The potential for spatial variations in exposure and infection risk within a typical UK classroom: COVID-19 as a case study

Carolanne Valérie Mathilde Vouriot

July 2022

Department of Environmental and Civil Engineering
Imperial College London

Submitted in part fulfilment of the requirements for the degree of Doctor of Philosophy of Imperial College London and the Diploma of Imperial College London

Declaration of Originality

I certify that any material in this thesis that is not my own has been properly acknowledged and referenced.

Carolanne Vouriot

Copyright Statement

Abstract

Indoor air quality in schools, and classrooms, is of critical importance for the health and well-being of pupils and staff. The recent COVID-19 pandemic highlighted the essential role that indoor air quality and ventilation systems play in limiting the spread of airborne diseases. Within this context, the research in this thesis seeks to determine how exposure varies within naturally ventilated classrooms, which are prevalent in the UK.

First, CO_2 data gathered from 45 classrooms in 11 different schools are analysed and estimates of the likelihood of airborne infection occurring within these classrooms is calculated. Results show that, assuming relatively quiet desk-based work, the number of secondary infections is likely to remain low; however, the risk can vary widely between classrooms of the same school even when the same ventilation system is present. Crucially, even without accounting for the variation in disease prevalence, the data highlight significant variation in infection risk between the seasons, with January being nearly twice as risky as July.

The second part of the thesis provides a detailed examination of the validity of using CO_2 as a proxy for far-field exposure. For this, a generic naturally ventilated UK classroom in wintertime is simulated using CFD. The ventilation provision is through high and low-level openings and driven by buoyancy. CO_2 is modelled as a passive scalar and shown not to be well-mixed within the environment. The ratio between actual exposure arising from a single infected individual and proxy exposure calculated from point measurements of CO_2 is also analysed. In doing so, the proxy exposure estimated from the CO_2 concentration is found to be within a factor of two of the actual far-field exposure. Whilst this factor of two might appear large, it is small relative to the typical uncertainties associated with factors required to model airborne disease transmission.

Although this thesis focuses on COVID-19, the tools presented herein are useful to study and improve indoor air quality more broadly.

Acknowledgements

I would like to thank my supervisor Henry Burridge for his guidance and support throughout my PhD. His enthusiastic backing of all my endeavours gave me the freedom and confidence to grow as a researcher. I would also like to thank Maarten van Reeuwijk for co-supervising me and for his help in delving into the intricacies of OpenFOAM. I learnt a lot from our conversations on numerical modelling and it was always a pleasure to participate in his group meetings.

I would like to acknowledge my collaborators from the CO-TRACE and MAGIC projects as well as from the TAPAS and FUVN networks, especially Cath Noakes and Paul Linden. I would also like to thank Laetitia Mottet for her patience and her willingness to share her knowledge with me. I am grateful to have joined the Fluids Dynamics Across Scales CDT and I would like to thank Clodagh Li for her help throughout the years. Finally, I want to acknowledge everyone I met in the Fluid Mechanics section and special thanks go to the people of 004c.

This work used the Imperial College London High Performance Computing facilities and ARCHER2 UK National Supercomputing Service (https://www.archer2.ac.uk) (project e01). It was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant EP/L016230/1 for the Imperial College London Centre for Doctoral Training in Fluid Dynamics Across Scales and grant EP/W001411/1 for the COvid-19 Transmission Risk Assessment Case Studies - education Establishments (CO-TRACE) project.

Finally, I would like to acknowledge all my friends for their support. I am grateful for the friendliness of Barnes Women RFC and Oxford Academicals RC who welcomed on the pitch and on the water. Special thanks go to Kaija for supporting me all the way from Canada, Emily for introducing me to sewing and Charlotte for tolerating me for many years as her housemate. I would also like to thank my family for their loving support and my brother for always asking interesting questions. None of this would have been possible without my parents, their resolute faith in me and their continual support in all my projects, all I can say is *Merci!* Last, but certainly not least, I would like to thank Ross who was a constant source of inspiration,

reassurance and love throughout these years, I look forward to many more adventures with him.

Contents

List of Figures

List of Tables

List of Abbreviations

ACH Air Changes per Hour

ASHRAE American Society of Heating, Refrigerating and Air-Conditioning Engineers

CFD Computational Fluid Dynamics

DNS Direct Numerical Simulation

HVAC Heating, Ventilation and Air Conditioning

IAQ Indoor Air Quality

LES Large Eddy Simulation

RANS Reynolds-averaged Navier-Stokes

VOCs Volatile Organic Compounds

Nomenclature

Roman symbols

A	Vent area	m^2
A^*	Effective vent area	m^2
B	Potentially infected breath scalar concentration	ppm
C	CO_2 concentration	ppm
C_C	Classroom CO_2 concentration	ppm
c_d	Discharge coefficient	-
C_p	CO_2 concentration added through exhaled breath	ppm
c_p	Specific heat capacity	J/kg/K
c_v	Coefficient of variation	%
E_a	Actual exposure to infected breath B	-
E_p	Proxy exposure calculated from CO_2 measurements	-
f	Rebreathed fraction	-
G	Individual CO_2 generation rate	m^3/s
g	Gravitational acceleration	m/s^2
H_c	Classroom height	m
H_D	Computational domain height	m
I	Number of infectors	-

N	Total number of classroom occupants	-
p	Breathing rate	m^3/s
P_A	Probability of infection	-
Q	Flow rate	m^3/s
q	Quanta generation rate	quanta/h
Q_p	Individual flow rate	L/s/person
R_c	Actual to proxy exposure ratio (the proxy exposure is calculated from the room average CO_2 concentration)	-
R_l	Actual to proxy exposure ratio (the proxy exposure is calculated from the the local CO_2 concentration at a point)	-
S_c	Classroom surface area	m^2
S_I	Secondary number of infections	-
T	Temperature	°C
T_A	Duration of exposure	s
V_c	Classroom volume	m^3
W_c	Classroom heat input	W
CF	Cross-flow ventilation configuration	-
SS	Single-sided ventilation configuration	-

Greek symbols

α	Thermal expansion coefficient	K^{-1}
ϵ	Error	%
η	Contaminant removal efficiency	-
λ	Infectivity rate	-
ϕ	Fraction of infectious occupants	-

ρ	Density of air	kg/m^3

Subscripts

a	Ambient	-
B	Bottom vent	-
c	Classroom	-
D	Computational domain	-
e	Exhaust	-
l	Local	-
T	Top vent	-
w	Near or at the classroom vertical walls	-

Chapter 1

Introduction

The built environment provides humans protection and shelter from their surroundings and buildings are often used to temper the effects of changing conditions outdoors due to natural cycles and their variations. With urbanisation, our time spent indoors has been increasing steadily and can now be as much as 90% (Klepeis et al., 2001). To ensure occupants' comfort and health, the indoor environment typically needs conditioning, frequently achieved by controlling the temperature and air quality through energy intensive Heating, Ventilation and Air Conditioning (HVAC). Buildings are ventilated by allowing or forcing air to flow in and around the space through windows or vents. By ensuring a relatively clean source of fresh incoming air, suitable ventilation rates ensure that any potentially harmful airborne pollutants are removed. In the past century, ventilation has increasingly been delivered using mechanical means such as driving fans, pumps or powered air handling units. Mechanical ventilation and air conditioning are appealing as they allow direct control over the ventilation rate, temperature and humidity, as well as the ability to filter the air. This perceived need for mechanical conditioning has led however to significant increases in building energy consumption; up to 20% of all the energy consumed in Europe and North America was used to condition buildings through HVAC at the start of the 21$^{\text{st}}$ century (Pérez-Lombard et al., 2008). The considerable energy burden of mechanical HVAC provides one of a number of significant motivations to transition to low energy ventilation strategies. For this, *natural ventilation* presents an alternative by harnessing natural forces created by pressure gradients due to temperature differences and the wind to drive a ventilation flow (see Chenvidyakarn, 2013, for a detailed overview). Despite the mass deployment of mechanical ventilation, natural ventilation has remained the main source of ventilation for certain building types located in temperate climates such as the United Kingdom. This is notably the case in schools with more than 90% of energy display certificates issued for school buildings in England and Wales from October 2008 to June 2021 describing the ventilation system as natural ventilation (Ministry of Housing, Communities & Local Government, 2021). Besides providing the potential for energy savings, natural ventilation can also provide cost savings, since one of the advantages of natural ventilation is that it relies on freely available resources to drive the ventilation flow, such as the wind or the heat generated by radiators, electronic equipment or occupants. The downside of exploiting these natural forces to drive indoor airflows is that they make the ventilation harder to predict and control. The successful deployment of a ventilation system therefore requires a more careful consideration of the building's design to ensure it provides a suitable environment to building occupants.

One of the requirements of any ventilation strategy is to provide an adequate indoor environment, including good air quality. An estimated 90% of the world population lives in areas where the World Health Organisation recommendations for healthy air are not met, thereby posing health risks caused by long term exposure to harmful pollutants (Health Effects Institute, 2020). With pollutant sources originating both indoors and outdoors, it is therefore essential for the design and operation of the ventilation system to ensure that the air breathed by building occupants is of adequate, or ideally high, quality.

