


Improving Estimates of the
Atmospheric Oxidative Capacity
and
Amazon Fire Emissions

Improving Estimates of the Atmospheric Oxidative Capacity
and Amazon Fire Emissions

Stijn Naus

Stijn Naus



Propositions

1. When source-specific emissions of methane are quantified, the role of the hydroxyl radical is too often ignored.
(this thesis)
2. Correct use of box models requires a deeper understanding of atmospheric dynamics than correct use of 3-D transport models.
(this thesis)
3. Debating with a climate skeptic should be mandatory for climate science students.
4. Overselling in science has important benefits.
5. The solution to climate change has long moved away from climate science.
6. The power of deadlines is severely underused in PhD projects.
7. The COVID pandemic demonstrates society's inability to anticipate crises that were forecast by scientific consensus.
8. The last sentence is more important than the first.

Propositions belonging to the thesis entitled

Improving estimates of the atmospheric oxidative capacity and Amazon fire emissions

Stijn Naus

Wageningen, 3 February 2021

Improving estimates of
the atmospheric oxidative capacity
and
Amazon fire emissions

Stijn Naus

Thesis committee

Promotors:

Prof. Dr Maarten C. Krol
Professor of Air Quality and Atmospheric Chemistry
Wageningen University & Research
Professor of Atmospheric Physics and Chemistry
Utrecht University

Prof. Dr Wouter Peters
Personal chair, Air Quality and Atmospheric Chemistry
Wageningen University & Research
Professor of Atmospheric Composition Modeling
University of Groningen

Co-promotors:

Dr Stephen A. Montzka
Research chemist
National Oceanic and Atmospheric Administration

Other members:

Dr R. Teuling, Wageningen University & Research
Dr M. E. Poppa, Utrecht University
Prof. Dr G. van der Werf, Vrije Universiteit Amsterdam
Dr M. Rigby, School of Chemistry, Bristol, UK

This research was conducted under the auspices of the Graduate School for Socio-Economic and Natural Sciences of the Environment (SENSE)

**Improving estimates of
the atmospheric oxidative capacity
and
Amazon fire emissions**

Stijn Naus

Thesis

submitted in fulfilment of the requirements for the degree of doctor
at Wageningen University

by the authority of the Rector Magnificus

Prof. Dr A.P.J. Mol,

in the presence of the

Thesis Committee appointed by the Academic Board

to be defended in public

on Wednesday, 3 February 2021

at 4.00 p.m. in the Aula.

Stijn Naus

Improving estimates of the atmospheric oxidative capacity and Amazon fire emissions
176 pages.

PhD thesis, Wageningen University, Wageningen, NL (2021)

With references, with summary in English

ISBN: 978-94-6395-642-0

DOI: 10.18174/536720

Improving estimates of
the atmospheric oxidative capacity
and
Amazon fire emissions

Stijn Naus

Summary

The composition of the air outside, the atmosphere, is the end-result of complex chemistry, dynamical transport and emissions of gases. Anthropogenic activity, such as fossil fuel combustion, refrigeration and large-scale deforestation, contributes to a rapidly changing composition of the atmosphere, often with adverse effects: air quality issues, global warming and depletion of stratospheric ozone are some of the most complex and impactful problems that society has to face. Many gases are removed from the atmosphere through the process of oxidation and the primary atmospheric oxidant is the hydroxyl radical (OH). Gases removed by OH include important pollutants, such as the potent greenhouse gas methane (CH_4) and urban pollutants such as CO and NO_x . Therefore, the oxidation capacity of the atmosphere, as determined by the "cleansing-agent" OH, is an important quantity that is central to this thesis.

Since OH is very reactive, it has an atmospheric lifetime of seconds and a low atmospheric abundance. This makes it difficult to measure OH concentrations directly and impossible to extrapolate those measurements to larger scales. Consequently, despite the importance and omnipresence of OH, there are still open questions concerning its distribution around the globe and how this distribution has varied over past decades. We know that OH concentrations are highest in the tropics, but the degree of hemispheric symmetry is uncertain. Similarly, earlier work has established that the atmospheric oxidative capacity is well-buffered, but year-to-year variations of a few percent are still possible, for example driven by climate variability. These uncertainties limit our ability to interpret atmospheric budgets of a variety of pollutants, such as that of CH_4 .

In Chapters 2 and 3 of this thesis, we investigate indirect observational constraints on the atmospheric oxidative capacity, mainly through the use of the trace gas methyl chloroform (MCF). MCF is an anthropogenically produced gas, that was used mainly as a solvent in paints and degreasers. The production of MCF was phased out in the Montreal protocol and subsequent amendments (1987–1999), because its emissions harm the stratospheric ozone layer. Fortunately, MCF is removed from the atmosphere mainly through oxidation by OH, resulting in an atmospheric lifetime of 5 to 6 years. Therefore, the production phase-out resulted in a rapid atmospheric decline in MCF abundance of approximately 20%/year, and variations in this rate of decline provide a proxy for large-scale OH varia-

tions.

The atmospheric decline of MCF has been monitored mainly from remote surface sites that well-represent the background atmosphere. Comprehensive interpretation of these observations requires the use of atmospheric models. Forward models simulate emissions of gases into the atmosphere, and their subsequent transport and chemistry. Such models can vary in complexity from a global one-box model to a complex 3D transport model that represents the global atmospheric circulation. In this thesis, we develop and use inverse models, which use observations of a gas such as MCF to estimate, for example, those MCF emissions and OH variations that result in the best agreement with observations.

In Chapter 2, we use a two-box model inversion of the troposphere to explore the constraints that MCF surface observations, from the National Oceanic and Atmospheric Administration (NOAA), place on global-scale OH variations over the 1994–2014 period. CH₄ is also included in the simulations, to help us understand the implications of OH variations for the atmospheric budget of this important greenhouse gas. The two-box model set-up incorporates important aspects of the real atmosphere, because the atmosphere is mixed significantly faster within hemispheres than between them, and because the majority of anthropogenic emissions enter the atmosphere in the Northern Hemisphere. This is why a similar set-up was applied to the MCF–OH problem in two previous studies (Turner et al., 2017; Rigby et al., 2017). However, because the real troposphere does not consist of two boxes, the two-box model approach is susceptible to biases that we quantify through the use of a 3D transport model: TM5. In the forward version of TM5, we simulate the 3D distribution of MCF and CH₄, and we reduce the 3D output of these simulations to a few quantities that can be used as input for a two-box model inversion. We identify time- and tracer-dependent biases in four of these quantities.

Interhemispheric exchange varies interannually because of physical drivers such as the El Niño Southern Oscillation (ENSO). We find, however, that the effective interhemispheric exchange that is used in two-box models is highly dependent on a tracer’s distribution within each hemisphere. The distribution of MCF specifically has changed strongly over the 1994–2014 period that we investigate, due to the rapid drop in its emissions. We find that this change resulted in large ($> 20\%$), multi-annual variations in the interhemispheric exchange of MCF, that are very different from the variations derived for CH₄. The same redistribution of MCF drove a rapid decrease in loss of tropospheric MCF to the stratosphere by up to 70%, as well as a shift in the effective interhemispheric OH ratio that MCF is exposed to. We additionally quantify observational uncertainties in hemispheric averages of MCF and CH₄, which reveals that changes in the hemispheric burdens of the troposphere are strongly constrained from surface sites.

To test the impact of the four identified biases on the OH variations and CH₄ emissions derived in a tropospheric two-box model inversion, we perform two sets of inversions: one in which we correct for the biases, and one in which we do not. Notably, we find that the

inversion that includes bias corrections produces OH variations that show a positive trend, while the standard inversion does not. However, we also find that the uncertainties driven by other two-box parameters that cannot be constrained from a 3D model simulation, for example related to MCF emissions, remain large. The implication for the CH₄ budget is that it remains difficult to attribute variations in the atmospheric CH₄ abundance to either emission or OH changes.

In Chapter 3, we investigate how MCF constraints on OH change if we move the inversion to the 3D transport model TM5. In a TM5-4DVAR inversion, we can include observed MCF gradients within hemispheres, the tropical maximum in OH and compare our results to aircraft observations: all advantages over the two-box model approach. We co-optimize MCF emissions and the latitudinal distribution of OH over the 1998–2018 time period, and we find that small interannual variations in OH ($< 3\%$) without a longterm trend already result in a good match with NOAA surface observations at most sites. The timing and sign of the interannual OH variations are found to be robust with respect to the choice of prior OH and MCF emission distribution, while the amplitude of the variations mostly depends on the degree of convergence and, relatedly, on assumed uncertainties in the observations. The interannual variations in OH show an improved correlation with ENSO ($r=-0.47$; $p=0.05$), compared to results from two-box model inversions, which likely indicates that our 3D model inversion produces the more realistic solution.

However, we also find that, to reproduce observed intrahemispheric gradients of MCF, large adjustments in the latitudinal OH distribution are required, the amplitude of which we consider to be physically unrealistic. From a series of sensitivity tests we conclude that the most likely explanation for these MCF gradients includes a changed ocean flux. Specifically, while the ocean is principally a sink of MCF, earlier work has hypothesized that, in response to the emission drop, oceans at high-latitudes have become a source of MCF, which fits the sign and approximate magnitude of the intrahemispheric bias we find for MCF. Positively, our study presents the first evidence from MCF surface observations for an ocean source. Negatively, the existence of an uncertain ocean source further complicates the derivation of OH variations from MCF observations. However, we consider that the most likely driver of interannual variations in MCF is still its OH oxidation sink. Therefore, we conclude that the reference twenty-year timeseries of OH that we have derived is worth including in future work, for example in global CH₄ inversions.

In Chapter 4, we move away from a global perspective to zoom in on one of Earth’s most precious ecosystems: the Amazon basin. The Amazon is home to the world’s largest rainforests and to a rich biodiversity, but large-scale deforestation and agricultural expansion threaten the ability of this ecosystem to survive in a rapidly changing climate. Every year, during the local dry season, fires burn through the Amazon forests and savanna. It is crucial to understand and monitor these fires, because they are driven by direct, local anthropogenic activity, but their extent is also sensitive to drought intensity and

frequency that might increase due to climate change. Fires emit vast amounts of pollutants into the atmosphere, and we use satellite observations of one such trace gas, CO, to constrain fire emissions over South-America. Specifically, we use the TM5-4DVAR inverse system, together with satellite-observed CO columns of the Measurements of Pollution in the Troposphere (MOPITT) instrument, to optimize reported fire emissions of CO (from the Global Fire Assimilation System; GFAS) over the 2003–2018 period.

MOPITT CO columns over South-America display strong seasonality and interannual variations that are matched in the optimized CO emissions. Additionally, a simulation with optimized fire emissions better reproduces independent aircraft profiles, which were sampled over five sites in the Brazilian Amazon, compared to a simulation with GFAS emissions. This confirms the skill of the inverse system to estimate emissions at sub-Amazon scales. Similarly, we find that we can firmly constrain emissions also at the level of individual Brazilian states, and interannual variations in emissions at state-level correlate well with local soil moisture anomalies. Superimposed on the interannual variations, we find that emissions have decreased between 2003 and 2012, and stabilized afterward. The decrease is especially strong over forest-covered areas (55%), and the timing and magnitude of this decrease is confirmed in deforestation rates reported by the Instituto Nacional de Pesquisas Espaciais (INPE). Optimized emissions are additionally found to be robust with respect to the input fields: for example, inversions based on a climatological fire prior retrieve largely the same interannual variations, even at state-level, as the standard inversions.

These results demonstrate that our inverse system provides strong constraints on Amazon fire emissions that are a product of local anthropogenic activity and natural variability. In principle, our approach can straightforwardly be adapted into an operational system that would be a valuable addition to the existing palette of fire monitoring systems. Due to the integrated signal that CO observations provide, such a system will be less hampered by cloud occurrence or missed understory fires than existing systems that are based on land remote sensing. Moreover, constraints from CO can help better quantify the carbon release during fires.

In conclusion, we present in this thesis how observational constraints in inverse models can improve our understanding of atmospheric oxidation and Amazon fire emissions. Chapter 2 and 3 focus on the former, and signify the respective strengths and weaknesses of simple versus complex, comprehensive models. Our study in Chapter 4 exemplifies an application where the inverse modeling approach works particularly well, because of the high information content of satellite observations and the strong local signal from fires. Future work on either application can likely benefit from the lessons and conclusions presented in this thesis.

Contents

	Page
Summary	vii
Contents	xi
Chapter 1 General Introduction	1
1.1 Outline	2
1.2 Background	2
1.3 Observing atmospheric composition	11
1.4 Modelling the atmosphere and interpretation of observations	17
1.5 Thesis outline	22
Chapter 2 Constraints and biases in a tropospheric two-box model of OH	27
2.1 Introduction	29
2.2 Methods	32
2.3 Results	40
2.4 Discussion	50
2.5 Summary and Conclusions	53
2.S1 Procedure to obtain hemispheric-average MCF mixing ratios	55
2.S2 The nudged simulation	56
2.S3 Interhemispheric asymmetries in correlations between k , OH and CH_4/MCF	58
2.S4 Robustness of derived interhemispheric exchange rates	64
Chapter 3 A 3D-model inversion of methyl chloroform to constrain the atmospheric oxidative capacity	69
3.1 Introduction	71
3.2 Methods	73
3.3 Results	77
3.4 Further discussion	88
3.5 Conclusions	91
3.S1 Tropospheric lifetimes of MCF and CH_4	92
3.S2 High-resolution simulations	93

3.S3 Simulations of SF ₆ and HFC-152a	94
3.S4 Ocean-atmosphere exchange of MCF in TM5	97
3.S5 Ten year (1998-2008) inversions	99
3.S6 The adjoint of OH chemistry	102
3.S7 Tables	104
Chapter 4 An atmospheric perspective on Amazon fires	107
4.1 Introduction	109
4.2 Methods	110
4.3 Results	112
4.4 Discussion	118
4.5 Conclusions	119
4.S1 Match with and sensitivity to assimilated observations	121
4.S2 Sensitivities in the inverse system	126
4.S3 CO columns over the Santarém aircraft site	134
Chapter 5 Synthesis	137
5.1 Introduction	138
5.2 Improving constraints on the atmospheric oxidative capacity	138
5.3 Constraints from CO on Amazonian fire activity	146
5.4 Final outlook	150
References	153
Acknowledgements	171
Graduate school certificate	174

Chapter 1

General Introduction

1.1 Outline

Human influence has changed the chemical composition of the atmosphere in many ways. When you drive any petrol or diesel car, a mix of gases is released from the exhaust pipes. Once emitted, these gases enter the atmosphere and become a part of the air we breathe. Some gases will only remain in the air briefly, barely having time to exit the vehicle exhaust. Others will be with us for decades or centuries, circling the Earth and affecting our environment. Eventually, however, there are processes through which gases are reduced, or removed from the air again. Depending on the gas, it might be absorbed into oceans, destroyed by sunlight, or even react with other gases to form something new. Most generally, the science central to this thesis concerns the emission of gases into the atmosphere, their eventual removal, and how the composition of the atmosphere is affected in between.

More specifically, we focus on one mechanism through which many gases are removed from the atmosphere: oxidation. The most important atmospheric oxidant is the hydroxyl radical (OH) (Leighton, 1961; Levy, 1971), which is dubbed the atmospheric cleansing agent for its integral role in pollutant removal. When large amounts of pollutants are released into the atmosphere, the amount of cleansing agent available for oxidation may in turn be affected. This means that pollutant emissions partly govern how long the emitted pollutants will remain in the air. This cyclical interaction between pollutant and oxidizing agent forms the core of the studies presented in this thesis.

In Section 1.2, we outline the key principles that govern the composition and the dynamics of the atmosphere. Next, we discuss how this dynamical atmosphere is observed and monitored (Section 1.3). Then, we move on to the role of models as tools to help interpret observations and test our understanding (Section 1.4). In the final section of this chapter, Section 1.5, we use the established theoretical framework to outline the research questions of this thesis.

1.2 Background

1.2.1 Atmospheric oxidation

Atmospheric oxidation is one of the most important pathways through which gases are removed and formed in the atmosphere (Jacob, 1999). Oxidation is a chemical reaction in which electrons are transferred from one gas to another: the oxidant. In the atmosphere, this often simply means that two gases exchange an oxygen atom. An intuitive example is the oxidation of carbon monoxide (CO):








Figure 1.1: Pie charts of global sources (left) and sinks (right) of atmospheric OH. Values were retrieved from Table 1 in (Lelieveld et al., 2016). NMVOC is an abbreviation for non-methane volatile organic compounds: a group that includes non-methane hydrocarbons (see Section 1.2.2). Note that in the sink diagram, many reactions are grouped together. For example, the grouped NMVOC sink is larger than the CH₄ sink, but more OH is lost to CH₄ than to any individual NMVOC.

This example shows already two important side-effects of oxidation. Firstly, when one harmful pollutant is removed by oxidation (CO), another might form (CO₂). In fact, a large portion of atmospheric CO itself is formed through oxidation of the greenhouse gas methane (CH₄) (Bergamaschi et al., 2000). Secondly, not only CO, but also OH is consumed in this reaction. Simply speaking, if too much CO enters the atmosphere, we might run out of OH (see sinks in Figure 1.1).

Luckily, the complexity of atmospheric chemistry ensures that there are balancing and buffering mechanisms in place to prevent this from happening (Atkinson, 1986). For example, a near-instantaneous follow-up reaction to 1.1 is:



with M an air molecule (mostly N₂ or O₂) that absorbs excess energy. HO₂ can be then be converted back to OH:



In fact, there are many pathways to convert HO₂ back to OH. For this reason, OH and HO₂ are often grouped as HO_x. The final reaction also exemplifies that while some

pollutants, notably CO, use up atmospheric oxidants, others, such as NO, can actually boost the abundance of OH.

In this way, atmospheric oxidation is buffered, but not perfectly so. There is a balance between OH and HO₂: if too much pressure is put on one side of the balance in the form of excess CO emissions, then not enough NO might be available to recycle HO₂ back to OH. Moreover, there are pathways through which HO₂ is lost without formation of OH. The frequency with which OH is recovered after it is used for oxidation is called its recycling efficiency. For example, a recycling efficiency of 70% indicates that for every ten CO molecules that are oxidized, three OH molecules will be lost. The recycling efficiency is highly dependent on local atmospheric conditions (e.g. the presence of NO), and so it will vary in space and time.

Given that OH is continuously lost due to imperfect recycling, why does it not run out? Or: where does OH come from in the first place? Primarily, OH is formed from the photolysis of ozone (O₃), and the subsequent reaction of the excited oxygen atom with water vapour (H₂O):



This shows that sunlight and H₂O are the main ingredients for the primary formation of OH. In the absence of light, for example during the night, OH concentrations quickly drop to levels near-zero (Hard et al., 1986). In the absence of water vapour, the O(^1D) atom readily reacts with other gases, without producing OH. In fact, even in the presence of water vapour, the O(^1D) atom reacts much more commonly with O₂ or N₂ than with H₂O (Levy, 1971).

Due to its high reactivity, the atmospheric lifetime of OH is on the order of seconds. For this reason, large-scale patterns in the OH distribution emerge only because of large-scale patterns in OH precursors. Consequently, since both solar input and H₂O abundance are highest in the tropics, OH concentrations are highest in the tropics too (e.g. Crutzen & Zimmermann (1991) and Figure 1.2). The role of the tropics in atmospheric oxidation is even more important than the OH distribution suggests, because many oxidation reactions take place more efficiently at higher temperatures. For example, oxidation of methane (see Section 1.2.2) by OH proceeds twice as fast at 33°C than at 0°C (Burkholder et al., 2020).

The combination of primary formation of OH and secondary formation in recycling reactions fuels the atmospheric oxidizing capacity. This manifests as a key property of the

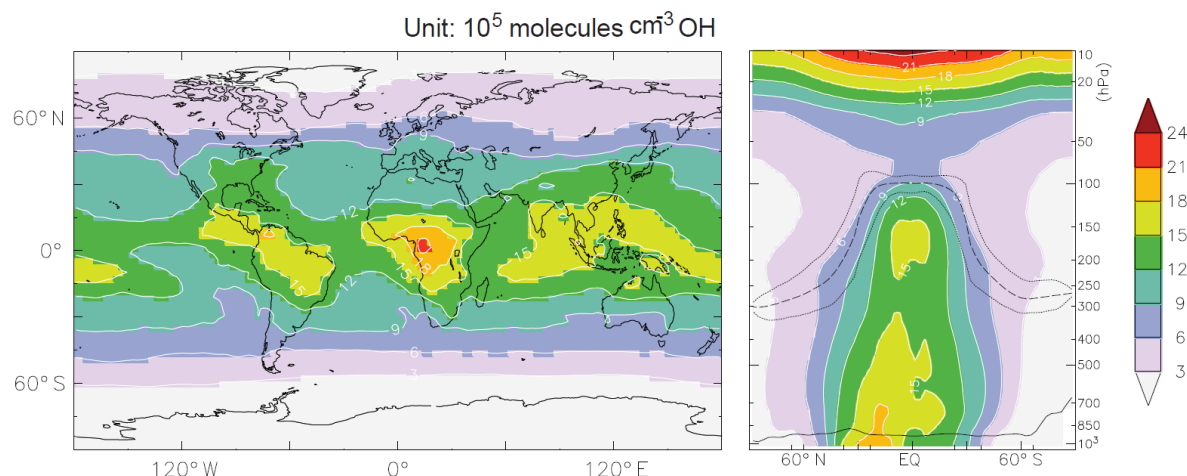


Figure 1.2: The global, annual mean distribution of OH. The left panel shows the horizontal distribution averaged over the troposphere (this concept is introduced in Section 1.2.3), while the right panel shows the vertical distribution OH at different longitudes. In the right panel, the lower solid line indicates the average boundary layer height, the upper dashed line the mean tropopause and the solid lines the annual minimum and maximum tropopause height. Figure taken from (Lelieveld et al., 2016).

atmosphere that determines the atmospheric lifetime of some of the most important pollutants, and consequently how much harm these species inflict on our environment. Figure 1.1 shows that only an estimated 33% of atmospheric OH is formed through primary formation (Reaction 1.4). Therefore, while the global OH distribution follows the distribution of the ingredients required for primary formation (water vapour and sunlight), the global atmospheric oxidative capacity is in large part determined and sustained by secondary formation. The high contribution from secondary formation in the OH budget corresponds to an OH recycling efficiency of approximately 67%. This indicates that on large scales, the atmospheric oxidizing capacity is well-buffered with respect to perturbations in, for example, CO or methane emissions (Lelieveld et al., 2002; Voulgarakis et al., 2013).

In summary, we have seen that the requirements of primary OH formation (sunlight and water vapour) to a large degree determine the atmospheric distribution of OH (Figure 1.2) and that OH concentrations on large scales are well-buffered. However, there are still important open questions surrounding the details of the global OH distribution. Firstly, the annually averaged OH distribution is approximately symmetric between the hemispheres, but the degree of this symmetry has been a topic of debate, with reported Northern to Southern hemisphere OH ratios ranging from 0.85 to 1.4 (Brenninkmeijer et al., 1992; Montzka et al., 2000; Naik et al., 2013; Patra et al., 2014). Secondly, even if OH is well-buffered, interannual variations of a few percent cannot be excluded (Montzka et al., 2011). Such OH variations can strongly impact how we interpret the atmospheric

record of, for example, the greenhouse gas methane (e.g. Rigby et al. (2017); see also Section 1.2.2).

These open questions are leading in Chapters 3 and 4, but first we move on to discuss some trace gases that are relevant to this thesis, and that are all removed through oxidation by OH.

1.2.2 Important atmospheric species

Methane

Methane (CH_4) is the second-most important greenhouse gas, after CO_2 , and it is responsible for up to 20% of direct global warming (Myhre et al., 2013). Anthropogenic emissions of methane are generally unintentional, since methane is a fossil fuel and can be used as a source of energy. Large anthropogenic sources include gas leaks (e.g., Kort et al., 2014; Pandey et al., 2019a) and incomplete digestion in cattle (Johnson & Johnson, 1995). Moreover, methane has an atmospheric lifetime of around nine years (Prather et al., 2012). This means that reducing anthropogenic emissions of CH_4 is both economically interesting and that such measures will quickly reduce atmospheric CH_4 concentrations. Contrastingly, CO_2 , the other major greenhouse gas, is an unavoidable end-product of fossil fuel combustion and is removed from the atmosphere against a typical timescale of a century (Jacob, 1999).

Methane is not only emitted anthropogenically, but has also large natural sources. The single largest natural source of methane to the atmosphere is the anaerobic decomposition of organic material in wetlands (Bartlett & Harriss, 1993). In addition, there are highly uncertain geological emissions (e.g. mud volcanoes) (Etiope & Klusman, 2002; Hmiel et al., 2020) and there are huge reservoirs of methane locked in marine hydrates (Kvenvolden, 1988) and beneath permafrost (Christensen et al., 2004). Both release of CH_4 from these reservoirs and wetland emissions are likely to be affected by increases in temperature and by changes in precipitation patterns that are associated with climate change (O'Connor et al., 2010). Therefore, there is a distinct possibility that climate change can trigger additional release of methane, which in turn will enhance climate change in a positive feedback loop (Dean et al., 2018).

The economic viability of reducing anthropogenic methane emissions, combined with the potential feedbacks between climate change and natural methane emissions have generated much scientific attention for the atmospheric methane budget (Kirschke et al., 2013). Despite this attention, the variations in the atmospheric abundance of methane (see Figure 1.4, bottom right panel) remain poorly understood. For example, in the 1990's, the growth of atmospheric methane leveled off and between 1999 and 2006 the atmospheric abundance of methane was approximately constant. These variations indicated an approach to steady-state: a balance between methane sources and sinks, with neither

changing much (Dlugokencky et al., 1998). Unexpectedly, in 2007, methane concentrations started increasing again (Rigby et al., 2008), and this renewed growth has even accelerated since 2014 (Nisbet et al., 2019). None of these variations were expected before they occurred and they remain difficult to interpret, even in hindsight. The renewed growth has been explained by, among others, agricultural or wetland emissions (Schaefer et al., 2016), fossil fuel emissions (Hausmann et al., 2016; Worden et al., 2017), and, of course, variations in the primary sink of methane: OH (Rigby et al., 2017; Turner et al., 2017). Understanding the main drivers of the renewed and accelerating methane growth is an important first step in developing mitigation strategies that help halt or reverse the methane growth.

In this thesis, we focus on the oxidation sink of atmospheric methane, rather than on its sources. We note also that, as OH is the primary sink of methane, methane is in turn a large sink of OH (see Figure 1.1). The strength of this link between methane and OH can be quantified as follows: the average lifetime of methane in the atmosphere is roughly nine years, but the effective lifetime of a new methane molecule released into the atmosphere can be up to fifteen years (Prather, 1996). This difference manifests because the new methane slightly reduces OH concentrations, which slightly increases the atmospheric lifetime of all the methane already in the atmosphere. The large role that OH plays in the methane budget, combined with the importance of understanding the methane budget, make that one of the primary applications of an improved understanding of the atmospheric oxidative capacity, as we aim to achieve in this thesis, is found in the methane budget. Currently, uncertainty in the oxidation sink of CH_4 translates to substantial uncertainties in global methane emissions (Holmes et al., 2013), as is evident from the plethora of potential explanations for atmospheric growth rate variations of methane.

Carbon Monoxide

We return to the gas from our first oxidation example: carbon monoxide (CO). Atmospheric CO is emitted and formed through two strongly related processes: incomplete combustion of carbon-containing material and atmospheric oxidation of hydrocarbons. Incomplete combustion occurs to some degree in all combustion of fossil fuels, such as in cars and in coal plants, but also in wildfires. Hydrocarbons are gases that contain both hydrogen and carbon atoms and the most obvious example is methane. In fact, oxidation of methane by OH accounts for 30% of all CO formed and emitted in our atmosphere (Bergamaschi et al., 2000). Other hydrocarbons, often referred to as non-methane hydrocarbons (NMHC), are individually less important to the CO budget than methane, but add up to another substantial source of CO (and a sink of OH; see Figure 1.1).

CO has an atmospheric lifetime of a few months (Logan et al., 1981). Since CO is removed primarily by OH, its lifetime in the tropics, where OH concentrations are highest, is

typically one month (Bowman, 2006), while in the dark polar winter the CO lifetime can be much longer. Where CH_4 is the second-largest sink of OH globally, CO is the largest (Figure 1.1). This means that even though CO does not contribute to global warming directly, it does act as an indirect greenhouse gas (Daniel & Solomon, 1998). CO reduces the atmospheric oxidative capacity and, consequently, increases the atmospheric lifetime of methane. Additionally, CO can detrimentally affect local air quality, especially indoors (Bernstein et al., 2008).

We discuss CO in this thesis in two contexts. Firstly, as one of the most important drivers of variations in the atmospheric oxidative capacity. Secondly, in Chapter 4, we use the large quantities of CO injected into the atmosphere during wildfires as a proxy for fire activity in the Amazon.

Methyl Chloroform

The final species we discuss here is methyl chloroform (chemical formula: CH_3CCl_3 ; hereafter referred to as MCF). Same as CH_4 and CO, MCF is removed primarily by reaction with OH, resulting in an atmospheric lifetime of 5-6 years. However, different from the other two gases, MCF does not impact the atmospheric OH budget much. Emissions of MCF have always been relatively low and so the amount of OH lost to oxidation of MCF is negligible: consequently, MCF does not show up in Figure 1.1. Yet of the three gases we discuss, MCF is most central in our approach to achieve the primary objective of this thesis: improving constraints on OH.

MCF is emitted exclusively anthropogenically, mainly as a solvent, for example in paint or degreasers of metals. MCF contributes to the depletion of stratospheric ozone and its production has therefore been phased out in compliance with the Montreal protocol and subsequent amendments, resulting in a sharp drop in MCF emissions between 1990 and '95 (McCulloch & Midgley, 2001; Montzka et al., 2000). After circa 2010, emissions of MCF have approached zero (Rigby et al., 2013). Consequently, the amount of MCF in the atmosphere has been steadily declining by around 20% per year (see bottom left panel in Figure 1.4), mainly because MCF is being oxidized by OH. However, this rate of decline varies a bit: it might be 19% one year, and 20% in the next. The guiding principle of our approach to constraining OH is that these small variations in the rate of decline of MCF can be used as a measure for variations in OH (see for example Equation 1.5).

In principle, the same approach would be applicable to CH_4 , CO, or any other gas removed by OH: since OH plays such a large role in their budgets, variations in OH will translate to variations in the atmospheric abundances of these species. The difference is that for all these gases, there are substantial uncertainties in their emissions. We already discussed that an important application for stronger constraints on OH is better quantification of CH_4 emissions. To use CH_4 as a proxy for OH however, requires better constraints on CH_4 emissions, i.e. it becomes a chicken-and-egg story. The same holds for CO and other




Figure 1.3: A schematic overview of the global atmospheric circulation in the troposphere. Figure adapted from Lutgens et al. (2001).

gases. MCF is currently the only gas that possesses these two properties: emissions play a small role in its atmospheric budget and it is removed primarily by OH. This makes MCF uniquely qualified for our purpose. We expand on the use of MCF to constrain OH in Section 1.5, after some other relevant aspects of this challenge have been introduced.

1.2.3 Atmospheric transport

In the previous section, we have introduced the budgets of atmospheric gases as if the atmosphere is a well-mixed box. However, in reality, the atmosphere is a giant, dynamical reservoir of air. Atmospheric mixing is important when tracer-gas budgets are concerned (e.g. Pandey et al., 2019b). A gas might be emitted into the atmosphere in one place, but, subsequently, the gas will be transported and dispersed by atmospheric motions. There are certain large-scale (> 1000 km) patterns in atmospheric dynamics that determine how, once emitted, gases are distributed globally. A schematic overview of the global-scale circulation is shown in Figure 1.3.

The primary driver of all atmospheric motion is uneven heating of the Earth's surface by solar energy. For example, the tropics are heated more than the poles and the resulting temperature gradient drives winds from the poles towards the tropics. In the tropics, air masses from the Northern and the Southern hemisphere collide and rise upwards, creating the moist Inter-Tropical Convergence Zone (ITCZ; marked in Figure 1.3). The ITCZ can be considered the meteorological border between air mass in the Northern and the Southern Hemisphere and it is one of the main features of atmospheric circulation. The movement in the latitudinal direction of solar-powered winds is deflected by Earth's rotation: towards the right in the Northern Hemisphere and towards the left in the

Southern Hemisphere. The balance between pole-to-tropics winds and their deflection by Earth’s rotation results in the large-scale transport cells shown in Figure 1.3. Relatedly, we find strong easterly trade winds at the equator and dominant westerlies at mid-latitudes, for example over the U.S. and Europe.

The atmosphere is not only separated in the horizontal, but also in the vertical. The lowest 9–17 km of the atmosphere is called the troposphere and the next-highest layer is the stratosphere. The troposphere contains 80% of the atmospheric air mass and most atmospheric pollution. The stratosphere contains most of the other 20% of air mass and the main stratospheric feature is the ozone layer, that protects life on Earth against harmful solar ultraviolet (UV) radiation. Interestingly, whereas solar radiation heats the troposphere at the surface, the stratosphere is heated at the top, because of the absorption of UV radiation by ozone. Heating at the surface makes tropospheric air unstable (warm air rises), which drives not only the dominant transport cells shown in Figure 1.3, but is also at the root of all weather. Conversely, the stratosphere is stably stratified (hence the name) and, as a consequence, air inside the stratosphere mixes only slowly (e.g. Krol et al., 2018). The troposphere and stratosphere are separated by the tropopause: by definition, the tropopause corresponds to the altitude where temperature first starts to increase with altitude. The tropopause is highest in the tropics, where vertical motion in the troposphere is most intense, and lowest at the poles (see also Figure 1.3).

The exact delineation between the different transport cells varies with time. Most importantly, the seasonal solar cycle drives a seasonal cycle in transport. For example, the ITCZ follows the Sun each year, moving northwards in anticipation of the Northern Hemispheric summer and southwards afterwards. In addition to the seasonal cycle, there are also interannual variations in large-scale transport. One summer might be dry and hot, the next wet and cold. It turns out that much of the interannual variability in large-scale transport patterns can be explained with a few atmospheric modes of variability. The most important mode is the El Niño Southern Oscillation (ENSO) (McPhaden et al., 2006). ENSO oscillates between El Niño and La Niña periods, which generally last months to a year. El Niño events affect weather patterns on large scales and are associated with, for example, floodings in South-America and droughts in Australia (Ward et al., 2014). La Niña events have an approximately opposite effect. In the context of this thesis these events are important, because the ENSO cycle influences the atmospheric abundance and distribution of gases relevant to OH chemistry. For example, El Niño years are associated with more wildfires (Van Der Werf et al., 2004) and more lightning NO_x production (Murray et al., 2013).

We summarize the most important mixing timescales of different transport cells in the atmosphere as follows:

1. Due to tropical tradewinds and mid-latitudinal westerlies, air is mixed in the longitudinal direction faster (\sim one to three months) than in the latitudinal direction

(\sim six months).

2. Within each hemisphere, air below 10-15 kilometres (the troposphere) is mixed much faster (\sim six months) than air above this altitude (the stratosphere; \sim one to five years).
3. Mixing of air between the two hemispheres is slow (\sim one to two years), compared to mixing within hemispheres (\sim one to six months). The difference is mainly caused by the blocking vertical motion in the ITCZ that separates the two hemispheres.

These atmospheric transport cells and their associated timescales determine the atmospheric distribution of trace gases. In the distribution of OH we can see the location of the ITCZ (Figure 1.2), as well as the tropopause. However, as OH has a lifetime of only seconds, the presence of solar input and water vapour is more important than transport (Section 1.2.1). In contrast, gases such as CH_4 and MCF, with lifetimes of more than a few years, can be considered well-mixed throughout at least the troposphere. There are still observable spatial gradients in their atmospheric distribution, for example between hemispheres (Figure 1.4), but the gradients are small relative to the atmospheric abundance. Contrastingly, CO has an atmospheric lifetime of a few months, so that even longitudinal gradients of CO can exceed 100% during, for example, wildfire events. Knowledge of the presence (or absence) of predictable large-scale patterns in the distribution of gases is integral to approaches that aim to observe atmospheric composition, as we will see in the next section.

1.3 Observing atmospheric composition

Up to now we have discussed the atmosphere, its composition and its dynamics in an abstract sense. However, key to the scientific enterprise is observation. There are many ways in which atmospheric composition is observed, and different perspectives on the atmosphere offer different strengths and weaknesses. If you are interested in local emissions from a factory, you might want to measure close to the factory. However, we are interested in global, multi-decadal questions (see Sections 1.2.1 and 1.5), and so we seek observational approaches that fit these scales.

The most intuitive way in which the atmospheric composition can be measured is by filling a flask with outside air (the sample), and measuring the composition of the sample in a lab. It might seem like you would need a large amount of samples to understand what is going on in a reservoir of air that is as large as our atmosphere. However, most species we investigate are well-mixed throughout the troposphere, because their atmospheric lifetime is long with respect to transport mixing timescales ($>$ one year; see also Section 1.2.3). If this is the case, then a single one-liter sample of air filled in the right location can already be representative for an area 1000's of kilometres in extent.

1.3.1 Surface networks

The surface network of the National Oceanic and Atmospheric Administration (NOAA) (Dlugokencky et al., 1994; Montzka et al., 2000) operates on the principles outlined above: by monitoring atmospheric composition at a few remote locations, much information on trends and variability in the global atmospheric composition can be inferred. The top panel in Figure 1.4 shows the distribution of the subset of NOAA sites that monitor MCF. Mostly, the surface sites are spread along the latitudinal direction, because latitudinal gradients in the atmospheric distribution of gases are larger than longitudinal gradients. Additionally, a few sites are placed more closely together, but at different altitudes, such that limited information on vertical gradients can be retrieved. For example, on Hawaii, atmospheric composition is monitored both on Mauna Loa (3397m altitude) and at Cape Kumakahi (15m altitude).

To illustrate the representativeness of these sites, Figure 1.4 also shows timeseries of MCF and CH₄ mole fractions observed at two NOAA sites located in the high-latitude Northern and Southern Hemisphere, respectively. In the timeseries of MCF we see that in the early 1990's, atmospheric abundance was relatively high and decreased exponentially afterwards. Initially, the North to South gradient of MCF was also large, which was driven by emissions that were located mainly in the Northern Hemispheric midlatitudes. The zoom-in of 2010–13 shows that the interhemispheric gradient of MCF has almost, but not quite disappeared, now that emissions have all but ceased.

The picture for CH₄ is very different. Over most of the 1985–2020 period, CH₄ mole fractions have increased. However, compared to the rapid decline of MCF (20%/year), CH₄ abundance has been rather stable: the largest yearly increase in the NOAA record is 14 ppb in 1991, or less than 1%. Notably, even though the two sites are spaced roughly 6500 kilometres apart, the annual CH₄ growth rates at both sites are very similar: this is testament to the representativeness of the surface sites. The smaller long-term variations in the CH₄ record also make that the seasonal cycle of CH₄, which is driven primarily by seasonal variations in OH, is much more dominant and visible. Especially at the South Pole, farthest removed from sources, the seasonal cycle is very regular. During the polar winter, there is no light to form OH (Section 1.2.1), so that atmospheric lifetimes of CH₄ and of MCF become very long, resulting in elevated mole fractions. The reverse happens in summer. The cycle at Alert is similar but opposite in timing, because it is located in the other hemisphere. The opposite seasonality at both sites can be seen especially well in the 2010–2013 zoom-in window shown for MCF.

This brief description of the MCF and CH₄ atmospheric records reveals that many of the dominant features in the distribution of a long-lived gas are already well-constrained from only two surface sites. Moreover, the trace gas distribution provides some information on emission history, atmospheric oxidation and transport. Additional surface sites further help to constrain and separate the different components in tracer budgets. For exam-

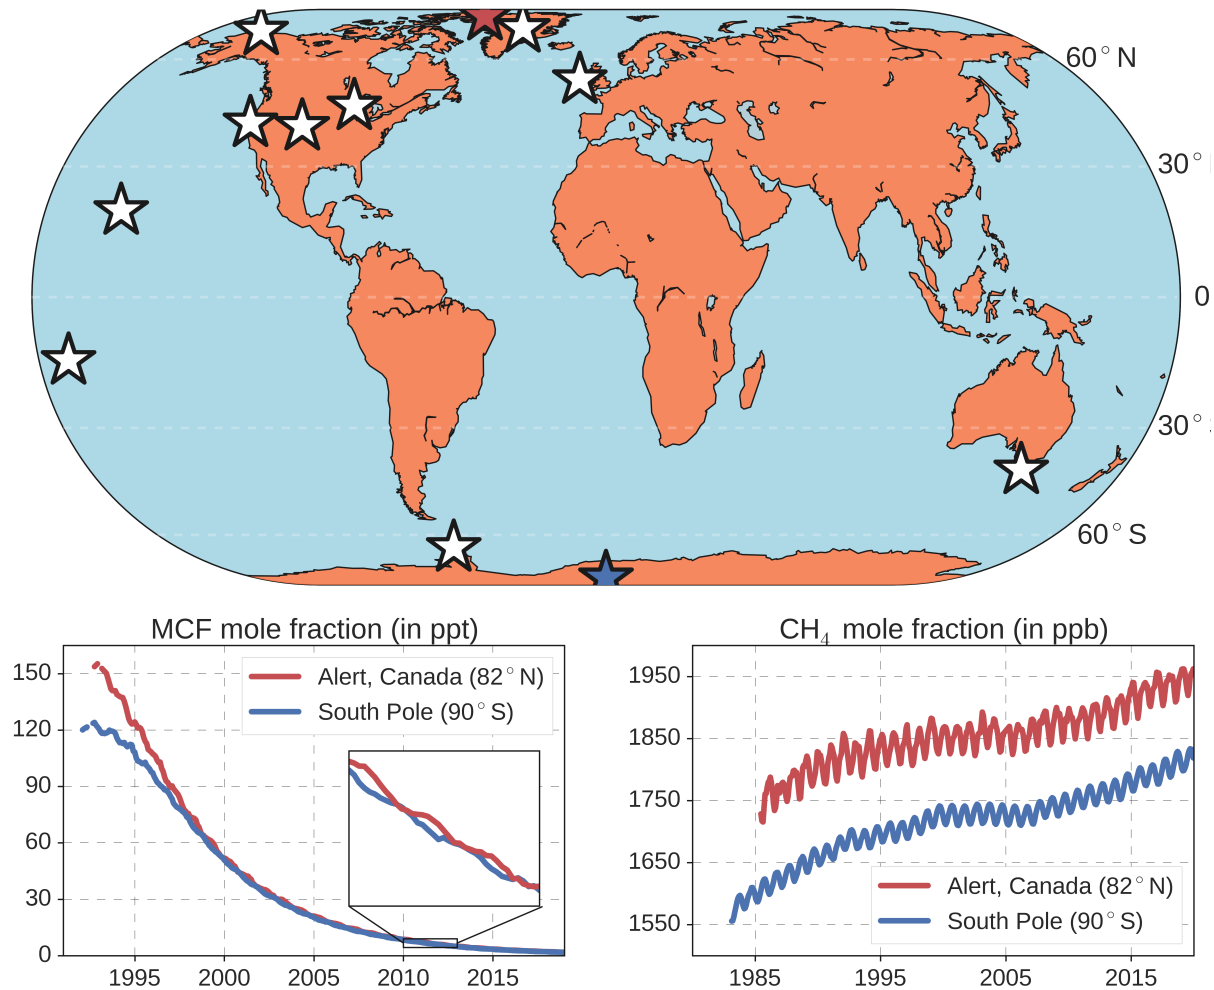


Figure 1.4: A map of the remote NOAA surface sites where MCF mole fractions are monitored (top) and timeseries of MCF (bottom left) and CH₄ (bottom right) mole fractions. Timeseries of mole fractions are shown for two sites: Alert, Canada (red) and South Pole (blue). These two sites are also marked with their respective colors on the map in the top panel. In the timeseries of MCF, a zoom-in of the mole fractions between 2010 and 2013 is shown as a visual aid. Note that the number of remote NOAA surface sites that monitor CH₄ mole fractions (>35) is higher than for MCF.

ple, observations from low-latitude sites show that, uniquely, MCF mole fractions after ~ 1998 are lowest in the tropics (Montzka et al., 2000): a feature that is revealing for the distribution of OH (Figure 1.2).

There are two surface networks that provide long-term monitoring of our species of interest: the NOAA network and the Advanced Global Atmospheric Gases Experiment (AGAGE) network Prinn et al. (1983, 2018). While both networks are set up using the principles outlined above, their exact approaches differ. In the NOAA network, flask samples are collected at different sites, transported to a central lab, and measured for atmospheric composition at this central lab. In the AGAGE network, each site has a measurement instrument that monitors atmospheric composition at that site. The NOAA network includes more surface sites, while the AGAGE network measures at a much higher temporal frequency. Therefore, it can be considered that the NOAA network better represents the average tropospheric distribution of a gas, while the AGAGE network is able to capture individual pollution plumes and meteorological events. Both sources of information are valuable and cross-comparison of observations from the two networks helps validate and better understand information contained in the records of each network (e.g. Rigby et al., 2013). However, in this thesis we will mostly use the NOAA surface network, since its strengths best fit our purpose.

Monitoring the atmosphere from a surface network has distinct strengths. Most importantly, data from surface networks have been available for a relatively long time: depending on the species, concentrations have been monitored since the late 1950's. Maintaining long-term measurement consistency in these data, mainly in the calibration scales, is a key objective and accomplishment of the surface networks. The long-term availability and consistency of surface measurements form the backbone of the work presented in this thesis. However, other observational approaches can complement surface observations, when we require additional information.

1.3.2 Aircraft observations

A monitoring method that is complementary to surface networks is the aircraft profile. Generally, aircraft profiles are measured using similar tools as those used for surface networks: either flasks of air are collected and later analyzed in a lab, or an in-situ instrument measures in-flight. Unlike observations from surface networks, however, aircraft profiles inform mainly on the vertical distribution of trace gases. Vertical transport of air is mostly driven by moist convection, which is difficult to capture correctly in models, compared to horizontal transport. Therefore, aircraft profiles add information to the spatial dimension that is most uncertain and that is largely missed by surface sites.

Another advantage of aircraft profiles is found in an area we investigate in Chapter 4: the Amazon rainforest. The Amazon is vast and diverse: large swaths remain pristine, while near the edges human influence is present nearly everywhere. Surface observations




Figure 1.5: The spectrum of outgoing longwave radiation (OLR) as emitted by Earth as a blackbody (smooth black line) and the spectrum of longwave radiation that actually escapes to space (blue/grey line). The difference between the two spectra is caused by absorption and emission of radiation by gases in the Earth’s atmosphere. The approximate absorption wavelengths of a few dominant greenhouse are shown. Additionally, the 4.7 μm absorption wavelength of CO is indicated in red: this absorption wavelength is exploited in the MOPITT-TIR and IASI products to estimate CO abundance. Adapted from a figure by Lukas Kluft, CC BY-SA 4.0, <https://commons.wikimedia.org/w/index.php?curid=79434392>.

off the coast of the Amazon can give some information on what goes on over land, but with very little detail. Surface observations inside the Amazon often only inform on very local conditions, because many sources and chemical processes occur near the Earth’s surface here. Aircraft profiles over the Amazon are able to bridge this gap (Gatti et al., 2014). Moving up into the atmosphere, air from the surface is mixed and so the air is representative for a larger area than air that is sampled near the surface. Moreover, the vertical distribution of gases can inform on the origin of certain plumes. This aspect of aircraft profiles is highly valuable and will be used in Chapter 4.

1.3.3 Satellite observations

Surface observations and aircraft observations follow essentially the same approach: to monitor the atmospheric composition at discrete points in the atmosphere. A vastly different method to monitor atmospheric composition is to observe it remotely from space, using satellites.

Any body emits blackbody radiation, with a spectrum that is a direct function of its temperature and therefore knowable. The Sun emits blackbody radiation that corresponds to

~ 5800 K, also known as shortwave radiation, which includes visible light and the harmful UV radiation. The Earth absorbs part of the solar radiation, and, because the Earth is also a body, the Earth emits its own radiation. Since the Earth is much cooler than the Sun (~ 290 K), the Earth emits longwave radiation that includes infrared (IR) radiation (solid black line in Figure 1.5). Both shortwave and longwave radiation travel through the atmosphere and, crucially, gases in the atmosphere absorb part of the radiation. The absorption in the atmosphere of outgoing longwave radiation is an important component of the greenhouse effect. Different gases absorb and re-emit radiation at different wavelengths (indicated in Figure 1.5), so that the spectrum of radiation that escapes to space is no longer smooth as that of a blackbody (blue/grey line in Figure 1.5). Satellite instruments monitor the spectrum of escaped radiation to infer the composition of the air that the radiation has traveled through.

In general, the advantage of satellites over aircraft profiles and surface networks is that satellites retrieve a much larger number of observations, also over very remote regions that might be difficult to access from the Earth's surface. The downside is that satellite observations are more difficult to interpret. The link between absorption lines in the observed radiation spectrum and the abundance of gas in an air column is not straightforward and not only do gases absorb radiation at certain wavelengths, they also emit their own black-body radiation that shifts the spectrum of radiation that escapes to space. Consequently, the accuracy of one satellite observation is usually much lower than that of one surface observation. In fact, satellite data need to be calibrated with observations from surface networks and aircraft campaigns. However, when tens of thousands of satellite observations are available per day and we do have surface observations to provide us with the ground truth, then satellite observations can provide us with a level of coverage that other observational approaches cannot match.

There is a wide variety of satellite instruments that observe atmospheric composition, each with their own strengths and weaknesses. The first distinction is the wavelength band that an instrument monitors, since different gases absorb light at different wavelengths (Figure 1.5). The choice of spectral band can also impact the vertical sensitivity of a satellite instrument. All satellite instruments principally measure the total amount of trace gas in a column of air, but measurements in one waveband might contain more information on the composition of air in the upper troposphere, while sensitivity in another waveband might be more vertically uniform. Aside from the waveband, other choices are available. For example, some satellites provide very detailed information on small scales, while others are less detailed, but cover a larger area. Some satellites always monitor the atmosphere above the same region of the Earth's surface, other satellites can be aimed at a specific area and yet others orbit the Earth in regular cycles.

