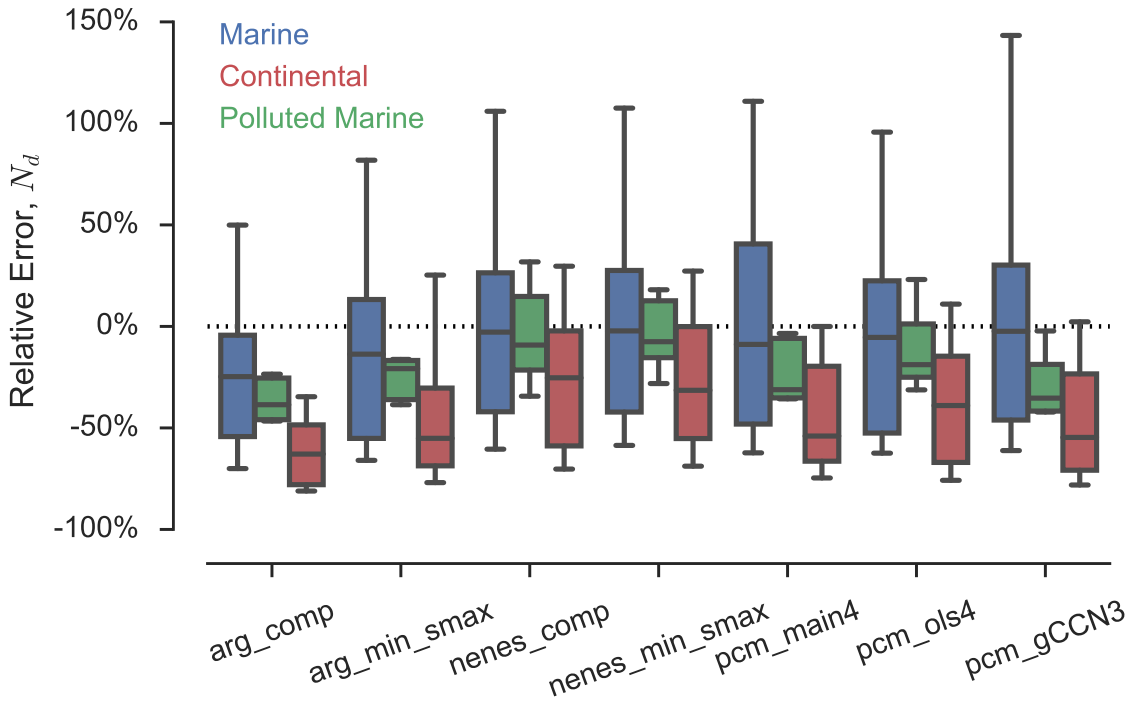

Figure 4-3: Distributions of relative error between observed and simulated CDNC (N_d) for each configuration of MARC, aggregated by region.

boundary layer, and thus we restrict our comparison to maritime regions between −60°S-60°N. Enhancement in CDNC downwind and in the vicinity of continents and anthropogenic emissions sources are clearly visible in the satellite dataset, particularly in the regions offshore of the United States and China. Averaged over the entire maritime region under consideration, MARC under-predicts CDNC by 45 % to 65 % depending on which activation scheme are used. However, CDNC is consistently too small in several regions regardless of that scheme, particularly in both the north and south Atlantic, in the portion of the Southern Ocean that lies south of the Indian Ocean, and in the north Pacific. The only oceanic region where the model over-estimates CDNC is in the equatorial upwelling region of the eastern Pacific. Enhancement of CDNC by anthropogenic aerosol in costal regions is best-captured by the nenes and emulation schemes.

In interpreting these comparisons between observed and simulated CDNC, it is critical to recognize that there is an important set of underlying parameters which is not constrained—the aerosol size distribution and local meteorology. Globally-sourced observations of local variability in aerosol size distribution and composition are far more difficult to ob-

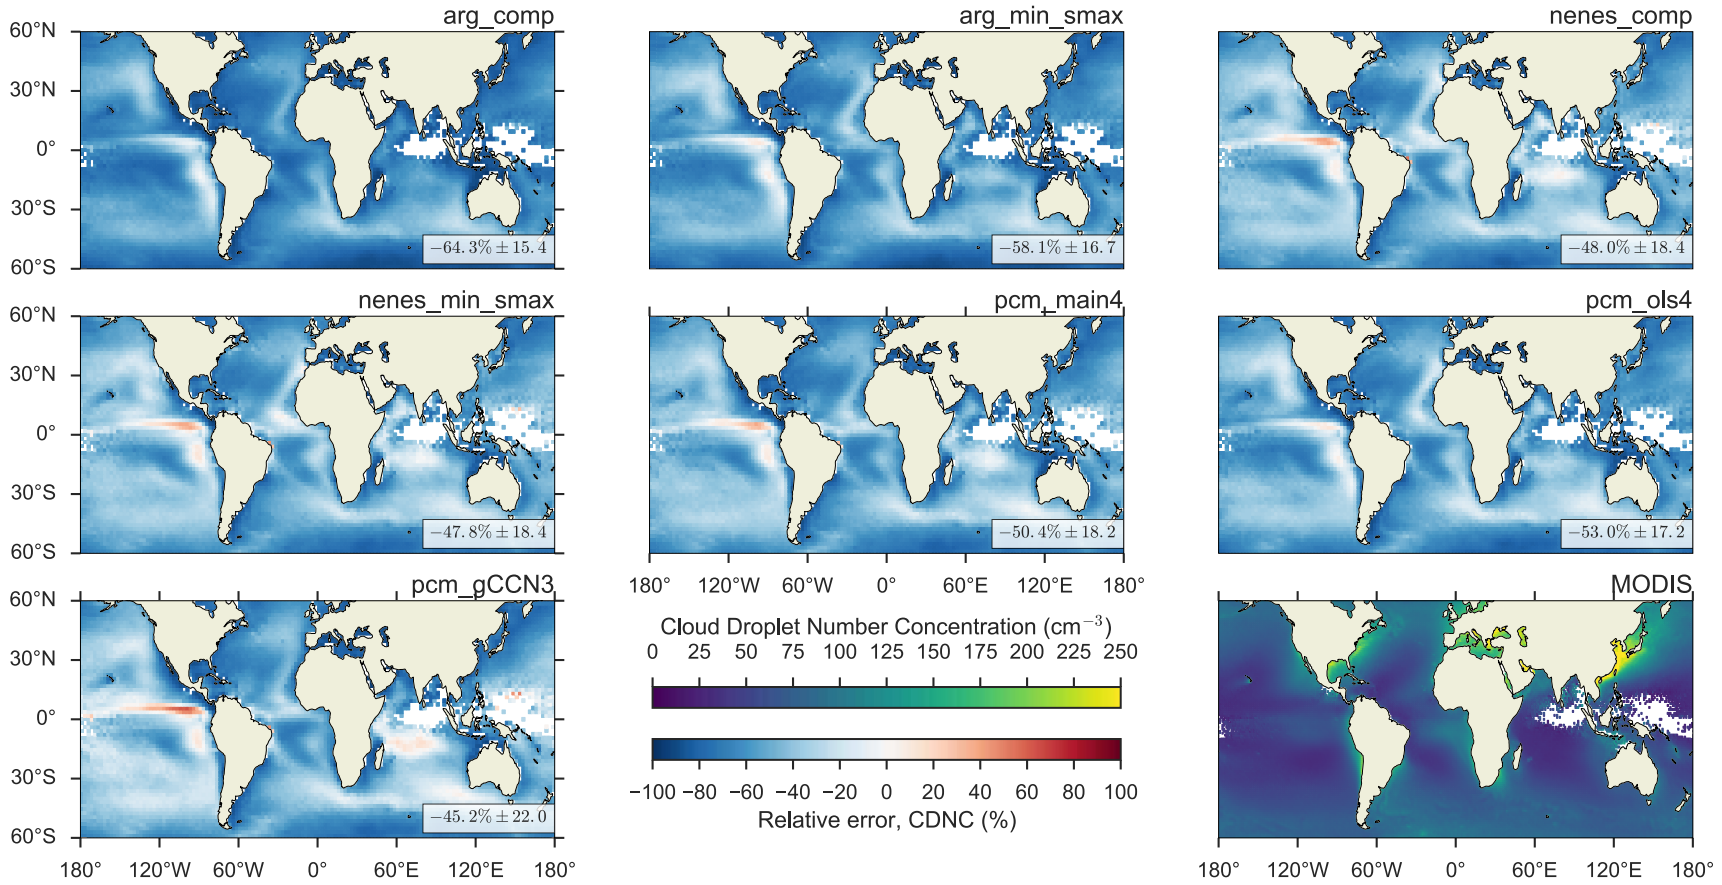

Figure 4-4: Global distribution in relative error of MARC-simulated CDNC versus MODIS-derived satellite observations (bottom-right)

tain than those of cloud microphysical properties. Without confirmation that a model is faithfully simulating the aerosol size distribution at a given location, one cannot expect the model to perfectly simulate cloud droplet number. For instance, Morales Betancourt and Nenes (2014b) highlighted the importance of the number concentration of particles in the accumulation mode aerosol as a critical factor in controlling cloud droplet number, a result which Chapter 3 further illustrated, specifically applied to the MARC model employed in this study. In cases such as this study, where we limit our focus to one global model and do not modify the underlying aerosol module, we are at least evaluating the CDNC from a consistent frame of reference. But in inter-model comparisons, differences in that aerosol module are an important caveat.

Although the use of different activation schemes does not directly perturb the simulated aerosol distributions in the model, the two-way coupling facilitated by nucleation scavenging can indirectly influence average aerosol number concentration. In these simulations, the present-day accumulation mode number concentration over the oceans is 31 % to 40 % smaller in the simulations using the nenes and pcm activation schemes versus the arg. This is likely because those schemes tend to nucleate slightly more droplets for the same aerosol population. The attendant increase in nucleation scavenging decreases accumulation mode number, which then depresses potential cloud droplet number. As a result, the difference in cloud droplet number concentration in the long-term average between the different schemes is not as large as it otherwise might be, hence the similar distributions of error relative to MODIS-derived CDNC. The region of too-high CDNC in the eastern equatorial pacific coincides with a region of enhanced, persistent deep convection and precipitation in the model.

4.3.2 Influence on Clouds and Radiation

Compared to original version of CESM/CAM5.3, the inclusion of an alternative aerosol formulation does not substantially change the model's simulated cloud and radiation fields, as illustrated in Figure 4-5. To demonstrate the extent to which altering the activation scheme can influence these fields on the large-scale, we have included in Figure 4-5 zonal averages computed from an ensemble of aerosol-climate models, four of which are themselves variants on the CAM5.3 with alternative microphysics (Gettelman and Morrison, 2015) and/or

moist turbulence scheme (Bogenschutz et al., 2013). Except in the northern hemisphere sub-tropics, MARC tends to under-predict total aerosol optical depth (AOD)relative to both observations and the reference models (Figure 4-5a). This is generally the case for all of the reference models as well, and MODIS estimates of AOD are thought to be biased high over oceanic regimes (Levy et al., 2013b). However, MARC consistently predicts smaller AOD than the other models considered here. As discussed in Section 4.3.1, there are differences in simulated accumulation-mode sulfate number concentration depending on the activation scheme used with MARC, but the fact that these differences do not show up in the zonal-average AOD suggests that MARC is instead consistently predict too-low coarse mode aerosol.

Present-day zonal average cloud macrophysical quantities are summarized in Figure 4-5b-d. MARC generally performs comparably with other models in re-producing zonal patterns in cloud fraction, cloud optical depth, and liquid water path. Using different activation schemes does not appreciably alter the models' performance with regard to these fields, although nenes and PCM schemes produce slightly higher cloud optical depth across latitudes and particularly in the tropics. As with its parent model and other reference models, MARC under-predicts cloud optical depth and liquid water path in polar regions. The preponderance of mixed-phase clouds greatly complicates these regions, and MODIS retrievals can become unreliable due to changing surface conditions (in particular, reflective surfaces such as snow which increase in frequency towards the poles). Notably, all or most of the reference models similarly under-predict cloud optical depth and liquid water path poleward of 60°, and MARC is well within the inter-model spread in these quantities across latitudes.

Following the discussion in Section 4.3.1 the largest activation-induced differences between simulations arises in cloud-top cloud droplet number (CDNC, Figure 4-5e). Poleward of 60°, CDNC using the nenes schemes can be twice that predicted by the arg ones. The differences are most pronounced in latitudes with significant anthropogenic aerosol emissions, particularly in the Northern Hemisphere. All the models shown here produce sharp declines in CDNC in the extra-tropics, but MODIS-derived estimates are highly uncertain in this region and provide a poor constraint on model performance. The combination of these differences in cloud properties yields small differences in the model-estimated shortwave cloud radiative effect (Figure 4-5f). Note that we apply the radiative decomposition of Ghan (2013) to account for the influence of aerosol scattering above clouds when computing

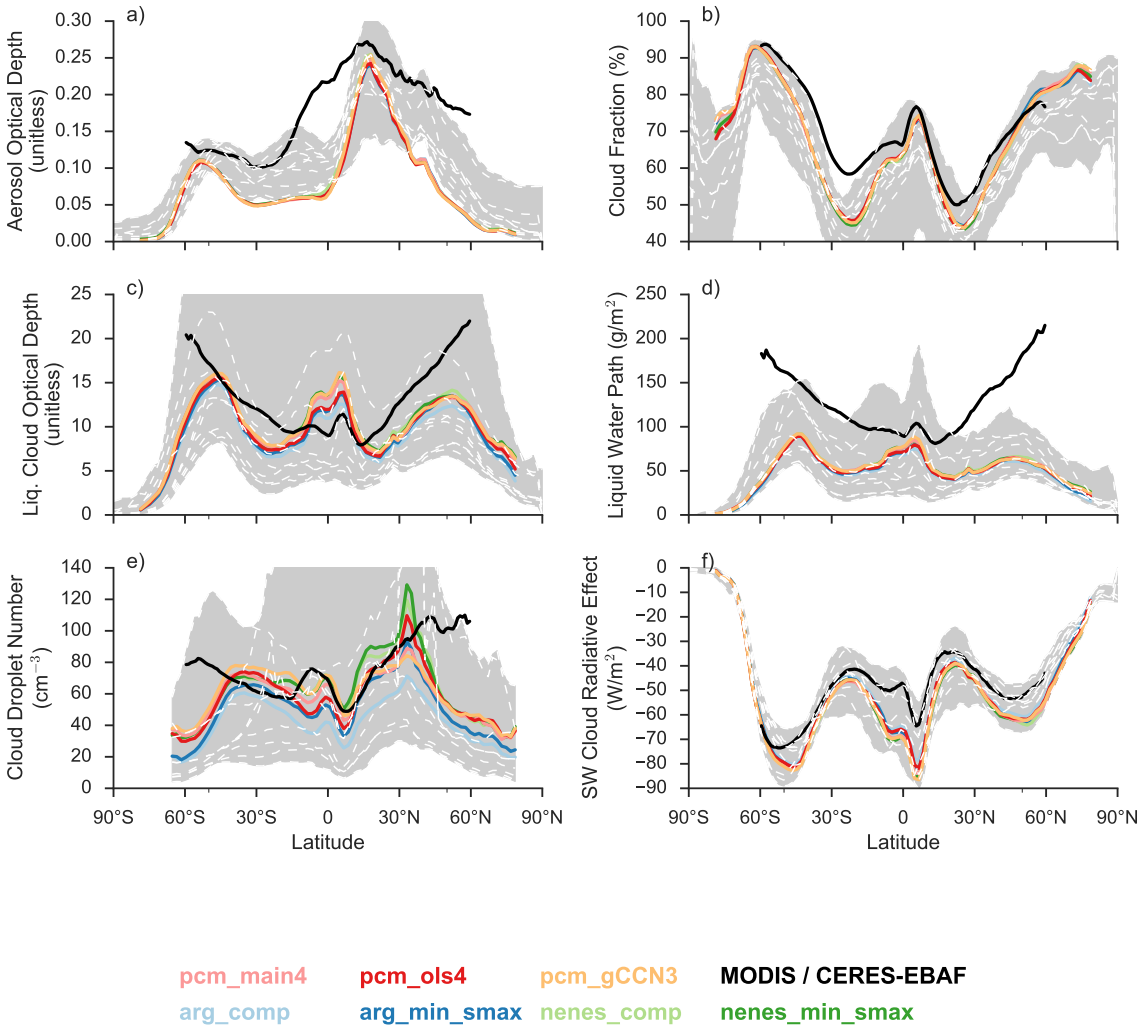

Figure 4-5: Zonal average aerosol, cloud, and radiation fields under present-day emissions scenario. Colored lines correspond to configurations of MARC using different activation schemes; black lines are derived from CERES-EBAF (Shortwave Cloud Radiative Effect) and MODIS (all other panels) observations. The shaded gray area corresponds to the inter-model spread for all available models participating in the AeroCom Indirect Effects Experiment; dashed white lines are the zonal averages for each participating model. Cloud droplet number is computed at cloud-top, using only grid cells over the ocean.

cloud radiative effects in the shortwave.

Measured in absolute terms, the inter-model difference in present-day shortwave cloud radiative effect is small. However, Figure 4-6 shows that the choice of activation scheme can influence the details of these differences. The change in cloud radiative effect between the pre-industrial and present day emissions simulations has a spread of nearly $2\,\mathrm{W\,m^{-2}}$ across

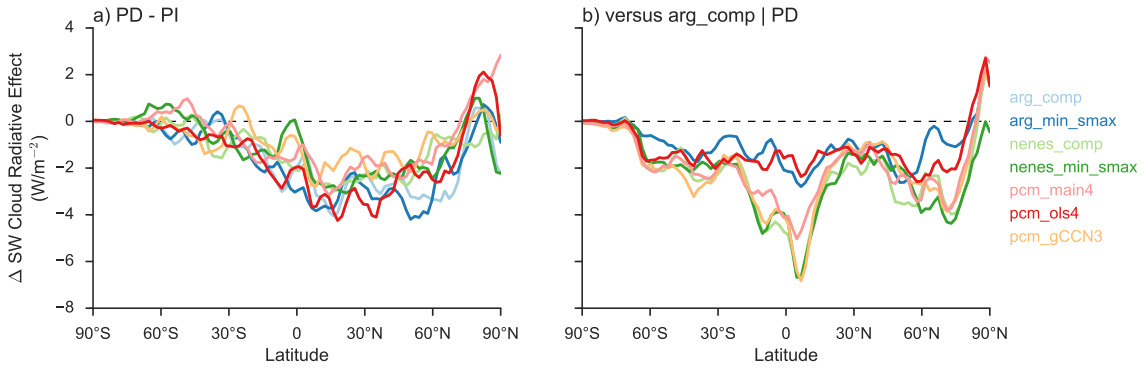

Figure 4-6: Difference in shortwave cloud radiative effect between pre-industrial and present-day emissions scenarios (a) and relative to the arg_comp simulation for present-day emissions (b).

all latitudes. These differences can be much larger on local scales. Activation schemes which produce the smallest cloud-top droplet number concentration generally produce largest differences in cloud radiative effect between the two emissions schemes. Relative to the arg_comp scheme, all other schemes produce a smaller cloud radiative effect in the present-day emissions case (Figure 4-6b), and the magnitude of that difference is comparable to the change versus the pre-industrial emissions. These largest differences arise in the tropics and in the mid-latitudes of the northern hemisphere, both regions influenced by anthropogenic aerosol emissions and where the largest differences in CDNC occur.

Figure 4-7 illustrates the change in aerosol and cloud microphysical fields between the pre-industrial and present-day emissions scenarios. In order to focus the following discussion on the influence of droplet activation, we opt to present differences in cloud condensation nuclei (CCN) which more directly relate to potential droplet number than aerosol optical depth, even though both fields are thought to be strongly correlated (Andreae, 2009). CCN increases the most in the northern hemisphere, where aerosol emissions are the largest (Figure 4-7a). These increases are co-located with the largest increases in cloud droplet number (Figure 4-7b), although there is a factor of 3-4 difference between activation schemes. At the same latitudes, droplet effective radius decreases (Figure 4-7c), optical depth increases (Figure 4-7d), and liquid water path increases (Figure 4-7f). The change in liquid water path differs the most between each simulation following cloud droplet number. Despite which activation scheme is used, compared to Figure 4-5c-d, liquid water path and cloud optical depth increase by up to 20 % co-located where the largest increases in droplet number oc-

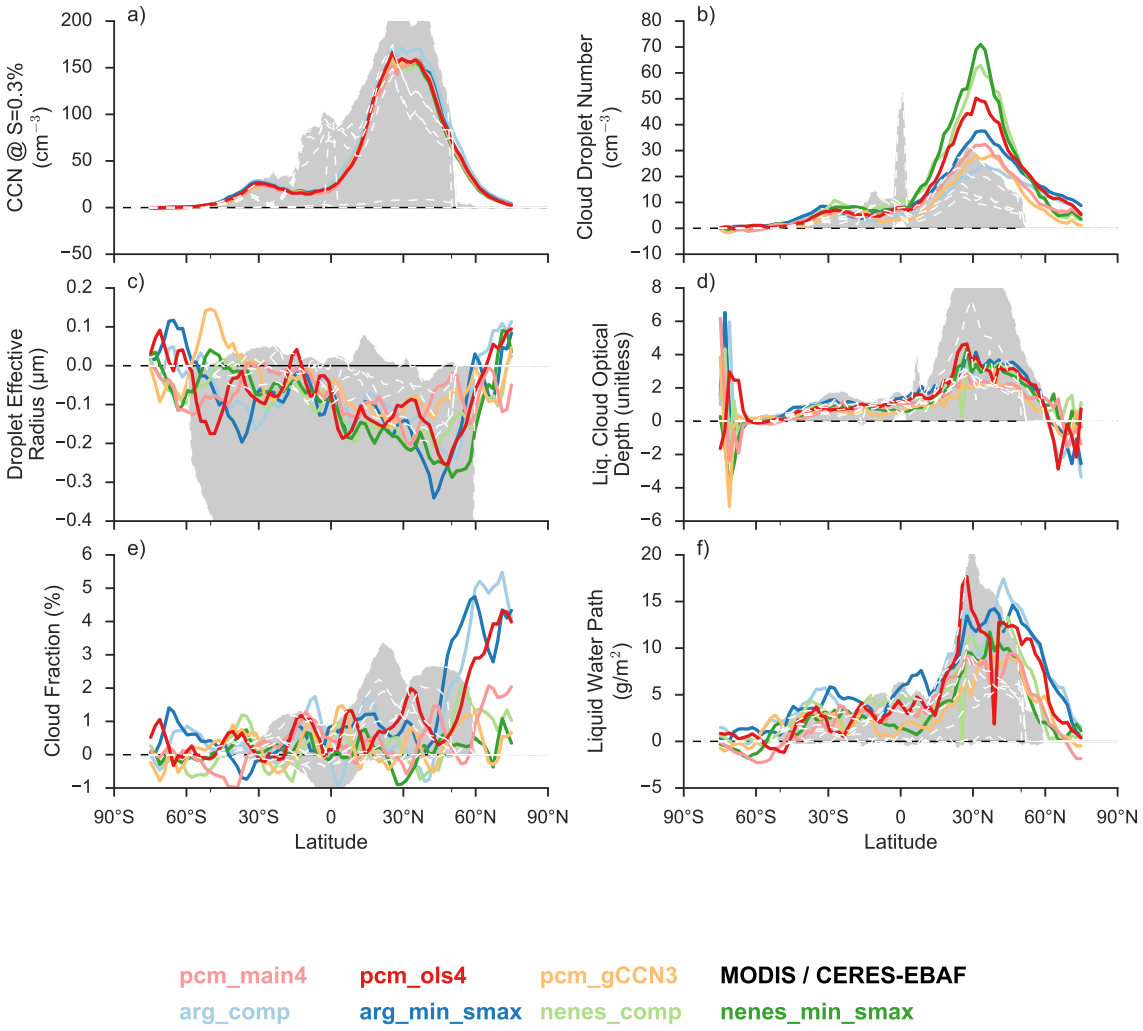

Figure 4-7: Same as Figure 4-5, except illustrating differences in indicated fields between pre-industrial (PI) and present-day (PD) emissions scenarios. CCN here is computed at an altitude of 1 kilometer in the model.

cur. This indicates the important role that activation plays in influencing simulated cloud properties.

4.3.3 Influence on the Global Aerosol Indirect Effect

We now turn our attention to global metrics for the aerosol indirect effect—those considered by the IPCC and others (Boucher et al., 2013). Figure 4-8 shows differences in globally-averaged aerosol direct and indirect effects computed from the pairs of pre-industrial and present-day runs for each activation scheme. The particulars of droplet activation in these simulations produce up to a 0.8 W m^{-2} difference between simulations, or a 100 % increase

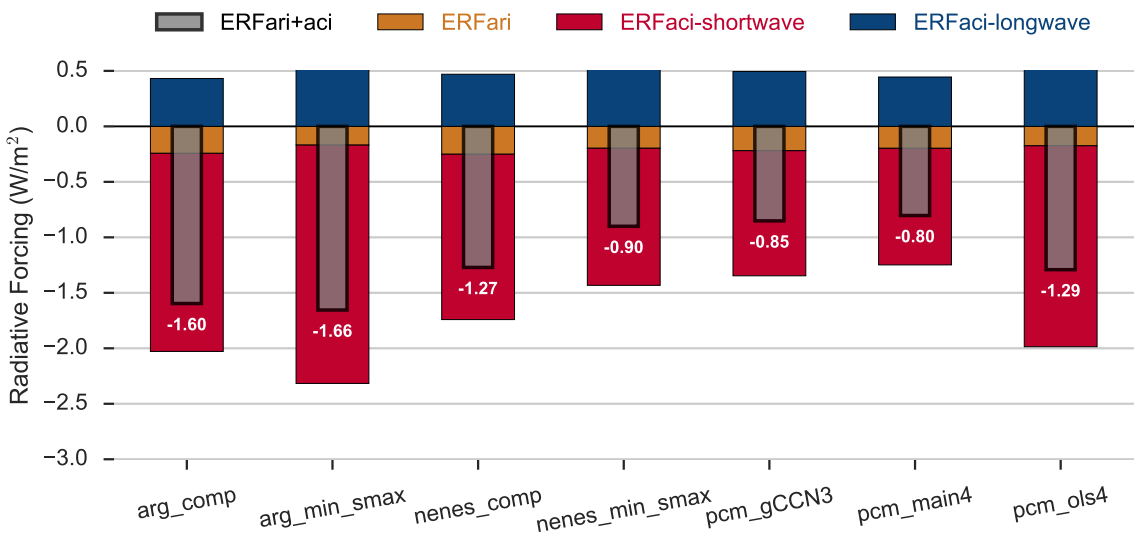

Figure 4-8: Global-average effective radiative forcing for aerosol direct radiative effects (ERFari) and indirect effects (ERFaci) in both the shortwave and longwave. The total effect is computed as the sum of the direct and both indirect components.

Activation Scheme	Δ R	Total CRE	Δ SW CRE	Δ LW CRE	DRF	Δ RF	Res
arg_comp	-1.66	-1.35	-1.79	0.43	-0.24	-1.6	0.07
arg_min_smax	-1.62	-1.49	-2.15	0.66	-0.17	-1.66	-0.04
nenes_comp	-1.35	-1.02	-1.49	0.47	-0.25	-1.27	0.07
nenes_min_smax	-0.9	-0.70	-1.24	0.53	-0.20	-0.9	-0.001
pcm_gCCN3	-0.72	-0.63	-1.13	0.50	-0.22	-0.85	-0.13
pcm_main4	-0.85	-0.61	-1.05	0.45	-0.20	-0.8	0.05
pcm_ols4	-1.32	-1.12	-1.81	0.70	-0.18	-1.29	0.03

Table 4.2: Aerosol direct and indirect effects (in $\mathrm{W\,m^{-2}}$) for the different activation schemes considered in this study. In all cases, we consider the change in the top-of-atmosphere radiative flux to compute these metrics (the net balance of which is given by Δ R). Our decomposition of the shortwave indirect effect (SW CRE) follows Ghan (2013) to account for above-cloud scattering and absorbing aerosol; similarly, the direct effect is computed diagnostically within each simulation. Following Gettelman (2015) we compute a residual between the top-of-atmosphere radiative imbalance and the direct and indirect effects such that Res = Total CRE + DRF - Δ R

by the strongest over the weakest estimate. Differences in the total aerosol influence on climate (ERFaci+ari) here are strongly modulated by perturbations to the shortwave cloud radiative effect (ERFaci_shortwave), which is broadly consistent with the changes in cloud microphysical properties noted in the previous section.