High air quality is especially important in schools and their classrooms. Children breathe in more air relative to their body size compared to adults and are still growing and developing which makes them particularly vulnerable to pollution (Landrigan et al., 2004). As children spend 30% of their time in school, 70% of which is in classrooms (Bakó-Biró et al., 2012), public health authorities have identified schools classrooms as a high-impact location to improve indoor air quality. There is a growing body of evidence linking indoor air quality to children's health and well-being: in their reviews, Salthammer et al. (2016), Fisk (2017) and Chatzidiakou et al. (2012), showed that school classrooms often have insufficient ventilation and high levels of pollutants which was found to directly impact the health and academic performance of the pupils. Demonstrating this correlation, the European-wide project SINPHONIE measured indoor air quality in schools across Europe and found higher levels of pollutants inside classrooms than outdoors (European Commission et al., 2014), with levels often above the recommended standards (Baloch et al., 2020). Bakó-Biró et al. (2012) and Clements-Croome et al. (2008) also showed that low levels of classroom ventilation negatively affects pupils' learning, lowering their performance on standardised assessments. Overall, the literature indicates that many classrooms in Europe are under ventilated and this lack of appropriate ventilation directly impacts pupils' learning and health (Haverinen-Shaughnessy et al., 2011; Corsi et al., 2021). These data reinforce the need to improve air quality in classrooms, which presents the opportunity to positively impact a large and vulnerable proportion of the population in a public and predominantly naturally ventilated setting, which can be directly regulated and monitored by public authorities.

The COVID-19 pandemic has brought the issue of air quality in schools and classrooms to the forefront of the attention of the wider public, with airborne disease transmission now a major concern. In early 2020, at the start of the pandemic, many countries resorted to unprecedented school closures to limit the spread of the disease, with UNESCO estimating at the time that 90% of the world's schoolchildren were impacted by school closures (UNESCO, 2020). These closures had a significant impact on children's education and development as reviewed by Ziauddeen et al. (2020), Corsi et al. (2021) and Engzell et al. (2021) who showed that despite online provisions children suffered measurable learning losses during that period. Reopening schools was therefore deemed a priority across the world, with a balance between mitigating the risk of transmission and continued teaching provision having to be reached. This notably called for a better understanding of the links between ventilation, indoor air quality and infection risk to ensure the health of pupils and staff in schools and, more generally, building occupants indoors.

Based upon existing knowledge, it is believed that SARS-CoV-2 transmission happens mostly indoors, as shown by Qian et al. (2021) through the analysis of hundreds of outbreaks in China. Long range airborne transmission was shown to play a significant role in some well documented cases with inadequate ventilation such as the Skagit choir (Miller et al., 2021), the Guangzhou restaurant (Lu and Yang, 2020) or Swiss courtrooms (Vernez et al., 2021) outbreaks. Li et al. (2007) showed that an improved ventilation system was necessary to reduce the spread of a number of airborne infectious diseases, highlighting the potential of improved ventilation beyond the COVID-19 pandemic. Aerosols exhaled by infected individuals indoors are responsible for long range airborne transmission and can be thought of as a source of contaminant indoors. As such, infection control is an integral part of indoor air quality and should be considered to provide a good indoor environment to building occupants.

Designing ventilation systems to reduce infection risks can require trade-offs to be made to ensure that the design and operation choices made to attain good air quality do not affect thermal comfort nor the building's energy performance and thus a holistic approach is needed. Although effective ventilation is a key solution known to positively impact pupils' health beyond limiting the risk of transmission (Corsi et al., 2021), it can be challenging to implement and operate successfully, especially in the case of naturally ventilated spaces. Often, a careful balance between good ventilation, thermal comfort, outdoor pollution, energy bills and now COVID-19 transmission risk has to be struck. Delivering this is difficult enough for trained HVAC engineers but in recent years it has now also been added to the growing list of responsibilities of teachers and school staff. The provision of CO_2 display monitors (i.e. monitors that provide an in-room display of CO_2 levels but do not record the levels) to UK schools (Department for Education, 2021) has provided a useful tool to address this issue by helping staff to practically manage ventilation and mitigate the risk of airborne transmission. However it does raise the question of whether a single CO_2 sensor can accurately represent the ventilating flow, the resulting exposure to contaminants indoors and potential infection risk of an entire classroom.

In this thesis, tools are developed to improve our understanding of the dispersion and ventilation of contaminants, such as infected breath, in canonical natural ventilation regimes. Although these methods are applicable to all ventilation systems, we specifically seek to understand the flows associated with buoyancy driven natural ventilation and how they interact with pollutants. This was achieved through the analysis of large datasets of CO_2 measurements in operational classrooms and from Computational Fluid Dynamics (CFD) simulations of a generic UK naturally ventilated classroom. These results provide a meaningful step towards ensuring the comfort and health of people in modern naturally ventilated built environments, with a specific concern on exposure to infected breath in naturally ventilated UK classrooms. This thesis aims to address the following research question: *how can CO_2 measurements be used to determine exposure to contaminants and infection risk in naturally ventilated classrooms?*

1.1 Indoor air quality and infection risk

The quality of the air we breathe has a direct effect on our health, both in the long and short term. Air quality is determined from the presence or lack of pollutants, an exposure to which might lead to adverse health outcomes, both via acute and chronic effects. In this thesis, the term contaminant will be used interchangeably with pollutant to define a substance found indoors above background concentrations. Although a distinction is sometimes made between contaminant and pollutant (Chapman, 2007), the focus herein is on exposure to a given substance and does not address directly the subsequent evaluation of toxicity and health impact. Poor Indoor Air Quality (IAQ) due to pollution indoors can lead to acute health effects (like respiratory infections), chronic conditions (i.e. asthma or allergies) and non-specific building-related symptoms such as headaches, respiratory problems, fatigue and irritation (European Commission et al., 2014).

Sources of indoor air pollution are described by BS ISO 16814:2008 (British Standard Institution, 2008) and include particulate matter, of different sizes such as PM10 or PM2.5, Volatile Organic Compounds (VOCs) or gases, like NO_2, CO or CO_2. Crucially, contaminants can originate both from indoors and outdoors (Morawska et al., 2013) and are directly influenced by a wide range of parameters such as the building materials, HVAC system and occupants' behaviour. In some cases, importantly airborne infection, the occupants themselves are the source of pollution. Although public concern has been focused on outdoor pollution, indoor air can be as or more polluted as outdoors and the risk caused by indoor pollution is increased by the time spent indoors: exposure is both a function of the pollutant concentration and the duration of exposure. Indoor air pollution has a significant impact on people's health worldwide, with household air pollution identified as the 9^{th} most significant cause of attributable global death (Health Effects Institute, 2020), despite decreases in exposure in recent years due to a move towards cleaner energy sources. The link between indoor air quality and occupants' health has been investigated by a number of studies in the past 30 years as reviewed by Tham (2016). This work showed a need for improved ventilation but studying indoor air is difficult, as demonstrated for instance by Lakey et al. (2021) who showed the different scales of variability encountered within buildings. These scales range in both time and space and emerge due to a number of causes including chemical reactions, surface interactions and transport caused by airflows, such as ventilating airflows. In cases where the source of pollution cannot be removed, ventilation is essential to reduce exposure to pollution indoors and manage the room's overall IAQ.

Assessing and measuring air quality within a room still remains a challenging issue. In particular, many countries do not have national guidelines for indoor air quality; if they do, they are often not regulatory and cannot be enforced (Morawska and Huang, 2022). Standards exist to ensure that concentrations of indoor pollutants do not rise above a level that can pose a risk to the health of the occupants or impact their comfort (summarised for instance by BS ISO 16814:2008, British Standard Institution, 2008). However, such metrics are set at the room level and as such might not be satisfied for every occupant (Ole Fanger, 2006) or take into account each individual person's exposure to pollutants. The existing standards for IAQ can be broadly divided in three categories; the first includes those determined by the exposure thresholds set by different health organisations considering both long and short term health effects. These thresholds are pollutant specific and in order to show compliance a large number of measurements have to be made, which is difficult to implement at a large scale and over long periods. The second category is based on how the majority of the occupants will perceive the quality and suitability of their environment, and are by design subjective and measured through surveys. Commonly, the third approach has been to ensure adequate IAQ indirectly by setting a minimum bulk ventilation rate, measured directly and expressed in Air Changes per Hour (ACH) or specified per occupant, or calculated indirectly through a room averaged pollutant concentration level.

Whilst known to have limits (as articulated for instance by ASHRAE, 2022), CO_2 measurements have notably been used to assess ventilation levels and IAQ. For instance, CO_2 levels have been used by Department for Education (2018), through maximum daily averages and maximum 20 minutes peak values, to provide guidance on ventilation design for schools in the UK. In France, the stuffiness index, ICONE, has been devised to take into account both the peaks of CO_2 measurements and their frequency. To quantify a room's stuffiness, ICONE gives a rating between 1 and 5, with 5 indicating a very stuffy room with high levels of CO_2 which has potentially dangerous levels of pollutants accumulated indoors, which could lead to an increased exposure for pupils and staff. This index was tested in a nationwide measurement campaign in France in schools and nurseries (Ramalho et al., 2013; Canha et al., 2016) showing a general high level of stuffiness in classrooms. The effectiveness of this metric resulted in it being adopted as part of the French IAQ regulations. Unlike other pollutants, CO_2 can be easily measured for long times, across years, using relatively cheap sensors that are becoming widespread. In their review, Lowther et al. (2021) showed that whilst it can be difficult to directly attribute adverse health effects to low levels of CO_2 (below 5,000 ppm) due to the number of confounding factors, CO_2 levels can however be used as an indication of the rates of ventilation and resulting pollutant concentrations indoors. This is useful information but also has limitations as the source of CO_2 differs from other commonly encountered pollutants indoors. Indeed, in the absence of major sources of combustion, CO_2 originates mostly from the occupants' exhalation. Using measurements in classrooms, Chatzidiakou et al. (2015a) showed that while CO_2 guidelines and temperature thresholds might suffice as a proxy for pollutants originating from indoors, they may not be adequate for traffic related pollution (Chatzidiakou et al., 2015b). Metrics based on CO_2 measurements should therefore be thought of as estimates of the room's ventilation, with high values likely to indicate poor ventilation and therefore a likely increase in the concentrations of other pollutants, including infected breath.