In this thesis, we use data provided by two satellite instruments that monitor atmospheric CO abundance: the Measurements Of Pollution In The Troposphere instrument

(MOPITT; Deeter et al. (2019)) and the Infrared Atmospheric Sounding Interferometer (IASI, MetopA; (Clerbaux et al., 2009)). Both instruments use the CO absorption line around $4.7\text{ }\mu\text{m}$ (marked in red in Figure 1.5) to estimate CO columns. MOPITT also provides a product based on the $2.3\text{ }\mu\text{m}$ absorption line, but we do not use in that product in this thesis. IASI and MOPITT both provide global coverage on a daily basis, but their orbits are not aligned. The two instruments primarily differ in the measurement techniques used to retrieve spectral data and in the retrieval method that converts spectral data to CO columns. By deploying different methods to monitor CO abundance in the same spectral band, the two instruments provide comparable estimates that can be cross-validated to better understand uncertainties in retrieved CO columns (George et al., 2015; Nechita-Banda et al., 2018).

We use satellite-observed CO columns in Chapter 4 to estimate CO emissions from wildfires in the Amazon rainforest. Other observational tools are less well suited for the task: CO is difficult to constrain from a sparse surface network, because it is not well-mixed due to its short lifetime. Moreover, the Amazon is largely inaccessible and wildfires are highly variable in space and time, which further complicates any ground-based assessment.

1.4 Modelling the atmosphere and interpretation of observations

In atmospheric science, we often cannot measure directly what we want to know. In these cases, models can fill the gap between what we can measure and what we want to know.

In the simplest case, we measure the methane concentration at a few surface sites, but we want to track how much methane is in the entire troposphere. One way is to calculate the average concentration over all sites and multiply this with the total amount of air mass in the troposphere: a one-box model. A more advanced model would take into account that surface sites are not spaced evenly: for example, there are more sites in the Northern than in the Southern Hemisphere (Figure 1.4). Additionally, the meteorological boundary between the two hemispheres moves between seasons (Section 1.2.3), so that a tropical surface site will sometimes represent Northern, and other times Southern Hemispheric air. Very quickly a seemingly simple problem becomes complicated.

The above example illustrates two things: reality is complex and interpretation of observations almost always requires a model. The complexity of the model need not match the complexity of reality. A one-box model can already provide a good first estimate of the amount of methane in the troposphere. The delineation between two hemispheres is a simple suggestion that can improve the estimate significantly. However, the seasonal movement of the ITCZ is much more complex to incorporate in models and perhaps not of vital importance. The modeler's choice is to select and include those processes that are

most important and to consider the impact of the processes that are not (fully) accounted for.

In this section, we distinguish between two approaches. In what we consider to be simple models, the intention is to include as few processes as possible, while still capturing the important characteristics of a problem. In more complex 3D transport models, an attempt is made to reproduce atmospheric transport to the highest possible accuracy, which is only limited by incomplete understanding and limited computational resources. The former approach is more efficient and transparent, the latter is more complete.

1.4.1 Simple models

First, we consider the the global one-box model. In a one-box model, the entire atmosphere is represented by one box. The evolution of a gas X in a one-box model is given by:

$$\frac{d[X]}{dt} = \text{Emissions} - \text{Loss}, \quad (1.5)$$

In words, Equation 1.5 indicates that the concentration $[X]$ increases through time when emissions exceed loss and vice versa. Whether the emissions are anthropogenic or natural, in the middle of the atmosphere or on the ground, they all enter the atmosphere in the same way in a one-box model. The one-box model also includes very little of the atmospheric dynamics we have discussed in Section 1.2.3.

To be able to include some more information, we can increase the resolution to two boxes. An intuitive boundary between the two boxes would be the ITCZ, so that we distinguish between the Northern and Southern Hemisphere. Equation 1.5 then becomes:

$$\frac{d[X_{NH}]}{dt} = \text{Emissions}_{NH} - \text{Loss}_{NH} - k_{IH}(X_{NH} - X_{SH}), \quad (1.6a)$$

$$\frac{d[X_{SH}]}{dt} = \text{Emissions}_{SH} - \text{Loss}_{SH} + k_{IH}(X_{NH} - X_{SH}). \quad (1.6b)$$

Most obviously, we now need two equations, one for each box. We also need to include transport between the two boxes. Here, we assume that transport is proportional to and in the direction of the gradient between the two hemispheres (i.e. $(X_{NH} - X_{SH})$). The interhemispheric exchange coefficient k_{IH} indicates how fast the transport between hemispheres is. We have already seen that a typical timescale for interhemispheric transport is one year.

The distinction between a one- and a two-box model might seem trivial, and why not include even more boxes? For many gases, there are some key characteristics that make

two-box models very suitable. Most importantly, in a two-box model, the interhemispheric gradient is included. Many gases of interest, such as methane, are emitted predominantly in the Northern Hemisphere, because that is where most anthropogenic emissions take place. In contrast, most loss processes, e.g. oxidation by OH, are relatively symmetric between the hemispheres. This means that the concentration of such gases is higher in the Northern than in the Southern Hemisphere (see Figure 1.4). The interhemispheric gradient can, in this way, be used as a measure for the magnitude of emissions.

Why is this additional constraint important? Because in our first two studies, we are interested primarily in constraining OH, a loss process, but we also have to account for uncertainties in emissions. To come back to the example at the end of Section 1.2.2, if MCF declines with 19% in one year and with 20% in the next, then this could indicate a variation in OH, but it could also be due to a variation in emissions. Here it helps that the interhemispheric gradient of MCF will be affected differently by emissions than by OH. In Section 1.3.1 and Figure 1.4, we have already seen that even in later years, there is still a small interhemispheric gradient in MCF mole fractions, which could be indicative of small, ongoing emissions. Therefore, the better we are able to distinguish between MCF emission variations and OH variations, the better we are able to constrain OH.

While attractive in its simplicity, there are shortcomings to simplifying the atmosphere into a two-box model. The simplest example is the interhemispheric exchange rate. While the corresponding timescale is one year, the exchange rate does vary interannually, for example driven by the ENSO. There are methods for deriving interannual variations in the interhemispheric exchange coefficient, for example by using observations of the nearly-inert trace gas SF₆ (Geller et al., 1997). However, such an approach is rather uncertain and, as we will see in Chapter 3, sometimes the exchange rate can even differ between gases. A more fundamental limitation of simple models is that there is information that cannot be included. For example, we know that OH concentrations are highest in the tropics, which drives a tropical minimum in the latitudinal MCF distribution, but it is not possible to use this information in a two-box model. The potential problems of two-box models, as well as some possible solutions, are the subject of Chapter 2.

Tropospheric two-box models are important for this thesis, but there are many other approaches to approximating the atmosphere. A slight modification to a tropospheric two-box model would be to include a third box for the stratosphere: the next major transport cell in large-scale atmospheric circulation (as was done in (Rigby et al., 2017)). In other studies, the problem has been approached in a 2D twelve-box model (Prinn et al., 2001, 2005). In such a model, there is more opportunity to include our understanding of atmospheric dynamics and of the OH distribution, but it is still relatively cheap computationally. In this thesis, when we do not use a two-box model, we use a 3D transport model, which is on the other end of the spectrum of model complexity. In the next section, we discuss this type of model in more detail.

1.4.2 3D chemistry-transport models

A 3D transport model is in essence not different from a two- or twelve-box model: it just includes many more boxes. This makes it computationally expensive, but physically most realistic. Similar to the two-box model each of the boxes of a 3D transport model can have emissions, loss and tracer transport to and from adjacent boxes. Different from the two-box model, the transport between boxes in a 3D transport model is driven by some form of meteorological data, for example winds observed from satellites. In a good 3D transport model, large-scale transport phenomena are captured, e.g. seasonal movement of the ITCZ, a dynamical tropopause and El Niño events.

How well smaller-scale transport is simulated primarily depends on the grid resolution, i.e. how large the model boxes are. In this thesis, we use one particular 3D transport model: TM5 (Krol et al., 2005). A typical horizontal grid resolution of TM5 is one hundred to a few hundred kilometres. This resolution is fine enough to simulate off-coast transport of emissions towards an observational site in the middle of the Atlantic. It is too coarse, however, to tell us if air originates from Amsterdam or Rotterdam. As mentioned in Section 1.3, observatories are sometimes placed on top of mountains. This is also something that is difficult to simulate in these models, because mountains are not well-represented by 100 kilometre grid cells. As with simpler models, the key is to exploit the strengths of 3D transport models, while staying aware of the weaknesses.

An important strength of the 3D transport model is that it can be used to simulate individual observations, especially when sampled in remote places. For example, in a two-box model, surface observations are aggregated to hemispheric averages. In reality, both hemispheres contain multiple surface sites and the gradients between these sites can contain important information. Additionally, 3D models allow us to simulate other types of observations, such as aircraft profiles and satellite observations. This we rely on heavily in Chapter 4.

1.4.3 Inverse modelling

The modelling approaches that have been discussed up to now, both box models and 3D transport models, are so-called forward models. Input to these models are emissions, atmospheric loss, transport, etc. Output can be, for example, concentration fields of some tracer, or surface observations sampled from the model. In practice we already know the atmospheric concentrations, because that is exactly what we are able to measure, and what we are interested in primarily are emissions and atmospheric loss. An inverse model framework allows us to find those emissions and OH variations that are most consistent with atmospheric observations.

This objective can be quantified as a cost J that needs to be minimized:

$$J(\vec{x}) = \frac{1}{2}(\vec{x} - \vec{x}_{prior})^T \mathbf{B}^{-1}(\vec{x} - \vec{x}_{prior}) + \frac{1}{2}(\mathbf{H}\vec{x} - \vec{y})^T \mathbf{R}^{-1}(\mathbf{H}\vec{x} - \vec{y}). \quad (1.7)$$

The cost function consists of two terms. The first term (marked in blue) contains prior knowledge. We might want to estimate MCF emissions and OH, but we also know something already about these quantities. We know that OH is located mainly in the tropics and we know that the OH distribution varies seasonally. For emissions, there are often inventories available, for example because countries have to report their emissions. While neither of those estimates is perfect, they do provide an anchor point to start from. There is a cost for deviating too much from this first estimate, which is quantified in the prior cost. The state vector \vec{x} contains all quantities that we want to estimate, such as OH or some tracer's emissions, and \vec{x}_{prior} denotes our first estimate for these quantities. The error covariance matrix \mathbf{B} contains the uncertainties in the prior estimate.

The second term (marked in red) quantifies how well observations are reproduced by the model. The \vec{y} -vector contains all observations, such as surface observations from the NOAA network or satellite observations. \mathbf{H} denotes the forward model, which could be e.g. a 3D transport model, or a two-box model. The forward model is applied to the state vector to calculate simulated observations ($\mathbf{H}\vec{x}$). The observational cost is then calculated as the difference between model-simulated and real-world observations, weighted by the observational error matrix \mathbf{R} . \mathbf{R} combines both the measurement and the model uncertainty of observations. The model uncertainty is driven by, for example, potential errors in the modeled transport and by sub-grid variations that are per definition not included in the model.

There are many statistical algorithms designed to find the state \vec{x} that minimizes J . In this thesis, we use the 4DVAR method (Figure 1.6 and Fisher (1995); Meirink et al. (2008)). For the 4DVAR method, the cost function gradient ∇J also needs to be calculated:

$$\nabla J(\vec{x}) = \mathbf{B}^{-1}(\vec{x} - \vec{x}_{prior}) + \mathbf{H}^T \mathbf{R}^{-1}(\mathbf{H}\vec{x} - \vec{y}). \quad (1.8)$$

The only new quantity in the cost function gradient, compared to the cost function, is the adjoint model \mathbf{H}^T . Where the forward model \mathbf{H} is applied to the state to calculate model-simulated observations, the adjoint model is applied to the difference between simulated and real-world observations ($\mathbf{H}\vec{x} - \vec{y}$) to calculate how the state needs to be adjusted to better fit observations: for example, by how much emissions or OH need to be adjusted to better reproduce MCF observations. The cost function gradient quantifies how we need to adjust \vec{x} to make the cost J smaller.

Figure 1.6 schematically shows the workflow in a 4DVAR inversion. We start from a first estimate x_{prior} that is used as input for the forward model \mathbf{H} to simulate observations and, by extent, to calculate the cost J . The model-measurement mismatch is then used as

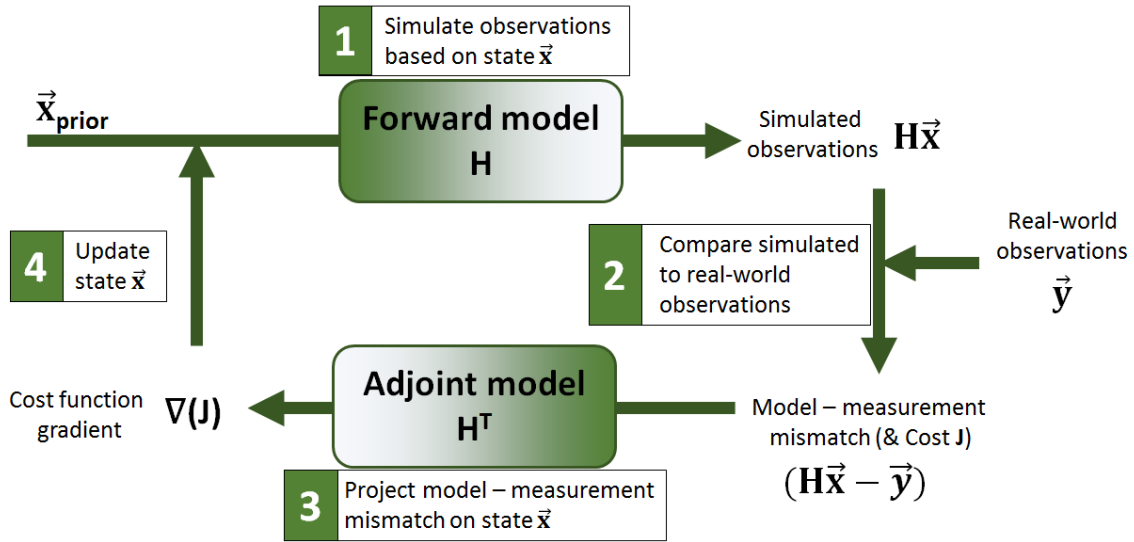


Figure 1.6: A schematic of the workflow in a 4DVAR inversion. Input to the algorithm is the initial estimate \vec{x}_{prior} , which is iteratively adjusted and optimized to result in the best fit with observations \vec{y} . Step 1 and 3 involve the most computational time and are the most work to set up, as in these steps the forward and adjoint version of the two-box or TM5 transport model are used. Step 2 is a simple subtraction, and for step 4 statistical optimization algorithms are available that, based on the cost function and its gradient, propose an intelligent adjustment to the state (e.g. the `m1qn3` solver (Gilbert & Lemaréchal, 1989) that we use for our TM5 inversions).

input for the adjoint model H^T to calculate how the state needs to be adjusted to reduce the cost, which is quantified as the cost function gradient ∇J . Based on the gradient, an adjustment in the state x is proposed, which can then be used to calculate new simulated observations and a new cost function. If the cost has decreased after the adjustment to the state, then the adjustment was successful, so that we have a new best estimate. This process typically continues until the cost function gradient becomes small, so that further adjustments to the state do not much reduce the cost anymore, which indicates that a balance has been reached between the blue and red term in Equation 2.2.

1.5 Thesis outline

Previous sections outline all concepts and tools that are used in this thesis. In this section, we build on this theoretical framework and summarize our approach in a set of research questions that are leading in the scientific studies presented in this thesis. A visualization of the thesis outline is provide in Figure 1.7.

The primary objective of Chapters 2 and 3 is to improve and better understand constraints on OH and, by extent, the atmospheric oxidative capacity, by using observations of MCF.

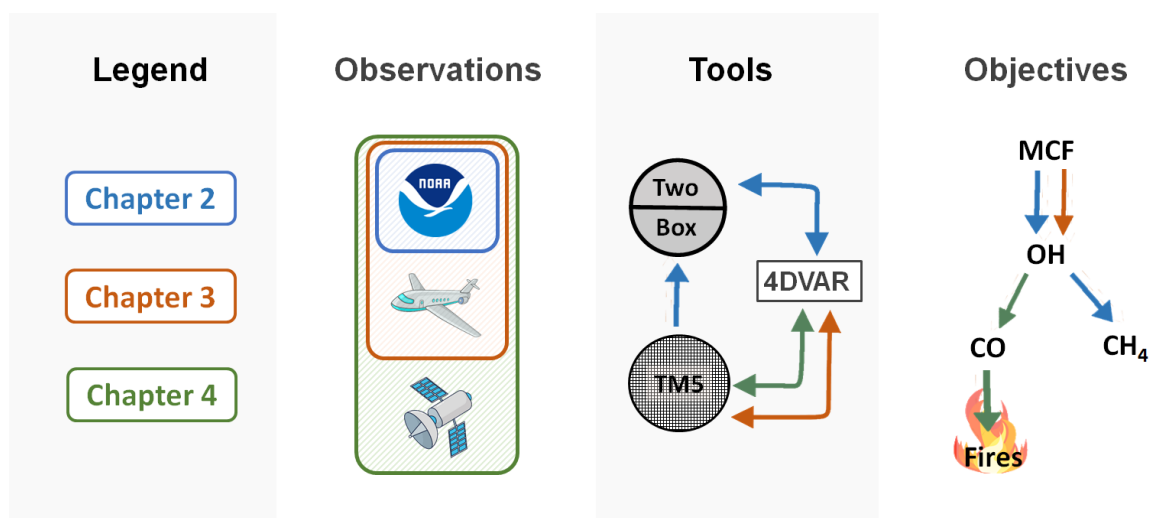


Figure 1.7: An overview of the type of observations and tools used in each chapter, along with a visualization of the corresponding research questions. Each chapter is marked with a different color (legend on the left). Types of observations, from top to bottom are: the NOAA surface network; aircraft observations; satellite observations. The 4DVAR algorithm is described in detail in Figure 1.6.

MCF was recognized early on as a good proxy for OH, because its industrial production was considered well-constrained (Loveloock, 1977; Prinn et al., 1987). Early studies that exploited these MCF production inventories found large interannual variations ($> 10\%$) and significant trends in the derived OH (Krol et al., 1998; Prinn et al., 2001). The large amplitude of these OH variations was found to be inconsistent with our understanding of OH chemistry (Lelieveld et al., 2004) and with variations in the CH₄ growth rate (Dentener et al., 2003). Later studies have shown that, if the MCF emissions and especially their timing are considered to be somewhat more uncertain, much smaller OH variations are needed to fit the MCF observational record (Krol & Lelieveld, 2003; Bousquet et al., 2005).

When emissions of MCF dropped, the constraints that MCF observations placed on OH much improved. In the absence of significant MCF emissions, it was possible to derive a typical interannual variability for global OH of $2.3 \pm 1.5\%$ and OH variations derived from MCF were found to be consistent with variations in the growth rate of other trace gases, such as CH₄ (Montzka et al., 2011). Likewise, in tropospheric two-box model inversions that included both MCF and CH₄, evidence was found for significant co-variations between variability in OH and the CH₄ growth rate (Rigby et al., 2017; Turner et al., 2017). The same two studies showed that even interannual variability in OH of a few percent can significantly modify the most probable source attribution of the CH₄ budget. However, significant uncertainties remained regarding the precise variability of OH. For example, it was not possible to exclude a scenario where OH had not varied at all.

The primary advantage of the box model approach is computational efficiency. This allows the user to set up a more comprehensive inverse estimation framework (e.g. a Markov chain Monte Carlo in (Rigby et al., 2017)) and it makes sensitivity tests accessible. Moreover, the two-box model is seemingly transparent due to its simplicity. However, the atmosphere does not consist of two or three boxes, and simplifying it as such carries some inherent assumptions. In Chapter 2, we more closely investigate the two-box model approach by tuning our own two-box model inversion of MCF and CH_4 with 3D model output of the transport model TM5. Specifically, we address the following questions:

1. How do the simplifying assumptions inherent to a two-box model affect the best estimate of interannual OH variations derived from MCF?
2. Can we reduce uncertainties in OH variability derived in a two-box model inversion by accounting for transport effects with a 3D transport model?

In Chapter 2, we complement our two-box model inversion of MCF as best as we can with 3D transport effects. However, in the end, we still approach the problem in a two-box model. This means that we cannot include some important information. For example, in a two-box model, MCF observations take the form of two hemispheric averages, even though ten to fifteen NOAA surface sites monitor MCF. Gradients between the sites can potentially inform on the distribution and variability of OH. Additionally, aircraft observations of MCF are available, but these cannot be compared to two-box model results. Therefore, in Chapter 3, we continue with a 3D model inversion of MCF in the framework of TM5-4DVAR (Meirink et al., 2008), which does allow us to use all information available. Leading in Chapter 3 is the question:

3. What constraints on OH can we derive in a 3D transport model inversion of MCF?

In Chapter 4 we move away from MCF and apply the TM5-4DVAR inverse system to a different problem, related to CO emissions from fires in the Amazon basin. The Amazon basin is home to the world's largest rainforest, which has a huge and unique biodiversity. Each year fires burn through the Amazon forest and the neighbouring savanna, particularly during the local dry season (Sanford et al., 1985), and these fires strongly influence the composition of the atmosphere (Crutzen & Andreae, 1990). A small part of the fires is a natural part of the ecosystem, sparked mainly by lightning (Latham & Williams, 2001). However, anthropogenic activities related to deforestation and agriculture have greatly enhanced fire frequency and intensity (Bowman et al., 2009). Extensive deforestation can additionally reduce the resilience of these biomes to fires and droughts (Malhi et al., 2009). Superimposed on changes induced by local anthropogenic activity, changes in the global climate might drive an increase in drought frequency and intensity in the Amazon (Cochrane & Barber, 2009), which can additionally enhance fire activity (Aragão et al., 2018). With the help of local deforestation abatement policies, deforestation in the Amazon has been reduced over the past few decades. However, it remains of crucial importance to monitor whether this trend continues and how a changing climate affects fire

activity.

There are many methods for monitoring fires. Most of these methods exploit space-borne products, because the area covered by the Amazon is vast and remote. Different methods can show different signals and diverse monitoring methods help build confidence in our understanding of fire dynamics. That is why we exploit one key property of fires to better constrain them: they emit large amounts of CO. Because CO has a typical lifetime of months, the large CO emissions associated with fires drive a large and local spike in the atmospheric abundance of CO. By leveraging satellite observations of CO over the Amazon in the TM5-4DVAR inverse system, we try to resolve the question:

4. What can a 3D transport model inversion of CO teach us about fire activity and deforestation in the Amazon?

In addition to satellite observations, we use a unique set of aircraft profiles sampled at five locations in the Amazon. As discussed in Section 1.3, where satellite observations provide us with many details, aircraft observations are better calibrated and provide a ground truth.

It can be expected that the strong local signal of CO emissions and the high level of detail contained in satellite observations offer very different challenges than resolving the diffuse signal of OH in sparse surface network observations of MCF. However, Figure 1.7 gives an overview of the approach adopted in each chapter, which shows that there is also significant overlap between the three studies presented in this thesis. For example, while we constrain CO fire emissions mainly with satellite observation, we also use surface CO observations to constrain the global, atmospheric background. Additionally, a key strength of the 3D transport model is that it is adaptable to provide an atmospheric perspective on any large-scale problem.

Chapter 2

Constraints and biases in a tropospheric two-box model of OH

This chapter is based on:

Naus, S., Montzka, S. A., Pandey, S., Basu, S., Dlugokencky, E. J., & Krol, M. (2019). Constraints and biases in a tropospheric two-box model of OH. *Atmospheric Chemistry and Physics*, 19, p. 407–424.

Abstract

The hydroxyl radical (OH) is the main atmospheric oxidant and the primary sink of the greenhouse gas CH₄. In an attempt to constrain atmospheric levels of OH, two recent studies combined a tropospheric two-box model with hemispheric-mean observations of methyl chloroform (MCF) and CH₄. These studies reached different conclusions concerning the most likely explanation of the renewed CH₄ growth rate, which reflects the uncertain and underdetermined nature of the problem. Here, we investigated how the use of a tropospheric two-box model can affect the derived constraints on OH due to simplifying assumptions inherent to a two-box model. To this end, we derived species- and time-dependent quantities from a full 3D transport model to drive two-box model simulations. Furthermore, we quantified differences between the 3D simulated tropospheric burden and the burden seen by the surface measurement network of the National Oceanic and Atmospheric Administration (NOAA). Compared to commonly used parameters in two-box models, we found significant deviations in the magnitude and time-dependence of the interhemispheric exchange rate, exposure to OH, and stratospheric loss rate. For MCF these deviations can be large due to changes in the balance of its sources and sinks over time. We also found that changes in the yearly-averaged tropospheric burden of CH₄ and MCF can be obtained within 0.96 ppb/yr and 0.14%/yr by the NOAA surface network, but that substantial systematic biases exist in the interhemispheric mixing ratio gradients that are input to two-box model inversions.

To investigate the impact of the identified biases on constraints on OH, we accounted for these biases in a two-box model inversion of MCF and CH₄. We found that the sensitivity of interannual OH anomalies to the biases is modest (1-2%), relative to the uncertainties on derived OH (3-4%). However, in an inversion where we implemented all four bias corrections simultaneously, we found a shift to a positive trend in OH concentrations over the 1994-2015 period, compared to the standard inversion. Moreover, the absolute magnitude of derived global mean OH and by extent that of global CH₄ emissions were affected much more strongly by the bias corrections than their anomalies (~10%). Through our analysis, we identified and quantified limitations in the two-box model approach, as well as an opportunity for full 3D simulations to address these limitations. However, we also found that this derivation is an extensive and species-dependent exercise, and that the biases were not always entirely resolvable. In future attempts to improve constraints on the atmospheric oxidative capacity through the use of simple models, a crucial first step is to consider and account for biases similar to those we have identified for the two-box model.

2.1 Introduction

For the interpretation of atmospheric observations in the context of, for example, atmospheric pollution, or in that of global warming, atmospheric models are often used. Atmospheric models vary in complexity from simple one box models to state-of-the-art 3D transport models. Different types of models are suitable for addressing different types of problems to different degrees of scrutiny. Therefore, there is no model category that fits all problems. Simple box models are easy to set up, computationally cheap, and transparent. For these and other reasons, their use in atmospheric studies is ubiquitous and has provided useful insights (e.g. Quay et al. (1999); Walker et al. (2000); Montzka et al. (2011); Schaefer et al. (2016); Schwietzke et al. (2016)). However, simple box models also put limitations on the derived results, as they are by definition less comprehensive than complex models. For example, box models do not explicitly contain much information on a species' spatial distribution, which can be important if interacting quantities (e.g. loss processes) are distributed non-homogeneously in space. Where exactly these limitations lie and what the gain is from increasing model complexity can be difficult to diagnose, and depends on the application.

A problem that has often been approached in box models is that of constraining the global atmospheric oxidizing capacity, which is largely determined by the tropospheric hydroxyl radical (OH) concentration (Montzka et al., 2000, 2011). OH is dubbed the detergent of the atmosphere for its dominant role in the removal of a wide variety of pollutants, including urban pollutants (CO, NO_x), greenhouse gases (CH₄, HFCs), and HCFCs, which are greenhouse gases, and also contribute to stratospheric ozone depletion. The budgets of many of these pollutants have been strongly perturbed since pre-industrial times, and it is important to understand what consequences this has had in the past, and could have in the future, for the atmosphere's oxidizing capacity.

Due to its high reactivity, OH has a lifetime of seconds, which inhibits extrapolation of direct measurements. Moreover, OH abundance is the net result of many different reactions and reaction cycles, and thus modelling it process-based in full-chemistry models is complex and dependent on uncertain emission inventories of the many gases involved. Therefore, the most robust observational constraints on OH on the larger scales are thought to be derived indirectly from its effect on tracers: gases that are predominantly removed by OH. Depending on how well the tracer emissions are known, the time evolution of the global mixing ratio of such a tracer can serve to constrain OH. The most well-established tracer for this purpose is methyl chloroform (MCF) (e.g. Montzka et al. (2000); Bousquet et al. (2005)). In part, this is because it was identified early on as a tracer with a well-defined production inventory that allowed emission estimates with small errors, relative to other gases (Lovelock (1977); Prinn et al. (1987)). Moreover, production of MCF was phased out in compliance with the Montreal Protocol, and the resulting rapid drop in emissions made loss against OH the dominant term in the MCF budget (Montzka et al.,

2011).

Research and debate surrounding OH (Krol & Lelieveld (2003); Krol et al. (2003); Reimann et al. (2005); Prinn et al. (2005); Rigby et al. (2013); McNorton et al. (2016)) has lead to considerable improvements in its constraints: for example, a likely upper bound on global interannual variability of OH of a few percent (Montzka et al., 2011). Two recent studies derived OH variations in a tropospheric two-box model through an inversion of atmospheric MCF and CH₄ observations (Rigby et al. (2017); Turner et al. (2017)). In such an inversion, a range of parameters is optimized (most prominently emissions of MCF and CH₄, and OH), so that the modelled mixing ratios best match atmospheric observations of the tracers involved.

Both studies found that constraints on OH in this set-up were weak enough that a wide range of OH concentration variations over time and, by extent, CH₄ emission scenarios were possible as an explanation for the post-2007 increase in its measured global mole fraction. This is an important conclusion, because the CH₄ growth rate, combined with the CH₄ lifetime (in turn dominated by MCF-derived OH), is generally assumed to provide the strongest top-down constraints on global CH₄ emissions and variations therein. We note that in Rigby et al. (2017) the two tropospheric boxes were supplemented by a single stratospheric box, making it technically a three-box model. However, due to our focus on the troposphere, we hereafter treat this type of model too as a two-box model, and where relevant we discuss the implication of the addition of a stratospheric box.

There are two important reasons to approach the problem of constraining OH in a model of exactly two tropospheric boxes. Firstly, through the focus on annual timescales and hemispheric spatial scales, the result is only sensitive to interannual variability in large-scale transport of the modelled tracers. Moreover, by focusing on interannual variability as opposed to absolute OH or emission levels, remaining systematic offsets are not thought to significantly affect the outcome.

Secondly, a crucial part of the optimization consists of disentangling the influence of OH and that of emission variations on observed MCF mixing ratios. Ideally, MCF emission variations would be prior knowledge. However, though MCF production is well documented, the emission timing is much more uncertain (McCulloch & Midgley, 2001). MCF was mainly used as a solvent in, for example, paint and degreasers of metals. In these applications, MCF is released only when used, rather than when produced, which results in uncertainty in the emission timing. Moreover, due to the continuing decline of the atmospheric MCF mixing ratios, small, ongoing MCF emissions could eventually become important. Observation-inferred emissions exceeding bottom-up emission inventories have been identified both from the U.S. (Millet & Goldstein, 2004) and from Europe (Krol et al., 2003), as well as from other processes, such as MCF re-release from the ocean (Wennberg et al., 2004). Therefore, in the absence of other constraints, emission uncertainties would limit the use of MCF for deriving interannual variability of OH. However, in a two-box

set-up, an additional constraint is provided by the IH mole fraction gradient of MCF. Emission inventories show that MCF emissions are predominantly located in the Northern Hemisphere (NH), whereas OH has a NH to SH ratio that is uncertain, but the ratio has a likely range of 0.80 to 1.10 (Montzka et al. (2000); Patra et al. (2014)). This means that emission variations have a strong effect on the IH mole fraction gradient of MCF, whereas the effect of large-scale OH variations is much weaker. Thus, the IH gradient is an important piece of information that can help to disentangle the influence of emissions from the influence of OH on MCF growth rate variations. This use of the IH gradient for constraining global emissions of anthropogenically emitted gases has also been recognized in previous research (Liang et al. (2017); Montzka et al. (2018)).

Despite the appealing degree of simplicity offered by the two-box model, its results still hinge on many simplifying assumptions, both explicit (e.g. interhemispheric transport) and implicit (e.g. intrahemispheric transport). In this context, the uncertain outcome of the two recent two-box model studies put forward an important question: how do the simplifying assumptions inherent to the two-box set-up affect the conclusions drawn from it? Or, conversely, would these conclusions change when moving the analysis to a 3D transport model? A recent study (Liang et al., 2017) partly explored these questions. The study investigated how to incorporate information from 3D transport models in a two-box model, to increase the robustness of two-box model derived constraints on OH. They found that there are key parameters in the two-box model that can be tuned to better represent the 3D simulation results, and thus ideally better represent atmospheric transport in general. For example, they found that IH transport rates can be species-dependent.

Here, we provide a different approach to the issue. In the first part of our study, we parametrized results from the 3D global transport and chemistry model TM5 into a two-box model. Through this parametrization, we explored difficulties in the translation from the 'reality' of a 3D transport model to a two-box model, and the assumptions made in the process. We focused on four aspects of the parametrization.

Firstly, we investigated the tracer-dependent nature of IH transport as reported by Liang et al. (2017). Secondly, we analysed the IH OH ratio. Previous research has shown that because of tracer-specific source-sink distributions, different tracers can be exposed to different global mean OH concentrations (Lawrence et al., 2001). We extended this observation to a species-dependent IH OH ratio. Thirdly, we looked at the stratospheric loss for MCF specifically. This net loss to the stratosphere might be slowing after its emissions dropped (Krol & Lelieveld (2003); Bousquet et al. (2005)). Fourthly, we used the 3D simulation to investigate differences between the burden seen by the surface measurement network of the National Oceanic and Atmospheric Administration Global Monitoring Division (NOAA-GMD) and the true tropospheric and hemispheric burden in our 3D model: a bias that was also discussed in Liang et al. (2017).

In the second part of this study, we assessed the impact of these four potential biases on derived OH variations in a two-box inversion set-up that is very similar to Rigby et al. (2017) and Turner et al. (2017). The objective was to provide a quantitative estimate of the impact of biases in a two-box inversion, and to explore if and how these can be accounted for. Though this study is focused on the problem of OH, it also serves as a case study of potential pitfalls in two-box models in general, when applied to interpreting global-scale atmospheric observations.

2.2 Methods

2.2.1 Two-box inversion

In this section, we discuss the set-up of our two-box model inversion. The model incorporated two tracers (MCF and CH₄) and consisted of two boxes (the troposphere in the NH and in the SH), which were delineated by a fixed equator. The stratosphere was implicitly included in the model through a first-order loss process that was taken to be equal for both hemispheres. The governing equations for a tracer mixing ratio X are given in Equation 2.1.

$$\frac{dX_{\text{NH}}}{dt} = E_{\text{NH}} - (k_{\text{OH}}[\text{OH}]_{\text{NH}} + l_{\text{strat}} + l_{\text{other}})X_{\text{NH}} - k_{\text{IH}}(X_{\text{NH}} - X_{\text{SH}}), \quad (2.1\text{a})$$

$$\frac{dX_{\text{SH}}}{dt} = E_{\text{SH}} - (k_{\text{OH}}[\text{OH}]_{\text{SH}} + l_{\text{strat}} + l_{\text{other}})X_{\text{SH}} + k_{\text{IH}}(X_{\text{NH}} - X_{\text{SH}}). \quad (2.1\text{b})$$

Thus, within each hemisphere, there were emissions (E), loss to OH ($k_{\text{OH}}[\text{OH}]X$), loss to the stratosphere ($l_{\text{strat}}X$), loss to other processes ($l_{\text{other}}X$; e.g. ocean deposition), and transport between the hemispheres ($k_{\text{IH}}(X_{\text{NH}} - X_{\text{SH}})$). The model ran at an annual timestep. The fundamentals of this model set-up are also found in Rigby et al. (2017) and Turner et al. (2017), though the exact treatment of the different budget terms can differ. For example, Turner et al. (2017) combined all tropospheric loss, including loss to the stratosphere, in one term, whereas Rigby et al. (2017) included a stratospheric box, so that stratospheric loss becomes a transport rather than a first-order loss term. Where relevant, we point out further differences with these previous studies.

Since the objective was to leverage observed mixing ratios to infer information on tropospheric OH, we also set up an inverse estimation framework, complementary to the above forward model. The objective of the inversion was to optimize a state x , such that the forward model best reproduced the observations, without straying too far from a first best guess: the prior. Therefore, the state is the vector which contains all parameters that needed to be optimized. The optimization objective is analogous to minimizing the cost function J , as defined in Equation 2.2:

$$J(\mathbf{x}) = \frac{1}{2}(\mathbf{x} - \mathbf{x}_{prior})^T \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_{prior}) + \frac{1}{2}(\mathbf{H}\mathbf{x} - \mathbf{y})^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}), \quad (2.2)$$

with \mathbf{B} and \mathbf{R} the prior and observation error covariance matrix respectively, \mathbf{H} the forward model, and \mathbf{y} the observations. In addition, we compute the cost function gradient ∇J (Equation 2.3).

$$\nabla J(\mathbf{x}) = \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_{pri}) + \mathbf{H}^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}), \quad (2.3)$$

with \mathbf{H}^T the transpose of the forward model, also known as the adjoint model. Note that because the forward model \mathbf{H} was non-linear (e.g. OH chemistry), we used the adjoint of the tangent-linear forward model. Calculation of the cost function gradient facilitates quicker convergence of the optimization. For the minimization we used the Broyden-Fletcher-Goldfarb-Shanno algorithm. In essence, this statistical inversion set-up is the same as that used in the 4DVAR system of ECMWF (Fisher, 1995) and TM5-4DVAR (Meirink et al., 2008).

For the optimization of MCF emissions, we used an extended version of the emission model from McCulloch & Midgley (2001). This emission model was adopted to account for the varying and uncertain release rates of MCF when used in different applications (e.g. degreasing agent, paint). This uncertainty results in a gap between the uncertainty in production, or integrated emissions ($\sim 2\%$), and the uncertainty in annual emissions (up to 40%) (McCulloch & Midgley, 2001). Therefore, production was distributed between four different categories with different release rates: rapid, medium, slow and stockpile. In the prior distribution, the bulk of production ($> 95\%$) was placed in the rapid category. To account for uncertainty in the production inventory, we also adopted an additional emission term superimposed on the production-derived emissions. The emissions in year i were then given by Equations 2.4 and 2.5. For each year i , we optimized four parameters for MCF emissions: three parameters that shifted emissions between the rapid production category and each of the other three categories (f_{Medium}^i , f_{Slow}^i and f_{Stock}^i in Equation 2.6), and the additional emissions term ($E_{Additional}^i$), which had an uncertainty constant through time. This emission model is similar to that used in (Rigby et al., 2017), though ours leaves more freedom with respect to the timing of emissions.

$$E^i = E_{Rap}^i + E_{Med}^i + E_{Slow}^i + E_{Stock}^i + E_{Additional}^i, \quad (2.4)$$

for the emissions in year i , where:

$$\begin{aligned}
 E_{Rap}^i &= 0.75P_{Rap}^i + 0.25P_{Rap}^{i-1}, \\
 E_{Med}^i &= 0.25P_{Med}^i + 0.75P_{Med}^{i-1}, \\
 E_{Slow}^i &= 0.25P_{Slow}^{i-1} + 0.75P_{Slow}^{i-2}, \\
 E_{Stock}^i &= \sum_{j=1}^{11} P_{Stock}^{i-j},
 \end{aligned} \tag{2.5}$$

and, in the optimization:

$$\begin{aligned}
 P_{Rap}^i &= (1 - f_{Med}^i - f_{Slow}^i - f_{Stock}^i)P_{Rap,prior}^i, \\
 P_{Med}^i &= P_{Med,prior}^i + f_{Med}^i P_{Rap,prior}^i, \\
 P_{Slow}^i &= P_{Slow,prior}^i + f_{Slow}^i P_{Rap,prior}^i, \\
 P_{Stock}^i &= P_{Stock,prior}^i + f_{Stock}^i P_{Rap,prior}^i.
 \end{aligned} \tag{2.6}$$

An important choice in the inversion set-up is which parameters to prescribe and which to optimize. Rigby et al. (2017) optimized all parameters, so as to explore the full uncertainty of the optimization within the inversion framework. Turner et al. (2017) only optimized hemispheric MCF and CH₄ emissions and hemispheric OH, while the remaining uncertainties were partly explored in sensitivity tests. We choose to optimize four end-products for each year: global OH, global MCF emissions, global CH₄ emissions, and the CH₄ emission fraction in the NH. Thus we had a closed system, as we also fitted to four observations: the global mean mixing ratio and the IH gradient of both MCF and CH₄. In addition to the 4DVAR inversion, we generated a Monte Carlo ensemble, where in each realization the prior and the observations were perturbed, relative to their respective uncertainties. Then, the new prior was optimized using the new observations. The Monte Carlo simulation quantified the sensitivity of the optimization to the prior choice and to the realization of the observations. The Monte Carlo set-up also allowed us to explore the sensitivity of the inversion to parameters that were not optimized, such as the fraction of MCF emissions in the NH. This approach had the added advantage that parameters that were perturbed in the Monte Carlo simulation, but not optimized in the 4DVAR system, did not need to have a Gaussian error distribution. Gaussian probability distributions are normally a prerequisite in a 4DVAR inversion. The specifics of our inversion set-up are given in Table 2.1.

2.2.2 TM5 set-up and two-box parametrizations

3D model set-up

For the 3D model simulations we used the atmospheric transport model TM5 (Krol et al., 2005). The model was operated at a $6^\circ \times 4^\circ$ horizontal resolution, at 25 vertical hybrid sigma pressure levels. The simulation period was 1988-2015, where we treated 1988

Table 2.1: The relevant settings we used in the inversion of our two-box model. The upper section contains the parameters optimized in the inversion, which were also perturbed in the Monte-Carlo ensemble. These parameters have Gaussian uncertainties, and their mean and 1- σ uncertainty are given. The middle section contains parameters that were perturbed in the Monte Carlo, but not optimized. The middle parameters have uniform uncertainties, of which the lower and upper bound are given. The bottom section contains parameters that were neither optimized nor perturbed. For these parameters, the left column gives the standard setting, whereas the alternative column indicates whether we also ran an inversion using a TM5-derived timeseries (see Section 2.2.2).

Parameters optimized in inversion and perturbed in the Monte Carlo ensemble (Gaussian)		
Parameter	Prior estimate	Uncertainty
Global MCF emissions	Based on McCulloch & Midgley (2001)	
- f_{Medium}	0%	5%
- f_{Slow}	0%	5%
- f_{Stock}	0%	5%
- Unreported emissions	0 Gg/yr	10 Gg/yr
Global CH ₄ emissions	550 Tg/yr	15%
Global OH	9×10^5 molec/cm ³	10%
Fraction NH CH ₄ emissions	75%	10%

Parameters not optimized in inversion, but perturbed in the Monte Carlo ensemble (Uniform)		
Parameter	Lower bound	Upper bound
Fraction NH MCF emissions	90%	100%

Parameters not optimized in inversion and not perturbed in the Monte Carlo ensemble		
Parameter	Standard	Alternative
Interhemispheric OH ratio	0.98	TM5-derived*
MCF lifetime w.r.t. oceanic loss	83 yr	-
MCF lifetime w.r.t. stratospheric loss	45 yr	TM5-derived*
CH ₄ lifetime w.r.t. stratospheric loss	150 yr	TM5-derived*
Interhemispheric transport	1 yr ⁻¹	TM5-derived*

*See Section 2.2.2

and 1989 as spin-up years. TM5 transport was driven by meteorological fields from the ECMWF ERA-Interim reanalysis (Dee et al., 2011). Convection of tracer mass was based on the entrainment and detrainment rates from the ERA-Interim dataset. This is an update from the previous convective parametrization used by, for example, Patra et al. (2011). The new convective scheme results in faster interhemispheric exchange of tracer mass, more in line with observations (Tsuruta et al., 2016).

We ran TM5 with three tracers: CH_4 , MCF and SF_6 . For CH_4 , we annually repeated the 2009-2010 a priori emission fields used by Pandey et al. (2016), and we also used the same fields for stratospheric loss to Cl and $\text{O}(^1\text{D})$. For MCF, we used emissions from the TransCom- CH_4 project (Patra et al., 2011). Since these emissions were available only up to 2006, we assumed a globally uniform exponential decay of 20%/year afterwards, similar to Montzka et al. (2011). MCF-specific loss fields (ocean deposition and stratospheric photolysis) were also taken from the TransCom- CH_4 project. Details of the MCF loss and emission fields can be found in the TransCom- CH_4 protocol. The OH loss fields we used were a combination of the 3D fields from Spivakovsky et al. (2000) in the troposphere and stratospheric OH as derived using the 2D MPIC chemistry model (Brühl & Crutzen, 1993). The OH fields were scaled by a factor 0.92, as described by Huijnen et al. (2010). For SF_6 , we used emission fields from the TransCom Age of Air project (Krol et al., 2018), with no loss process implemented.

Since the above set-up is simplistic in some aspects (e.g. annually repeating CH_4 emissions), we also ran a 'nudged' simulation. In the nudged simulation, we scaled the mixing ratios of a tracer up or down in latitudinal bands, depending on the mismatch of the model with NOAA observations (analogous to Bândă et al. (2015)), with a relaxation time of 30 days. This method ensured that the model followed the long-term trend in observations, without requiring a full inversion. The nudged simulation provided a test of the sensitivity of our results to the source-sink distributions we used in the 3D simulation.

Parametrizing 3D model output to two-box model input

Here we outline how we used the TM5 simulations to derive two-box model parametrizations for stratospheric loss (l_{strat}) and for interhemispheric exchange (k_{IH}). Firstly, the 3D fields were divided in 3 boxes: the troposphere in the NH and in the SH, and the stratosphere. The border between the hemispheres was taken as the equator, fixed in time. Where relevant we discuss the sensitivity of our results to this demarcation. We defined a dynamical tropopause as the lowest altitude where the vertical temperature (T) gradient is smaller than 2 K/km, clipped at a geopotential height of 9 and 18 km. Our analysis was found to be insensitive to the exact definition of the tropopause. Next, we computed an annual budget for each box. For the two tropospheric boxes, this was done as in Equation 2.1. This was supplemented by Equation 2.7 for the stratospheric box.

$$\frac{dX^{\text{Strat}}}{dt} = -L_{\text{local}}^{\text{Strat}} + l_{\text{strat}}(X_{\text{SH}} + X_{\text{NH}}), \quad (2.7)$$

Emissions, local loss and mixing ratios per box could be derived from the 3D model input and output, and thus l_{strat} and k_{IH} could be inferred from these equations. Note that we did not strictly need the stratospheric budget equation to resolve two parameters, but we used it to resolve numerical inaccuracies. Resolving the budget of each species in this manner provided the necessary input of the tropospheric two-box model defined in Section 2.2.1, such that on the hemispheric and annual scale, identical results were obtained with the 3D and the two-box models.

An additional parameter that we derived from the TM5 simulations was the IH OH ratio to which each tracer was exposed. We quantified this parameter as the ratio between hemispheric lifetimes with respect to OH (τ_{OH}), as in Equation 2.8. Note that this might differ from the physical IH OH ratio, because of correlations between the tracer distribution, the OH field and the temperature distribution.

$$r_{\text{LOH}} = \frac{\tau_{X,\text{OH}}^{\text{SH}}}{\tau_{X,\text{OH}}^{\text{NH}}}. \quad (2.8)$$

Model-sampled observations

The standard in tracking global trends in atmospheric trace gases are surface measurement networks: for CH_4 and MCF most notably the NOAA-GMD (Dlugokencky et al. (2009); Montzka et al. (2011)) and the AGAGE (Advanced Global Atmospheric Gases Experiment) (Prinn et al., 2018) networks. By selecting measurement sites far removed from sources, the theory is that a small number of sites already puts strong constraints on the global growth rate (Dlugokencky et al., 1994). In general, quantification of the robustness of the derived growth rates based solely on observations can be difficult, since there are likely systematic biases inherent to sampling a small number of surface sites. When assimilated into a 3D transport model, these biases will largely be resolved (if transport is correctly simulated). However, when the data is aggregated to two hemispheric averages, as in a two-box model, quantification of the potential biases is crucial.

We explored the resulting bias in our model framework. By subsampling the TM5 output at the locations of NOAA stations, at NOAA measurement instances, we generated a set of model-sampled observations. These model-sampled observations were intended to be as representative as possible for the real-world observations of the NOAA network. To aggregate the station data to hemispheric averages, we used methods similar to those deployed by NOAA (for MCF: Montzka et al. (2011), with further details on our adaption in Supplement 2.S1; for CH_4 : Dlugokencky et al. (1994)). Hemispheric averages for CH_4 were derived from 27 sites, and MCF averages from 12 sites. By comparison of

the resulting products with the calculated tropospheric burden, as derived from the full tropospheric mixing ratios, we could assess how well the burden derived from the NOAA network represents the model-simulated tropospheric burden. The two end-products we investigated for each tracer were the rate of change of the global mean mixing ratio and that of the IH gradient. Note that by mixing ratio we mean the dry air mole fraction. These two parameters best reflect the information as it is used in a two-box model: the global mean mixing ratio is used to constrain the combined effect of OH and emissions, while the IH gradient is used to distinguish between the two. Note that in previous box-model studies of MCF, often only global growth rates were derived (Montzka et al., 2000, 2011).

2.2.3 Potential biases in the two-box model

By concentrating on the budget of MCF, we identified three parameters that need attention in the two-box model: IH transport, the IH OH ratio, and loss of MCF to the stratosphere. In addition, we investigated the potential bias in converting station data to hemispheric averages (see Section 2.2.2 and Supplement 2.S1). We quantified these biases and propagated them in two-box model inversions, as discussed in Section 2.3.2, to quantify their impact on derived quantities related to OH.

Interhemispheric transport

IH transport of tracer mass can vary because of variations in IH transport of air mass (e.g. influenced by ENSO, particularly at Earth’s surface (Prinn et al. (1992); Francey & Frederiksen (2016); Pandey et al. (2017))), or because of variations in the source-sink distribution, and thus of the tracer’s concentration distribution itself. Generally, interannual variability in IH transport is considered to be in the order of 10% (Patra et al., 2011). Two-box model studies typically assume time-invariant IH exchange (Turner et al., 2017) and/or similar exchange rates for different tracers (Rigby et al., 2017). Here we investigated whether such assumptions hold for a tracer which undergoes strong source-sink redistributions over time, such as MCF. The IH transport variations we derived for each tracer are discussed in Section 2.3.1.

Surface sampling bias

As discussed in Section 2.2.2, we explored the bias that results from representing hemispheric averages using sparse surface observations. Surface networks are a valuable resource, because they provide high-quality, long-term measurements of a growing variety of tracers. However, temporal, horizontal, and vertical coverage of surface networks is limited. In Section 2.3.1 we discuss how these limitations can result in biases in two-box model observations.

The interhemispheric OH ratio

The IH ratio of OH concentrations is an uncertain parameter. This is mostly because of a mismatch between results from full-chemistry models (1.13–1.42 (Naik et al., 2013)) and from MCF-derived constraints (0.80–1.10 Brenninkmeijer et al. (1992); Montzka et al. (2000); Patra et al. (2014)). The latter is generally the loss ratio considered in two-box models (1.0 in Turner et al. (2017) and 0.95–1.20 in Rigby et al. (2017)), and is similar to the ratio we used in the TM5 simulations (0.98 (Spivakovsky et al., 2000)). However, the bias we consider here is of a different nature: it is the difference between the physical IH OH ratio and the IH loss ratio a particular tracer is exposed to. It is known that different tracers can be exposed to different oxidative capacities (Lawrence et al., 2001). Therefore, different tracers might similarly be influenced by different IH ratios in OH. We explore this bias in Section 2.3.1.