Each perturbed component in the top-of-atmosphere (TOA) radiative budget are decom-

posed in Table 4.2. Consistent with expectations, there is relatively little variance in the direct radiative forcing diagnosed for each activation case, even accounting for the feedback of increased nucleation scavenging depressing aerosol number. Similarly, there is not much difference in the longwave cloud radiative effect, which is dominated by glaciated clouds not necessarily modified by tweaking activation. The longwave indirect effect in these models is net positive; dust and large sulfate particles in the model can nucleate ice crystals (Liu et al., 2007), and the increase in aerosol between the two emissions scenarios yields ice clouds with modestly higher crystal number concentration, higher ice water path, and increased longwave influence relative shortwave. Using a similar modeling setup, Gettelman et al. (2012) showed a similar influence of aerosol on ice cloud-longwave radiation interactions. Furthermore, we note that the indirect effect in the longwave is critically sensitive to the baseline ice crystal number burden simulated in the model. Additional tests using an alternative, aerosol-coupled ice nucleation scheme (DeMott et al., 2010) decreased the longwave cloud radiative effect in MARC in a manner that scaled nearly linearly with ice crystal number. Using this alternative scheme produced much higher cloud-top ice crystal concentrations and ice water path, as well as a larger change between the two fields in the pre-industrial and present day simulations.

The majority of the difference in the indirect effect and net TOA radiative flux thus arises from changes in cloud interactions with shortwave radiation via cloud brightness. For the shortwave cloud radiative effect alone, the spread from the different activation schemes is larger than the net effect at $1.1\,\mathrm{W\,m^{-2}}$. Both the longwave cloud radiative effect and aerosol direct effect act to minimize the net effect; in the simulations with the largest shortwave contribution, both the direct and cloud longwave effects are smaller and larger, respectively. The small spread in direct effect in these simulations correlates very strongly with the change in global-average aerosol optical depth, but changes in that field are only loosely related to changes in the available CCN.

4.3.4 Sensitivity of Indirect Effect to Aerosol-Cloud Perturbation

In the previous sections, we demonstrated that differences in the treatment of droplet activation have a large influence on model estimates of global cloud and radiation fields, as

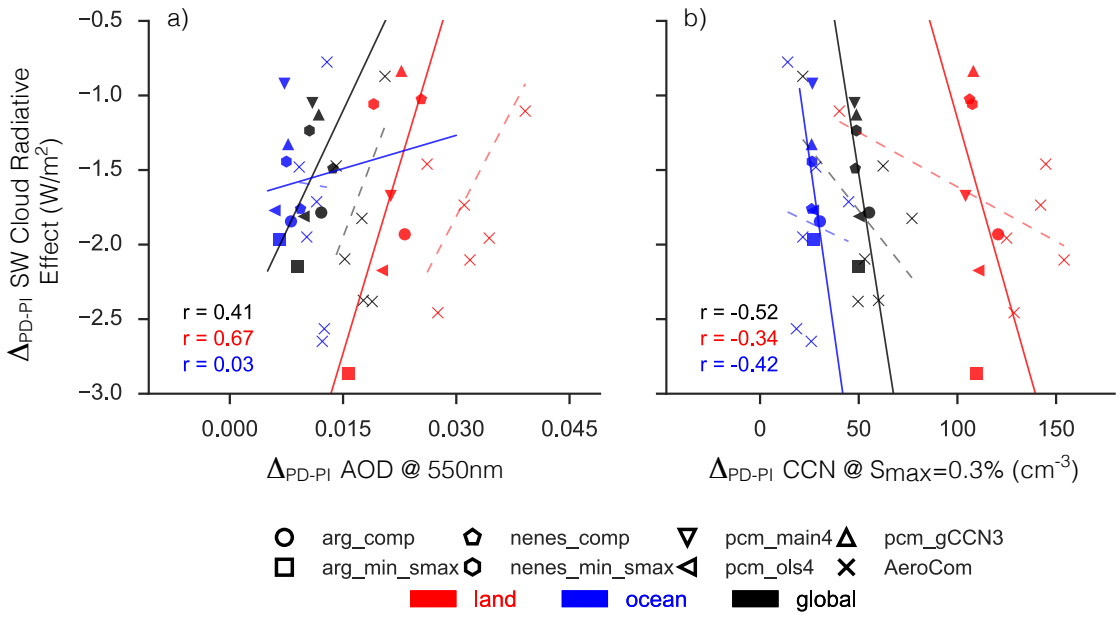

Figure 4-9: Relationship between regionally-averaged PD-PI change in shortwave cloud radiative effect and aerosol optical depth (a) or CCN concentration (b). Colors denote averages over all land area (red), ocean (blue), or whole globe (black). Glyphs denote MARC simulations with different activation schemes; X denotes AeroCom model. Linear regressions for MARC simulations are represented by solid lines, with corresponding correlation coefficient indicated on plot; linear regressions for AeroCom models are given by the dashed lines.

well as their perturbations to change. However, we seek to more rigorously connect these differences in sensitivity to aerosol perturbation with variance in the indirect effect. To accomplish this, we have used our high-temporal frequency output to consider grid cells where only liquid water clouds are present, between −60°S and 60°N. We then compute time and spatial averages of cloud and radiation fields over these grid cells, both across the global and by differentiating between oceanic and continental (land) regimes. The calculations were also computed for the AeroCom models previously considered here, where data was available. To quantify the aerosol-cloud perturbation/indirect effect (approximated as the change in clean-sky, shortwave cloud radiative effect between pre-industrial and present day) sensitivity, we regress differences in cloud and radiation fields obtained from the different models and model configurations.

Figure 4-9 details this relationship for two metrics of the change in the global aerosol burden. The increase in direct aerosol and precursor gas emissions in the present-day emissions scenario should tend to increase both aerosol optical depth (AOD) and the availability

of cloud condensation nuclei (CCN). Every model captures this change. However, the exact magnitude of the increase is dependent on the formulation of the aerosol module in each model, in particular assumptions about the size distribution. Inter-model diversity in present-day AOD and sensitivity to perturbation from the pre-industrial climate has been associated with up to a $0.5\,\mathrm{W\,m^{-2}}$ spread in estimates of the direct effect (Shindell et al., 2013), but it also has implications for the indirect effect following Figure 4-9-a. The largest AOD increases occur over land, and these are also associated with larger perturbations to the indirect effect. The MARC simulations and most of the AeroCom models do not simulate major differences in AOD over the ocean, even while there is considerable spread in the sensitivity of cloud radiative forcing in that regime. One explanation for the small response in AOD over these ocean regimes could be that AOD is dominated by large natural aerosol such as sea salt in these regions, which would not directly increase in response to anthropogenic emissions.

CCN also directly increases due to anthropogenic emissions. However, in contrast with AOD, larger changes in CCN are associated with a larger indirect effect (Figure 4-9-b). Different aerosol metrics have previously been shown to have different relationships with model- and satellite-derived estimates of the indirect effect, but they usually at least have the same sign (Penner et al., 2011). This suggests that each metric is capturing a different facet of the aerosol size distribution which may or may not be relevant to changes in the indirect effect, depending on how they strongly they influence potential cloud droplet number.

To assess this influence, we plot similar relationships between cloud-top droplet number concentration (CDNC), liquid water path (LWP) and liquid cloud fraction in Figure 4-10. The spread in the change in CDNC very weakly correlates with the strength of the indirect effect (Figure 4-10-a). Within the AeroCom models, this is likely due to the diverse representations of cloud microphysics and macrophysics utilized by the ensemble (Ghan et al., 2016). Instead, a much better predictor of the indirect effect is the pre-industrial CDNC, which is strongly positively correlated with the change in CDNC between pre-industrial and present day. This is evident in Figure 4-10-b, which shows a negative relationship between the pre-industrial CDNC and the indirect effect. The models which produce higher CDNC for the same background or natural aerosol tend to have weaker indirect effects.

Liquid water path and cloud fraction exhibit a different relationship with the indirect effect (Figure 4-10-c-f). Using different activation schemes in MARC directly influences the

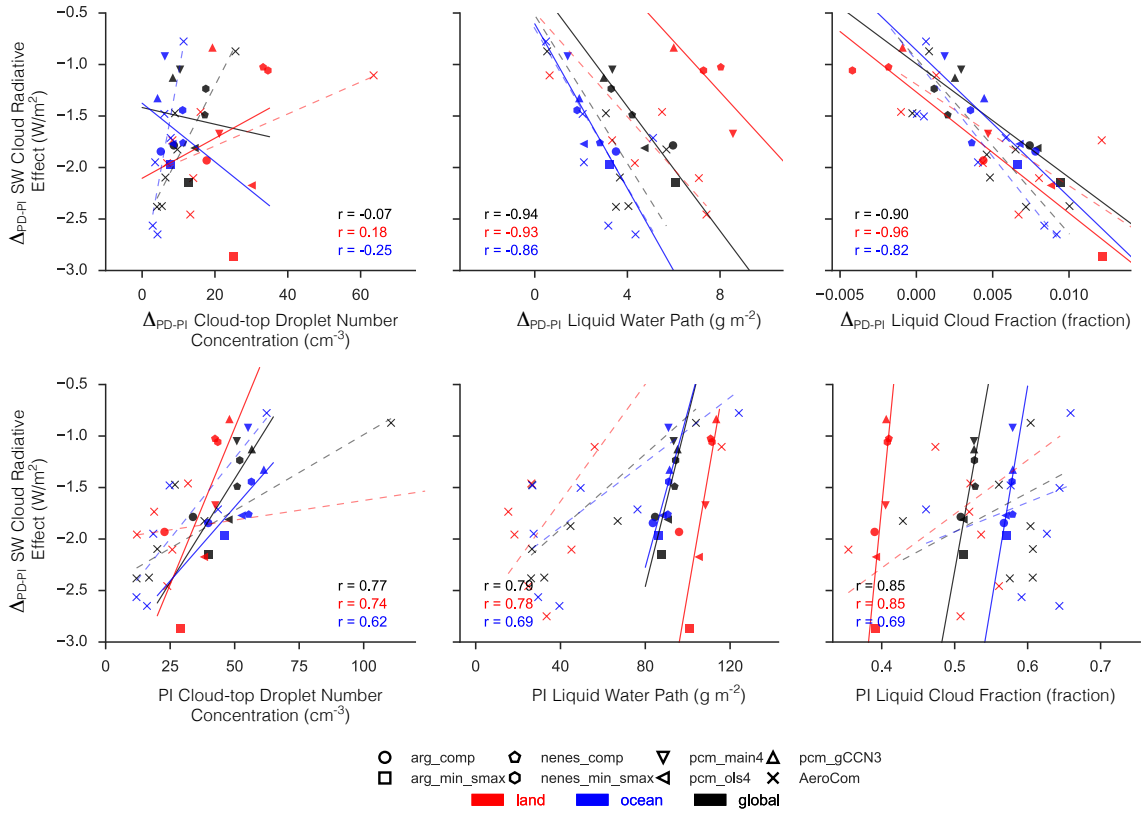

Figure 4-10: Similar to Figure 4-9; top row (a-c) denotes relationship between PD-PI changes in shortwave cloud radiative effect and cloud properties, bottom row (d-f) shows regression versus averages from PI emissions scenario.

sensitivity of LWP to aerosol perturbation, which strongly increases the indirect effect. The same is true of the liquid cloud fraction, which is correlated with LWP in both MARC and the ensemble of AeroCom models, particularly for oceanic regimes. The AeroCom models simulate far more diversity in LWP and cloud fraction under pre-industrial emissions, but tend to agree on the magnitude of change between pre-industrial and present day. The same is true of MARC, even across oceanic and continental regimes. That the pre-industrial average LWP and cloud fraction are similar regardless of the activation scheme used with MARC suggests that the large-scale cloud properties are less insensitive to aerosol and instead to the combination of aerosol and cloud macrophysics. This is somewhat striking, given the relative lack of variance in the modeled perturbations in LWP and cloud fraction, and points to the critical role of changes in droplet number due to aerosol to further influence cloud macrophysics.

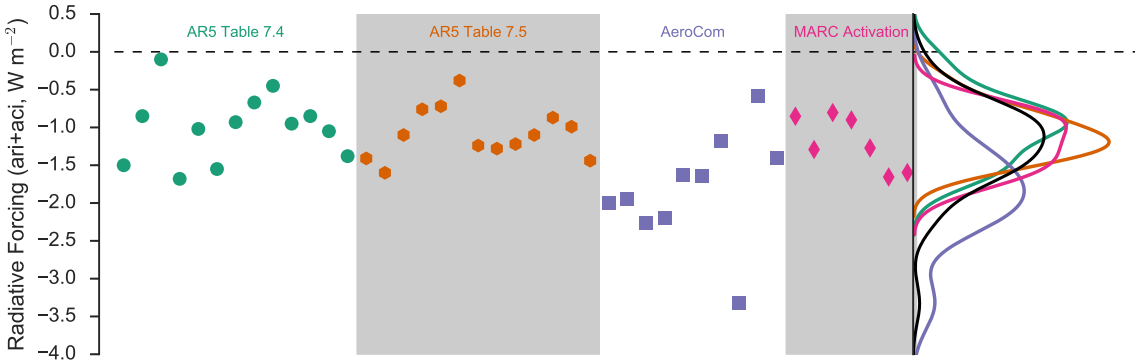

Figure 4-11: Comparison between estimates of ERFari+aci derived here and a subset of those previously reported by the IPCC AR5. Corresponding kernel density estimates of the distribution of ERFari+aci are given on the right-hand panel; the solid black curve shows the distribution accounting for all the estimates on the plot.

4.3.5 Summarizing the Influence of Droplet Activation

In order to contextualize the influence of droplet activation on the indirect effect in the simulations presented here, we plot estimates of the indirect effect (ERFari+aci) reported by Boucher et al. (2013) in Figure 4-11. These include a highlighted subset of models and results combining satellite observations with model analysis (Figure 4-11, AR5 Table 7.4), results from a previous model inter-comparison using CMIP-class models (Figure 4-11, AR5 Table 7.5 Shindell et al., 2013), and new estimates derived from the AeroCom models considered here and the various configurations of MARC with different activation schemes. The estimates presented here span a wide variety of potential model physics and aerosol couplings, and therefore different aerosol indirect effects.

In our simulations with MARC, differences in droplet activation produce a spread in estimates of the indirect effect similar in magnitude to the total inter-model diversity. Furthermore, our estimates–especially for the configurations with lower cloud droplet number concentration–tend to cluster in the higher-end of estimates compared to previous inter-comparisons. The same is true for the AeroCom models considered here, although we note that four of the AeroCom models are closely-related variants of the same parent model as MARC (the NCAR CAM5.3) and therefore the estimates are not totally independent of one another.

Our range of indirect effects induced by different activation treatments is much larger than the few others reported in the literature. By re-ordering the droplet activation cal-

culation in each model timestep, Gettelman (2015) induced a $0.43\,\mathrm{W\,m^{-2}}$ decrease in the magnitude of the indirect effect; Ghan et al. (2011) reduced it by just $0.16\,\mathrm{W\,m^{-2}}$ when switching between two different activation schemes. This range is much smaller than the $0.86\,\mathrm{W\,m^{-2}}$ spread we report for the experiments conducted here. We note that both of those previous estimates of the sensitivity of the indirect effect to activation used nearly identical global models (early versions of the NCAR CAM5 with the same aerosol module). Our use of a unique aerosol model could contribute to some of the difference in the range of estimates of the indirect effect. This possibility can indirectly be tested using the suite of model results presented in this work, since one included AeroCom model is the NCAR CAM5.3 in its default configuration, which should be nearly identical to the arg_comp MARC configuration here save for the different aerosol module. The difference in ERFari+aci between these two simulations is $0.45\,\mathrm{W\,m^{-2}}$, which is half of the total range reported here for MARC with different activation schemes.

4.4 Discussion and Conclusions

In this study, we have quantified the influence of the representation of droplet activation in global models on the aerosol indirect effect. Using a suite of state-of-the-science activation parameterizations incorporated into our global aerosol-climate model, the MARC, we performed simulations under both pre-industrial and present-day aerosol emissions to estimate the indirect effect and relationship to changes in both cloud and radiation fields. Previously, few studies exploring the indirect effect focused explicitly on the role of droplet activation, instead concentrating on either the processes that load the ambient aerosol (emissions and atmospheric chemistry) or the results of changes purely occurring in cloud droplet number concentration (such as imposed minimum values for cloud droplet number or in microphysical processes which modify this number). Beyond assessing three unique activation schemes, we supplement our analysis by considering three additional, idealized droplet activation schemes which use a heuristic to simplify accounting for competition between different aerosol modes for moisture during the nucleation process.

The relationship between cloud droplet number concentration and aerosol in MARC is critically influenced by the representation of droplet activation. Estimates of CDNC in the present-day climate are up to 40 % higher in polluted regimes when using the most-sensitive

activation scheme, and the increase from pre-industrial to present day is up to twice as large. CDNC in regimes dominated by natural aerosol, especially remote marine regions with prevalent sea salt, is also impacted by the activation scheme. Using the advanced droplet activation schemes included here, which explicitly account for biases due to giant CCN particles, helps reduce the under-prediction in maritime regimes compared to satellite observations and in polluted regimes compared to in situ observations. However, MARC systematically produces too few CDNC in most parts of the globe. While this could be due to cloud processes, we emphasize that it could also be fundamentally related to the simulated aerosol size distribution within MARC and how it apportions aerosol number and mass in the ranges where likely-CCN are culled.

Compared to available satellite measurements and the models participating in the AeroCom inter-comparison, though, MARC does well at capturing the present-day climatology of cloud and radiation fields, likely because its parent model, the NCAR CESM, itself is well-tuned towards this end. However, the details of activation and how it influences cloud microphysics plays a major role in setting the shortwave cloud radiative effect. Under present-day emissions, the differences between that effect for each of the different activation schemes is as large as the change from pre-industrial for each scheme. This leads to large differences in the modeled indirect effect in each model, almost entirely occurring due to the shortwave cloud radiative effect. The resulting spread in indirect effect estimates is twice as large as that previously reported by studies considering activation, and about as large as the inter-model spread from both historical and recent model inter-comparisons, which consider models including a variety of different aerosol effects.

We note that the pre-industrial CDNC burden is a very strong predictor of the strength of the indirect effect, but not necessarily the change between pre-industrial and present day; this hints at the previously-hypothesized buffering effect of clouds on aerosol perturbations (Stevens and Feingold, 2009). Our analysis strongly supports this notion; our case which produces–by far–the largest change in CDNC yields nearly the smallest indirect effect, but simultaneously produces the highest values for pre-industrial CDNC. While consistent with the results of Hoose et al. (2009), this is exactly the opposite of those reported by Storelvmo et al. (2009), although their model uses prescribed aerosol fields so there is no interaction between cloud and aerosol processes. Changes in liquid water path and cloud fraction correlate much more strongly with the magnitude of the indirect effect than changes in

CDNC in our simulations. Both of these relationships hold for the broader sample of models provided by the AeroCom inter-comparison.

The weight of these results suggests an important role for activation in setting the sensitivity of the indirect effect. However, we caution that our approach is not able to disentangle the influence of activation from that due to the underlying aerosol model and its implicit aerosol size and CCN distributions. This is not meant to diminish the influence of cloud microphysical treatments on the indirect effect; Gettelman (2015) illustrates the importance of the implicit cloud lifetime effects arising from liquid water path changes associated with aerosols in contributing to the indirect effect. But since these relationships are themselves highly sensitive to simulated cloud droplet number, the influence of the aerosol size distribution and activation is somewhat more fundamental and just as poorly constrained with available observations. Furthermore, because of the chain of sensitivities initiating with the aerosol size distribution and activation particulars, estimates of the indirect effect produced from models with highly-simplified aerosol-CDNC relationships (such as explicitly prescribed CDNC or empirical fits to aerosol mass or volume) are likely significantly biased.

To test this idea, it would be valuable to carry out a similar set of alternative-activation simulations using a global climate model with a different underlying aerosol physics and chemistry model. Such work should be broadened to include chemical influences on activation where suitable. The influence of organic surfactants can enhance the sensitivity of droplet activation to variability in updraft velocity (Lance et al., 2004); this effect is compounded by the fact that co-condensation of organic vapors increases as particles grow (Topping and McFiggans, 2012). The quantification of secondary organic aerosol formation and its influence on climate is a major of focus of current modeling efforts (Jimenez et al., 2009) and could have important implications on activation dynamics and the indirect effect Probing the influence of organic vapors and aerosols on the indirect effect will likely require the development and assessment of new activation schemes.

We suggest that future sensitivity analyses in the vein of Carslaw et al. (2013) include perturbations to the fundamental activation or CCN-CDNC relationship to account for this source of uncertainty. Constraining this uncertainty is a different matter altogether. Current observations of aerosol-cloud interactions from satellite-based platforms are critically limited since both fields cannot be measured simultaneously. But even if they could overcome this issue, satellite observations can not currently constrain the aerosol size distribution—a

lynchpin in setting the sensitivity of cloud droplet number to aerosol perturbations. New in situ data from aircraft sampling clouds in regimes across the globe could play a key role in addressing this limitation. In particular, future observation campaigns should continue to focus on measuring CCN spectra around clouds, as well as sampling CDNC in both nascent and mature cloud formations.

In order to better understand contemporary and account for future climate change, the community must continue to seek constraints on the aerosol indirect effect. Although epistemic uncertainty due to unknown pre-industrial emissions complicates this task, the role of droplet activation illustrated in this work highlights an additional path that the community may explore to provide indirect or emergent constraints on AIE via the basic aerosol-CDNC relationship.

Chapter 5

Aerosol-Cloud Interactions and the Indirect Effect

5.1 Introduction

The physical processes which give rise to the aerosol indirect effect on climate (AIE) tangle amongst themselves to produce a complex, inter-dependent network of pathways whereby perturbations in the aerosol burden alter cloud radiative properties. A reductionist approach is often employed to tackle this problem by isolating individual pathways, subjecting them to rigorous analysis via observations or measurements, and quantifying the uncertainty that arises from unknown or poorly constrained processes or factors. Such an approach was employed in the first two chapters of this dissertation, which focused explicitly on droplet activation. It has also been employed more broadly to incorporate observational constraints on the indirect effect by expanding a high-level expression of AIE into constituent components which encode relationships between clouds and aerosols (Quaas et al., 2008, 2009). Each of these relationships, or aerosol-cloud interactions (ACI), can then be related directly to theoretical considerations or explored through available observational data. Over the past decade, such an approach has provided an alternative, physically-based and empirically-

This chapter is being prepared (along with the previous chapter) for submission to *Atmospheric Chemistry and Physics*

constrained perspective on AIE wholly distinct from the endeavor to capture myriad, intricate complexities of ACI in a consistent fashion within a global modeling framework. However, it has not successfully provided persuasive or rigorous constraints on the magnitude of AIE (Boucher et al., 2013). Much has been written on why this is the case, and it has been postulated that there exists a certain amount of epistemic uncertainty in such an approach arising from our inability to directly measure aerosol emissions and chemistry in the pre-industrial atmosphere (Carslaw et al., 2013), compounded by the idea that constraints we derive using the present day state of the atmosphere may not actually apply to the pre-industrial one (Penner et al., 2011). Furthermore, the relationship between changes in aerosol emissions and resulting indirect forcing are not necessarily linear. Instead, it may be state-dependent, and the magnitude of the indirect effect is likely closely intertwined with climatology of aerosol-cloud interactions in the pre-industrial climate (Carslaw et al., 2013).

While observational or modeling constraints on ACI may not resolve the overarching question of the magnitude of the aerosol indirect effect on climate, they still provide a powerful way to understand which processes impose parametric uncertainty on AIE. Targeting those uncertain processes with renewed efforts in observing the aerosol-cloud climate system, then, could help reduce this parametric uncertainty. This effort would be extremely relevant towards efforts to understand transient climate change, both in the recent and coming decades. For instance, a re-distribution in the emissions of anthropogenic aerosols from the United States and Europe to East and South Asia is potentially associated with strong, regional imprints on surface temperature, an attendant relaxation of the meridional temperature gradient and weakened and expanded circulation in the tropics (Wang et al., 2015). This re-distribution is extremely likely to continue and enhance in the future, but swiftly change character at some point in the future as global efforts to reduce aerosol emissions take over from regionally-disparate prerogatives (Myhre et al., 2013). However, it could also provide new constraints for the magnitude of the indirect effect, especially in light of its non-linear nature. The re-distribution of aerosol emissions since the 1970's coincide with increased observational capacity to measure the climate system; presumably, the change in cloud radiative forcing associated with these changes in aerosol emissions is captured by this data, and could provide some form of global constraint on the indirect effect. However, because of the potential state-dependence of the indirect effect, these constraints may not

be applicable to the pre-industrial to present day change in aerosol emissions.

We may not be able to travel back in time to observe the pre-industrial atmosphere, but we can certainly study the transient changes in clouds, radiation, circulation and precipitation produced by simultaneous changes in the regional emissions of aerosols. Although it does not eliminate the problem entirely (the observations of the atmosphere necessary to constrain do not extend even a few decades into the past), it may be possible to more tightly constrain atmospheric chemistry from the mid 20^{th}-century, lending to improved estimates of ACI and the indirect effect. Resolving the magnitude and diversity of ACI then, and understanding its contribution to the success in global modeling of these transient changes could then be a useful systems-level check on the performance of global models. Improving said performance would undoubtedly yield insight into AIE over the 20^{th} century.

This chapter considers topics pertinent towards this end. In Section 5.2 we review the physical basis for measuring and evaluating ACI and relate it to measures of AIE. Following this development, we assess the importance of the direct coupling between aerosol and cloud droplet number in Sections 5.3 to 5.4, first by considering metrics for evaluating the relationship and then by studying the influence of droplet activation on it and its attendant impact on the indirect effect. We close with a short synopsis and outlook for how this work may be extended, focusing on resolving the discrepancy between modeled and observed ACI.

5.2 Theory

5.2.1 Aerosol-Cloud Interactions in Warm Clouds

Of primary concern here is the problem of how the chain of microphysical changes initiated by a perturbation to the ambient aerosol size distribution influences a cloud's radiative properties—particularly, cloud optical depth, τ_d. For simplicity, we begin by restricting our analysis to warm, stratiform boundary layer clouds. Air parcels entering the cloud at its base will cool adiabatically, driving activation on any aerosol particles which grow, via condensation, larger than some critical radius. The activation process yields a nascent, narrow cloud droplet distribution which continues to grow by quenching any excess water vapor, approximately restricting the air parcels to saturation and following the moist adiabatic lapse rate. This yields an idealized, adiabatic model for the vertical structure of boundary

layer clouds,

$$w(h) = c_w h \tag{5.1}$$

where $w(h)$ is the cloud liquid water content in $\mathrm{kg\,m^{-3}}$, h is the altitude above the base of the cloud (defined as the lifting condensation level or otherwise where activation first initiates) and c_w is the condensation rate in $\mathrm{kg/m^4}$, related to the change in the parcel's saturation vapor pressure by the Clausius-Clapeyron relationship. As a consequence, c_w depends approximately on temperature, but in an adiabatic ascent over the typical depth of a boundary-layer cloud, the change in temperature is small, so we assume it to be a constant.

Strictly speaking, no cloud follows exactly this idealized model, but observations have shown that it closely approximates the structure of maritime boundary layer clouds such that $w(z_{\mathrm{top}})$ is often about 80 % of the purely-adiabatic value (Pawlowska and Brenguier, 2000; Bennartz, 2007). Notably, entrainment of warm, dry air above the inversion capping many boundary layer includes yields mixing which further depresses behavior from this adiabatic model. However, integrating over a cloud depth H, we can write the total cloud liquid water path W in this model as

$$W = \frac{1}{2} c_w H^2 \tag{5.2}$$

Within warm boundary layer clouds, the droplet number N is approximately constant throughout the depth of the cloud and closely related to the ambient aerosol number (Pawlowska and Brenguier, 2000). This follows from the assumption that all cloud droplets are formed at cloud base; condensation only slowly contributes to growth for droplets above a few microns in size, and collision/coalescence is not efficient for hydrometeors that small (Pruppacher and Klett, 1997).