As current standards are primarily designed to ensure the comfort of the majority of occupants with minimal consideration of adequate IAQ, they do not, for most public spaces (the exception being hospital settings) take into account infection transmission risk. However, infection risk has been shown to be an integral part of indoor air quality, as exposure to virus particles indoors can directly affect the health of building occupants. In this thesis, the focus is on COVID-19 transmission which is thought to happen mainly indoors via three main routes: the touch (or contact) route, the short range droplet route and the far-field airborne route. Herein, we focus on far-field airborne transmission and as such do not consider the complex droplet and flow dynamics taking place near an infected individual. This route of transmission happens through the generation of viral aerosols (droplets of a range of sizes and composition created when talking, coughing, sneezing or breathing) by an infected person, the spread of these particles across an indoor space and their inhalation by a susceptible occupant. Airborne transmission of infectious diseases has been shown to occur within an indoor space, between indoor spaces and between buildings (Ai et al., 2020). The typical mitigating measures of social distancing and enhanced cleaning have little effect on long range airborne transmission and the most effective measure is ventilation. The COVID-19 pandemic has led experts to call for a 'paradigm shift' in how ventilation is designed (Morawska et al., 2021) so that it includes minimising the risk of airborne infection indoors in addition to considering thermal comfort. Focusing on the far-field airborne route implies that the methods and results developed in this thesis are not specific to COVID-19 but can be applied to all airborne diseases as well as other indoor and outdoor pollutants found in buildings.

In the context of airborne disease transmission, a space with good IAQ is one where the risk of infection via the airborne route is low. Assessing the risk of a space is also key both to understand what contributes most to the spread of the disease and to determine what mitigating measures can be put in place to limit further infections. Typically, for airborne infections, this has been done using either a Wells-Riley model (Riley et al., 1978) or a dose response model, as reviewed by Sze To and Chao (2010). These models have been applied extensively in recent years to COVID-19. Zhang and Lin (2021) notably extended the typical Wells-Riley and rebreathed fraction (Rudnick and Milton, 2003) models to allow for spatial and temporal variations by using a dilution model. Burridge et al. (2021b) further extended the work of Rudnick and Milton (2003) to account for variable occupancy and to calculate the number of secondary infections. Jones et al. (2021) and Peng et al. (2022) used similar models to estimate the varying risk of transmission in different indoors setting using either a relative exposure index or risk parameters. Su et al. (2021) investigated the effect of different ventilation methods on transmission risk in an office by combining CFD simulations and a revised Wells-Riley model. CO_2 measurements have also been used to estimate the flow rates and resulting infection risks in New York City classrooms (Pavilonis et al., 2021) as well as to design a ventilation strategy to minimise the spread of the virus (Stabile et al., 2021).

Often, these studies rely on assuming that the air within a room is well mixed, or that the room can be divided into well mixed zones as done for instance by Noakes and Sleigh (2009) in their study of hospital ventilation. More widely, this has been used to study contaminants indoors, where the air within a room is taken to be well-mixed such that it can be assumed that pollutants will be quickly diluted and their concentration can be assumed to be uniform in the room. This thesis aims to investigate the validity of this assumption in naturally ventilated classrooms. Although assuming a well mixed space when considering most mechanically ventilated rooms can be appropriate, in the case of natural ventilation, which can include stratification, it is problematic. Previous measurements in an experimental chamber (Mahyuddin and Awbi, 2010) and classrooms (Mahyuddin et al., 2014) showed how this assumption breaks down in practice with considerable CO_2 variations observed in the space and highlighted their link to the chosen ventilation system. Nicas (1996) also showed the impact of using the well-mixed assumption when the room presents stratified layers. By using simple mathematical models and a two zone room with different temperatures, they demonstrated that using a well-mixed assumption can lead to an underestimation for up to 40% of the exposure in the zone of occupancy, despite the use of a mixing factor supposed to account for an imperfectly mixed room. If IAQ and transmission risk are to be accurately quantified in naturally ventilated spaces, the assumption of a well mixed space cannot be relied upon and a thorough understanding of ventilating airflows and their interaction with contaminants indoors is required.

1.2 Exposure in naturally ventilated spaces

To study infection risk and more widely indoor air quality, it is necessary to quantify exposure to a certain pollutant. Assessing risk for IAQ has several components such as identifying the contaminant, the human response to a certain dose of this contaminant and the exposure someone might suffer from it, the latter part being the focus of this thesis. Exposure depends both on concentration and exposure duration, which requires a knowledge of the pollutants' sources, how they are transported and mixed through the space and any potential removal processes. This in effect means that the flow field across the space needs to be known so that its interaction with the pollutant can be determined.

This is particularly challenging for naturally ventilated spaces, where the flow is not mechanically forced but occurs naturally and will therefore be harder to predict. Natural ventilation has been widely investigated analytically, experimentally and computationally, as reviewed by Linden (1999) and Chenvidyakarn (2013). These studies have focused mostly on flows driven solely by buoyancy, or stack-driven ventilation. This is relevant as it is the case where wind does not assist ventilation: a worst case if buildings are designed such that any wind assists the ventilation flow.

A well studied canonical ventilation regime is displacement ventilation where the heat sources are localised and the flow is unidirectional through the openings resulting in a two layer stratification. By contrast, in the case of mixing ventilation, fresh cool air is supplied through a single opening and mixes with the warm indoor air producing a uniform temperature distribution within the room and an exchange flow at the vent. The localised release of heat in displacement ventilation gives rise to a buoyant turbulent plume that, in an enclosed space with openings at different heights, will drive a bulk ventilating flow and internal temperature stratification. Consideration of a single room provides a starting point to develop understanding and simple extensions enable the examination of buildings in their entirety via so called 'flow loop' models (Axley, 2007).

By studying displacement ventilation analytically and comparing it to water-bath experiments, Linden et al. (1990) showed the establishment of a steady state where the height of the stratified layer is independent of the strength of the heat source and is, instead, a function only of the geometry of the room and the entrainment created by the plume. This work was then extended to account for multiple openings at different heights (Fitzgerald and Woods, 2004), the presence of doorways (Higton et al., 2021) and transient behaviour (Kaye and Hunt, 2004). Gladstone and Woods (2001) also extended this work to look at a distributed heat source which represents underfloor heating or, in the limit, a large number of localised heat sources. Distributed heating also gives rise to a steady displacement flow, although temperature stratification within the room is less distinct and the environment is more mixed. Most of these studies compare analytical results to water-bath experiments, where a small-scale model of a room or building is submerged in fresh water, with a source of saline water representing the buoyant heat source. These analogue small-scale models have been shown to replicate the key features of the full-scale flows (Linden, 1999) and as such are extremely valuable to study buoyancy driven ventilating flows.

Different approaches have been followed to study the behaviour of contaminants in a naturally ventilated environment. Hunt and Kaye (2006) developed a theoretical model, validated by water-bath experiments, to determine the time needed for passive pollutants to be flushed from a room with different localised heat sources. This was later extended by Zhong et al. (2012) to find the pollutant concentration in each layer. Using a similar method, Bolster and Linden (2007) looked at the efficiency of contaminant removal and Bolster and Linden (2009a,b) considered the behaviour of particles in a displacement ventilated space. This notably showed that considering average concentrations of pollutants will not accurately represent the exposure to pollutants suffered by occupants in a room. By using particle measurements in a test room with displacement ventilation and different sources of contaminants, Mundt (2001) also showed that contaminant removal effectiveness, that is how quickly the contaminant is removed, was influenced by the ventilating flow rate and the location of the source of the pollutants, including its height and relative distance to the heat source driving the flow. Other methods to estimate exposure in naturally ventilated spaces have included the use of tracer gas to determine experimentally the flow rate and ventilation efficiency, probabilistic methods to assess the risk of contamination (Bolster and Tartakovsky, 2008) or field measurements in apartments buildings (Park et al., 2014). However, these studies remain limited by experimental constraints with a limited number of available measurement points to determine exposure from.

CFD, on the other hand, gives access to the full three dimensional fields of velocity, buoyancy and pollutant concentrations, and allows the study of complex geometries. Different methods are available to resolve turbulence and each offers different advantages. For instance, Direct Numerical Simulation (DNS) resolves all scales of motion for given boundary and initial conditions. As such it is accurate, but also computationally expensive and impracticable for high Reynolds flows in real life applications, with scales ranging from 10 m to less than 1 cm in the building environment. On the other hand, Reynolds-averaged Navier-Stokes (RANS) simulations use a parametrisation of turbulence to determine the Reynolds stresses, while the mean fields are resolved using the Reynolds averaged equations. Large Eddy Simulation (LES) looks to strike a balance between the two by filtering out eddies smaller than the computational grid and parameterise them with a subgrid-scale model, like the Smagorinsky model (Smagorinsky, 1963). Due to their relative low cost, RANS simulations have been used for a range of engineering applications including indoor and urban flows as reviewed by Blocken (2018). In the case of natural ventilation flows specifically, the standard isotropic closure models need to be extended to account for buoyancy effects (Kumar and Dewan, 2014).

In the context of natural ventilation, a number of cases have been considered computationally, primarily using RANS simulations. For example Kaye et al. (2009) used RANS simulations to study transient ventilation in the case of the canonical displacement ventilation and considered different openings. Yang et al. (2015) also performed RANS simulations to look at this scenario and the results are validated against experimental measurements and theoretical predictions.

More complex cases have also been investigated, such as a building with an atrium (Ji et al., 2007) or a room ventilated both by buoyancy and wind (Cook et al., 2003). On the other hand, studies using LES have been limited, considering canonical displacement ventilation flows (Abdalla et al., 2007, 2009) or comparing its performance against RANS simulations (Durrani et al., 2012, 2015), where LES was shown to be more accurate than RANS simulations, but also more computationally expensive. In this work, it was decided to employ RANS simulations so that a number of ventilation configurations could be studied rather than investigating a single case in depth.