MCF loss to the stratosphere

The second-most important loss process of MCF is stratospheric photolysis. In our TM5 set-up, this loss process resulted in an in-stratosphere lifetime (stratospheric burden/stratospheric loss) of 4 to 5 years. It is generally assumed that this in-stratosphere loss translates to a global lifetime of MCF with respect to the stratosphere (global burden/stratospheric loss) of 40 to 50 years (Naik et al. (2000); Ko et al. (2013a)), which corresponds to $\sim 10\%$ of global MCF loss. Rigby et al. (2017) assumed a time-invariant in-stratosphere lifetime, but due to the inclusion of a stratospheric box the global lifetime with respect to stratospheric loss could vary somewhat due to changes in the troposphere-stratosphere gradient. These variations were tuned to result in a global lifetime with respect to stratospheric loss of 40 (29–63) years. Turner et al. (2017) incorporated this loss process in the OH loss term. Due to the rapid drop in MCF emissions and the relatively slow nature of troposphere-stratosphere exchange, this lifetime could vary through time (Montzka et al. (2000); Krol & Lelieveld (2003); Bousquet et al. (2005)). We will investigate this possibility in Section 2.3.1.

2.2.4 Standard two-box inversion and bias correction

To assess the impact of the biases discussed in Section 2.2.3 on a two-box model inversion, we ran our inversion (see Section 2.2.1) using different settings. In the standard, default inversion we did not consider any of the four biases discussed above. Thus, we used constant IH exchange (1 year), constant stratospheric loss of MCF (45 years), and a constant IH OH ratio (0.98) (see Table 2.1). The first three potential bias corrections were then straightforwardly implemented by replacing these constant values with the timeseries we derived for each parameter from the full 3D simulations (details in Section 2.2.2). As




Figure 2.1: Hemispheric, annual mean timeseries of CH_4 (left) and MCF (right), as derived from the NOAA surface sampling network (for CH_4 , 27 sites were used; for MCF, 12 sites were used). Solid lines denote averages as derived directly from the NOAA surface sampling network (which are used in our standard inversion). Dashed lines denote the same timeseries, but adjusted by correction factors that were derived from our TM5 simulations. The correction factors reflect the differences between hemispheric averages based on model-sampled observations, and hemispheric averages derived from the full TM5 troposphere. Figure 2.3 shows the ratios between the standard and corrected timeseries.

mentioned in Section 2.2.1, the inversion did not include uncertainties in these three parameters. We did this because conventional uncertainties tend to be large, and therefore, including them would have attenuated the impact of the bias corrections, while the corrections were the main interest of this comparison. For the surface sampling bias, we first computed a correction between the hemispheric means as derived from the model-sampled observations and the calculated (TM5) hemispheric, tropospheric means (with demarcation at the equator). Then, we applied this correction to the real-world NOAA hemispheric means we used in the standard inversion. This gave a new set of observations, which we used in the inversion (discussed in Section 2.3.1). Both the standard and the corrected set of observations are shown in Figure 2.1. Through comparison of the results of the standard inversion and of an inversion with one or more biases implemented simultaneously, we can evaluate the individual and cumulative impact of the biases on derived OH and CH_4 emissions.

2.3 Results

2.3.1 Biases

Interhemispheric transport

The IH exchange coefficients, derived for the three different tracers as described in Section 2.2.2, are shown in Figure 2.2. Clearly, the exchange rates differ between tracers both in mean value, as well as in interannual variability. MCF is the clear outlier, but SF_6




Figure 2.2: The IH exchange rate for MCF, CH₄ and SF₆, as derived from a TM5 simulation (see Section 2.2.2).

and CH₄ also show different variations. The drivers of these differences are differences in intrahemispheric tracer distributions, and in the underlying source and sink distributions. The three tracers differ strongly in this respect: SF₆ and MCF are emitted almost exclusively in the NH mid-latitudes, whereas CH₄ has significant emissions in the tropics and in the SH. SF₆ has no sink implemented in our simulations, whereas MCF and CH₄ have a sink with a distinct tropical maximum in OH. This all affects how IH transport of air mass translates to IH transport of tracer mass.

Most notable is the minimum in the IH exchange rate for MCF in the 2000-2005 period. The timing of the 1989-2003 decline in k_{IH} coincides with the initial drop in MCF emissions. An important shift in the distribution of the MCF mixing ratio is that the global minimum shifts from the South Pole to the tropics. In the same period, there is a strong vertical redistribution which has also likely impacted IH exchange. It is not obvious that these changes should result in slower IH exchange, but in the end, in TM5, they do.

Another notable feature is the positive trend in the IH exchange rate for CH₄ ($+0.35 \pm 0.05$ %/yr; $p=0.00$) and for SF₆ ($+0.50 \pm 0.01$ %/yr; $p=0.00$). For CH₄, we used annually repeating sources, whereas for SF₆ we did include emission variations (see Section 2.2.2). This means that for CH₄, changes in the source-sink distribution did not contribute to the trend or to the variability. Indeed, in a simulation with annually repeating meteorology, we found near-zero variability in k_{IH} for CH₄ (see Supplement 2.S4). Therefore, there is something in the combination of the meteorological data, the treatment of this data in TM5 and the source-sink distribution of both CH₄ and SF₆ which resulted in a significantly positive trend in the IH exchange rate of both gases. This trend could either indicate an acceleration of IH transport of air mass, or a shift in the pattern of IH transport which favours IH exchange of CH₄ and SF₆. It is unclear from this analysis what the underlying mechanism is exactly, except that it is driven by temporal variations in transport and thus that there are parameters in the meteorological fields which also show a trend: otherwise this final product cannot exhibit a trend. However, it might be that the sensitivity of

Table 2.2: Mean observational errors as derived from TM5 simulations over the 1994-2015 period. The errors were quantified as the mean difference between annual means derived from model-sampled observations and annual means derived from the full tropospheric grid. CH₄ uncertainties are given both in ppb/yr and relative to the global mean mixing ratio. Uncertainties for MCF are only given relative to the global mean, because of its strong temporal decline.

	Global growth rate	IH gradient rate of change
CH ₄	0.96 ppb/yr / 0.05 %/yr	2.56 ppb/yr / 0.13 %/yr
MCF	- / 0.14 %/yr	- / 0.33 %/yr

TM5 transport to these parameters is biased.

To test the sensitivity of the derived IH exchange rates to the source-sink distribution, we compared k_{IH} derived from the standard simulation to the nudged simulation (the nudging procedure is explained in Supplement 2.S2). IH transport of CH₄ as derived from the nudged simulation showed higher interannual variations than in the standard simulation (more discussion in Supplement 2.S4), which can be expected, as the source-sink distribution becomes more variable. However, the general characteristics were conserved: most notably, the positive trend over the entire period persisted, for CH₄ and for SF₆. For MCF, we find that the general characteristics of derived k_{IH} are similarly insensitive to nudging, with the main change being a deeper 2000-2005 minimum in the nudged simulation. In the end, we deem the anomalies presented in Figure 2.2 quite robust with respect to the spatio-temporal source-sink distribution.

When the hemispheric interface is shifted from the equator to 8°N, which is more representative of the average position of the ITCZ, the IH exchange rate increases for all tracers, but the variability in IH exchange of CH₄ and SF₆ remains largely unaffected (see Supplement 2.S4). However, for MCF, the variability shifts completely. Rather than decreasing after the emission drop, the IH exchange rate now increases. This sensitivity reflects that for a tracer with a relatively small IH gradient which minimizes in the tropics, it becomes difficult to define an IH transport rate in a two-box model. By extension, care should be taken when interpreting the IH gradient of MCF in later years, since the influence of IH transport is difficult to isolate. Sensitivities in the derivation of the IH exchange rate are discussed in more detail in Supplement 2.S4.

Surface sampling bias

Figures 2.1 and 2.3 show the surface network bias in the global mean mixing ratios and in the IH gradient. In Figure 2.3, the bias is quantified as the ratio between values derived from the model-sampled observations (see Section 2.2.2) and values derived from the




Figure 2.3: The surface sampling bias in the global mixing ratio (left) and in the IH gradient (right) of MCF and of CH₄. The bias was quantified as the ratio between values derived from the NOAA surface sampling network and values derived from the full (TM5) troposphere. The biases were derived from 27 and 12 sites for CH₄ and for MCF, respectively. Figure 2.1 visualizes the impact of correcting for the sampling bias in real-world NOAA observations.

hemispheric (TM5) tropospheres. A comparison with global mean mixing ratios derived from real-world NOAA observations is given in Supplement 2.S2.

The bias in the IH gradient was particularly large, because averages based on NOAA surface stations systematically overestimated the tropospheric burden in the NH and underestimated the burden in the SH. Two important effects contributed to this bias. Firstly, in the NH, where most emissions were located, mixing ratios tended to decrease with altitude, while in the SH vertical gradients were much smaller, or even reversed. Secondly, latitudinal gradients of both MCF and CH₄ tended to be highest in the tropics, where few or no measurement sites were available. Again, due to high emissions in the NH, mixing ratios in the NH decreased towards the equator, while mixing ratios increased towards the equator in the SH. Both biases were of opposite sign in each hemisphere. Thus, in a global average, these biases largely cancelled, and only a small overestimate remained (left panel in Figure 2.3). For the IH gradient, however, these biases added up, which resulted in an overestimate of the IH gradient by surface stations of up to 20-40% (right panel in Figure 2.3). For MCF before 1995 and for CH₄ throughout the analysis period, the bias from the vertical gradient dominated. The shift in the bias for MCF was driven by a shift in the latitudinal gradient. The IH gradient of MCF got a minimum in the tropics and apparently this exacerbated the effect of the lack of tropical stations, combined with the simple, linear latitudinal interpolation we adopted for MCF (see Supplement 2.S1).

We note that the derived bias in the IH gradient is sensitive to the demarcation of the two tropospheric boxes. When we shifted the IH interface from the equator to 8°N, the bias was reduced to 15% for CH₄, and varied between 15 and 25% for MCF. The trend in the IH bias of MCF became smaller, but persisted.

Liang et al. (2017) performed a similar analysis for MCF. They reported a similar low to absent bias in the global mean and a more significant bias in the IH gradient of MCF ($\sim 10\%$). This is smaller than the bias we found, even if we demarcated the hemisphere at 8°N . However, an important difference is that in Liang et al. (2017) model-sampled observations were compared to the surface grid, instead of to the full troposphere. Thus, their bias estimate did not include vertical effects. When we used the surface grid as a reference, the IH bias for CH_4 was reduced to -10% : i.e. it reversed. For MCF the bias shift persisted, and the maximum bias was only slightly reduced to 15% , indicating a dominant influence from the latitudinal dimension. We emphasize that for a tropospheric two-box model the comparison with the full troposphere is most relevant.

This analysis also provided an estimate of uncertainties in the rate of change of the global mixing ratio and in that of the IH gradient: the relevant observational parameters in a two-box inversion. Table 2.2 gives the differences between the quantities derived from model-sampled observations and from the full troposphere, i.e. the "true" (TM5) error. We can compare this TM5-derived uncertainty to uncertainties derived only from observations, which we used in the two-box inversions. For CH_4 , we used uncertainties as reported by NOAA. These were obtained by generating an ensemble of surface network realizations, where in each realization different sites are excluded or double-counted randomly (bootstrapping). For each realization, aggregated quantities such as the global mean growth rate can be derived. The spread within the ensemble then provides a measure for the uncertainty. For MCF no such uncertainties are reported. Therefore, we developed our own method, which is described in Supplement 2.S1.

Following these methods, we found observation-derived uncertainties in the global mean growth rate of around 0.60 ppb/yr and $0.6\%/yr$ for CH_4 and for MCF respectively. NOAA does not report an uncertainty in the IH gradient of CH_4 , but error propagation from hemispheric means gave an uncertainty of 1.1 ppb/yr . For MCF, we found a time-dependent uncertainty in the rate of change of the IH gradient of $1.0 - 1.5\%$.

The CH_4 errors we derived from the TM5 simulation were slightly higher than the uncertainties reported by NOAA. Furthermore, since we used annually repeating CH_4 emissions, variations in CH_4 emissions can further increase the error. Indeed, the nudged run (see Supplement 2.S2) resulted in 20% higher uncertainties. However, it is important to note that the CH_4 uncertainties reported by NOAA are intended to reflect the match with the marine boundary layer (MBL), rather than with the full troposphere. Therefore, it is not surprising that the errors we find are somewhat higher.

For MCF, we adopted observation-derived uncertainties that were significantly lower than those used by Rigby et al. (2017) and Turner et al. (2017): both studies reported uncertainties of around 5% in hemispheric averages. Both studies used different methods, that were grounded on different observational information. In Rigby et al. (2017), temporal variability dominated the uncertainty estimate, while in Turner et al. (2017) spatial varia-




Figure 2.4: The ratio between tracer lifetime with respect to OH loss in the SH troposphere and NH troposphere (see Equation 2.8). Additionally, the IH ratio in OH concentrations is shown.

tions were used. Our method is more similar to Rigby et al. (2017), but with modifications that averaged out some of the temporal variability, under the assumption that variability at different measurement sites was largely uncorrelated (details in Supplement 2.S1). This shows that observation-derived uncertainties in MCF averages are uncertain quantities, in large part due to the relatively low number of available surface sites. Therefore, the uncertainty derived from TM5 is an especially useful addition for MCF.

Table 2.2 shows that TM5-derived uncertainties in MCF averages are significantly lower than all observation-derived estimates. This result indicates that even the use of a simple averaging algorithm and a small number of surface sites, relative to what is available for CH_4 , already results in well-constrained hemispheric and global growth rates for MCF. The TM5-derived estimate thus supports the use of our observation-derived uncertainty estimates, rather than the higher estimates used in previous studies.

Interhemispheric OH ratio

In the TM5 simulations from which the global loss rates were derived, the prescribed tropospheric OH fields were taken from Spivakovsky et al. (2000). In these fields, the IH OH ratio is 0.98, when the IH interface is considered to be the equator. One might expect a similar ratio between OH loss in the NH and in the SH, which we quantified through the IH ratio in tracer lifetime with respect to OH loss (Equation 2.8). We found that this is not the case (see Figure 2.4).

The loss ratio was up to 7% higher than the physical OH ratio. Moreover, the ratio was not the same for MCF and CH_4 , and the ratio that corresponded to MCF showed a trend. The IH asymmetry in temperature in our model was small, so that it didn't explain the different between the IH loss and the IH OH ratio. Instead, we found that the systematic positive offset was largely driven by an IH asymmetry in the spatio-temporal correlations




Figure 2.5: The tropospheric loss rate to the stratosphere, as derived from the TM5 simulations (see Section 2.2.2).

between OH and temperature. Mostly, this was because the OH maximum in the NH was located at lower altitude than in the SH in our 3D model. Since at low altitudes temperatures are higher, and higher temperatures correspond to higher reaction rates, this asymmetry resulted in relatively high NH loss rates. As such, the ratio bias was sensitive to the OH distribution used in the 3D model simulation.

The trend in the ratio for MCF was driven by the change in the spatial distribution of MCF after the emission drop in the mid-90s. Before the drop, the IH gradient of MCF was emission-driven and high (25%). This resulted in a negative correlation between OH/temperature and MCF in the NH, which drove the initially lower loss ratio. After the emission drop, the IH gradient became largely sink-dominated, and dropped to 3%. The ratio then became similar, though not identical, to that of CH₄, which also has a relatively low IH gradient (5%). The exact reasons for the IH asymmetry in the OH loss rate were complex: further details are discussed in Supplement 2.S3.

The derived IH OH ratio was sensitive to the demarcation of the two tropospheric boxes. When we shifted the position from the equator to 8°N, all IH OH ratios were reduced by 10 to 15%. However, the offset between the physical IH OH ratio and the actual loss ratio remained similar, as did the trend in the loss ratio for MCF.

Loss to the stratosphere

Figure 2.5 shows the stratospheric loss rate, as derived from Equations 2.1 and 2.7. Most notably, the stratospheric loss rate showed a significant negative trend for MCF, decreasing by 68% from 1991 to 1997. The MCF lifetime with respect to stratospheric loss in 1990 as calculated from TM5 was similar to the range reported in literature: 40 to 50 years (Naik et al. (2000); Ko et al. (2013a)). Afterwards however, the corresponding timescale for stratospheric loss quickly increases. As loss to the stratosphere is a secondary loss process, it is generally assumed that variability in MCF loss is driven predominantly by OH variations (Montzka et al. (2011); Turner et al. (2017); Rigby et al. (2017)). Here, we

found that this is not necessarily the case. The decline in loss to the stratosphere was not an artefact resulting from treating a transport process as a loss process: when taking the exchange proportional to the troposphere-stratosphere gradient, we still found a decrease in the exchange rate of 63%. Previous research has identified that the tropospheric lifetime with respect to stratospheric loss could be decreasing (Krol & Lelieveld (2003); Prinn et al. (2005); Bousquet et al. (2005)), but not to the degree that we found here, and not relative to the troposphere-stratosphere gradient. This is important, because it means that also a three-box model with an explicit stratospheric box, such as in Rigby et al. (2017), would not capture the decline.

The explanation we suggest for the increase in MCF lifetime with respect to stratospheric loss has to do with the nature of troposphere-stratosphere exchange, which consists of an upward and a downward flux. In practice, as MCF emissions decreased, the troposphere started to transport air to the stratosphere which was exposed to lower MCF emissions, while the stratosphere was still transporting older air back to the troposphere (in the downward branch of the Brewer-Dobson circulation (Butchart, 2014)) that was exposed to higher MCF emissions. Therefore, the delay between the two opposed fluxes resulted in a reduced net upward flux rate in an atmosphere with decreasing emissions compared to an atmosphere with increasing or constant emissions. Consistent with this hypothesis, we found that the stratospheric loss rate did not decrease in a TM5 simulation with MCF emissions fixed at 1988 levels, and that stratospheric loss did decrease, but recovered, when we fixed emissions at 2005 levels over the entire analysis period (results not shown). This also implies that the troposphere-stratosphere exchange will slowly recover when MCF emissions stop decreasing.

For CH_4 , we found a stratospheric lifetime of 160-170 years, similar to the range reported in Chipperfield & Liang (2013). For SF_6 , there was no loss process implemented in our model. However, storage of SF_6 in the stratosphere acted as an effective sink to the troposphere, with a lifetime of 100-160 years.

2.3.2 Two-box inversion results

In this section, we present a comparison between the results of the standard inversion and an inversion that incorporated the four bias corrections (referred to as "four-biases"). The inversion set-ups are described in Section 2.2.4. The OH and CH_4 emission anomalies of both inversions are presented in Figure 2.6, along with uncertainty envelopes of one standard deviation. The envelopes are wide, and with respect to these envelopes there were no significant differences between our two inversions. Interestingly, differences between the two inversions were the smallest in the 1998-2007 period, during which MCF is thought to provide the strongest constraint on OH (Montzka et al., 2011). Note that the final analysis period started from 1994 (rather than from 1990), because we only had sufficient NOAA coverage of MCF available from 1994 onwards.




Figure 2.6: The results of two inversions of the two-box model: tropospheric OH anomalies (left) and CH_4 emission anomalies (right). In the standard inversion, we kept IH transport, NH/SH OH ratio and stratospheric loss of MCF constant, and we used NOAA observations. In the second inversion, we implemented all four bias corrections instead (as described in Section 2.2.4). Both the mean anomalies and the 1-standard deviation envelopes are shown, where anomalies were taken relative to the time-averaged mean in each respective ensemble member. Plotted in grey are the anomalies as derived by Rigby et al. (2017) (from the NOAA dataset) and by Turner et al. (2017) (from a combined NOAA+AGAGE dataset), adjusted so that they, too, average to zero. The 1-standard deviation envelope from the Rigby et al. (2017) estimate is hatched in grey.

Shown in grey in Figure 2.6 are the anomalies derived by Rigby et al. (2017) (from the NOAA dataset) and by Turner et al. (2017). The four inversions showed qualitatively similar time-dependencies, and differences generally fell within one standard deviation, and always within two standard deviations. Differences with Turner et al. (2017) are largest, most notably after 2010, which can be expected since they use a combined AGAGE+NOAA dataset, whereas we only use NOAA data. In Rigby et al. (2017) it was shown that the use of a different dataset can result in different OH anomalies, though these differences were insignificant with respect to their uncertainty envelopes. Also visible is the uncertainty envelope of one standard deviation from Rigby et al. (2017), which is notably larger than our envelopes. This is likely due to a combination of the higher observational uncertainties and the higher number of optimized parameters adopted in Rigby et al. (2017). Further discussion of differences with these two studies is provided in Section 2.4.

It is illustrative to further investigate how the identified biases impact the results. For this purpose, Table 2.3 presents five metrics for each of the two inversions, as well as for inversions where we implemented the bias corrections one-by-one (taking standard settings for the other parameters).

The first metric is the mean absolute error (MAE) in the OH anomalies between each respective inversion and the standard inversion. The MAE provides an estimate of how much the OH estimate in a given year is affected by accounting for the bias. The highest MAE of 1.3% is small compared to the full envelope of each individual OH inversion

Table 2.3: Five metrics that describe the outcome of the two-box inversions. The two-box inversions listed are the standard set-up, four inversions with one bias implemented, and one inversion with all biases implemented. From left to right: 1) Mean absolute error (MAE) in OH anomalies between the standard inversion and each respective inversion. 2) Trend in OH over the 1994-2015 period. 3) Mean lifetime of MCF with respect to OH (tropospheric burden MCF/loss to OH). 4) Mean total tropospheric lifetime of CH₄ (tropospheric burden CH₄/total loss CH₄). 5) Mean annual CH₄ emissions (with soil sink).

Implemented bias(es)	MAE OH [%]	OH trend [%/yr]	τ_{OH} MCF [yr]	τ_{trop} CH ₄ [yr]	CH ₄ emissions [Tg/yr]
None/Standard	-	-0.02 ± 0.15	5.7	9.2	522
Interhemispheric transport	1.07	0.05 ± 0.14	5.9	9.4	510
Surface sampling	0.85	0.09 ± 0.15	6.0	9.6	501
OH ratio	0.68	0.00 ± 0.15	5.5	8.7	546
MCF stratospheric loss	0.68	0.04 ± 0.14	5.3	8.6	555
All four	1.28	0.18 ± 0.15	5.5	8.8	539

(3 – 4%). This means that in terms of interannual variability over the entire period, the outcome was not much affected by the biases. However, as most biases showed their strongest trends over short periods, the peak values of the differences between inversions even out somewhat when averaging over the entire period.

Secondly, we derived an OH trend for each inversion set-up. As described in Section 2.2.1, we mapped the uncertainty of each inversion set-up in a Monte Carlo ensemble of inversions. We fitted a linear trend to the derived OH timeseries of each ensemble member. From the resulting collection of linear fit coefficients, we derived a mean linear fit coefficient and its standard deviation. Differences between the OH trends derived from the different inversions are insignificant. However, it is interesting to see that when all four biases are combined, we derived a shift to more positive OH trends. In the standard inversion, 43% of the ensemble shows a positive trend, whereas in the four-bias inversion 88 % of the ensemble shows a positive trend.

The final three metrics are the tropospheric lifetime of MCF with respect to OH ($(k_{MCF+OH}[OH])^{-1}$, as in Equation 2.1), the total tropospheric lifetime of CH₄ ($(k_{CH_4+OH}[OH] + l_{other})^{-1}$, as in Equation 2.1) and the derived global mean CH₄ emissions, averaged over the 1994–2015 period. For global CH₄ emissions, we added the soil sink (32 [26–42] Tg/yr (Kirschke et al., 2013)), which was not included in the two-box model set-up. Naturally, these three are strongly correlated. When we compare the relative differences in for example the lifetime of MCF with respect to OH between different inversion set-ups to the MAE in anomalies, it is clear that the systematic offset between the different inversions (up to 10%) was much higher than the differences in anomalies (up to

1.3%). This is similar to what was seen for the biases themselves, where the systematic component tended to be much higher than the temporal variations (e.g. the bias in the IH OH ratio, shown in Figure 2.4). We discuss this offset in more detail in Section 2.4.

2.4 Discussion

A first point that deserves discussion is the low global CH₄ emissions (1994-2015) we derived compared to those reported in literature. Our best estimate corresponds to 539 Tg/yr (Table 2.3), which is significantly lower than the 580-600 Tg/yr estimates reported by the two-box inversions of Turner et al. (2017) and Rigby et al. (2017). Our estimate is also on the low end of 3D modelling studies: Saunio et al. (2016) derived CH₄ emissions from 30 3D model inversions, and found emissions of 558 [540-570] Tg/yr over the 2000-2012 period. Bousquet et al. (2006) performed a full 3D inversion of CH₄, using OH fields that were optimized against MCF in a separate 3D model inversion (Bousquet et al., 2005). In their study, CH₄ emissions of 525 ± 8 Tg/yr were found over the 1984-2003 period, so that their estimate does not include the renewed CH₄ growth.

We found that several factors contribute to the differences. Firstly, in the model used by Turner et al. (2017) the atmospheric mass was taken as the global atmospheric mass ($5.15 \cdot 10^{18}$ kg), whereas we used the tropospheric mass ($4.4 \cdot 10^{18}$ kg). When we ran our two-box inversion with the global atmospheric mass, we also found emissions close to 600 Tg/yr. Secondly, we could close the gap with Rigby et al. (2017) by adjusting our a priori two-box model parameters. Specifically, when we adopted an IH exchange rate and an IH OH ratio similar to theirs (1.4 yr^{-1} and 1.07 respectively) in our standard inversion, we found global CH₄ emissions of 595 Tg/yr. This points to a strong sensitivity of the derived CH₄ emissions to these parameters of the two-box model, which in our case are derived from full 3D TM5 model simulations.

In our standard 3D simulation, the IH gradient of MCF tended to be overestimated compared to observations from the NOAA network up to 2005, while global mean mixing ratios were captured much better. Translated into our two-box model, an inversion would tend to reduce MCF emissions to efficiently bring down the IH MCF gradient. To subsequently close the global MCF budget, OH will be reduced, resulting in lower global CH₄ emissions in the two-box model inversion. The lower CH₄ emissions we derived in our two-box model inversion are thus in line with the overestimated MCF latitudinal gradient in TM5. There are several possible explanations for this overestimate.

Firstly, MCF emissions that we used in our 3D simulation were too high. In our two-box inversion, we found significantly lower MCF emissions ($\sim 10\text{--}30\%$) than the prior estimate based on emission inventories, with the exception of the 2010–2014 period. Liang et al. (2017) also derived MCF emissions from the IH gradient, and found these to be systematically lower than those based on bottom-up industrial inventories. Secondly,

the NH/SH OH ratio might be higher than 0.98 (Spivakovsky et al. (2000); Patra et al. (2011)) and more in line with higher estimates from atmospheric chemistry simulations (Naik et al., 2013). Thirdly, a higher fraction of MCF emissions could be located in the SH (15-20% instead of 5-10%). Finally, IH exchange in TM5 could be too slow. A combination of the last two points would also arise if MCF emissions moved from NH mid-latitudes to NH low-latitudes (e.g. India), since low-latitude emissions will be exchanged more rapidly with the SH. At this point it is not clear which of these explanations is most likely.

As is acknowledged in the previous two-box inversion studies of OH (Rigby et al. (2017); Turner et al. (2017)), the problem of deriving OH from MCF and to a lesser degree from CH₄ is strongly underconstrained. Therefore, many solutions fit the problem almost equally well. Moreover, a best estimate, or most likely solution, derived from a two-box model is a function of uncertain input parameters. For example, if a priori it is assumed that OH can only vary within a small band of 2%, then a most likely solution with small OH variations will be found. In this study, we have identified a number of parameters which show variations outside of conventionally assumed bounds. As such, for these parameters, the variations we find are never fully explored in a conventional two-box model inversion: even if done as comprehensively as in Rigby et al. (2017). A clear example is stratospheric loss of MCF, which is generally assumed to have only small variability (10 to 20%). Here, we found a persistent 68% drop in loss of MCF to the stratosphere. Potentially, this loss rate can recover if MCF emissions stop decreasing. Similarly, we find variations in transport of MCF of up to 20% that persist for multiple years, compared to a conventional uncertainty in IH exchange of 10%. In the 1994-1998 period, during a period of strong redistribution of MCF, the individual impact of each of the four biases was quite high, though when combined in one inversion the biases partly cancelled. During the 1998-2007 period, derived OH was less sensitive to the derived biases, likely due to a combination of a small role of uncertain emissions in the MCF budget (Montzka et al., 2011) and a period of relatively small redistribution of MCF. After this period, as MCF abundance continued to decline, we saw a growing impact from the IH exchange bias, as increasingly MCF emissions were constrained from the IH gradient, rather than from the emission inventory.

Another crucial parameter in the two-box inversion is the uncertainty in the global mean mixing ratios and in the IH gradient, as these uncertainties quantify the information content of the observational records. We provided an independent estimate of the uncertainty using 3D model output in Section 2.3.1, summarized in Table 2.2. We can compare the uncertainties we find to observational uncertainties as derived from bootstrapping by Turner et al. (2017). They find uncertainties in hemispheric means of 6-8 ppb for CH₄ and of 5-6% for MCF. Clearly, this is much higher than what we find and their uncertainties seem an overestimate considering the limited sensitivity of our result to a different source-sink distribution. In their most likely solution, derived OH variations were such that the

observed post-2007 renewed growth of CH_4 coincides with a decrease in CH_4 emissions. This solution does not fall within the uncertainty envelope we derived here (right panel in Figure 2.6). The difference in observational uncertainties is likely an important reason for this: their solution corresponds to a statistical inversion framework where less weight is given to the observations.

In the end, conclusions from our study and those drawn by Rigby et al. (2017) and Turner et al. (2017) remain qualitatively similar. The post-2007 renewed growth of CH_4 need not be caused by a sudden increase in emissions in 2007. Rather, emissions could have increased more gradually over the 1994-2007 period, while CH_4 growth was suppressed temporarily by elevated OH levels. The lack of sensitivity of the inversion to the bias corrections and the large remaining uncertainty envelope in the final inversion both indicate that there are other parameters that result in significant uncertainties. Examples are the emission fraction in the NH, observational uncertainties and uncertainty in emission timing of MCF. Thus, while a first step can be made through the incorporation of 3D model information, we confirm the conclusion drawn in Rigby et al. (2017) and Turner et al. (2017) that the current state of the problem is still strongly underdetermined.

In another recent study, an effort was made to find tracer alternatives to MCF (Liang et al., 2017). For this, their suggested method was to use 3D model output to improve the results of a two-box model through intelligent parametrizations. Clearly, this is similar to the work described here. For example, similar to us, they found different IH exchange time scales for different tracers. However, we explicitly resolved the two-box model in the 3D framework, while their study focused mostly on fitting parameters empirically to find a match between two-box and 3D model results. Additionally, for the parametrization, Liang et al. (2017) used hemispheric mean mixing ratios derived from the surface network, whereas we based mixing ratios on the full (hemispheric) troposphere in TM5. We identified a trend and strong, persistent variations in IH transport (CH_4 and MCF) and in the surface sampling bias (MCF) which were not identified in Liang et al. (2017). Additionally, they described a two-box strategy in which two tracers are used to derive the IH OH ratio, which can then also be used for other tracers. Our work suggests that there should be careful consideration of different IH OH ratios seen by different tracers, and potential trends therein. A two-box inversion is sensitive to the IH OH ratio, and we have shown that the effective IH OH ratio a tracer is exposed to depends strongly on that tracer's source-sink distribution. Some of the differences between their findings and ours may be explained by the definition of hemispheric mean mixing ratio (surface-based versus full troposphere), but further reconciliation of the two approaches in future research is necessary.

It is worth noting that the TM5 model, on which the two-box parametrization is based, has its own limitations, and so has treating TM5 as 'the truth'. For example, our simulations were done on the coarse horizontal resolution of $6^\circ \times 4^\circ$. This would have impacted how

well NOAA background sites were actually situated in the background. We checked that the TM5-derived observational timeseries were not systematically more polluted than the real-world NOAA-GMD observations. For this, we detrended and deseasonalized the CH₄ and MCF timeseries per surface site, and quantified the spread in the residuals. At most sites, we found no offset between residual spread in the TM5-derived versus the real-world timeseries. At a small number of sites, TM5-derived timeseries showed more spread in residuals, while at others the spread was less. Therefore, we found no evidence for systematic biases in TM5-sampled observations. Additionally, any transport model is susceptible to some form of transport errors, and using a different 3D model for the two-box parametrization will likely result in different parameters. Therefore, we are careful in suggesting quantitative interpretation of our results. Certain aspects of the biases, such as a slow-down of MCF loss to the stratosphere and the strong variations in IH transport of MCF, are likely to also be found in other 3D transport models, as they are a direct consequence of the MCF emissions drop. Other aspects, such as the exact interannual variations of IH transport of CH₄, or the 7% offset between the physical OH ratio and the effective OH ratio, should be interpreted with more care, as these more strongly depend on the input emission and loss fields, and on the exact treatment of transport in the 3D model. Additional sensitivity tests done in multiple transport models can help in identifying sensitivities of the derived bias corrections. However, our analysis does show a potential for these biases to arise and TM5 is a good starting point for exploring them, as TM5 has provided a strong basis for a wide variety of studies in the past (e.g. Alexe et al., 2015; van der Laan-Luijkx et al., 2015; Bândă et al., 2016).

2.5 Summary and Conclusions

In this study, we investigated variations in the global atmospheric oxidizing capacity, in conjunction with variations in the global CH₄ budget. We specifically revisited the use of two-box models to infer information about these quantities using global observations of MCF and CH₄.

We identified four two-box model parameters that can benefit from 3D model-derived information. Two of these are known and obvious (IH transport; surface sampling bias), while the other two are less so (stratospheric loss; IH OH ratio). Two-box model parameters for these processes that were quantified from full 3D model output showed strong temporal trends mainly for MCF, which have not been identified in any previous research. In general, the biases resulted from a combination of variations in transport and in the spatio-temporal source-sink distributions of each tracer.

We tested the impact of each of the biases in a two-box model inversion. As expected, we found that absolute OH and thus absolute CH₄ emissions show large deviations between the different inversions ($\sim 10\%$). Given that large parts of these deviations were constant through time, they do not necessarily impact conclusions of past two-box modelling studies

that focused on interannual variations.

Compared to the absolute differences, we found only small differences in OH anomalies (up to 1.3%, averaged over 1994-2015) relative to the full uncertainty envelope found here (3-4%) or in Rigby et al. (2017) (8%). This indicates that significant uncertainties in parameters unrelated to the identified biases remain. As such, we confirm in large part the conclusions drawn by Rigby et al. (2017) and Turner et al. (2017) regarding the underdetermined state of the problem. In the end, we did find that the conclusions one can draw from each individual inversion could be strongly affected by the bias corrections: in the standard inversion only 43% of our Monte-Carlo ensemble showed a positive trend in OH over the 1994-2015 period, compared to 88% in the four-bias inversion.

The identified two-box model biases contribute to the already significant uncertainty in derived OH, and properly accounting for them can be a piece in the puzzle of improving constraints on OH. Moving forward, a likely next step is to incorporate more tracers in an effort to further tighten constraints on OH. In such a scenario, the tracer-dependent nature of the biases will likely increase the bias impact, and a proper 3D model analysis for each tracer becomes even more important. Already, efforts have been made to do so (Liang et al., 2017), and in this study we provide further suggestions for such an approach. A distinct advantage in this approach is that information from multiple 3D transport models can be used to tune the two-box inversion, making the inversion outcome less reliant on transport parametrizations of any single 3D transport model. Additionally, computational efficiency of simple models allows for complex statistical inversion frameworks, incorporating, for example, hierarchical uncertainties (Rigby et al., 2017).

On the other hand, the biases are often dependent on the sources and sinks used in the 3D model simulation. As such, a feedback loop between the two-box inversion and the 3D transport models might be necessary to correctly derive bias corrections, which makes such an analysis cumbersome. Additionally, a bias such as that in IH exchange of MCF might be difficult to resolve at all, because IH exchange of MCF is ill-defined in a two-box model (see Section 2.3.1 and Supplement 2.S4). Therefore, we deem it important that a multi-tracer inversion in a full 3D model should also be performed, similar to the 3D inversion of MCF performed by Bousquet et al. (2005), but extended to more recent years. As an added advantage, a 3D model inversion would increase the pool of potential tracers that can be implemented to constrain OH. For example, the short-lived tracer ^{14}CO has been identified as a potential tracer to constrain OH (Brenninkmeijer et al., 1992; Quay et al., 2000; Krol et al., 2008)), but would not be implementable in a two-box model.

Supporting Information to Chapter 2

2.S1 Procedure to obtain hemispheric-average MCF mixing ratios

To aggregate individual measurements of atmospheric air samples to a consistent record of hemispheric average mixing ratios requires correct interpretation of the information contained in the samples. For CH_4 , CO_2 and in principle also for other gases, NOAA has developed a comprehensive interpolation algorithm that approximates these averages. For example, the algorithm functions with gaps in the data and with a changing station network. For MCF, such an algorithm is not available. Instead, we applied an averaging scheme similar to Montzka et al. (2011). In this scheme, we first generated monthly means for each station that cover the whole time period of interest. Next, we aggregated the individual stations to hemispheric averages.

Aggregating measurement data to monthly means requires several steps. Firstly, we linearly interpolated between measurements to obtain daily data. These daily averages were then straightforwardly averaged to monthly means. This worked reasonably well when each month has one or more measurements available. However, for most stations there were multi-month gaps. If more than one consecutive month was missing, we instead substituted values of nearby stations, corrected by an offset assumed to be constant through time. As an example, the offset between Cape Grim and South Pole is relatively constant, so if either station misses multiple months, a good approximation is obtained by substituting the other's values corrected for by the offset. This worked well when multiple stations that represent a similar air mass were available. The exception was the American Samoa Observatory (SMO), for which we only linearly interpolated in time. The implications of this exception were small, as the SMO record had few multi-month gaps in data coverage. Finally, to combine the monthly means of individual stations into hemispheric averages, we computed a weighted average. We combined sites that represent similar air masses in groups. A group's weight reflected the fraction of air mass that the group represented. Mostly, we used the cosine of the average latitude of the sites included in each group. Global mean mixing ratios that were derived using this procedure are shown in Figure 2.S1 (red line in the right panel).

A crucial question in this procedure is how to estimate uncertainties in derived averages. A common technique for estimating uncertainties is a bootstrapping procedure, used, for example, in Turner et al. (2017). In bootstrapping, different surface network realizations are randomly generated by counting some stations multiple times, while others are left out entirely. The spread in averages derived from the different network realizations then gives some measure for the uncertainties of these averages. For CH_4 , this is how the uncertainties reported by NOAA (and used in this study) are derived. For MCF how-

ever, fewer stations are available, so that bootstrapping becomes less robust. Moreover, timeseries of annual means per station, as we derived them here, can be interdependent because of the substitution-technique described above. For these reasons, we adopted a different approach to derive uncertainties for MCF.

The procedure we adopted to derive uncertainties for MCF was based on the method presented by Rigby et al. (2017). In Rigby et al. (2017), the annual mean standard deviation (SD) of the MCF timeseries per surface site was used as a measure of the uncertainty in hemispheric means. We also used this measure, but we first removed an exponential trend from the timeseries at each station, before deriving the annual mean SD. This trend removal prevented the long-term trend in the MCF measurements from dominating the derived uncertainties. The SD in the residuals mainly resulted from source-sink variations, meteorological variability and measurement uncertainty. We deem this a more appropriate measure for uncertainty than the long-term trend. The trend removal reduced the annual mean SD per station from around 5% (as in Rigby et al. (2017)) to 1.5%. Next, we used a Monte-Carlo simulation to propagate the uncertainty per station to an uncertainty in the global and hemispheric means. In the Monte-Carlo simulation, we perturbed the annual means per station proportional to the annual mean SD per station, and derived new global and hemispheric means. The SD of each quantity derived from the resulting ensemble was used as a measure for uncertainty. This Monte-Carlo approach accounts for the fact that stations are not expected to be influenced by the same variability at the same time, so that their errors should be largely uncorrelated. The time-dependent uncertainties in the global mean growth rate and in the rate of change of the IH gradient varied between 0.5–0.7% and 1.0–1.5% respectively. As can be expected from the modifications we have made to their method, this is significantly lower than the uncertainties adopted by Rigby et al. (2017), though still higher than the errors derived from the TM5 simulations (see Table 2).

The timeseries of MCF we have derived, along with the corresponding uncertainties, are presented in a supplemental dataset, along with the same parameters for CH₄.

2.S2 The nudged simulation

In this supplement, we discuss the nudged simulation and some of its results. In the nudged simulation, we nudged the model to observations from NOAA surface sites, following a procedure analogous to Bândă et al. (2015). In essence the nudging approach is a crude mass-balance inversion. We compared modelled mixing ratios at ground level at the dateline to latitudinally interpolated, monthly mean observations from the NOAA network, under the assumption that both of these are representative of the background atmosphere at that latitude. Then, we scaled the entire latitude band according to the model-to-observation ratio using a relaxation time of 100 days, to ensure that the model follows the long-term trend in real-world observations. Since the nudging does not de-

pend on the a priori emission and loss fields, the nudged run provides a good test for the sensitivity of our conclusions to the source-sink distribution used.




Figure 2.S1: Global, annual mean mixing ratios for CH_4 (left) and MCF (right) as derived from model-sampled observations from the standard and nudged simulation, and from real-world NOAA observations.

Figure 2.S1 shows the global, annual mean mixing ratios of MCF and CH_4 as derived from model-sampled observations in the standard simulation and in the nudged simulation, and as derived from real-world NOAA data. For all three timeseries, the same averaging algorithms were used. The standard simulation for CH_4 performed poorly, as can be expected from a simulation with annually repeating emissions and constant loss fields. The nudged run captures the rough characteristics of CH_4 growth rate variations much better: the 2000-2006 stagnation and the renewed growth after 2007 are reproduced quite well. For MCF, the standard run already captured the important aspects of the global atmospheric decline. However, from 2008 onwards the standard simulation increasingly underestimated MCF mixing ratios. Therefore, the nudged run tends to add MCF towards the end of the simulation.

Observations sampled from the nudged run and real-world NOAA observations still differed for two reasons. Firstly, for the nudging, we did not compare mixing ratios at the site location, but at the dateline. Because we scaled the entire latitude band in the lower atmosphere, also the site location is nudged, but only indirectly. This difference is likely driving the systematic offset between annual means derived from the nudged runs and from the real-world annual means of CH_4 (left panel in Figure 2.S1). At the dateline itself, the latitudinal fields used for nudging and the its modelled counterpart overlapped near perfectly. Secondly, only a subset of the measurement sites that were used for averaging, were used for nudging. We want to stress that the resulting differences are not important for our purposes, as our main objective was to test the sensitivity to a different, somewhat more realistic emission distribution, rather than provide a full inversion.

To give an impression of the nudging strength, Figure 2.S2 shows the global emissions both from the standard and from the nudged TM5 simulation, if the derived nudging




Figure 2.S2: Global, annual emissions of CH_4 (left) and MCF (right). Standard emissions are the emissions as used in all our TM5 simulations, while the standard+nudged emissions are the result of adding the global, annual mean nudging to the standard emissions.

term (positive or negative) is assumed indicative of variations in emissions. Note that this interpretation is arbitrary, as nudging can also indicate variations in the loss fields. For CH_4 , nudging can be quite extreme, with interannual variations of up to 40 Tg/yr. Though part of these emission variations may be real, the extreme extent of the variations likely resulted from overfitting or incorrect assignment of emissions to certain latitudes. This reveals the limitations of the nudging approach. For MCF, there was a clear need for additional emissions in later years with respect to our initial assumption that emissions declined with 20%/year after 2006. Since differences in MCF mixing ratios between the standard run and real-world observations were especially high at Mace Head (not shown), this would indicate some ongoing European emissions.

2.S3 Interhemispheric asymmetries in correlations between k , OH and CH_4/MCF

2.S3.1 Introduction

In Section 3.1.3 it is outlined that the IH OH ratio to be used in a two-box model is not necessarily the same as the physical IH OH ratio. For example, in our 3D simulations runs, the IH ratio of the OH fields was 0.98, while the OH ratio seen by CH_4 was 1.05: a difference of 7%. The differences are a result of first averaging k , OH, and CH_4/MCF per hemisphere, per year individually, and subsequently calculating the hemispheric loss (as in a box model), compared to a 3D model, where the loss is integrated over the cells that make up the hemisphere. Because of spatio-temporal correlations between the parameters involved (OH, k_{OH} , tracer distribution) and from hemispheric asymmetries in these correlations, differences between these two approaches arise. In this section, we outline in more detail which correlations predominantly drive the differences.




Figure 2.S3: The annually averaged rr ratios, as defined in Equation 2.10, for the different components of chemical loss of CH_4 (left) and MCF (right).

First, we define the ratio r in equation 2.9. The overline denotes the average over spatial and/or temporal dimensions. The ratio r is unity for uncorrelated parameters, and $r > 1$ for positive and $r < 1$ for negative correlations. The ratio rr denotes the IH ratio of r (Equation 2.10). If $rr > 1$, the correlation is stronger in the NH than in the SH, and the reverse is the case for $rr < 1$. Thus, while r is a measure for the strength of the correlation in one hemisphere, rr is a measure for the IH asymmetry in this correlation.

$$r[a, b, \dots] = \frac{\overline{a \, b \, \dots}}{\overline{a} \, \overline{b} \, \dots} \quad (2.9)$$

$$rr[a, b, \dots] = \frac{r[a, b, \dots]_{NH}}{r[a, b, \dots]_{SH}} \quad (2.10)$$

Ratios $rr[k, OH, \text{CH}_4]$ and $rr[k, OH, \text{MCF}]$ were both found to deviate significantly from 1 in the 3D simulations we used to tune our two-box model. This means that multiplying the hemispheric, annual averages of each parameter in a two-box model does not reproduce the loss rate calculated in the full 3D simulation. We emphasize that this result and the following analysis are specific to the OH field described in Spivakovsky et al. (2000) and the implementation in TM5. There are three correlations that can drive the difference: between k_{OH} and OH, between k_{OH} and CH_4/MCF and between OH and CH_4/MCF . The contributions by each of these three to the asymmetry is shown in Figure 2.S3. From this figure, it can already be inferred that the main driver of both $rr[k_{OH}, OH, \text{CH}_4]$ and $rr[k_{OH}, OH, \text{MCF}]$ being greater than 1 is the correlation between k_{OH} and OH, while the time-dependent behaviour of the IH OH ratio of MCF is driven predominantly by the correlation between OH and MCF. In the following, we discuss each of the three correlations in more detail.




Figure 2.S4: The annually averaged evolution of the $rr[k_{CH_4}, OH]$ ratio, as defined in Equation 2.10. The contribution from each dimension is isolated by only considering that dimension for the rr ratio. Note that the black line corresponds to the blue line in the left panel of Figure 2.S3

2.S3.2 Correlation between reaction coefficient k_{OH} and OH

As the reaction coefficient k_{OH} of both MCF and CH_4 are a direct function of only temperature (T), the spatial distribution of k_{OH} and T are identical. For this reason, the correlation between k_{OH+CH_4} and OH and k_{OH+MCF} and OH are similar. Here, we discuss only k_{OH+CH_4} . OH and k_{OH} are strongly positively correlated: spatially, with maxima in the tropics, and minima at the poles; and temporally, with maxima in the summer hemisphere and minima in the winter hemisphere. However, this correlation is stronger in the NH than in the SH, resulting in $rr[k_{OH}, OH] > 1$ (blue lines in Figure 2.S3; black line in Figure 2.S4). This effect turns out to be the main driver of both $rr[k_{OH}, OH, CH_4]$ and $rr[k_{OH}, OH, MCF]$ being greater than 1. The asymmetry is the integrated effect of correlations in the latitudinal, longitudinal, vertical and seasonal dimension. To assess the contribution of each dimension to the correlation asymmetry, we first averaged k_{OH+CH_4} and OH over all but one dimension before computing rr . In this way, we isolate the influence of one dimension on the asymmetry. The result is shown in Figure 2.S4.

The largest effect is seen when isolating the vertical dimension, which gives an rr close to the full rr (red line in Figure 2.S4). Further analysis reveals that this is because the vertical distribution of OH in the NH and SH is different. More precisely, the maximum in OH is located at lower altitude in the NH than in the SH (Figure 2.S5). This is likely driven by higher NO_x emissions in the NH, which increases OH recycling close to the surface (Spivakovsky et al. (2000)). At lower altitudes temperatures are higher, and as a consequence the correlation between k_{OH} and OH is strongest in the NH.

An effect in the same direction is seen for the seasonal component. Figure 2.S6 shows the seasonal cycle of OH and k_{OH+CH_4} in each hemisphere. In the NH, the two seasonal




Figure 2.S5: The average vertical distribution of tropospheric OH (solid lines, bottom y-axis) and $k_{\text{OH}+\text{CH}_4}$ (dashed lines, top y-axis) NH (blue) and in the SH (red).

cycles are very well in line, while in the SH less so. Moreover, the seasonal cycle in temperature in the NH is stronger than in the SH. Both effects are likely a consequence of the higher land fraction in the NH, which results in a lower surface heat capacity. This drives a stronger and more direct response of temperature to seasonal variations in solar irradiance.

Finally, there is an opposing asymmetry in the latitudinal correlation. Figure 2.S7 shows that tropical OH is highest in the SH, whereas extra-tropical OH is highest in the NH. As temperatures are highest in the tropics, the correlation between temperature and OH is highest in the SH.

As expected, the dimension of longitude has only a small effect: both temperature and OH are relatively uniform zonally (not shown).

2.S3.3 Correlation between OH and CH_4/MCF

Figure 2.S3 shows that the trend in the IH OH ratio of MCF is largely driven by the $rr[\text{OH}, \text{MCF}]$ component of the correlation. It follows intuition that when the distribution of MCF strongly changes, then also the MCF distribution with respect to OH changes. Similar to $rr[k, \text{OH}]$, we can isolate the influence of each dimension on the asymmetry in the correlation (Figure 2.S8).

The largest contribution to $rr[\text{OH}, \text{MCF}]$ comes from the latitudinal component (green line in Figure 2.S8). This is driven by the decreasing IH MCF gradient through time (Figure 2.S9). Initially, MCF decreases towards the tropics in the NH, and increases towards the tropics in the SH. As OH is highest in the tropics, MCF and OH are negatively




Figure 2.S6: The seasonal cycle of tropospheric OH (left axis, solid lines) and k_{OH} (right axis, dashed lines) in the NH (blue) and in the SH (red).