In order to relate the cloud's structure to its optical properties, we need some description of the droplet size distribution (DSD). For an arbitrary DSD, we can compute a volume-average droplet radius

$$
\begin{aligned}
r_v^3 &= \frac{1}{N} \int_0^\infty r^3 N(r)\, dr \\
&= \frac{3}{4} \frac{W}{N \pi \rho_l}
\end{aligned} \tag{5.3}
$$

Since $W = w(h)$, This directly implies a vertical structure in r_v,

$$
\begin{aligned}
r_v(h) &= \left(\frac{3}{4} \frac{w(h)}{N \pi \rho_l} \right)^{1/3} \\
&= \left(\frac{3}{4} \frac{c_w}{\pi \rho_l} \right)^{1/3} h^{1/3}\, N^{-1/3}
\end{aligned} \tag{5.4}
$$

Precipitation processes such as the formation of drizzle lead to gross departures from these idealized relationships by reducing both liquid water content and droplet number, fundamentally altering the DSD and the dynamical lifecycle of the cloud. For a detailed discussion of the influence of drizzle and its relationship to H and N, we refer the reader to Bennartz (2007) and citations therein.

The vertical microstructure of the adiabatic cloud allows us to predict its optical properties. For instance, the cloud optical depth τ_d can be computed

$$
\tau_d = \int_0^H \int_0^\infty Q\, \pi r^2 N(r)\, dr\, dz \tag{5.5}
$$

where Q is the extinction efficiency. Although $Q = Q(r/\lambda)$ where r is the size of a droplet and λ is the wavelength of radiation interacting with the droplet, for solar radiation and the range of droplet sizes we consider, $Q = 2$ is a good approximation (Twomey, 1977). Then, integrating Equation (5.5) by assuming some vertically-weighted mean size distribution, $\tilde{r}$, yields the simplified cloud extinction equation

$$
\tau_d = 2\pi N \tilde{r}^3 H \tag{5.6}
$$

which can be combined with a vertically-weighted mean version of Equation (5.3) to relate liquid water path to the droplet size distribution

$$W = \frac{2}{3}\rho_L \tau_d \frac{\tilde{r}^3}{\tilde{r}^2} = \frac{2}{3}\rho_L \tau_d r_{\text{eff}} \tag{5.7}$$

where r_{eff} is the droplet effective radius, defined as the ratio of the third and second moments of the DSD:

$$r_{\text{eff}} = \frac{\overline{r}^3}{\overline{r}^2} \tag{5.8}$$

Using the volume-mean radius previously derived, the droplet effective radius can alternatively be expressed as $r_{\text{eff}} = r_v k^{-1/3}$ for some droplet dispersion constant k, which is less than unity. However, we can explicitly incorporate the vertically-varying microstructure of our adiabatic cloud model from Equation (5.4) in the integral and show that

$$\tau_d = \frac{3}{5}\pi Q \left(\frac{3}{4}\frac{c_w}{\pi \rho_L}\right)^{2/3} k^{1/3}\, N^{1/3}\, H^{5/3} \tag{5.9}$$

From these equations, relationships between cloud microphysical and optical properties are readily apparent, such as

$$\tau_d \propto N^{1/3}$$

and

$$r_{\text{eff}} \propto \frac{W}{\tau_d}$$

Following Twomey (1977), Feingold et al. (2001) extended this analysis to account for the direct influence of aerosol on droplet number by using a power-law relationship, $N_d \propto N_a^{a_1}$. Within this extended framework, we define "aerosol-cloud interactions" (ACI) as the sensitivity of cloud microphysical or optical properties to perturbations in some given aerosol metric, α:

$$\text{ACI}_{\tau_d} = \left.\frac{\partial \ln \tau_d}{\partial \ln \alpha}\right|_W \tag{5.10a}$$

$$\text{ACI}_{r_{\text{eff}}} = \left.\frac{\partial \ln r_{\text{eff}}}{\partial \ln \alpha}\right|_W \tag{5.10b}$$

$$\text{ACI}_N = \frac{d \ln N}{d \ln \alpha} \tag{5.10c}$$

Note that ACI_N need not calculated for constant W as we previously assumed that

activation alone dictated N and the boundary layer thickness set W. Aerosol can be parameterized in the ACI formulation a number of different ways, for instance using the aerosol optical depth (AOD; e.g. Feingold et al., 2001, 2003; Quaas et al., 2006; McComiskey and Feingold, 2008; Quaas et al., 2009; Penner et al., 2011, among many others) or a direct measure of the aerosol size distribution such as CCN at a given supersaturation. In the adiabatic cloud model, it can be further shown that these ACI metrics are algebraically related,

$$\mathrm{ACI}_{\tau_d} = -\mathrm{ACI}_{r_{\mathrm{eff}}} = \frac{1}{3}\mathrm{ACI}_N \tag{5.11}$$

The power-law relationship $N_d \propto N_a^{a_1}$ implies that $0 < \mathrm{ACI}_{\tau_d} < 1/3$ since a_1 must be no greater than unity. However, observational analyses using a variety of different measurement approaches produce a much larger range of estimates, including many values which exceed the limits placed by the adiabatic cloud model (Table 1 McComiskey and Feingold, 2008). Additionally, global modeling studies produce even larger diversity in estimates of ACI, either due to inter-model variance (Quaas et al., 2009) or regional and seasonal variability in aerosol and meteorology fields in a given model (Penner et al., 2011).

There are many explanations for the divergence between observational and model-derived estimates of ACI, particularly ACI_N. From the perspective of dynamics, the adiabatic cloud model is only suitable for a small portion of the total diversity of cloud regimes encountered in nature. As previously noted, precipitation processes quickly disrupt the assumptions invoked about the vertical structure of the droplet size distribution function and act to sink liquid water content. Importantly, droplet activation is not simply a function of the aerosol number concentration, but also depends on the aerosol size distribution and the strength of the updraft feeding a given cloud. Feingold et al. (2001) expanded the derivative in ACI_N to account for the contribution of the size, location and shape parameters of an idealized, single lognormal mode aerosol size distribution, but further advances in the modeling of droplet activation can account for a much larger and more realistic aerosol parameter space.

There is also a critical scale problem in deriving these quantitative estimates of ACI (McComiskey and Feingold, 2012). Equations (5.10a) to (5.10c) as written directly describe a process-level relationship. Aggregations of the terms in the numerator or denominator of each equation over large spatial scales, as done in satellite retrievals, may not be the

most appropriate, particularly if they neglect to account for the heterogeneity of clouds in a given region or scene, and could lead to strong bias in dampening the estimate of ACI. Furthermore, at large spatial scales, the meaning of each term fundamentally changes; considering larger scenes with diverse cloud types and regimes implicitly includes additional microphysical processes which act to disrupt the idealized relationships derived here between aerosol and cloud microphysical and radiative properties.

5.2.2 Relation with the Aerosol Indirect Effect

Throughout this discussion, we have explicitly avoiding using the phrase "aerosol indirect effect" (AIE). In the literature, definitions of AIE and ACI are sometimes muddied or used interchangeably. For instance, in the parlance of the IPCC, AIE is related to the effective radiative forcing (ERF) *arising* from aerosol perturbations to cloud properties—the definition of ACI which we have adopted here. The analytical device recently employed by Ghan et al. (2016) makes explicit this relationship:

$$\frac{d\ln\overline{R}}{d\ln\overline{E}} = \frac{d\ln\overline{R}}{d\ln\overline{N_d}}\frac{d\ln\overline{N_d}}{d\ln\overline{CCN}}\frac{d\ln\overline{CCN}}{d\ln\overline{E}} \tag{5.12}$$

where the overbar indicates a quantity time-averaged over a sufficiently long period that clouds have adjusted to a perturbation in the aerosol. The lefthand side of Equation (5.12) is a statement of AIE (given as a radiative forcing, R) normalized by the aerosol perturbation (in emissions, E); the righthand side expresses a chain of interdependencies which exposes how cloud macro and microphysical processes are influenced by the aerosol perturbation and ultimately produce a radiative forcing. Although Ghan et al. (2016) evaluated each term in Equation (5.12) by studying differences in its time- and spatial-averaged field between present day and pre-industrial emissions scenarios, one might instead consider retaining the equation's differential form. In that case, we would ultimately apply an analysis which closely parallels the framework for studying ACI that was developed at the beginning of this section.

Critically, this is the underlying framework employed in deriving observational estimates of the indirect effect. Quaas et al. (2008) explicitly wrote the short-wave radiative forcing due to anthropogenic aerosols as

$$\Delta F = \Delta\tau_d \times \overline{F}^{\downarrow} \tag{5.13}$$

where $\Delta\tau_d \propto \mathrm{ACI}_N$ from before, which is further expanded to include other ACI metrics such as the influence of aerosol on N_d, W, and cloud fraction. Each relationship of the form $d\ln x/d\ln\alpha$ involved in the resulting set of equations can be derived from suitable satellite observations and used to construct, with uncertainty, a value for the (first) aerosol indirect effect. Other secondary indirect effects, such as the cloud lifetime effect, can be further derived in a similar manner (see Equations 9-11 of Quaas et al. (2008) for more details).

Both Equation (5.12) and Equation (5.13) encode the same idea, using different levels of abstraction to tweak the involved differentials into forms which can be quantitatively evaluated using available observational or modeling data. Of course, therein lies the difficulty in constraining the indirect effect: the many different terms in expansions following from Equation (5.12) and Equation (5.13) yield a complex analysis requiring many different observations, potentially made using different instruments or techniques on different temporal and spatial scales. One could seemingly add endless complexity to these relationships, raising the specter of interacting or compensating errors. To circumvent this problem, Feingold et al. (2016) recently recommended taking a higher-level approach, focusing on cloud field properties (such as albedo and cloud fraction) that may be closely related to cloud-radiative effects but easier to measure. For instance, Feingold et al. (2016) developed an example of using process modeling to study the relationship between the albedo of a given mixed cloudy/clear-sky scene to the cloud fraction with that scene. Fundamental, ACI-like relationships are implicitly contained in such a problem formulation, but they are balanced with other perspectives which may yield emergent phenomena which could provide additional, indirect constraints on the aerosol-cloud-climate system amenable to simple modeling and theoretical insight.

In the remaining sections of this chapter, we will eschew this emerging perspective on the problem in favor of the traditional one. We will attempt to add detail on the factors that can influence ACI, particularly focusing on droplet activation and its role in setting ACI_N. Through this effort, we will seek ways to constrain ACI_N for regimes where it cannot be directly measured, such as in the pre-industrial climate system. Ultimately, we will link our assessment of ACI_N back to uncertainty in AIE.

5.3 Analysis of ACI_N

5.3.1 Choice of Aerosol Proxy on ACI_N Constraints

Applying the previously discussed ACI formulas to either model- or satellite-derived data requires a subjective choice about how to describe the presence of aerosol. In Figure 5-1 we show scatterplots of the relationship between aerosol optical depth (AOD), cloud condensation nuclei concentration (CCN), and cloud droplet number concentration (CDNC). AOD is commonly used in the literature, particularly in applications which employ comparisons with satellite-based measurements since it is readily computable from these retrievals. But according to Equation (5.5), AOD should not necessarily correlate with the key factors of the aerosol size distribution relevant for influencing droplet nucleation.

Neglecting the influence of the vertical variation in aerosol (which could de-couple the aerosol from the cloud formation process, for instance if there is an aerosol layer aloft, above a boundary layer cloud), AOD is strongly influenced by the aerosol surface area distribution. However, this influence is skewed, as particles much smaller than the wavelength of incident solar radiation are inefficient scatterers and thus contribute little to estimates of AOD. On the other hand, the contribution of particles much larger than these wavelengths is weighted much more heavily via Q, which tends to a value of 2 in the limit of geometric optics. Although Andreae (2009) showed a correlation between AOD and CCN for in-situ measurements, they note that in any given circumstance there could be up to a factor of 4 deviation between the relationship. At least one other work has called into doubt the use of AOD as a predictor of CDNC altogether (Romakkaniemi et al., 2012), although it used a very limited set of observations.

As Figure 5-1 illustrates, within MARC there are strong but highly uncertain relationships between AOD, CCN, and CDNC (ACI_N). Panels (a) and (d) show correlations between AOD and CCN over both land and ocean regimes, respectively, with the land being much stronger in both the present day (PD) and pre-industrial (PI) cases. The shift between emissions cases leads to small shift in the distribution of AOD towards larger values, but a much greater one in the CCN. This is not surprising; anthropogenic emissions lead to much larger number burdens in the accumulation mode sulfate, which contains particles mostly large enough to be captured by the CCN metric used here (diagnosed at a $S = 0.3\%$). However, the variance in the AOD-CCN relationship is extremely large. Over land, CCN

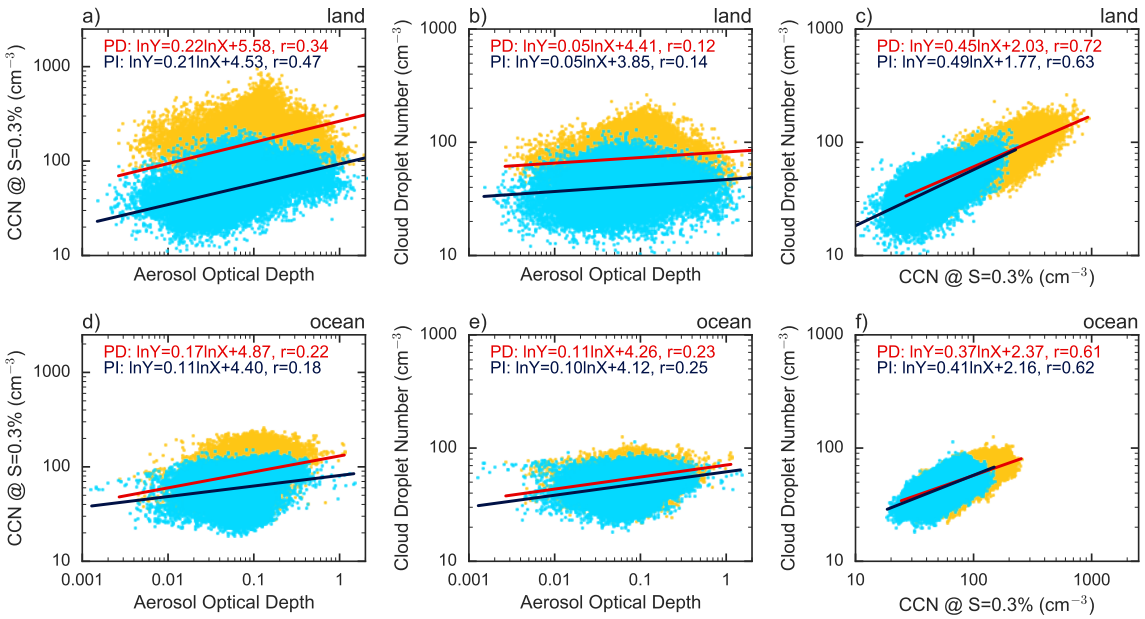

Figure 5-1: Relationship between AOD and CCN (a), AOD and CDNC (b) and CCN and CDNC (c) for a reference MARC simulation, across geographical regions and seasons. Yellow and blue markers correspond to present day and pre-industrial emissions cases, respectively. Plotted regression is fit to the logarithm of the data in each plot.

can vary by nearly an order of magnitude for a given AOD measurement. While the variance remains large in either emissions case, the overall AOD-CCN relationship does not generalize between them. The AOD-CCN relationship derived for, say, the PD case would systematically over-predict CCN. Other modeling studies have noted that AOD and CCN are at best weakly correlated in the present-day climate, which raises concerns about the possibility of deriving comprehensive estimates of the global CCN burden (Stier, 2016). The relationships between AOD and CCN discussed here are broadly consistent with those other results, despite using a wholly different aerosol-climate model.

As panel Figure 5-1b further shows, the AOD-CDNC relationship leads to a systematic over-prediction of pre-industrial CDNC, at least over land. In ocean regimes, the problem is not quite as bad. The bias of AOD weighting heavily towards the fewer coarse mode particles captures a fundamental limitation of giant CCN on droplet activation in aerosol-starved clouds, which is expected to occur in areas far removed from anthropogenic pollution, where sea salt aerosol are prevalent (such as in Figure 5-1e). However, CCN is a much stronger predictor of CDNC than AOD as highlighted in Figure 5-1c/f. The linear relationship

between CCN and CDNC explains far more of the observed variance in CDNC, with r^2 ranging from 0.61 to 0.72. The relationships are extremely similar in both the PD and PI cases for either the land or ocean regime, although we caution that the regression slopes are statistically significantly different (outside each other's 95% confidence interval) owing to the difference in domain of the CCN for each case. To first order, CCN sets the value for CDNC, particularly in the clean ocean regime. Thus, it is reasonable to find that CCN performs so much better at predicting CDNC than AOD, since AOD could be influenced by structure in the aerosol size distribution which does not necessarily impose constraints on activation, especially in polluted regimes or regimes with very strong updrafts.

The role of activation in setting CDNC raises an additional possibility, which is further hinted at by the similarity in the PD and PI relationships between CCN-CDNC. Plotted in Figure 5-2 are distributions of relative error in CDNC when using the PD-derived relationships of AOD-CDNC and CCN-CDNC in land and ocean regimes to predict CDNC in the PI case. In the cleaner ocean regimes, the distribution of relative errors are relatively similar, although there is a fatter tail to the AOD-CDNC errors, suggesting that under certain cases, it will more significantly over-predict CDNC on average. But over land there is a stark contrast. Here, and as seen in the scatter plot in Figure 5-1b, the AOD-CDNC relationship breaks down and systematically predicts twice as many CDNC on average. Similar to what Penner et al. (2011) showed, the AOD-CDNC relationship derived for the present-day only poorly constrains the pre-industrial era, which complicates the use of satellite-derived AOD-CDNC relationships in estimating the indirect effect. But oceanic regimes may not be as heavily impacted by this issue, so such estimates focusing on maritime regimes may still be useful.

The potential for using CCN as an aerosol proxy to improve constraints on ACI_N is further highlighted when assessing regional relationships, too, as in Figure 5-3, which depicts relationships for the summer and winter in North America (left two columns) and the spring season in Asia (far-right column). The AOD distributions in both regions and all seasons are much narrower in both the PD and PI cases, and do not appreciably shift between the two, even while CDNC is statistically significantly greater under PD emissions. As a result, the AOD-derived relationship for the PD case would poorly constrain ACI_N for the PI, although such an exercise would be problematic due to the very low r^2 values in each region. Using CCN in deriving ACI_N greatly improves both the statistical significance of the

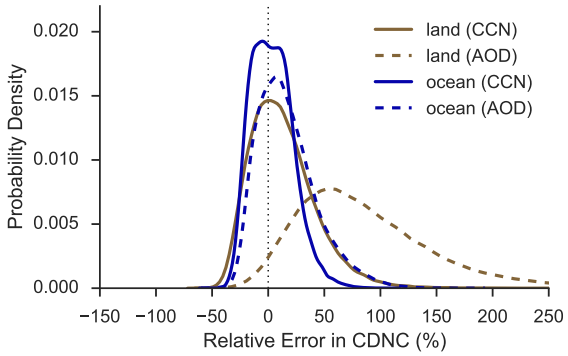

Figure 5-2: Relative error in CDNC predicted for PI case using PD regressions shown in Figure 5-1 for land (blue) and (ocean), using either CCN (solid) or AOD (dashed) as a predictor.

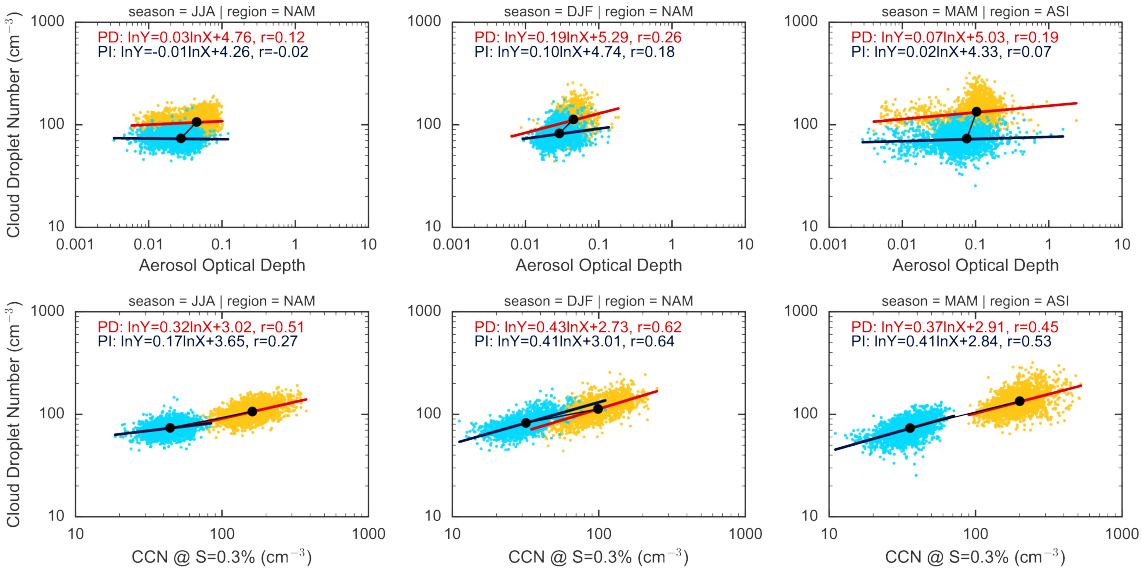

Figure 5-3: Following (Penner et al., 2011), regressions of cloud droplet number against aerosol optical depth (top row) and CCN (bottom row; in lieu of aerosol index) for the pre-industrial (blue/blue) and present-day (yellow/red) simulations, with corresponding slope and intercept overlaid. From left-to-right, the columns show regressions for North America in JJA, North America in DJF, and Asia in MAM. The regressions were fit to the natural logarithm of both the exogenous and endogenous variables.

relationship and its consistency between the PD and PI cases. Even though the individual CCN distributions in both the PD and PI cases do not explicitly overlap except for DJF in North America, the regression slopes and offsets are similar for each.

To more thoroughly explore the role that droplet activation is playing in these diverse estimates of ACI_N, we consider our additional simulations with MARC, conducted using

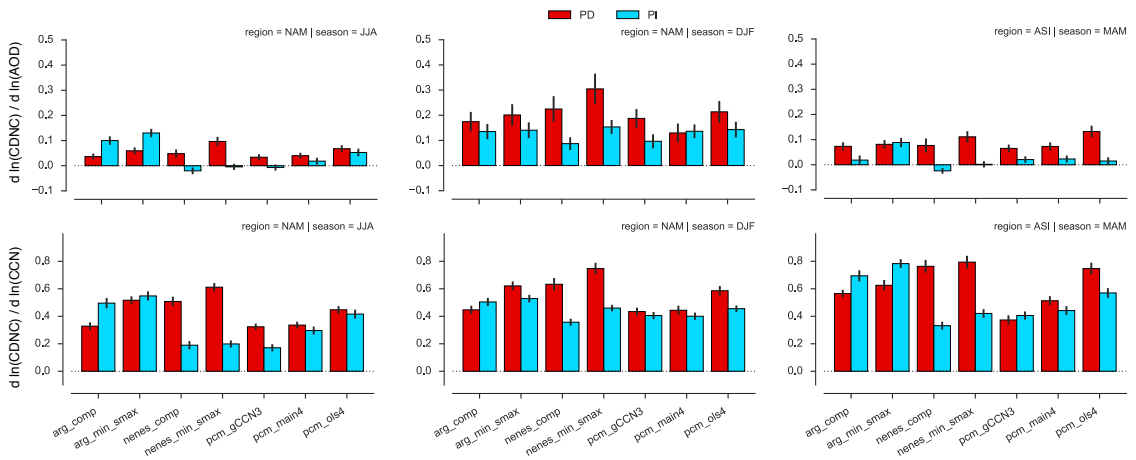

Figure 5-4: Similar to Figure 5-3, but showing uncertainty in the regression slopes computed for each simulation for the PD and PI cases. Uncertainty bars correspond to the 95% confidence interval surrounding the slope estimate.

alternative activation parameterizations. In particular, we supplement Figure 5-3 by repeating the same calculations for these other MARC simulations, and summarize the resulting regression slopes and their uncertainty in Figure 5-4. Using alternative formulations for activation, which modify the fundamental, underlying aerosol-CDNC relationship that ACI_N encodes, adds diversity to our AOD-based regressions, but never offers an example where the PD case agrees with the PI. However, in most cases, the CCN-based regressions still tend to agree. The only cases where this is consistently not the case are the Morales Betancourt and Nenes (2014a)-based schemes (nenes), which yields grossly different PD and PI estimates of ACI_N. This scheme is far more sensitive to aerosol perturbation than the others, though, and these results point to a highly non-linear response in aerosol-CDNC relationships simulated by it. Under the PD emissions scenario, ACI_N derived with this scheme is nearly twice as large when using the PI emissions scenario, something which none of the other schemes exhibit. In fact, most of the other schemes do not produce a statistically-significant different estimate ACI_N given the two different emissions scenarios.

5.3.2 Regional and Seasonal Variation in ACI_N

Because of variations in the dominant sources of aerosol transport and meteorology and the land-ocean contrast demonstrated in the previous section, the different regions depicted in Figure 5-5 should show diversity in their relative ACI_N. In congruence with this notion,

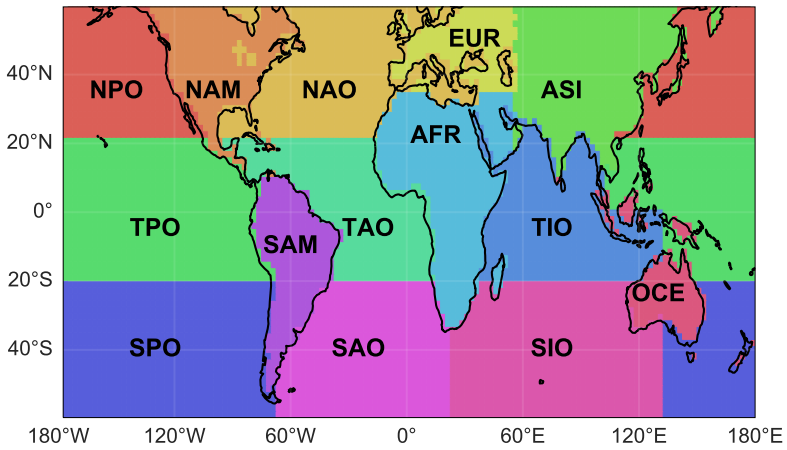

Figure 5-5: Map of geographical regions considered in this study, adapted from Quaas et al. (2009).

several studies have shown a strong regional and seasonal variation in ACI_N (Quaas et al., 2009; Penner et al., 2011; Ma et al., 2014). We add to this discussion with the computed ACI_N shown in Figure 5-6, which shows results using one of the emulator-based activation schemes coupled with MARC. Using AOD as an aerosol proxy, we generally find estimates of ACI_N on the lower end of model-derived values, although these are generally better in line with satellite-derived values. For instance, Quaas et al. (2009) found an average ACI_N of 0.07-0.08 over land regions and of over ocean regions of 0.25 over ocean regions. We find averages of 0.1 over land regions, but just 0.04 over ocean regions using annual data which is significantly smaller than satellite data (Quaas et al., 2009; Ma et al., 2014). However, we note that aggregating, instead, all the ocean and land grid cells (as in Figure 5-1) in lieu of breaking them down into regions produces a result more consistent with satellite observations, with ACI_Ns of 0.05 over land and 0.1 over the ocean. In some individual regions, we actually compute a negative ACI_N, suggesting that CDNC is *decreasing* as the aerosol increase; this is the case in some of the ocean regions (although usually not statistically significant), and over Africa in most seasons, which points to a role for giant CCN particles such as dust or sea salt in this relationship, which may be activation-scheme specific. However, the model may also be systematically under-sensitive to aerosol perturbation in ocean regions, as satellite estimates of ACI_N are expected to be damped relative to their real-world values due to spatial co-variance which is not adequately captures in their calculation depending on the resolution of the data used (McComiskey and Feingold, 2012).