CFD has also been used to study pollutant transport inside naturally ventilated spaces. Studies have investigated a range of scenarios, using RANS, LES (Zhang et al., 2007; Liu and Novoselac, 2014) and DNS (Yang et al., 2022). However, the flow is rarely solely buoyancy driven and the inflow is often forced at the inlet. Since the start of the pandemic, interest has grown to investigate indoors transmission with CFD. Long range airborne transmission through infected aerosols inside rooms has been simulated both with particle and tracer gas simulations. For example, Pei et al. (2021) and Mariam et al. (2021) both investigated infection between two people in a small room using CFD, highlighting a variation due to the choice of ventilation system. A key issue highlighted notably by Ai et al. (2020) Pourfattah et al. (2021) and Coldrick et al. (2022), is the choice of boundary conditions at the source that will accurately represent the range of infected particles emitted by the population and likely to be encountered in realistic settings. These studies have been quite specific, often replicating an experimental set-up, and are applicable to a limited number of rooms. Tan et al. (2022) considered this issue and created a database of 174 simulations to study the inhomogeneous distribution of infected droplets in a range of mechanically ventilated locations such as classrooms or buses. This study found that the standard deviation of the concentration field was comparable to its mean and, when coupled to a dose-response model, infection probability was found to differ by an order of magnitude from the well-mixed prediction. The interaction between buoyancy driven ventilation and contaminant transport indoors however, has not been widely studied in a generic setting. Simulating a canonical buoyancy driven set-up allows for a better understanding of what parameters govern exposure and the resulting infection risk in a space. The expected variations of both within a room can also be easily quantified using CFD which can then be used to assess the accuracy of other methods.

1.3 Scope

This thesis aims to study indoor air quality in classrooms and its link to exposure and infection risk through the use of two methods. The first takes advantage of the growing datasets of CO_2 measurements indoors, where, unlike punctual measurement campaigns run for a few weeks, data can be gathered for longer periods revealing long-term trends. CO_2 measurements originally used to manage the ventilation system of a classroom and collected over a period of 5 years are used to infer the risk of airborne infection arising from a single infected individual. Implicitly, the available measurements are assumed to be representative of the overall classroom and, by using CFD simulations in the second part of this thesis, the consequences and validity of the use of a single point measurement are explored. The link between actual far-field exposure to infected breath and proxy exposure calculated from CO_2 measurements is also investigated.

Due to their occupant numbers and their vulnerability to indoor contaminants, classrooms are used as a case study. Most UK classrooms are naturally ventilated and this thesis does not consider explicitly mechanical ventilation, although the methodology used could be used to investigate a wider range of ventilation systems, beyond natural ventilation. The cases studied herein are buoyancy driven flows that represent the limiting case where wind provides no assistance in driving the ventilating flow. A particular focus is drawn to the transmission of COVID-19 but the methods developed in this thesis can be applied to any other airborne disease, as well as pollutants found indoors that interact with ventilation within a space: this thesis is concerned with improving indoor air quality as a whole. In this view, the cases studied here are fundamental and do not represent specific classrooms in the hope that the conclusions of this thesis will remain applicable to the widest range of spaces possible.

As a first approximation, it is assumed that occupants are solely passive sources of CO_2 and heat, their movements, breathing and how this might affect the flow, are not considered. This work also does not take into account how occupants might interact with their environment and impact the performance of the ventilation system. This is because the focus is on the ventilation system and how it affects indoor air quality. This is crucial for its design but for a ventilation set-up to be successfully implemented in a building, its built, operation and maintenance also have to be carefully considered.

1.4 Aims and Objectives

The aim of this thesis is to investigate contaminant behaviour in naturally ventilated classrooms and determine the extent to which exposure, and airborne infection risk, can be evaluated from CO_2 measurements. This work uses both measurements in operational classrooms and CFD to explore these issues; including a) providing estimates of the scale of airborne infection risk in classrooms for COVID-19, b) determining how valid the well mixed assumption might be in these naturally ventilated spaces, and c) assessing how CO_2 measurements might best be used to infer airborne infection risk.

Although CO_2 measurements have been shown to be appropriate to assess ventilation, being able to determine whether CO_2 can be used as a proxy for exposure in a naturally ventilated indoor space remains an open problem that this thesis aims to address. As such, the following objectives for this thesis have been identified:

O1. Demonstrate the ability to estimate the risk of airborne infection from CO_2 measurements in operational classrooms.

O2. Determine how CFD, and specifically OpenFOAM, can be used to model and simulate generic naturally ventilated UK classrooms.

O3. Examine what CO_2 variations can be expected in generic naturally ventilated classrooms.

O4. Determine how CO_2 can be used as a proxy for far-field exposure to infected breath.

1.5 Structure of the thesis

The general structure of this thesis and an overview of the different chapters can be summarised as follows:

Chapter 2: Describes how airborne infection risk can be estimated from CO_2 measurements and demonstrates application by determining the probability of infection in UK classrooms based on data gathered in 45 classrooms over a 5-year period (O1, O4). This work was published in Vouriot et al. (2021).

Chapter 3: Details how a generic buoyancy driven naturally ventilated classroom can be simulated using CFD and specifically OpenFOAM (O2).

Chapter 4: Shows what CO_2 variations can be expected in generic naturally ventilated classrooms (O3).

Chapter 5: Details what variations are expected when calculating the risk of airborne infection from actual far-field exposure to infected breath and compared to a proxy obtained through CO_2 measurements (O4).

Chapter 6: Presents the conclusions of this work.

Chapter 2

Risk of airborne infection in classrooms from monitored CO_2

This chapter focuses on estimating airborne transmission risk, in our case for COVID-19, in UK classrooms using CO_2 data collected in 45 classrooms and 11 schools across the UK. Measurements were collected between 2015 and 2020 as part of the classroom's ventilation management system. The proposed approach follows the work of Riley et al. (1978) and Rudnick and Milton (2003) to estimate the number of secondary airborne infections arising from a single infected individual attending a classroom for a period of 5 working days. The results highlight the variation in infection risk across classrooms, schools and the seasons due to differences in environmental conditions and ventilation provision.

The work in this chapter is based on research published in:

- C. V. M. Vouriot, H. C. Burridge, C. J. Noakes, and P. F. Linden. Seasonal variation in airborne infection risk in schools due to changes in ventilation inferred from monitored carbon dioxide. *Indoor Air*, 31(4):1154–1163, 2021

2.1 Background

Estimating the likelihood of airborne infection in schools became critical as the COVID-19 outbreak turned into a global pandemic throughout 2020 and 2021. Most schools were initially closed, along with many other public spaces, as part of lockdown measures used to contain the spread of the virus. With the exception of children of 'key-workers' and children within certain vulnerable groups, this resulted in pupils learning from home, with only some receiving remote provision for the initial few months. As schools reopened it was crucial to reach a balance between providing a suitable space for learning and one with limited risks of COVID-19 transmission. Within UK primary schools, it is typical that classes of pupils and adults remain located within the same classroom for the majority of the school week. During the reopening of schools, the government recommendation that schools form so called 'bubbles' (classes or year groups remaining relatively isolated from others within the school) resulted in most secondary school pupils remaining within the same classroom for much of the school week. As such, classrooms form an interesting case for which the same space is attended over long durations by the same occupants (both staff and students) over much of the working week, making them a viable space in which to study the risk of COVID-19 transmission via relatively simple models. The cohort of pupils and classroom staff covers around 11 million people in the UK (BESA, 2021) with transmission to their families and social groups potentially affecting many more millions. It is especially important to understand how this risk might vary under different conditions; for example, changes in the seasons and the approach of winter (where time spent indoors can be expected to increase), so that observation made during one season might be adjusted to inform planning for another season. Although children seem to be less likely to become seriously ill from COVID-19, it is known that they can still get infected (WHO, 2020) and so increase the chance of spreading the virus to more susceptible members of the population, including teachers and elderly relatives.

COVID-19 is thought to transmit mostly indoors (as shown, for example, by Qian et al., 2021, with the analysis of hundreds of outbreaks in China) through three routes: the droplet (or spray) route, the contact (or touch) route and airborne (or aerosol) route. A number of papers discuss these transmission routes and how they can be mitigated including for example Burridge et al. (2021a). The focus of this chapter and the overall thesis is on relatively far-field (typically taken to occur over distances greater than 2 m) airborne transmission where small infected respiratory droplets and aerosols remain suspended in the air where they can be transported by indoor air currents and inhaled. This route is thought to be significant in the spread of COVID-19, with the study of well documented outbreaks such as the Skagit Valley Chorale event for which Miller et al. (2021) show that aerosol transmission has to be considered in order to account for the extensive resulting number of cases. While measures such as social distancing and appropriate cleaning procedures will reduce the risk linked to the droplet and contact routes, infections due to airborne transmission may be harder to contain. They require, along with other measures (Morawska et al., 2020), the appropriate use of ventilation as recommended by the WHO (2020).

But, as described notably by Bhagat et al. (2020), the physics of indoor flows are complex and thus the consistent provision of clean air at appropriate rates to all desired locations within an indoor spaces remains a scientific and engineering challenge.

To analyse the likelihood of airborne infection from COVID-19 caused by attendance at schools, this chapter adopts the approach described in the seminal works of Riley et al. (1978) and Rudnick and Milton (2003), as further discussed in §2.2. This method predicts the absolute risk of airborne infection over a period of time in a given space, by taking CO_2 measurements and occupancy profiles and calculating the number of secondary infections, or the number of infections due to a single originally infected individual. This assumes that anyone showing any symptoms ceases attending school, as it has been recommended throughout the pandemic. We further assume that staff and students spend the vast majority of their school day in the classroom and thus infer that the risk of airborne infection is dominated by their time in the classroom — which is likely to be the case in most schools, with break times being taken outdoors (minimal airborne infection risk) or within the classroom in inclement weather. We therefore calculate the likely consequences of airborne infection, the number of secondary infections, for pre/asymptomatic periods where the original infector continues to regularly attend the space; in this instance, a classroom. The pre/asymptomatic period may be especially significant, as it corresponds to the period where infectivity is thought to reach a peak (He et al., 2020).