Figure 2.S7: The latitudinal distribution of tropospheric OH (left axis, solid lines) and k_{OH} (right axis, dashed lines) in the NH (blue) and in the SH (red).

correlated in the NH and positively correlated in the SH. This drives $rr[OH, MCF] < 1$. The IH gradient largely disappears over time, so that $rr[OH, MCF]$ approaches 1.

A weaker, opposite effect is seen in the vertical. Initially, NH MCF decreases with altitude strongly, while SH MCF slightly increases with altitude (Figure 2.S10). As OH maximizes at low altitude, this drives a positive correlation in the NH, compared to the weak negative correlation in the SH. This results in an $rr[OH, MCF] > 1$. As the gradients disappear over time, so too does the asymmetric correlation.

For CH_4 , spatio-temporal gradients are smaller. However, it too has an $rr[OH, CH_4] < 1$, mostly due to the IH gradient in CH_4 .




Figure 2.S8: The annually averaged evolution of the $rr[\text{MCF}, \text{OH}]$ ratio, as defined in Equation 2.10. The contribution from each dimension is isolated by only considering that dimension for the rr ratio. Note that the black line corresponds to the green line in the right panel of Figure 2.S3

2.S3.4 Correlation between reaction coefficient k_{OH} and CH_4/MCF

As indicated earlier, correlations between k and OH are very high, and thus the correlation between k_{OH} and CH_4/MCF behaves largely similar to the correlation between OH and CH_4/MCF . However, interestingly, the net effect on rr can be opposite. In the early 90s, we find a positive trend in $rr[\text{OH}, \text{MCF}]$, while the trend in $rr[k_{\text{OH}+\text{MCF}}, \text{MCF}]$ is negative (right panel in Figure 2.S3). This is because temperature, and thus k , have stronger and more unidirectional gradients in the vertical, compared to OH. Therefore, the correlation between mixing ratios and k is dominated by the vertical dimension, whereas the correlation between OH and mixing ratios is dominated by the latitudinal dimension. This detail reveals the intricacies of the compensating effects in the final rr ratio, and thus of the resulting IH OH ratio seen by different tracers.

2.S3.5 Conclusions on the IH OH ratio

In conclusion, the symmetry between the NH and the SH that is implicitly assumed in the typical two-box set-up does not necessarily hold. This means that different tracers potentially require different IH OH ratios. That different tracers might be exposed to different atmospheric oxidative capacities has been noted before Lawrence et al. (2001). However, we quantify exactly the contributions of different parameters and dimensions on the ratio. The most important implication of this analysis is the sheer number of parameters that the final OH ratio hinges on (e.g. IH gradient; vertical profile; temperature dependence of k_{OH}). This makes it very difficult to predict how the ratio would vary under changing conditions and for different tracers. Of course, our results depend strongly on the specific OH distribution and 3D transport model that is used, but the




Figure 2.S9: The latitudinal distribution of MCF in 1990 and in 2015. MCF mixing ratios are scaled by MCF mixing ratios at the equator.

OH distribution and 3D model we use has been shown to result in realistic CH_4 mixing ratio fields (Patra et al. (2011); Huijnen et al. (2010)). Moreover, the observation that such large deviations between tracers and in time are possible, will likely hold for any OH distribution and 3D transport model.

2.S4 Robustness of derived interhemispheric exchange rates

In Section 3.1.1, we present IH exchange rates (k_{IH}) for CH_4 , MCF and SF_6 as derived from a TM5 simulation. Here, we explore the robustness of the derived IH exchange rates in three scenarios. Firstly, we repeated the standard simulation, but with annually repeating meteorological fields (hereafter referred to as the RM simulation). We repeated meteorological fields from the year 2012, but results were qualitatively similar if a different year was chosen. Secondly, we redid the TM5 to two-box parametrization (see Section 2.2.2), but with the two hemispheres demarcated at 8°N instead of at the equator, which is more representative of the average position of the inter-tropical convergence zone (ITCZ). Thirdly, we investigated the exchange rate in the nudged simulation (Supplement 2.S2), which gives a measure for the sensitivity of derived rates to realistic variations in the source-sink distributions. The results are shown in Figure 2.S11.

In the RM simulation, we isolated the influence of source-sink variations on the exchange rate. Since we used annually repeating source and sink fields for CH_4 , the derived exchange rate showed nearly no variability in the RM simulation. This shows that all interannual variability and the trend we found in the standard simulation for CH_4 were driven by transport, and not, for example, by any spin-up effects. For MCF and SF_6 ,




Figure 2.S10: The vertical distribution of MCF in 1990 and in 2015. MCF mixing ratios are scaled by ground-level MCF.

interannual variations in the emission fields were implemented, and this was reflected in the IH exchange rate derived from the RM simulation. For SF_6 , the emissions variations were small and gradual, and so were the variations in the IH exchange rate. For MCF, we found a decrease in IH exchange of 25%, that coincided with the drop in MCF emissions. In the standard simulation this decrease was masked by transport variability, so that the exchange rate minimized in 2000-2005 and then recovered. This indicates that the derived IH exchange rate of MCF is especially sensitive to interannual variations in transport.

Next, we tested the influence of the position of the hemispheric demarcation on the IH exchange rate. For all three tracers, moving the division to 8°N resulted in faster exchange rates. For CH_4 and for SF_6 the offset was nearly constant through time at 35% and 33% respectively. Thus, variability in the exchange rate was found to be robust with respect to the choice of hemispheric demarcation. For MCF, this was not the case. In fact, in the simulation with annually repeating meteorology, the decrease in IH exchange became an increase instead. The choice of hemispheric demarcation is somewhat arbitrary, since there is not one meteorological dividing line which sharply isolates the two hemispheres. Yet the IH exchange variations of MCF were highly sensitive to the demarcation. The result is not so surprising, when considering that the IH exchange rate for MCF became ill-defined post-1998. After 1998, the latitudinal MCF mixing ratio gradient started to show a minimum in the tropics, where most exchange takes place, so that the net IH gradient is not well representative of the gradient against which transport takes place. Thus, if the IH gradient of MCF, as derived from surface networks, is to be used to disentangle OH from emissions, then a method needs to be found that correctly accounts for (variations in) the IH exchange.

Finally, the sensitivity to source-sink variations was quantified in the nudged simulation.

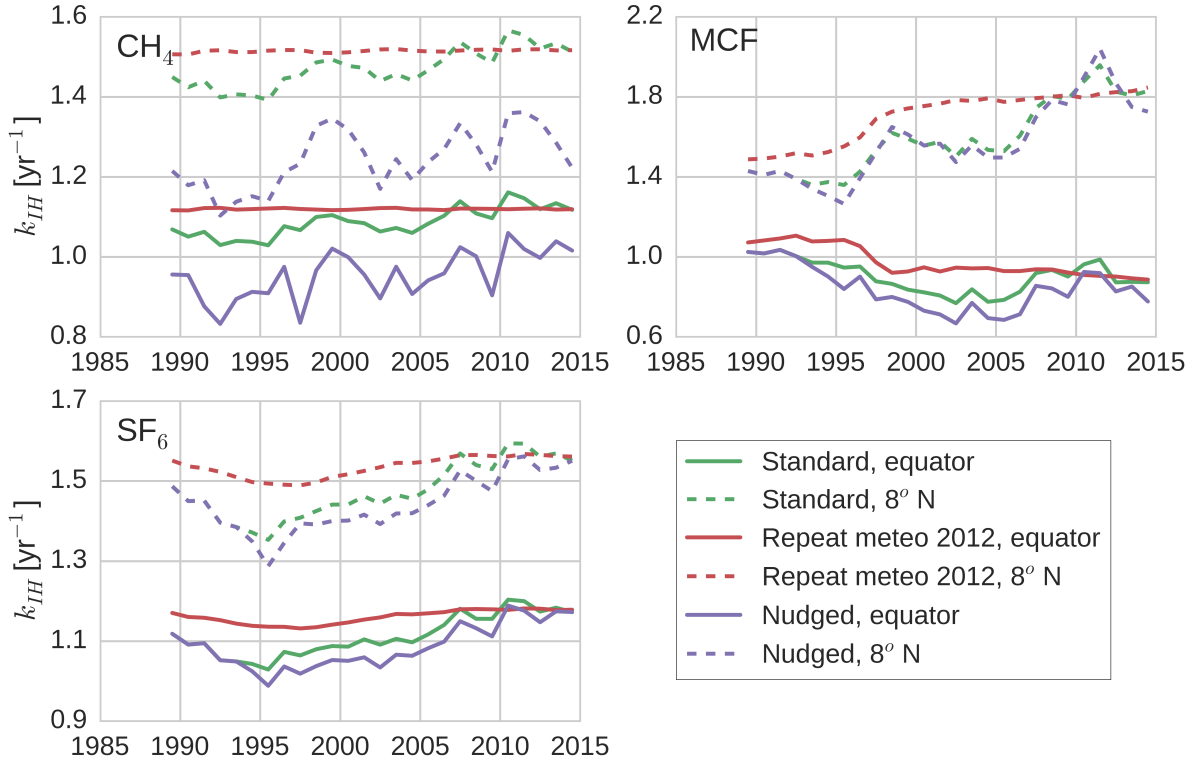


Figure 2.S11: The IH exchange coefficients for CH_4 (top left), MCF (top right) and SF_6 (bottom left) as derived from a set of sensitivity tests. Rates derived from the standard simulation (green), a simulation with annually repeating meteorology (red) and the nudged simulation (purple) are shown. Additionally, for each of the three simulations, we show the sensitivity of the IH exchange rate to the hemispheric demarcation in the derivation, by shifting it from the equator (solid lines) to 8°N (dashed lines). Note that the range on the axes differ for the three tracers.

We found that all tracers were somewhat sensitive to nudging, but that the overall tendencies were conserved. CH_4 was most sensitive, which can be expected, as it was nudged most strongly due to a poor fit to surface observations in the standard simulation. Most notably, we found a systematic shift to slower exchange rates for CH_4 in the nudged simulation. However, in general, years with positive anomalies in the standard simulation still have positive anomalies in the nudged simulation, albeit with a higher amplitude. Additionally, the positive trend for CH_4 and SF_6 persists, as does the minimum for MCF, even if it is slightly modified. Therefore, we conclude that the interannual variability of the derived IH exchange rates is quite robust with respect to source-sink variations of realistic magnitude.

In conclusion, we found that the IH exchange rates presented in Section 3.1.1 can be sensitive to prior assumptions in the 3D transport model, such as the source-sink distribution and the hemispheric demarcation. Especially for MCF we found a high sensitivity to hemispheric demarcation, which indicates that defining a two-box exchange rate for

MCF is difficult. Moreover, we found that IH exchange of MCF was especially sensitive to transport variations. For CH_4 and SF_6 , interannual source-sink variations were not as extreme, which resulted in a small sensitivity of their IH exchange rates to prior assumptions. Thus, the parametrization is in principle robust, unless a tracer's global distribution changed considerably.

Chapter 3

A 3D-model inversion of methyl chloroform to constrain the atmospheric oxidative capacity

This chapter is based on:

Naus, S., Montzka, S. A., Patra, P. K., & Krol, M. C. (2020). A 3D-model inversion of methyl chloroform to constrain the atmospheric oxidative capacity. *Atmospheric Chemistry and Physics Discussions, in review*, p. 1–23.

Abstract

Variations in the atmospheric oxidative capacity, largely determined by variations in the hydroxyl radical (OH), form a key uncertainty in many greenhouse and other pollutant budgets, such as that of methane (CH_4). Methyl chloroform (MCF) is an often-adopted tracer to indirectly put observational constraints on large-scale variations in OH. We investigated the budget of MCF in a 4DVAR inversion using the atmospheric transport model TM5, for the period 1998-2018, with the objective to derive information on large-scale, interannual variations in atmospheric OH concentrations.

We derived interannual variations in the global oxidation of MCF that bring simulated mole fractions of MCF within 1-2% of the assimilated observations from the NOAA-GMD surface network at most sites. Additionally, the posterior simulations better reproduce aircraft observations used for independent validation, compared to the prior simulations. The derived OH variations showed robustness with respect to the prior MCF emissions and the prior OH distribution. The interannual variations were typically small ($<3\%$ /year), with no significant longterm trend in global mean OH concentrations.

The inverse system found strong adjustments of the latitudinal distribution of OH, with systematic increases in tropical and decreases in extra-tropical OH concentrations (both up to 30%). These spatial adjustments were driven by intrahemispheric biases in simulated MCF mole fractions, which have not been identified in previous studies. Given the unexpectedly large amplitude of these adjustments and a residual bias in intrahemispheric gradients, we suggest a reversal in the extratropical ocean sink of MCF in response to declining atmospheric MCF abundance (as hypothesized in Wennberg et al. (2004)). This ocean source provides a more realistic explanation for the biases, possibly complimentary to adjustments in the OH distribution.

While we identified significant added value in the use of a 3D transport model over simpler box models, we also found a trade-off in computational expense and convergence problems. While the effect of the derived temporal OH variations on MCF mole fractions is small, these variations do result in an improved match with MCF observations relative to an interannually repeating prior for OH. Therefore, we consider the derived variations relevant for studying the budget of e.g. CH_4 .

3.1 Introduction

The hydroxyl radical (OH) is the main atmospheric oxidant and plays an integral role in atmospheric chemistry. OH is involved in the removal of a wide variety of toxic pollutants (carbon monoxide, nitrogen oxides), greenhouse gases (e.g. methane (CH₄) and HFC's) and gases that contribute to stratospheric ozone depletion (various HCFC's). Therefore, in order to better understand and constrain these different pollutant budgets, robust constraints on OH are required.

Due to its high reactivity, OH has a short atmospheric lifetime of seconds and is present at low abundances in the atmosphere. This makes direct measurements of the OH concentration (denoted as [OH]) difficult, and extrapolation of these measurements to global scales near impossible. Bottom-up modeling of OH in full-chemistry models is a useful tool to better understand the OH budget. However, this modeling relies on accurate understanding of the complex chemistry involved and of the emissions of many species, and it lacks observational constraints. Variations in atmospheric oxidation can also be estimated indirectly from observed variations in the growth rate of a tracer removed by OH. This "proxy"-method can provide independent observational constraints on OH on large spatio-temporal scales, against which the understanding implemented in full-chemistry models can be tested.

Different species have been proposed and used to constrain OH, with a particular focus on global mean OH concentrations ([OH]_{GM}) (Dentener et al., 2003; Krol et al., 2008; Liang et al., 2017). The most widely adopted of these is methyl chloroform (MCF). MCF was used as a solvent up to the early 1990s, after which its production was phased out in the Montreal Protocol and its subsequent amendments (McCulloch & Midgley, 2001; Rigby et al., 2013). MCF is predominantly removed through oxidation by OH, with secondary sinks by stratospheric photolysis and oceanic hydrolysis, resulting in an atmospheric lifetime of 5-6 years (Chipperfield & Liang, 2013). As MCF is emitted purely anthropogenically and its production is thought to be well-constrained, MCF was identified early on as a potential tracer for OH (Lovelock, 1977; Prinn et al., 1987). This potential improved after the production phase-out and the rapid emission decline, as the MCF budget became dominated by its OH sink. The combination of a period with well-constrained production, followed by a period with a small role of its emissions, makes MCF a good tracer for OH.

Study of the MCF budget has provided important insights into the atmospheric distribution of OH and temporal variations therein. For example, observations of MCF were used to quantify a likely upper limit to interannual variations in [OH]_{GM} of 2-3% (Montzka et al., 2011), which is in line with full-chemistry models. On the other hand, whereas chemistry models typically find higher OH concentrations in the Northern than in the Southern Hemisphere, a modeling study found MCF observations to be most consistent

with interhemispheric parity (Patra et al., 2014).

Use of MCF to constrain $[\text{OH}]_{\text{GM}}$ and the OH distribution is not without challenges, however. Through improvements in measurement techniques, measurement quality has mostly kept pace with the atmospheric decline of MCF, but artifact-free sampling has become more difficult. Additionally, a decline of the atmospheric burden implies an increasingly important role of any small, persistent emissions in the MCF budget. As a further complication, atmospheric decline of MCF has been suggested to result in a reversal of the small but significant ocean sink of MCF at high latitudes (Wennberg et al., 2004). Finally, the surface measurement networks that monitor MCF are relatively sparse and thus they provide limited capacity to distinguish one budget term from another.

These limitations were emphasized in two recent studies that constrained interannual variations of $[\text{OH}]_{\text{GM}}$ in a two-/three-box model of MCF and CH_4 (Turner et al., 2017; Rigby et al., 2017). In their most likely solution, both studies suggested the possibility for an important contribution from OH to CH_4 growth rate variations. However, the looseness of derived constraints on OH variations also allowed for a solution with no variations in OH. In an extension of these studies, we investigated in previous work how the use of a relatively simple box model, rather than a more sophisticated 3D transport model, could have affected these conclusions (Naus et al., 2019). We found that large changes in the MCF budget over time (i.e. the sudden drop in its emissions) resulted in significant changes in, for example, interhemispheric transport of MCF and the stratospheric MCF sink. However, accounting for these changes in our two-box model did not alter the conclusion that MCF-derived constraints on multi-annual variations of OH are too uncertain to determine the exact contribution of OH to the relatively small but important CH_4 growth rate variations.

In this study, we present an inversion of MCF in the 3D chemistry-transport model TM5 aimed at constraining $[\text{OH}]_{\text{GM}}$ and the OH distribution. The advantage of approaching the problem in a 3D transport model, instead of in a box model, is two-fold. Firstly, by explicitly resolving transport, we avoid the transport biases that hamper simple box models. Secondly, we can fully exploit the available observations and the gradients between surface sites. For example, since MCF is predominantly removed in the tropics, the latitudinal distribution of MCF has had a tropical minimum since 1998 (Spivakovsky et al., 2000; Montzka et al., 2000). There is potential for this type of information to provide constraints on OH that has not been exploited by box-model approaches.

The last comprehensive 3D model OH-inversion of MCF investigated the 1980-2000 period (Bousquet et al., 2005). Here, we instead cover the 1998-2018 period. An important difference between these two periods concerns the role of emissions. The 1980-2000 period includes a period with high MCF emissions, as well as the subsequent strong decrease in emissions from 1990 onward, which complicated the interpretation of MCF observations. In contrast, the situation over 1998-2018 was simpler, with a near-constant exponential

decline of atmospheric MCF and a small role for its emissions.

The objective of this study is to investigate information on large-scale variations in OH concentrations contained in measurements of the most promising tracer identified to date, MCF, during the period that follows its drop in emissions, using the most comprehensive tools available to us, in the form of a state-of-the-art inverse system built around a 3D transport model.

3.2 Methods

3.2.1 Inverse method

3.2.1.1 TM5-4DVAR

To simulate atmospheric transport, we used the transport model TM5 (Krol et al., 2005). We operated TM5 on a 6° by 4° horizontal resolution, with 25 vertical layers. Transport in TM5 was driven by offline meteorological fields from the ERA-Interim reanalysis from the European Centre for Medium Range Weather Forecasts (ECMWF) (Dee et al., 2011).

For optimization of OH concentrations and other parameters, we used the inversion framework TM5-4DVAR (Meirink et al., 2008). TM5-4DVAR is an inverse system based on the variational optimization technique 4DVAR. TM5-4DVAR has been applied successfully in previous studies to a variety of problems, such as to constrain the global CH_4 budget (Bergamaschi et al., 2009) and to investigate the Indonesian wildfires (Nechita-Banda et al., 2018). The objective of our set-up of the TM5-4DVAR inverse system is to find the optimal configuration of MCF emissions and OH variations that best reproduce atmospheric observations of MCF. Formally, this objective is quantified as minimization of the cost function J (Equation 3.1).

$$J(\vec{x}) = \frac{1}{2}(\vec{x} - \vec{x}_{prior})^T \mathbf{B}^{-1}(\vec{x} - \vec{x}_{prior}) + \frac{1}{2}(\mathbf{H}\vec{x} - \vec{y})^T \mathbf{R}^{-1}(\mathbf{H}\vec{x} - \vec{y}). \quad (3.1)$$

J is a function of the state \mathbf{x} , which contains all the parameters to be optimized, such as OH concentrations. The cost consists of two terms. First is the deviation from the first guess x_{prior} , weighted by the prior error covariance matrix \mathbf{B} . Second is the difference between simulated MCF mole fractions, calculated in the forward version of TM5 (denoted as \mathbf{H}), and the real-world observations \mathbf{y} , weighted by the observational covariance matrix \mathbf{R} . Additionally, in the 4DVAR optimization, the gradient of the cost function ∇J is calculated and used (Equation 3.2).

$$\nabla J(\mathbf{x}) = \mathbf{B}^{-1}(\mathbf{x} - \mathbf{x}_{pri}) + \mathbf{H}^T \mathbf{R}^{-1}(\mathbf{H}\mathbf{x} - \mathbf{y}), \quad (3.2)$$

with \mathbf{H}^T the adjoint of the forward model \mathbf{H} . The adjoint of TM5 is extensively described in Meirink et al. (2008) and Krol et al. (2008). Since OH chemistry is non-linear and since we optimized emissions non-linearly (see Section 3.2.1.2), \mathbf{H}^T is actually the adjoint of the forward tangent linear model. The derivation of the adjoint OH chemistry is described in Supplement 3.S6.

3.2.1.2 Inversion set-up

In this section we discuss the set-up of the three inversions we performed. First we describe the set-up of the standard inversion (hereafter referred to as REF); next we describe the corresponding \mathbf{B} matrix; finally we describe the two variations of the REF inversion which we performed. Note that the \mathbf{R} matrix is discussed in Section 3.2.2.

In the REF inversion, we used the MCF source and sink fields from the TransCom-CH4 project (Patra et al., 2011). Loss fields for OH, stratospheric photolysis and ocean uptake, as well as initial fields of MCF, are described in the TransCom-CH₄ protocol. Briefly, the OH fields are a combination of tropospheric OH fields from Spivakovsky et al. (2000), scaled by a factor 0.92, and stratospheric OH fields derived with the 2D MPIC chemistry model (Brühl & Crutzen, 1993). The ocean flux is a first-order sink proportional to MCF mole fractions in the lowest model layer and a spatially variable uptake rate, which maximizes in the tropics. Stratospheric photolysis fields were generated with simulations of the ACTM model (Patra et al., 2009). TransCom MCF emissions were available only up to 2006. In subsequent years, we repeated the 2006 spatial distribution. We pre-optimized annual total MCF emissions to approximately reproduce global mean MCF mole fractions, as derived from the NOAA surface network, assuming no interannual variations in OH. The procedures to obtain hemispheric and global mean MCF mole fractions from the NOAA surface network are extensively described in Supplement 3.S1 of Naus et al. (2019).

We included scaling factors for both OH concentrations and MCF emissions in the state vector for optimization. Emissions were optimized monthly in each grid box with a horizontal correlation length of 500 kilometres and a temporal correlation of 9.5 months. We adopted a grid-box error in emissions of 50% up to 2005 that increases with 15% per year, up to 200% in 2015. This increase ensures that when prior emissions become low, the absolute error on emissions remains large (around 4 Gg/year), which reflects the lack of prior constraints on MCF emissions in later years. In the construction of the \mathbf{B} -matrix, OH concentration fields were optimized monthly, in 45 latitudinal bands (i.e. on our native latitudinal resolution), with an error per band of 10%. We imposed high correlations between the latitudinal bands of 0.8 (corresponding to a correlation length of 1500-2000 km), and a temporal correlation length of 12 months. Since there are fewer state elements for OH than for emissions, we increased the cost of adjusting OH relative to that of adjusting emissions by an additional factor 5, similar to e.g. Maasakkers et al.

(2019). This adjustment ensures that the cost of adjusting total emissions or $[\text{OH}]_{\text{GM}}$ by $1\text{-}\sigma$ of the global uncertainty is comparable.

Both OH concentrations and emissions were optimized with non-linear scaling factors, which preclude negative emissions and OH concentrations. This method was introduced in Bergamaschi et al. (2009) and we describe it more extensively in Supplement 3.S5. Due to the non-linear nature of this inverse problem, we use the M1QN3 solver (Gilbert & Lemaréchal, 1989). The downside of this choice is that we cannot retrieve the posterior covariance matrix. We instead tested the robustness of derived solutions with respect to the OH and the emission distribution, two very likely sources of uncertainty, in two additional inversions.

In our second inversion, referred to as POP, we redistributed the same annual total MCF emissions as in the REF inversion proportional to population density (as retrieved from CIESIN, Columbia University (2018)). In the third inversion (referred to as TM5OH), we used the same emissions as in the REF inversion, but adopted a tropospheric OH distribution based on a simulation of the year 2006 performed with the full-chemistry version of TM5 (Huijnen et al., 2010), combined with the same stratospheric distribution as in the REF inversion. Differences between the two OG distributions are typically 10-15%, depending on the latitude. The TM5OH distribution has relatively higher OH concentrations in the Northern hemisphere (see Figure S8).

On a final note, while we optimized scaling factors for OH concentrations per latitudinal band, we prefer to discuss the change in global mean oxidation in further sections, rather than the change in $[\text{OH}]_{\text{GM}}$. We quantified the change in oxidation as the atmospheric mass-weighted average of $k(T) \cdot [\text{OH}]$, with $k(T)$ the temperature-dependent reaction rate between OH and MCF (Burkholder et al., 2015). This is necessary, because we allowed for adjustments in the latitudinal distribution of OH. Since latitude and temperature are strongly correlated, a latitudinal redistribution of OH that conserves global total oxidation often implies a change in $[\text{OH}]_{\text{GM}}$. In that case, we consider the conservation of global total oxidation the relevant quantity, rather than a change in $[\text{OH}]_{\text{GM}}$ (similar to the recommendation in Lawrence et al. (2001)). We calculate the variations in oxidation relative to the prior, so that for example interannual variations of temperature will not affect these variations in oxidation, since temperature variations remain the same between prior and posterior simulations. Where relevant, we make note of this distinction. We also present results for the latitudinal adjustments in OH concentrations.

3.2.2 Observations

3.2.2.1 Surface observations

We used MCF observations from the surface network of the National Oceanic and Atmospheric Administration (NOAA) Global Monitoring Laboratory (GML) as the only

observational constraints in the inversion. The network consists of a core set of 7 surface sites that have monitored MCF since 1992, and additional sites have been added since: data from a total of 12 sites are available since 1998. The sites we used, including site abbreviations, are described in Table S2. At each site, paired flask samples are collected at weekly to monthly frequency, following a sampling protocol that typically favors sampling under meteorological conditions that correspond to clean background air. Flask samples are then collected and measured on one central measurement system against the NOAA calibration scale for MCF. The measurement uncertainties we used are those reported by NOAA, which are based on the difference between the mole fractions measured for each flask in a flask pair. Up to 2018, short-term measurement repeatability remained consistently around 0.5% of the measured mole fraction. On top of the measurement error, we also included a model representativeness error for each observation. This error is calculated as an absolute average over the mole fraction gradients between the model grid cell that contains an observation and horizontally and vertically adjacent grid cells (Bergamaschi et al., 2005). Typically, the model error was up to 1% at Northern midlatitude sites (e.g. LEF), and as low as 0.1% at e.g. SPO or ALT. The addition of these two error sources, with no correlations in between, constitute the **R** matrix (see Section 3.2.1).

3.2.2.2 Aircraft campaigns

For validation of the inversion results, we used two sets of aircraft campaigns: the HIAPER Pole-to-Pole Observations (HIPPO, 2009-2012, Wofsy et al. (2011)) and the Atmospheric Tomography Missions (ATom, 2016-2018, Wofsy et al. (2018)). The five HIPPO campaigns were flown from 85°N to 67°S over the Pacific, with vertical profiles being sampled up to 8500 meters at approximately 2.2° intervals. The horizontal extent of the ATom campaigns was wider, with global coverage over the Pacific, the Atlantic and the Arctic oceans.

While a number of different measurement systems provided measurements of MCF during HIPPO, only results from the NOAA Whole Air Sampler during H-1, -3, -4, and -5 were included in this analysis because they are most closely tied to NOAA’s surface network results (i.e., in methodology, precision, and calibration consistency). ATom results from deployments A-2 and A-3 were also included in this work, as those flask samples were analyzed on the same NOAA instrument as the surface network flasks. We exclude from our analysis a subset of results from HIPPO-1, as well as all of HIPPO-2 and the ATom-1 missions, since samples from these deployments suffered from a deployment-specific measurement interference (a portion of H-1 and all of H-2) or were analyzed on a different instrument in NOAA that exhibited poorer precision (A-1).




Figure 3.1: Monthly variations in global mean $k(T) \cdot [\text{OH}]$, derived in the three different inversion set-ups. Variations in $k(T) \cdot [\text{OH}]$ are calculated relative to the prior simulation, thus removing the influence of interannual variations in temperature. Gray bars indicate a spin-up and spin-down period of 1.5 years.

3.3 Results

3.3.1 Variations in the atmospheric oxidative capacity and in MCF emissions

In Figure 3.1, we show the monthly anomalies in global oxidation (different from $[\text{OH}]_{\text{GM}}$: see Section 3.2.1.2), as derived in the three inversion set-ups. We have shown the entire twenty-year inversion period, which will include a spin-up and spin-down period of 1-2 years (indicated by the gray bars). For example, even though our initial MCF mole fraction fields are realistic, the strong positive oxidation anomaly in 1998 might be linked to errors in the initial field. The corresponding variations in the tropospheric lifetimes of MCF and of CH_4 are described in Supplement 3.S1.

Interannual variations in global oxidation are typically small ($\sim 2\%$). In this, there is consistency between the different inversion set-ups. Variability in global oxidation derived in the POP and the TM5OH inversions are strongly positively correlated ($R^2 = 0.92$), while the correlation with the REF inversion is weaker ($R^2 = 0.51$). It is counter intuitive that the REF inversion deviates from the other two, as it shares its emission distribution with the TM5OH inversion, and its OH distribution with the POP inversion. Contrastingly, the POP and TM5OH inversion share neither. This indicates convergence problems, rather than a significant influence from prior distributions. For reasons outlined in Section 3.3.5, we attribute differences in variations mostly to differences in the degree




Figure 3.2: Hovmöller diagram of the spatio-temporal adjustments made to $[\text{OH}]$ in the REF inversion set-up. The prior $[\text{OH}]$ distribution, with respect to which changes are shown, is based on Spivakovsky et al. (2000), scaled down by a factor 0.92. Gray bars indicate a spin-up and spin-down period of 1.5 years.

of convergence. Because the REF inversion resulted in the best match with observations (see Section 3.3.2), and because the REF inversion is most consistent with a set of ten-year inversions (see Supplement 3.S5), we consider it to be the solution that converged best.

As outlined in Section 3.2, we optimized OH concentrations in 45 latitude bands of 4° each. Figure 3.2 shows the resulting adjustments per latitude band through time, for the REF inversion. Clearly, adjustments to zonal mean OH concentrations can be much larger (up to 30%) than adjustments to annual $[\text{OH}]_{\text{GM}}$ (up to 5%). Moreover, there is a strong systematic tendency to increase tropical OH concentrations, and decrease extra-tropical OH concentrations, especially in the Southern Hemisphere. This tendency was observed in each of the three inversion set-ups, i.e. also when a different OH distribution was used. We further investigate this tendency in Sections 3.3.2 and 3.3.4. The prior and posterior latitudinal OH distributions are compared to a range of literature estimates in Figure S8.

The MCF emissions that result from the three inversions are shown in Figure 3.3. In general, emissions are increased relative to the prior, especially in later years. The REF inversion shows the largest adjustments and pronounced interannual variations, with even an increase in some years superimposed on the downward trend. Firstly, we note that the small emission totals in later years of around 2 Gg/yr, with interannual variations of 0.2 Gg/yr, would be hard to exclude based on prior knowledge of emissions. Secondly, the emission variations are largely uncorrelated with derived variations in atmospheric oxida-




Figure 3.3: Timeseries of annual global total MCF emissions, derived in the three different inversion set-ups (solid lines) and the prior emissions (black dashed). To visualize variations in the low emissions in later years, we provide emissions in the second decade on a different y-scale (right axis).

tion ($r = 0.13$ in the REF inversion), which indicates that the MCF observations, to some degree, allow for independently derived $[\text{OH}]_{\text{GM}}$ and MCF emission variations. However, the contribution of the emissions to the MCF budget increases over the inversion period. Emissions are equivalent to 12% of the loss to OH from 2000 to 2006, but afterwards this ratio increases to up to 22% in 2017-18. Notably, we find that the emissions stabilize around 2 Gg/year in the final 5 years. Admittedly, this is partly driven by the assumed prior emissions (based on global mean MCF mole fractions), to which the inversion makes only small adjustments in later years. A growing relative contribution of emissions in the MCF budget would make derived OH variations more uncertain, especially if emission distributions had changed substantially over time. Therefore, even if high-quality measurements remain possible, residual MCF emissions could make it increasingly difficult to derive interannual changes in $[\text{OH}]_{\text{GM}}$ from MCF. Additionally, the uncertain decay rate of MCF emissions since 2001 (the last year with reported emissions) complicates the derivation of a trend in $[\text{OH}]_{\text{GM}}$.

3.3.2 Comparison to observations

3.3.2.1 NOAA surface network

Simulated global mean MCF mole fractions and the interhemispheric gradients in the posterior simulation match well with observed gradients (Figure 3.4). This illustrates the skill of the inverse framework to adjust emissions and $[\text{OH}]$ such that large-scale gradients

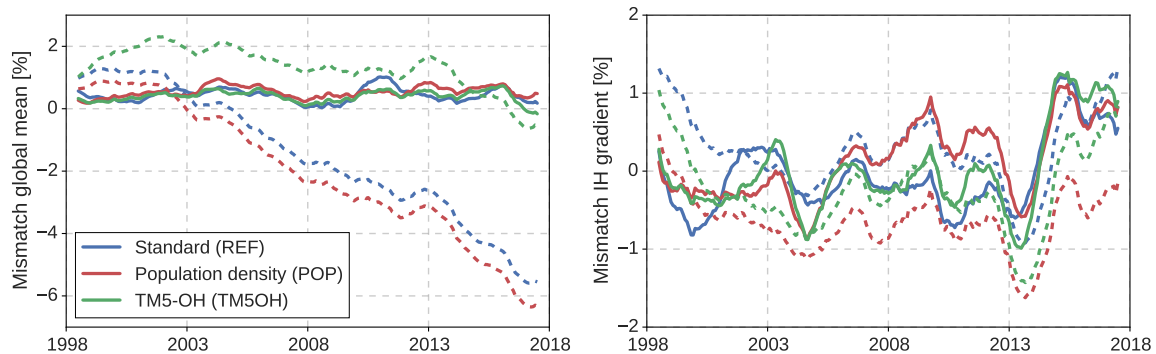


Figure 3.4: Fractional mismatch between simulated and observed MCF mole fractions in the global mean (left) and the interhemispheric gradient (right). Mismatches are given relative to global mean MCF and averaged using a twelve-month running mean. Dashed lines indicate prior mismatches, while solid lines show posterior mismatches. Global and hemispheric mean mole fractions were calculated from the NOAA surface network following the methods outlined in S1 of Naus et al. (2019).

of MCF are very well ($< 1\%$) reproduced.

However, the inversion cannot reproduce some of the observed gradients between stations, especially gradients within hemispheres. In Figure 3.5, gradients between three pairs of NOAA surface sites are shown. Firstly, the interhemispheric gradient between ALT and CGO is captured well in all inversions. This interhemispheric gradient is strongly affected by emissions and therefore the inversion framework can adjust it with relative ease.

In contrast, intrahemispheric gradients are less well captured. In Figure 3.5 it can be seen that both the gradient within the Northern Hemisphere, between ALT and MLO, and the gradient within the Southern Hemisphere, between SMO and CGO, are underestimated. More precisely, MCF mole fractions simulated at tropical sites are systematically too high ($1\text{-}2\sigma$, with σ the total error), with a smaller, opposite bias at high-latitude sites ($0.5\text{-}1\sigma$). The large latitudinal adjustments of OH concentrations (see Figure 3.2) are an attempt by the inversion to reduce this bias, but the adjustments only reduce the bias partly (Figure 3.5). As an example, the SMO-CGO gradient is, in the REF inversion, increased from -1.5% to -2% , i.e. an increase of 30% , which corresponds well with a 30% latitudinal adjustment in $[\text{OH}]$. However, the observed gradient is larger still at -3% . We investigate this residual bias in more detail in Section 3.3.4.

We have quantified the skill of a simulation that uses optimized OH and MCF emission distributions from the REF inversion to reproduce observed MCF mole fraction in a root-mean-squared error (RMSE) per site, averaged over the 1998-2018 period (top panel in Figure 3.6). We distinguish between the RMSE of individual observations (as used in the optimization, in red), the RMSE of the monthly mean values (cyan), and the RMSE of twelve-month running averages (green). This helps to disentangle the contribution of

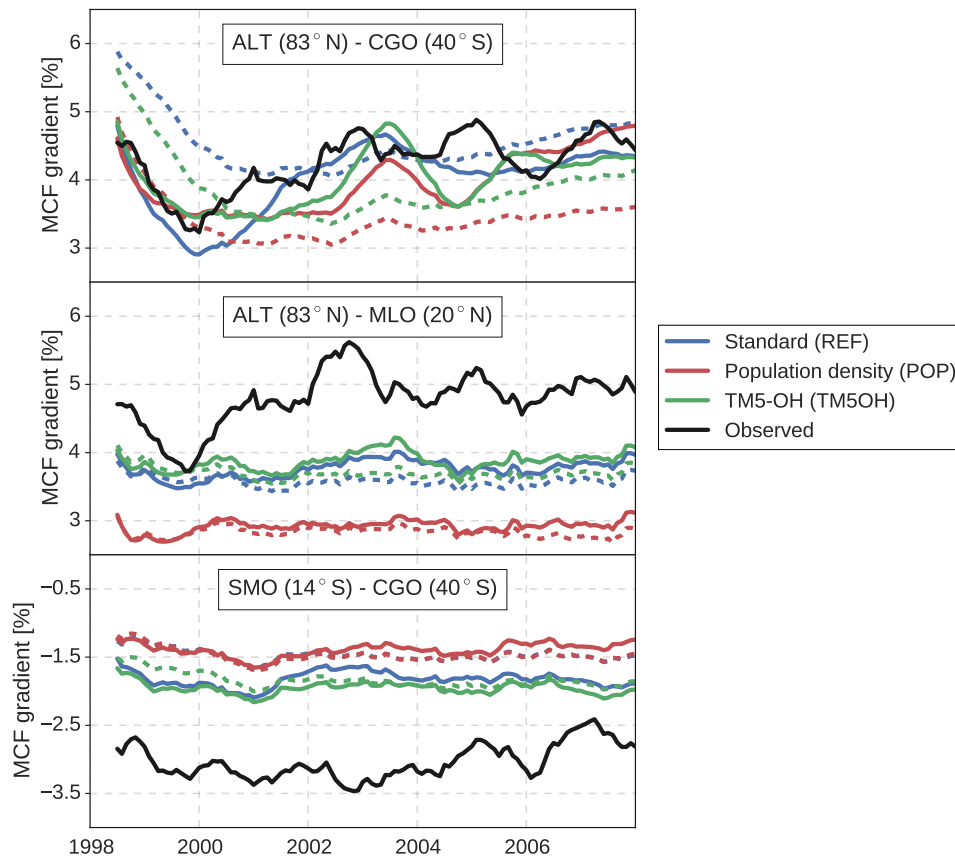


Figure 3.5: Timeseries of observed and modelled gradients between three sets of NOAA surface sites. Gradients are given relative to global mean MCF and averaged using a twelve-month running mean. Dashed lines indicate prior mismatches, while solid lines show posterior mismatches. Site latitudes are also shown and further details on assimilated surface sites are given in Table S2.

short-term versus long-term variations to the model-measurement mismatch. Also shown is the pre-defined total observational error (gray), which we used in all inversions. The total error consists of a model error based on modelled spatial gradients and a measurement error based on short-term measurement repeatability (see Section 3.2.2.1). Additionally, in the lower two panels, we visualize two examples of these mismatches for MHD, a somewhat polluted site with no residual bias, and for SMO, a relatively clean site with residual bias.

At all sites, the posterior RMSE of individual observations exceeds the total error. However, at most sites, the RMSE comes more in line with the total error for monthly means, and especially for twelve-month running averages. This implies that, largely, the RMSE is related to short-term variations. Our inverse system, which employs relatively smooth and stiff OH and emission distributions, has limited capability to fit short-term variations of MCF. Short-term variations in MCF are likely related to errors in the emission distri-

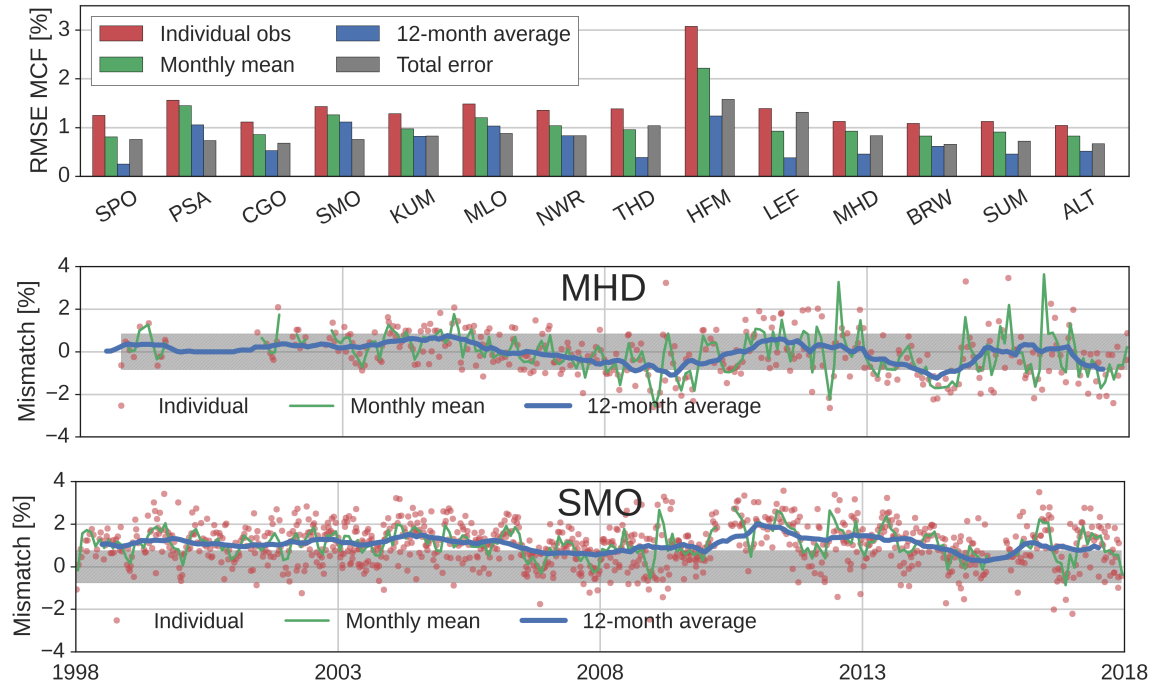


Figure 3.6: Top panel: Posterior root-mean-squared errors (RMSE) per surface site, for the REF inversion. RMSE are shown for individual observations, for monthly mean and for twelve-month running mean MCF mole fractions. Additionally, the site-dependent average total error that was used in the inversions, combined from a model and measurement error, is shown. RMSE are averages over the 1998-2018 inversion period and are given relative to the measured MCF mole fraction. Site abbreviations are explained in Table S2. Bottom two panels: Timeseries of the posterior measurement-model mismatch at MHD (top) and SMO (bottom) from individual observations, from monthly means and from twelve-month running averages, with the average total error shaded in gray.

bution and in small-scale transport, since OH has an integrated, slow effect on MCF mole fractions. Therefore, we do not expect these residuals to affect our OH estimate significantly. Importantly, we point out that our initial error estimate might have been overly conservative (mostly $< 1\%$, see Figure 3.6), which is supported by error estimates used in previous MCF inversions (1-2% for individual observations (Bousquet et al., 2005); 5% for monthly, hemispheric averages (Turner et al., 2017; Rigby et al., 2017)). We conclude that, as long as we capture long-term variations of MCF at each site, the unresolved residuals on a sample-to-sample basis are unlikely to impact our conclusions. For the same reason, we would recommend a follow-up inversion to optimize monthly means per site, rather than individual observations. At some sites (notably SMO, MLO and PSA), we have identified systematic biases which are of more concern, and these result in relatively large RMSE even in twelve-month running averages. We further discuss systematic offsets in Section 3.3.4.




Figure 3.7: The mismatch of simulated minus observed MCF mole fractions from the HIPPO and ATom aircraft campaigns. Differences are shown relative to the observed mole fraction. We aggregated results from the HIPPO 1, 3- 5 and ATom 2 and 3 campaigns. Left: A plot of the model-measurement mismatches for each of the aircraft samples. Modelled mole fractions were sampled from the posterior REF simulation. Right: A probability-density distribution of the prior (red) and posterior (blue) mismatches in the REF inversion. Numbers denote the median mismatch and its standard deviation. Mismatches that fall outside the range on the x-axis were added to the outermost bars.

3.3.2.2 HIPPO and ATom aircraft campaigns

As the inversions were driven by observations from the NOAA surface network, we used the HIPPO and the ATom aircraft campaigns as independent data sources for validation. The main added value of the aircraft over surface observations is that the former provide snapshots of vertical gradients. When observed MCF mole fractions from all campaigns are compared to model-sampled mole fractions (Figure 3.7), a few features emerge.

Firstly, we find that the optimized REF simulation overestimates MCF mole fractions in the lower stratosphere at high latitudes ($> 50^\circ$) in both hemispheres (left panel in Figure 3.7). This bias points to limited ability of TM5 to capture vertical gradients in the downward branch of the Brewer-Dobson circulation. A similar bias was identified in TM5 simulations of CH_4 , which has a stratospheric sink similar to MCF (Houweling et al., 2014).

In addition to the large overestimation at high altitudes, we also find a weak but significant positive correlation between altitude and model-measurement mismatch ($R^2=0.12$; $p<0.001$). Close to the surface, the average model-measurement mismatch is close to 0%, which increases to 1.5% at 10km. In other words, TM5 increasingly overestimates MCF mole fractions at higher altitudes, resulting in an underestimate of vertical gradients in TM5.

Positively, we find that the inversion improves the agreement between simulations and aircraft observations (right panel in Figure 3.7). While simulated mole fractions in the prior simulation are too low, in the posterior simulation MCF mole fractions compare well with observed mole fractions. The right panel in Figure 3.7 does show a distribution skewed towards model overestimation, which is linked to the overestimation in the lower

stratosphere discussed above. While the average model-measurement difference decreases in the inversion, the standard deviation of the differences remains around 5%. This shows that vertically uniform adjustments in OH concentrations and in MCF emissions do not significantly modify simulated vertical gradients.

In conclusion, assimilation of surface observations improves agreement of our simulations with aircraft observations. However, modelled vertical gradients of MCF remain slightly smaller than those observed, although the maximum model bias is 1.5%, which is small compared to the 5% random error. Moreover, since these biases are consistent between the aircraft campaigns, we deem the impact of the biases on derived interannual and multi-annual variability of $[\text{OH}]_{\text{GM}}$ or MCF emissions small. Estimates of the total atmospheric oxidizing capacity are more likely to be affected by a systematic underestimate of vertical MCF gradients. This underestimate could be driven by an underestimate of the vertical OH gradients or by too-fast vertical mixing in TM5. While TM5 typically compares well to other transport models in terms of large-scale transport features, for example in the Age of Air experiment (Krol et al., 2018), this comparison does highlight the crucial role of aircraft campaigns in helping to identify remaining transport model biases.

3.3.3 Physical drivers of variations in atmospheric oxidation

The El Niño Southern Oscillation (ENSO), a dominant mode of natural atmospheric variability, has previously been suggested to influence interannual variations in $[\text{OH}]_{\text{GM}}$ (Prinn et al., 2001; Turner et al., 2018). ENSO affects many processes that are linked to OH, such as temperature, atmospheric moisture, lightning, wildfires and atmospheric transport.

Figure 3.8 shows variations in ENSO, quantified in the Multivariate ENSO Index (MEI), together with the annual, global mean $k(T) \cdot [\text{OH}]$ variations derived in the REF inversion. We find a negative correlation of -0.47 ($p = 0.05$) between the MEI and the derived $k(T) \cdot [\text{OH}]$ anomalies, if we exclude one year of spin-up and spin-down. The largest exception to the negative correlation is the positive $k(T) \cdot [\text{OH}]$ anomaly in 2012 that is not explained by a coincident variation in the MEI, which indicates that there are other controlling processes that can drive $[\text{OH}]_{\text{GM}}$ variations.

That variability in $[\text{OH}]_{\text{GM}}$ correlates with a dominant driver of atmospheric variability seems logical. However, attribution of the negative correlation to specific processes is difficult, given the large number of processes that are affected by ENSO and that in turn could affect OH. Nonetheless, we can hypothesize. For example, El Niño years (high MEI) are associated with more wildfires, resulting in higher CO emissions which could suppress OH concentrations (Zhao et al., 2020; Nguyen et al., 2020). La Niña years (low MEI) are associated with increased convection over the Pacific, increased lightning NO_x production and by extent increased OH recycling (Turner et al., 2018). The ENSO cycle additionally influences tropospheric, tropical O_3 abundance, although the sign of this effect can vary




Figure 3.8: Variations in the annual mean MEI (right axis) and variations in global, annual mean $k(T) \cdot [\text{OH}]$ (left axis), as derived in the REF inversion. Gray bars indicate a spin-up and spin-down period of 1.5 years.

regionally (Oman et al., 2011).

The correlation between MEI and the $[\text{OH}]_{\text{GM}}$ variations derived in a two-box model inversion (Naus et al., 2019) is -0.01 if we do not tune the box model with a 3D transport model, and -0.34 if we do. This increase after accounting for some of the box model biases, and the even higher correlation if we move the inversion completely to a 3D transport model, shows that the correlation becomes apparent only when realistic transport of a 3D transport model is included. For example, transport variations related to ENSO have been shown to strongly affect interhemispheric differences of CO_2 (Francey & Frederiksen, 2016) and of CH_4 (Pandey et al., 2019b), as well as mole fractions of MCF at SMO (Prinn et al., 1992).