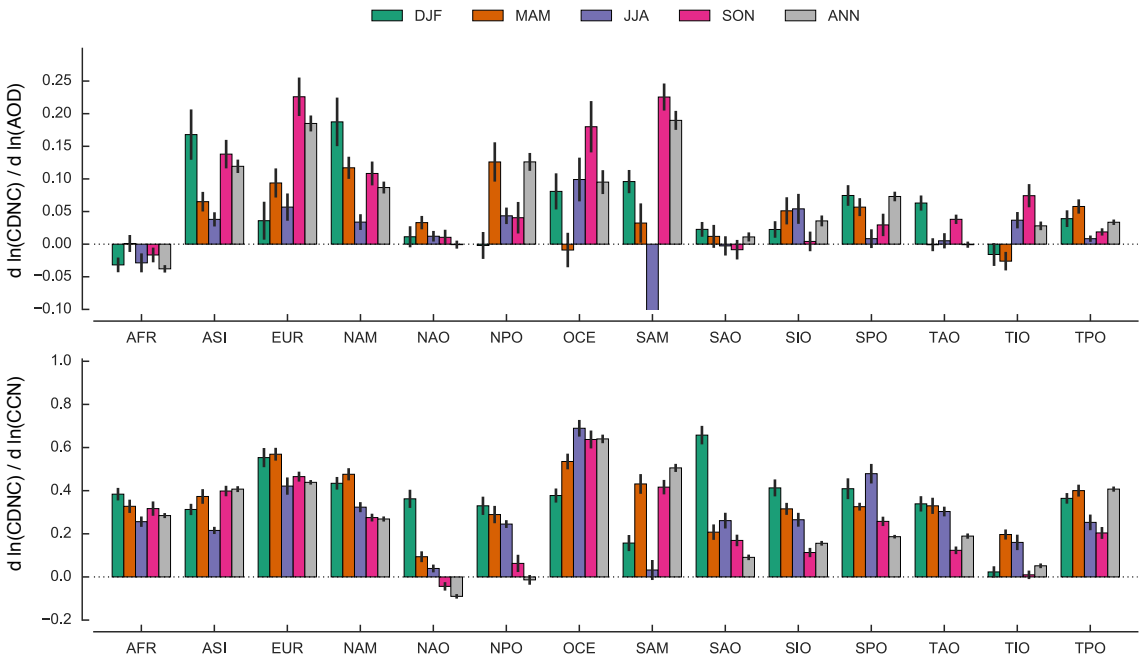

Figure 5-6: Regression slopes between CDNC and AOD (top) and CCN (bottom) for all regions and seasons, computed for the pcm_gCCN3 simulation with present-day emissions. Regressions were computed for the log-transform of each variable using the spatial average timeseries computed for each region. Error-bars indicate the 95% confidence interval on the regression slope.

There is also significant seasonal diversity in the estimates of ACI_N using AOD, although these differences are not systematic or consistent between regions. This is in agreement with previous modeling and satellite studies. Inter-seasonal deviations in the strength of ACI_N point strongly to localized factors which influence the strength of aerosol-cloud interactions, such as regional modes of circulation or meteorology regimes which exhibit different sensitivities to aerosol perturbation. Congruent with this diversity would also be a strong sensitivity to the particular aerosol composition in a given region and season. Regions where the dominant aerosol modes shift from, say, anthropogenic sulfate to natural sea salt in different seasons would be expected to produce a seasonal variation in ACI_N which may strongly co-vary with the seasonal changes in meteorology, complicating such an analysis.

The bottom panel of Figure 5-6 re-computes ACI_N using CCN instead of AOD as a proxy for aerosol. This dramatically changes the analysis, as there is far less regional and seasonal diversity in these estimates. This is particularly the case for land regions, with the notable exception of the boreal summer months in South America (SAM). In nearly all other cases,

there is less variance in ACI_N. Ocean regions generally have smaller ACI_N when using CCN than their land counterparts, and have larger inter-seasonal variability. In only one region do we compute negative ACI_N when using CCN—the northern Atlantic ocean (NAO) in the boreal fall. Negative values of ACI_N occur mostly in regions with a preponderance of a giant CCN, such as desert dust in AFR and sea salt in TIO, NAO, and NPO. In particular, there is a negative relationship between CDNC and CCN in NAO during the Fall which arises from outlier scenarios with very low CCN number, and this causes the annual average to be negative as well.

For reference, we have included the results of regressions applied to every region, season, and activation case coupled with MARC in Figures 5-14 to 5-15 for the PD emissions case. Modifying the underlying aerosol-CDNC relationship by using an alternative activation scheme does not grossly change our assessment of ACI_N, but we do note that there are some regions where activation leads to significant variability. In a handful of regions (SAM, OCE for the land and TIO, TAO, and SPO for the ocean), switching between activation schemes changes the sign of the AOD-CDNC relationship (Figure 5-14). There is not a particular reason why this should be the case, as the inter-scheme spread is not consistent between the regions where the differences occur. The inter-scheme differences to tend to occur consistently despite season, though, suggesting that there is a fundamental bias in how some schemes handle the different aerosol and meteorological scenario in some regions.

The differences persist in Figure 5-15, but the sign of ACI_N never changes from positive to negative no matter what region, season, or activation scheme is considered. That said, there can be very large diversity between schemes for given regions. For instance, in remote marine conditions, the pcm_gCCN3 systematically produces smaller ACI_N in ocean regimes no matter the season. This is most dramatically the case for the TIO region. On the other hand, the arg schemes tend to produce stronger ACI_N in the Southern Ocean regions (SAO, SPO, TIO), especially in the MAM and SON seasons. This suggests that they are more sensitive to changes in smaller-size sulfate particles which may be arising from the background emissions of sulfate predecessors such as DMS. The nenes schemes produce much larger ACI_N in heavily anthropogenic-polluted regions, particularly ASI and NPO, its downwind oceanic region.

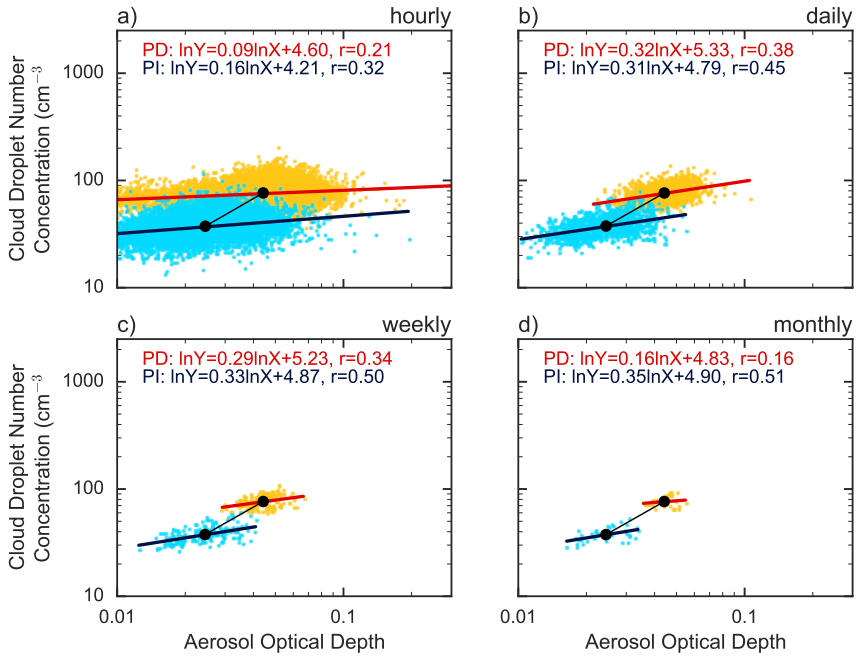

Figure 5-7: Calculation of ACI_N using AOD as the aerosol parameter, and timeseries with hourly instantaneous data (a), daily averages (b), weekly averages (c), and monthly averages (d). Regression formulae for both the PD and PI cases are overlaid on each plot. Note that for the monthly averages, we used an additional two years of simulation data, in order to increase the number of samples available.

5.3.3 Influence of Time Averaging

In the discourse initiated by Penner et al. (2011), Quaas et al. (2011) raised a concern over the influence of the averaging of aerosol and meteorology fields. Satellite fields used in diagnosing ACI usually reflect an instantaneous atmospheric state, and some works in the literature have inconsistently employed time averaging in their comparisons. For instance, time-averaging CDNC fields could substantially decrease variability and mute the signal in ACI.

Figures 5-7 to 5-8 reproduce the analysis in previous sections, but taking progressively coarser time averages over a given regionally-averaged timeseries of aerosol and CDNC. Time-averaging does not change the mean of either AOD, CCN, or CDNC for a given region and season, but it can have a large influence on the relationship between the fields. This is not necessarily the case in Figure 5-7; AOD is relatively stable with respect to averaging. However, the character of ACI_N is heavily influenced by the time averaging when CCN is the aerosol parameter, as in Figure 5-8. Coarser averaging progressively separates

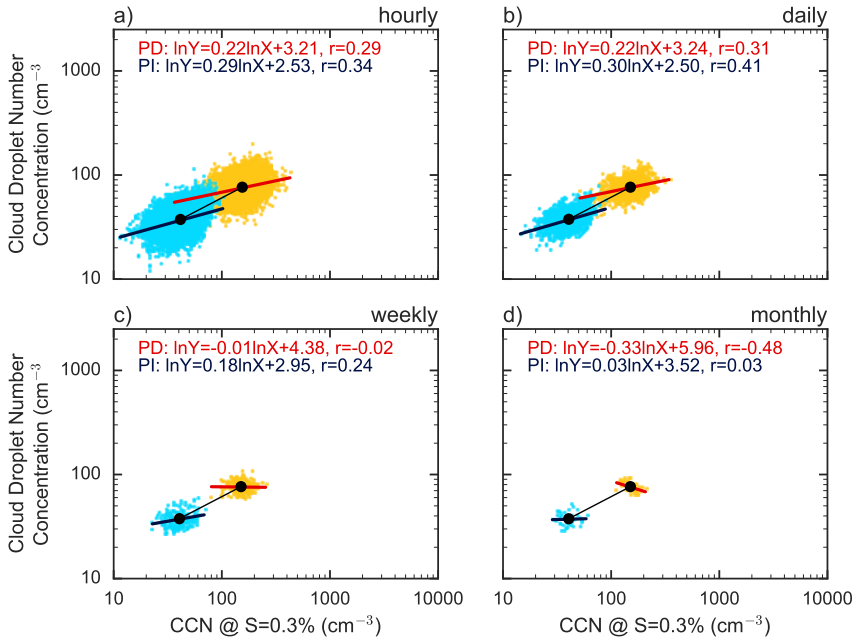

Figure 5-8: Like Figure 5-7, but using CCN as the aerosol parameter.

the distributions of CCN for the PD and PI cases, and also greatly reduces the variance. At weekly and monthly scales, the sign of ACI_N is actually flips, although the statistical significance of the relationship becomes dubious and difficult to detect owing to the very small sample sizes available (we include an additional two years of simulation data in order to boost the statistical analysis). At sub-daily averaging, though, the relationship does not statistically significantly change for the CCN case.

An overview of the influence of time averaging across regions for a given activation scheme, region, and season is shown in Figure 5-9. AOD-derived estimates of ACI_N are highly unstable when averaging over the time-series. Here, monthly averaging dramatically degrades the statistical significance of the regression slopes such that in most regions, they are indistinguishable from zero. The same is true for many other regions. Additionally, estimates using averaged time-series of different resolution are generally not statistically distinguishable from one another. The alternative estimates of ACI_N based on CCN are more resilient to time-averaging. In most regions, daily and monthly values capture within their uncertainty bounds the instantaneous relationship between CCN and CDNC, particularly over land regions. Over the ocean, the monthly averages tend not to produce statistically-significant trends.

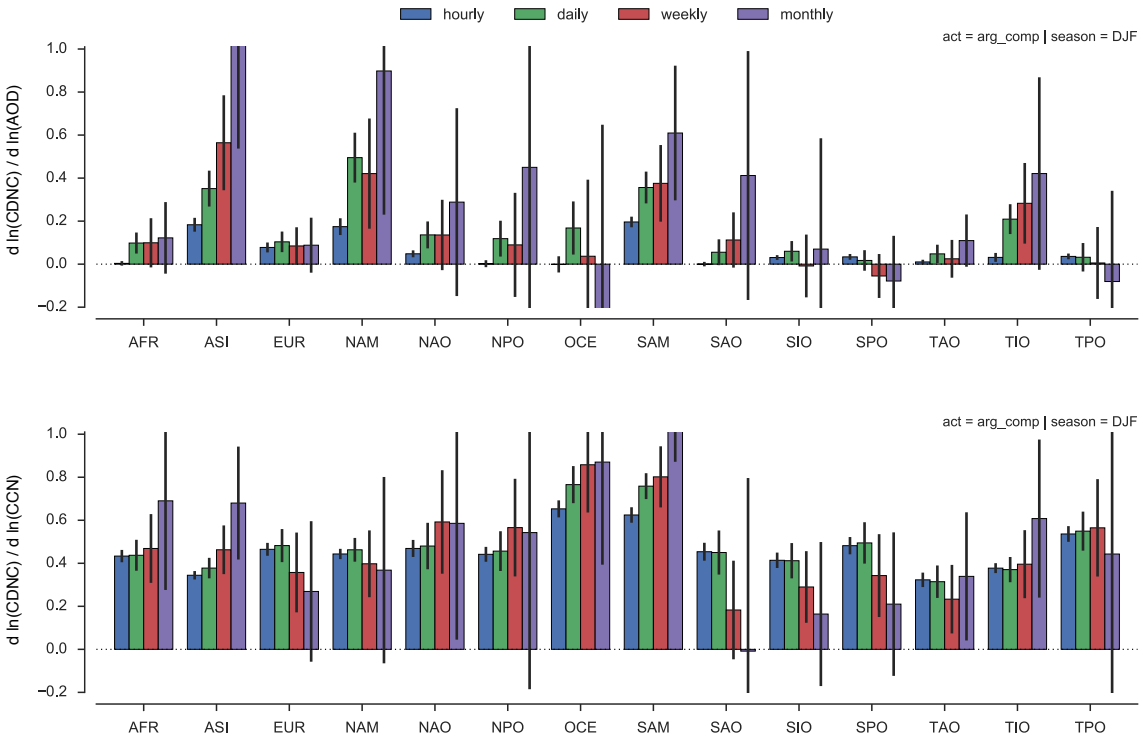

Figure 5-9: Regional comparison of the influence of time-averaging (colors) on the strength of ACI_N for winter DJF season in North America, using the arg_comp MARC simulation. The top panel uses AOD as the independent aerosol parameter; the bottom panel subsitutes CCN. For each bar, an error bar is superimposed which indicates the 95% confidence interval on the regression slope. The monthly average values use an additional two years of simulation data.

5.3.4 Global Variation in ACI_N

In addition to analyzing regionally-averaged timeseries of aerosol and cloud fields, we apply the method of deriving ACI_N on a global scale, to each column in our suite of MARC simulations with differing activation schemes. This produces maps of global variation in ACI_N for both the pre-industrial and present day, from which we can estimate the change in ACI_N over the period of anthropogenic aerosol emissions. The results of these analyses are summarized in Figure 5-16 and Figure 5-17 for ACI_N computed with AOD and CCN, respectively.

Under the influence of different activation schemes, the global spatial pattern of ACI_N does not appreciably change. However, regional variation is greatly influenced by the aerosol metric used to calculate ACI_N. Using AOD, we note that ACI_N is generally negative over land areas. This is consistent between the PD and PI cases, but ACI_N tends to become

larger (more positive) under the influence of increased anthropogenic aerosol emissions in the PD case, particularly in the United States and eastern Asia. The picture over the global oceans is much more complex. Near the equator, particularly in regions where the ITCZ is strongest, ACI_N is more often negative than positive. However, this could be spurious; these regions have very little local aerosol input and high precipitation frequency due to active convection which acts to diminish AOD. The enhanced ACI_N over these regions could likely be due to higher CDNC associated with seeding of stratiform clouds with convectively detraining water within the model. Straddling these regions of higher ACI_N are ones with anomalously large (in the global sense) positive ACI_N. This is particularly the case in the southern Atlantic and Indian Oceans.

However, negative values of ACI_N are somewhat misleading. They do not necessarily imply that CDNC robustly decreases with increasing aerosol loading. Instead, they indicate the presence here of strong feedbacks to suppress cloud droplet number. This is particularly the case over land regions, where the CDNC can become large, but is fundamentally limited due to precipitation processes whereas the AOD increases nearly without bound as emissions continue to add to the ambient aerosol number. Furthermore, computing ACI_N using single grid cells introduces far more variance into the statistic than the regionally-averaged approach used in the previous section does, as that tends to damp out differences due to alternative aerosol and cloud regimes.

The difference between the PD and PI cases is more strongly influenced by the choice of activation scheme than either the PD or PI case itself. All the schemes show a consistent increase in ACI_N over the US and China, as well as decrease in southern Africa and South America. They also show increases in ACI_N over the northern Atlantic. Where they differ the most is in the southern Atlantic and Indian Oceans. Some of the simulations indicate a weakening of ACI_N in these regions, whereas others show a more complex pattern similar to a dipole off the coast of South America. This could be associated with circulation changes arising from changes in the aerosol indirect effect in these regions between the two emissions cases; a shift in the distribution of clouds could result in an apparent shift in ACI_N.

Shifting to use CCN as the aerosol proxy in ACI_N produces a similar picture, as in Figure 5-17. The global distribution of ACI_N is more muted in this case, with smaller values in nearly all regions except for Africa. The region of very strong, positive ACI_N in the Atlantic and Indian Oceans is gone here, replaced instead with a zone of relative weaker

or even negative ACI_N. This could relate to the preponderance of sea salt in these regions which might skew the AOD estimate and its relationship to the aerosol size distribution, since it weights aerosol by their cross-sectional area rather than their number as CCN does. There is a general decrease in ACI_N over oceanic regions where it is moderately positive in the pre-industrial case, relative to the present day.

Differences in ACI_N when using CCN between the PD and PI cases are more difficult to understand. Some regions show a very strong decrease (even flipping from positive to negative), such as in the mid-Atlantic and over Eurasia. However, there is a robust increase in ACI_N over central Asia across all schemes, as well as potentially in the Southern Ocean. This is also the case in the northwestern Pacific.

The large changes in ACI_N are generally downwind for regions with either dust export or anthropogenic aerosol emissions. This suggests that the increase in aerosol emissions going towards the present day emissions scenario is influencing the relationship between CDNC and CCN. In particularly, because these regions show a large decrease in ACI_N, it would suggest that the relationship is saturating as the background CDNC becomes larger. However, as we previously noted, these calculations are intrinsically much noisier than the regional averages previously discussed, and totally decouple individual grid cells from larger-scale changes happening in the regions around them.

5.4 ACI and the Indirect Effect

Following the development in Section 5.2, ACI is closely related to the strength of the indirect effect (AIE) by setting the fundamental sensitivity of cloud microphysical fields and consequently their optical properties to aerosol perturbation. Using the regional breakdown from Figure 5-5, we compare the relationship between cloud droplet number and ACI using both AOD and CCN under different activation schemes in Figure 5-10. In the pre-industrial climate, there is a strong anti-correlation between droplet number and ACI over both land and ocean regions. Put in more concrete terms, cloud with lower droplet number concentrations are more sensitive to perturbations in aerosol number. The AOD-based relationships for ACI capture this relationship over the oceans, but with far more uncertainty over land regimes. This is related to the very different aerosol composition over the two regions; except for one outlier, the CDNC-AOD relationship derived over land tends to show very little

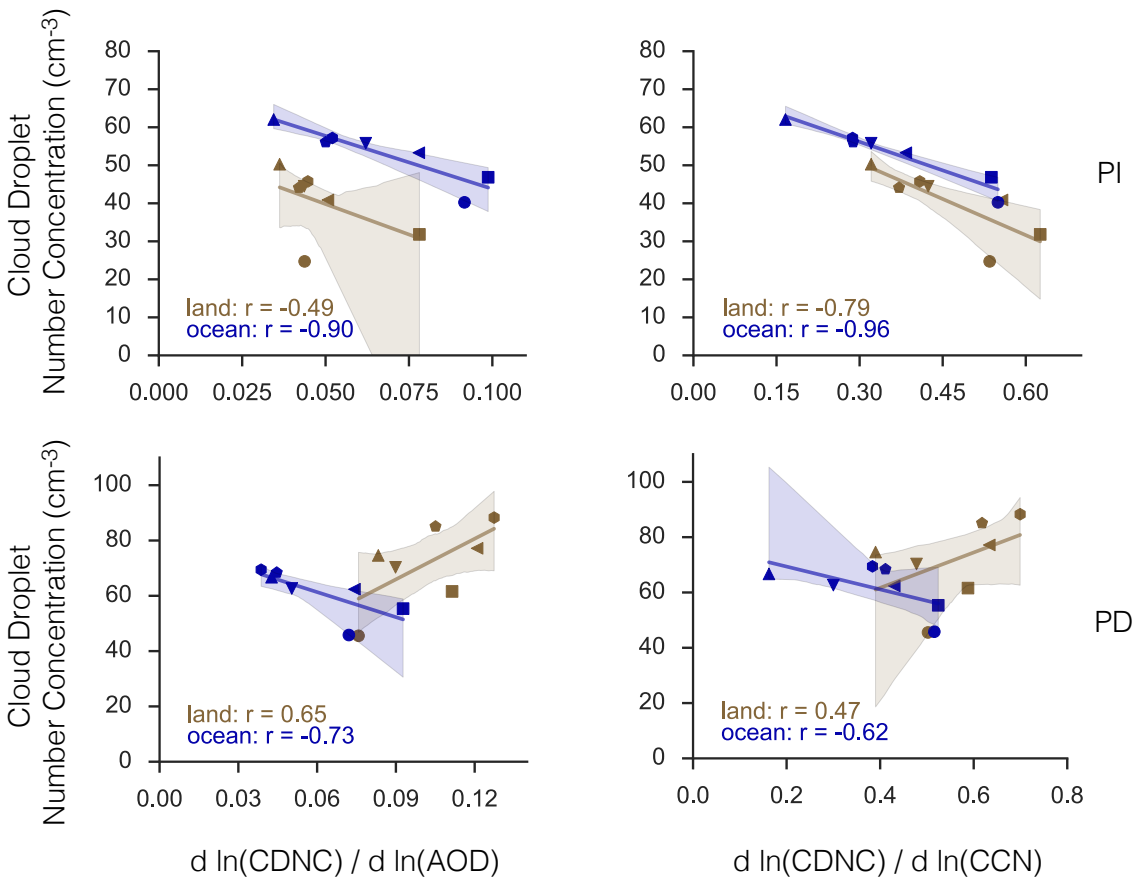

Figure 5-10: Regressions of area-weighted averages of ACI and CDNC computed over land and ocean using regions from Figure 5-5, for each of the different MARC simulations with different activation cases (different glyphs). A 95% confidence interval on the computed regression slope is shown by the shaded area, with the best-fit depicted as a solid line. The top row uses ACI and CDNC computed from the pre-industrial simulations; the bottom roe from the present day.

variability compared to that over the oceans, which spans a range of a factor of four over the smallest estimate.

However, the relationship ACI and CDNC is remarkably robust in the pre-industrial climate simulated by MARC. What differs is the background CDNC levels, or the baseline set by the different activation schemes. The ACI relationships plotted here relate the logarithms of CDNC and either AOD or CCN in a linear fashion:

$$\ln \mathrm{CDNC} = a \ln \mathrm{CCN} + b \tag{5.14}$$

where a is the ACI metric considered here. But this can be re-written as

$$\mathrm{CDNC} = b'\mathrm{CCN}^a \tag{5.15}$$

taking $b' = e^b$. In the formulation of our linear regression, b represents a constant offset between CDNC and CCN. But considering the process of droplet activation alone, b should have to be strictly greater than 0 while a should be less than 1, because there cannot be more droplets than available CCN. If $b > 0$, then for sufficiently small CCN concentrations $\mathrm{CCN} < b'^{-1/(a-1)}$, CDNC will be greater than CCN, which is unphysical considering the activation process alone. To have higher CDNC burdens than CCN, one would need droplet multiplication or enhancement processes. Droplet multiplication could occur if the droplets were large enough to become raindrops which can break apart due to hydrodynamic stresses; alternatively, the CDNC number could increase due to mixing with cloud parcels that already have a higher number concentration.

In the pre-industrial case, fitted estimates of b range from about 0.5 to 3.5, and a ranges from 0.15 to 0.6. Within this range, a typical critical CCN value (below which predicted CDNC is greater than CCN) is about $50\,\mathrm{cm}^{-3}$. The modeled average CCN is generally above this value. However, for any given region or any given instantaneous estimate of CDNC, these relationships may not be the most appropriate, because CCN can be vanishingly small under certain meteorological conditions which may deplete and fail to replenish the ambient aerosol number burden. Physically-based activation parameterizations obviate this problem to some extent, but based on Figure 5-10 do not firmly constrain estimates of ACI arising from simulations using them.

Additionally, estimates of b are very strongly negatively correlated with a with r equal to -0.97 and -1.0 over land and ocean, respectively. Unsurprisingly, b is also strongly correlated with pre-industrial average CDNC burden (r equal to 0.92 and 0.98 over land and ocean). This relationship re-iterates the sensitivity of droplet burden to aerosol perturbation in the pre-industrial atmosphere in way that is closely connected to activation; where activation sets a higher average baseline droplet number concentration, it also yields a relative insensitivity between CDNC and the ambient CCN.

The relationship between background CDNC and ACI is more complicated in the present day climate (bottom panel, Figure 5-10). Using both the AOD and CCN proxies for aerosol, CDNC tends to increase with ACI over the land, but decrease over the ocean. At least over

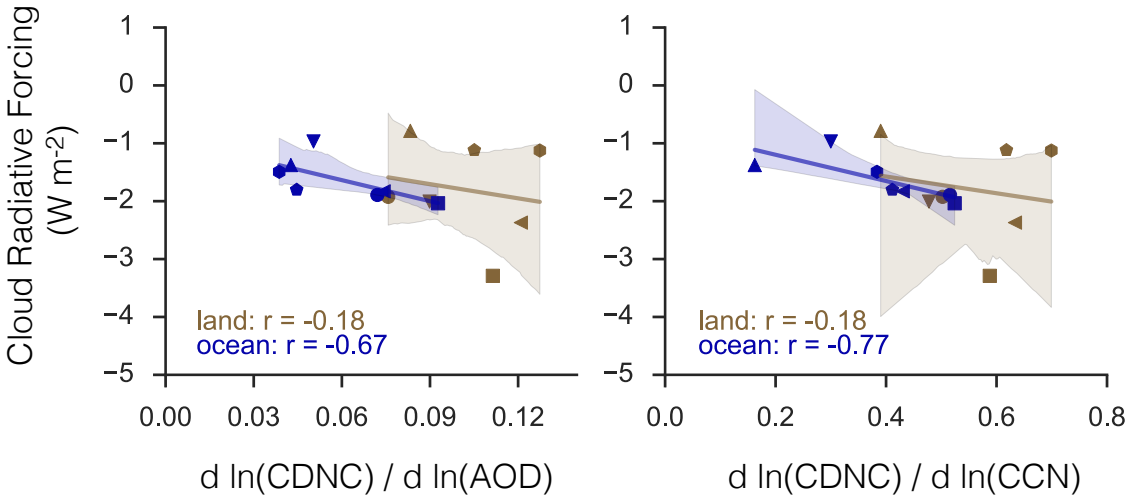

Figure 5-11: Similar to Figure 5-10, but depicting the relationship between ACI and cloud radiative forcing

the ocean, average CDNC does not change much between the pre-industrial and present day simulations, hence the similar relationship with ACI. But this is not the case over the land regions, which show very large increases in CDNC due to anthropogenic aerosols. This does not necessarily mean that the land regions are becoming increasingly sensitive to aerosol perturbations; instead, the increase in available CCN is so large that very high CDNC concentrations are being realized more often and the tail of the CDNC-CCN distribution (toward smaller values for CCN) is not constraining the ensuing relationship very well. In fact, the power-law formulation related CDNC and CCN implicit in this analysis grows unbounded. Whether this can occur in the real atmosphere is a different story. While sufficiently strong updrafts could presumably activate up to the entire aerosol population, this likely does not occur outside of deep convection. Within MARC, the average updraft speed used to activate aerosol in stratiform clouds is on the order tens of centimeters per second, but outside of regimes where a few large particles suppress activation, this unbounded behavior can still be recovered.