This approach is particularly applicable to primary schools (5–11 years old) where the same group of students can be assumed to attend the same classroom every day. It also remains suitable for secondary schools in which strategies have been implemented to reduce mixing between students, introducing for example fixed bubbles or groups of students in one classroom.

In addition, this methodology is useful because it does not require measurements nor estimates of ventilation flow rates, which are seldom recorded. Instead, it relies on the use of monitored CO_2 which is becoming increasingly widespread in newly built schools in the UK in an effort to assess the performance of ventilation systems along with showing compliance to guidelines (e.g. Department for Education, 2018). In response to the pandemic, there has also been a nationwide roll-out of CO_2 monitors in schools (Department for Education, 2021) that, whilst only displaying CO_2, show the ability and willingness to deploy these relatively affordable monitors at scale – perhaps generating significant amounts of data in the future. The focus on airborne transmission makes the work presented here relevant to a wider range of airborne diseases, such as measles, influenza or SARS, their spread being closely linked to ventilation as evidenced by Li et al. (2007). Finally, this work provides a tool to assess indoor air quality in schools more generally and thus ensure that pupils are provided with a suitable healthy learning environment.

The application of the methodology to monitored CO_2 data is described in §2.2. §2.3 shows how the infection risk is estimated for a range of UK schools and how the resulting number of secondary infections varies between classroom and with the seasons. Finally, conclusions are drawn and discussed in §2.4.

2.2 Approach

We calculate the likelihood of airborne infection using a Wells-Riley model based on the works of Riley et al. (1978) and the later extension of Rudnick and Milton (2003). This is one of two approaches, along with the dose response mode, typically used to study the probability of airborne infection of respiratory diseases in indoor environments as reviewed by Sze To and Chao (2010). In the Wells-Riley model, a quantum of infection is used to quantify the required number of infectious particles to cause infections in 63% of susceptible people (Wells, 1955). An infectivity rate λ is estimated from $\lambda = Ipq/Q$, where I is the number of infectors, p the breathing rate, Q the ventilating flow rate and q the quanta generation rate. This allows the calculation of the probability of escaping the infection and the complementary, the risk of infection, can be found (Riley et al., 1978). Both the quantum and related quanta generation rate are modelling constructs which implicitly take into account the biochemical and virological processes associated both with the emission of viral particles and potential subsequent infection. The quanta generation rate is measured indirectly, often based on the analysis of outbreaks as further detailed in §2.2.4. The dose response model, on the other hand, uses a dose response relationship obtained experimentally to estimate the quantity of pathogen, the dose, required for infection. Although this approach avoids the approximations introduced by the calculation of a quanta generation rate, dose response data are difficult to measure accurately for humans, and even more difficult to obtain in the case of new diseases such as COVID-19. This encouraged the use of a Wells-Riley approach which has also been successfully applied to study the far-field transmission of other airborne diseases such as measles, rhinovirus and influenza (Bueno de Mesquita et al., 2020) making the model presented in this chapter widely applicable.

2.2.1 Calculating the risk of airborne infection in a regularly attended space

Focusing only on transmission via the airborne route, we wish to determine the likely number of secondary infections that might arise within a given classroom should an infected individual attend the space. We assume here that the number of infectors within the population of 'regular attendees' is constant, and that anyone 'exposed' does not become infectious during the duration of interest. For many indoor spaces, a wide variety of people come and go with varying frequency, which makes the prediction of risk via the Wells-Riley approach invalid. However, other indoor spaces are attended on a regular basis, day-in-day-out, and for significant durations each day by the same (or similar) group of people. We term these spaces 'regularly attended spaces'. Examples of these spaces include open-plan offices and school classrooms, the latter being the focus of this work.

In order to determine the likely number of secondary infections, S_I, from attendance of a classroom we must first specify some duration to investigate. This duration is often arbitrary and a number of reasonable choices could be made, each leading to a different likelihood. However, for regularly attended spaces, e.g. a school classroom, it is reasonable to assume that once an infected individual exhibits symptoms they cease attending. Thus one can assess the likelihood of infection occurring during the pre/asymptomatic infectious period, assuming the infector ceases to attend once symptoms develop. For COVID-19 the pre/asymptomatic infectious period is estimated to be 5–7 days, therefore allowing us to examine the likelihood of infection occurring over five consecutive working days, i.e. a duration $T_A = 5$ weekdays.

For infection to occur via the airborne route, a susceptible occupant of the classroom must breathe in infectious particles that are being carried by indoor air currents. The infectious particles, or aerosols, originate in the exhaled breathe of an infector and we assume that they remain in the 'rebreathed air', i.e. air that has already been breathed by another individual, which we take as a suitable surrogate for estimating airborne infection risk. The resulting probability of infection P_A was described by Riley et al. (1978), and then extended by Rudnick and Milton (2003), as

$$P_A = 1 - \exp\left(-\int_0^{T_A} \lambda \, \mathrm{d}t\right) = 1 - \exp\left(-\frac{I}{n} q \int_0^{T_A} f \, \mathrm{d}t\right), \tag{2.1}$$

where λ is the infectivity rate, I is the number of infectors, n is the number of occupants in the space, q is the emission rate of infectious doses, known as the quanta generation rate (see §2.2.4) and f is the fraction of rebreathed air. This in turn is defined as $f = (C - C_a)/C_p$, with C the monitored CO_2 within the space, C_a indicating the outdoor ambient level and C_p the concentration of CO_2 added through exhaled breath. This assumes that the main source of CO_2 is occupants, which is pertinent in classrooms and other spaces without combustion sources. In both the formulation of Riley et al. (1978) and the later form developed by Rudnick and Milton (2003), Wells-Riley models calculate the complement of the probability that no-one becomes infected within the duration (in our case T_A); hence the exponential terms in (2.1) are by no means representative of the response to a cumulative dose. The method presented here assumes implicitly that susceptible people attending the classroom will not become infected elsewhere or by another route than the airborne one. This assumption allows us to determine the contribution of a specific setting (in our case the classroom) and transmission route (in our case airborne) to the spread of the disease, and thus quantify how it might vary with environmental factors as the seasons change.

2.2.2 Accounting for spaces with varied occupancy

Equation (2.1) cannot be directly applied to durations which cover varied occupancy, such as a classroom with students coming in during the day and empty at night. Herein, we assumed that each classroom was regularly attended by $N = 32$ people (typical in UK state schools) and the occupancy n varied between zero and N according to the school timetable (which was typically accessed through online records). The number of infectious individuals I was set to be a constant fraction of the current occupants, i.e. $I = \phi n$, which can be linked to prevalence levels in the population. This was set to lead to $I = 1$ at full occupancy (with $\phi = I/n = 1/N$), and $I = 0$ when empty. With these assumptions, the likelihood P_A of airborne infection occurring due to a single infectious person attending school during a pre/asymptomatic period can then be calculated using:

$$P_A = 1 - \exp\left(-\frac{I}{n}\int_0^{T_A} \sigma(t)\; f\, q \;\mathrm{d}t\right), \tag{2.2}$$

where the term $\sigma(t)$ is a Heaviside operator based on whether the space is occupied or unoccupied, defined by

$$\sigma(t) = \begin{cases} 1\,, & \text{if the room is occupied } (n \neq 0), \\ 0\,, & \text{otherwise.} \end{cases}$$

This likelihood then provides the expected number of secondary infections that might arise via the airborne route from an infectious pre/asymptomatic person regularly attending school, via

$$S_I = (N - 1)\, P_A\,. \tag{2.3}$$

Both P_A and S_I were calculated for all 5 weekday periods within the data gathered on a rolling basis (a total of 15,000 pre/asymptomatic periods in total). The rolling period did not include weekends, for instance, a 5 weekday period starting on a Friday would run until the following Thursday and not take into account measurements for the Saturday and Sunday. In the rolling absolute number of secondary infections reported herein, we exclude all values for which the 5 weekday period contained any other unoccupied days, for example those that included weekdays that fell in the school holidays.

2.2.3 Calculating the rebreathed fraction from CO_2 measurements

The rebreathed fraction f was found from datasets of CO_2 levels monitored in schools, described in §2.3, where we took C_a to be the average CO_2 within the space between the hours of 05:00 and 06:00 each day (which corresponds to an unoccupied space which has had time to reach a baseline CO_2 level after occupation of the previous day). The levels above C_a each day correspond to the excess CO_2. Furthermore, we take $C_p = 42{,}250\,\mathrm{ppm}$, this value being determined by taking the averaged CO_2 generation rates (Persily and de Jonge, 2017) and pulmonary rate (Office of Research and Development's National Center for Environmental Assessments (NCEA), 2011) for children aged 6-11 for an activity level of 1.4 met which is between someone sitting quietly and light office work. The choice of a suitable value for the quanta generation rate q is discussed in §2.2.4.

We note that no assumption as to the distribution of CO_2 within the classroom has been made in deriving (2.2). As such, (2.2) provides the probability of airborne infection occurring at the sensor location, under the above assumptions, with the additional assumption that infectious airborne material is uniformly mixed but only within the rebreathed air. Hence, in order to employ (2.2) it is not required to assume that all of the air within the classroom is well-mixed but it is required to assume that the infected breath is evenly mixed within the rebreathed air. Highlighting this emphasises that when multiple CO_2 sensors are present within a single indoor space, differing CO_2 levels recorded by those sensors are not inherently contradictory and may indicate variations of risk within the space. CO_2 variations within an indoor space such as a classroom and the related judicious placement of monitoring sensors remains a challenging problem and is addressed in the following chapters.