3.3.4 Explaining the underestimated intrahemispheric gradients of MCF

The systematic underestimation of intrahemispheric gradients in both hemispheres deserves further elaboration, since the inversion framework has difficulties to adjust OH concentrations and emissions in such a way that intrahemispheric gradients are reproduced. MCF mole fractions are overestimated in the tropics and underestimated at high latitudes (Figure 3.5). To resolve these biases, the inversion introduces large adjustments in the latitudinal distribution of OH (up to 30%, see Figure 3.2). A set of inversions that only covered the 1998-2008 period were run to higher convergence (see Supplement 3.S5). These inversions show that more extreme adjustments in the latitudinal OH distribution

(up to 60%) better reproduce intrahemispheric gradients. However, intrahemispheric gradients were still not quite captured, and substantially higher MCF emissions were required to reproduce the gradients. The amplitude of these adjustments seem physically unlikely (e.g. see Figure S8) and in none of the inversions the biases were fully resolved. Therefore, we consider here alternative, or complimentary explanations for the biases, that were not explored in the inverse framework. We emphasize that the biases are quite constant over the twenty-year period, and while especially pronounced in the Southern Hemisphere, also present in the Northern Hemisphere. These aspects make anthropogenic emissions as a sole explanation unlikely.

One explanation for underestimated intrahemispheric gradients is that intrahemispheric mixing is too fast in TM5. Even after extensive participation of TM5 in intermodel comparisons (e.g., Krol et al., 2018), transport biases might persist. We have tested the effect of both horizontal (down to $1^\circ \times 1^\circ$) and vertical (up to 60 layers) model resolution on the simulated mole fractions (see Supplement 3.S2). While we did find some sensitivity to resolution, the intrahemispheric biases persisted. Additionally, when we simulated SF_6 , a chemically inert tracer commonly used for diagnosing large-scale transport (e.g., Patra et al., 1997; Denning et al., 1999), we did not find any comparable bias. Finally, we also simulated HFC-152a, a gas that is removed by OH against a shorter lifetime than MCF (~ 1.5 years (Ko et al., 2013b)). For HFC-152a, we also found none of these biases. The results for SF_6 and HFC-152a are discussed more extensively in Supplement 3.S3. We conclude that the biases seem particular to MCF and are therefore unlikely to be related to either transport or OH.

This leaves stratospheric photolysis, ocean-atmosphere exchange, another unknown MCF budget term, or a site-dependent measurement bias as potential explanations. Of these four, we consider the oceanic flux the likeliest. Wennberg et al. (2004) suggested that low MCF hydrolysis rates in cold, high-latitude oceans could result in a build-up of oceanic MCF. This reservoir would be released again when atmospheric mole fractions of MCF dropped, following production reductions demanded by the Montreal Protocol. They calculated and forecast the associated change in the ocean flux in five latitude bands for the period of 1961 to 2020 in a global ocean circulation model. We implemented their ocean flux in a forward simulation of MCF, and found that the adjusted ocean flux reduced the intrahemispheric bias by around 50% (see Supplement 3.S4). This reduction is larger than the adjustments brought about by a 30% adjustment in tropical OH concentrations, and provides an arguably more likely scenario.

We also sampled aircraft observations in the simulation with an adjusted ocean flux. Release of MCF from high-latitude oceans would enhance vertical gradients, which combines well with the observation that TM5 underestimates aircraft-observed vertical gradients of MCF (see Section 3.3.2.2). However, we find that vertical gradients of MCF are only very slightly modified by the change in ocean flux, so that the signal of oceanic release of

MCF in aircraft observations is too small to identify using the tools available to us.

Some new issues arise when we include the new ocean flux. Global mean MCF is overestimated, as the oceanic flux is much reduced, or even reversed. Additionally, the MCF gradient between ALT and CGO is captured less well. However, these are offsets that our inverse framework can resolve, as was found in Section 3.3.2.1. Therefore, we propose that variations in the oceanic flux of MCF very likely provide part of the solution to the residual intrahemispheric biases in our inversions. Further investigation of this issue, however, is beyond the scope of this study.

We note the dichotomy in the evidence for oceanic release of MCF. Positively, it emphasizes the resolving power of a sparse network of internally-consistent, low-frequency observations: we unexpectedly found a likely constraint on the oceanic uptake of MCF, in addition to the constraints on OH and emissions. Negatively, the existence of an ocean-atmosphere exchange of MCF, as opposed to a first-order sink, further increases uncertainty on any MCF-derived constraints on OH. Under the assumption that the exchange from Wennberg et al. (2004) is accurate, the impact of a different oceanic sink on derived interannual variations of OH has been shown to be small (Prinn et al., 2005). However, the ocean-atmosphere flux is rather uncertain, since experiments that measured hydrolysis rates of MCF are scarce, especially at low ($< 20^{\circ}\text{C}$) temperatures (see Supplement 3.S4), and the hypothesis of an ocean source of MCF has not been substantiated with field measurements. Therefore, we consider that the similar spatial signature of the OH sink and the ocean uptake (both high in the tropics) will make it difficult to resolve between the two using only the surface network. Still, given the continuous decline of MCF abundance over our inversion period, no strong interannual variations in the ocean flux are expected. Therefore, variations in OH concentrations are still the most likely driver of interannual variations in MCF mole fractions, which we consider therefore a robust feature of our inversion. Contrastingly, uncertainty in the oceanic flux does complicate the derivation of a trend in $[\text{OH}]_{\text{GM}}$ from MCF observations, as well as spatial adjustments in the OH distribution. In other words, while we deem the adjustments to the OH distribution of 30% unlikely, we cannot exclude some underestimation of tropical OH concentrations as a partial explanation for the intrahemispheric biases.

3.3.5 The problem of convergence

Conventionally, the convergence criterion of a 4DVAR inversion is defined as a reduction in the cost function gradient. Ideally, the resulting converged state reproduces observations within the pre-defined total error. Due to pragmatic constraints, this procedure did not provide us with the desired result. Firstly, computational expense (~ 3 iterations in two days) limited the number of iterations. Secondly, while further convergence was possible, the derived adjustments to OH concentrations became less realistic physically. Two underlying problems have been pointed out: a difficulty in defining observational

and model uncertainties and difficult-to-resolve intrahemispheric biases. However, while each of the three inversions faced these two problems, they did not reach the same level of convergence. For example, an emission distribution based on population density is more diffuse, and therefore contains more degrees of freedom than the TransCom emission distribution, but the TransCom distribution resulted in better convergence.

The importance of considering convergence when comparing the inverse results is illustrated by a set of inversions with the same set-up described in Section 3.2.1.2, but covering only 1998-2008. These ten-year inversions required less time per iteration and converged more consistently and in fewer iterations. The three ten-year inversions result in much more consistent interannual variations and spatial adjustments in OH concentrations, as described in more detail in Supplement 3.S5. In turn, the interannual $[\text{OH}]_{\text{GM}}$ variations and the spatial OH adjustments derived in the ten-year inversions were most similar to those derived in the twenty-year REF inversion. This suggests that the converged solutions of the three inversions are likely to be similar, but that for unclear reasons the twenty-year REF inversion managed to converge furthest, while the other two either required more iterations, or reached a false minimum.

In the end, we consider that the twenty-year REF inversion reproduces NOAA surface observations to a satisfying degree (except for the intrahemispheric biases) and so we present the twenty-year REF inversion as our reference dataset. We do note that the amplitude of OH adjustments in the better-converged ten-year REF inversion is larger than that in the twenty-year REF inversion. For instance, interannual variability in $k(T) \cdot [\text{OH}]$ averaged over 1999-2008 (i.e. excepting the spin-up period) was 2.9% and 1.9% in the ten- and twenty-year REF inversions, respectively. Therefore, we note that the fully converged solution of our inverse set-up will likely have higher interannual variability than the twenty-year inversion: whether this higher variability would be real or a result of overfitting of the available data (e.g. due to an underestimated model error) remains a topic of discussion. In both scenarios, however, the variability is typically smaller than 3%.

3.4 Further discussion

An important question concerns the robustness of the derived $[\text{OH}]_{\text{GM}}$ variations. Our inversion procedure did not provide formal posterior uncertainties, but we did test the robustness of derived variations in sensitivity tests. Firstly, while we observed differences between the three inversion set-ups, we generally identified similar tendencies in each. Interannual variations in global oxidation and latitudinal adjustments in OH concentrations seem only marginally affected by the prior OH and emission distributions, and most differences were driven by different degrees of convergence (see Section 3.3.5). Secondly, in a previous test inversion where we only optimized one global monthly scaling factor

for OH, interannual variations turned out similar. Thirdly, in a forward simulation, we applied the interannual variations in oxidation to the prior OH distribution, i.e. without adjustments in the spatial distribution. This resulted in a significantly improved match with observations compared to the interannual constant prior, which means that the temporal variations in OH are important for reproducing the observations. These results indicate that the solution we have derived is robust and consistent with observed gradients in MCF.

Key to our inversions and the derived OH variations is the use of MCF observations from the NOAA-GML surface network. MCF has additionally been monitored by the Advanced Global Atmospheric Gases Experiment (AGAGE) surface network and its predecessors since 1978 (Prinn et al., 1983, 2018). Both networks maintain their own calibration scales for MCF, and previous comparisons have revealed site-dependent scale drifts and a possible phase difference between the two networks that exceed measurement repeatability (Holmes et al., 2013; Rigby et al., 2013). Due to these differences, and because of different site locations and measurement approaches, an inversion incorporating AGAGE instead of NOAA observations will result in different OH variations. It is difficult to assess the impact without performing additional inversions, but in a three-box inverse system differences between interannual OH variations derived from each network were shown to be insignificant in the presence of other uncertainties in the MCF budget (Rigby et al., 2017).

The only OH inversion studies that cover a similar period as this work were done in box models. As expected, the box model inversion that gave interannual $[\text{OH}]_{\text{GM}}$ variations most similar to those presented here is from our previous work (Naus et al., 2019), where we partly accounted for transport-related biases using TM5 ($R^2 = 0.49$; significant at $p=0.01$). This can be compared to the two-box inversion where we did not account for these biases ($R^2 = 0.21$; not significant at $p=0.01$). We find similarly low correlations with the OH variations derived in Turner et al. (2017) and in Rigby et al. (2017). In Section 3.3.3, we already noted that the two-box inversion where we accounted for transport biases also correlated better with the ENSO cycle, while the OH variations derived in this work correlate best with ENSO. This result, combined with the explicit simulation of transport and the more extensive use of available data (e.g. site-to-site gradients) give us confidence that the $[\text{OH}]_{\text{GM}}$ variations derived in this work are physically more realistic than those derived in two-box models. As such, we recommend the use of the $[\text{OH}]_{\text{GM}}$ variations derived in our twenty-year REF inversion as an alternative to interannually repeating OH fields, or to OH fields derived in full-chemistry simulations, in for example CH_4 inversions. We do note that the $[\text{OH}]_{\text{GM}}$ variations derived in this work are statistically consistent with these previous studies, as uncertainties remain large relative to the amplitude of interannual variations in $[\text{OH}]_{\text{GM}}$.

It remains difficult to derive a multi-annual trend in $[\text{OH}]_{\text{GM}}$ from an inversion of MCF.

To facilitate convergence of our inversion, we already pre-fitted MCF emissions to the global MCF trend. In principle spatially distinct patterns of emissions and OH still allow for an $[\text{OH}]_{\text{GM}}$ trend in the posterior solution, but the signal of such an $[\text{OH}]_{\text{GM}}$ trend in MCF observations would be small. Moreover, we have found evidence for a possible role of MCF release from oceans at high latitudes, as well as a growing relative contribution of MCF emissions in the MCF budget. While we can leverage $[\text{OH}]_{\text{GM}}$ variations as the most likely source of interannual variations in the MCF growth rate, variations in these other budget terms become more important on decadal timescales. Therefore, on the basis of the inversions presented in this work, we conclude that a trend in $[\text{OH}]_{\text{GM}}$ cannot be excluded, though we have found no evidence for such a trend in the MCF budget.

While MCF has been the most widely used tracer for observation-derived constraints on $[\text{OH}]_{\text{GM}}$ over the past decades, promising alternatives have been explored. ^{14}CO is produced from cosmic radiation in the stratosphere and subsequently transported to tropospheric surface sites (Quay et al., 2000). Observed variations of ^{14}CO could provide additional constraints on $[\text{OH}]_{\text{GM}}$ and the OH distribution, with a higher sensitivity to the tropics (Krol et al., 2008). Currently, long-term monitoring of ^{14}CO is limited to one site (Baring Head, New Zealand), but efforts are being made to expand the network (Petrenko et al., 2019). Recently, it was proposed that high-resolution satellite data of CH_4 can be used to constrain both CH_4 sources and sinks, which would be possible with a combination of satellite products with different vertical sensitivities (Zhang et al., 2018). Future research will reveal how feasible this approach is, but it will not provide information retroactively, for past decades: here, surface data still provides the strongest constraints. All potential tracers for OH are in principle complementary and the implementation of multiple tracers in a comprehensive inversion framework could offer a more robust method for deriving OH variations (Liang et al., 2017). Already, we have shown that tracers such as SF_6 and HFC-152a can help in understanding the role of OH in the MCF budget. However, also in light of the results presented in this study, we deem it likely that each tracer will bring its own complications, which should be considered with care. This makes a multi-tracer approach a promising, but complicated and involving exercise.

Observational constraints have also been implemented in full-chemistry models. Assimilation of satellite data of those gases that dominate the OH budget, such as CO and O_3 , is likely to bring full-chemistry simulations closer to reality and makes them less reliant on uncertain emission inventories (e.g., Miyazaki et al., 2012; Flemming et al., 2017). For example, assimilation of CO satellite data in a full-chemistry simulation resulted in an enhanced positive trend in $[\text{OH}]_{\text{GM}}$, compared to an unconstrained simulation, which was driven predominantly by a decreasing atmospheric CO burden (Gaubert et al., 2017). An approach complimentary to budget-based inverse modelling and full-chemistry simulations has been suggested in formaldehyde (HCHO) (Wolfe et al., 2019). HCHO is well-monitored, both in-situ and by satellites, and the production of HCHO is tightly

linked to OH abundance, especially in the remote troposphere. This approach provides high-resolution information on the OH distribution. All these methods benefit from the availability of each of the other methods as independent validation. Therefore, we consider that constraining the atmospheric oxidative capacity, a traditionally difficult problem at the heart of atmospheric chemistry, potentially has a solution in the convergence of a growing number of independent lines of evidence that need combined exploration.

3.5 Conclusions

In this study, we present a set of global 3D model inversions of MCF that cover the 1998–2018 period, performed within the TM5-4DVAR inverse framework. In the inversions, MCF emissions and OH concentrations were optimized simultaneously, with the objective to derive interannual variations in $[\text{OH}]_{\text{GM}}$ and in the OH distribution. The inverse system was typically able to reproduce observed MCF mole fractions within 1%. The main conclusions of this work can be summarized as follows:

1. Small interannual variability (2–3%) of $[\text{OH}]_{\text{GM}}$ without a significant trend allows for a relatively good match with MCF observations over the 1998–2018 period.
2. We have identified an underestimate of intrahemispheric gradients of MCF in the TM5 simulations that is not easily reduced using only adjustments in OH concentrations and anthropogenic MCF emissions. Instead, we propose that an explanation involving oceanic release of MCF at high latitudes can help to at least partly correct for these biases.
3. Most derived metrics, such as interannual $[\text{OH}]_{\text{GM}}$ variability and the absence of a trend in $[\text{OH}]_{\text{GM}}$, are relatively robust with respect to the prior $[\text{OH}]$ and emission distribution. However, we put the most confidence in interannual variations of $[\text{OH}]_{\text{GM}}$, as trends and spatial adjustments of $[\text{OH}]$ will be most strongly affected by remaining uncertainties in the system, such as the exact magnitudes of the oceanic flux and of anthropogenic emissions.
4. We find a significant and robust negative correlation between the derived variations in $[\text{OH}]_{\text{GM}}$ and the ENSO cycle. Such a correlation was not found in box model studies that did not fully account for transport variations.
5. We suggest inclusion of our optimized $[\text{OH}]_{\text{GM}}$ variations in global CH_4 inversions, since the timeseries of optimized $[\text{OH}]_{\text{GM}}$ better reproduces MCF observations than annually repeating $[\text{OH}]_{\text{GM}}$. However, we emphasize that the solution remains very uncertain.

Supporting Information to Chapter 3

3.S1 Tropospheric lifetimes of MCF and CH₄




Figure 3.S1: Prior (dashed lines) and posterior (solid lines) lifetimes of tropospheric MCF and CH₄ with respect to oxidation by OH. Posterior lifetimes were derived from the 20-year REF inversion. Lifetimes were calculated assuming an interannually constant distribution of MCF and CH₄ that did include a seasonal cycle. The tracer distributions were taken from the forward TM5 simulations performed in Naus et al. (2019), which use input fields described in the Transcom-CH4 protocol Patra et al. (2011). The tropopause was defined, analogous to Naus et al. (2019), as the lowest altitude where the vertical temperature gradient is smaller than 2 K km⁻¹, with a minimum and maximum geopotential height of 9 and 18 km, respectively.

In the main text we have discussed mainly the anomalies we derived for $k(T) \cdot [\text{OH}]$. Here we translate those anomalies to tropospheric lifetimes of MCF and CH₄, with respect to oxidation by OH (Figure 3.S1). Variations in the lifetimes of MCF and CH₄ due to variations in $[\text{OH}]_{\text{GM}}$ are similar (see Figure 3.S1). This shows that slight differences in their atmospheric distributions and in the temperature-dependence of their reaction rates with OH do not significantly affect variability in their tropospheric lifetimes.

The posterior lifetimes we derive, with respect to oxidation by OH, are 5.33 ± 0.11 and 8.98 ± 0.19 years for tropospheric MCF and CH₄, respectively. These posterior lifetimes do not differ significantly from the prior lifetimes (MCF: 5.29 ± 0.03 ; CH₄: 8.93 ± 0.05). Interannual variations in the prior lifetimes are driven by interannual variations in temperature and transport, because the prior OH and tracer distributions were interannually repeating. In Krol & Lelieveld (2003), comparable lifetimes of 5.47 and 5.38 years were derived for tropospheric MCF using two different models. The MCF lifetime of $6.1^{+0.5}_{-0.4}$ years, reported in Prinn et al. (2005), represents the global lifetime of MCF with respect to oxidation by tropospheric OH. The corresponding posterior estimate from our REF inversion would be 5.76 ± 0.11 years, i.e. in statistical agreement with the estimate from Prinn et al. (2005).




Figure 3.S2: Mismatch between simulated and observed MCF mole fractions in the global mean (left) and the interhemispheric gradient(right), when TM5 is operated at different vertical and horizontal resolutions. The mismatches are given relative to global mean observed MCF and averaged using a twelve-month running mean. Monthly mean hemispheric and global averages were calculated from individual observations of the NOAA surface network, following the procedure described in Naus et al. (2019).

3.S2 High-resolution simulations

We performed the inversions of MCF at the coarsest resolution supported in TM5: 6° longitude by 4° latitude, with 25 vertical levels. TM5 can also be operated at 3° by 2° and at 1° by 1° , with up to 60 vertical levels. The resolution at which a model is operated affects the simulated transport, with a higher resolution generally being associated with more accurate transport. Therefore, we performed forward simulations of MCF at different resolutions to determine the impact of resolution on our inversion results. We performed simulations at 1° by 1° with 25 vertical layers for the 1998-2001 period, and at 3° by 2° with 25 and with 60 vertical layers for the 1998-2008 period. Since the amplitude of the effect was apparent from the first few years of the simulations, we did not simulate the whole period. In these forward simulations, we used the prior settings of the REF inversion.

Figure 3.S2 shows the effect of model resolution on the mismatch between global mean mixing ratio and the interhemispheric gradient of MCF, as determined from surface sites. We find that with a higher horizontal resolution the global lifetime and the interhemispheric gradient of MCF are slightly reduced. The reduced lifetime is driven by an increase in the secondary sinks, both oceanic loss and stratospheric photolysis. However, the increase in the secondary sinks corresponds to only 0.1% to 0.5% of the total OH sink and is largely time-invariant. Therefore, we deem these effects not important for our derivation of OH variations. Not shown here are the site-to-site gradients, because these were only marginally affected by model resolution.

In conclusion, we find the impact of increased resolution to be small, so that the conclusions derived from the coarse-resolution inversions remain unaffected.

3.S3 Simulations of SF₆ and HFC-152a

We used the TM5 set-up described in Section to also simulate SF₆ and HFC-152a mole fractions, in a forward simulation from 1998 to 2018. The objective of simulating additional tracers was to diagnose whether the intrahemispheric biases that are apparent for MCF (as discussed in Section 3.4) are specific to MCF, or can also be seen for other tracers. To do this, we sampled a simulation with HFC-152a and SF₆ at the same NOAA surface sites as the MCF simulations. Both HFC-152a and SF₆ are measured in the NOAA HATS network, following the procedures similar to those for MCF, which are described in Section 2.2.1.

3.S3.1 SF₆

Similar to MCF, SF₆ is emitted only anthropogenically, and its emissions are therefore predominantly located in the Northern mid-latitudes. Atmospheric loss of SF₆ occurs in the upper stratosphere, slowly, resulting in a global lifetime of >1000 years Kovács et al. (2017). Atmospheric gradients of SF₆ are therefore solely determined by emissions and transport. Given reasonable constraints on SF₆ emissions, the atmospheric distribution of SF₆ is a good proxy for large-scale atmospheric transport Levin & Hesshaimer (1996) Gloor et al. (2007). In our simulations, we used emission fields described in the TransCom Age of Air project Krol et al. (2018), with no loss process implemented. In Krol et al. (2018) it was shown that these emissions, when implemented in TM5, realistically reproduce the global evolution of SF₆, as well as its interhemispheric gradient.

Figure 3.S3 shows the simulated and observed intrahemispheric SF₆ gradients between ALT and MLO in the Northern Hemisphere, and between SMO and SPO in the Southern Hemisphere. These gradients can be compared to Figure 5 for MCF. We are able to reproduce the intrahemispheric gradients of SF₆. For MCF, we found intrahemispheric gradients to be underestimated by 20-30%. If this were related to a transport bias in TM5, then mixing within the hemispheres would be too fast, and intrahemispheric gradients of SF₆ should also be systematically underestimated. We do find some multi-annual differences between observed and simulated SF₆ mole fraction gradients, but these differences are of opposite sign in both hemispheres and do not persist over the whole twenty-year period, which points to other error sources (e.g. the emission distribution and magnitude). Therefore, we deem it unlikely that a transport bias is the dominant driver of the underestimate of MCF intrahemispheric gradients.




Figure 3.S3: Simulated and observed SF₆ mole fraction gradients between two site pairs. We applied a twelve-month moving average to the monthly mean mole fractions per site.

3.S3.2 HFC-152a

HFC-152a is an anthropogenically emitted gas, used mainly as an aerosol and foam-blowing agent. The dominant atmospheric loss process of HFC-152a is its reaction with OH, resulting in a global lifetime of 1.6 years (Ko et al. (2013b)). In Liang et al. (2017), HFC-152a was proposed as a tracer for [OH]_{GM}, alternative (or complimentary) to MCF. HFC-152a was shown to be especially sensitive to the seasonality and to the interhemispheric ratio of OH, partly due to its shorter lifetime compared to MCF. Most countries report HFC-152a emissions, but often aggregated with emissions of other HFC's. Additionally, atmospheric observations indicate that HFC-152a emissions, as reported in these bottom-up inventories, are significantly too low (Simmonds et al. (2016)). We used emission maps provided in the EDGAR v4.2 inventory, which are available up to 2008. For later years, we repeated the 2008 distribution. We adopted annual emission totals from Simmonds et al. (2016), which approximately reproduced the global evolution of HFC-152a as observed by the AGAGE network up to 2014. After 2014, we repeated the 2014 emission totals. We implemented loss to OH and ran one simulation with the prior OH fields from Spivakovsky et al. (2000) and one with the optimized OH fields from the REF inversion. We found that when we reduced the emissions reported in Simmonds et al. (2016) by 10%, which is well within their uncertainty bounds, the simulation best reproduced the global growth rate of HFC-152a.

Figure 3.S4 shows site-to-site gradients of HFC-152a, as in Figure 5. We find that the gradients relative to the global mean mixing ratio, which we derived from the surface




Figure 3.S4: Simulated and observed HFC-152a mole fraction gradients between three site pairs. We applied a twelve-month moving average to the monthly mean mole fractions per site and we calculated the fractional gradients relative to the global mean mole fraction of HFC-152a.

network using the same methods as for MCF, are largely insensitive to the emission totals we adopt. Given significant uncertainties in the emission distributions, it is expected that we do not capture all the gradients perfectly. However, we do expect that the SMO - CGO gradient is least sensitive to the emission distribution, and most sensitive to the OH distribution, under the assumption that it is unlikely that a large fraction of HFC-152a emissions is located in the Southern Hemisphere. Here, we see that the HFC-152a gradient between SMO and CGO is underestimated in the simulation with prior OH concentrations. This underestimate increases significantly if we impose the optimized OH fields. We note that the adjustments to the OH distribution that are needed to fully correct for the biases in the simulations of MCF are even larger than the ones implemented here, which would further deteriorate agreement with observed HFC-152a gradients .

We acknowledge that the HFC-152a budget is uncertain, as is its emission distribu-




Figure 3.S5: The latitudinal distribution of the annual total ocean flux from Wennberg et al. (2004) and that of the first-order ocean sink from our posterior REF simulation, at three points in time. The fluxes are given for five latitudinal bands visualized in shaded areas. Fluxes are a function of ocean-atmosphere exchange rates and area covered by oceans in each latitudinal band.

tionSimmonds et al. (2016) and we do not suggest that any underestimate of tropical OH concentrations is excluded with this HFC-152a analysis. However, an explanation that attributes the underestimate of intrahemispheric biases of MCF completely to an underestimate of tropical OH concentrations seems to be inconsistent with simulations of HFC-152a. Therefore, we conclude that the intrahemispheric biases seem to a large degree specific to MCF, which makes it more likely that errors in the ocean-atmosphere exchange of MCF play a role (see Section 3.4).

3.S4 Ocean-atmosphere exchange of MCF in TM5

Figure 3.S5 illustrates the latitudinal distribution of the Wennberg ocean flux at three points in time, compared to the first-order ocean sink used in our inversions. The absolute flux is largest and persistently negative in the tropics in both estimates, whereas at high latitudes the estimates diverge. Since the first-order ocean sink is based on optimized atmospheric mole fractions from the REF inversion, agreement in the tropics between the two estimates indicates that the forecasted MCF mole fractions in Wennberg et al. (2004) were at least qualitatively realistic. For most of the period, the Wennberg ocean flux at high latitudes is positive, i.e. the oceans release MCF into the atmosphere. The difference at high latitudes is most pronounced in the Southern Hemisphere, mostly due to a larger ocean surface area. The latitudinal distribution of the Wennberg flux is somewhat time-variant, but persistently shows an enhanced latitudinal gradient compared to the first-order ocean sink over the 1998-2018 period. These observations are consistent




Figure 3.S6: Simulated and observed MCF mole fraction gradients, treated analogous to Figure 5. Here, we compare the optimized (red solid) and prior (red dashed) gradients from the REF inversion to a forward simulation that includes the ocean-atmosphere exchange from Wennberg et al. (2004).

with a relatively constant intrahemispheric bias that is higher in the Southern than in the Northern Hemisphere. For reference, the REF inversion found anthropogenic MCF emissions of 2.1 Gg in 2017, which can be compared to the 0.5 Gg difference between the Wennberg flux and the first-order sink in the Southern Hemisphere in the same year.

To test the potential impact on our derived results, we replaced the first-order ocean sink with the Wennberg ocean flux in a forward simulation of MCF, while we kept the other fields the same as in the prior of the REF inversion. The effect on the intrahemispheric gradients of MCF is shown in Figure 3.S6. For comparison, the prior and optimized simulations that result from the REF inversion are also shown. Predictably, adding MCF preferably at high latitudes compared to tropical latitudes deepens the tropical minimum of MCF. We find that this adjustment in the ocean sink provides strong leverage on the intrahemispheric gradient: implementation of the Wennberg flux improves the gradients more than the $\sim 30\%$ adjustments in the latitudinal OH distribution that were derived in the REF inversion. Further implications of this result are discussed in Section 3.4.

While the air-sea exchange from Wennberg et al. (2004) provides a viable alternative to a first-order sink, the estimated exchange does depend strongly on simulated ocean mixing. A second source of uncertainty are the oceanic hydrolysis rates. These rates are based on measurements performed at temperatures above 25°C , and were extrapolated to colder temperatures following the linear Arrhenius relation. While this relation was shown to hold at a range of temperatures above 25°C Gerken & Franklin (1989) Jeffers et al. (1989), the only experiment performed at 10°C Pearson & McConnell (1975) provided much faster hydrolysis rates than those based on extrapolation. As concluded in Wennberg et al. (2004), in-situ oceanic measurements of MCF would be needed to confirm and better constrain the oceanic release of MCF.

Based on these uncertainties, we emphasize that the actual air-sea exchange of MCF could differ significantly from Wennberg et al. (2004). Better quantification and further confirmation of an ocean source would be greatly helped by both lab measurements of MCF hydrolysis rates at low temperatures and by field measurements that investigate the ocean flux of MCF. Admittedly, the latter would be hampered by the currently low atmospheric and oceanic abundance of MCF. In the absence of additional measurements, the surface network of MCF is our most substantial source of information on the ocean flux. However, given the spatially similar signature of the proposed ocean flux and the OH sink, we consider it unlikely that both processes can be simultaneously constrained in an inversion based on surface network observations. Partly for this reason, we choose not to further investigate the issue in additional inversions. However, we do find evidence in the underestimation of MCF intrahemispheric gradients that makes oceanic release of MCF more likely, even if we can't exactly quantify its magnitude. Such information is important for interpretation of the atmospheric MCF record.

3.S5 Ten year (1998-2008) inversions

In addition to the three twenty-year inversions, we performed three ten-year inversions, from 1998 to 2008. Since a ten-year inversion takes less time per iteration and the inverse problem is smaller, we found that the inversions converged more consistently and in fewer iterations. Considering the significant convergence problems we have found in the twenty-year inversions (see Section 3.5), this exercise served as a sanity check. In principle, if everything works correctly, then the twenty and ten-year inversions should converge to the same solution, excepting a spin-up and spin-down period.

Figure 3.S7 shows the global $k\cdot[\text{OH}]$ variations derived from the ten and twenty-year inversions. It is clear that whereas there was quite some spread between the twenty-year inversions, the ten-year inversions converged to a similar solution that is most consistent with the standard twenty-year inversion, although adjustments in the ten-year inversions are larger in amplitude.

Similarly, the spatial adjustments in the OH distribution of the three ten-year inversions are very similar (Figure 3.S8), even if the posterior distribution of the TM5-OH inversion still retains some of the differences present in its prior distribution, such as a two-pronged maximum in OH concentrations. In turn these adjustments are most similar in shape to the twenty-year REF inversion, although more extreme in amplitude.

We have also included in Figure 3.S8 OH distributions reported in literature. Except for the Spivakovsky et al. (2000) distribution, these estimates cover a range of full-chemistry simulations. Two of these (the CAMS-OH and the Gaubert et al. (2017) estimate) additionally assimilated satellite observations of trace gases such as CO, for an improved atmospheric state. The amplitude of OH adjustments necessary to fit intrahemispheric




Figure 3.S7: Monthly variations in global mean $k(T) \cdot [\text{OH}]$, derived in the three different inversion set-ups. Solid lines indicate the results from inversions from 1998 to 2018: these are the same lines as those shown in Figure 1 of the main text. Dashed lines indicate the results from the ten-year inversions, which ran from 1998 to 2008.

gradients of MCF is best quantified in the posterior solutions from the ten-year inversions. Here we see that the resulting tropics to extratropics OH ratio falls well outside the range of literature estimates that represent our best understanding of OH chemistry. Of course, this could indicate incomplete understanding. However, we consider the more likely explanation that part of the intrahemispheric bias we find for MCF is best explained by a high-latitude ocean source of MCF (e.g. Section 3.4 of the main manuscript).

In all cases, the ten-year inversions better reproduced MCF observations than the twenty-year inversions over the 1998-2008 time period. Mostly as a consequence of larger adjustments in the latitudinal OH distribution, the intrahemispheric gradients are reproduced better in the ten-year inversion, although still not quite captured (Figures 3.S9). The inversion in which emissions were distributed according to population density (POP) consistently performs most poorly. Interestingly, the ten-year inversion that converges best uses an OH field based on a full-chemistry simulation of TM5, rather than the often-adopted distribution from Spivakovsky et al. (2000). However, this solution does require significantly elevated MCF emissions over the 1998-2008 period: a total of 21.4 Gg over ten years, or an average increase of 12.0% relative to the prior. This increase is likely a consequence of relatively high OH concentrations in the Northern Hemisphere, so that higher MCF emissions are required to balance the interhemispheric MCF gradient. Additionally, in the ten-year inversions, variations in MCF emissions and variations in OH adjustments do become significantly negatively correlated (typically $r = -0.6$), which was not the case for the twenty-year inversions. This indicates that the two cannot be fully




Figure 3.S8: The latitudinal distribution of tropospheric OH concentrations derived in the ten and twenty year inversions (dashed and dotted lines, respectively) and the same distribution for a set of literature estimates (solid lines). All literature estimates, except for Spivakovsky et al. (2000), have been derived in full-chemistry simulations. We have scaled each distribution to the corresponding tropospheric mass-averaged OH concentration to focus on the distribution of OH between tropics and extratropics. Additionally, we have grouped OH estimates from the Atmospheric Chemistry and Climate Model Intercomparison Project (ACCMIIP)(Lamarque et al., 2013) as gray lines, to distinguish better between literature estimates and inversion results. Note that different literature estimates correspond to different time periods. Details of the literature estimates are outlined in Table 3.S1.

disentangled at this level of detail.

By reducing the problem to one decade, we have more confidence that the inverse system managed to find a statistically correct solution than in the twenty-year inversion. This is supported by the consistency of the derived solutions. Since the solutions derived in the ten-year inversions most closely match the twenty-year REF inversion, we consider that the REF inversion managed to converge furthest. We are unsure why this particular inverse set-up managed to converge further than the other two. Notably, the interannual $k(T) \cdot [\text{OH}]$ variations between the REF 20 year and 10 year inversions match quite well, except in 2006-2007, which can be considered spin-down years of the ten-year inversion. This lends confidence to this aspect of our twenty-year REF estimate. For the latitudinal adjustments in OH concentrations, it can be observed that more extreme adjustments result in incremental improvements in intrahemispheric gradients. I.e., given strong enough adjustments in the OH distribution, in combination with compensatory elevated MCF emissions, the intrahemispheric gradients can be reproduced. However, given the extremity of the necessary OH adjustments of at least 60% and the need for substantially increased emissions, we doubt the physical realism of this solution.

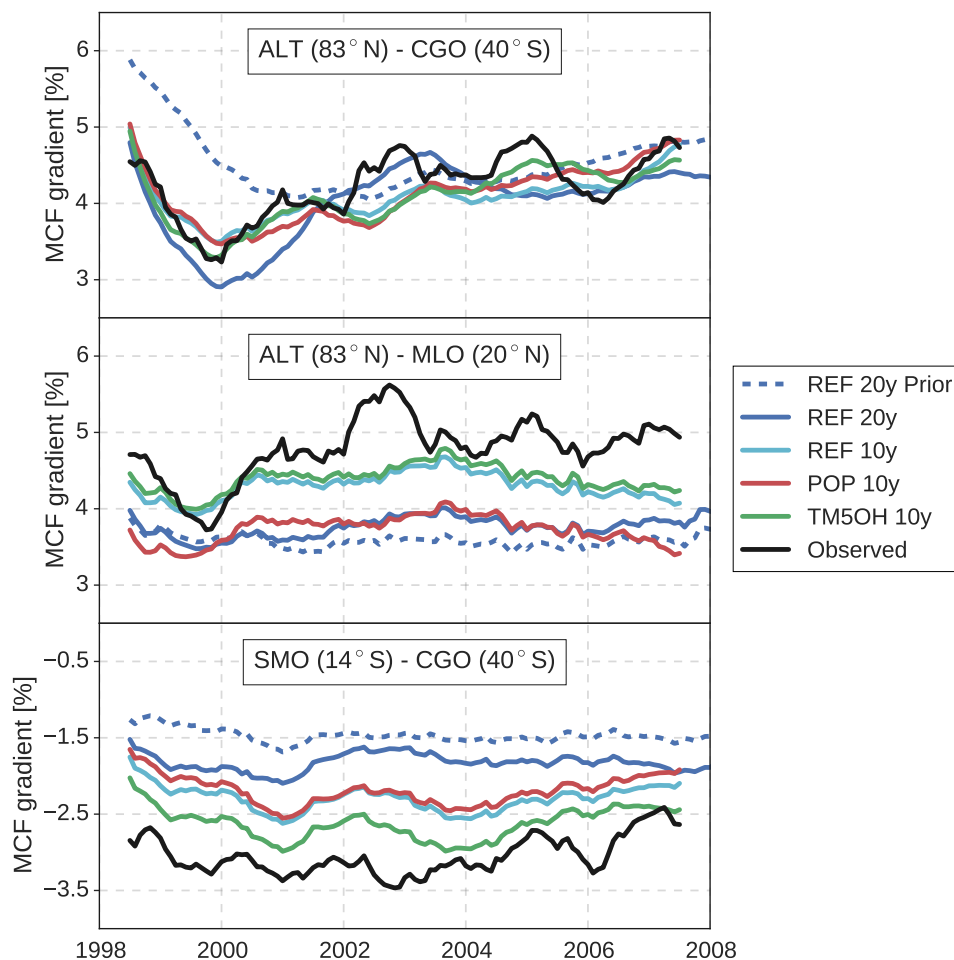


Figure 3.S9: Simulated and observed MCF mole fraction gradients, treated analogous to Figure 5 in the main text. Here, we compare all three ten-year inversions and the REF inversion. We only show the prior gradient of the REF inversion (red dashed).

3.S6 The adjoint of OH chemistry

A 4DVAR inversion requires the adjoint of the forward version of TM5 to calculate the cost function gradient. The adjoint of TM5 has been described in previous publications Meirink et al. (2008) Krol et al. (2008), here we describe the additions that are required for optimization of OH concentrations.

The adjoint of the forward model \mathbf{H} is equivalent to the transpose of \mathbf{H} , which is straightforward to find if \mathbf{H} is linear. However, as loss to OH depends on both MCF mole fractions and OH concentrations, an adjustment in OH concentrations has a non-linear impact on loss of MCF. In that case, the adjoint is defined as the transpose of the tangent linear version of \mathbf{H} . Equations 3.3 describe the forward OH chemistry in TM5.

$$\text{MCF}(t + dt) = \text{MCF}(t) (1 - k_{OH} [\text{OH}](t)), \quad (3.3a)$$

$$[\text{OH}](t + dt) = [\text{OH}](t). \quad (3.3b)$$

$[\text{OH}]$ is the model OH concentration in molecules cm^{-3} , which does not vary within a month, dt is the model timestep in seconds and k_{OH} the reaction rate between MCF and OH in $\text{cm}^3 \text{s}^{-1} \text{molecules}^{-1}$. The resulting tangent linear model, with respect to adjustments in OH concentrations, is given in Equations 3.4.

$$d\text{MCF}(t + dt) = d\text{MCF}(t) (1 - k_{OH} [\text{OH}])(t) - \text{MCF} k_{OH} d[\text{OH}](t), \quad (3.4a)$$

$$d[\text{OH}](t + dt) = d[\text{OH}](t). \quad (3.4b)$$

Since both MCF and OH are affected by $[\text{OH}]$ adjustments, a non-linearity results. From the tangent linear model, the adjoint model can be derived and is given in Equations 3.5.

$$\text{adjMCF}(t) = \text{adjMCF}(t) (1 - k_{OH} \text{OH}(t)) \quad (3.5a)$$

$$\text{adj}[\text{OH}](t) = \text{adj}[\text{OH}](t) - \text{MCF}(t) k_{OH} \text{adjMCF}(t). \quad (3.5b)$$

As in the forward model, the adjoint fields of MCF mole fractions are 3-hourly, whereas the adjoint fields of OH concentrations are monthly. For the calculation of the adjoint of OH (Equation 3.5b), we need to save the 3-hourly, 3D MCF mole fraction fields from the forward run: effectively, we linearize around these fields. Note that if OH concentrations were not optimized, only Equation 3.5a would be needed.

Secondly, we need to translate between $[\text{OH}]$ and scaling factors of $[\text{OH}]$. The conditional scaling of $[\text{OH}]$ is defined in Equation 3.6.

$$[\text{OH}](t) = \begin{cases} [\text{OH}]_{\text{prior}}(t) (1 + f_{OH}(t)), & \text{if } f_{OH} \geq 0, \\ [\text{OH}]_{\text{prior}}(t) \exp(f_{OH}(t)), & \text{if } f_{OH} < 0, \end{cases} \quad (3.6)$$

with $[\text{OH}]_{\text{prior}}$ the monthly prior OH field and f_{OH} the monthly scaling factor that enters the state vector. Since $[\text{OH}]_{\text{prior}}$ is fixed, the only non-linearity comes from the exponent, and the tangent linear scaling becomes:

$$d[\text{OH}](t) = \begin{cases} [\text{OH}]_{\text{prior}}(t) df_{OH}(t), & \text{if } f_{OH} \geq 0, \\ [\text{OH}]_{\text{prior}}(t) df_{OH}(t) \exp(f_{OH}(t)), & \text{if } f_{OH} < 0. \end{cases} \quad (3.7)$$

Finally, the adjoint of the scaling factor can be described in terms of the adjoint field of OH as:

$$adjf_{OH}(t) = \begin{cases} [OH]_{prior}(t) adj[OH], & \text{if } f_{OH} \geq 0, \\ [OH]_{prior}(t) adj[OH] \exp(f_{OH}(t)), & \text{if } f_{OH} < 0. \end{cases} \quad (3.8)$$

Aggregation of scaling factors per grid box to zones, e.g. over the longitudinal dimension, is done through summation. Emissions are scaled following an analogous version of Equations 3.6-3.8, as in Bergamaschi et al. (2009).

3.S7 Tables

Table 3.S1: Model names and references that correspond to the different OH estimates in Fig. 3.S8. We have also included references to the dataset locations. Note that we only included ACCMIP participants that had the data we required (OH fields) easily available, since this subset sufficed for our comparison. For ACCMIP models we have shown also which version we retrieved, denoted by the rXiYpZ notation.

Model	Years	Model ref	Data ref	Further notes
CESM-CAM-Superfast	2000	Lamarque et al. (2012)	Shindell et al.	From ACCMIP, r1i1p1
CMAM	2000	Scinocca et al. (2008)	Shindell et al.	From ACCMIP, r1i1p1
EMAC	2001	Jöckel et al. (2006)	Shindell et al.	From ACCMIP, r1i1p1
MIROC-CHEM	2000	Watanabe et al. (2011)	Shindell et al.	From ACCMIP, r1i1p1
UM-CAM	2000	Zeng et al. (2010)	Shindell et al.	From ACCMIP, r1i1p1
GFDL-AM3	2000	Donner et al. (2011)	Shindell et al.	From ACCMIP, r1i1p1
MOCAGE	2000	Teyssède et al. (2007)	Shindell et al.	From ACCMIP, r1i1p1
GISS-E2-R	2000	Koch et al. (2006)	Shindell et al.	From ACCMIP, r1i1p3
CICERO-OsloCTM2	2000	Skeie et al. (2011)	Shindell et al.	From ACCMIP, r1i1p1
STOC-HadAM3	2000	Stevenson et al. (2004)	Shindell et al.	From ACCMIP, r1i1p1
LMDzORINCA	2000	Szopa et al. (2013)	Shindell et al.	From ACCMIP, r0i1p1
CAMS	2010-2011	Inness et al. (2019)	CAMS reanalysis - OH	-
CESM-CAM-Chem	2003-2013	Gaubert et al. (2017)	Gaubert & Worden (2017)	-
EMAC/MOM	2013-2014	Lelieveld et al. (2016)	Gromov et al. (2020)	-

Table 3.S2: Description of the NOAA surface sites used in this study. Shown are abbreviation, full name and location, latitude, altitude and longitude, and the start and end year of measurements at a site. The instances where a site does not cover the full inversion period are shown in bold.

Site short	Location	Lat. (°)	Long. (°)	Alt. (m)	Measurement period
SPO	South Pole, Antarctica	-90.0	-24.8	2810	1992 – Present
PSA	Palmer Station, Antarctica	-64.9	-64.0	10	1997 – Present
CGO	Cape Grim, Tasmania, Australia	-40.7	144.7	164	1991 – Present
SMO	Tutuila, American Samoa, USA	-14.2	-170.6	42	1991 – Present
KUM	Cape Kumukahi, Hawaiï, USA	19.5	-154.8	8	1995 – Present
MLO	Mauna Loa, Hawaiï, USA	19.5	-155.6	3397	1991 – Present
NWR	Niwot Ridge, Colorado, USA	40.0	-105.6	3526	1991 – Present
THD	Trinidad Head, California, USA	41.0	-124.2	107	2002 – Present
HFM	Harvard Forest, Massachusetts, USA	42.5	-72.2	340	1995 – Present
LEF	Park Falls, Wisconsin, USA	46.0	-90.3	472	1996 – Present
MHD	Mace Head, County Galway, Ireland	53.3	-9.9	26	1998 – Present
BRW	Barrow, Alaska, USA	71.3	-156.6	11	1992 – Present
SUM	Summit, Greenland	72.6	-38.5	3210	2004 – 2018
ALT	Alert, Nunavut, Canada	82.5	-62.5	205	1991 – Present

Chapter 4

An atmospheric perspective on Amazon fires

This chapter is based on:

Naus, S., Domingues L. G., Krol, M. C., Luijkx, I. T., Gatti, L V., Banda, N., Koren, G., Basu, S., Miller, J. B., Gloor, M., Peters, W. et al. (2020). An atmospheric perspective on Amazon fires. *To be determined, in preparation*

Abstract

Despite consensus on the overall downward trend in Amazon forest loss since the previous decade, data from local logging activity and remotely sensed biomass-, fire-, and burned-area changes still vary widely in their estimated yearly carbon emissions from deforestation. Here, we use sixteen years of satellite-derived carbon monoxide (CO) columns to constrain fire emissions from the Amazon basin over the 2003-2018 period. Through data assimilation, we produce 3-daily maps of fire carbon emissions from the Amazon that we verified to be consistent with a long-term monitoring program of aircraft CO profiles over five sites in the Amazon. Our new product independently confirms a long-term decrease of close to 55% in deforestation-related CO emissions over the study period. Interannual variability is large, with known anomalously dry years showing a more than two-fold increase in basin-wide fire emissions. At the level of individual Brazilian states, we find that both soil moisture anomalies and human ignitions determine fire activity, suggesting that future carbon release from fires depends on drought intensity as much as on continued forest protection. Our study shows that the atmospheric perspective on deforestation has become a valuable additional monitoring instrument that complements existing methods.

4.1 Introduction

The role of Amazon forests in supporting biodiversity, regional ecosystem services, and carbon storage (Gloor et al., 2012) is threatened by human activities, particularly large-scale deforestation (Davis et al., 2020). In Brazil specifically, various studies suggest that, in recent years, deforestation rates and associated fire activity are once again accelerating (INPE, 2020; Pereira et al., 2020), after having reached a minimum around 2012 (Yin et al., 2020). Moreover, recent droughts in 2010 and 2015/16 led to maxima in biomass burning (Silva Junior et al., 2019). Fire monitoring is key to detect such changes, and to gain the deeper understanding of anthropogenic and environmental influences on tropical deforestation that society needs to limit its adverse impacts.

Monitoring of Amazonian fires primarily relies on remote sensing products such as fire counts (Wiedinmyer et al., 2011), albedo changes (van der Werf et al., 2017), and Fire Radiative Power (FRP; Kaiser et al. (2012)), which are partly related (Fanin & van der Werf, 2015). Rapid and continuous processing of vast amounts of such data allowed recent unexpectedly high fire activity in 2019 to be detected and reported rapidly (Lizundia-Loiola et al., 2020; Brando et al., 2020). However, fire dynamics are complex, and products based on land remote sensing data are prone to miss small fires (Randerson et al., 2012), hampered by cloud occurrence (Schroeder et al., 2008), and might be poorly able to detect understory fires (Morton et al., 2013). Understory fires in particular contribute strongly to forest fragmentation and mortality, and feed back on vulnerability to burning (Nepstad et al., 2001; Alencar et al., 2004; Balch et al., 2011). In addition to direct detection of fire activity, information is needed to quantify the corresponding carbon loss to the atmosphere on scales from decades to seasons, and from the entire Amazon basin down to individual Brazilian states.

Here, we report on the use of atmospheric remote sensing data of Carbon Monoxide (CO) columns to constrain Amazon fire emissions over the 2003–2018 period. CO provides integrated information about fire occurrence and its imprint on atmospheric composition. We use a sixteen-year record of CO satellite data from the Measurement Of Pollution In The Troposphere (MOPITT) instrument (Deeter et al., 2019), and combine it with the data assimilation system TM5-4DVAR (Krol et al., 2005; Meirink et al., 2008) to optimize surface CO emissions from Amazon fires. We independently assess our emissions, as well as those from several existing products, using a multi-year aircraft sampling network of CO in the basin (Gatti et al., 2014). We explore the longterm trends and the interannual variability in the optimized fire emissions on various spatial scales to better understand the role of anthropogenic and natural influences on resulting fire emissions.




Figure 4.1: A map of the South-American zoom regions used in the TM5 simulations. The green and red grids indicate the 3° by 2° and 1° by 1° zoom regions, respectively. Over both zoom regions satellite observations were assimilated. Filled stars indicate NOAA surface sites that are additionally assimilated to constrain the global background of CO. Filled circles indicate sites where aircraft profiles were sampled that we use for independent validation.

4.2 Methods

4.2.1 Inverse system TM5-4DVAR

We operate the atmospheric transport model TM5 (Krol et al., 2005) at a global resolution of 6° longitude by 4° latitude, with nested zoom regions over South-America. Specifically, the outer South-American zoom region has a resolution of 3° longitude by 2° longitude and the inner region a resolution of 1° by 1° , in a set-up similar to e.g. van der Laan-Luijkx et al. (2015) (Figure 4.1). All regions are operated with 25 vertical layers. Transport in TM5 is driven by offline meteorological fields from the ERA-Interim reanalysis from the European Centre for Medium Range Weather Forecasts (ECMWF) (Dee et al., 2011).