Despite this potential unbounded growth in CDNC as CCN rise, it is unlikely that cloud radiative forcing (CRF) would become very large. Processes such as precipitation formation reduce the droplet number concentration and fundamentally limit the maximum values that CDNC can take integrated over the life-cycle of a given cloudy element. In Figure 5-11 we plot the same metrics of ACI but now against the cloud radiative forcing computed as the

difference between the average pre-industrial and present day simulated climates. Over land there is very weak correlation between ACI and CRF due to the large variance in regional CRF, but there is a much stronger direct correlation over ocean regions (the strength of the indirect effect is taken as the absolute magnitude of CRF). Comparing between Figure 5-10 and Figure 5-11, activation cases which produce lower baseline CDNC values tend to have strong indirect effects.

This is consistent with the development of ACI from Section 5.2, which can be extended to relate cloud optical depth and albedo to changes in the droplet number concentration, and ultimately produces a relationship which Platnick and Twomey (1994) named the "albedo susceptibility", or S. Critically, S implies that the unit change in albedo for a given change in droplet number concentration is greater for clouds with lower initial values of CDNC, and furthermore saturates for clouds with an initial albedo close to 1 (purely reflective). Larger changes in albedo could be associated with greater changes in cloud radiative forcing, a proxy for the indirect effect. But the magnitude of those albedo changes depends on the initial cloud droplet number concentration. Regions with higher background CDNC might have a higher cloud albedo and saturate with respect to aerosol changes, but cleaner regimes may yield large radiative forcing changes due to the same aerosol changes. This extends beyond just regions; activation cases which bias the CDNC towards low or high background values invoke the same effect, although they may not accurately reflect the relative sensitivity of indirect effects in a given region. Additionally, this result is not dependent on ones choice of aerosol proxy (AOD or CCN).

To disentangle the relationship between CDNC, radiative forcing, and ACI, Figure 5-12 and Figure 5-13 plot changes in cloud properties averaged over the studied regions versus their corresponding, diagnosed ACI, using AOD and CCN. In neither case are there strong relationships between changes in liquid cloud fraction or liquid water path and ACI; these are driven by the larger scale meteorology in each model that dictate climatologies of clouds. The relative insensitivity here suggests that activation-induced changes in droplet number are not producing very large second aerosol indirect effects (e.g. increasing cloud lifetime). Changes over the ocean are somewhat more strongly related to those over land, but the magnitude of the average changes in the cloud fraction and liquid water path are very small. This is robust to the choice of aerosol proxy in deriving ACI, although using AOD produces a stronger relationship (measured in the unit change of either cloud property per unit change

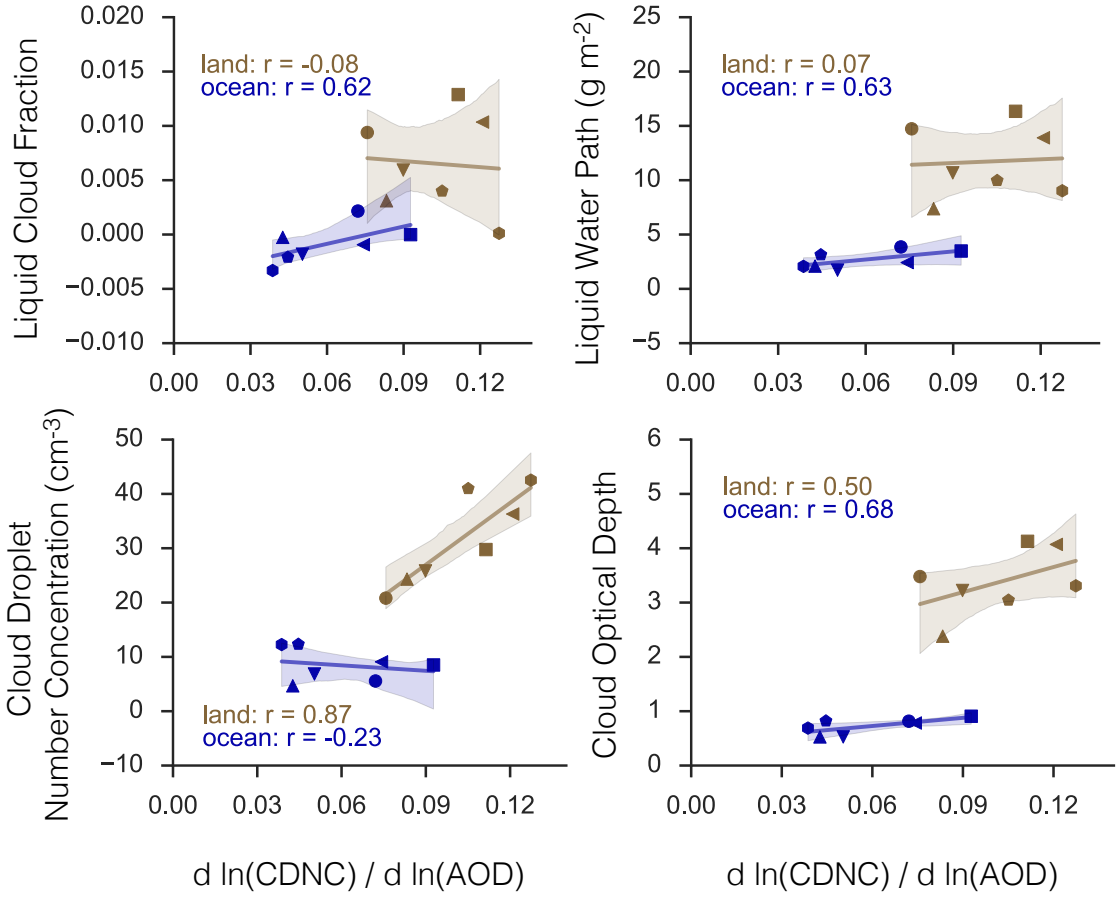

Figure 5-12: Similar to Figure 5-10, but depicting the relationships between ACI and changes in cloud field (liquid cloud fraction, liquid water path, CDNC, and optical depth) between the pre-industrial and present day. All of the ACI estimates in this figure are derived using AOD.

in ACI) because of the reduced various in ACI it entails. There is also very little consistency between the pre-industrial or present day cloud droplet number concentrations and changes in these cloud fields. This is also true of the relationship between cloud optical depth and ACI, although over ocean the two are much more strongly correlated when using CCN to diagnose ACI.

Unsurprisingly, the change in CDNC itself is very strongly related to the metrics of ACI, particularly over land. This is likely due to the previously-discussed potential for unboundedness between CDNC and CCN. Highly elevated background CCN and attendant increases in CDNC skew calculations of ACI over land very high by virtue of over-weighting the large tail of the CDNC-CCN relationship. But activation schemes which enable the

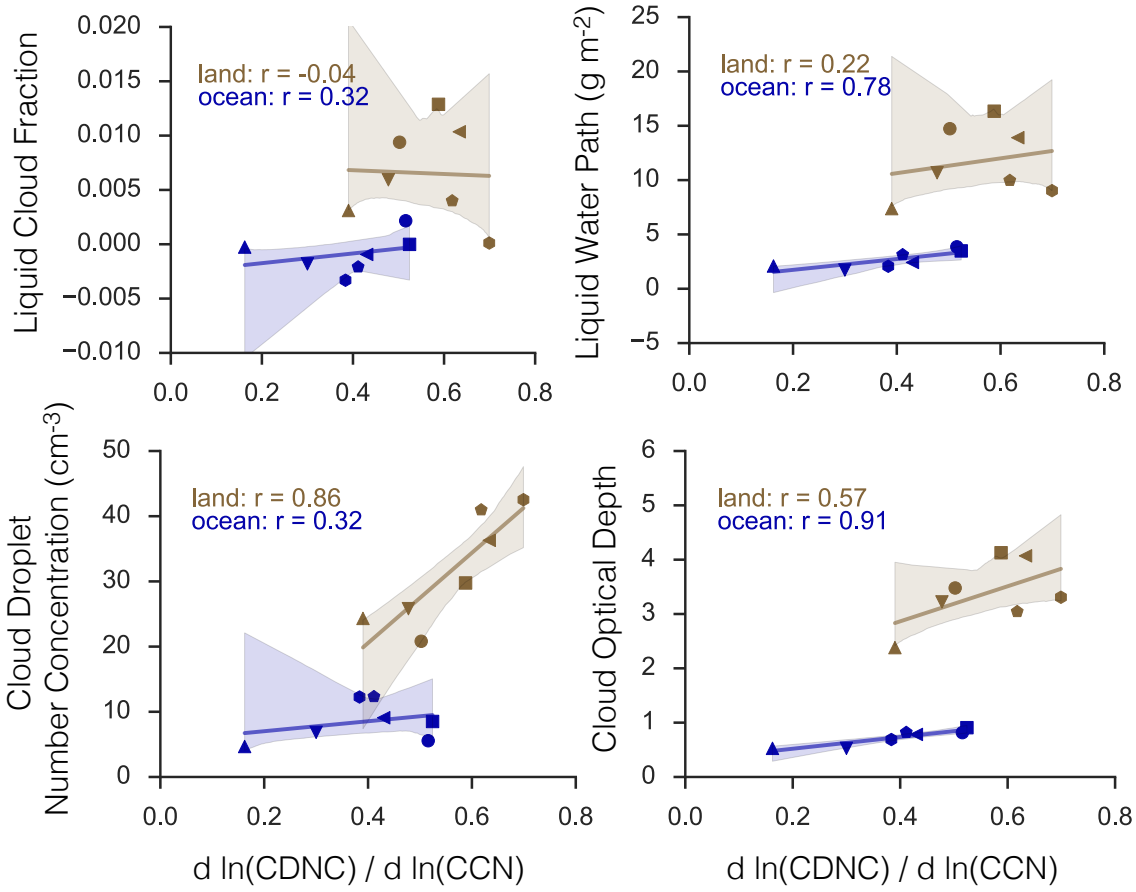

Figure 5-13: Similar to Figure 5-12, but using CCN to derive ACI.

model to explore these regimes are not necessarily subject to producing a very large indirect effect, following Figure 5-11. In fact, there is not a statistically significant relationship between the change in CDNC and the diagnosed change in CRF from those simulations. What is occurring in the ocean regimes, then, is more illuminating. In the ocean regions, the change in CDNC between pre-industrial and present day is both small and independent of ACI. But, over these regions we see about $1\,\mathrm{W\,m^{-2}}$ difference in diagnosed indirect effect.

5.5 Conclusions and Outlook

Here, we have considered a bottom-up approach to understanding the influence of activation on simple models of aerosol-cloud interactions and the indirect effect. To do so, we considered a part of the chain of microphysical processes which must respond to aerosol perturbations

in order to ultimately modify cloud optical properties and yield a radiative forcing. For suitably simple cloud systems, analytical expressions can be derived to show relationships between the sensitivity of various cloud properties to aerosol perturbations. These measures of aerosol-cloud interactions (ACI) can then be assessed using more sophisticated modeling tools or with certain kinds of satellite data. Ultimately, they can be linked directly to the magnitude of the indirect effect (AIE) and provide a toolkit for analyzing the sources of its uncertainty and sensitivity.

Deploying satellite data to assess ACI and AIE is fraught with difficulty. In particular, satellite retrievals do not wholly constrain the ambient aerosol size distribution, do not resolve co-located aerosol and cloud properties. Satellite retrievals of aerosol properties are limited to metrics related to the aerosol optical depth. As Equation (5.5) shows, optical depth is related to the surface area distribution of a given aerosol population, which is heavily weighted towards larger particles. However, droplet activation is much more strongly related to total number concentration of particles and the number size distribution. Thus, AOD-based metrics of ACI may provide a biased perspective of the potential activation dynamics for a given aerosol population. Using metrics like aerosol index have been proposed to account for this inconsistency between AOD and droplet activation potential, as they more heavily weight the presence of smaller particles (Penner et al., 2011).

We demonstrated this using output from an aerosol-climate model run with aerosol emissions scenarios for both the pre-industrial and present day. In most regions, the transition between these two emissions scenarios produces a step change in the distribution of AOD. As a result, even if the relationship between CDNC and AOD remains constant between the two emissions scenarios, one would not be able to use the PD-derived relationship to constrain CDNC in the pre-industrial era. Using CCN-based aerosol metrics alleviates this problem. CCN is related to CDNC explicitly by a physics module (an activation parameterization) in the global model used here, and CDNC shows similar sensitivity to CCN in both the pre-industrial and present day simulations. The present day emissions scenario dramatically increases the CCN burden, particularly in continental regimes, but in a smooth fashion that overlaps with its pre-industrial distribution; there is no abrupt shift like we see in the AOD distribution.

The activation scheme which fundamentally links CCN to CDNC in our global model remains the same in both the pre-industrial and present day emissions scenarios, and thus

it is not surprising that ACI_N is similar when using CCN. However, we stress two important caveats. First, the aerosol composition is significantly changing between the different emissions scenarios. Each activation scheme considered here has its own idiosyncrasies, and some schemes may have greater biases in either emissions scenario. The sensitivity of the CDNC-CCN relationship can thus shift on a regional basis, as new aerosol sources begin to dominate the total size distribution. But more importantly, activation alone does not solely account for CDNC in the relationships modeled here. Other processes act to reduce CDNC and modify other cloud microphysical properties, for instance precipitation. These processes arise in the analysis of ACI_N as the conditional distribution of CDNC for a given level of AOD or CCN. The analysis here does not separate out meteorological versus aerosol contributions to variation in CDNC, but other work has at least considered the context of these relationships in different dynamical regimes and shown that aerosol perturbations in regions with the most large-scale precipitation contribute the most to the indirect effect (Zhang et al., 2016).

The consistency of the ACI_N derived using CCN in both PD and PI emissions scenarios suggests an important role for *in situ* observations to constrain the indirect effect. Present day observations of CDNC and CCN measured across regimes would help validate the performance of present-day models. Unfortunately, CCN is not easy (if not impossible) to measure from satellite-based instrument packages using current methods, and *in situ* measurements from, say, aircraft campaigns are expensive and scarce. Nevertheless, having such data would provide powerful constraints on the behavior of the aerosol-clouds-climate system in the present day, and potentially that of the pre-industrial. If satellite methods are not viable for this purpose, then there may be a path forward using atmospheric composition re-analysis (e.g. Inness et al., 2013) to produce highly detailed, accurate assessments of the ambient aerosol size distribution for use in conjunction with satellite-derived estimates of cloud properties.

Additionally, the strong correlation between ACI and the indirect effect under the influence of different activation schemes demonstrates the importance of critically evaluating the climatologies of cloud properties simulated by global models. There is no *a priori* reason to choose a particular set of physical parameterizations (such as activation) over another. However, the interplay between the simulated aerosol size distribution (which may differ dramatically between different models even while they simulate similar bulk mass and number

concentrations) and the activation scheme dictates both the baseline cloud microphysical state and its sensitivity to aerosol perturbation. Although most modern models do not place unphysical restrictions on cloud droplet number minima (as Hoose et al. (2009) investigated), the subjective choice of physics components and their linkage accomplishes the same imposition. These implicit restrictions could explain much of the variance in modeled indirect effects between models, and metrics including ACI might provide a way to sort through and better compare them.

As a result, we recommend extending work such as Quaas et al. (2009) and Ma et al. (2014) to rigorously quantify the sensitivity of aerosol-cloud interactions in global aerosol-climate models. With the development of new, CCN-based constraints on ACI, such efforts could provide novel ways to constrain global model estimates of the indirect effect.

5.6 Appendix A: Activation-ACI Summary Figures

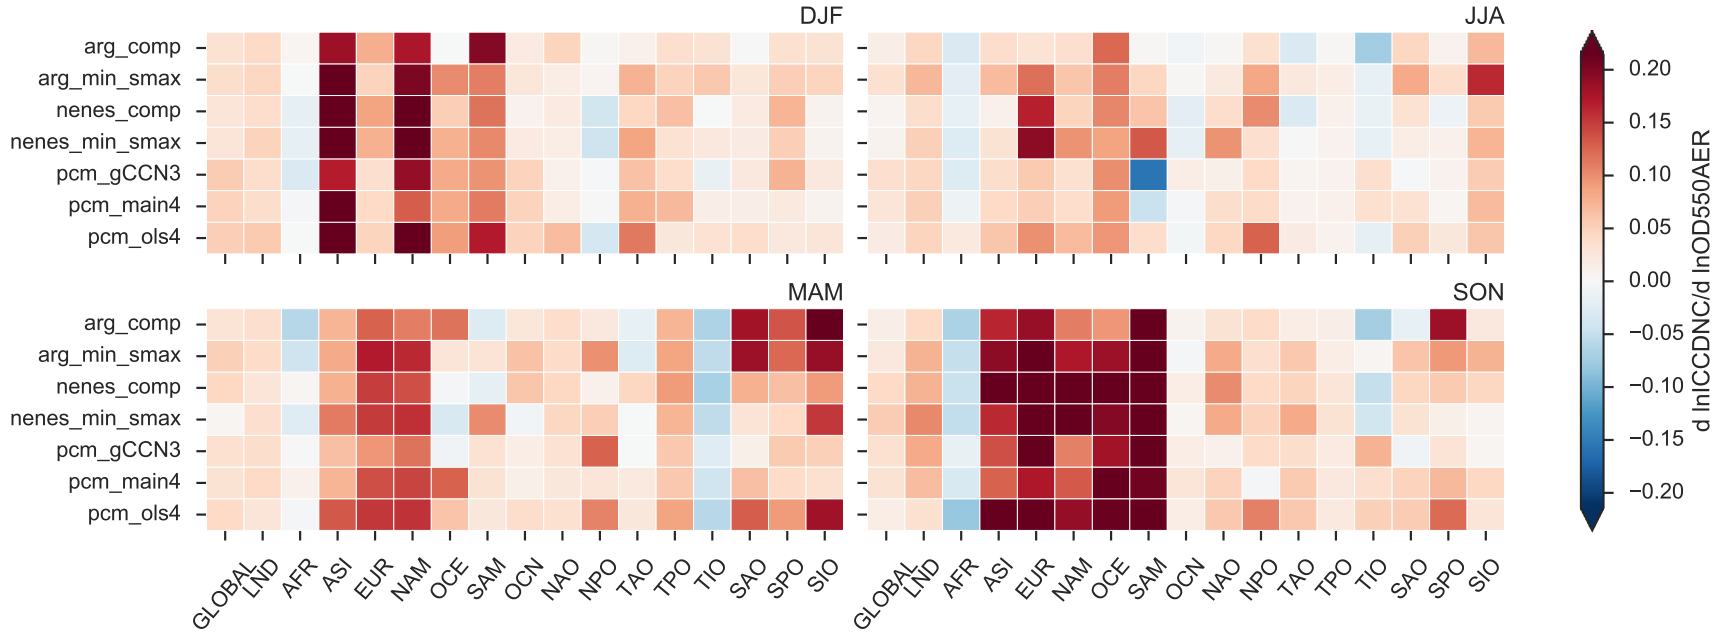

Figure 5-14: Heatmap depicting regression slopes between CDNC and AOD for all regions and seasons for all simulations. Regression slopes were computed using the logarithm of each variable.

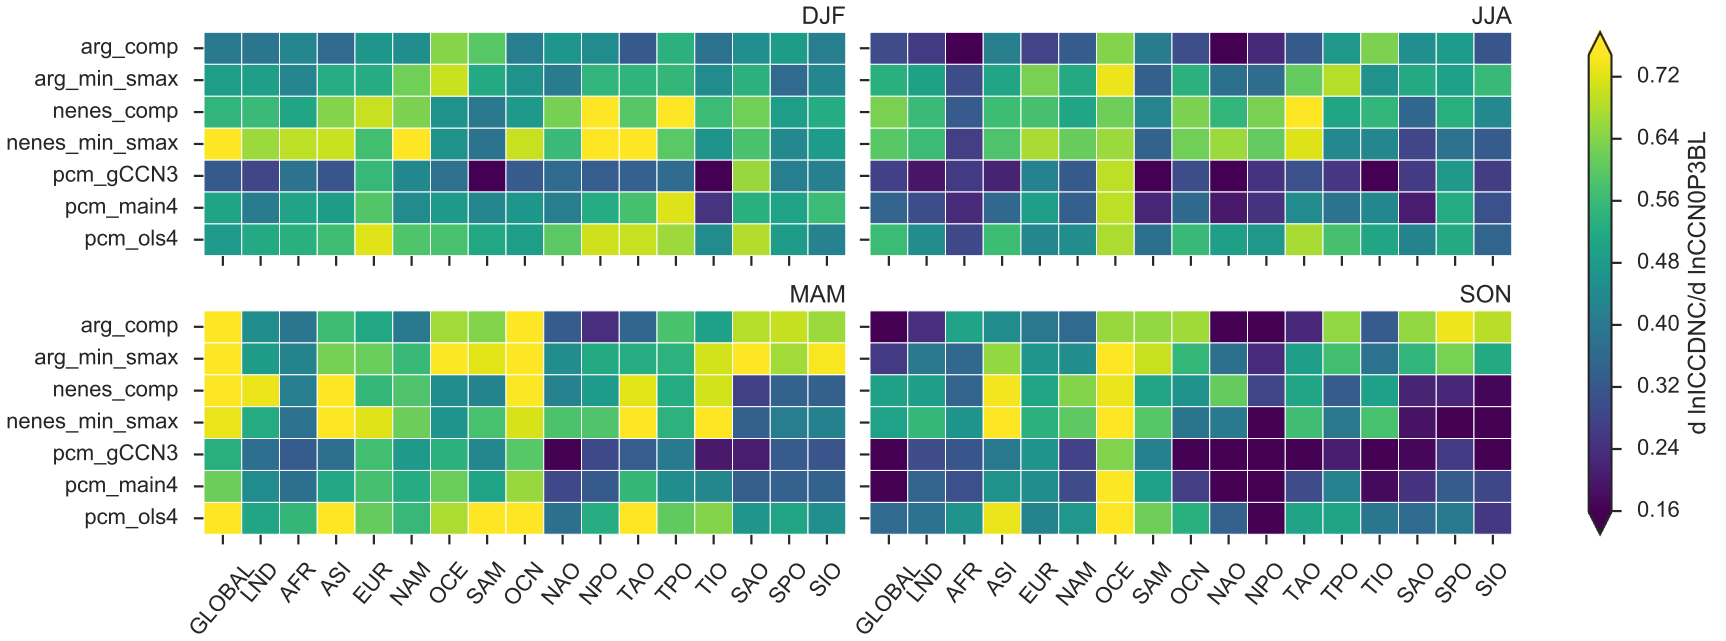

Figure 5-15: Same as Figure 5-14 but showing the relationship between CDNC and CCN.

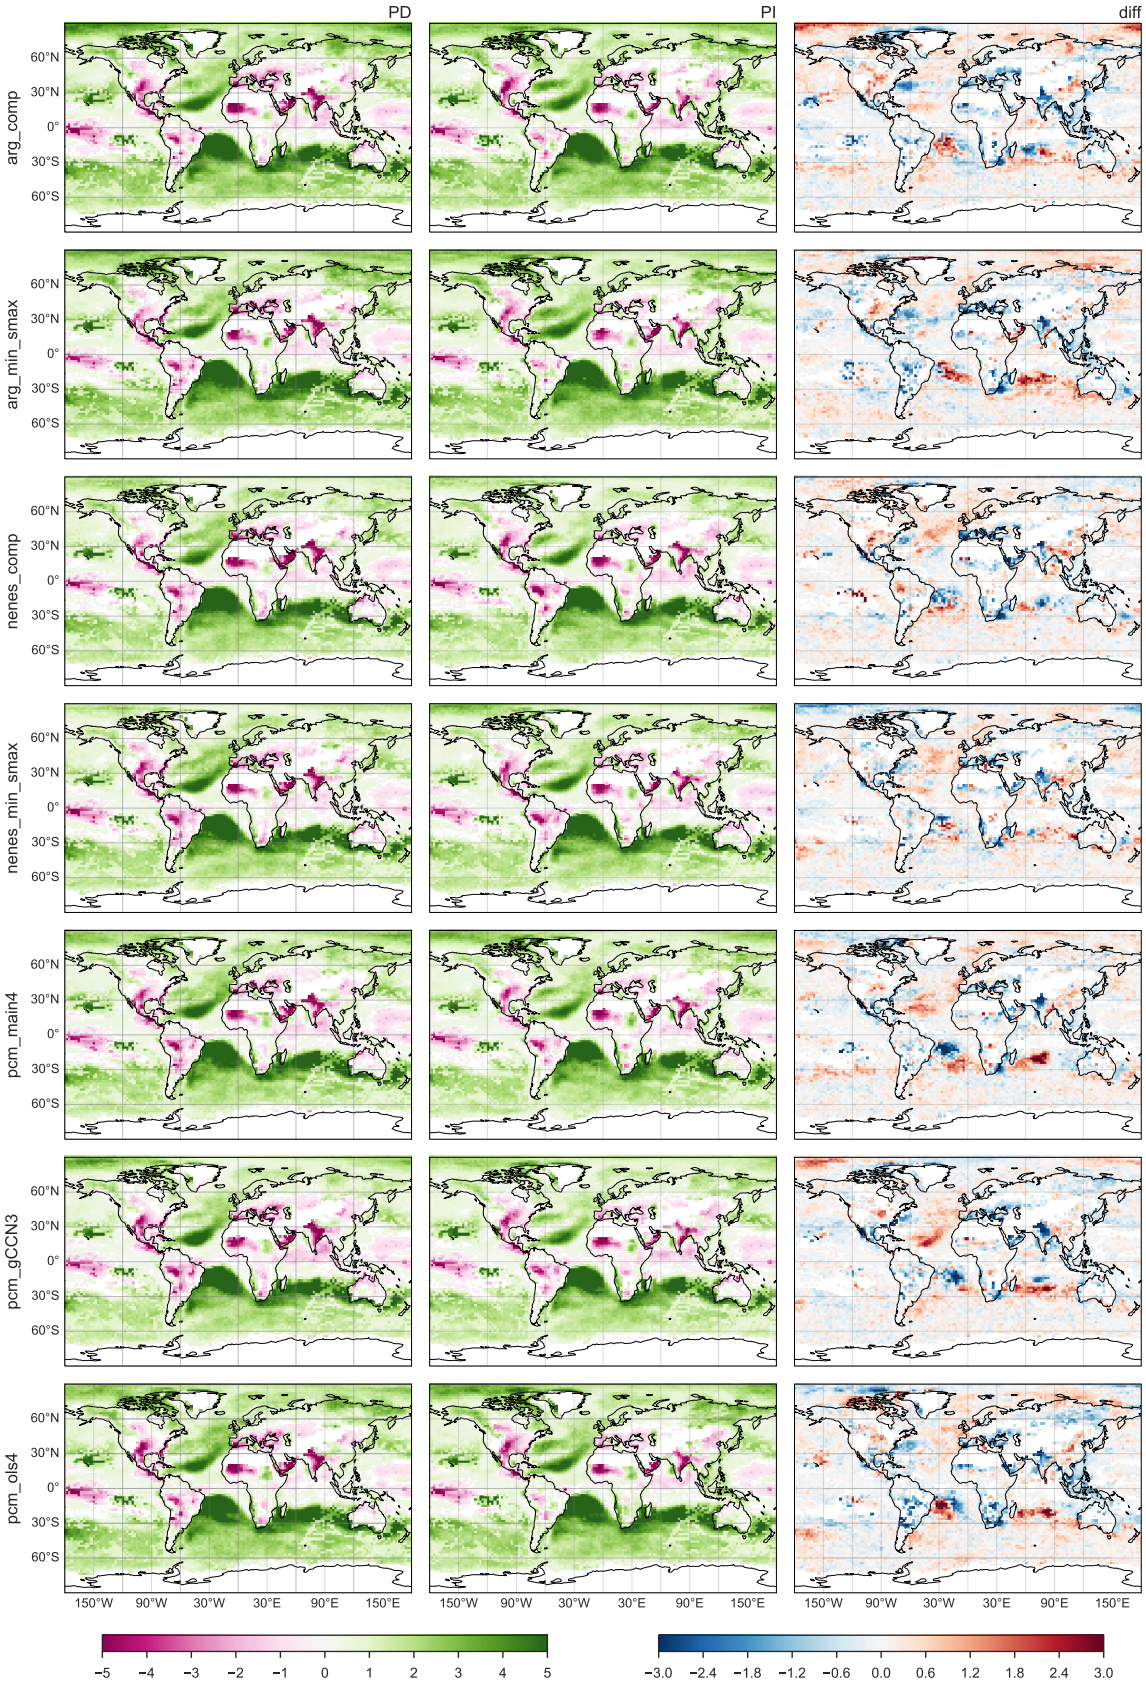

Figure 5-16: Global distributions of ACI_N computed using aerosol optical depth as a proxy for the aerosol concentration. Each rows corresponds to annual-average relationships for the indicated activation case; each column indicates the emissions scenario, where "diff" is the difference between the two. Regression coefficients which were not statistically significant or used less than $n = 1000$ samples were excluded