In most settings CO_2 measurements must remain unobtrusive and thus placing sensors on walls is a frequent choice. This means that direct measurements of the exhaled breath are avoided along with their complications, detailed for example by Melikov and Kaczmarczyk (2007) and Kierat et al. (2018). Sensors should ideally be placed far away from windows or other openings and, crucially to study infection risk, they should also remain within the height where occupants breathe in. In this study, our data come from classrooms in which a single CO_2 sensor was placed at a height within the breathing zone, which was deemed suitable to ensure ventilation control. The impact of the exact location of the sensor is further compounded by the choice of a constant and uniform concentration of CO_2 and the calculation of secondary infections. Indeed, when estimating the number of secondary infections, via (2.3), one is assuming that the airborne infection risk is appropriate for all susceptible occupants; an assumption which is valid in the case that all the air within the space is well-mixed. Although this is a significant assumption, it also prevents making further assumptions about the geometry of the classroom or the ventilation design and is alleviated by a number of other unknowns, a key one being the quanta generation rate discussed in the following section.

2.2.4 Determining appropriate quanta generation rates for school classrooms

For all infection risk modelling based on the Wells-Riley approach, the quanta generation rate, q, is invariably the hardest input parameter to quantify and the resulting uncertainty is typically a dominant factor, as it varies with disease, individuals and occupant activity levels (Bueno de Mesquita et al., 2020; Noakes and Sleigh, 2009). Given this uncertainty, we discuss the changes in our results that would arise from other feasible choices of the quanta generation rate. Estimates of quanta generation rates rely on indirect measures and have been made via a variety of methods for a host of airborne diseases. Understandably, for SARS-CoV-2, choices are limited but we apply the values reported in the work of Buonanno et al. (2020) who consider the likely values of q based on viral load and respiratory activity. As a value which we deem to be typically appropriate for school classrooms, we take a value of $q = 1\,\text{quanta/h}$ — this is obtained by taking $c_x = c_i\, c_v \approx 7 \times 10^6\,\text{RNA/ml}$, where $c_i = \{0.1, 0.01\}$ is the ratio between infectious quantum and the infectious dose expressed in viral RNA copies, and $c_v = \{7 \times 10^7, 7 \times 10^8\}\,\text{RNA/ml}$ is the viral load measured in sputum. These values assume that, for most of the time within school classrooms, the students are sitting breathing with perhaps a small number vocalising: the data for whispered counting fall between these two activities and are closer to breathing. As such, for our base case, we take data for whispered counting from Buonanno et al. (2020) and use their results to map our selected values of c_x to values of quanta generation rates q. We note that all of the work underlying the results of Buonanno et al. (2020) concerned adults and, for COVID-19, it is not known if appropriate values for the quanta generation rates differ significantly between adults and children. The quanta generation rate might also be expected to vary for a single individual during the course of the infection, but once again, information is limited in the case of SARS-CoV-2 and its variants and we take q to be constant in time.

Throughout our results we comment on the changes that would arise if the students within the classroom were (on average) all vocalising/talking (e.g. a noisy classroom), taking again $c_x \approx 7 \times 10^6$ RNA/ml gives $q \approx 5$ quanta/h. For our key findings, i.e. the seasonal variation in airborne infection risk, we also comment on the results of assuming much higher quanta generation rates of $q \approx 20$ quanta/h and $q \approx 100$ quanta/h which may represent higher viral loads or higher levels of activity.

2.3 Results

The data presented originate from recently built or renovated school classrooms, since these are more likely to have existing CO_2 monitoring provision installed. Overall, 45 spaces are monitored from within 11 different schools (8 primary and 3 secondary) and span the period November 2015 to March 2020 (when schools were subject to a UK wide lockdown), see Table 2.1 for details. The data originate from schools in England, as far north as Yorkshire, as far south and west as Somerset, and as far east as Kent; the data are sourced from schools in a mix of urban and rural settings. The ventilation system varied between classrooms but was, in all cases, a hybrid ventilation system which switched between naturally driven ventilation and mechanically driven ventilation modes depending on the conditions. The ventilation provision was controlled by automatic operation of louvres, vents, and fans. These classrooms were also typically fitted with additional windows or vents that could be opened manually by the classroom occupants.

School	Type	County	Rooms	Data span
1	Primary	Yorkshire	22	Nov/15 – Mar/19
2	Secondary	Berkshire	1	Nov/19 – Mar/20
3	Primary	Somerset	1	May/17 – Mar/18
4	Primary	Surrey	1	Dec/17 – May/18
5	Primary	Cambridgeshire	2	Aug/17 – Jan/18
6	Primary	Not disclosed	3	Dec/18 – Feb/19
7	Primary	Essex	4	Oct/16 – Dec/17
8	Secondary	Kent	1	Mar/18 – Apr/19
9	Primary	Surrey	4	Aug/17 – Aug/18
10	Primary	Kent	1	Aug/17 – Jul/18
11	Secondary	Hertfordshire	5	Sep/18 – Mar/20

Table 2.1: Schools where the monitored CO_2 data originate from.

2.3.1 The airborne infection risk in one classroom in January and in July

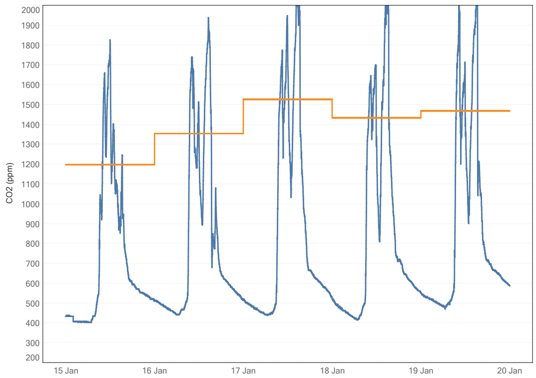

(a) CO_2 levels in January

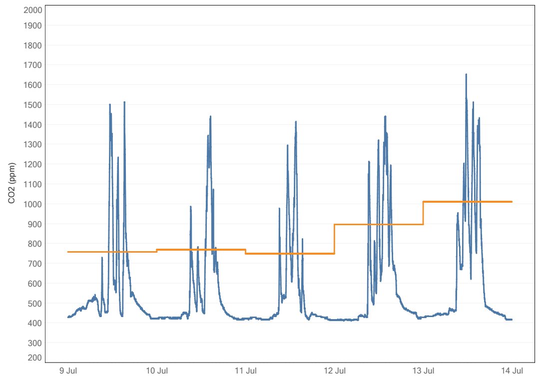

(b) CO_2 levels in July

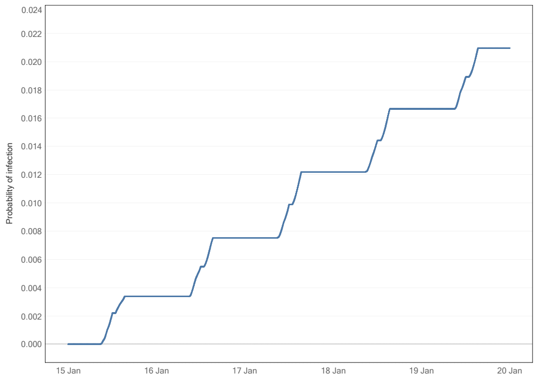

(c) Cumulative likelihood of infection in January

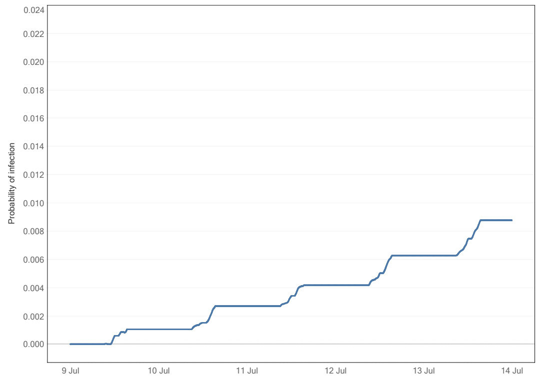

(d) Cumulative likelihood of infection in July

Figure 2.1: The variation in CO_2, top panes a) and b), over a pre/asymptomatic period within classroom Y1-2 and, bottom panes c) and d), the corresponding probability of infection (assuming $q = 1\,\text{quanta/h}$). The orange horizontal lines in a) and b) show the daily averaged CO_2 during the occupied periods (defined for school classrooms as between 09:00 and 16:00, see Department for Education, 2018). The left-hand panes a) and c) show data for a pre/asymptomatic period in January 2018; the right-hand panes, b) and d), show data for a similar period in July 2018.

To determine the airborne infection risk in schools we begin by taking one classroom (selected at random) from within our set (45 classrooms in total). For the year 2018 we plot the variation in CO_2 level in the upper panes of Figure 2.1 for two periods (of 5 school days each) one in January and one in July. The data from January show that, in adherence with the Department for Education (2018) BB101 guidance for a naturally ventilated classroom the daily average CO_2 level is kept below around 1,500 ppm, and the CO_2 levels only spike above 2,000 ppm for very short periods. This was the case for all classrooms that we examined throughout almost all the periods we had data for (spanning November 2015 to March 2020), which enables us to make some comments on whether adherence to the existing guidance might provide adequate airborne infection risk for COVID-19.