We employ the inverse framework TM5-4DVAR (Meirink et al., 2008) to optimize CO emissions with CO observations. Optimization in the 4DVAR framework is quantified as minimization of the cost function J (Equation 4.1a), and its gradient ∇J (Equation

4.1b).

$$J(\vec{x}) = \frac{1}{2}(\vec{x} - \vec{x}_{prior})^T \mathbf{B}^{-1}(\vec{x} - \vec{x}_{prior}) + \frac{1}{2}(\mathbf{H}\vec{x} - \vec{y})^T \mathbf{R}^{-1}(\mathbf{H}\vec{x} - \vec{y}), \quad (4.1a)$$

$$\nabla J(\vec{x}) = \mathbf{B}^{-1}(\vec{x} - \vec{x}_{prior}) + \mathbf{H}^T \mathbf{R}^{-1}(\mathbf{H}\vec{x} - \vec{y}). \quad (4.1b)$$

Here, \vec{x} denotes the state, with the emission parameters to be optimized, and \vec{x}_{prior} is the corresponding prior estimate. \mathbf{B} denotes the error matrix of the prior estimate, i.e. uncertainties in the prior emission fields. Similarly, \vec{y} is the vector containing all observations, with \mathbf{R} the error matrix of the observations. \mathbf{H} denotes the forward model, which is applied to the state vector to calculate the difference between simulated and observed CO. This difference is then used in the adjoint model \mathbf{H}^T to calculate the cost function gradient.

4.2.2 Input fields and prior errors

For chemical production of CO from methane (CH_4) and from non-methane volatile organic compounds (NMVOCs), we use fields generated in a 2006 simulation in the full-chemistry version of TM5 (Huijnen et al., 2010). The prior estimate of biomass burning is based on the Global Fire Assimilation System (GFAS) (Kaiser et al., 2012). Fire emissions are distributed vertically in the simulations following vertical emission profiles derived in the Integrated System for wild-land Fires (IS4FIRES) (Soares et al., 2015), which include a diurnal cycle. We use anthropogenic CO emissions from the Monitoring Atmospheric Composition and Climate Cityzen (MACCCity) inventory Lamarque et al. (2010). The hydroxyl radical (OH) concentration fields are a combination of tropospheric OH fields from (Spivakovsky et al., 2000), scaled by 0.92 as recommended in (Patra et al., 2011), and stratospheric OH fields derived in the 2D Max Planck Institute for Chemistry (MPIC) chemistry model (Brühl & Crutzen, 1993).

Inside both zoom domains, we optimize emissions from biomass burning, with a grid-box error of 250%, a horizontal correlation length of 200 km and a temporal correlation of three days. Total emissions in the global domain are optimized with a prior uncertainty of 250%, and with a horizontal and temporal correlation of 1000 km and 9.5 months, respectively. Emissions are optimized non-linearly to prevent negative emissions. This set-up is similar to that employed in Nechita-Banda et al. (2018).

Our study period is 2003–2018, which is explored in 16 separate inversions that each cover the April–December period of one year. Since we use the same initial field each year, and since we do not optimize initial conditions, April and December can be considered spin-up and spin-down months respectively. We note that emissions in these months are mostly insignificant, so that they have been included in our emission estimates.

4.2.3 Observations

As primary constraint on biomass burning emissions in South-America, we assimilated CO columns from the Measurements Of Pollution In The Troposphere (MOPITT) instrument, version 8 (Deeter et al., 2019). CO columns were only assimilated over the two zoom domains (Figure 4.1), and we only included columns retrieved from the thermal infrared (TIR) waveband. We inflated errors reported by MOPITT with a factor $\sqrt{50}$ to compensate for the high number of satellite data ($\sim 10000/\text{day}$), relative to the number of surface observations ($\sim 150/\text{month}$).

In addition to the satellite data, we assimilated CO mole fraction observations from the surface network of the National Oceanic and Atmospheric Administration (NOAA) (Petron et al., 2019). We used a fixed observational error of 2 ppb CO, with no model error. Compared to the satellite CO columns, surface observations have high precision and are well-calibrated to an absolute scale, so that they provide a high-quality constraint on the global background of CO.

We include aircraft profiles in our analysis that were not assimilated in the inverse system, so that they can be used for independent validation. Aircraft profiles were sampled at five sites in the Brazilian Amazon: Alta Floresta (ALF; 8.80°S, 56.75°W), Rio Branco (RBA; 9.38°S, 64.8°W), Santarém (SAN; 2.86°S, 54.95°W), Tabatinga (TAB; 5.96°S, 70.06°W) and Tefé (TEF; 5.96°S, 70.06°W). The record started in 2003 with measurements only in SAN. In 2010, operation on ALF, RBA and TAB started, and in 2013 the TAB site was replaced by TEF. The sampling flights were performed using small aircraft, typically two times per month, between 12–13 hr local time, in a descending spiral profile that avoids emissions from the aircraft. Profiles typically include 12–17 air samples between 300 m and 4500 m height above sea level (Gatti et al., 2010, 2014).

In addition to the standard inversions outlined above, we present sensitivity tests in the Supplemental Material. These include a set of 2010–2018 inversions in which satellite CO columns from the Infrared Atmospheric Sounding Interferometer (IASI) instrument Clerbaux et al. (2009) are assimilated, instead of MOPITT columns (Supplement 4.1.2), and two sets of 2003–2018 inversions that start from different priors for fire emissions (Supplement 4.2.1).

4.3 Results

4.3.1 Basin-wide Amazon fire emissions

MOPITT column-mean mole fractions of CO reveal distinct temporal and spatial patterns over the Amazon, which are largely driven by fire emissions. The bottom panel in Figure 4.2 shows that years with high fire activity are interspersed by years with smaller fire activity in the biomass burning season (August–December). Yearly basin-wide averages of

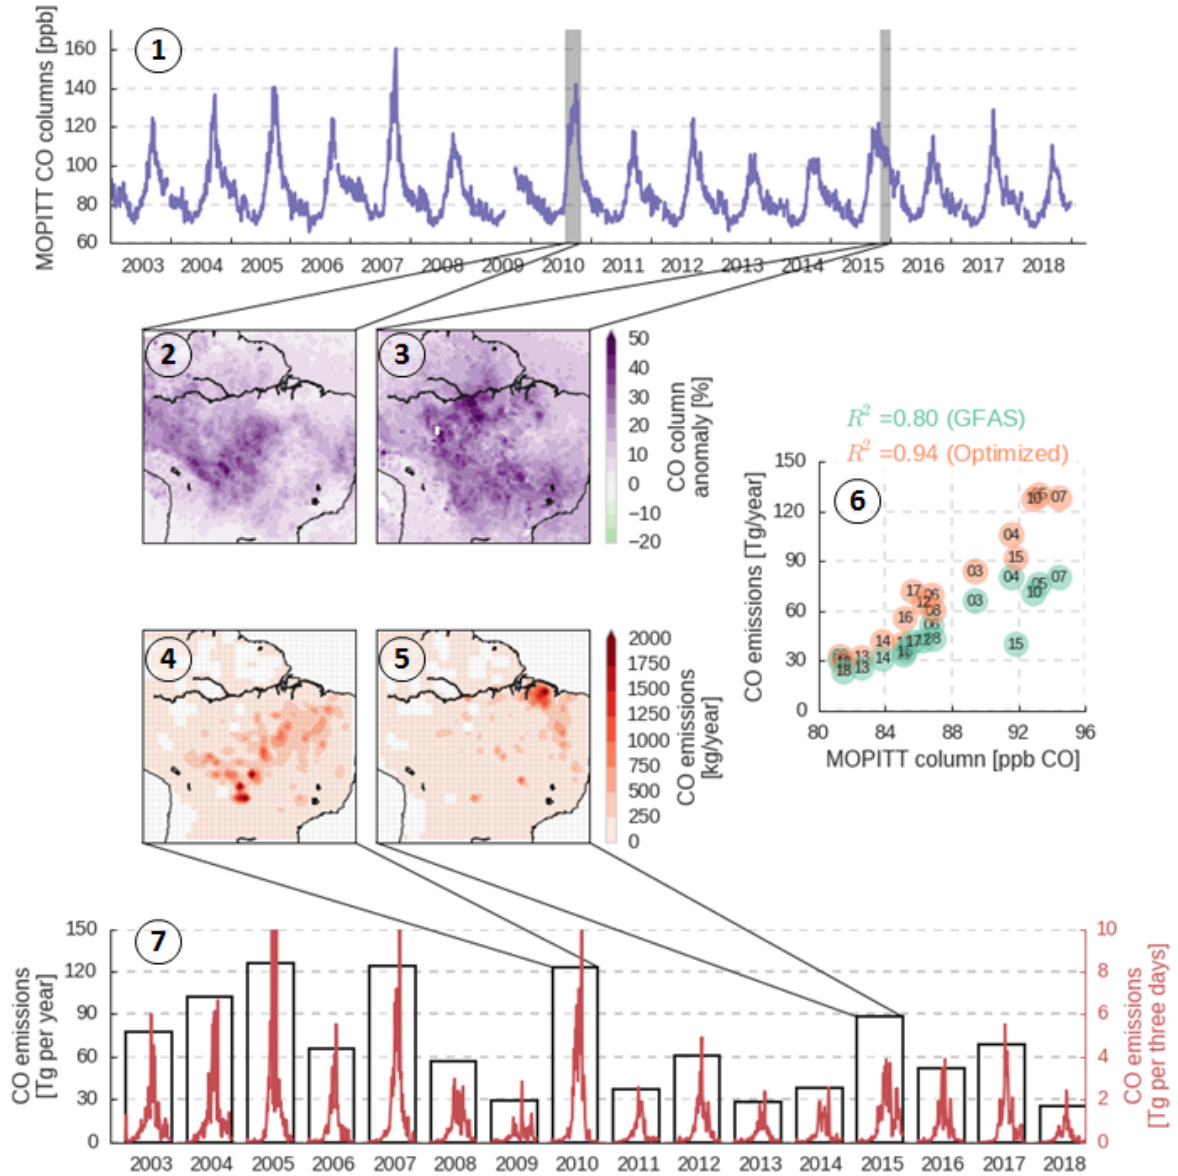


Figure 4.2: Satellite-observed CO columns from the MOPITT instrument and CO fire emissions derived from these satellite observations. **Panel 1:** MOPITT CO columns averaged over our 1° by 1° , South-American zoom domain, shown at three-daily time resolution. **Panels 2-3** Two maps of the anomalies in MOPITT-observed CO columns in 2010 (left) and 2015 (right), relative to the 2003-2018 average. Note that shaded bars in the top figure indicate the specific time-period covered by either map. **Panels 4-5** Two maps of (April – December) total CO emissions, after optimization with MOPITT-observed CO columns, for 2010 (left) and 2015 (right). **Panel 6** Correlation plot between Apr-Dec total CO emissions from biomass burning and the Apr-Dec averaged observed MOPITT CO columns, both for the 1° by 1° zoom domain. We show GFAS emissions (green) and emissions after assimilation of MOPITT CO columns in TM5-4DVAR (orange). Years corresponding to each data point are shown inside the points (20XX). **Panel 7** April – December total (black bars) and three-daily (red lines) total CO emissions from biomass burning in South-America after optimization with MOPITT CO columns. Emissions are summed over our 1° by 1° , South-American zoom domain.

MOPITT CO columns correlate strongly ($R^2=0.80$; $p=0.00$; $N=16$) with the prior GFAS emission estimates. After MOPITT CO columns have been assimilated, the correlation between inferred yearly basin-wide emissions and MOPITT CO columns, predictably, increases ($R^2=0.94$; $p=0.00$; $N=16$) (Figure 4.2).

The link between spatial patterns in CO columns and derived emissions is less direct. Figure 4.2 (panels 4 and 5) shows that fires in different years, here 2010 and 2015, show markedly different spatial patterns, that do not directly match the spatial distribution of CO columns (panels 2 and 3). The capacity of our system to constrain such regional-scale fire emissions originates from the inverse mapping of CO columns on emissions through transport in TM5 and its adjoint (e.g. Meirink et al., 2008; Kopacz et al., 2009), which provides a considerable advantage over studies that analyse satellite CO columns directly (e.g. Aragão et al., 2018). While the resulting accurate simulation of the assimilated MOPITT CO columns (Supplement 2.1.1) and of NOAA surface observations (Supplement 2.1.3) provides evidence for the correct implementation of our emission optimization procedure. Next, we will show that our MOPITT-based emission estimates also lead to an improved comparison with independent aircraft profiles.

4.3.2 Validation with aircraft profiles

Aircraft profiles were sampled over five sites in the Amazon between 2010 and 2017. We compare our simulated CO mole fractions to those measured from samples of vertical profiles from the network of Gatti et al. (2014) in Figure 4.3. This independent validation clearly shows an improved overall match after the assimilation of the MOPITT CO columns. In comparison to the simulation with prior emissions, the bias reduces from -62 ppb to -19 ppb. For Santarém, the residual bias is relatively large, but we additionally found that the CO columns are well matched over the same region (Supplement 4.3). We do find a significant absolute residual error between simulated and observed aircraft profiles, which can be explained by the relatively coarse resolution of the transport model that puts a lower limit on how well individual aircraft samples can be represented in TM5. However, we find that, in general, the MOPITT-derived emissions much improve the agreement with independent aircraft profiles at five different locations across the Amazon, compared to the GFAS prior. This highlights the capability of our system to estimate CO mole fractions and CO emissions across different states of Brazil.

4.3.3 Fire emissions in Brazilian states

Our inverse system demonstrates high skill in estimating CO fire emissions on the level of individual Brazilian states, where we find a strong link in interannual variability of fire CO emissions with local soil moisture anomalies. In Figure 3 we show our annual fire CO emissions together with soil moisture from the Global Land Evaporation Amsterdam Model (GLEAM, v3a; Martens et al. (2017)) for 5 states in Brazil. Interannual variations

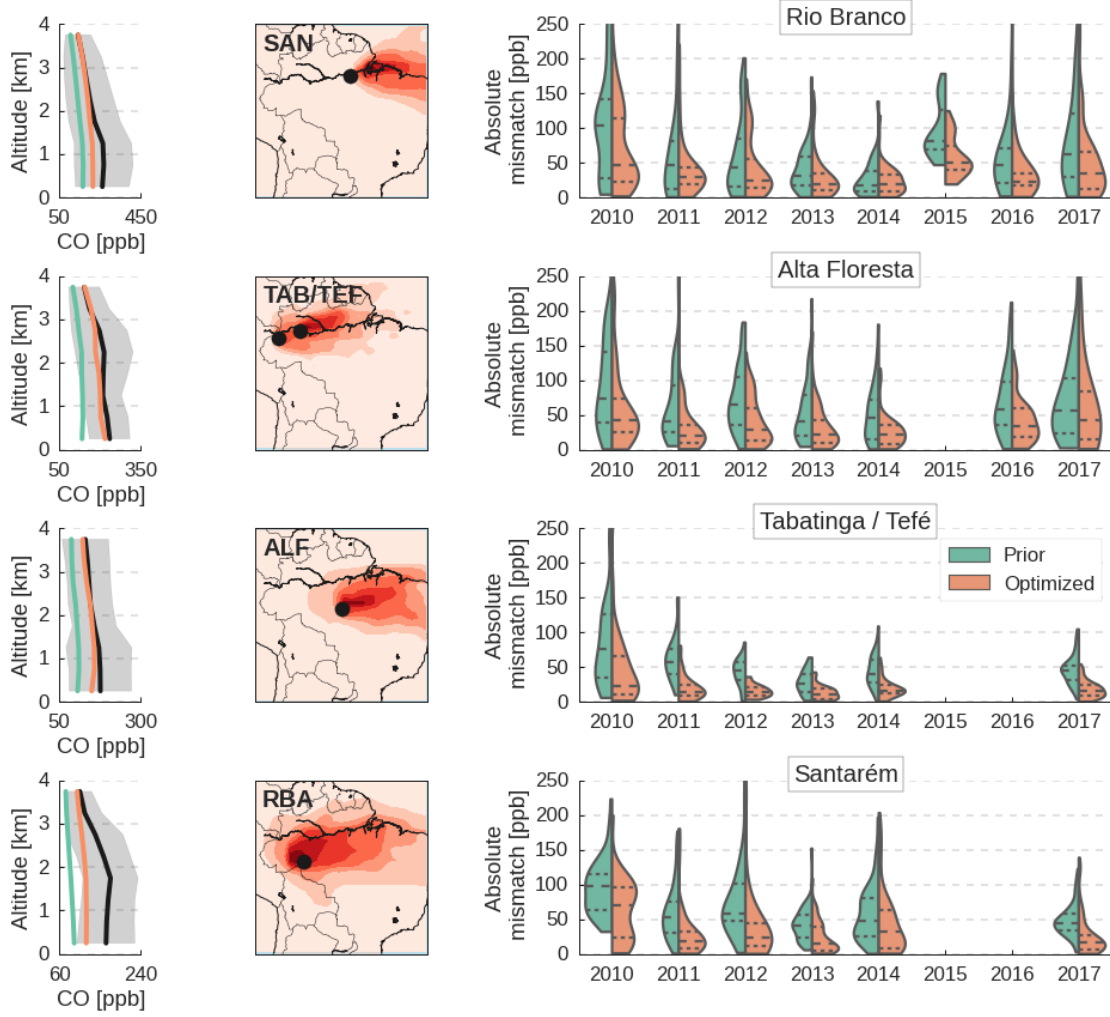


Figure 4.3: Comparison between simulated and observed aircraft profiles over five sites in the Amazon. The profiles cover the 2010-17 period, and we have included only profiles sampled between August-November. **Left column:** Time-averaged (2010-2017) simulated (prior in green; posterior in orange) and observed (black) aircraft profiles, binned in 500 meter intervals. Grey shaded areas show one standard deviation of the variability between the individual, observed aircraft profiles. **Center column:** Maps of the influenced area of each site, or site combination. Black dots indicate site locations, and the red area indicates the origin of air at the site location. Red areas are proportional to the logarithm of the number of back-trajectories that originate at the sampling location and altitude, and then pass through a grid cell, as determined from simulations in the HYSPLIT model. Further details are provided in Gatti et al. (2010). Backtrajectories from the Lagrangian grid in the HYSPLIT model were interpolated to the TM5 $1^\circ \times 1^\circ$ grid. The number of backtrajectories can be considered a rough proxy for typical air mass origin during the fire season. **Right column:** Violin plot of the absolute difference between observed and simulated aircraft samples of CO, in a simulation with GFAS (green) and in a simulation with MOPITT-optimized biomass burning emissions (orange). Dashed lines inside each violin indicate the median and the two inner quartiles.



Figure 4.4: Interannual variability in soil moisture and fire emissions in five Brazilian states. The four outside figures show timeseries of April–December total CO emissions from biomass burning (left axis, in red) and annual averaged soil moisture from the GLEAM product (right axis, in blue), for five Brazilian states. A year is defined here as June–December, i.e. centered on the dry fire season. CO emissions shown have been optimized with MOPITT observations. Also shown at the top are the correlation coefficients r between CO emissions and soil moisture. Maranhão and Tocantins are combined, since they are smaller states that represent similar regimes in terms of climate and anthropogenic activity. Note the different y-scales between the four figures, with lowest emissions in Amazonas. The centre figure shows a map of these five Brazilian states.

in CO emissions are markedly different per state and, in general, years with low soil moisture levels show high fire emissions. For example, the 2010 fires are mostly located in Mato Grosso, while the 2015/16 El Niño fires occur in Pará and Maranhão/Tocantins. These years coincide with local, negative anomalies in soil moisture. Furthermore, we find that posterior emission allocation between states is largely independent of assumptions about the prior emissions distribution (Supplement 4.2.1), except in strongly anomalous cases (e.g. in 2015), where information from the prior improves the obtained location and timing of the fires. Previous work had already established a strong link between fire emissions and soil moisture (e.g. Asner & Alencar, 2010; Silva et al., 2018), but the strength and consistency of the anti-correlation we find at the level of individual states is revealing. In contrast to the interannual variability, the long-term decreasing trend in fire CO emissions appears not to be driven by soil moisture, as is especially visible in the state of Pará.

4.3.4 Long-term trends

Our MOPITT-derived biomass burning emission patterns allow independent verification of reported long-term downward trends in deforestation rates, confirming the large role of humans in shaping Brazil’s fire regime. We find a longterm decrease in fire CO emissions over the 2003-2018 inversion period for 5 year running means, that varies between landcover types (Figure 4.5, panel 2). The decrease is strongest in forest-covered areas

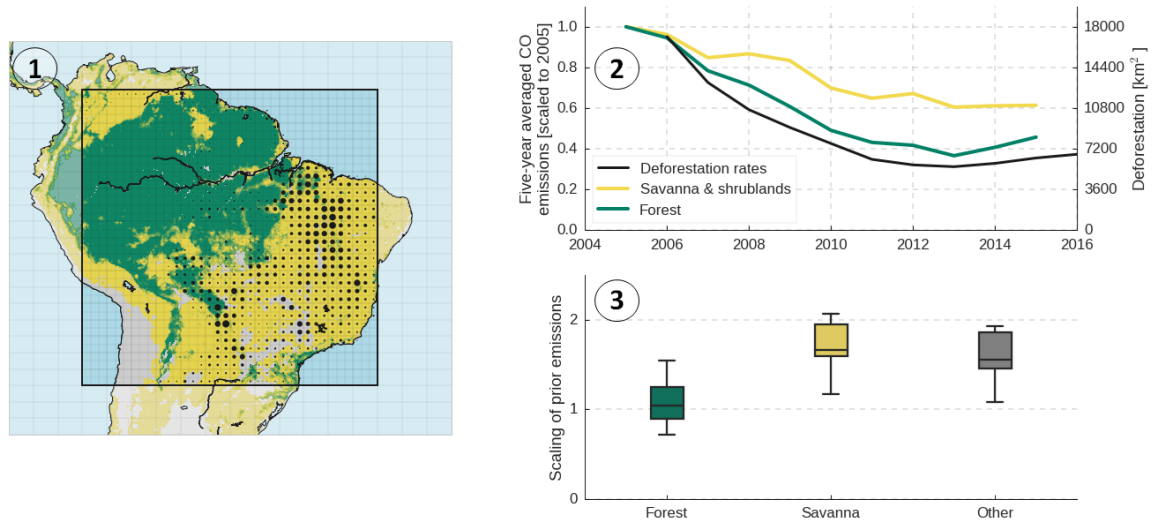


Figure 4.5: Spatial allocation of emissions between landcover types, and longterm trends in the emissions. **Panel 1:** A land-cover map of South-America, with landcover types retrieved from the MODIS product, at 0.1° by 0.1° resolution. We distinguish between forest (green), and savanna and shrublands (yellow). Grid cells covered by less than 70% of both are shown in gray. The map covers the 3° by 2° zoom region used in TM5, and the black outline indicates the 1° by 1° zoom region, which is the focus of our study. The area of the black circles is proportional to how much TgCO is added in the inversion to the GFAS prior emissions, over the 2003-2018 study period. **Panel 2:** Timeseries of five-year averaged annual total CO emissions from biomass burning over the $1^\circ \times 1^\circ$ TM5 zoom domain for two landcover types: forest (in green) and savanna (in yellow). Emission totals are scaled to the emissions in the first five-year window (see left y-axis), which is centered on 2005. Also shown are five-year averaged deforestation totals (black, right axis), as retrieved from the National Institute for Space Research (INPE). **Panel 3:** Box-plot that shows how much annual total emissions over each landcover type are adjusted through optimization with MOPITT satellite observations. This can be compared to the black circles in Panel 1.

($\sim 55\%$) and matches a decrease in deforestation rates reported by the Brazilian National Institute for Space Research (INPE, INPE (2020)) in both timing and magnitude. In the areas dominated by savanna and shrublands we also observe a smaller downward trend. Notably, both deforestation rates and fire emissions have stopped decreasing since 2012.

GFAS CO emissions are too low to reproduce MOPITT CO columns in all years, for all landcover types, but the underestimate is much stronger over savanna than over forests (Figure 4.5, Panel 3 and black circles in Panel 1). A strong systematic underestimate over savanna is indicative of underestimated CO emission factors, since other explanations, such as missed small fires or understory fires, are more likely to impact emission estimates from forests. In addition to the systematic 2003–2018 underestimate, we find that emissions in dry years, such as 2010 and 2015, are underestimated more strongly than emissions in wet years (e.g. Panel 6 in Figure 4.2). We find a comparable underestimate in emissions from other inventories (Supplement 4.2.1).

4.4 Discussion

The magnitude and spatial distribution of our posterior CO emissions are robust with respect to prior emission inventory or other CO budget terms. We find that the largest uncertainty in the optimized emissions derives from uncertainty in CO production from non-methane hydrocarbons (NMHCs; see Supplement 4.2.2), and we conservatively estimate a combined uncertainty in basin-wide emissions of ± 15 Tg/year (Supplement 4.2.5). We also repeated the 2003–18 inversions with climatological, interannually repeating prior fire emissions, which resulted in similar emission allocation between states, and similar interannual variability and longterm trends (Supplement 4.2.1). In 2015 in particular, we derive high emissions over the Santarém region in November and December, i.e. during the onset of the wet season that is normally not associated with high fire activity. These emissions differ significantly from climatology and from the emissions in GFAS (Figure 4.2, panel 6). Interestingly, even though GFAS emissions are strongly adjusted in 2015, we find that they also contain information that is necessary for the inverse system to derive these large adjustments, since these are not found to the same degree in an inversion based on a climatological prior (Supplement 4.2.1). The anomalous case of 2015 exemplifies that each component of the inverse system is necessary to derive the most realistic posterior emissions.

The robustness of derived emissions signifies the detail provided by the MOPITT-TIR product, but it also makes our estimates highly reliant on the quality of this product. MOPITT CO has been validated extensively (Deeter et al., 2019) and in our work we again present a good agreement with independent aircraft profiles. A comparable CO satellite data product is provided by the Infrared Atmospheric Sounding Interferometer (IASI) instrument (Clerbaux et al., 2009; George et al., 2015) and we have used this prod-

uct in a set of inversions over the 2010–2018 period (Supplement 4.1.2). A comparison with MOPITT-derived emissions over the whole period is not possible, because of inconsistencies in the reprocessing used for IASI, but the interannual variability derived from the two products within the 2010–2013 and 2015–2018 periods is similar. A consistently processed IASI product would help better validate the MOPITT-derived emissions.

An operational framework that estimates CO fire emissions based on satellite-observed CO columns and prior information (e.g. FRP) can provide unique and valuable information about regional variability and trends in fires. Such an estimation framework would be timely, because there are indications that pressures on the Amazon have increased in recent years (INPE, 2020; Pereira et al., 2020). The Copernicus Atmospheric Monitoring Service (CAMS) already provides an operational data assimilation framework in which, among other data products, MOPITT-TIR CO columns are assimilated (Flemming et al., 2017). However, in their approach the atmospheric abundance of CO is optimized, instead of surface CO emissions. A next development of CAMS should therefore consider emission optimization, since spatio-temporal allocation of CO emissions is a prerequisite for a good fire proxy. In a sensitivity study we find that their fire emissions, contained in the GFAS product, might be biased low, which, in the CAMS system, is partly compensated by OH concentration fields that are significantly too low over the Amazon (Supplement 4.2.3).

While we have derived a robust timeseries of CO emissions, it remains difficult to translate CO emissions to other valuable quantities, such as total carbon emissions or fire occurrence. For example, if the underestimate we find in GFAS CO emissions is related to errors in the CO emission factors, then the total carbon emissions reported in the GFAS inventory could still be accurate. If, on the other hand, the underestimate is related to missed understory fires, then it would also impact the estimate of total carbon emission. Likely, the differences between GFAS and MOPITT-optimized CO emissions are caused by a combination of these two error sources, and attribution of the differences to either requires further investigation. By distinguishing between landcover types, and between systematic versus time-dependent differences, we have already shown that such an approach is feasible (see Section 4.3.4). Moreover, the fire emissions we have derived are already used in a CO₂-focused analysis of drought response in the Amazon (Ch. 5 of Koren, 2020, paper in prep.)

4.5 Conclusions

We demonstrate an excellent capability to translate satellite-observed CO columns to CO emissions from fires in the Amazon domain, anchored in a unique set of aircraft profiles. Our inverse approach starts from the valuable spatiotemporal distribution of Amazon fires contained in bottom-up inventory estimates (e.g. GFAS), and sharpens these with constraints from atmospheric observations. The robust timeseries of CO emissions that

result are a rich proxy for fire variability and deforestation that is highly valuable, since changes in the natural and political climate surrounding the Amazon necessitate active fire monitoring by as many independent proxies as possible. In addition to the showcase results, our sensitivity analyses (Supplement 4.2) reveal the crucial components in an atmospheric CO inversion of the Amazon.

Variations in CO emissions over our 2003–2018 study period are a combination of strong interannual variations and a long-term decrease, mostly between 2003–2012. Interannual variations are closely correlated with variations in the Amazonian water balance, evident from a strong link with soil moisture even at state-level. In contrast, the long-term decline in CO emissions over the 2003–2012 period mirrors a decrease in deforestation rates, especially so in forested regions. These results emphasize the positive effect of deforestation abatement policies, as well as the potential impact of increased drought frequency in a changing climate. As such, sustained efforts to reduce deforestation can reduce the impact of climate change on fire risk, while a return to deforestation rates of the early 2000’s in a drier climate likely results in enhanced fire risks.

Supporting Information to Chapter 4

4.S1 Match with and sensitivity to assimilated observations

In this section we assess how well the prior and optimized simulations reproduce observed MOPITT CO columns and the NOAA surface observations. Since both sets of observations are assimilated, a good posterior match is expected. In addition, in Supplement 4.S1.2, we present a series of inversions in which satellite CO columns from the Infrared Atmospheric Sounding Interferometer (IASI) instrument (Clerbaux et al., 2009) were assimilated, instead of the MOPITT CO columns.

4.S1.1 Comparison between observed and simulated MOPITT-TIR CO columns




Figure 4.S1: The relative difference between observed and simulated MOPITT CO columns, in a simulation with the prior GFAS emissions (green) and after the MOPITT CO columns have been assimilated to optimize fire emissions (orange). The differences shown are three-daily averages over the 1° by 1° South-American zoom domain. The difference between simulated and observed columns is quantified as a percentage of the observed CO column.

In the prior simulation, we find significant differences between the domain-averaged simulated and observed CO columns (green line in Fig. 4.S1). After assimilation these differences are much reduced to less than 2% of the observed columns, except for a spin-up period of approximately one month (orange line in Fig. 4.S1). This result confirms that with adjustments in the GFAS biomass burning emissions inside the zoom domains, combined with adjustments in total emissions outside, we can reproduce observed MOPITT CO columns in our simulations. Differences between simulated and observed columns in the prior simulation are already indicative of the types of adjustments in CO emissions that are needed to reproduce observed columns. Firstly, simulated CO columns outside the dry season are systematically too high, by about 8-15%. Secondly, superimposed on this systematic overestimate, we find that the peaks in CO columns during dry season are underestimated in the prior simulations. The overestimate we attribute largely to too-high secondary production of CO (mostly from NMVOCs): an issue that is further




Figure 4.S2: Total CO emissions from biomass burning summed over the $1^\circ \times 1^\circ$ South-American zoom domain, and over the April – December inversion period. Results for three sets of inversions are shown, which each started from the GFAS fire prior. The first is the default inversion that used MOPITT satellite data, the second used IASI satellite data and the third used MOPITT satellite data, but adopted an observational error of 0.2 ppb CO on surface observations (instead of 2.0 ppb in the other two inversions). Purple bars are discussed in Supplement 4.S1.3.

explored in Supplement 4.S2.2. Since we do not optimize secondary production in our inversion, the overestimate is largely fixed by adjusting the boundary conditions of the domain, which are only loosely determined by the NOAA surface sites (see Supplement 4.S1.3). The underestimate of the dry-season peak in CO is most likely related to fires, since this is the CO source with the largest seasonal cycle. Indeed, in the inversions biomass burning emissions of CO are increased significantly compared to the GFAS prior, especially in dry years (middle right panel in Figure 4.2 of the main text and Supplement 4.S2.1).

4.S1.2 Inversions driven by IASI satellite data

In addition to the reference inversions that used MOPITT CO columns, we have also performed inversions with CO columns from the IASI instrument (Clerbaux et al., 2009). Similar to the MOPITT-TIR product, the IASI retrieval exploits the thermal infrared waveband to estimate atmospheric CO columns. This similarity in retrieval wavelength means that results based on IASI and on MOPITT-TIR CO columns are comparable, although differences in vertical sensitivities and retrieval method exist. We used the IASI-MetopA dataset, which is available from 2010 to present.

Two aspects stand out when comparing biomass burning emissions derived from IASI to those derived from MOPITT (Fig. 4.S2). Most importantly, before 2014, MOPITT-derived emissions are slightly higher than IASI-derived emissions (difference of 6 to 12 Tg CO/year), while after 2014 IASI-derived emissions are significantly higher than MOPITT-derived emissions (difference of 15 to 30 Tg/year). Secondly, we note that within each of these two time periods the difference between interannual variability in IASI- and MOPITT-derived emissions is small (~ 10 Tg/year) relative to this jump (~ 30 Tg/year).




Figure 4.S3: The difference between observed MOPITT CO columns and MOPITT CO columns sampled in a simulation that used CO emissions optimized with IASI CO columns, averaged over the 1° by 1° South-American zoom domain. The difference is quantified as a percentage of the observed CO columns. The date on which IASI switches between two meteorological datasets (2014/09/30) is indicated, and on this date a jump in the difference between simulated and observed MOPITT CO columns occurs.

This is important because it means that, barring the offset, both products give a comparable signal for interannual variability in CO emissions.

To investigate the timing of the jump between IASI- and MOPITT-derived emissions, we sampled MOPITT columns in a simulation with IASI-optimized biomass burning emissions (Fig. 4.S3). Analogous to the effect on derived emissions, over 2010–2013 simulated MOPITT CO columns are slightly lower than those observed (i.e. assimilating MOPITT columns would result in higher emissions), with a larger, reverse offset over 2015–2018. This switch in offset of around 8% occurs instantaneously on 2014/09/30, as indicated in Fig. 4.S3. Over the 2010–2018 period, several changes have been made to the IASI retrieval that can cause inconsistencies (e.g. Table 2 in (Bouillon et al., 2020)), and a major update to the processing algorithm that was implemented on 2014/09/30 apparently has a particularly large impact on the retrieved CO columns. Based on the coincidence of these two events, we consider the switch in IASI–MOPITT offset to be an artifact in the IASI data record.

For the nature of this study, long-term consistency of the driver satellite data is crucial, and we have therefore focused on the MOPITT-based inversions. In a previous study of the 2015 fires in Indonesia, the TM5-4DVAR system was used both with MOPITT and with IASI CO columns (Nechita-Banda et al., 2018). There, the spatio-temporal signals in MOPITT and IASI were largely consistent, with IASI-derived emissions being somewhat higher than MOPITT-derived emissions, which is consistent with what we find here. Their work included a more detailed comparison between MOPITT and IASI, including the influence of selecting daytime or nighttime data and the influence of higher coverage by IASI on derived emissions. Both effects were small relative to the MOPITT–IASI offset, indicating that this offset is likely driven by instrument sensitivities and retrieval methods. This is confirmed by the large influence of switching between meteorological

datasets that we identify here.

We conclude that as long as the IASI retrieval does not use a consistent meteorological dataset, the retrievals before and after 2014/09/30 are best treated as two separate data records. The IASI team is currently working on a full reprocessing of the CO Metop-A record, using the ERA5 reanalysis as input for temperature profiles in order to generate a homogeneous record. A consistent IASI product would provide excellent grounds for an uncertainty estimate of the driver satellite data, which is currently more difficult to perform. Additionally, the two satellite products have complimentary strengths. Namely, since in the MOPITT retrieval more weight is given to a climatological prior, compared to the IASI retrievals, the MOPITT product performs best under typical conditions, while the IASI product is better able to capture extreme events that deviate from climatology (George et al., 2015; Nechita-Banda et al., 2018). Note, however, that these different sensitivities do not play a role when averaging kernels are used in the model-satellite comparison, as is done in this study.

The impact of the MOPITT-IASI differences on this study’s conclusions is limited. The systematic difference between MOPITT and IASI is a measure for systematic uncertainty in the satellite data and its impact on derived emissions, which amounts to around 10–30 Tg/year. This is still small compared to the adjustments made to the GFAS prior (up to 60 Tg/year). Additionally, the relative consistency in interannual variability between IASI and MOPITT, excepting the 2014 break, is evidence for robustness of interannual variability in CO emissions derived from either satellite product.

4.S1.3 Comparison between simulated and observed NOAA surface observations

Since we assimilate satellite observations only over the South-American zoom domains, global CO emissions and, by extent, South-American boundary conditions are constrained by flask observations sampled at NOAA background sites. We assimilated all NOAA surface sites for which CO data are available, which amounts to 37–44 sites, depending on the year. Typically, 3–4 observations per site, per month are available, adding up to a global total of 100–150 observations per month. We used a fixed measurement uncertainty of 2 ppb CO for surface observations in the inverse system.

In Fig. 4.S4 the difference between simulated and observed CO mole fractions is shown at four representative surface sites, for 2010–2016. The comparison looks similar for the entire 2003–2018 period, but because the sensitivity inversion only includes 2010, 2015 and 2016, we zoom in on the 2010–2016 period. Alert (ALT) is a high-altitude Northern Hemispheric site, and, after optimization, mole fractions at ALT are very well reproduced, except for a spin-up period. The optimization works less well for the other three sites, which are located closer to the zoom domains. In the optimization, there is a trade-off between improving the match with surface observations and improving the match




Figure 4.S4: Timeseries of the difference between simulated and observed monthly mean CO mole fractions at four NOAA surface sites. Shown are results before optimization (green), after optimization in our standard inversion set-up with an observational error of 2.0 ppb per sample (orange), and after optimization in an inversion where the observational error was reduced to 0.2 ppb (black; only for 2010, 2015 and 2016).

with satellite observations that is particularly important near the South-American zoom domains, where coverage of surface observations and satellite CO columns partly overlap. Because there are many more satellite observations (around 10000 per day) compared to surface observations (3-4 per month, per site), satellite observations are weighted more heavily than surface observations. We have tried to reduce this effect by inflating the error on satellite observations by a factor $\sqrt{50}$ (as in (Nechita-Banda et al., 2018)), but near the zoom domain satellite observations still seem to dominate. This can be seen from the relative good match at ALT (and most other sites; not shown) compared to the other three sites shown in Fig. 4.S4. Especially Ragged Point, Barbados (RPB) and Ascension Island (ASC) are important for the inversion, since these sites roughly represent the inflow conditions of the Amazon domain (e.g. Gatti et al. (2014)). The posterior differences between simulated and observed mole fractions at these sites fall outside the uncertainty bounds set for the inversions and, especially at RPB, are not always an improvement on the prior simulation.

To investigate the importance of our difficulty in reproducing mole fractions at surface sites near the zoom domain, we performed inversions for 2010, 2015 and 2016, in which we reduced the error on CO surface observations from 2 ppb to 0.2 ppb. As expected, in these new inversions the agreement with observations at all sites improves compared to our standard inversions (black lines in Fig. 4.S4), while MOPITT CO columns in these three new inversions are reproduced equally well as in our standard inversions (not shown). Biomass burning emissions derived in these three inversions are lower than those derived in our standard inversions (Fig. 4.S2), but the difference is small relative to the

total adjustment to GFAS prior emissions, and consistent between years (7–11 Tg/year). We note that the sign of this difference is consistent with the posterior mismatch at RPB (orange line in Fig. 4.S4). Namely, the simulated mole fractions at RPB are consistently too low in our standard inversions, while in the inversions with a 0.2 ppb error simulated mole fractions at RPB are higher, which results in more inflow of CO into the Amazon domain and consequently lower local emissions derived by the inversion.

Notably, RPB is located inside the $3^\circ \times 2^\circ$ zoom domain, i.e. satellite observations are assimilated over this site. Therefore, the interannual consistency of the posterior bias at RPB (orange line in Fig. 4.S4) can be compared to the aircraft profiles at Santarém, where simulated mole fractions are also too low, even though MOPITT CO columns are reproduced well over the same area (see Supplement 4.S3 and Fig. 4.S11). The similarity between these offsets suggest either a systematic bias in the MOPITT CO columns, or systematic uncertainties in the vertical transport of TM5, since the vertical sensitivity of MOPITT CO columns, of surface observations and of aircraft profiles are all different. In previous work this issue was partly addressed by including a latitudinally dependent bias correction to MOPITT inversions in the optimization (Hooghiemstra et al., 2012). However, as long as the offsets are relatively constant interannually, the conclusions we draw in the main text are not significantly affected.

Assimilating NOAA surface observations in our inversion helps to derive a consistent global picture and the poor match between simulated and observed mole fractions at sites near the Amazon domain seems worrying. However, the inversions with a 0.2 ppb observational error show that we can simultaneously reproduce surface observations and MOPITT CO columns with approximately the same biomass burning emissions as in the standard inversions. This indicates a limited sensitivity of derived emissions to the boundary conditions. Intuitively, this can be understood, since, during the dry season, advection of CO into the Amazon domain is less important than local emissions. In conclusion, we find some inconsistencies in the Amazon boundary conditions in our standard inversions, but we also find limited sensitivity of derived emissions to these inconsistencies. In that sense, satellite observations put strong constraints on the derived emissions over the Amazon basin.

4.S2 Sensitivities in the inverse system

In our main text, we have focused on one series of inversions and most of our conclusions are based on this single set of inversions. We have not explicitly covered sensitivities and uncertainties inherent to this specific inverse set-up. The 4DVAR inverse system does not provide us with posterior uncertainties in the derived emissions, because our inverse set-up is non-linear (to prevent negative CO emissions, see Methods). Therefore, in this supplement, we discuss sensitivities of the derived CO emissions to the most important components of the inverse system, primarily the prior inventory for fire emissions, CO pro-




Figure 4.S5: Total CO emissions from biomass burning summed over the $1^\circ \times 1^\circ$ South-American zoom domain, and over the April – December inversion period. Results for three sets of inversions are shown, which each used a different fire prior. The first is the default inversion that used the GFAS fire prior, the second used the FINN inventory as prior and the third used a climatological fire prior based on GFAS.

duction from NMHC and OH chemistry. In general, we find that the absolute magnitude of the derived CO emissions, and the spatio-temporal patterns therein, are robust.

4.S2.1 Fire emission inventory used as prior

The inversions presented in the main text all started from the GFAS emission inventory as a prior for biomass burning. The inverse system works by scaling the prior emissions up or down relatively, which means that the prior emission distribution affects how MOPITT CO columns are translated to optimized CO emissions. To quantify the influence of the prior emission distribution, we present results from a set of inversions that started from three different fire emission inventories: the standard GFAS inventory (Kaiser et al., 2012), the Fire INventory from NCAR (FINN) (Wiedinmyer et al., 2011) and a climatological (i.e. annually repeating) prior based on the average emission distribution in GFAS over the 2003–2018 period (hereafter referred to as CLIM).

FINN is based on satellite-observed active fire counts, different from the fire radiative power used for GFAS. Both inventories overlay these proxies with landcover maps from the MODIS instrument, and use landcover-specific emission factors for CO (and other species) to produce the emission estimates we use. However, the landcover classifications and emission factors used differ between the two inventories. For the CLIM prior, we averaged the three-daily CO emission fields from GFAS over the 2003–2018 period. Next, if in a three-day period CO emissions in a grid cell were less than 0.03 Tg, the emissions in that grid cell, in that three-day window, were set to zero. Finally, three-daily total fire emissions inside the 1° by 1° zoom domain were divided uniformly over those grid cells that initially had emissions higher than 0.03 Tg. We chose this approach to prevent spatial gradients in the GFAS estimate from influencing the result, but the 0.03 Tg limit was needed to retain some potential to recover high, localized emissions from the CLIM prior. Outside the $1^\circ \times 1^\circ$ zoom domain, the CLIM prior still uses interannually varying

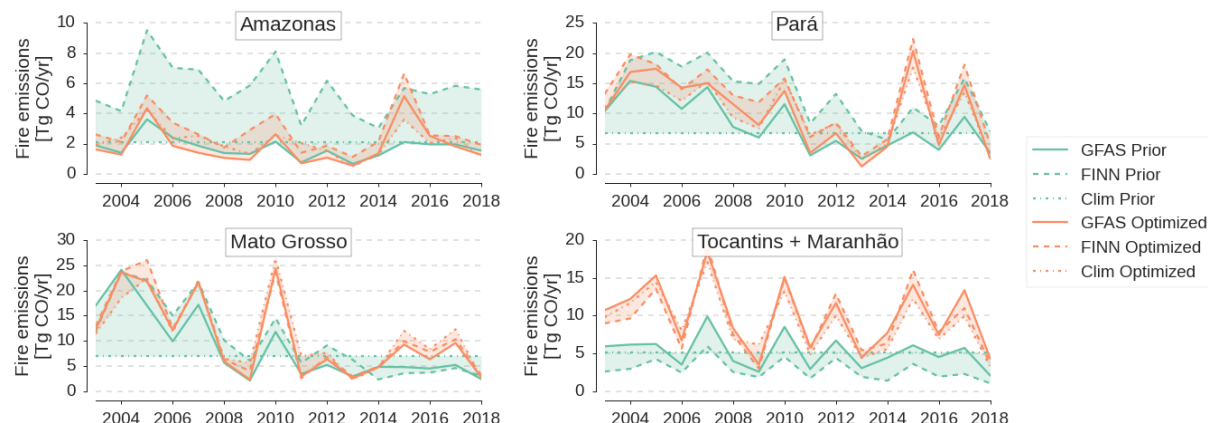


Figure 4.S6: Timeseries of prior (green) and optimized (orange) annual total CO emissions from biomass burning for five Brazilian states. Inversions based on three different emission inventories are indicated with solid (GFAS), dashed (FINN) and dotted (climatological) lines. As a visual help, the spread between the different inversions is shaded in the prior or posterior color. A year is defined here as June-December, i.e. centered on the fire season. Maranhão and Tocantins are combined, since they are smaller states that represent similar regimes in terms of climate and anthropogenic activity. Note that 2011 is missing from the climatological inversions. For a map of the geographic extent of these five states see Fig. 3 of the main text.

GFAS emissions.

The inversions based on the FINN inventory inform on the sensitivity of our derived emissions to starting from a different, but equally good and defensible emission distribution. The CLIM inversions provide a more extreme test for the ability of our system to retrieve interannual variability in fires when no prior variability is included.

On the Amazonian scale of our $1^\circ \times 1^\circ$ zoom domain, we find little sensitivity to the prior emission inventory that is used (Fig. 4.S5). This means that the MOPITT column observations, combined with TM5 transport, provide a strong control on domain-total emissions. The largest difference we find in 2015, during which lower emissions are derived in the CLIM inversion than in the other two inversions, even though the prior emissions are highest in the CLIM inversion. This is related to the anomalous emission timing in 2015, with significant emissions in November and December that are present in the FINN and GFAS priors, but not in the CLIM prior.

On a state-level scale, too, we find similar CO emissions between the three inversions (Fig 4.S6). Annual total emissions per state can be quite different between the prior emission inventories, and after optimization most of these differences disappear. This goes both for the long-term averaged allocation of CO emissions between states and for allocation of fires in specific dry years. For example, in Amazonas, FINN systematically suggests higher CO emissions than GFAS, but after assimilation of MOPITT observations this difference disappears.

In most years, the variability in the GFAS and FINN prior inventories is similar, and matches quite well with the variability in our posterior estimates. However, in some cases the inverse system proposes significant adjustments. For example, the 2010 emissions in Mato Grosso and the 2015 emissions in Pará have to be increased significantly compared to the prior. The latter adjustment is more difficult in the CLIM inversion, due to the anomalous timing of the 2015 fires discussed earlier.

In conclusion, we find only a small effect of the prior inventory used in the inversion on the optimized emissions, on any of the spatial and temporal scales explored in this study. This shows that MOPITT CO columns, combined with the TM5-4DVAR system, provide a strong control on regional attribution of CO emissions that supersedes the influence of the prior emission distribution. We do recognize the importance of starting our inversion from a realistic emission distribution especially in anomalous years, where inversions based on a climatological inventory cannot fully capture fire variability.

4.S2.2 CO production from NMVOC

In this section we investigate the influence of natural CO emissions on our derived fire emissions. The aggregated category of natural emissions as defined here includes secondary production of CO (from CH₄ and NMVOCs), and direct biogenic CO emissions. This source category partly correlates with fire emissions, as CO production from NMVOCs in particular peaks during the dry season, because of significant NMVOC emissions from fires, as well as higher biogenic emissions of natural emitted NMVOCs like isoprene.

The natural emissions used in our standard emissions were interannually repeating and retrieved from a full-chemistry simulation of TM5 for 2006. While we consider the spatio-temporal patterns of this distribution realistic, the absolute amount of natural CO emissions was very high. Globally, 1750 Tg CO/year is emitted in this aggregated category, compared to, for example, 1400 Tg/year suggested in Huijnen et al. (2010). When we compare our natural CO emissions with CO emissions calculated based on the Model of Emissions of Gases and Aerosols from Nature (MEGANv2.1)(Sindelarova et al., 2014), we find confirmation that our reference estimate of natural emissions is too high (Fig. 4.S7). Therefore, as a sensitivity test, we have performed inversions with natural emissions scaled down by 30% (green dashed line in Fig. 4.S7), which is more in line with other estimates.

In the inversions where natural emissions have been scaled down, we derive higher biomass burning emissions (orange bars in Fig. 4.S8). The difference between biomass burning emissions derived in the two inversions is 23-27 Tg/year, while natural emissions are scaled down by 47 Tg/year. In other words, a systematic bias in natural emissions of 1 Tg translates to approximately a 0.5 Tg bias in biomass burning emissions. We can compare this sensitivity to the interannual variability in CO emissions derived from




Figure 4.S7: Timeseries of monthly total biogenic CO emissions inside the $1^\circ \times 1^\circ$ South-American zoom domain. The climatological natural emissions used in this study (green solid), as well as the same estimate uniformly scaled to 70% (green dashed) are both shown. Additionally, natural emissions based on the MEGANv2.1 inventory (Sindelarova et al., 2014) are shown as one standard deviation spread around the 2003–2018 mean (shaded orange). Since the MEGANv2.1 inventory only provides direct biogenic CO emissions and emissions of CO precursors, we have used CO production calculated in simulations described in Zheng et al. (2019), that used precursor emissions from MEGANv2.1 to calculate secondary production of CO (see also Supplement 4.S2.4).

MEGANv2.1. The difference between natural emissions in the highest and lowest year is 30 Tg/year, which translates to a maximum uncertainty in biomass burning emissions of ~ 15 TgCO/year.

Given these results, we consider natural emissions to be one of the most important uncertainties in our inverse system. Since we used high natural emissions in our standard inversions, the biomass burning emissions we have derived might be on the low side. We note that our sensitivity inversions show that systematic errors in natural emissions mostly translate to systematic errors in biomass burning emissions, i.e. the interannual variability of our estimates remains unaffected. However, natural emissions do undergo interannual variability, albeit less so than biomass burning. We conservatively estimate an error in the interannual variability of derived CO emissions of 10–15 TgCO/year related to natural emissions. In contrast, our analysis in Supplement 4.S2.4 shows that biomass burning emissions derived in a different inverse system that used natural emissions from the MEGANv2.1 inventory are very similar to those derived in this study. This suggests that, ultimately, the uncertainty related to natural emissions can be cancelled out by other uncertainties.