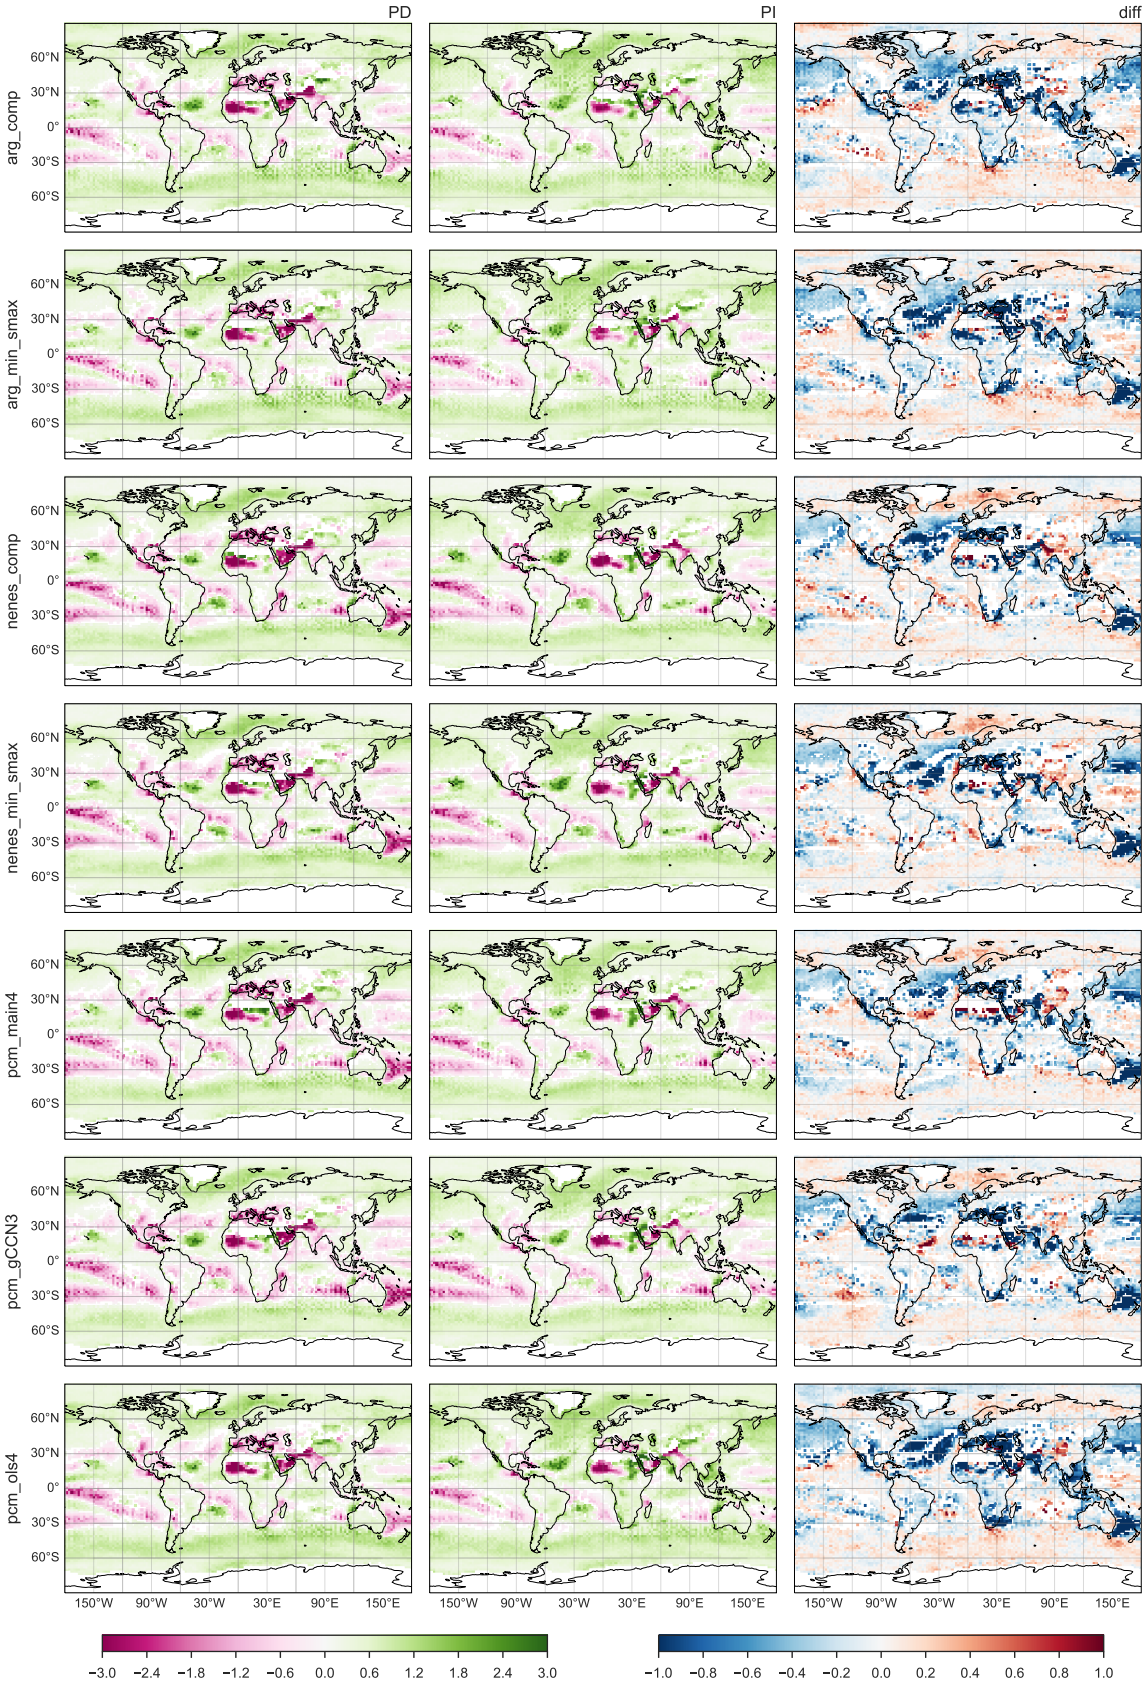

Figure 5-17: Same as Figure 5-16, but looking at relationships computed using CCN as the proxy for earosol concentration.

Chapter 6

Conclusions

6.1 Summary

In this dissertation, we have considered climatic phenomena at scales of the opposite extremes in both time and space.

We began our study on the scale of small particles just a few microns in size, growing via condensation over the course of seconds and minutes. Widening our perspective, we considered how collections of particles interact during this process, competing for moisture and limiting each other's growth into cloud droplets. We immediately recognized an important consequence of this competition: changes in the ambient particle population could change the concentration of droplets inside clouds. Initially, we limited our analysis of this consequence to a simple aerosol case: a single, lognormal aerosol mode. We showed that if the particles in such a mode had a homogeneous chemical composition, then assessing the sensitivity of the activation process (which transforms particles into cloud droplets) was straightforward; we showed that the average size of particles played an important, controlling role buffeted by the meteorology driving the process. Acknowledging that simplified parameterizations must substitute for detailed activation cases in many applications, we then analyzed the performance of several existing schemes with the idealized aerosol model, and noted that they were relatively accurate, with a few caveats. Those caveats included situations where there was heavy pollution (high aerosol number concentration) or weak meteorological forcing (shallow convection driving stratiform clouds).

The simple aerosol population we considered at first is wholly insufficient to describe

what happens in the real atmosphere, which features a staggering complexity of particles of many different sizes and compositions. To connect our initial investigation into droplet activation to the real world, we broadened our perspective to consider an aerosol population with multiple size modes of different species and mixing between some of them. Using a model of aerosol physics and chemistry built on such a description of the ambient aerosol, we simulated distributions of aerosol across the globe, subject to both natural and anthropogenic emissions sources of both particles and their gaseous precursors. We then asked a similar line of questioning aimed at assessing which particles were the most influential in the activation process. The ensuing results of this pursuit contained both familiar and surprising answers. On the one hand, we showed that activation in the complex case was very similar to the idealized one. However, in our formulation of the aerosol size distribution, one mode was clearly the most important—the accumulation mode sulfate. This mode encompasses a size range where particles have a low barrier to activation, and is often the most numerous by number concentration. The inclusion of giant condensation nuclei such as dust or sea salt did not appreciably change the importance of this mode, except in the cases where it was depleted of mass and number. Most surprisingly, activation was so sensitive to the details of the accumulation mode that in the majority of cases, one could neglect the other modes entirely in prognostic activation calculations and still derive cloud droplet concentrations within 10% of the expected value from more-detailed assessments.

While we studied the nuances of droplet activation, we began developing an alternative approach to parameterizing the process for inclusion in global models. Modeling the process of droplet activation itself, directly, is a relatively straightforward task involving the numerical integration of a set of coupled ordinary differential equations; a modern laptop can perform such a calculation in a dozen seconds at most. Including detail—like coupling the simulated aerosol size distribution to physical and chemical processes which change it—is not particularly difficult. But this is still too complex to directly include in global models, and thus the physical linkage between aerosol, clouds, and climate in such models requires parameterization. The first such parameterizations began appearing over two decades ago, and have since been generalized for use in nearly any modern climate model.

But generalizability does not necessarily imply accuracy, in a global sense. Different parameterizations have quirks or biases related to the input parameters supplied to them. Whereas these idiosyncrasies may not matter much for one given formulation of the ambient

aerosol and its attendant interfacial parameter space, they could be quite significant for another. We observed this for two schemes from the literature in both our idealized and complex aerosol cases. Particularly in the complex aerosol cases, different activation schemes could lead to fundamentally different climatologies of cloud droplet number concentration and its sensitivity to aerosol perturbations.

With this in mind, we sought to develop a framework for a new class of activation parameterizations based on statistical emulation of a detailed, process model. Such a framework could have several significant advantages over traditional parameterizations. First, emulators minimize the set of assumptions necessary to deploy a parameterization, since they target the functionality of a full-complexity model. Second, emulators can be extremely computationally efficient, both in terms of the number operations necessary to evaluate a given formulation and in the amount of data necessary to reserve in cache for the evaluation. But most importantly, by emulating a detailed model, they retain their generalizability. As long as the reference model is suitable for a given problem, the emulator should produce high-accuracy, reliable results.

Using a technique from the uncertainty quantification literature, we built emulators of a detailed parcel model and showed their superiority at reproducing droplet activation statistics in both idealized and complex aerosol scenarios. While not perfect in their application to the complex aerosol formulation, they offered an important check on the performance of the other parameterizations we considered. Furthermore, they offer a vehicle for important future work (to be addressed later in this chapter).

With our suite of old and new droplet activation schemes in hand, we turned our attention to a climate phenomenon at the opposite extreme of the time and space scales: the aerosol indirect effect on climate. Incorporating our different activation schemes into the same parent climate model, we simulated present day and pre-industrial climate states which were very familiar, except in the details of their cloud microphysical climatology. Actually, we stratified our models based on the average cloud droplet number concentration they simulated for identical emissions scenarios. Attendant with this stratification is a change in the modeled shortwave cloud radiative effect in each climate state, and its sensitivity to the pre-industrial-to-present-day increase in anthropogenic aerosol emissions. As a consequence, the diagnosed indirect effect in each version of the model is different, and between the seven activation schemes we considered here, the range of indirect effects is nearly as large as the

ensemble of model-derived estimates previously reported in the literature.

This result immediately suggests that the relationship between aerosols and clouds in the global models plays an important role in the indirect effect. To study this in more detail, we studied metrics of aerosol-cloud interactions using our suite of models and activation schemes. Historically, these metrics have been used to estimate and provide constraints for the indirect effect based on satellite measurements, as well as provide a metric to evaluate the performance of different global models. However, these metrics are often defined using aerosol optical depth as a measure of the change in aerosol burden. We compared this definition versus one using the cloud condensation nuclei concentration, which is more directly related to the activation process which lies at the heart of aerosol-cloud interactions. Such re-defined metrics provide a more consistent way to potentially develop constraints on aerosol-cloud interactions in the pre-industrial climate. Furthermore, they correlate more strongly with traditional, top-down estimates of the indirect effect derived using global models. Unfortunately, cloud condensation nuclei concentration are extremely difficult to measure from satellite observation platforms, limiting their application to *in situ* data from aircraft and other intensive measurement campaigns.

6.2 Implications for Future Work

Throughout this dissertation, we have highlighted many questions of particular research interest for future work. Here, we review some of them in an over-arching framework more closely aligned with the fundamental questions initially pursued in this work.

6.2.1 Emergent and Novel Constraints on Aerosol-Cloud Interactions

In Chapter 5, we developed the notion that traditional metrics of aerosol-cloud interactions (ACI) derived using a simplified, adiabatic cloud model in order to evaluate satellite retrievals could be problematic in a few ways. For instance, the adiabatic cloud model is really only applicable to maritime boundary layer clouds, which entirely precludes studying polluted, continental regimes. Even if we restrict our analysis to the cleanest regimes—which should be the most sensitive to aerosol perturbations—the uncertainties involved in such

satellite-based estimates of cloud properties could be very large. More importantly, satellites cannot be used to infer CCN concentrations, which we argued were more closely related to the physics of activation and could be used to place ACI in the pre-industrial climate, a major improvement over AOD-based metrics.

Despite their intrinsic difficulties, ACI-based approaches to constraining the indirect effect could still be powerful tools for probing the present-day climate. In particular, stratification of model and observational data by meteorological regime afford a way to compare the two one a more granular basis, where the underlying assumptions in the satellite retrievals may be better satisfied (Zhang et al., 2016). Alternatively, since we have far more freedom in how we process and analyze model output, it is entirely reasonable to simulate satellite-based retrievals online within a particular model run, creating an apples-to-apples comparison between the two (Ban-Weiss et al., 2014).

The bigger problem, though, is that what aerosol data we can glean from satellites may simply not provide a strong enough constraint on the physics of activation. As we showed in Chapter 5, CCN-based metrics are far more consistent with physically-based approaches for computing droplet activation. Unfortunately, data on the global distribution of CCN is woefully incomplete. Satellite retrieval methods do not yet exist for CCN, and even if they did, Stier (2016) suggests that available measurement techniques (e.g. observing aerosol radiative properties) would provide very poor constraints on it. Although some *in situ* CCN data is available, it is limited to a small selection of cloud regimes from a small portion of the globe, under a small variety of meteorological conditions. Yet, because of their value in addressing this problem, there should be a renewed focus on collating existing CCN-cloud datasets and including provisions in future measurement campaigns to broaden this archive. Stier (2016) recommends one such path forward, using space-borne lidars to feed a global aerosol model assimilation program.

This particular approach to evaluating ACI is but one tact available to the research community. Were it a cheap and easy path forward, one might be content to leave the discussion here. But, the main reason we do not have such *in situ* datasets is because acquiring this data is difficult, time-consuming, and costly, and the barriers to its collection will not be easily overcome. As a result, it is worthwhile to postulate alternative approaches and data sources which could help constrain ACI.

Feingold et al. (2016) offers one potential route which would require the research com-

munity to shift its focus from fundamental measures of ACI—correlations between aerosol and cloud droplet number or liquid water path—to other relationships more directly linked to the cloud radiative effect. Specifically, Feingold et al. suggest considering the relationship between the albedo and cloud fraction of a cloudy scene, which may be expressed as

$$A = A_c f_c + A_s(1 - f_c) \tag{6.1}$$

where A_c and A_s are the cloud and surface albedo, respectively, and F_c is the cloudy fraction of the scene. A_c is directly a function of cloud micro- and macro-physical properties, but is far simpler to measure from both space-borne and surface-based instruments. Reproducing observed, complex patterns in time and space between A_c and f_c in a data-driven, process modeling framework involving both aerosol and meteorology fields could yield powerful insight into the co-variation of aerosols and clouds.

But this represents an advanced, high-level fusion of approaches tackling ACI. We should also pursue the development of new, simple models of the aerosol indirect effect which can take advantage of new observational data or refined historical data, as it becomes available. Two such recent efforts taking similar (yet very distinct) approaches to this problem either reproduced the behavior of more complex global models (Ghan, 2013) or provided grossly revised and smaller estimates of the indirect effect (Stevens, 2015). Despite their differences, what the simple modeling approach provides is a short-circuit on the extremely complex, convoluted physics, which make studying indirect effects with global models so difficult.

Our conception of emergent constraints should not be limited to just explicit aerosol-cloud interactions and related measurements. Ekman (2014) noted that CMIP5 models which included detailed treatments of aerosol-cloud influences better resolved transient regional warming trends over the 20^{th} century. Another study by Wang (2015) suggests that aerosols have dominated the distribution of precipitations changes in the tropics and northern hemisphere extra-tropics with reference to the pre-industrial climate over the past few decades; these precipitation changes are closely related to the gradients in surface temperature anomaly due to aerosol-induced cooling in the northern hemisphere offsetting greenhouse gas warming. Observations such as those provide a powerful way to link aerosol direct and indirect effects to components of the climate system (temperature and precipitation) which are far better observed than cloud microphysical properties or aerosol details. Cre-

ative use of these observed climate-scale relationships, then, could offer a novel way to tune models or evaluate ensembles and inter-comparisons.

6.2.2 Climate Response to Aerosol Perturbation

Studies of the aerosol indirect effect on climate heavily focus on quantifying the radiative forcing associated with aerosol perturbations. As we have stressed many times in this dissertation, that radiative forcing is a major component of uncertainty in present and future climate change, and as a result is a high-profile research target. But a critical observer of the zonal distribution of changes in cloud radiative effects presented in Chapter 4 would note that aerosols impart a very large, regional forcing on the climate system. That forcing has been associated with large, local impacts on the climate system, particularly with respect to precipitation (Bollasina et al., 2011).

However, forcing in and of itself does not represent the complete picture of aerosol impacts on climate. Missing from this framework is an assessment of the climate system response to our diagnosed aerosol forcing. Unlike greenhouse gases, aerosol perturbations can be very large, but heterogeneous in both time and space. Clouds must respond to these perturbations on very fast timescales owing to the tight physical coupling between aerosols and their underlying processes. As a result, the attendant forcing from the indirect effect can be large, but spatiotemporally heterogeneous. One might expect, then, that the climate system should respond on a spectrum of temporal and spatial scales to the forcing due to aerosol indirect effects.

We might further expect these responses to manifest themselves most strongly in precipitation fields, which are sensitive to both the microphysical changes in clouds that aerosols induce, but also to regional changes in the atmospheric energy balance. Some recent work in the literature develops this notion in more detail, noting that fast responses scale very strongly with changes in atmospheric absorption, while slow responses scale approximately with surface temperature changes (Samset et al., 2016). However, this work has not explicitly looked at aerosol indirect effects.

Decomposing the fast and slow responses of the climate system to aerosol perturbations could help elucidate the type of emergent constraints on aerosol-cloud interactions we discussed in the previous section. However, such an exercise could also yield insight on climate

responses to forcings at a much more fundamental level. For instance, transient climate responses in models correlate very strongly with their equilibrium response, but at the same time, modeling of 20^{th} century climate change suggests that the transient response climate response to ozone and aerosols is much larger than that of carbon dioxide (Shindell, 2014). This is in part due the asymmetric hemispheric distribution of ozone and aerosol forcing over the past century, which is overwhelmingly in the northern hemisphere and as a result incurs rapid land responses and initiates stronger feedback mechanisms, particularly in the middle to high latitudes. The linkage between forcing agents, transient climate response and equilibrium climate sensitivity provide a method for applying contemporary observations to better constrain future climate change on very long timescales.

Responses on intermediate time scales also deserve careful consideration. Wang et al. (2015) showed that the spatial re-distribution of anthropogenic aerosol emissions over just the past few decades had large implications for the general circulation and regional warming and precipitation trends. What is the relationship between changes on these decadal scales to the spatial re-distribution of aerosol emissions and the global perspective we take when studying fast and slow responses? Tackling questions like these could also help generate the emergent constraints on the indirect effect which we outlined in the previous section, since they provide connections to robust, long-observed aspects of the climate system (specifically temperature, precipitation, and the top-of-atmosphere radiation balance).

6.2.3 Uncertainty in Future, Transient Climate Change

A little less than a year ago, the international political community negotiated an agreement on a new approach to mitigating climate change. This "Paris Agreement" represents an evolution of a mitigation principle first codified in the 1992 UN Framework Convention on Climate change, which asserted that each country should act "in accordance with their common but differentiated responsibilities and respective capabilities." This principle serves two purposes. First, it provides a way for developing countries to take advantage of cheap (but carbon-intensive) pathways for economic growth, and enables them to focus on improving welfare and fighting poverty in the short-term. Second, it acknowledges the trouble with one-size-fits-all mitigation policies such as universal benchmarks on emissions reductions (such as a mandated decrease to $X\%$ of 2005 emissions levels by 2050). Instead,

it offers countries the latitude and flexibility to adopt a plan best-adapted for their energy and economic situation.

As a result, the Paris Agreement encodes a mechanism where each signatory country commits to an "intended nationally determined contribution" (INDC). INDCs are mitigation targets; for benchmarking, we can convert them to reductions in aggregate emissions or emissions intensity, but most importantly, we can sum up the portfolio of INDCs at any given point in time to quantitatively assess the likelihood that we will avoid certain warming scenarios. To place this dissertation in the historical context of our civilization's struggle to grasp the challenges posed by climate change, while drafting this very chapter, the European Union backed the ratification of the Agreement, providing the impetus necessary for it to go into force on November 4, 2016.

Although comprehensive in its pursuit of mitigating catastrophic climate change, the Paris Agreement has a glaring omission: any mention of aerosols or short-lived climate pollutants.

A key implication of the work in this dissertation is that the specter of uncertainty haunting our understanding of 20^{th} climate change will persist with us in a perverse fashion over the next few decades. The observed increase in atmospheric aerosol pollution since the pre-industrial era has generally cooled the climate system. Since 1970, the sum total of that aerosol influence, derived using observational constraints on Earth's fundamental energy balance, is likely -1.1 ± 0.4 W m^{-2} (Murphy et al., 2009). Other estimates place the present-day, instantaneous negative forcing due to aerosols at between 0.5 W m^{-2} to 1 W m^{-2} (Westervelt et al., 2015). This constitutes a masking of climate change; the cooling effects of aerosol emissions, even in their redistribution across the globe (Wang et al., 2015), acts to offset warming which we would otherwise expect from increasing greenhouse gas emissions.

Over the next few decades, as part of their published INDCs, many nations will continue an aggressive pursuit to reduce their aerosol and precursor gas emissions. Such plans are a sort of side-step around the politics of climate change; particulate emissions are implicated with negative impacts on human health, so they can be justified on terms wholly distinct from climate. But at the same time, these sorts of policies have climate co-benefits. Many approaches taken to reduce particulate emissions would necessarily reduce carbon dioxide emissions, thus producing a positive side-effect towards mitigating climate change. Under

this guise and others, aerosol emissions are expected to decrease in the global average over the next several decades (Myhre et al., 2013).

Unfortunately, a rapid decline in global aerosol emissions in the future will necessarily involve a "climate penalty" as greenhouse gas forcing is unmasked by a decreased aerosol direct and indirect effect. A critical near-term research question is the rigorous quantification of this climate penalty, both from first principles and within the context of the repository of published INDCs and aerosol reduction commitments. Relatively few publications have focused on quantifying this climate penalty, and much of the existing work focuses explicitly on the emissions timelines assumed within the RCP projections (Westervelt et al., 2015). But the impact of these emissions reductions on climate change over the next century can be quite large. For instance, Levy et al. (2013a) noted that using one global model, relative to holding aerosol emissions at their present day levels, the emissions reductions in RCP4.5 lead to an increase in global average temperature of about 1 °C and increase in global average precipitation of about $0.1\,\mathrm{mm\,d^{-1}}$ by 2100. The regional response to the reduction in aerosol emissions could be much greater (Westervelt et al., 2015). Furthermore, the estimated range of radiative forcing from aerosol direct and indirect effects alone calculated using models evaluating the RCP projections ranges from $0.4\,\mathrm{W\,m^{-2}}$ to $2\,\mathrm{W\,m^{-2}}$ (Westervelt et al., 2015). Few (if any) studies have re-assessed these estimates using up-to-date emissions scenarios which more accurately reflect the global corpus of INDCs, and no simple integrated assessment model used to quickly analyze the impact of a given climate policy adequately accounts for aerosol indirect effects or their uncertainties.

As a result, these estimates are highly uncertain, and fundamentally rely on the magnitude of the modeled indirect effect in whatever global model is used to compute them. But this uncertainty implies an important and deeply uncomfortable question from both a science policy perspective: if the world has already experienced about a degree of warming, and reductions in aerosol emissions alone will inevitably unmask up to an additional degree of warming, is a 1.5 °C warming target even remotely feasible, let alone a 2 °C one? If not, then how do we reconcile this scientific reality with the entrenched rhetoric advocating for climate policy calibrated to avoid a "disastrous" 2 °C threshold?

The previously enumerated lines of future research will hopefully reduce the uncertainty surrounding the indirect effect, and greatly improve our ability to anticipation this future source of climate change. But a dedicated effort should target these questions directly.

Owing to their complexity and the uncertainty in their attendant parameterizations, fully-coupled atmosphere-climate models may not be the best way to study these questions. Instead, tools from the domain of simplified models may be a far better choice. One approach could involve evaluating a wide swath of potential emissions scenarios using intermediate complexity models like those of Ghan et al. (2013) or Stevens (2015). The computational efficiency of these models would allow the application of Monte Carlo techniques (such as the variance-based decomposition used in Chapters 2 and 3 of this dissertation) to rigorously quantify the uncertainty in both the emissions scenarios and the construction of the simplified models themselves (e.g. the use of tuning parameters). Such tools would be efficient enough to be developed into online or interactive policy evaluation tools, which could be provided to policymakers studying and negotiating future iterations of the Paris Agreement. But more importantly, they would provide an exhaustive, probabilistic perspective on future climate change.

6.3 Closing Thoughts

We end this lengthy discussion with a reflection back to Twomey's simple climate equation. Although it obeys fundamental relationships derived from natural laws, the climate system as encoded in Equation (1.1) is incredibly complex and involves physics and chemistry interacting along a breathtakingly vast continuum of spatial and temporal scales. As scientists practicing reductionism, we are tempted to break down such a complicated system into smaller constituents, rigorously test them, and then iteratively re-construct it from its parts.

This approach initially worked very well in understanding the aerosol indirect effect on climate, as both Twomey (1977) and Albrecht (1989) showed in a qualitative sense. But in a quantitative sense, it has yielded little progress in constraining the indirect effect. As a result, we have resorted to reductionism in an extreme form, attempting to simultaneously simplify the relationships between aerosol, clouds, and climate, but also to quantify the contribution of many progressively detailed elements of the system to uncertainty in the indirect effect itself. This dissertation is, in more ways than one, a culmination of this trend in analysis, and its primary scientific result—that activation itself (by means of setting the cloud droplet number concentration) contributes as much uncertainty as all the other

components of the system—is an inevitable result of this line of pursuit.

However, the work in this dissertation belongs in a much larger context and is part of a fundamental sea-change in the dialogue surrounding investigation into the indirect effect (Feingold et al., 2016). Quantifying the sources of uncertainty in our climate models will clearly only get us so far in solving this problem. Now, more than ever, we should embrace a renewed focus on improving observations of the climate system, particularly with regards to the details of the hydrological cycle and the microphysical properties of clouds. These new observations, coupled with the dramatic advances made in computational modeling of the atmosphere at all scales over the past decade, will provide an entirely new framework for understanding the interaction of aerosols and clouds. Successfully levying these developments into scientific understanding will require a liberal dose of creativity, and a willingness to re-evaluate established notions. I hope that the work discussed in this dissertation demonstrates the value of this outlook and approach, and positively contributes to our understanding of the complexities of our planet's climate.