What is stark from Figure 2.1 is the decrease in CO_2 levels, both in terms of daily mean and peak values, in July compared with January. It can be inferred from the gradual drop towards ambient CO_2 levels during January nights that this classroom can be relatively well-sealed when required by the thermal conditions and there is no need for ventilation. This is in contrast to the July data when, presumably, warmer outdoor conditions did not require all openings to be shut at the end of the day and CO_2 levels reached ambient levels for most of the night. Daytime excess CO_2 levels in July typically attain values around half those in January despite occupancy and activity levels which are not expected to vary significantly. This is clear evidence of increased ventilation levels in July in response to warmer weather that are either provided automatically by the ventilation system and/or by occupants' intervention, e.g. opening windows. Indeed, using the daytime average CO_2 concentration as a steady state concentration, the average ventilation flow rate can be calculated to be 8 L/s/person for the July data, more than twice as much as in January, when it is equal to 3.4 L/s/person. In any event, the change in risk is significant with the cumulative likelihood of airborne infection of COVID-19 being more than twice as high in January compared with July. As expected, in both cases the risk increases only during occupied hours producing a 'staircase' of increase likelihood over the pre/asymptomatic period. Accounting for an accuracy of ± 50 ppm in the CO_2 measurements leads to an uncertainty of ± 0.0012 in the predicted likelihood of infection after 5 working days shown in Figure 2.1. We note that the absolute probabilities of airborne infection remain low in both cases, implying that changes in the likelihood will exhibit an approximately linear response to changes in the underlying parameters, e.g. the quanta generation rate, and that results for the relative risk will be robust while the probabilities remain low.

2.3.2 The variation in airborne infection risk within one school

Figures 2.2a) and 2.2b) show the average number of secondary infections for 15 different classrooms within a typical primary school over January 2018 and July 2018, respectively. These are detailed in Table 2.2 and were chosen because they allowed a direct comparison of similar spaces within the same school for the chosen months. It is apparent from the figure that, as one would expect, the number of airborne secondary infections in winter (January) greatly exceed those of summer (July). For the classrooms shown the average S_I in January is 0.60, falling to 0.28 in the July of that year; this is assuming a quanta generation rate of $q = 1$ quanta/h and we return to a discussion of these values in §2.3.3. The change in risk between January and July is not easily classified; many of the relatively risky classrooms (defined as those with the highest number of

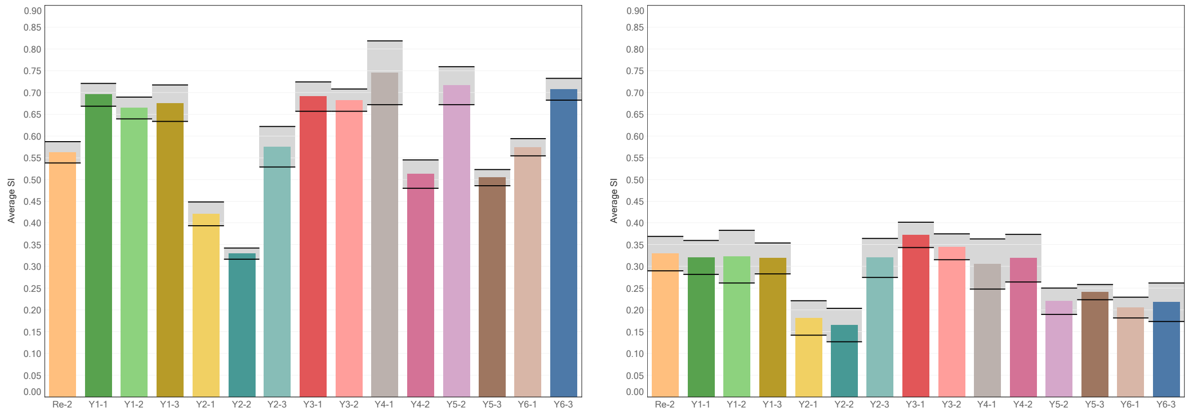

(a) January 2018

(b) July 2018

Figure 2.2: Variation in average number of secondary infections (S_I) between classrooms of a primary school for the most and least risky months (with the highest and lowest number of secondary infections respectively: January and July) in the school year 2017–2018.

Classroom	Elevation	Floor area (m^2)	Year	Age group
Re-2	Ground floor	67	Reception	4 to 5
Y1-1	Ground floor	66	Year 1	5 to 6
Y1-2	Ground floor	67	Year 1	5 to 6
Y1-3	Ground floor	66	Year 1	5 to 6
Y2-1	Ground floor	67	Year 2	6 to 7
Y2-2	Ground floor	67	Year 2	6 to 7
Y2-3	Ground floor	66	Year 2	6 to 7
Y3-1	First floor	60	Year 3	7 to 8
Y3-2	First floor	60	Year 3	7 to 8
Y4-1	First floor	60	Year 4	8 to 9
Y4-2	First floor	60	Year 4	8 to 9
Y5-2	First floor	61	Year 5	9 to 10
Y5-3	First floor	60	Year 5	9 to 10
Y6-1	First floor	60	Year 6	10 to 11
Y6-3	First floor	60	Year 6	10 to 11

Table 2.2: Description of the 17 primary school classrooms considered.

secondary infections) in January remain so in July (e.g. Re-2, Y1-1, Y1-2, Y1-3, Y3-1 and Y3-2), some of the relatively low risk classrooms in January remain so in July (e.g. Y2-1 and Y2-2), but others change from being relatively risky to relatively low risk (Y4-1, Y5-2 and Y6-3) or from being relatively low risk to being relatively risky (Y3-3). Interestingly, we found no obvious correlation between these differences and the age group, the location of the classroom in the school, e.g. whether located on the ground or first floor, or the proximity of one classroom to another.

What is also noticeable within Figure 2.2 is the relatively small variation in S_I within any given classroom within the same month. We choose to quantify this by the coefficient of variation (calculated as the ratio of the unbiased estimate of the standard deviation relative to the mean) which is, on average, 6% in January 2018 and 15% in July 2018 for the classrooms shown (see the grey bars in Figure 2.2). However, the variation between classrooms even within the same month is significant. The coefficient of variation of S_I for the classrooms shown is 20% in January and 27% in July, with the most risky classroom being at least twice as risky as the least in both January and in July, in spite of the fact that all 15 classrooms from this particular school have the same ventilation system installed by the same contractor.

Finally, we note that variations are also observed between school years: a risky classroom one year being relatively less risky in the following academic year for instance. This could be explained by changes in the ventilation system, with adjustments being made to the building management system or interventions from the occupants, or could be due to a different use of the space in a different academic year; for example, a class with either more or less pupils or different activities, or could be the result of differing weather conditions, which highlights the complexity in assessing the risk of airborne transmission.

2.3.3 The seasonal variation in predicted airborne secondary infections in schools

The variation in the absolute number of predicted secondary infections with season is presented in Figure 2.3a). For each of the 45 classroom the rolling absolute S_I was calculated for all available data and values pertaining to each calendar month were then averaged (August is not represented as it falls during the school holidays). As can be seen, at the start of each academic year (September) the risk of airborne infection rises as the UK enters winter weather, with the risk peaking in January and February after which the risk gradually subsides as more temperate spring and summer weather is approached. Taking a value for the quanta generation rate of $q = 1\,\text{quanta/h}$ (Buonanno et al., 2020) yields the values for the absolute airborne number of secondary infections shown in Figure 2.3. These lie reassuringly below unity (with a mean of $S_I = 0.38$), and may be representative for the airborne spread of COVID-19 in relatively quiet classrooms. However, classrooms in which children are expected to be more vocal or active (corresponding to $q = 5\,\text{quanta/h}$) then one could expect the airborne S_I to increase by a factor of approximately five, i.e. $S_I \approx 2.0$. There is uncertainty as to appropriate values for the quanta