Figure 4.S8: Total CO emissions from biomass burning summed over the $1^\circ \times 1^\circ$ South-American zoom domain, and over the April – December inversion period. Results for three sets of inversions are shown, which each started from the GFAS fire prior, and they all were optimized with MOPITT satellite data. The first is the default inversion, the second used reduced all natural emissions (e.g. CO production from NMHC) by 30% and the third used CAMS fields for OH, instead of the default OH fields which were interannually repeating and based on Spivakovsky et al. (2000) and Brühl & Crutzen (1993). The influence of OH on the inversions is discussed in Supplement 4.S2.3

4.S2.3 OH chemistry fields

The final uncertainty in the CO budget that we consider is loss to oxidation by OH. Since OH has a lifetime of seconds, its atmospheric abundance is determined by local atmospheric conditions, which can be disturbed by high emissions from fires: precisely our regime of interest. Our default OH field (from Spivakovsky et al. (2000) in the troposphere, scaled by a factor 0.92; from Brühl & Crutzen (1993) in the stratosphere) is climatological, and it therefore does not include the interannual disturbances in fire emissions. Since CO has a typical atmospheric lifetime of one month, variations in OH will drive an integrated, slow response in atmospheric CO abundance, which suggests that the impact of varying OH on emission localization and timing will be small. However, if we systematically over- or underestimate OH abundance, then this can cause a significant, possibly interannually varying bias in our CO emission estimates. To test the sensitivity of our results to the OH fields used, we performed inversions for 2010 and 2011 with daily OH fields from the reanalysis product released by the Copernicus Atmosphere Monitoring Service (CAMS) (Inness et al., 2019).

In inversions with CAMS OH fields, we find lower biomass burning emissions by 35 Tg and 19 Tg in 2010 and 2011 respectively (purple bars in Fig. 4.S8). This is one of the largest sensitivities in our inverse system with significant interannual variability. It is driven by the large difference between the two different OH fields we have tested: in some regions CAMS OH is lower than our default OH fields by a factor 100 (Fig. 4.S9). This difference is largely consistent between months and between years.

The gap in OH fields over remote forests is a known attribute of OH fields derived in some

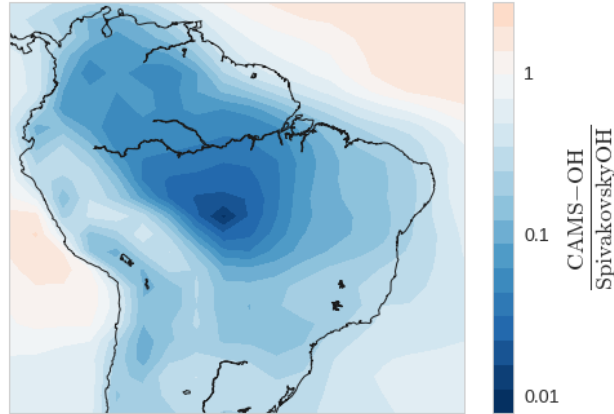


Figure 4.S9: The ratio between OH fields from the CAMS reanalysis product for 2010, and climatological OH fields from Spivakovsky et al. (2000), scaled by 0.92. OH fields were first averaged over our April–December inversion window and over the lowest 10 model layers, corresponding to approximately 500 hPa, or 5 km. The color scale is logarithmic.

full-chemistry models and it is likely related to missing recycling mechanisms for OH. In-situ measurements of OH over remote forests have revealed higher OH concentrations, even under low- NO_x conditions (Lelieveld et al., 2008). One explanation is that under these pristine conditions, isoprene oxidation can sustain a high OH recycling efficiency (e.g. Lelieveld et al., 2008; Taraborrelli et al., 2012), although more recent work has suggested that natural NO_x emissions have been underestimated over these regions (Wells et al., 2020). Whatever the driving mechanism, the gap in OH over the Amazon in CAMS-OH could explain why, in the CAMS reanalysis for CO (Flemming et al., 2017), GFAS emissions do not result in an underestimate of CO over the Amazon: too-low OH and too-low fire emissions cancel out.

Given these considerations, we deem our climatological OH fields more realistic over the Amazon region than the OH fields from the CAMS reanalysis. The large sensitivity of derived emissions to the OH field used is likely a reflection of the extreme difference between the OH fields we have tested, rather than a reflection of a large intrinsic sensitivity in the inverse system. We note that OH variations have a relatively diffuse, slow impact on CO, which means that its contribution to the sharp dry season peak in atmospheric CO abundance is limited. However, loss to OH does contribute substantially to the Amazon CO budget, and efforts to understand and constrain OH over remote regions (e.g. Fu et al., 2019; Nölscher et al., 2016) will translate to better biomass burning estimates. For example, in Nechita-Banda et al. (2018) it was found that emissions from fires reduced OH concentrations, and including this effect in the inversion reduced the estimated CO fire emissions by $\sim 7\%$. We do note that the fire source in their study, which was focused on Indonesia, included large peat fires that emit relatively low amounts of NO_x , so that the impact on OH concentrations is relatively high, compared to savanna or forest fires.




Figure 4.S10: Total CO emissions from different source categories summed over the $1^\circ \times 1^\circ$ South-American zoom domain, and over the April – December inversion period. Emission from our standard inversion are shown (i.e. GFAS fire emissions optimized with MOPITT CO columns), as well as emissions from a global MOPITT inversion presented in Zheng et al. (2019). Emission categories from Zheng et al. (2019) were merged to obtain emission categories comparable to ours.

4.S2.4 Comparison to Zheng et al. (2019)

Finally, we have compared our derived CO emissions to those derived by Zheng et al. (2019). Their emission estimate, covering 2000–2017, was derived in an inversion that also used MOPITT-TIR CO columns. However, that is the only shared aspect of our two inverse set-ups. Their inversion assimilated MOPITT CO columns globally, which will result in different boundary conditions for the Amazon domain than assimilation of surface observations. Additionally, their inversion was performed at a different spatial resolution (3.75° longitude by 1.9° latitude, with 39 vertical layers), in a different transport model (LMDz-SACS (Pison et al., 2009)). Prior OH fields used in Zheng et al. (2019) were the same as ours, but theirs were optimized with methyl chloroform observations. Additionally, they used biogenic CO emissions from the MEGANv2.1 inventory, which, as discussed in Supplement 4.S2.2, differ significantly from those used in our inversions. Further details on their inverse set-up are provided in Table 1 of Zheng et al. (2019). We limit our comparison to emissions derived in Inversion 1, as described in Zheng et al. (2019), in which satellite observations of formaldehyde and of methane were not assimilated.

We find that biomass burning emissions derived in Zheng et al. (2019) are comparable to ours, both in absolute magnitude and in interannual variability (red bars in Figure 4.S10), with an average annual total difference of -3.0 ± 6.7 TgCO/year (one standard deviation). This difference is small compared to interannual variability, and compared to individual sensitivities discussed in previous supplements. Notably, their emission estimates of CO from biomass burning are also systematically higher than those from GFAS (and FINN). The large, systematic difference between the natural emissions in their work and in ours does not seem to significantly affect the biomass burning emission estimates.

The excellent agreement between these two independent estimates provides much confidence to our final emission estimates. When many components in the inverse system are perturbed at once, rather than one-by-one, the individual impacts on biomass burning




Figure 4.S11: Simulated and observed monthly mean MOPITT CO columns over the Santarém aircraft site. Columns are averaged over a 1° by 1° area centered on Santarém (56°E ; 2.8°S). CO columns sampled from the prior (green) and from the posterior (orange) simulation are shown.

emission estimates either cancel out, or turn out to be smaller than the extreme cases we have tested in previous supplements.

4.S2.5 Conclusions on the uncertainty in our biomass burning emission estimate

Through these sensitivity tests, we have mapped the uncertainty in the assumptions made in our standard inversion on the final emission estimate. We conservatively estimate the uncertainty in the interannual variability of the MOPITT-derived CO emissions of biomass burning at 10–15 TgCO/year, and the systematic uncertainty at 30 TgCO/year. We consider this uncertainty estimate conservative, since the integrated comparison with Zheng et al. (2019) suggests an uncertainty of 7 TgCO/year (Supplement 4.S2.4). Our final uncertainty estimate is dominated by secondary production of CO from NMVOCs (Supplement 4.S2.2). Uncertainty in the MOPITT satellite data also contributes, but its effect is difficult to quantify in the absence of a comparable, long-term consistent satellite record for CO (Sup 4.S1.2). Another large budget term is loss of CO to oxidation by OH, but due to the slow effect of OH variations we deem this uncertainty of secondary importance, unless vastly different OH fields are used (Supplement 4.S2.3). We have also shown excellent skill in regional (e.g. state-level) attribution of fire emissions that is largely independent of the prior emission distribution (Supplement 4.S2.1). In short, we consider the longterm trends and interannual variability in CO emissions from biomass burning that we have presented in the main text robust features of our analysis.

4.S3 CO columns over the Santarém aircraft site

Out of the five sites for which aircraft profiles are compared to model results, we only found substantial differences between observed and TM5-simulated aircraft profiles at the Santarém aircraft site. Specifically, aircraft profiles sampled in a simulation with

MOPITT-optimized emissions were systematically too low (Figure 3 in the main text). In contrast, MOPITT CO columns over Santarém are well reproduced (Figure 4.S11). This effect is similar to what we observed at the NOAA surface site of Ragged Point, Barbados, where simulated CO mole fractions in the standard inversion were also lower than those observed (Supplement 4.S1.3). This could point to a systematic bias in the MOPITT-observed CO columns. However, surface observations, aircraft profiles and MOPITT CO columns all have different vertical sensitivities and therefore systematic uncertainties in the vertical transport of TM5 also affect this comparison. Again, we point out the potential value of an independent, long-term consistent record of satellite-observed CO columns for an uncertainty analysis of MOPITT CO columns, which is not yet available (Supplement 4.S1.2).

Chapter 5

Synthesis

5.1 Introduction

In this chapter, we synthesize the lessons and discoveries of previous chapters, in particular with reference to the research questions posed in Chapter 1.5. First, we consider how our research has affected constraints on the atmospheric oxidative capacity, as well as the best way to proceed. Next, we discuss our research into Amazonian fires and the possible ways in which our findings can be extended to an atmospheric monitoring system of fire emissions.

5.2 Improving constraints on the atmospheric oxidative capacity

5.2.1 Methyl chloroform in a two-box model

In our first study, we incorporated information from the 3D transport model TM5 in a two-box model inversion of MCF to constrain OH. The objective was to extend on two previous studies that used a similar box-model approach (Turner et al., 2017; Rigby et al., 2017). The results from the two-box model inversions can be viewed from three perspectives.

Firstly, as a reflection on atmospheric (two-)box models in general. In Chapter 3 we quantified potential biases in two-box models from full 3D model results. The extreme amplitude of the derived two-box model biases are mostly specific to MCF, because the atmospheric distribution of MCF has varied strongly over the 1994-2014 study period. For CH_4 , which was approximately in steady-state in our simulations, we find comparatively small biases. For example, the modest variations in the interhemispheric exchange rate of CH_4 in Figure 2.2 are unlikely to fundamentally change our understanding of the CH_4 budget, compared to an assumed constant exchange rate. This shows that, predictably, two-box model biases are most important for gases that vary significantly in space and time, either due to emission changes or due to a short lifetime.

Two-box model biases remain, however, conceptually important for any trace gas. For instance, superficially, an emission distribution only enters the two-box model as a distribution between the two hemispheres. In reality, variations in the emission distribution within and between hemispheres will also affect the interhemispheric exchange rate and other two-box model parameters. In addition, while we find that the global tropospheric burden is well represented by the NOAA surface network, we also find that an interhemispheric gradient derived directly from surface observations is significantly lower than the true tropospheric interhemispheric gradient (around 25% for steady-state CH_4 ; right panel in Figure 2.3). The effect of such systematic two-box model biases can change the best estimate for the tropospheric lifetime of CH_4 and for absolute CH_4 emissions by up

to 5%, although, in our study, the impact of the four identified biases on the CH₄ lifetime partly cancel (e.g. see Table 2.3). The potential of these hidden uncertainties requires consideration in the context of a specific box model application.

Our findings can also be viewed as a two-box model perspective on a 3D transport model. By reducing the extensive 3D model output to a few insightful quantities, important lessons about large-scale transport can be derived. For example, we have identified a significant and rather large ($+0.35 \pm 0.05\%/year$) multi-decadal increase in the inter-hemispheric exchange rate. Sensitivity runs revealed that this trend is driven by input meteorological fields from the Era-Interim reanalysis (Dee et al., 2011). It would be interesting to repeat the two-box parametrization exercise for other 3D transport models and for other sources of meteorological input, as the variations in such a comparison can help understand this positive trend in relation to transport uncertainties. This perspective could be further explored in a model-intercomparison project (e.g. Krol et al., 2018), which could offer valuable insights.

The third and final perspective relates specifically to the two research questions that we started out from:

1. How do the simplifying assumptions inherent to a two-box model affect the best estimate of interannual OH variations derived from MCF?

Significant adjustments in two-box model parameters driven by redistribution of MCF did affect the best estimate we derived for interannual OH variations. Most importantly, OH variations derived in the standard two-box inversion that did not include bias corrections, showed a near-zero trend for OH over the 1994-2014. Correcting for any individual bias did not strongly affect the derived OH trend, but an inversion that corrected all four biases did show a positive OH trend (see Table 2.3). This is an example in which the bias corrections drive the solution in the same direction, so that the combined effect exceeds that of individual biases. Despite the shift in trend and other adjustments, we found that the solutions of the standard and of the bias-corrected two-box model inversions remain statistically consistent. This, in turn, shows that uncertainties in a two-box model inversion remain large, and firm conclusions about OH variability and trends are difficult to draw.

2. Can we reduce uncertainties in OH variability derived in a two-box model inversion by accounting for transport effects with a 3D transport model?

A two-box model inversion in which the interhemispheric exchange rate is assumed to be known, because it is based on results from a full 3D model, results in a solution with less uncertainty than an inversion in which the exchange rate is varied or co-optimized. The same goes for the other two-box model parameters that we parametrized from 3D model output. However, in our two-box inversions we have diagnosed that the effect of fixing and adjusting these parameters on the uncertainty in the final solution is small.

The uncertainties that remain, mainly in the MCF emission magnitude and distribution, and in MCF observations, still drive significant uncertainties in the posterior solution. Turner et al. (2017) and Rigby et al. (2017) stated in their box model studies that OH variability could have contributed significantly to growth rate variations of atmospheric CH₄ over the 2000-2010 period, but it could not be excluded that OH had not contributed at all. This is consistent with our results. Despite the small effect of the identified biases on this aspect of our study's conclusion, we emphasize that box model studies that do not account for 3D transport effects are susceptible to the types of biases that we have identified.

We recognise that uncertainties in MCF emissions and observations are inherent limitations to any MCF-OH inversion, also in full 3D inversions. However, in a 3D transport model, observational uncertainties change, because individual observations per site replace hemispheric averages. Moreover, the distinction between OH and emissions possibly becomes clearer when the tropical maximum of OH is explicitly resolved in the model. By accounting for detailed transport, and for OH-oxidation between emission locations and observations, we hope to reduce uncertainties in OH. These were the key motivations for the 3D transport model inversions presented in Chapter 3.

5.2.2 Methyl chloroform in a 3D transport model inversion

Setting up a 3D model inversion of two decades of MCF observations to constrain the atmospheric oxidative capacity was challenging, but, in the end, rewarding. Initially, we performed an inversion with a global, monthly scaling factor for OH. However, because we could not reproduce observed intrahemispheric gradients of MCF, we decided to increase the degrees of freedom in the state and to optimize OH in 45 latitudinal bands instead. Even with this addition, it remained difficult to fit the intrahemispheric gradients of MCF. After numerous attempts with different settings in the inverse modelling framework, we concluded that the only other adjustment that significantly affected intrahemispheric MCF gradients was the addition of a high-latitude ocean source of MCF, as hypothesized in Wennberg et al. (2004).

In hindsight, this finding seems predictable. However, we note that MCF has been included 3D transport model simulations that did not include an ocean source, and that in these studies the intrahemispheric biases were not diagnosed (e.g. Patra et al., 2011). The main reason that we identified the biases was because our inverse system started to make adjustments that we consider physically unrealistic. This signifies the value of a 3D transport model inversion. If a problem is perfectly understood, then the important parameters can be selected and the simplest model that includes all these parameters can be set up. However, such a set-up is carefully curated to existing understanding and is therefore unlikely to reveal gaps in this understanding. A 3D transport model inversion can be complex and slow, but it also includes far fewer implicit prior assumptions than simple

models and is therefore more likely to provide surprising results. Moreover, the process we went through to reach our final inversion shows that it is worthwhile to carefully set up a statistical framework: first starting from simple prior assumptions and only gradually increasing the complexity, so that the inverse system remains understandable.

In addition to the interesting evidence for an ocean source of MCF, we addressed the question:

3. What constraints on OH can we derive in a 3D transport model inversion of MCF?

We derived a twenty-year timeseries of variations in the atmospheric oxidative capacity that fit observations of MCF at most NOAA surface sites within 1% of the observed mole fraction, especially when we considered monthly mean mole fractions and removed systematic biases. We did not retrieve posterior error estimates, because the inverse set-up was non-linear. Instead, we attempted to assess the robustness of the results by performing a set of sensitivity inversions. While it remains difficult to attribute the derived variations in global mean OH concentrations to specific atmospheric processes, we observe a significant negative correlation with ENSO ($r=-0.47$; $p=0.05$). This correlation, which is much weaker or even absent in inversions using box models, indicates that accounting for variable transport in 3D transport models provides a more credible relation between global OH abundance and large-scale climatic factors that affect this abundance. Examples of these factors are emissions from large-scale biomass burning, lightning, and large-scale perturbation of the water vapour distribution.

In Chapter 3, we have extensively discussed the convergence problems that hampered our inverse system. For example, three ten-year inversions that included different OH and MCF emission distributions converged to a similar posterior solution. This would indicate limited sensitivity of our results to prior assumptions for these distributions. However, the three twenty-year inversions did not converge to the same solution, because they reached different degrees of convergence.

Therefore, we conclude that we have derived a twenty-year timeseries for variations in the atmospheric oxidative capacity that is consistent with the MCF record as simulated in a 3D transport model. However, based on our experience, improvements in the inverse approach are possible, which we discuss in the next section.

5.2.3 Future use of methyl chloroform

In Chapters 2 and 3, we have closely investigated the constraints on OH that are contained in surface and aircraft observations of methyl chloroform. Atmospheric mole fractions of MCF continue to decline and, as a result, the potential for sustained emissions, the ocean flux, or measurement issues to interfere with the constraints that MCF observations place on OH grows. A valid concern is then whether any follow-up study can expect to learn more from methyl chloroform than we have.

In Chapter 3 we have outlined key points that could be implemented in a follow-up study. Most importantly, due consideration should be given to the possibility of a high-latitude ocean source of MCF. In our inversions, we found that separating anthropogenic MCF emissions from the influence of OH already required a careful weighting of the cost function. The spatial distribution of the hypothesized MCF ocean source is much more similar to the OH distribution than to the distribution of anthropogenic MCF emissions. Therefore, we remain pessimistic about the prospect of co-optimizing an ocean source in an MCF inversion, and, fundamentally, about the ability of the NOAA surface network to separate variations in OH from the MCF ocean flux. However, we did not present a definitive proof that the two cannot be separated, and a thorough follow-up study might reach a more optimistic conclusion.

To separate the three components (emissions, ocean flux and OH), the follow-up study would need to be performed in a 3D transport model. Careful consideration should be given to the number of degrees of freedom given to the state relative to the observations. We suggested in Chapter 3 that perhaps it is better to use monthly mean mole fractions per site in the optimization, instead of individual observations. This suggestion was based on the limited ability of our inverse system to fit the observed sub-monthly variations in MCF mole fractions. Similarly, the degrees of freedom in the state need to be considered. It is unlikely that the influence of OH and of the ocean flux can be separated if both are optimized in 45 latitudinal bands, as we have done for OH. These detailed technical choices will determine the outcome of such an inversion. Given that we suggest to reduce the degrees of freedoms in the inverse system, other inverse approaches than 4DVAR become viable. For example, in Bousquet et al. (2005) the sensitivity of observations to each state element was calculated in forward simulations, which is only computationally viable with a limited number of degrees of freedom in the state. This approach could result in better convergence than the non-linear 4DVAR approach.

Would we recommend such a follow-up study? Based on our findings, we estimate that the return-on-investment will be small. Historically, the primary application of MCF-derived variations in the atmospheric oxidative capacity has been in the atmospheric CH_4 budget. A global CH_4 inversion will not benefit much from a new MCF study that derives slightly modified OH variations, if, for example, this MCF study still cannot exclude an interannually constant timeseries of OH. While MCF has provided valuable top-down constraints on OH, we deem it unlikely that a timeseries of annual mean OH concentrations with uncertainties smaller than a few percent can be derived from only MCF surface observations. However, a comprehensive 3D model inversion that incorporates these recommendations would still be a valuable exercise to give credit to decades of careful measurement efforts. Such an inversion could be complemented with measurements of MCF hydrolysis, to derive stronger prior constraints on the MCF ocean flux.

5.2.4 The use of other tracers to constrain OH

MCF is the most ideal tracer for OH that we are likely to find. The only better option would be emission of a trace gas with a known release rate: an idea that was put forward on multiple occasions (e.g. Jöckel et al., 2003). However, we consider it unlikely that such a synthetic tracer will ever be released, if only because it hasn't been released yet. Therefore, to further improve constraints on the atmospheric oxidative capacity, we see two possibilities. Firstly, multiple tracers could be coupled together in an inverse modelling framework to constrain OH: in this case, MCF could still serve a valuable role as one of these tracers. Secondly, if satellite observations were used instead of only surface observations, the sheer volume of available data opens up new possibilities. We will discuss these prospects one by one.

Alternatives to MCF have been proposed in Huang & Prinn (2002), and one variation of a multi-tracer framework was put forward in Liang et al. (2017). They suggested that OH could be constrained in a two-box model inversion that includes multiple anthropogenically produced tracers. The underlying premise of their approach was similar to the two-box model studies of MCF: leverage the interhemispheric gradient of a trace gas to separate emission from OH variations. Additionally, tracer-specific growth rate variations are more likely to be driven by emissions, while growth rate variations found in all tracers suggest an influence from variations in the shared OH sink. They found that different tracers are sensitive to different aspects of the OH distribution and in such an integrated framework these sensitivities can complement each other.

It is important to note that MCF has been a much better tracer for OH than the other suggested tracers are likely to be, because it is the only gas for which emissions play a minor role in its atmospheric budget. Therefore, key to this approach would be separation of emissions from OH. This requires a very careful inverse set-up. For this reason, we think that, as suggested in Liang et al. (2017), a simple two-box model inversion would work better than a 3D model inversion, because the computational efficiency of a two-box model allows it to be incorporated in a more comprehensive statistical framework. Our findings in Chapter 2 suggest that it is important to account for tracer-dependent transport effects as much as possible. This could be through a two-box parametrization of a 3D transport model as in Chapter 2, but other options are available. For example, in a recent study (Stell et al., 2020) the global CH₄ budget was explored through Gaussian process emulators that ingest the same input parameters as a 3D transport model, with hemispherically averaged CH₄ mole fractions (and $\delta^{13}\text{CH}_4$) as output. Such an emulator is more computationally efficient than a 3D transport model, but, unlike (untuned) box models, the emulator output implicitly accounts for transport effects. Regardless of the exact set-up, a valuable addition to such inversions would be to critically evaluate the derived solution in a forward simulation of a 3D transport model.

A different tracer for OH that has been put forward is ¹⁴CO (e.g. Volz et al., 1981; Jöckel

et al., 2002). ^{14}CO is formed through oxidation of the cosmogenically generated ^{14}C in the stratosphere and upper troposphere (MacKay et al., 1963). ^{14}CO has a tropospheric lifetime of approximately two months, mostly due to reaction with OH. The combination of a source region at high altitudes, a relatively short atmospheric lifetime and a pressure-dependent reaction rate with OH gives ^{14}CO a unique spatio-temporal sensitivity to atmospheric OH (Krol et al., 2008). Due to its short lifetime, ^{14}CO is sensitive mainly to regional variations in OH. The primary driver of variability in the cosmogenic source of ^{14}CO is the 11-year solar cycle. This means that, simplistically, variability in ^{14}CO on an 11-year timescale is driven mainly by longterm variations in OH, because the source magnitude repeats every 11 years. In other words, ^{14}CO observations can provide constraints on long-term, but regional variations in OH. Currently, extended timeseries of ^{14}CO observations are only available for very few surface sites. Strong constraints on large-scale OH variations would require expanded atmospheric monitoring of ^{14}CO , but could provide valuable new insights into the spatial distribution of OH and variations therein.

Next we discuss promising research into the use of satellite observations to constrain large-scale variations in the OH concentrations.

A satellite-based approach that directly integrates constraints on OH into the CH_4 budget is the use of satellite-observed CH_4 columns to constrain not only CH_4 emissions, but also OH variations. In Zhang et al. (2018) an observing system simulation experiment (OSSE) was performed, in which two synthetic satellite products with different vertical sensitivities were assimilated. The difference in vertical sensitivities combined with the high spatial resolution of the two products creates potential to separate variations in CH_4 columns driven by CH_4 emissions, mostly located at the surface, from variations in OH, mostly located in the middle troposphere. In Maasakkers et al. (2019) an inverse framework is presented in which existing satellite data of CH_4 is assimilated and one global, annual scaling factor for OH is optimized along with grid-scale CH_4 emissions. They find skill to independently estimate absolute OH and absolute CH_4 emissions, but posterior uncertainties on OH variations remain significant and OH trend estimation was not possible. Although the use of satellite instruments with different vertical sensitivities (Zhang et al., 2018) sounds promising, co-location of observations and cross-calibration of the instruments are major challenges. Moreover, the optimistic conclusion of both studies (Zhang et al., 2018; Maasakkers et al., 2019) is based on optimizing global scaling factors of OH using a fixed OH distribution. Zhang et al. (2018) showed that optimizing the OH distribution is a much harder problem. Future studies will have to demonstrate that this promising approach can be put into practice to independently estimate both CH_4 emissions and variations in the global OH distribution.

Constraints on OH can also be derived in full-chemistry simulations of 3D transport models that assimilate satellite data of those tracers that impact the atmospheric OH

budget (Miyazaki et al., 2012; Flemming et al., 2017). Traditionally, the accuracy of OH estimates derived in full-chemistry simulations has been limited by uncertainties in the emissions of the many trace gases that affect the atmospheric abundance and distribution of OH, for example CO. By assimilating large quantities of satellite data, the atmospheric state in the simulation is nudged closer to the true atmosphere, which likely reduces the impact of emission errors. The potential impact of this modification was shown in a study that assimilated only satellite-observed CO, as a decreasing CO burden drove a strong, positive trend in global mean OH concentrations over the 2002–2013 period (Gaubert et al., 2017). Such an approach remains dependent on correct representation of all relevant gases: a trend in OH driven by CO can be offset by variations in other gases. In addition, details of the complex chemical pathways can never be fully implemented in large-scale models, and rate constants and photolysis frequencies also have their uncertainties. In Chapter 4 we have shown that over South-America the data assimilation system of the Copernicus Atmospheric Monitoring Service (CAMS) (Flemming et al., 2017) produces OH concentration fields that are likely biased too low. Since, according to our estimates, their estimate for biomass burning is also biased low, these compensating errors can still result in a good match with the assimilated CO columns. We discuss this issue more thoroughly in connection to estimating biomass burning emissions of CO at the end of this synthesis. Conclusively, we consider that assimilation of satellite data in full-chemistry simulations provides a significant step forward, but independent top-down constraints from other sources remain highly relevant.

A study that used satellite and aircraft observations of formaldehyde (HCHO) to constrain OH in the remote troposphere has shown promising results (Wolfe et al., 2019). HCHO is formed as an intermediate in the oxidation of hydrocarbons and has an atmospheric lifetime of a few hours. The distribution of HCHO in the remote troposphere is mainly determined by the distribution of OH, because the main source of HCHO in remote locations is oxidation of well-mixed CH₄ by OH. Therefore, the HCHO distribution can be used as a proxy for the OH distribution. The focus of Wolfe et al. (2019) was on constraining the large-scale distribution and seasonality of OH. Currently, the method hinges on the limited availability of aircraft data that is required to validate the satellite retrievals. However, since the OMI satellite retrieval that was used has been available from 2004 (González Abad et al., 2015), there is also potential to explore the constraints that HCHO places on OH over longer time periods.

In (Wells et al., 2020), satellite retrievals of HCHO and isoprene were combined to constrain both isoprene emissions and atmospheric oxidation. They found that, under low-NO_x conditions, OH is not recycled efficiently. In Lelieveld et al. (2008), OH concentrations over the pristine Amazon rainforest were reported that were significantly higher than those simulated in full-chemistry models. In Lelieveld et al. (2008) and in a number of follow-up studies (e.g. Taraborrelli et al., 2012; Fuchs et al., 2013), these elevated OH concentrations were attributed to efficient recycling of OH during isoprene oxidation,

even under pristine, low- NO_x conditions. In contradiction to this explanation, Wells et al. (2020) argue that OH concentrations over tropical forests are indeed elevated, but that these concentrations are more likely sustained by NO_x emissions, rather than by missing recycling mechanisms under low- NO_x conditions. Such contradicting findings are likely to result in a quickly improved understanding of atmospheric oxidation, which will in turn improve our ability to interpret the atmospheric budgets of many trace gases of interest. For example, this particular issue relates directly to OH fields over the Amazon that we use to constrain fire emissions of CO in Chapter 4.

In summary, multiple independent approaches to constrain the atmospheric oxidative capacity are being developed. The recent release of the isoprene retrieval (Fu et al., 2019) exemplifies the potential of new atmospheric monitoring tools to offer a fresh perspective on atmospheric oxidation. While each of the methods discussed in this section still has strong limitations, parallel development will likely result in a better understanding of large-scale OH variations and the patterns therein. These promising new developments can complement what we have learned from the benchmark MCF-centered approach.

5.3 Constraints from CO on Amazonian fire activity

In Chapter 4 we applied the TM5-4DVAR inverse framework to the Amazon domain to answer the question:

3. What can a 3D transport model inversion of CO teach us about fire activity and deforestation in the Amazon?

The high detail contained in satellite data, combined with emission allocation provided by the realistic transport in TM5, resulted in a robust, sixteen year timeseries of CO emissions from Amazon fires. At the level of individual Brazilian states, we found that interannual variations in fire emissions were closely correlated with local anomalies in soil moisture. Such a link can be expected if the derived fire emissions are realistic, since drought is a large driver of Amazon fires, but the strength of the correlations and the consistency between states are noteworthy. In some states, for example Pará, CO emissions showed a longterm decrease, superimposed on the interannual variations. Zooming out over the wider Amazon basin, we found this decrease to be particularly strong over forested regions. Moreover, the decrease in five-year averaged CO emissions was quantitatively similar to a decrease in deforestation rates reported by Instituto Nacional de Pesquisas Espaciais (INPE). The strong control of soil moisture on interannual variability in emissions and the close connection between deforestation rates and a long-term trend in emissions signify that natural and anthropogenic influences each play an important role in the Amazon domain.

An interesting question concerning the results presented in Chapter 4 is: what exactly is the added value of our inverse system to the input products? Figure 4, panel 6 in Chapter 4 shows that basin-wide fire emissions can be estimated with high accuracy from basin-wide anomalies in MOPITT CO columns. Moreover, the longterm decrease in CO emissions is already largely contained in the GFAS emission inventory, as is most of the interannual variability at state-level and the corresponding correlation with local soil moisture anomalies. Therefore, up- and downscaling of GFAS emissions with basin-wide MOPITT CO columns could have resulted in similar emissions, compared to the emissions derived by our involved inversion procedure.

The answer is two-fold. Firstly, if we do not start from GFAS fire emissions, but from a different inventory or even from a climatological prior, we retrieve largely the same trends and variability. In other words, we do not retrieve similar trends and variability as GFAS because we start from GFAS, but rather because the independent observational constraints, as interpreted by our inverse system, are consistent with the GFAS prior, thereby providing independent verification of the GFAS inventory. Secondly, in anomalous cases, we derive corrections to the GFAS prior, for example the November–December emissions in 2015. Similarly, we have shown that the adjustments to prior emissions can vary significantly between landcover types. These examples show that the inverse system facilitates a more nuanced and comprehensive comparison between fire emission inventories and satellite data.

The systematic underestimate of CO emissions in fire emission inventories such as GFAS that we identify in Chapter 4 deserves further elaboration. In Van Leeuwen et al. (2013), the potential uncertainty in CO emission factors for dry matter burned was explored, and they found that estimated CO fire emissions vary significantly between scenarios where assumptions for the CO emission factors are varied. However, in general, their findings do not align with ours. In particular, over South-America, they report that scenarios with a variety of alternative CO emission factors result in more emissions from forested regions, but lower emissions from savanna, and in lower CO mixing ratios at the heights where the MOPITT instrument is most sensitive: both signals are opposite of what we need to explain the difference between inventories and our inverse estimate. The typical amplitude of the adjustments they find over South-America is small ($\sim 3\%$), compared to the required increase that we have identified in particular over the savanna ($\sim 65\%$), although their MCE-SEASON scenario, which is referred to as most promising, produced $\sim 11\%$ higher emissions over South-America, relative to the reference GFED3 set-up. The results from Van Leeuwen et al. (2013) are based on the GFED3 inventory, whereas our standard inversions start from the GFAS inventory, but we find that the differences between emission inventories are generally smaller than the difference between inventories and our posterior estimate. In conclusion, further interpretation of the large gap between inventories and our posterior emission estimate requires additional research, for example a cross-comparison of a wider variety of emission inventories and inverse estimates, and

a revisitation of emission factors, particularly over the Cerrado.

Given the robustness of derived emissions, we are highly optimistic about the prospect of an operational system that estimates CO fire emissions, so that these emissions can be used as a fire proxy alongside other products, such as Fire Radiative Power, or burned area. Here we zoom in on a few of the components of the inverse system that can benefit from additional constraints.

The largest sensitivity in our inverse framework is the production of CO from oxidation of Non-Methane Hydrocarbons (NMHCs). In our inversions, we used a climatological prior for CO production that was not optimized. In reality, production from NMHCs can vary interannually, and, since NMHCs are also emitted in fires and by vegetation, this variability can correlate with direct CO fire emissions. Therefore, an operational framework should incorporate a better CO production estimate. One option is to use a more advanced prior product, for example based on the Model of Emissions of Gases and Aerosols from Nature (MEGAN, Sindelarova et al. (2014)). Additionally, production from NMHCs could be co-optimized in the inverse system. This would be difficult to do based on only CO observations, but, for example, satellite data of HCHO and isoprene could prove useful in separating the two budget terms. For HCHO, this was done already in Zheng et al. (2019). Correct integration of different satellite products in one inverse system is not straightforward, and care should be taken that the information contained in the products is projected correctly on CO emission and production adjustments. In principle, such an approach could result in an estimate of CO fire emissions that is less reliant on prior assumptions for CO production.

Another uncertainty in the inverse system is the choice of OH concentration fields. We found a large sensitivity to this component in inversions based on the CAMS-OH fields, but this we attributed to unrealistically low OH concentrations over the Amazon in those fields. A better test of this sensitivity was presented in Nechita-Banda et al. (2018), where OH fields from a full-chemistry simulation of TM5 were included to quantify its effect on derived CO fire emissions over Indonesia. They reported that OH concentrations were lower than climatology during the large fires of 2015, and accounting for this decrease resulted in significantly lower derived CO emissions, by $\sim 7\%$. This sensitivity is likely an overestimate for the effect of Amazon fires on OH, because this particular event included large peat fires, which emit a composition of gases (e.g. relatively low amounts of NO_x) that is more likely to suppress OH concentrations than savanna or forest fires. However, the result does signify that better constraints on OH can help CO fire estimates. In this thesis, much attention has been devoted to improving constraints on atmospheric oxidation as a global, integrated quantity, but MCF surface observations do not strongly constrain OH variations on the smaller scale of the Amazon domain. However, the work on combining isoprene and HCHO retrievals to constrain OH concentrations over the Amazon (Wells et al., 2020), as discussed in Section 5.2.4, is highly relevant here. Unfortunately,

the isoprene retrievals are not available for the entire 2003–2018 period, but they could help to establish a more accurate OH climatology and these retrievals could be employed in future inversions that investigate more recent years. The availability of such new satellite products additionally makes it more viable to simulate the OH budget in full-chemistry simulations that assimilate satellite data (e.g. Miyazaki et al., 2012; Flemming et al., 2017) to optimize the atmospheric state of a variety of trace gases, including CO.

On the observational side, we have found the MOPITT-TIR CO product to be an excellent fit to our research question, but other CO satellite products are available. In Chapter 4, Supplement 1.2, we explore the use of CO columns from the Infrared Atmospheric Sounding Interferometer (IASI) over the 2010–2018 period. We found this product less suitable than MOPITT, because of inconsistencies in the reprocessing of spectral data. However, the IASI team is working on a consistent reprocessing, and, when made available, such a product can help better quantify uncertainties in an emission estimate based on MOPITT. While the CO retrievals from these two products are best-established, the newer Tropomi instrument currently retrieves CO columns with a higher spatial resolution. In the Tropomi CO product the 2.3 μm waveband is used: the same waveband is exploited in the MOPITT-NIR product and makes retrieved CO columns more sensitive to surface CO than TIR products. Initial work on Tropomi CO has shown promising results (Borsdorff et al., 2018), and in Fu et al. (2016), potential was shown for a combined Tropomi-CrIS CO product, which exploits the different vertical sensitivities of the two instruments to inform on the vertical CO distribution, in an approach very similar to what has been suggested for CH_4 (Zhang et al., 2018) (see also Section 5.2.4). The greatest value of the MOPITT-TIR product for our study, however, is its longterm consistency and availability. While improved resolution and complementary vertical sensitivities could prove helpful, we show in Chapter 4 that these are not strictly necessary to robustly estimate CO emissions over the Amazon. Finally, we note that CO aircraft profiles over the Amazon are an important source of independent data to validate our largely satellite-based approach, and we therefore consider the continued availability of such profiles of high value.

Output of our inverse system is an estimate of CO fire emissions. Emission of CO is an integrated quantity that is not equivalent to deforestation, fire incidence, or total carbon release, yet it provides information on all of these. The influence of deforestation is likely best visible as longterm, slow variations in CO emissions, since reported deforestation rates show much smaller variability than CO emissions. Consequently, interannual variations in CO emissions likely have other drivers than man-induced deforestation, such as drought occurrence. In Chapter 4, we show examples of both links, but it is difficult to interpret CO emissions more directly. Where exactly in our CO emissions is the information on understory fires that are harmful to forest resilience, but hard to observe from land-remote sensing? Does the underestimate in GFAS CO emissions translate to an underestimate in GFAS CO_2 emissions? Answers to such questions remain difficult to pin down, but a follow-up study that investigates our derived emissions in more detail

could help. For example, the answers to these questions are likely to be very different over remote, pristine rainforests (e.g. in Amazonas) than over savanna (e.g. in Tocantins), and future work could take a more focused approach than we have done. In this way, the CO emissions we have derived provide a new reference point for research into Amazon fires and deforestation, rather than definitive answers.

5.4 Final outlook

We have already discussed how the different research methods in this thesis can be extended upon. Despite the different research focus in Chapters 2–3 versus Chapter 4, similar challenges and potential new research approaches appear in both. Most notably, a growing fleet of satellite instruments, as well as increasingly longterm timeseries provided by existing instruments, provide a level of detail that facilitates new research approaches. Integration of different datastreams of a variety of tracers in one data assimilation system can potentially help more conclusively separate different budget terms. We have discussed examples of such systems, which has also revealed existing challenges.

Most generally, if a variety of datasets is included in a data assimilation system, it is essential, but often difficult, to understand the contribution of individual components to the final result. A first example is the inverse system presented in Yin et al. (2015) that was later used in Zheng et al. (2019). In their set-up, the oxidation chain of CH_4 , HCHO , CO and finally CO_2 is simulated, as well as MCF to provide additional constraints on OH . While this is an innovating and interesting approach, it is unclear to what degree each component informs the inverse system. This makes it difficult to compare their derived OH variations to ours, or to understand to what degree assimilation of HCHO satellite data helps separate natural production from fire emissions in the CO budget. Another example is the CAMS data assimilation system, in which satellite data for a variety of gases is assimilated, including MOPITT CO . While their results indicate that a simulation with GFAS CO emissions results, on average, in a good match with MOPITT CO columns over the Amazon, we find that errors in their OH distribution might bias this finding.

While setting up a multi-tracer data assimilation system is also a technical challenge, we consider that the main challenge of this approach would be to develop a comprehensive understanding of the results. What is the contribution of individual components? How can different satellite products with potentially different biases be combined in one system? How do different components of the state need to be weighted, and what is the best method for combining satellite data and surface observations, when these two information sources have such different qualities? These questions and more determine whether the system just reproduces assimilated observations, or is also able to help our understanding of the atmosphere.

The ultimate objective of inverse modelling is to provide the best interpretation of observations, in terms of quantities that are most relevant to us. The product of such an approach is not only a timeseries of global mean OH variations, or of CO fire emissions: equally important is the understanding that we have developed of the information content of observations regarding these quantities. This abstract objective is reflected in our leading research questions. Not: "*How has OH varied?*", but the more convoluted: "*What constraints on OH can we derive in a 3D transport model inversion of MCF?*". While future research that uses our best estimates will find sufficient information in the answer to the first question, future research that needs to interpret and improve on the quality of our estimates requires an accurate mapping of the uncertainties and considerations that we have discovered. Although such details have often ended up in the expansive supplements to Chapters 2–4, I hope the information is there for the next PhD student to explore.

References

- Alencar, A. A., Solórzano, L. A., & Nepstad, D. C. (2004). Modeling forest understory fires in an eastern Amazonian landscape. *Ecological Applications*, *14*, 139–149.
- Alexe, M. et al. (2015). Inverse modelling of CH₄ emissions for 2010–2011 using different satellite retrieval products from GOSAT and SCIAMACHY. *Atmospheric Chemistry and Physics*, *15*, 113–133.
- Aragão, L. E. O. C. et al. (2018). 21st Century drought-related fires counteract the decline of Amazon deforestation carbon emissions. *Nature communications*, *9*, 536.
- Asner, G. P., & Alencar, A. (2010). Drought impacts on the Amazon forest: the remote sensing perspective. *New phytologist*, *187*, 569–578.
- Atkinson, R. (1986). Kinetics and mechanisms of the gas-phase reactions of the hydroxyl radical with organic compounds under atmospheric conditions. *Chemical Reviews*, *86*, 69–201.
- Balch, J. K., Nepstad, D. C., Curran, L. M., Brando, P. M., Portela, O., Guilherme, P., Reuning-Scherer, J. D., & de Carvalho Jr, O. (2011). Size, species, and fire behavior predict tree and liana mortality from experimental burns in the Brazilian Amazon. *Forest Ecology and Management*, *261*, 68–77.
- Bândă, N., Krol, M., Noije, T., Weele, M., Williams, J. E., Sager, P. L., Niemeier, U., Thomason, L., & Röckmann, T. (2015). The effect of stratospheric sulfur from Mount Pinatubo on tropospheric oxidizing capacity and methane. *Journal of Geophysical Research: Atmospheres*, *120*, 1202–1220.
- Bândă, N., Krol, M., Van Weele, M., Van Noije, T., Le Sager, P., & Röckmann, T. (2016). Can we explain the observed methane variability after the Mount Pinatubo eruption? *Atmospheric chemistry and physics*, *16*, 195–214.
- Bartlett, K. B., & Harriss, R. C. (1993). Review and assessment of methane emissions from wetlands. *Chemosphere*, *26*, 261–320.
- Bergamaschi, P., Frankenberg, C., Meirink, J. F., Krol, M., Villani, M. G., Houweling, S., Dentener, F., Dlugokencky, E. J., Miller, J. B., Gatti, L. V., Engel, A., & Levin, I. (2009). Inverse modeling of global and regional CH₄ emissions using SCIAMACHY satellite retrievals. *Journal of Geophysical Research: Atmospheres*, *114*.

- Bergamaschi, P., Hein, R., Brenninkmeijer, C. A. M., & Crutzen, P. J. (2000). Inverse modeling of the global CO cycle: 2. Inversion of $^{13}\text{C}/^{12}\text{C}$ and $^{18}\text{O}/^{16}\text{O}$ isotope ratios. *Journal of Geophysical Research: Atmospheres*, 105, 1929–1945.
- Bergamaschi, P., Krol, M., Dentener, F., Vermeulen, A., Meinhardt, F., Graul, R., Ramonet, M., Peters, W., & Dlugokencky, E. J. (2005). Inverse modelling of national and European CH_4 emissions using the atmospheric zoom model TM5. *Atmospheric Chemistry and Physics*, 5, 2460.
- Bernstein, J. A. et al. (2008). The health effects of nonindustrial indoor air pollution. *Journal of Allergy and Clinical Immunology*, 121, 585–591.
- Borsdorff, T., Aan de Brugh, J., Hu, H., Aben, I., Hasekamp, O., & Landgraf, J. (2018). Measuring carbon monoxide with tropomi: First results and a comparison with ecmwf ifs analysis data. *Geophysical Research Letters*, 45, 2826–2832.
- Bouillon, M., Safieddine, S., Hadji-Lazaro, J., Whitburn, S., Clarisse, L., Doutriaux-Boucher, M., Coppens, D., August, T., Jacquette, E., & Clerbaux, C. (2020). Ten-year assessment of IASI radiance and temperature. *Remote Sensing*, 12, 2393.
- Bousquet, P. et al. (2006). Contribution of anthropogenic and natural sources to atmospheric methane variability. *Nature*, 443, 439.
- Bousquet, P., Hauglustaine, D. A., Peylin, P., Carouge, C., & Ciais, P. (2005). Two decades of OH variability as inferred by an inversion of atmospheric transport and chemistry of methyl chloroform. *Atmospheric Chemistry and Physics*, 5, 2635–2656.
- Bowman, D. M. J. S. et al. (2009). Fire in the earth system. *science*, 324, 481–484.
- Bowman, K. P. (2006). Transport of carbon monoxide from the tropics to the extratropics. *Journal of Geophysical Research: Atmospheres*, 111.
- Brando, P., Macedo, M., Silvério, D., Rattis, L., Paolucci, L., Alencar, A., Coe, M., & Amorim, C. (2020). Amazon wildfires: Scenes from a foreseeable disaster. *Flora*, (p. 151609).
- Brenninkmeijer, C. A. M., Manning, M. R., Lowe, D. C., Wallace, G., Sparks, R. J., & Volz-Thomas, A. (1992). Interhemispheric asymmetry in OH abundance inferred from measurements of atmospheric ^{14}CO . *Nature*, 356, 50.
- Brühl, C., & Crutzen, P. J. (1993). MPIC two-dimensional model. *NASA Ref. Publ*, 1292, 103–104.
- Burkholder, J. B., Sander, S. P., Abbatt, J. P. D., Barker, J. R., Cappa, C., Crounse, J. D., Dibble, T. S., Huie, R. E., Kolb, C. E., Kurylo, M. J., Percival, C. J., & Wine, P. H. (2020). *Chemical kinetics and photochemical data for use in atmospheric studies; evaluation number 19*. Technical Report, Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space.
- Burkholder, J. B., Sander, S. P., Abbatt, J. P. D., Barker, J. R., Huie, R. E., Kolb,

- C. E., Kurylo, M. J., Orkin, V. L., Wilmouth, D. M., & Wine, P. H. (2015). *Chemical kinetics and photochemical data for use in atmospheric studies: evaluation number 18*. Technical Report, Pasadena, CA: Jet Propulsion Laboratory, National Aeronautics and Space
- Butchart, N. (2014). The Brewer-Dobson circulation. *Reviews of geophysics*, 52, 157–184.
- CAMS reanalysis - OH (). <https://atmosphere.copernicus.eu/catalogue#/product/urn:x-wmo:md:int.ecmwf::copernicus:cams:prod:an:oh:pid443>. Online; accessed 23 Sep 2020.
- Chipperfield, M. P., & Liang, Q. e. a. (2013). Model estimates of lifetimes. In S. Reimann, M. K. W. Ko, P. A. Newman, & S. E. Strahan (Eds.), *SPARC, 2013: SPARC Report on the Lifetimes of Stratospheric Ozone-Depleting Substances, Their Replacements, and Related Species* chapter 5. WCRP-15/2013.
- Christensen, T. R., Johansson, T., Åkerman, H. J., Mastepanov, M., Malmer, N., Friberg, T., Crill, P., & Svensson, B. H. (2004). Thawing sub-arctic permafrost: Effects on vegetation and methane emissions. *Geophysical research letters*, 31.
- CIESIN, Columbia University (2018). Gridded Population of the World, Version 4 (GPWv4): Population Density, Revision 11.
- Clerbaux, C., Boynard, A., Clarisse, L., George, M., Hadji-Lazaro, J., Herbin, H., Hurtmans, D., Pommier, M., Razavi, A., Turquety, S., Wespes, C., & Coheur, P.-F. (2009). Monitoring of atmospheric composition using the thermal infrared IASI/MetOp sounder. *Atmospheric Chemistry and Physics*, 9, 6041–6054.
- Cochrane, M. A., & Barber, C. P. (2009). Climate change, human land use and future fires in the Amazon. *Global Change Biology*, 15, 601–612.
- Crutzen, P. J., & Andreae, M. O. (1990). Biomass burning in the tropics: Impact on atmospheric chemistry and biogeochemical cycles. *Science*, 250, 1669–1678.
- Crutzen, P. J., & Zimmermann, P. H. (1991). The changing photochemistry of the troposphere. *Tellus A: Dynamic Meteorology and Oceanography*, 43, 136–151.
- Daniel, J. S., & Solomon, S. (1998). On the climate forcing of carbon monoxide. *Journal of Geophysical Research: Atmospheres*, 103, 13249–13260.
- Davis, K. F. et al. (2020). Tropical forest loss enhanced by large-scale land acquisitions. *Nature Geoscience*, 13, 482–488.
- Dean, J. F. et al. (2018). Methane Feedbacks to the Global Climate System in a Warmer World. *Reviews of Geophysics*, 56, 207–250.
- Dee, D. P. et al. (2011). The ERA-Interim reanalysis: Configuration and performance of the data assimilation system. *Quarterly Journal of the royal meteorological society*, 137, 553–597.