Bibliography

Abdul-Razzak, H., 2002: A parameterization of aerosol activation 3. Sectional representation. *Journal of Geophysical Research*, **107 (D3)**, 4026, doi:10.1029/2001JD000483, URL http://www.agu.org/pubs/crossref/2002/2001JD000483.shtmlhttp://doi.wiley.com/10.1029/2001JD000483.

Abdul-Razzak, H., and S. J. Ghan, 2000: A parameterization of aerosol activation 2. Multiple aerosol types. *Journal of Geophysical Research*, **105 (D5)**, 6837, doi:10.1029/1999JD901161, URL http://www.agu.org/pubs/crossref/2000/1999JD901161.shtmlhttp://dx.doi.org/10.1029/1999jd901161.

Abdul-Razzak, H., and S. J. Ghan, 2004: Parameterization of the influence of organic surfactants on aerosol activation. *Journal of Geophysical Research: Atmospheres*, **109 (D3)**, n/a–n/a, doi:10.1029/2003JD004043, URL http://doi.wiley.com/10.1029/2003JD004043.

Abdul-Razzak, H., S. J. Ghan, C. Rivera-Carpio, H. A. Razzak, and C. R. Carpio, 1998: A parameterization of aerosol activation: 1. Single aerosol type. *Journal of Geophysical Research*, **103 (D6)**, 6123, doi:10.1029/97JD03735, URL http://www.agu.org/pubs/crossref/1998/97JD03735.shtmlhttp://doi.wiley.com/10.1029/97JD03735.

Adams, B. M., and Coauthors, 2014: DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 6.0 User's Manual. Tech. rep., Sandia National Laboratories, Albuquerque, New Mexico.

Albani, S., and Coauthors, 2014: Improved dust representation in the Community Atmosphere Model. *Journal of Advances in Modeling Earth Systems*, n/a–n/a, doi:10.1002/2013MS000279, URL http://doi.wiley.com/10.1002/2013MS000279.

Albrecht, B. A., 1989: Aerosols, cloud microphysics, and fractional cloudiness. *Science*, **245 (4923)**, 1227–1230, doi:10.1126/science.245.4923.1227, URL http://www.sciencemag.org/content/245/4923/1227.abstracthttp://www.ncbi.nlm.nih.gov/pubmed/17747885.

Altaratz, O., I. Koren, L. Remer, and E. Hirsch, 2014: Review: Clouds invigoration by aerosols - coupling between microphysics and dynamics. *Atmospheric Research*, doi:10.1016/j.atmosres.2014.01.009, URL http://linkinghub.elsevier.com/retrieve/pii/S0169809514000106.

Andreae, M. O., 2009: Correlation between cloud condensation nuclei concentration and aerosol optical thickness in remote and polluted regions. *Atmospheric Chemistry and Physics*, **9 (2)**, 543–556, doi:10.5194/acp-9-543-2009, URL http://www.atmos-chem-phys.net/9/543/2009/.

Anttila, T., and V.-M. Kerminen, 2007: On the contribution of Aitken mode particles to cloud droplet populations at continental background areas – a parametric sensitivity study. *Atmospheric Chemistry and Physics*, **7 (17)**, 4625–4637, doi:10.5194/acp-7-4625-2007, URL http://www.atmos-chem-phys.net/7/4625/2007/.

Ban-Weiss, G. A., and Coauthors, 2014: Evaluating clouds, aerosols, and their interactions in three global climate models using satellite simulators and observations. *Journal of Geophysical Research: Atmospheres*, **119 (18)**, 10,876–10,901, doi:10.1002/2014JD021722, URL http://doi.wiley.com/10.1002/2014JD021722.

Barahona, D., and A. Nenes, 2007: Parameterization of cloud droplet formation in large-scale models: Including effects of entrainment. *Journal of Geophysical Research*, **112 (D16)**, D16 206, doi:10.1029/2007JD008473, URL http://doi.wiley.com/10.1029/2007JD008473.

Barahona, D., R. Sotiropoulou, and A. Nenes, 2011: Global distribution of cloud droplet number concentration, autoconversion rate, and aerosol indirect effect under diabatic droplet activation. *Journal of Geophysical Research*, **116 (D9)**, D09 203, doi:10.1029/2010JD015274, URL http://doi.wiley.com/10.1029/2010JD015274.

Barahona, D., R. E. L. West, P. Stier, S. Romakkaniemi, H. Kokkola, and A. Nenes, 2010: Comprehensively accounting for the effect of giant CCN in cloud activation parameterizations. *Atmospheric Chemistry and Physics*, **10 (5)**, 2467–2473, doi:10.5194/acp-10-2467-2010, URL http://www.atmos-chem-phys.net/10/2467/2010/.

Barth, M. C., P. J. Rasch, J. T. Kiehl, C. M. Benkovitz, and S. E. Schwartz, 2000: Sulfur chemistry in the National Center for Atmospheric Research Community Climate Model: Description, evaluation, features, and sensitivity to aqueous chemistry. *Journal of Geophysical Research: Atmospheres*, **105 (D1)**, 1387–1415, doi:10.1029/1999JD900773, URL http://doi.wiley.com/10.1029/1999JD900773.

Bennartz, R., 2007: Global assessment of marine boundary layer cloud droplet number concentration from satellite. *Journal of Geophysical Research*, **112 (D2)**, D02 201, doi:10.1029/2006JD007547, URL http://doi.wiley.com/10.1029/2006JD007547.

Blatman, G., and B. Sudret, 2011: Adaptive sparse polynomial chaos expansion based on least angle regression. *Journal of Computational Physics*, **230 (6)**, 2345–2367, doi:10.1016/j.jcp.2010.12.021, URL http://linkinghub.elsevier.com/retrieve/pii/S0021999110006856.

Bogenschutz, P. A., A. Gettelman, H. Morrison, V. E. Larson, C. Craig, and D. P. Schanen, 2013: Higher Order Turbulence Closure and Its Impact on Climate Simulations in the Community Atmosphere Model. *Journal of Climate*, 130806142946000, doi:10.1175/JCLI-D-13-00075.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JCLI-D-13-00075.1?af=R.

Bollasina, M. A., Y. Ming, and V. Ramaswamy, 2011: Anthropogenic Aerosols and the Weakening of the South Asian Summer Monsoon. *Science*, **334 (6055)**, 502–505, doi:10.1126/science.1204994, URL http://www.sciencemag.org/content/334/6055/502.abstract.

Boucher, O., and U. Lohmann, 1995: The Sulfate-Ccn-Cloud Albedo Effect - a Sensitivity Study With 2 General-Circulation Models. 281–300 pp., doi:10.1034/j.1600-0889.47.issue3.1.x.

Boucher, O., and Coauthors, 2013: Clouds and Aerosols. *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley, Eds., Cambridge University Press, Cambridge, United Kingdom, and New York, NY, USA, chap. 7, 571–657.

Brent, R. P., 1973: An Algorithm with Guaranteed Convergence for Finding a Zero of a Function. *Algorithms for Minimization without Derivatives*, Prentice-Hall, Englewood Cliffs, New Jersey, chap. 4, 195.

Brown, P. N., G. D. Byrne, and A. C. Hindmarsh, 1989: VODE: A Variable-Coefficient ODE Solver. *SIAM Journal on Scientific and Statistical Computing*, **10 (5)**, 1038–1051, doi:10.1137/0910062, URL http://epubs.siam.org/doi/abs/10.1137/0910062.

Budyko, M. I., 1969: The effect of solar radiation variations on the climate of the Earth. *Tellus*, **21 (5)**, 611–619, doi:10.1111/j.2153-3490.1969.tb00466.x, URL http://tellusa.net/index.php/tellusa/article/view/10109.

Calbó, J., W. Pan, M. Webster, R. G. Prinn, and G. J. McRae, 1998: Parameterization of urban subgrid scale processes in global atmospheric chemistry models. *Journal of Geophysical Research*, **103 (D3)**, 3437, doi:10.1029/97JD02654, URL http://doi.wiley.com/10.1029/97JD02654.

Carslaw, K. S., and Coauthors, 2013: Large contribution of natural aerosols to uncertainty in indirect forcing. *Nature*, **503 (7474)**, 67–71, doi:10.1038/nature12674, URL http://www.ncbi.nlm.nih.gov/pubmed/24201280.

Cohard, J.-M., J.-P. Pinty, and C. Bedos, 1998: Extending Twomey's Analytical Estimate of Nucleated Cloud Droplet Concentrations from CCN Spectra. *Journal of the Atmospheric Sciences*, **55 (22)**, 3348–3357, doi:10.1175/1520-0469(1998)055<3348:ETSAEO>2.0.CO;2.

Cohard J-M., and J.-P. Pinty, 2000: A comprehensive two-moment warm microphysical bulk scheme. I: Description and tests. *Quart. J. Roy. Meteor. Soc.*, **126**, 1815–1842, doi:10.1002/qj.49712656613, URL http://doi.wiley.com/10.1002/qj.49712656614.

Cohen, J. B., and R. G. Prinn, 2011: Development of a fast, urban chemistry metamodel for inclusion in global models. *Atmospheric Chemistry and Physics*, **11 (15)**, 7629–7656, doi:10.5194/acp-11-7629-2011, URL http://www.atmos-chem-phys.net/11/7629/2011/.

Crook, J. A., L. S. Jackson, and P. M. Forster, 2016: Can increasing albedo of existing ship wakes reduce climate change? *Journal of Geophysical Research: Atmospheres*, **121 (4)**, 1549–1558, doi:10.1002/2015JD024201, URL http://doi.wiley.com/10.1002/2015JD024201.

DeMott, P. J., and Coauthors, 2010: Predicting global atmospheric ice nuclei distributions and their impacts on climate. *Proceedings of the National Academy of Sciences of the United States of America*, **107 (25)**, 11 217–22, doi:10.1073/pnas.0910818107, URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=2895116{&}tool=pmcentrez{&}rendertype=abstract.

Ekman, A. M. L., 2014: Do sophisticated parameterizations of aerosol-cloud interactions in CMIP5 models improve the representation of recent observed temperature trends? *Journal of Geophysical Research: Atmospheres*, **119 (2)**, 817–832, doi:10.1002/2013JD020511, URL http://doi.wiley.com/10.1002/2013JD020511.

Ekman, A. M. L., A. Engström, and A. Söderberg, 2011: Impact of Two-Way Aerosol–Cloud Interaction and Changes in Aerosol Size Distribution on Simulated Aerosol-Induced Deep Convective Cloud Sensitivity. *Journal of the Atmospheric Sciences*, **68 (4)**, 685–698, doi:10.1175/2010JAS3651.1, URL http://journals.ametsoc.org/doi/abs/10.1175/2010JAS3651.1.

Ekman, A. M. L., A. Engström, and C. Wang, 2007: The effect of aerosol composition and concentration on the development and anvil properties of a continental deep convective cloud. *Quarterly Journal of the Royal Meteorological Society*, **133**, 1439–1452, doi:10.1002/qj, URL http://onlinelibrary.wiley.com/doi/10.1002/qj.108/abstract.

Ekman, A. M. L., M. Hermann, P. Gro, J. Heintzenberg, D. Kim, and C. Wang, 2012: Submicrometer aerosol particles in the upper troposphere/lowermost stratosphere as measured by CARIBIC and modeled using the MIT-CAM3 global climate model. *Journal of Geophysical Research Atmospheres*, **117 (11)**, 1–18, doi:10.1029/2011JD016777.

Ekman, A. M. L., C. Wang, J. Ström, and R. Krejci, 2006: Explicit Simulation of Aerosol Physics in a Cloud-Resolving Model: Aerosol Transport and Processing in the Free Troposphere. *Journal of the Atmospheric Sciences*, **63 (2)**, 682–696, doi:10.1175/JAS3645.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JAS3645.1.

Ekman, A. M. L., C. Wang, J. Wilson, and J. Ström, 2004: Explicit simulations of aerosol physics in a cloud-resolving model: a sensitivity study based on an observed convective cloud. *Atmospheric Chemistry and Physics*, **4 (3)**, 773–791, doi:10.5194/acp-4-773-2004, URL http://www.atmos-chem-phys.net/4/773/2004/acp-4-773-2004.pdf.

Engström, A., A. M. L. Ekman, R. Krejci, J. Ström, M. de Reus, and C. Wang, 2008: Observational and modelling evidence of tropical deep convective clouds as a source of mid-tropospheric accumulation mode aerosols. *Geophysical Research Letters*, **35 (23)**, L23 813, doi:10.1029/2008GL035817, URL http://doi.wiley.com/10.1029/2008GL035817.

Ervens, B., G. Feingold, and S. M. Kreidenweis, 2005: Influence of water-soluble organic carbon on cloud drop number concentration. *Journal of Geophysical Research D: Atmospheres*, **110**, 1–14, doi:10.1029/2004JD005634.

Fan, J., L. R. Leung, Z. Li, H. Morrison, H. Chen, Y. Zhou, Y. Qian, and Y. Wang, 2012: Aerosol impacts on clouds and precipitation in eastern China: Results from bin and bulk microphysics. *Journal of Geophysical Research*, **117 (null)**, D00K36, doi:10.1029/2011JD016537, URL http://dx.doi.org/10.1029/2011JD016537http://www.agu.org/pubs/crossref/2012/2011JD016537.shtml.

Feingold, G., W. L. Eberhard, D. E. Veron, and M. Previdi, 2003: First measurements of the Twomey indirect effect using ground-based remote sensors. *Geophysical Research Letters*, **30 (6)**, 19–22, doi:10.1029/2002GL016633, URL http://doi.wiley.com/10.1029/2002GL016633.

Feingold, G., and A. J. Heymsfield, 1992: Parameterizations of Condensational Growth of Droplets for Use in General Circulation Models. *Journal of the Atmospheric Sciences*, **49 (23)**, 2325–2342, URL http://journals.ametsoc.org/doi/abs/10.1175/1520-0469(1992)049{\T1\textless}2325:POCGOD{\T1\textgreater}2.0.CO;2.

Feingold, G., A. McComiskey, T. Yamaguchi, J. S. Johnson, K. S. Carslaw, and K. S. Schmidt, 2016: New approaches to quantifying aerosol influence on the cloud radiative effect. *Proceedings of the National Academy of Sciences*, **113 (21)**, 201514 035, doi:10.1073/pnas.1514035112, URL http://www.pnas.org/lookup/doi/10.1073/pnas.1514035112.

Feingold, G., L. Remer, J. Ramaprasad, and Y. J. Kaufman, 2001: Analysis of smoke impact on clouds in Brazilian biomass burning regions: An extension of Twomey's approach. *Journal of Geophysical Research*, **106 (D19)**, 22 907, doi:10.1029/2001JD000732, URL http://doi.wiley.com/10.1029/2001JD000732.

Fountoukis, C., and A. Nenes, 2005: Continued development of a cloud droplet formation parameterization for global climate models. *Journal of Geophysical Research*, **110 (D11)**, D11 212, doi:10.1029/2004JD005591, URL http://www.agu.org/pubs/crossref/2005/2004JD005591.shtml.

Fountoukis, C., and Coauthors, 2007: Aerosol–cloud drop concentration closure for clouds sampled during the International Consortium for Atmospheric Research on Transport and Transformation 2004 campaign. *Journal of Geophysical Research*, **112 (D10)**, D10S30, doi:10.1029/2006JD007272, URL http://doi.wiley.com/10.1029/2006JD007272.

Gantt, B., J. He, X. Zhang, Y. Zhang, and A. Nenes, 2014: Incorporation of advanced aerosol activation treatments into CESM/CAM5: Model evaluation and impacts on aerosol indirect effects. *Atmospheric Chemistry and Physics*, **14 (14)**, 7485–7497, doi:10.5194/acp-14-7485-2014, URL http://www.atmos-chem-phys.net/14/7485/2014/.

Garimella, S., Y. W. Huang, J. S. Seewald, and D. J. Cziczo, 2014: Cloud condensation nucleus activity comparison of dry- and wet-generated mineral dust aerosol: The significance of soluble material. *Atmospheric Chemistry and Physics*, **14 (12)**, 6003–6019, doi:10.5194/acp-14-6003-2014.

Gautschi, W., 1994: Algorithm 726 - ORTHPOL: A package of routines for generating orthogonal polynomials and Gauss-type quadrature rules. *ACM Transactions on Mathematical Software*, **20 (1)**, 21–62, doi:10.1145/292395.292467, URL http://arxiv.org/abs/math/9307212, 9307212.

Gettelman, A., 2015: Putting the clouds back in aerosol–cloud interactions. *Atmospheric Chemistry and Physics*, **15 (21)**, 12 397–12 411, doi:10.5194/acp-15-12397-2015, URL http://www.atmos-chem-phys.net/15/12397/2015/.

Gettelman, A., J. E. Kay, and K. M. Shell, 2012: The Evolution of Climate Sensitivity and Climate Feedbacks in the Community Atmosphere Model. **25 (5)**, 1453–1469, doi:10.1175/JCLI-D-11-00197.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JCLI-D-11-00197.1.

Gettelman, A., and H. Morrison, 2015: Advanced Two-Moment Bulk Microphysics for Global Models. Part I: Off-Line Tests and Comparison with Other Schemes. *Journal of Climate*, **28 (3)**, 1268–1287, doi:10.1175/JCLI-D-14-00102.1, URL http://journals.ametsoc.org/doi/10.1175/JCLI-D-14-00102.1.

Gettelman, a., H. Morrison, C. R. Terai, and R. Wood, 2013: Microphysical process rates and global aerosol–cloud interactions. *Atmospheric Chemistry and Physics*, **13 (19)**, 9855–9867, doi:10.5194/acp-13-9855-2013, URL http://www.atmos-chem-phys.net/13/9855/2013/.

Gettelman, A., and Coauthors, 2010: Global simulations of ice nucleation and ice supersaturation with an improved cloud scheme in the Community Atmosphere Model. *Journal of Geophysical Research*, **115 (D18)**, D18 216, doi:10.1029/2009JD013797, URL http://doi.wiley.com/10.1029/2009JD013797.

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*, 201514036, doi:10.1073/pnas.1514036113, URL http://www.pnas.org/lookup/doi/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, doi:10.5194/acp-13-9971-2013, URL http://www.atmos-chem-phys.net/13/9971/2013/.

Ghan, S. J., L. R. Leung, R. C. Easter, and H. Abdul-Razzak, 1997: Prediction of cloud droplet number in a general circulation model. *Journal of Geophysical Research: Atmospheres*, **102 (D18)**, 21 777–21 794, doi:10.1029/97JD01810, URL http://doi.wiley.com/10.1029/97JD01810.

Ghan, S. J., and Coauthors, 2011: Droplet nucleation: Physically-based parameterizations and comparative evaluation. *Journal of Advances in Modeling Earth Systems*, **3 (10)**, M10 001, doi:10.1029/2011MS000074, URL http://dx.doi.org/10.1029/2011MS000074.

Ghan, S. J., and Coauthors, 2013: A simple model of global aerosol indirect effects. *Journal of Geophysical Research: Atmospheres*, **118 (12)**, 6688–6707, doi:10.1002/jgrd.50567, URL http://doi.wiley.com/10.1002/jgrd.50567.

Golaz, J.-C., M. Salzmann, L. J. Donner, L. W. Horowitz, Y. Ming, and M. Zhao, 2011: Sensitivity of the Aerosol Indirect Effect to Subgrid Variability in the Cloud Parameterization of the GFDL Atmosphere General Circulation Model AM3. *Journal of Climate*, **24 (13)**, 3145–3160, doi:10.1175/2010JCLI3945.1, URL http://journals.ametsoc.org/doi/abs/10.1175/2010JCLI3945.1.

Griffin, R. J., D. R. Cocker, J. H. Seinfeld, and D. Dabdub, 1999: Estimate of global atmospheric organic aerosol from oxidation of biogenic hydrocarbons. *Geophysical Research Letters*, **26 (17)**, 2721–2724, doi:10.1029/1999GL900476, URL http://doi.wiley.com/10.1029/1999GL900476.

Haywood, J., and O. Boucher, 2000: Estimates of the direct and indirect radiative forcing due to tropospheric aerosols: A review. *Reviews of Geophysics*, **38 (4)**, 513, doi:10.1029/1999RG000078, URL http://doi.wiley.com/10.1029/1999RG000078.

Hoose, C., J. E. Kristjánsson, T. Iversen, a. Kirkevåg, Ø. Seland, and a. Gettelman, 2009: Constraining cloud droplet number concentration in GCMs suppresses the aerosol indirect effect. *Geophysical Research Letters*, **36 (12)**, L12 807, doi:10.1029/2009GL038568, URL http://doi.wiley.com/10.1029/2009GL038568.

Inness, a., and Coauthors, 2013: The MACC reanalysis: an 8 yr data set of atmospheric composition. *Atmospheric Chemistry and Physics*, **13 (8)**, 4073–4109, doi:10.5194/acp-13-4073-2013, URL http://www.atmos-chem-phys.net/13/4073/2013/.

Jimenez, J. L., and Coauthors, 2009: Evolution of Organic Aerosols in the Atmosphere. *Science*, **326 (5959)**, 1525–1529, doi:10.1126/science.1180353, URL http://www.sciencemag.org/cgi/doi/10.1126/science.1180353.

Jones, A., D. L. Roberts, and A. Slingo, 1994: A climate model study of indirect radiative forcing by anthropogenic sulphate aerosols. *Nature*, **370 (6489)**, 450–453, doi:10.1038/370450a0, URL http://www.nature.com/doifinder/10.1038/370450a0.

Jones, A., D. L. Roberts, M. J. Woodage, and C. E. Johnson, 2001: Indirect sulphate aerosol forcing in a climate model with an interactive sulphur cycle. *Journal of Geophysical Research: Atmospheres*, **106 (D17)**, 20 293–20 310, doi:10.1029/2000JD000089, URL http://doi.wiley.com/10.1029/2000JD000089.

Karydis, V. A., S. L. Capps, A. G. Russell, and A. Nenes, 2012: Adjoint sensitivity of global cloud droplet number to aerosol and dynamical parameters. *Atmospheric Chemistry and Physics*, **12 (19)**, 9041–9055, doi:10.5194/acp-12-9041-2012, URL http://www.atmos-chem-phys.net/12/9041/2012/.

Karydis, V. a., P. Kumar, D. Barahona, I. N. Sokolik, and A. Nenes, 2011: On the effect of dust particles on global cloud condensation nuclei and cloud droplet number. *Journal of Geophysical Research: Atmospheres*, **116 (D23)**, n/a–n/a, doi:10.1029/2011JD016283, URL http://doi.wiley.com/10.1029/2011JD016283.

Khairoutdinov, M., and Y. Kogan, 2000: A New Cloud Physics Parameterization in a Large-Eddy Simulation Model of Marine Stratocumulus. *Monthly Weather Review*, **128 (1)**, 229–243, doi:10.1175/1520-0493(2000)128<0229:ANCPPI>2.0.CO;2.

Khvorostyanov, V. I., and J. a. Curry, 2006: Aerosol size spectra and CCN activity spectra: Reconciling the lognormal, algebraic, and power laws. *Journal of Geophysical Research*, **111 (D12)**, D12 202, doi:10.1029/2005JD006532, URL http://www.agu.org/pubs/crossref/2006/2005JD006532.shtml.

Khvorostyanov, V. I., and J. A. Curry, 2008: Kinetics of Cloud Drop Formation and Its Parameterization for Cloud and Climate Models. *Journal of the Atmospheric Sciences*, **65 (9)**, 2784–2802, doi:10.1175/2008JAS2606.1, URL http://journals.ametsoc.org/doi/abs/10.1175/2008JAS2606.1.

Kim, D., C. Wang, A. M. L. Ekman, M. C. Barth, and D.-I. Lee, 2014: The responses of cloudiness to the direct radiative effect of sulfate and carbonaceous aerosols. *Journal of Geophysical Research: Atmospheres*, n/a–n/a, doi:10.1002/2013JD020529, URL http://doi.wiley.com/10.1002/2013JD020529.

Kim, D., C. Wang, A. M. L. Ekman, M. C. Barth, and P. J. Rasch, 2008: Distribution and direct radiative forcing of carbonaceous and sulfate aerosols in an interactive size-resolving aerosol–climate model. *Journal of Geophysical Research*, **113 (D16)**, D16 309, doi:10.1029/2007JD009756, URL http://www.agu.org/pubs/crossref/2008/2007JD009756.shtml.

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*, **117 (D23)**, D23 204, doi:10.1029/2012JD018588, URL http://doi.wiley.com/10.1029/2012JD018588.

Koren, I., and G. Feingold, 2011: Aerosol-cloud-precipitation system as a predator-prey problem. *Proceedings of the National Academy of Sciences of the United States of America*, **108 (30)**, 12 227–32, doi:10.1073/pnas.1101777108, URL http://www.pnas.org/cgi/content/abstract/108/30/12227http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3145706{&}tool=pmcentrez{&}rendertype=abstract.

Kravitz, B., and Coauthors, 2013: Climate model response from the Geoengineering Model Intercomparison Project (GeoMIP). *Journal of Geophysical Research: Atmospheres*, n/a–n/a, doi:10.1002/jgrd.50646, URL http://doi.wiley.com/10.1002/jgrd.50646.

Kumar, P., I. N. Sokolik, and a. Nenes, 2009: Parameterization of cloud droplet formation for global and regional models: including adsorption activation from insoluble CCN. *Atmospheric Chemistry and Physics*, **9 (7)**, 2517–2532, doi:10.5194/acp-9-2517-2009, URL http://www.atmos-chem-phys.net/9/2517/2009/.

Lamarque, J.-F., G. P. Kyle, M. Meinshausen, K. Riahi, S. J. Smith, D. P. van Vuuren, A. J. Conley, and F. Vitt, 2011: Global and regional evolution of short-lived radiatively-active gases and aerosols in the Representative Concentration Pathways. *Climatic Change*, **109 (1-2)**, 191–212, doi:10.1007/s10584-011-0155-0, URL http://link.springer.com/10.1007/s10584-011-0155-0.

Lance, S., A. Nenes, and T. A. Rissman, 2004: Chemical and dynamical effects on cloud droplet number: Implications for estimates of the aerosol indirect effect. *Journal of Geophysical Research*, **109**, D22 208, doi:10.1029/2004JD004596, URL http://onlinelibrary.wiley.com/doi/10.1029/2004JD004596/fullhttp://doi.wiley.com/10.1029/2004JD004596.

Latham, J., 1990: Control of global warming? *Nature*, **347 (6291)**, 339–340, doi:10.1038/347339b0, URL http://www.nature.com/doifinder/10.1038/347339b0.

Latham, J., and Coauthors, 2012: Marine cloud brightening. *Philosophical transactions. Series A, Mathematical, physical, and engineering sciences*, **370 (1974)**, 4217–62, doi:10.1098/rsta.2012.0086, URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3405666{&}tool=pmcentrez{&}rendertype=abstract.

Leaitch, W. R., J. W. Strapp, G. A. Isaac, and J. G. Hudson, 1986: Cloud droplet nucleation and cloud scavenging of aerosol sulphate in polluted atmospheres. *Tellus B*, **38B (5)**, 328–344, doi:10.1111/j.1600-0889.1986.tb00258.x, URL http://www.tellusb.net/index.php/tellusb/article/view/15141.

Leibensperger, E. M., and Coauthors, 2011: Climatic effects of 1950–2050 changes in US anthropogenic aerosols – Part 1: Aerosol trends and radiative forcing. *Atmospheric Chemistry and Physics Discussions*, **11 (8)**, 24 085–24 125, doi:10.5194/acpd-11-24085-2011, URL http://www.atmos-chem-phys-discuss.net/11/24085/2011/.

Levy, H., L. W. Horowitz, M. D. Schwarzkopf, Y. Ming, J.-C. Golaz, V. Naik, and V. Ramaswamy, 2013a: The roles of aerosol direct and indirect effects in past and future climate change. *Journal of Geophysical Research: Atmospheres*, **118 (10)**, 4521–4532, doi:10.1002/jgrd.50192, URL http://doi.wiley.com/10.1002/jgrd.50192.