- Deeter, M. N., Edwards, D. P., Francis, G. L., Gille, J. C., Mao, D., Martínez-Alonso, S., Worden, H. M., Ziskin, D., & Andreae, M. O. (2019). Radiance-based retrieval bias mitigation for the MOPITT instrument: the version 8 product. *Atmospheric Measurement Techniques*, 12, 4561–4580.
- Denning, A. S. et al. (1999). Three-dimensional transport and concentration of SF₆: A model intercomparison study (TransCom 2). *Tellus B: Chemical and Physical Meteorology*, 51, 266–297.
- Dentener, F., Peters, W., Krol, M., van Weele, M., Bergamaschi, P., & Lelieveld, J. (2003). Interannual variability and trend of CH₄ lifetime as a measure for OH changes in the 1979–1993 time period. *Journal of Geophysical Research: Atmospheres*, 108.
- Dlugokencky, E., Masarie, K., Lang, P., & Tans, P. (1998). Continuing decline in the growth rate of the atmospheric methane burden. *Nature*, 393, 447–450.
- Dlugokencky, E. J., Bruhwiler, L., White, J. W. C., Emmons, L. K., Novelli, P. C., Montzka, S. A., Masarie, K. A., Lang, P. M., Crotwell, A. M., Miller, J. B., & Gatti, L. V. (2009). Observational constraints on recent increases in the atmospheric CH₄ burden. *Geophysical Research Letters*, 36.
- Dlugokencky, E. J., Steele, L. P., Lang, P. M., & Masarie, K. A. (1994). The growth rate and distribution of atmospheric methane. *Journal of Geophysical Research: Atmospheres*, 99, 17021–17043.
- Donner, L. J. et al. (2011). The Dynamical Core, Physical Parameterizations, and Basic Simulation Characteristics of the Atmospheric Component AM3 of the GFDL Global Coupled Model CM3. *Journal of Climate*, 24, 3484–3519.
- Etiope, G., & Klusman, R. W. (2002). Geologic emissions of methane to the atmosphere. *Chemosphere*, 49, 777–789.
- Fanin, T., & van der Werf, G. R. (2015). Relationships between burned area, forest cover loss, and land cover change in the brazilian amazon based on satellite data. *Biogeosciences*, 12, 6033–6043.
- Fisher, M. (1995). Estimating the covariance matrices of analysis and forecast error in variational data assimilation. *ECMWF Tech. Mem.* 220, .
- Flemming, J. et al. (2017). The cams interim reanalysis of carbon monoxide, ozone and aerosol for 2003–2015. *Atmospheric Chemistry and Physics*, 17, 1945–1983.
- Francey, R. J., & Frederiksen, J. S. (2016). The 2009–2010 step in atmospheric CO₂ interhemispheric difference. *Biogeosciences*, 13, 873.
- Fu, D., Bowman, K. W., Worden, H. M., Natraj, V., Worden, J. R., Yu, S., Veefkind, P., Aben, I., Landgraf, J., Strow, L., & Han, Y. (2016). High-resolution tropospheric carbon monoxide profiles retrieved from CrIS and TROPOMI. *Atmospheric Measurement Techniques*, 9, 2567–2579.

- Fu, D., Millet, D. B., Wells, K. C., Payne, V. H., Yu, S., Guenther, A., & Eldering, A. (2019). Direct retrieval of isoprene from satellite-based infrared measurements. *Nature communications*, *10*, 1–12.
- Fuchs, H., Hofzumahaus, A., Rohrer, F., Bohn, B., Brauers, T., Dorn, H., Häseler, R., Holland, F., Kaminski, M., Li, X. et al. (2013). Experimental evidence for efficient hydroxyl radical regeneration in isoprene oxidation. *Nature Geoscience*, *6*, 1023–1026.
- Gatti, L. V. et al. (2014). Drought sensitivity of Amazonian carbon balance revealed by atmospheric measurements. *Nature*, *506*, 76.
- Gatti, L. V., Miller, J. B., D’Amelio, M. T. S., Martinewski, a., Basso, L. S., Gloor, M. E., Wofsy, S., & Tans, P. (2010). Vertical profiles of CO₂ above eastern Amazonia suggest a net carbon flux to the atmosphere and balanced biosphere between 2000 and 2009. *Tellus B*, *62*, 581–594. doi:doi: 10.1111/j.1600-0889.2010.00484.x.
- Gaubert, B., & Worden, H. (2017). NCAR/MOPITT Reanalysis. <https://doi.org/10.5065/D6SB44GZ>. Online; accessed 24 Sep 2020.
- Gaubert, B., Worden, H. M., Arellano, A. F. J., Emmons, L. K., Tilmes, S., Barré, J., Martinez Alonso, S., Vitt, F., Anderson, J. L., Alkemade, F., Houweling, S., & Edwards, D. P. (2017). Chemical feedback from decreasing carbon monoxide emissions. *Geophysical Research Letters*, *44*, 9985–9995.
- Geller, L., Elkins, J., Lobert, J., Clarke, A., Hurst, D., Butler, J., & Myers, R. (1997). Tropospheric sf₆: Observed latitudinal distribution and trends, derived emissions and interhemispheric exchange time. *Geophysical research letters*, *24*, 675–678.
- George, M., Clerbaux, C., Bouarar, I., Coheur, P.-F., Deeter, M. N., Edwards, D. P., Francis, G., Gille, J. C., Hadji-Lazaro, J., Hurtmans, D., Inness, D., A. Mao, & M., W. H. (2015). An examination of the long-term CO records from MOPITT and IASI: comparison of retrieval methodology. *Atmospheric Measurement Techniques*, *8*, 4313–4328.
- Gerken, R. R., & Franklin, J. A. (1989). The rate of degradation of 1, 1, 1-trichloroethane in water by hydrolysis and dehydrochlorination. *Chemosphere*, *19*, 1929–1937.
- Gilbert, J. C., & Lemaréchal, C. (1989). Some numerical experiments with variable-storage quasi-Newton algorithms. *Mathematical programming*, *45*, 407–435.
- Gloor, M., Dlugokencky, E., Brenninkmeijer, C., Horowitz, L., Hurst, D. F., Dutton, G., Crevoisier, C., Machida, T., & Tans, P. (2007). Three-dimensional SF₆ data and tropospheric transport simulations: Signals, modeling accuracy, and implications for inverse modeling. *Journal of Geophysical Research: Atmospheres*, *112*.
- Gloor, M. et al. (2012). The carbon balance of south america: a review of the status, decadal trends and main determinants. *Biogeosciences*, *9*, 5407–5430.
- González Abad, G., Liu, X., Chance, K., Wang, H., Kurosu, T. P., & Suleiman, R. (2015).

- Updated smithsonian astrophysical observatory ozone monitoring instrument (sao omi) formaldehyde retrieval. *Atmospheric Measurement Techniques*, 8, 19–32.
- Gromov, S., Taraborrelli, D., Pozzer, A., & Lelieveld, J. (2020). Atmospheric hydroxyl distribution from EMAC/MOM model v1. <https://doi.org/10.5281/zenodo.4044501>. Online; accessed 24 Sep 2020.
- Hard, T. M., Chan, C. Y., Mehrabzadeh, A. A., Pan, W. H., & O'Brien, R. J. (1986). Diurnal cycle of tropospheric OH. *Nature*, 322, 617–620.
- Hausmann, P., Sussmann, R., & Smale, D. (2016). Contribution of oil and natural gas production to renewed increase in atmospheric methane (2007–2014): top-down estimate from ethane and methane column observations. *Atmospheric Chemistry and Physics*, 16, 3227–3244. doi:doi: 10.5194/acp-16-3227-2016.
- Hmiel, B., Petrenko, V., Dyonisius, M., Buizert, C., Smith, A., Place, P., Harth, C., Beaudette, R., Hua, Q., Yang, B. et al. (2020). Preindustrial $^{14}\text{CH}_4$ indicates greater anthropogenic fossil CH_4 emissions. *Nature*, 578, 409–412.
- Holmes, C. D., Prather, M. J., Søvde, O. A., & Myhre, G. (2013). Future methane, hydroxyl, and their uncertainties: key climate and emission parameters for future predictions. *Atmospheric Chemistry and Physics*, 13, 285–302.
- Hooghiemstra, P. B., Krol, M. C., van Leeuwen, T. T., van der Werf, G. R., Novelli, P. C., Deeter, M. N., Aben, I., & Röckmann, T. (2012). Interannual variability of carbon monoxide emission estimates over south america from 2006 to 2010. *Journal of Geophysical Research: Atmospheres*, 117.
- Houweling, S. et al. (2014). A multi-year methane inversion using SCIAMACHY, accounting for systematic errors using TCCON measurements. *Atmospheric Chemistry and Physics*, 14, 3991–4012.
- Huang, J., & Prinn, R. G. (2002). Critical evaluation of emissions of potential new gases for OH estimation. *Journal of Geophysical Research: Atmospheres*, 107, ACH–18.
- Huijnen, V. et al. (2010). The global chemistry transport model TM5: description and evaluation of the tropospheric chemistry version 3.0. *Geoscientific Model Development*, 3, 445–473.
- Inness, A. et al. (2019). The cams reanalysis of atmospheric composition. *Atmospheric Chemistry and Physics*, 19, 3515–3556.
- INPE (2020). Amazon deforestation monitoring project (PRODES). <http://www.obt.inpe.br/OBT/assuntos/programas/amazonia/prodes>. Online; accessed 29 Sep 2020.
- Jacob, D. J. (1999). *Introduction to atmospheric chemistry*. Princeton University Press.
- Jeffers, P. M., Ward, L. M., Woytowitch, L. M., & Wolfe, N. L. (1989). Homogeneous hydrolysis rate constants for selected chlorinated methanes, ethanes, ethenes, and propanes. *Environmental Science & Technology*, 23, 965–969.

- Jöckel, P., Brenninkmeijer, C. A. M., & Crutzen, P. J. (2003). A discussion on the determination of atmospheric OH and its trends. *Atmospheric Chemistry and Physics*, 3, 107–118.
- Jöckel, P., Brenninkmeijer, C. A. M., Lawrence, M. G., Jeuken, A. B. M., & van Velthoven, P. F. J. (2002). Evaluation of stratosphere–troposphere exchange and the hydroxyl radical distribution in three-dimensional global atmospheric models using observations of cosmogenic ^{14}CO . *Journal of Geophysical Research: Atmospheres*, 107.
- Jöckel, P. et al. (2006). The atmospheric chemistry general circulation model ECHAM5/MESy1: consistent simulation of ozone from the surface to the mesosphere. *Atmospheric Chemistry and Physics Discussions*, 6, 6957–7050.
- Johnson, K. A., & Johnson, D. E. (1995). Methane emissions from cattle. *Journal of animal science*, 73, 2483–2492.
- Kaiser, J. W., Heil, A., Andreae, M. O., Benedetti, A., Chubarova, N., Jones, L., Morcrette, J. J., Razinger, M., Schultz, M. G., Suttie, M., & van der Werf, G. R. (2012). Biomass burning emissions estimated with a global fire assimilation system based on observed fire radiative power. *Biogeosciences*, 9, 527–554.
- Kirschke, S. et al. (2013). Three decades of global methane sources and sinks. *Nature geoscience*, 6, 813.
- Ko, M. K. W., Newman, P. A., Reimann, S., Strahan, S., Plumb, R. A., Stolarski, R. S., Burkholder, J. B., Mellouki, W., Engel, A., Atlas, E. L. et al. (2013a). Lifetimes of stratospheric ozone-depleting substances, their replacements, and related species. *SPARC Report*, 20.
- Ko, M. K. W., Newman, P. A., Reimann, S., & Strahan, S. E. (2013b). SPARC report on the lifetimes of stratospheric ozone-depleting substances, their replacements, and related species. *SPARC Report No. 6*, .
- Koch, D., Schmidt, G. A., & Field, C. V. (2006). Sulfur, sea salt, and radionuclide aerosols in GISS ModelE. *Journal of Geophysical Research: Atmospheres*, 111.
- Kopacz, M., Jacob, D. J., Henze, D. K., Heald, C. L., Streets, D. G., & Zhang, Q. (2009). Comparison of adjoint and analytical Bayesian inversion methods for constraining Asian sources of carbon monoxide using satellite (MOPITT) measurements of CO columns. *Journal of Geophysical Research: Atmospheres*, 114.
- Koren, G. (2020). *Constraining the exchange of carbon dioxide over the Amazon*. Wageningen University.
- Kort, E. A., Frankenberg, C., Costigan, K. R., Lindenmaier, R., Dubey, M. K., & Wunch, D. (2014). Four corners: The largest US methane anomaly viewed from space. *Geophysical Research Letters*, 41, 6898–6903.
- Kovács, T. et al. (2017). Determination of the atmospheric lifetime and global warming

- potential of sulfur hexafluoride using a three-dimensional model. *Atmospheric Chemistry and Physics*, 17, 883–898.
- Krol, M. et al. (2018). Age of air as a diagnostic for transport timescales in global models. *Geoscientific Model Development*, 11, 3109–3130.
- Krol, M., Houweling, S., Bregman, B., Broek, M., Segers, A., Velthoven, P. v., Peters, W., Dentener, F., & Bergamaschi, P. (2005). The two-way nested global chemistry-transport zoom model TM5: algorithm and applications. *Atmospheric Chemistry and Physics*, 5, 417–432.
- Krol, M., van Leeuwen, P. J., & Lelieveld, J. (1998). Global OH trend inferred from methylchloroform measurements. *Journal of Geophysical Research: Atmospheres*, 103, 10697–10711.
- Krol, M. C., & Lelieveld, J. (2003). Can the variability in tropospheric OH be deduced from measurements of 1,1,1-trichloroethane (methyl chloroform)? *Journal of Geophysical Research: Atmospheres*, 108.
- Krol, M. C., Lelieveld, J., Oram, D. E., Sturrock, G. A., Penkett, S. A., Brenninkmeijer, C. A. M., Gros, V., Williams, J., & Scheeren, H. A. (2003). Continuing emissions of methyl chloroform from Europe. *Nature*, 421, 131–135.
- Krol, M. C., Meirink, J. F., Bergamaschi, P., Mak, J. E., Lowe, D., Jöckel, P., Houweling, S., & Röckmann, T. (2008). What can ^{14}CO measurements tell us about OH? *Atmospheric chemistry and physics*, 8, 5033–5044.
- Kvenvolden, K. A. (1988). Methane hydrates and global climate. *Global biogeochemical cycles*, 2, 221–229.
- van der Laan-Luijkx, I. T. et al. (2015). Response of the Amazon carbon balance to the 2010 drought derived with CarbonTracker South America. *Global Biogeochemical Cycles*, 29, 1092–1108.
- Lamarque, J.-F. et al. (2010). Historical (1850–2000) gridded anthropogenic and biomass burning emissions of reactive gases and aerosols: methodology and application. *Atmospheric Chemistry and Physics*, 10, 7017–7039.
- Lamarque, J.-F. et al. (2012). CAM-chem: description and evaluation of interactive atmospheric chemistry in the Community Earth System Model. *Geoscientific Model Development*, 5, 369–411.
- Lamarque, J.-F. et al. (2013). The Atmospheric Chemistry and Climate Model Intercomparison Project (ACCMIP): overview and description of models, simulations and climate diagnostics. *Geoscientific Model Development*, 6, 179–206.
- Latham, D., & Williams, E. (2001). Lightning and forest fires. In *Forest Fires* (pp. 375–418). Elsevier.
- Lawrence, M. G., Jöckel, P., & Kuhlmann, R. v. (2001). What does the global mean OH

- concentration tell us? *Atmospheric Chemistry and Physics*, 1, 37–49.
- Leighton, P. A. (1961). *Photochemistry of Air Pollution*. Academic Press: New York.
- Lelieveld, J., Butler, T. M., Crowley, J. N., Dillon, T. J., Fischer, H., Ganzeveld, L., Harder, H., Lawrence, M. G., Martinez, M., Taraborrelli, D., & Williams, J. (2008). Atmospheric oxidation capacity sustained by a tropical forest. *Nature*, 452, 737–740.
- Lelieveld, J., Dentener, F. J., Peters, W., & Krol, M. C. (2004). On the role of hydroxyl radicals in the self-cleansing capacity of the troposphere. *Atmospheric Chemistry and Physics*, 4, 2337–2344.
- Lelieveld, J., Gromov, S., Pozzer, A., & Taraborrelli, D. (2016). Global tropospheric hydroxyl distribution, budget and reactivity. *Atmospheric Chemistry and Physics*, 16, 12477–12493.
- Lelieveld, J., Peters, W., Dentener, F. J., & Krol, M. C. (2002). Stability of tropospheric hydroxyl chemistry. *Journal of Geophysical Research: Atmospheres*, 107, ACH-17.
- Levin, I., & Hesshaimer, V. (1996). Refining of atmospheric transport model entries by the globally observed passive tracer distributions of ⁸⁵krypton and sulfur hexafluoride (SF₆). *Journal of Geophysical Research: Atmospheres*, 101, 16745–16755.
- Levy, H. (1971). Normal atmosphere: Large radical and formaldehyde concentrations predicted. *Science*, 173, 141–143.
- Liang, Q. et al. (2017). Deriving Global OH Abundance and Atmospheric Lifetimes for Long-Lived Gases: A Search for CH₃CCl₃ Alternatives. *Journal of Geophysical Research: Atmospheres*, 122, 11,914–11,933.
- Lizundia-Loiola, J., Pettinari, M. L., & Chuvieco, E. (2020). Temporal Anomalies in Burned Area Trends: Satellite Estimations of the Amazonian 2019 Fire Crisis. *Remote Sensing*, 12, 151.
- Logan, J. A., Prather, M. J., Wofsy, S. C., & McElroy, M. B. (1981). Tropospheric chemistry: A global perspective. *Journal of Geophysical Research: Oceans*, 86, 7210–7254.
- Lovelock, J. E. (1977). Methyl chloroform in the troposphere as an indicator of OH radical abundance. *Nature*, 267, 32–32.
- Lutgens, F. K., Tarbuck, E. J., & Tasa, D. (2001). *The atmosphere, 8th Edition*. Prentice-Hall Englewood Cliffs, NJ.
- Maasakkers, J. D. et al. (2019). Global distribution of methane emissions, emission trends, and OH concentrations and trends inferred from an inversion of GOSAT satellite data for 2010–2015. *Atmospheric Chemistry and Physics*, 19, 7859–7881.
- MacKay, C., Pandow, M., & Wolfgang, R. (1963). On the chemistry of natural radiocarbon. *Journal of Geophysical Research (1896-1977)*, 68, 3929–3931.

- Malhi, Y., Aragão, L. E. O. C., Galbraith, D., Huntingford, C., Fisher, R., Zelazowski, P., Sitch, S., McSweeney, C., & Meir, P. (2009). Exploring the likelihood and mechanism of a climate-change-induced dieback of the Amazon rainforest. *Proceedings of the National Academy of Sciences*, 106, 20610–20615.
- Martens, B., Gonzalez Miralles, D., Lievens, H., Van Der Schalie, R., De Jeu, R. A. M., Fernández-Prieto, D., Beck, H. E., Dorigo, W., & Verhoest, N. (2017). GLEAM v3: Satellite-based land evaporation and root-zone soil moisture. *Geoscientific Model Development*, 10, 1903–1925.
- McCulloch, A., & Midgley, P. M. (2001). The history of methyl chloroform emissions: 1951–2000. *Atmospheric Environment*, 35, 5311–5319.
- McNorton, J. et al. (2016). Role of OH variability in the stalling of the global atmospheric CH₄ growth rate from 1999 to 2006. *Atmospheric Chemistry and Physics*, 16, 7943–7956.
- McPhaden, M. J., Zebiak, S. E., & Glantz, M. H. (2006). ENSO as an integrating concept in earth science. *science*, 314, 1740–1745.
- Meirink, J. F., Bergamaschi, P., & Krol, M. C. (2008). Four-dimensional variational data assimilation for inverse modelling of atmospheric methane emissions: method and comparison with synthesis inversion. *Atmospheric chemistry and physics*, 8, 6341–6353.
- Millet, D. B., & Goldstein, A. H. (2004). Evidence of continuing methylchloroform emissions from the United States. *Geophysical research letters*, 31.
- Miyazaki, K., Eskes, H. J., Sudo, K., Takigawa, M., van Weele, M., & Boersma, K. F. (2012). Simultaneous assimilation of satellite no₂, o₃, co, and hno₃ data for the analysis of tropospheric chemical composition and emissions. *Atmospheric Chemistry and Physics*, 12, 9545–9579.
- Montzka, S. A. et al. (2018). An unexpected and persistent increase in global emissions of ozone-depleting CFC-11. *Nature*, 557, 413.
- Montzka, S. A., Krol, M., Dlugokencky, E. J., Hall, B., Jöckel, P., & Lelieveld, J. (2011). Small interannual variability of global atmospheric hydroxyl. *Science*, 331, 67–69.
- Montzka, S. A., Spivakovsky, C. M., Butler, J. H., Elkins, J. W., Lock, L. T., & Mondeel, D. J. (2000). New observational constraints for atmospheric hydroxyl on global and hemispheric scales. *Science*, 288, 500–503.
- Morton, D. C., Le Page, Y., DeFries, R., Collatz, G. J., & Hurtt, G. C. (2013). Understorey fire frequency and the fate of burned forests in southern amazonia. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368, 20120163.
- Murray, L. T., Logan, J. A., & Jacob, D. J. (2013). Interannual variability in tropical tropospheric ozone and OH: The role of lightning. *Journal of Geophysical Research: Atmospheres*, 118, 11–468.

- Myhre, G. et al. (2013). Anthropogenic and natural radiative forcing. *IPCC AR5: WG1: Climate change*, 423, 658–740.
- Naik, V., Jain, A. K., Patten, K. O., & Wuebbles, D. J. (2000). Consistent sets of atmospheric lifetimes and radiative forcings on climate for CFC replacements: HCFCs and HFCs. *Journal of Geophysical Research: Atmospheres*, 105, 6903–6914.
- Naik, V. et al. (2013). Preindustrial to present-day changes in tropospheric hydroxyl radical and methane lifetime from the atmospheric chemistry and climate model inter-comparison project (accmip). *Atmospheric Chemistry and Physics*, 13, 5277–5298.
- Naus, S., Montzka, S. A., Pandey, S., Basu, S., Dlugokencky, E. J., & Krol, M. (2019). Constraints and biases in a tropospheric two-box model of OH. *Atmospheric Chemistry and Physics*, 19, 407–424.
- Nechita-Banda, N., Krol, M., Van Der Werf, G. R., Kaiser, J. W., Pandey, S., Huijnen, V., Clerbaux, C., Coheur, P., Deeter, M. N., & Röckmann, T. (2018). Monitoring emissions from the 2015 Indonesian fires using CO satellite data. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373, 20170307.
- Nepstad, D., Carvalho, G., Barros, A. C., Alencar, A., Capobianco, J. P., Bishop, J., Moutinho, P., Lefebvre, P., Silva Jr, U. L., & Prins, E. (2001). Road paving, fire regime feedbacks, and the future of Amazon forests. *Forest ecology and management*, 154, 395–407.
- Nguyen, N. H., Turner, A. J., Yin, Y., Prather, M. J., & Frankenberg, C. (2020). Effects of Chemical Feedbacks on Decadal Methane Emissions Estimates. *Geophysical Research Letters*, 47.
- Nisbet, E. G., Manning, M., Dlugokencky, E., Fisher, R., Lowry, D., Michel, S., Myhre, C. L., Platt, S. M., Allen, G., Bousquet, P. et al. (2019). Very strong atmospheric methane growth in the 4 years 2014–2017: Implications for the paris agreement. *Global Biogeochemical Cycles*, 33, 318–342.
- Nölscher, A. C., Yañez-Serrano, A. M., Wolff, S., De Araujo, A. C., Lavrič, J. V., Kesselmeier, J., & Williams, J. (2016). Unexpected seasonality in quantity and composition of Amazon rainforest air reactivity. *Nature communications*, 7, 1–12.
- O'Connor, F. M., Boucher, O., Gedney, N., Jones, C., Folberth, G., Coppel, R., Friedlingstein, P., Collins, W., Chappellaz, J., Ridley, J. et al. (2010). Possible role of wetlands, permafrost, and methane hydrates in the methane cycle under future climate change: A review. *Reviews of Geophysics*, 48.
- Oman, L. D., Ziemke, J. R., Douglass, A. R., Waugh, D. W., Lang, C., Rodriguez, J. M., & Nielsen, J. E. (2011). The response of tropical tropospheric ozone to ENSO. *Geophysical Research Letters*, 38.
- Pandey, S. et al. (2019a). Satellite observations reveal extreme methane leakage from

- a natural gas well blowout. *Proceedings of the National Academy of Sciences*, 116, 26376–26381.
- Pandey, S., Houweling, S., Krol, M., Aben, I., Chevallier, F., Dlugokencky, E. J., Gatti, L. V., Gloor, E., Miller, J. B., Detmers, R., Machida, T., & Röckmann, T. (2016). Inverse modeling of GOSAT-retrieved ratios of total column CH₄ and CO₂ for 2009 and 2010. *Atmospheric chemistry and physics*, 16, 5043–5062.
- Pandey, S. et al. (2017). Enhanced methane emissions from tropical wetlands during the 2011 la niña. *Scientific reports*, 7, 45759.
- Pandey, S., Houweling, S., Krol, M., Aben, I., Nechita-Banda, N., Thoning, K., Röckmann, T., Yin, Y., Segers, A., & Dlugokencky, E. J. (2019b). Influence of atmospheric transport on estimates of variability in the global methane burden. *Geophysical Research Letters*, 46, 2302–2311.
- Patra, P. K. et al. (2011). Transcom model simulations of ch₄ and related species: linking transport, surface flux and chemical loss with ch₄ variability in the troposphere and lower stratosphere. *Atmospheric Chemistry and Physics*, 11, 12813–12837.
- Patra, P. K. et al. (2014). Observational evidence for interhemispheric hydroxyl-radical parity. *Nature*, 513, 219–223.
- Patra, P. K., Lal, S., Subbaraya, B. H., Jackman, C. H., & Rajaratnam, P. (1997). Observed vertical profile of sulphur hexafluoride (SF₆) and its atmospheric applications. *Journal of Geophysical Research: Atmospheres*, 102, 8855–8859.
- Patra, P. K. et al. (2009). Growth rate, seasonal, synoptic, diurnal variations and budget of methane in the lower atmosphere. *Journal of the Meteorological Society of Japan. Ser. II*, 87, 635–663.
- Pearson, C. R., & McConnell, G. (1975). Chlorinated C₁ and C₂ hydrocarbons in the marine environment. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 189, 305–332.
- Pereira, E. J. A. L., de Santana Ribeiro, L. C., da Silva Freitas, L. F., & de Barros Pereira, H. B. (2020). Brazilian policy and agribusiness damage the Amazon rainforest. *Land Use Policy*, 92, 104491.
- Petrenko, V., Murray, L., Smith, A., Crosier, E., Colton, A., Hua, Q., Yang, B., Kazemi, R., Neff, P., Etheridge, D., Usoskin, I., & Poluianov, S. (2019). Using atmospheric ¹⁴CO to provide additional constraints for global OH: results from a new approach and potential for future measurements. In *Geophysical Research Abstracts*. volume 21.
- Petron, G., Crotwell, A. M., Dlugokencky, E., & Mund, J. W. (2019). Atmospheric Carbon Monoxide Dry Air Mole Fractions from the NOAA ESRL Carbon Cycle Cooperative Global Air Sampling Network, 1988-2018. Version: 2019-08.
- Pison, I., Bousquet, P., Chevallier, F., Szopa, S., & Hauglustaine, D. (2009). Multi-

- species inversion of CH₄, CO and H₂ emissions from surface measurements. *Atmospheric Chemistry & Physics*, 9.
- Prather, M. J. (1996). Time scales in atmospheric chemistry: Theory, GWPs for CH₄ and CO, and runaway growth. *Geophysical Research Letters*, 23, 2597–2600.
- Prather, M. J., Holmes, C. D., & Hsu, J. (2012). Reactive greenhouse gas scenarios: Systematic exploration of uncertainties and the role of atmospheric chemistry. *Geophysical Research Letters*, 39.
- Prinn, R., Cunnold, D., Rasmussen, R., Simmonds, P., Alyea, F., Crawford, A., Fraser, P., & Rosen, R. (1987). Atmospheric trends in methylchloroform and the global average for the hydroxyl radical. *Science*, 238, 945–950.
- Prinn, R., Cunnold, D., Simmonds, P., Alyea, F., Boldi, R., Crawford, A., Fraser, P., Gutzler, D., Hartley, D., Rosen, R., & Rasmussen, R. (1992). Global average concentration and trend for hydroxyl radicals deduced from ALE/GAGE trichloroethane (methyl chloroform) data for 1978–1990. *Journal of Geophysical Research: Atmospheres*, 97, 2445–2461.
- Prinn, R. G. et al. (2005). Evidence for variability of atmospheric hydroxyl radicals over the past quarter century. *Geophysical Research Letters*, 32.
- Prinn, R. G. et al. (2001). Evidence for Substantial Variations of Atmospheric Hydroxyl Radicals in the Past Two Decades. *Science*, 292, 1882–1888.
- Prinn, R. G., Simmonds, P. G., Rasmussen, R. A., Rosen, R. D., Alyea, F. N., Cardelino, C. A., Crawford, A. J., Cunnold, D. M., Fraser, P. J., & Lovelock, J. E. (1983). The atmospheric lifetime experiment: 1. Introduction, instrumentation, and overview. *Journal of Geophysical Research: Oceans*, 88, 8353–8367.
- Prinn, R. G. et al. (2018). History of chemically and radiatively important atmospheric gases from the Advanced Global Atmospheric Gases Experiment (AGAGE). *Earth System Science Data*, 10, 985–1018.
- Quay, P., King, S., White, D., Brockington, M., Plotkin, B., Gammon, R., Gerst, S., & Stutsman, J. (2000). Atmospheric ¹⁴CO: A tracer of OH concentration and mixing rates. *Journal of Geophysical Research: Atmospheres*, 105, 15147–15166.
- Quay, P., Stutsman, J., Wilbur, D., Snover, A., Dlugokencky, E. J., & Brown, T. (1999). The isotopic composition of atmospheric methane. *Global Biogeochemical Cycles*, 13, 445–461.
- Randerson, J., Chen, Y., Van Der Werf, G., Rogers, B., & Morton, D. (2012). Global burned area and biomass burning emissions from small fires. *Journal of Geophysical Research: Biogeosciences*, 117.
- Reimann, S. et al. (2005). Low European methyl chloroform emissions inferred from long-term atmospheric measurements. *Nature*, 433, 506–508.

- Rigby, M. et al. (2017). Role of atmospheric oxidation in recent methane growth. *Proceedings of the National Academy of Sciences*, 114, 5373–5377.
- Rigby, M., Prinn, R. G., Fraser, P. J., Simmonds, P. G., Langenfelds, R., Huang, J., Cunnold, D. M., Steele, L. P., Krummel, P. B., Weiss, R. F. et al. (2008). Renewed growth of atmospheric methane. *Geophysical research letters*, 35.
- Rigby, M. et al. (2013). Re-evaluation of the lifetimes of the major cfcs and CH_3CCl_3 using atmospheric trends. *Atmospheric Chemistry and Physics*, 13, 2691–2702.
- Sanford, R. L., Saldarriaga, J., Clark, C., K. E. and Uhl, & Herrera, R. (1985). Amazon rain-forest fires. *Science*, 227, 53–55.
- Saunio, M. et al. (2016). The global methane budget 2000-2012. *Earth System Science Data*, 8, 697.
- Schaefer, H. et al. (2016). A 21st-century shift from fossil-fuel to biogenic methane emissions indicated by $^{13}\text{CH}_4$. *Science*, 352, 80–84.
- Schroeder, W., Csiszar, I., & Morisette, J. (2008). Quantifying the impact of cloud obscuration on remote sensing of active fires in the Brazilian Amazon. *Remote Sensing of Environment*, 112, 456–470.
- Schwietzke, S., Sherwood, O. A., Bruhwiler, L. M. P., Miller, J. B., Etiope, G., Dlugokencky, E. J., Michel, S. E., Arling, V. A., Vaughn, B. H., White, J. W. C., & Tans, P. (2016). Upward revision of global fossil fuel methane emissions based on isotope database. *Nature*, 538, 88–91.
- Scinocca, J. F., McFarlane, N. A., Lazare, M., Li, J., & Plummer, D. (2008). The CCCma third generation AGCM and its extension into the middle atmosphere. *Atmospheric Chemistry and Physics Discussions*, 8, 7883–7930.
- Shindell, D., Lamarque, J., Collins, W., Eyring, V., Nagashima, T., Naik, V., Szopa, S., & Zeng, G. (). The model data outputs from the Atmospheric Chemistry & Climate Model Intercomparison Project (ACCMIP). <http://catalogue.ceda.ac.uk/uuid/ded523bf23d59910e5d73f1703a2d540>. Online; accessed 23 Sep 2020.
- Silva, C. V. J. et al. (2018). Drought-induced Amazonian wildfires instigate a decadal-scale disruption of forest carbon dynamics. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373, 20180043.
- Silva Junior, C. H., Anderson, L. O., Silva, A. L., Almeida, C. T., Dalagnol, R., Pletsch, M. A., Penha, T. V., Paloschi, R. A., & Aragão, L. E. (2019). Fire responses to the 2010 and 2015/2016 Amazonian droughts. *Frontiers in Earth Science*, 7, 97.
- Simmonds, P. G. et al. (2016). Global and regional emissions estimates of 1,1-difluoroethane (HFC-152a, CH_3CHF_2) from in situ and air archive observations. *Atmospheric Chemistry and Physics*, 16, 365–382.
- Sindelarova, K., Granier, C., Bouarar, I., Guenther, A., Tilmes, S., Stavrakou, T., Müller,

- J.-F., Kuhn, U., Stefani, P., & Knorr, W. (2014). Global data set of biogenic VOC emissions calculated by the megan model over the last 30 years. *Atmospheric Chemistry and Physics*, *14*, 9317–9341.
- Skeie, R. B., Berntsen, T., Myhre, G., Pedersen, C. A., Ström, J., Gerland, S., & Ogren, J. A. (2011). Black carbon in the atmosphere and snow, from pre-industrial times until present. *Atmospheric Chemistry and Physics*, *11*, 6809–6836.
- Soares, J., Sofiev, M., & Hakkarainen, J. (2015). Uncertainties of wild-land fires emission in aqmeii phase 2 case study. *Atmospheric Environment*, *115*, 361–370.
- Spivakovsky, C. M., Logan, J. A., Montzka, S. A., Balkanski, Y. J., Foreman-Fowler, M., Jones, D. B. A., Horowitz, L. W., Fusco, A. C., Brenninkmeijer, C. A. M., Prather, M. J., Wofsy, S. C., & McElroy, M. B. (2000). Three-dimensional climatological distribution of tropospheric OH: Update and evaluation. *Journal of Geophysical Research: Atmospheres*, *105*, 8931–8980.
- Stell, A. C., Western, L. M., & Rigby, M. (2020). Atmospheric methane source and sink sensitivity analysis using gaussian process emulation. *Atmospheric Chemistry and Physics Discussions*, *2020*, 1–29.
- Stevenson, D. S., Doherty, R. M., Sanderson, M. G., Collins, W. J., Johnson, C. E., & Derwent, R. G. (2004). Radiative forcing from aircraft NO_x emissions: Mechanisms and seasonal dependence. *Journal of Geophysical Research: Atmospheres*, *109*.
- Szopa, S. et al. (2013). Aerosol and ozone changes as forcing for climate evolution between 1850 and 2100. *Climate dynamics*, *40*, 2223–2250.
- Taraborrelli, D., Lawrence, M. G., Crowley, J. N., Dillon, T. J., Gromov, S., Groß, C. B. M., Vereecken, L., & Lelieveld, J. (2012). Hydroxyl radical buffered by isoprene oxidation over tropical forests. *Nature Geoscience*, *5*, 190–193.
- Teyssède, H. et al. (2007). A new tropospheric and stratospheric Chemistry and Transport Model MOCAGE-Climat for multi-year studies: evaluation of the present-day climatology and sensitivity to surface processes. *Atmospheric Chemistry and Physics*, *7*, 5860.
- Tsuruta, A. et al. (2016). Development of CarbonTracker Europe-CH₄—Part 2: Global methane emission estimates and their evaluation for 2000–2012. *Geosci. Model Dev. Discuss*, .
- Turner, A. J., Frankenberg, C., Wennberg, P. O., & Jacob, D. J. (2017). Ambiguity in the causes for decadal trends in atmospheric methane and hydroxyl. *Proceedings of the National Academy of Sciences*, (p. 201616020).
- Turner, A. J., Fung, I., Naik, V., Horowitz, L. W., & Cohen, R. C. (2018). Modulation of hydroxyl variability by ENSO in the absence of external forcing. *Proceedings of the National Academy of Sciences*, *115*, 8931–8936.

- Van Der Werf, G. R., Randerson, J. T., Collatz, G. J., Giglio, L., Kasibhatla, P. S., Arellano, A. F., Olsen, S. C., & Kasischke, E. S. (2004). Continental-scale partitioning of fire emissions during the 1997 to 2001 El Niño/La Niña period. *Science*, *303*, 73–76.
- Van Leeuwen, T., Peters, W., Krol, M., & Van Der Werf, G. (2013). Dynamic biomass burning emission factors and their impact on atmospheric CO mixing ratios. *Journal of Geophysical Research: Atmospheres*, *118*, 6797–6815.
- Volz, A., Ehhalt, D. H., & Derwent, R. G. (1981). Seasonal and latitudinal variation of ^{14}CO and the tropospheric concentration of OH radicals. *Journal of Geophysical Research: Oceans*, *86*, 5163–5171.
- Voulgarakis, A. et al. (2013). Analysis of present day and future OH and methane lifetime in the ACCMIP simulations. *Atmospheric Chemistry and Physics*, *13*, 2563–2587.
- Walker, S. J., Weiss, R. F., & Salameh, P. K. (2000). Reconstructed histories of the annual mean atmospheric mole fractions for the halocarbons CFC-11 CFC-12, CFC-113, and carbon tetrachloride. *Journal of Geophysical Research: Oceans*, *105*, 14285–14296.
- Ward, P. J., Jongman, B., Kummu, M., Dettinger, M. D., Weiland, F. C. S., & Winsemius, H. C. (2014). Strong influence of El Niño Southern Oscillation on flood risk around the world. *Proceedings of the National Academy of Sciences*, *111*, 15659–15664.
- Watanabe, S. et al. (2011). MIROC-ESM 2010: model description and basic results of CMIP5-20c3m experiments. *Geoscientific Model Development*, *4*, 845–872.
- Wells, K. C., Millet, D. B., Payne, V. H., Deventer, M. J., Bates, K. H., de Gouw, J. A., Graus, M., Warneke, C., Wisthaler, A., & Fuentes, J. D. (2020). Satellite isoprene retrievals constrain emissions and atmospheric oxidation. *Nature*, *585*, 225–233.
- Wennberg, P. O., Peacock, S., Randerson, J. T., & Bleck, R. (2004). Recent changes in the air-sea gas exchange of methyl chloroform. *Geophysical research letters*, *31*.
- van der Werf, G. R., Randerson, J. T., Giglio, L., van Leeuwen, T. T., Chen, Y., Rogers, B. M., Mu, M., van Marle, M. J. E., Morton, D. C., Collatz, G. J., Yokelson, R. J., & Kasibhatla, P. S. (2017). Global fire emissions estimates during 1997–2016. *Earth System Science Data*, *9*, 697–720.
- Wiedinmyer, C., Akagi, S. K., Yokelson, R. J., Emmons, L. K., Al-Saadi, J. A., Orlando, J. J., & Soja, A. J. (2011). The fire inventory from near (finn): a high resolution global model to estimate the emissions from open burning. *Geoscientific Model Development*, *4*, 625–641.
- Wofsy, S. C., Afshar, S., Allen, H. M., Apel, E. C., Asher, E. C., Barletta, B., Bent, J., Bian, H., Biggs, B. C., Blake, D. R. et al. (2018). ATom: Merged atmospheric chemistry, trace gases, and aerosols. *ORNL DAAC*, .
- Wofsy, S. C., the HIPPO science team, cooperating modellers, & satellite teams (2011). HIAPER Pole-to-Pole Observations (HIPPO): fine-grained, global-scale measurements

- of climatically important atmospheric gases and aerosols. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 369, 2073–2086.
- Wolfe, G. M. et al. (2019). Mapping hydroxyl variability throughout the global remote troposphere via synthesis of airborne and satellite formaldehyde observations. *Proceedings of the National Academy of Sciences*, 116, 11171–11180.
- Worden, J. R., Bloom, A. A., Pandey, S., Jiang, Z., Worden, H. M., Walker, T. W., Houweling, S., & Röckmann, T. (2017). Reduced biomass burning emissions reconcile conflicting estimates of the post-2006 atmospheric methane budget. *Nature communications*, 8, 1–11.
- Yin, Y., Bloom, A., Worden, J., Saatchi, S., Yang, Y., Williams, M., Liu, J., Jiang, Z., Worden, H., Bowman, K., Frankenberg, C., & Schimel, D. (2020). Fire decline in dry tropical ecosystems enhances decadal land carbon sink. *Nature communications*, 11, 1–7.
- Yin, Y., Chevallier, F., Ciais, P., Broquet, G., Fortems-Cheiney, A., Pison, I., & Saunois, M. (2015). Decadal trends in global CO₂ emissions as seen by MOPITT. *Atmospheric Chemistry and Physics*, 15, 13433–13451.
- Zeng, G., Morgenstern, O., Braesicke, P., & Pyle, J. A. (2010). Impact of stratospheric ozone recovery on tropospheric ozone and its budget. *Geophysical Research Letters*, 37.
- Zhang, Y., Jacob, D. J., Maasakkers, J. D., Sulprizio, M. P., Sheng, J.-X., Gautam, R., & Worden, J. (2018). Monitoring global tropospheric OH concentrations using satellite observations of atmospheric methane. *Atmospheric Chemistry and Physics*, 18, 15959–15973.
- Zhao, Y. et al. (2020). On the role of trend and variability of hydroxyl radical (OH) in the global methane budget. *Atmospheric Chemistry and Physics Discussions*, 2020, 1–28.
- Zheng, B., Chevallier, F., Yin, Y., Ciais, P., Fortems-Cheiney, A., Deeter, M. N., Parker, R. J., Wang, Y., Worden, H. M., & Zhao, Y. (2019). Global atmospheric carbon monoxide budget 2000–2017 inferred from multi-species atmospheric inversions. *Earth System Science Data Discussions*, .

Acknowledgements

It has been a very nice and fulfilling four years, and for that, believe it or not, atmospheric oxidation was not the only reason. Indeed, I would like to thank here the many people and institutes with whom I shared some aspect of this journey that culminated in a book.

Thank you, Maarten for showing me how science is done, how to walk the line between scientific integrity and a healthy dose of pragmatism. Your ideas are at the core of this thesis, and your guidance and confidence have allowed me to develop and expand on them. Our talks always helped me get back some perspective when I got lost in the details. Perhaps less well known than your scientific qualities are your karaoke, dancing and stepping skills, but know that, there too, I am left deeply impressed. Thank you for first selecting me and for being on my team every step of the way.

Thank you, Steve, my second supervisor, for being the highest barrier in my publishing process. Thank you for being precise and critical, but also kind and attentive. Our scientific discussions sometimes meandered in no particular direction, toward no particular endpoint, which brought me to a deeper understanding of what knowledge is. Or maybe not, but in any case, I still very much enjoyed them. Your skepticism of models always laid bare the open ends in our story, for which, in the end, I really do am grateful. I hope to visit again someday to enjoy your grilling and your jokes.

Rare is the idea in this thesis that I can call mine alone. Thank you Wouter and Ingrid, for the steady rhythm of weekly group meetings that helped beat out the pages of this thesis, and for our joint final sprint on the Amazon chapter that helped push me over the finish line. Thank you Sourish, for the warm welcome in Boulder and for being the first to help me towards some perceived understanding of TM5. Thank you Narcisa and Sudhanshu for your code, and for guidance in the interpretation of the more difficult parts – if I ever did find spelling errors at least they were consistent. Thank you, Lucas, for your insights and for our enjoyable collaboration: my thesis stands all the stronger for it. And thank you to all the other people with whom I had the pleasure to discuss and share science.

Even on the rare occasion where work was not a pleasure, my Wageningen colleagues

made it a pleasure to go to work. In particular I want to thank my mentor Alba, for the warm welcome and for everything after (especially the sunny desk); Anja, for the distractions that made sure I did not finish my thesis too early, but much more happily; Aris, for being the most eloquent one-man drinking committee that I am aware of; Auke, for sharing the stress of getting into the EGU convener's party, which turned out to be quite *cool*; Imme, for your boundless enthusiasm that lends itself well to finishing a PhD, as well as to parties and soil drilling; Gerbrand, for your pleasant helpfulness – I can only hope to one day comment code like you do; Xabi, for your relaxing influence such as your sage advice to watch chickens more often. Thank you, Wei, for a happy and positive year, and for making sure I finished my plate on our Norway trip; Jordi, for your cheerful pessimism and for a memorable guest appearance in the party committee; Laurens, for giving me a taste of good teaching and for a chance to *be inspired*. It has been a pleasure to work at MAQ, a great group of people working on great science.

On Fridays, I happily made the shorter trip to the 6th floor of Buys Ballot, to IMAU. Happily, of course, because of the people. Thank you Elena, for being so confident in my abilities that I'm now thinking that, maybe, I can actually do this; Hossein, for your impressive push-up skills – mental as well as physical; Jin, for your friendliness (although I do have to ask that, now that it's all finalized, you find no more bugs in my code); Getachew, for your enthusiasm and for your dance moves: these I've seen rivaled only by a Juhi on crutches – thanks for that, too. And thanks to all the rest: all the times you did not forget me for lunch more than made up for the very few times you did.

I was lucky to spend a few months in Boulder at the NOAA research institute, where I had the opportunity to discuss science with a great group of researchers. The same trip was made much better even by my stay in Lorna's house. Thank you Yunsong, for all the trips, hikes and the daily lunches: it's pronounced *Stijn* but we can work on it; Xiangyu for your endless energy, for the hotpot and dumplings, and for your driving skills; Eunsol, Kai, Matthijs, Ronnie, for all the hikes, parties, trips and other wonderful memories.

A PhD project is incomplete without a few conferences, workshops and summer schools. Here I met, again, some wonderful people. In particular, I will not forget the nicely organized IGAC summer school and conference in Takamatsu, for which I was happy to be selected, even more so in hindsight. Thank you, Max, for - with characteristic confidence - identifying me already at the airport as a fellow participant, for semi-reluctantly grabbing the karaoke mike that first night and for being a great and knowledgeable travel companion. Thank you, Sonya, for the beers, laughs and for weathering the storm: the whole experience really did turn out to be pretty good; arigato gozaimasu Ryo, for showing me the onsen etiquette: I am so sorry we do not have them in Europe, we're still learning!

Als aanvulling op werk was daar altijd lacrosse, bij de Devils en – natuurlijk – in Team Ba-naam. Integraal onderdeel van mijn verwerking van werkfrustraties was af en toe heel hard

tegen iemand aanrennen, en daarna (of daarvoor) misschien één biertje. Maar soms had ik een langere pauze nodig, en tijdens dat jaarlijkse intermezzo waren er altijd de Travel Boyz. Bedankt Erik, Roelant en Ramon, voor de vele mooie reizen en alles daartussenin, misschien het best getypeerd door onze artistieke interpretatie van het legendarische blikje in het water. Snel weer!

Tot slot was mijn thesis er natuurlijk niet zonder jullie, mam, pap, met jullie nuchtere optimisme en hoog-meelevend vermogen, onder andere in de vorm van vele gebrande kaarsjes. Ik vertrouw erop dat, na het lezen van dit boekje, jullie het allemaal begrijpen, en anders is het ook goed. Dank je Suus en Gijs, voor jullie aandacht, de gezelligheid en voor het meedenken over de mooie thesis cover.



*Netherlands Research School for the
Socio-Economic and Natural Sciences of the Environment*

D I P L O M A

for specialised PhD training

The Netherlands research school for the
Socio-Economic and Natural Sciences of the Environment
(SENSE) declares that

Stijn Naus

born on 1 June 1993 in Dordrecht, The Netherlands

has successfully fulfilled all requirements of the
educational PhD programme of SENSE.

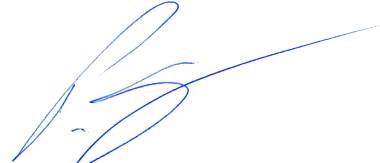
Wageningen, 3 February 2021

Chair of the SENSE board



Prof. dr. Martin Wassen

The SENSE Director



Prof. Philipp Pattberg

The SENSE Research School has been accredited by the Royal Netherlands Academy of Arts and Sciences (KNAW)



K O N I N K L I J K E N E D E R L A N D S E
A K A D E M I E V A N W E T E N S C H A P P E N



The SENSE Research School declares that **Stijn Naus** has successfully fulfilled all requirements of the educational PhD programme of SENSE with a work load of 34.0 EC, including the following activities:

SENSE PhD Courses

- o Environmental research in context (2016)
- o Research in context activity: 'Engaging the measurement community with my model-focused results, National Oceanic and Atmospheric Administration Global Monitoring Laboratory, Boulder, Colorado, United States' (2020)

Other PhD and Advanced MSc Courses

- o Atmospheric Chemistry summer school, Forschungszentrum Jülich (2017)
- o Early Career Short Course IGAC, International Global Atmospheric Chemistry (2018)
- o Scientific Artwork, Wageningen Graduate Schools (2019)
- o Infographics and Iconography, Wageningen Graduate Schools (2019)
- o Career Perspectives, Wageningen Graduate Schools (2020)

Management and Didactic Skills Training

- o Assisting practicals of the BSc course 'System Earth' (2016-2019)

Oral Presentations

- o A two box perspective on 3D modelling, TransCom Workshop, 17-20 September 2018, Lund, Sweden
- o Interannual variability in biomass burning over the Amazon, European Geosciences Union, 3-5 May 2019, Vienna, Austria

SENSE coordinator PhD education

Dr. ir. Peter Vermeulen

This work was carried out on the Dutch National e-Infrastructure with the support of SURF Cooperative. This work was funded through the Netherlands Organisation for Scientific Research (NWO), project number 824.15.002.

Financial support from Wageningen University & Research for printing this thesis is gratefully acknowledged.

This thesis was printed by ProefschriftMaken.nl

Cover design by Katarina Naskovski, based on ideas provided by Gijs Sanders, Stijn Naus and Suzanne Naus.