Levy, R. C., S. Mattoo, L. A. Munchak, L. A. Remer, A. M. Sayer, F. Patadia, and N. C. Hsu, 2013b: The Collection 6 MODIS aerosol products over land and ocean. *Atmospheric Measurement Techniques*, **6 (11)**, 2989–3034, doi:10.5194/amt-6-2989-2013, URL http://www.atmos-meas-tech.net/6/2989/2013/.

Liu, X., J. E. Penner, S. J. Ghan, and M. Wang, 2007: Inclusion of ice microphysics in the NCAR Community Atmospheric Model version 3 (CAM3). *Journal of Climate*, **20**, 4526–4547, doi:10.1175/JCLI4264.1.

Liu, X., and Coauthors, 2012: Toward a minimal representation of aerosols in climate models: description and evaluation in the Community Atmosphere Model CAM5. *Geoscientific Model Development*, **5 (3)**, 709–739, doi:10.5194/gmd-5-709-2012, URL http://www.geosci-model-dev.net/5/709/2012/.

Lohmann, U., and J. Feichter, 1997: Impact of sulfate aerosols on albedo and lifetime of clouds: A sensitivity study with the ECHAM4 GCM. *Journal of Geophysical Research: Atmospheres*, **102 (D12)**, 13 685–13 700, doi:10.1029/97JD00631, URL http://doi.wiley.com/10.1029/97JD00631.

Lohmann, U., and J. Feichter, 2004: Global indirect aerosol effects: a review. *Atmospheric Chemistry and Physics Discussions*, **4 (6)**, 7561–7614, doi:10.5194/acpd-4-7561-2004, URL http://hal.archives-ouvertes.fr/hal-00295633/.

Lohmann, U., J. Feichter, C. C. Chuang, and J. E. Penner, 1999: Prediction of the number of cloud droplets in the ECHAM GCM. *Journal of Geophysical Research*, **104 (D8)**, 9169–9198, URL http://onlinelibrary.wiley.com/doi/10.1029/1999JD900046/full.

Lohmann, U., and S. Ferrachat, 2010: Impact of parametric uncertainties on the present-day climate and on the anthropogenic aerosol effect. *Atmospheric Chemistry and Physics*, **10 (23)**, 11 373–11 383, doi:10.5194/acp-10-11373-2010, URL http://www.atmos-chem-phys.net/10/11373/2010/.

Lucas, D. D., and R. G. Prinn, 2005: Parametric sensitivity and uncertainty analysis of dimethylsulfide oxidation in the clear-sky remote marine boundary layer. *Atmospheric Chemistry and Physics*, **5 (6)**, 1505–1525, doi:10.5194/acp-5-1505-2005, URL http://www.atmos-chem-phys.net/5/1505/2005/.

Ma, X., F. Yu, and J. Quaas, 2014: Reassessment of satellite-based estimate of aerosol climate forcing. *Journal of Geophysical Research: Atmospheres*, **119 (17)**, 10,394–10,409, doi:10.1002/2014JD021670, URL http://doi.wiley.com/10.1002/2014JD021670.

Mann, G. W., and Coauthors, 2014: Intercomparison and evaluation of global aerosol microphysical properties among AeroCom models of a range of complexity. *Atmospheric Chemistry and Physics*, **14 (9)**, 4679–4713, doi:10.5194/acp-14-4679-2014, URL http://www.atmos-chem-phys.net/14/4679/2014/.

Mayer, M., C. Wang, M. Webster, and R. G. R. Prinn, 2000: Linking Local Air Pollution to Global Chemistry and Climate. *Journal of Geophysical Research*, **105 (63)**, 22 869, doi:10.1029/2000JD900307, URL http://doi.wiley.com/10.1029/2000JD900307http://www.agu.org/pubs/crossref/2000.../2000JD900307.shtml.

McComiskey, A., and G. Feingold, 2008: Quantifying error in the radiative forcing of the first aerosol indirect effect. *Geophysical Research Letters*, **35 (2)**, L02 810, doi:10.1029/2007GL032667, URL http://doi.wiley.com/10.1029/2007GL032667.

McComiskey, A., and G. Feingold, 2012: The scale problem in quantifying aerosol indirect effects. *Atmospheric Chemistry and Physics*, **12 (2)**, 1031–1049, doi:10.5194/acp-12-1031-2012, URL http://www.atmos-chem-phys.net/12/1031/2012/.

Mcfiggans, G., and Coauthors, 2006: The effect of physical and chemical aerosol properties on warm cloud droplet activation. *Atmospheric Chemistry and Physics*, **6 (9)**, 2593–2649, doi:doi:10.5194/acp-6-2593-2006.

Meskhidze, N., 2005: Evaluation of a new cloud droplet activation parameterization with in situ data from CRYSTAL-FACE and CSTRIPE. *Journal of Geophysical Research*, **110 (D16)**, D16 202, doi:10.1029/2004JD005703, URL http://doi.wiley.com/10.1029/2004JD005703.

Ming, Y., V. Ramaswamy, L. J. Donner, and V. T. J. Phillips, 2006: A New Parameterization of Cloud Droplet Activation Applicable to General Circulation Models. *Journal of the Atmospheric Sciences*, **63 (4)**, 1348–1356, doi:10.1175/JAS3686.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JAS3686.1.

Moore, R. H., V. A. Karydis, S. L. Capps, T. L. Lathem, and A. Nenes, 2013: Droplet number uncertainties associated with CCN: an assessment using observations and a global model adjoint. *Atmospheric Chemistry and Physics*, **13 (8)**, 4235–4251, doi:10.5194/acp-13-4235-2013, URL http://www.atmos-chem-phys.net/13/4235/2013/acp-13-4235-2013.html.

Morales, R., and a. Nenes, 2010: Characteristic updrafts for computing distribution-averaged cloud droplet number and stratocumulus cloud properties. *Journal of Geophysical Research*, **115 (D18)**, D18 220, doi:10.1029/2009JD013233, URL http://doi.wiley.com/10.1029/2009JD013233.

Morales, R., a. Nenes, H. Jonsson, R. C. Flagan, and J. H. Seinfeld, 2011: Evaluation of an entraining droplet activation parameterization using in situ cloud data. *Journal of Geophysical Research: Atmospheres*, **116 (15)**, 1–9, doi:10.1029/2010JD015324.

Morales Betancourt, R., and a. Nenes, 2014a: Droplet activation parameterization: the population splitting concept revisited. *Geoscientific Model Development Discussions*, **7 (3)**, 2903–2932, doi:10.5194/gmdd-7-2903-2014, URL http://www.geosci-model-dev-discuss.net/7/2903/2014/.

Morales Betancourt, R., and A. Nenes, 2014b: Understanding the contributions of aerosol properties and parameterization discrepancies to droplet number variability in a global climate model. *Atmospheric Chemistry and Physics*, **14 (9)**, 4809–4826, doi:10.5194/acp-14-4809-2014, URL http://www.atmos-chem-phys.net/14/4809/2014/.

Morrison, H., G. de Boer, G. Feingold, J. Harrington, M. D. Shupe, and K. Sulia, 2011: Resilience of persistent Arctic mixed-phase clouds. *Nature Geoscience*, **5 (1)**, 11–17, doi:10.1038/ngeo1332, URL http://www.nature.com/ngeo/journal/v5/n1/full/ngeo1332.html{#}ref69.

Morrison, H., and A. Gettelman, 2008: A New Two-Moment Bulk Stratiform Cloud Microphysics Scheme in the Community Atmosphere Model, Version 3 (CAM3). Part I: Description and Numerical Tests. *Journal of Climate*, **21 (15)**, 3642–3659, doi:10.1175/2008JCLI2105.1, URL http://journals.ametsoc.org/doi/abs/10.1175/2008JCLI2105.1.

Murphy, D. M., S. Solomon, R. W. Portmann, K. H. Rosenlof, P. M. Forster, and T. Wong, 2009: An observationally based energy balance for the Earth since 1950. *Journal of Geophysical Research*, **114 (D17)**, D17 107, doi:10.1029/2009JD012105, URL http://doi.wiley.com/10.1029/2009JD012105.

Myhre, G., and Coauthors, 2013: Anthropogenic and Natural Radiative Forcing. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, URL www.climatechange2013.org, 659–740 pp.

Neale, R. B., and Coauthors, 2012: Description of the NCAR Community Atmosphere Model (CAM 5.0). NCAR Technical Notes. *Ncar/Tn-464+Str*, 214.

Nenes, A., R. J. Charlson, M. C. Facchini, M. Kulmala, A. Laaksonen, and J. H. Seinfeld, 2002: Can chemical effects on cloud droplet number rival the first indirect effect? *Geophysical Research Letters*, **29 (17)**, 1848, doi:10.1029/2002GL015295, URL http://www.agu.org/pubs/crossref/2002/2002GL015295.shtml.

Nenes, A., S. Ghan, H. Abdul-Razzak, P. Y. Chuang, and J. H. Seinfeld, 2001: Kinetic limitations on cloud droplet formation and impact on cloud albedo. *Tellus Series B-Chemical and Physical Meteorology*, **53 (2)**, 133–149, doi:10.1034/j.1600-0889.2001.d01-12.x, URL http://www.tellusb.net/index.php/tellusb/article/view/16569http://www.tellusb.net/index.php/tellusb/article/view/16569$\delimiter"026E30F$n{\T1\textless}GotoISI{\T1\textgreater}://WOS:000169810700003.

Nenes, A., and J. H. Seinfeld, 2003: Parameterization of cloud droplet formation in global climate models. *Journal of Geophysical Research*, **108 (D14)**, 4415, doi:10.1029/2002JD002911, URL http://www.agu.org/pubs/crossref/2003/2002JD002911.shtml.

Ovtchinnikov, M., and S. J. Ghan, 2005: Parallel simulations of aerosol influence on clouds using cloud-resolving and single-column models. *Journal of Geophysical Research D: Atmospheres*, **110 (15)**, 1–13, doi:10.1029/2004JD005088.

Paine, R. T., 1966: Food Web Complexity and Species Diversity. *The American Naturalist*, **100 (910)**, 65–75, doi:10.1086/282400, URL http://www.journals.uchicago.edu/doi/10.1086/282400.

Pan, W., M. A. Tatang, G. J. McRae, and R. G. Prinn, 1997: Uncertainty analysis of direct radiative forcing by anthropogenic sulfate aerosols. *Journal of Geophysical Research*, **102 (D18)**, 21 915, doi:10.1029/97JD01653, URL http://doi.wiley.com/10.1029/97JD01653.

Park, S., and C. S. Bretherton, 2009: The University of Washington shallow convection and moist turbulence schemes and their impact on climate simulations with the community atmosphere model. *Journal of Climate*, **22 (12)**, 3449–3469, doi:10.1175/2008JCLI2557.1.

Pawlowska, H., and J.-L. Brenguier, 2000: Microphysical properties of stratocumulus clouds during ACE-2. *Tellus B*, **52 (2)**, 868–887, doi:10.1034/j.1600-0889.2000.00076.x, URL http://www.blackwell-synergy.com/links/doi/10.1034/j.1600-0889.2000.00076.x.

Penner, J. E., L. Xu, and M. Wang, 2011: Satellite methods underestimate indirect climate forcing by aerosols. *Proceedings of the National Academy of Sciences of the United States of America*, **108**, 13 404–13 408, doi:10.1073/pnas.1018526108.

Penner, J. E., and Coauthors, 2006: Model intercomparison of indirect aerosol effects. *Atmospheric Chemistry and Physics*, **6 (11)**, 3391–3405, doi:10.5194/acp-6-3391-2006, URL http://www.atmos-chem-phys.net/6/3391/2006/.

Petters, M. D., and S. M. Kreidenweis, 2007: A single parameter representation of hygroscopic growth and cloud condensation nucleus activity. *Atmospheric Chemistry and Physics*, **7 (8)**, 1961–1971, doi:10.5194/acp-7-1961-2007, URL http://www.atmos-chem-phys.net/7/1961/2007/.

Pierrehumbert, R. T., 2010: *Principles of Planetary Climate*. Cambridge University Press, 674 pp.

Platnick, S., and S. Twomey, 1994: Determining the susceptibility of cloud albedo to changes in droplet concentration with the Advanced Very High Resolution Spectrometer. 334–347 pp., doi:10.1175/1520-0450(1994)033<0334:DTSOCA>2.0.CO;2.

Pruppacher, H. R., and J. D. Klett, 1997: *Microphysics of Clouds and Precipitation (2 ed.)*. Kluwer Academic Publishers.

Quaas, J., O. Boucher, N. Bellouin, and S. Kinne, 2008: Satellite-based estimate of the direct and indirect aerosol climate forcing. *Journal of Geophysical Research: Atmospheres*, **113 (D5)**, n/a–n/a, doi:10.1029/2007JD008962, URL http://doi.wiley.com/10.1029/2007JD008962.

Quaas, J., O. Boucher, N. Bellouin, and S. Kinne, 2011: Which of satellite- or model-based estimates is closer to reality for aerosol indirect forcing? *Proceedings of the National Academy of Sciences of the United States of America*, **108 (46)**, E1099; author reply E1100–1, doi:10.1073/pnas.1114634108, URL http://www.pubmedcentral.nih.gov/articlerender.fcgi?artid=3219127{&}tool=pmcentrez{&}rendertype=abstract.

Quaas, J., O. Boucher, and U. Lohmann, 2006: Constraining the total aerosol indirect effect in the LMDZ and ECHAM4 GCMs using MODIS satellite data. *Atmospheric Chemistry and Physics*, **6 (4)**, 947–955, doi:10.5194/acp-6-947-2006, URL http://www.atmos-chem-phys.net/6/947/2006/acp-6-947-2006.html?FrameEngine=falsehttp://www.atmos-chem-phys.net/6/947/2006/.

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, doi:10.5194/acp-9-8697-2009, URL http://www.atmos-chem-phys.net/9/8697/2009/.

Raatikainen, T., and Coauthors, 2013: Worldwide data sets constrain the water vapor uptake coefficient in cloud formation. *Proceedings of the National Academy of Sciences of the United States of America*, **110 (10)**, 3760–4, doi:10.1073/pnas.1219591110, URL http://www.pnas.org/content/110/10/3760.short.

Romakkaniemi, S., and Coauthors, 2012: Effect of aerosol size distribution changes on AOD, CCN and cloud droplet concentration: Case studies from Erfurt and Melpitz, Germany. *Journal of Geophysical Research: Atmospheres*, **117 (D7)**, n/a–n/a, doi:10.1029/2011JD017091, URL http://doi.wiley.com/10.1029/2011JD017091.

Rothenberg, D., and C. Wang, 2016: Metamodeling of Droplet Activation for Global Climate Models. *Journal of the Atmospheric Sciences*, **73 (3)**, 1255–1272, doi:10.1175/JAS-D-15-0223.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JAS-D-15-0223.1http://journals.ametsoc.org/doi/10.1175/JAS-D-15-0223.1.

Rotstayn, L. D., and Y. Liu, 2005: A smaller global estimate of the second indirect aerosol effect. *Geophysical Research Letters*, **32 (5)**, 1–4, doi:10.1029/2004GL021922.

Saleeby, S. M., and W. R. Cotton, 2004: A Large-Droplet Mode and Prognostic Number Concentration of Cloud Droplets in the Colorado State University Regional Atmospheric Modeling System (RAMS). Part I: Module Descriptions and Supercell Test Simulations. *Journal of Applied Meteorology*, **43 (1)**, 182–195, doi:10.1175/1520-0450(2004)043<0182:ALMAPN>2.0.CO;2, URL http://journals.ametsoc.org/doi/abs/10.1175/1520-0450(2004)043{\T1\textless}0182:ALMAPN{\T1\textgreater}2.0.CO;2.

Saleeby, S. M., and S. C. van den Heever, 2013: Developments in the CSU-RAMS Aerosol Model: Emissions, Nucleation, Regeneration, Deposition, and Radiation. *Journal of Applied Meteorology and Climatology*, **52 (12)**, 2601–2622, doi:10.1175/JAMC-D-12-0312.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JAMC-D-12-0312.1.

Saltelli, A., P. Annoni, I. Azzini, F. Campolongo, M. Ratto, and S. Tarantola, 2010: Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity

index. *Computer Physics Communications*, **181 (2)**, 259–270, doi:10.1016/j.cpc.2009.09.018, URL http://linkinghub.elsevier.com/retrieve/pii/S0010465509003087.

Samset, B. H., and Coauthors, 2016: Fast and slow precipitation responses to individual climate forcers: A PDRMIP multi-model study. *Geophysical Research Letters*, n/a–n/a, doi:10.1002/2016GL068064, URL http://doi.wiley.com/10.1002/2016GL068064.

Scanza, R. A., N. Mahowald, S. Ghan, C. S. Zender, J. F. Kok, X. Liu, and Y. Zhang, 2014: Modeling dust as component minerals in the Community Atmosphere Model: development of framework and impact on radiative forcing. *Atmospheric Chemistry and Physics Discussions*, **14 (12)**, 17 749–17 816, doi:10.5194/acpd-14-17749-2014, URL http://www.atmos-chem-phys-discuss.net/14/17749/2014/.

Seinfeld, J. H., and S. N. Pandis, 2006: *Atmospheric Chemistry and Physics: From Air Pollution to Climate Change*, Vol. 2nd. Wiley, 1203 pp., URL http://www.knovel.com/knovel2/Toc.jsp?BookID=2126.

Sellers, W. D., 1969: A Global Climatic Model Based on the Energy Balance of the Earth-Atmosphere System. *Journal of Applied Meteorology*, **8 (3)**, 392–400, doi:10.1175/1520-0450(1969)008<0392:AGCMBO>2.0.CO;2, URL http://journals.ametsoc.org/doi/abs/10.1175/1520-0450(1969)008{\T1\textless}0392:AGCMBO{\T1\textgreater}2.0.CO;2.

Shindell, D. T., 2014: Inhomogeneous forcing and transient climate sensitivity. *Nature Climate Change*, **4 (4)**, 18–21, doi:10.1038/NCLIMATE2136.

Shindell, D. T., and Coauthors, 2013: Radiative forcing in the ACCMIP historical and future climate simulations. *Atmospheric Chemistry and Physics*, **13 (6)**, 2939–2974, doi:10.5194/acp-13-2939-2013, URL http://www.atmos-chem-phys.net/13/2939/2013/.

Shipway, B. J., 2015: Revisiting Twomey's approximation for peak supersaturation. 3803–3814, doi:10.5194/acp-15-3803-2015.

Shipway, B. J., and S. J. Abel, 2010: Analytical estimation of cloud droplet nucleation based on an underlying aerosol population. *Atmospheric Research*, **96 (2-3)**, 344–355, doi:10.1016/j.atmosres.2009.10.005, URL http://linkinghub.elsevier.com/retrieve/pii/S0169809509002798.

Simpson, E., P. Connolly, and G. McFiggans, 2014: An investigation into the performance of four cloud droplet activation parameterisations. *Geoscientific Model Development*, **7 (4)**, 1535–1542, doi:10.5194/gmd-7-1535-2014, URL http://www.geosci-model-dev.net/7/1535/2014/.

Sobol', I., 2001: Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, **55 (1-3)**, 271–280, doi:10.1016/S0378-4754(00)00270-6, URL http://dx.doi.org/10.1016/S0378-4754(00)00270-6http://linkinghub.elsevier.com/retrieve/pii/S0378475400002706.

Sorjamaa, R., and A. Laaksonen, 2007: The effect of H 2 O adsorption on cloud drop activation of insoluble particles : a theoretical framework. **(2006)**, 6175–6180.

Squires, P., and S. Twomey, 1961: The relation between cloud drop numbers and the spectrum of cloud nuclei. *Physics of Precipitation, Mongraph, No. 5*, American Geophysical Union, Washington, DC, 211–219.

Stevens, B., 2015: Rethinking the lower bound on aerosol radiative forcing. *Journal of Climate*, 150319112745003, doi:10.1175/JCLI-D-14-00656.1, URL http://journals.ametsoc.org/doi/abs/10.1175/JCLI-D-14-00656.1.

Stevens, B., and G. Feingold, 2009: Untangling aerosol effects on clouds and precipitation in a buffered system. *Nature*, **461 (7264)**, 607–613, doi:10.1038/nature08281, URL http://www.ncbi.nlm.nih.gov/pubmed/19794487.

Stier, P., 2016: Limitations of passive remote sensing to constrain global cloud condensation nuclei. *Atmospheric Chemistry and Physics*, **16 (10)**, 6595–6607, doi:10.5194/acp-16-6595-2016, URL http://www.atmos-chem-phys.net/16/6595/2016/.

Storelvmo, T., U. Lohmann, and R. Bennartz, 2009: What governs the spread in shortwave forcings in the transient IPCC AR4 models? *Geophysical Research Letters*, **36 (1)**, 1–5, doi:10.1029/2008GL036069.

Sudret, B., 2008: Global sensitivity analysis using polynomial chaos expansions. *Reliability Engineering & System Safety*, **93 (7)**, 964–979, doi:10.1016/j.ress.2007.04.002, URL http://dx.doi.org/10.1016/j.ress.2007.04.002http://linkinghub.elsevier.com/retrieve/pii/S0951832007001329.

Tao, W.-k., J.-p. Chen, Z. Li, C. Wang, and C. Zhang, 2011: Impact of Aerosols on Convective Clouds and Precipitation. *Reviews of Geophysics*, **50 (2)**, 1–162, doi:10.1029/2011RG000369, URL http://www.agu.org/pubs/crossref/2012/2011RG000369.shtml.

Tatang, M., W. Pan, R. Prinn, and G. McRae, 1997: An efficient method for parametric uncertainty analysis of numerical geophysical models. *Journal of Geophysical Research*, **102**, 925–932, URL http://www.agu.org/pubs/crossref/1997/97JD01654.shtml.

Tibshirani, R., 2011: Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **73 (3)**, 273–282, doi:10.1111/j.1467-9868.2011.00771.x, URL http://doi.wiley.com/10.1111/j.1467-9868.2011.00771.x.

Tibshirani, R., I. Johnstone, T. Hastie, and B. Efron, 2004: Least angle regression. *The Annals of Statistics*, **32 (2)**, 407–499, doi:10.1214/009053604000000067, URL http://projecteuclid.org/Dienst/getRecord?id=euclid.aos/1083178935/.

Topping, D., P. Connolly, and G. McFiggans, 2013: Cloud droplet number enhanced by co-condensation of organic vapours. *Nature Geoscience*, **6 (6)**, 443–446, doi:10.1038/ngeo1809, URL http://www.nature.com/doifinder/10.1038/ngeo1809.

Topping, D. O., and G. McFiggans, 2012: Tight coupling of particle size, number and composition in atmospheric cloud droplet activation. *Atmospheric Chemistry and Physics*, **12**, 3253–3260, doi:10.5194/acp-12-3253-2012.

Twomey, S., 1959: The nuclei of natural cloud formation part II: the supersaturation in natural clouds and the variation of droplet concentration. *Pure and Applied Geophysics*, URL http://www.springerlink.com/index/Q0N7L477832V7170.pdf.

Twomey, S., 1974: Pollution and the planetary albedo. *Atmospheric Environment (1967)*, **8 (12)**, 1251–1256, doi:10.1016/0004-6981(74)90004-3.

Twomey, S., 1977: The Influence of Pollution on the Shortwave Albedo of Clouds. *Journal of the Atmospheric Sciences*, **34 (7)**, 1149–1152, doi:10.1175/1520-0469(1977)034<1149:TIOPOT>2.0.CO;2.

Twomey, S., 1991: Aerosols, clouds and radiation. *Atmospheric Environment Part A, General Topics*, **25 (11)**, 2435–2442, doi:10.1016/0960-1686(91)90159-5.

Vehkamäki, H., 2002: An improved parameterization for sulfuric acid–water nucleation rates for tropospheric and stratospheric conditions. *Journal of Geophysical Research*, **107 (D22)**, 4622, doi:10.1029/2002JD002184, URL http://doi.wiley.com/10.1029/2002JD002184.

Wang, C., 2004: A modeling study on the climate impacts of black carbon aerosols. *Journal of Geophysical Research*, **109 (D3)**, D03 106, doi:10.1029/2003JD004084, URL http://www.agu.org/pubs/crossref/2004/2003JD004084.shtml.

Wang, C., 2005a: A modeling study of the response of tropical deep convection to the increase of cloud condensation nuclei concentration: 1. Dynamics and microphysics. *Journal of Geophysical Research*, **110 (D21)**, 1–16, doi:10.1029/2004JD005720, URL http://www.agu.org/pubs/crossref/2005/2004JD005720.shtml.

Wang, C., 2005b: A modeling study of the response of tropical deep convection to the increase of cloud condensation nuclei concentration: 2. Radiation and tropospheric chemistry. *Journal of Geophysical Research*, **110 (D22)**, 1–15, doi:10.1029/2005JD005829, URL http://www.agu.org/pubs/crossref/2005/2005JD005829.shtml.

Wang, C., 2015: Anthropogenic Aerosols and the Distribution of Past Large-Scale Precipitation Change. *Geophysical Research Letters*, n/a–n/a, doi:10.1002/2015GL066416, URL http://doi.wiley.com/10.1002/2015GL066416.

Wang, M., and Coauthors, 2011: The multi-scale aerosol-climate model PNNL-MMF: model description and evaluation. *Geoscientific Model Development*, **4 (1)**, 137–168, doi:10.5194/gmd-4-137-2011, URL http://www.geosci-model-dev.net/4/137/2011/.

Wang, Y., J. H. Jiang, and H. Su, 2015: Atmospheric responses to the redistribution of anthropogenic aerosols. *Journal of Geophysical Research: Atmospheres*, **120 (18)**, 9625–9641, doi:10.1002/2015JD023665.

Ward, D. S., T. Eidhammer, W. R. Cotton, and S. M. Kreidenweis, 2010: The role of the particle size distribution in assessing aerosol composition effects on simulated droplet activation. *Atmospheric Chemistry and Physics*, **10 (12)**, 5435–5447, doi:10.5194/acp-10-5435-2010, URL http://atmos-chem-phys-discuss.net/10/4189/2010/acpd-10-4189-2010.pdfhttp://www.atmos-chem-phys.net/10/5435/2010/.

Weirs, V. G., J. R. Kamm, L. P. Swiler, S. Tarantola, M. Ratto, B. M. Adams, W. J. Rider, and M. S. Eldred, 2012: Sensitivity analysis techniques applied to a system of hyperbolic conservation laws. *Reliability Engineering & System Safety*, **107**, 157–170, doi:10.1016/j.ress.2011.12.008, URL http://linkinghub.elsevier.com/retrieve/pii/S0951832011002717.

West, R. E. L., P. Stier, A. Jones, C. E. Johnson, G. W. Mann, N. Bellouin, D. G. Partridge, and Z. Kipling, 2014: The importance of vertical velocity variability for estimates of the indirect aerosol effects. *Atmospheric Chemistry and Physics*, **14 (12)**, 6369–6393, doi:10.5194/acp-14-6369-2014, URL http://www.atmos-chem-phys.net/14/6369/2014/.

Westervelt, D. M., L. W. Horowitz, V. Naik, J. C. Golaz, and D. L. Mauzerall, 2015: Radiative forcing and climate response to projected 21st century aerosol decreases. *Atmospheric Chemistry and Physics*, **15 (22)**, 12 681–12 703, doi:10.5194/acp-15-12681-2015.

Whitby, K. T., 1978: The physical characteristics of sulfur aerosols. *Atmospheric Environment (1967)*, **12 (1-3)**, 135–159, doi:10.1016/0004-6981(78)90196-8, URL http://linkinghub.elsevier.com/retrieve/pii/S1352231007009739http://linkinghub.elsevier.com/retrieve/pii/0004698178901968.

Wilson, J., C. Cuvelier, and F. Raes, 2001: A modeling study of global mixed aerosol fields. *Journal of Geophysical Research*, **106 (D24)**, 34 081–34 108, doi:10.1029/2000JD000198, URL http://www.agu.org/pubs/crossref/2001/2000JD000198.shtml.

Zhang, S., and Coauthors, 2016: On the characteristics of aerosol indirect effect based on dynamic regimes in global climate models. *Atmospheric Chemistry and Physics*, **16 (5)**, 2765–2783, doi:10.5194/acp-16-2765-2016, URL http://www.atmos-chem-phys.net/16/2765/2016/.