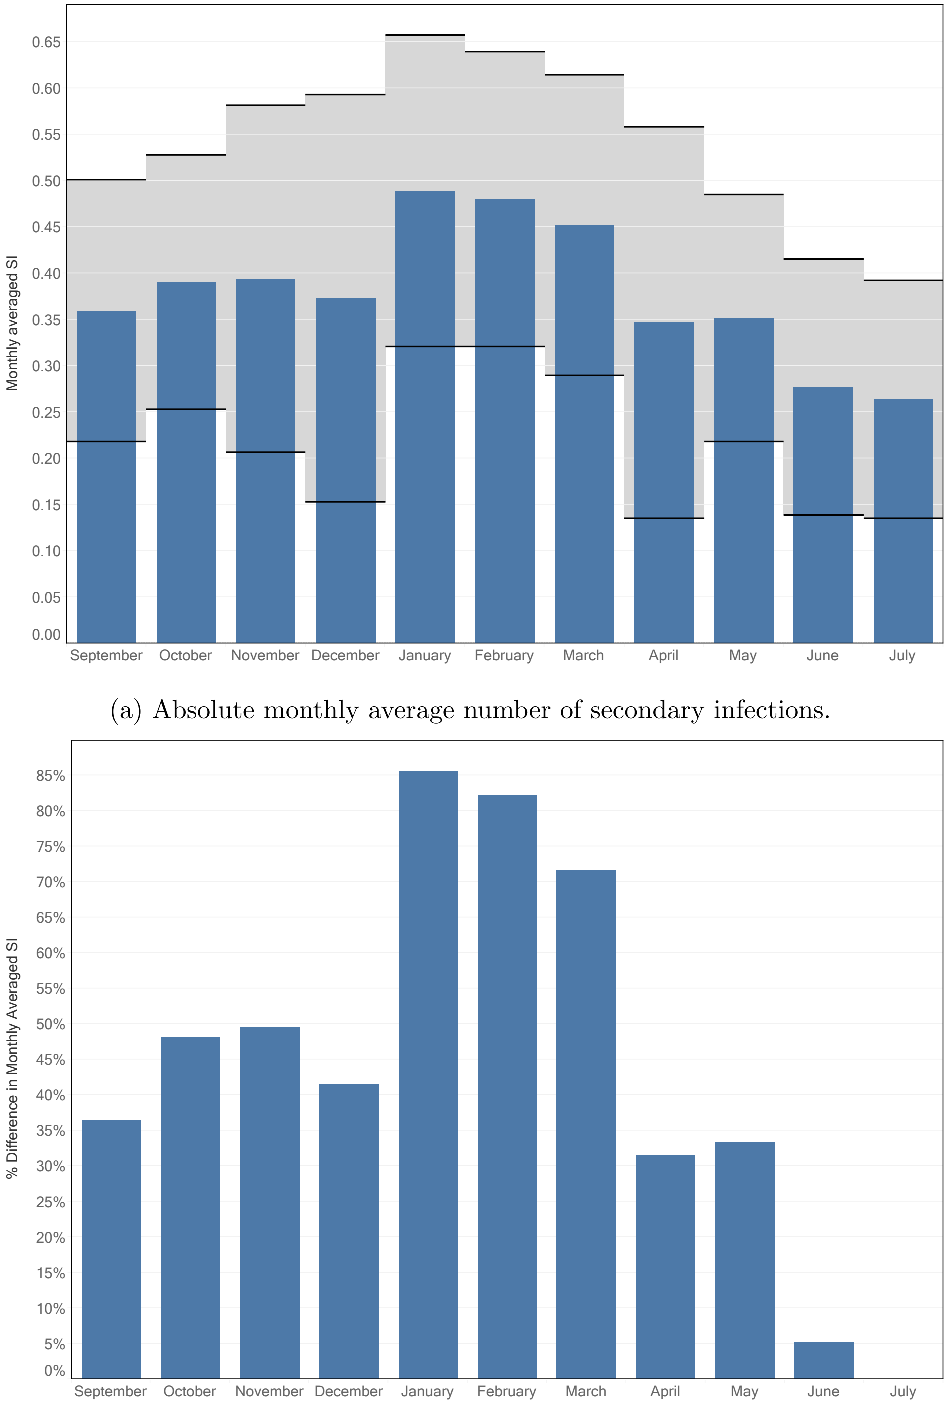

(a) Absolute monthly average number of secondary infections.

(b) Percentage difference with respect to a summer month (July).

Figure 2.3: The variation in a) the absolute, and b) the relative, average monthly number of secondary infections across the seasons. Data shown represent CO_2 monitoring in 45 classrooms (within 11 different schools) and span the period November 2015 to March 2020 (see Table 2.1).

generation rate for COVID-19; with current evidence suggesting that a relatively low number of high spreading events make an unusually important contribution to the disease spread (Adam et al., 2020). Consequently, these absolute risk results should be treated with caution and focus should be drawn, instead, on the relative risk.

Before examining the seasonal variation of the relative risk it is worth considering the variations in risk within each month. The average coefficient of variation of S_I in all 45 classrooms within each single month is approximately 44%. This is approximately twice as large as the coefficient of variation for S_I within each single month from a single school with 15 classrooms suggesting that, in addition to the variations intra school (see §2.3.2), significant variations also exist between the different schools. We note also that all the schools had broadly similar ventilation strategies (i.e. hybrid ventilation systems) and demonstrated adherence to existing ventilation guidance. Should schools with very different ventilation provision be included then one might expect wider variations: mechanically ventilated schools classrooms are set more demanding ventilation provision within the guidance, while schools with uncontrolled ventilation are likely to vary more widely due to the greater influence of occupants' interventions.

The relative number of airborne secondary infections is presented in Figure 2.3b), in which the data from Figure 2.3a) are shown normalised by the risk in summer (here taken to be July). As the academic year begins we see that average S_I in classrooms might be expected to be about 30–40% above the levels that would be seen in summer. Broadly speaking, levels then steadily increase towards the Christmas holidays and in January and February, average S_I peak at 80–90% above July levels. Levels remain high in March and then fall sharply back towards summer levels. This seasonal pattern in S_I is consistent with intermediate ventilation provision in autumn and spring weather, highest ventilation in summer, and minimum ventilation in winter to avoid heat loss.

We tested the seasonal variation in number of airborne secondary infections in schools for varied levels of quanta generation rate. For quanta generation rates $0.1 \leq q \leq 5\,$quanta/h the quantitative predictions illustrated in Figure 2.3b) remain almost unchanged. For much higher quanta generation rates $20 \leq q \leq 100\,$quanta/h the qualitative trends in the data remain but the precise values change somewhat, for example for $q = 100\,$quanta/h, S_I in January would be predicted to be 43% higher than the July value.

2.4 Conclusions

Based on CO_2 monitored within 45 classrooms from across England we have shown that airborne infection risks in schools are likely to vary significantly with the season. Purely based on changes in environmental conditions the expected levels of secondary infections in winter (e.g. January and February) are nearly double those in summer (e.g. July). We suggest that, while our results for schools with relatively recently installed ventilation provision indicate airborne infection risk for COVID-19 may be low, ventilation provision for wintertime should be assessed and reviewed. It is important to note that the analysis in this chapter is based on historical CO_2 data measured before significant restrictions were introduced in the UK to control COVID-19 transmission. The study therefore models the likely risk under normal operating conditions for the monitored schools. The UK Department for Education has issued guidance to schools recommending additional ventilation, and hence it is likely that airborne risks in following winters may be lower than those predicted here. We have also shown that airborne infection risk can vary widely within a school, even in the case that the same ventilation system is installed by the same contractor. Where these differences might increase airborne infection risk, occupants should be encouraged to review their behaviours to ensure that these are appropriate. These differences may be even more exaggerated in schools with no controlled ventilation provision, which represents the majority of the existing school stock within the UK. Moreover, our assessment of the airborne infection risk as low may not apply to these schools but monitoring of CO_2 would provide the evidence required to make an assessment.

The seasonal variations in airborne infection risk described in this chapter account only for those due to changes in the indoor environment which arise due to ventilation and/or occupants' behaviours due to varying outdoor temperatures during the year. It is likely that the virus and the human response to it will also exhibit some seasonal variations and, as such, the risks of airborne infection might be compounded. If these increase the risk of transmission, e.g. due to a weaker immune response in winter shown by Dopico et al. (2015) or to changes in humidity as described for example by Marr et al. (2019) for influenza, our results would in effect be a lower bound for seasonal variation in airborne infection risk. In addition, the risk presented herein refers to the airborne route of transmission only. The two other major disease transmission routes might also be shown to exhibit variations in infection risk, which would contribute to observed changes in total number of infections throughout the year. The estimates of quanta generation rates were also done early in the COVID-19 pandemic and might not represent the evolving behaviour of the virus transmission as new variants appear.

Although the focus is on COVID-19, the method we present here is applicable to all airborne diseases. We have shown that, in particular, the variations in relative airborne infection risk with ventilation conditions, and hence seasons, hold over a wide range of quanta generation rates and so should be directly applicable to a wide range of diseases.

The approach deployed within this chapter relies on the measurement of CO_2 at a single location in classrooms. With CO_2 sensors becoming widespread in schools and beyond, in part as a response to the pandemic, it is likely that a large amount of measurement data will soon become available. It presents a great opportunity to gather information on the state of ventilation in schools, the potential exposure of pupils in classrooms, the risk of airborne infection and the resulting indoor air quality. Single point measurements were used here to represent the concentration in the entire room and, implicitly by calculating the number of secondary infections, the classroom was assumed to be well mixed. In naturally ventilated spaces particularly, this is unlikely to be the case as significant temperature variations can be expected. This poses the question of the accuracy of a single point measurement to assess exposure and thus risk in naturally ventilated classrooms. This issue is the focus of the following chapters and is addressed using Computational Fluid Dynamics (CFD) simulations. These aim to quantify the uncertainties introduced by the use of CO_2 measurements to assess risk although other factors, chiefly the quanta generation rate (predicted to vary by at least three orders of magnitude by Mikszewski et al., 2021), are likely to remain sources of significant uncertainty in the prediction of airborne infection risk. Caution is therefore advised when providing absolute estimates of the likelihood of infection and estimates of the uncertainty should be provided. In a well-defined scenario however, this method is extremely valuable because, apart from those mentioned previously, it addresses most other uncertainties. The predicted value of the absolute risk can therefore be meaningful and could for instance be considered along the infection likelihood arising from the two other major transmission routes.

Chapter 3

Computational set-up of a classroom simulation

In Chapter 2, CO_2 measurements were used as a proxy to estimate the risk of far-field airborne infection in classrooms. Implicitly via the calculation of the number of secondary infections, this assumed that a single point measurement of CO_2 was representative of the entire CO_2 level across the whole classroom. In effect, this result assumed that the air within the room was well-mixed. Whilst this assumption is often used in Indoor Air Quality (IAQ) and airborne infection studies, it may not hold in naturally ventilated classrooms where the flow is buoyancy-driven and variations of temperature are expected in the space. It was then deemed pertinent to investigate how valid the well-mixed assumption is in practice so as to gain a deeper understanding of the uncertainty in the risk predicted in Chapter 2.

To examine how exposure varies within a naturally ventilated room in more detail, Computational Fluid Dynamics (CFD) simulations were used. In this chapter, the details of how the CFD simulations were formulated (encompassing the choice of solver, computational domain and boundary conditions) are described. CFD simulations were used to model a generic classroom in which both the general CO_2 levels and a potential infected breath source were modelled. Their distribution was calculated within the classroom and was used to assess the uncertainty introduced by using a single point measurement in the risk predicted in Chapter 2. The focus of this chapter is directed at a scenario where mixing is limited and the spatial variations in concentration within the classroom are accentuated, thus providing an upper bound for the error introduced by a well-mixed assumption. When setting up the CFD simulations, a main concern was to find a trade-off between solution accuracy and computational complexity.

The objective was then to set up the simulations such that they were both precise enough to study exposure within the room but also not overly computationally costly so that they could run in a reasonable time for a number of scenarios encompassing a range of ventilating flow rates.

The chapter is structured as follows: the selection and design of a generic UK classroom is described in §3.1. The chosen numerical set-up is then detailed in §3.2 and the choice of a transient solver is discussed in §3.3. §3.4 and §3.5 respectively describe the chosen computational domain and grid used. Finally, conclusions are drawn and discussed in §3.6.

3.1 Characterising a generic UK classroom environment

The set-up investigated via CFD simulations aimed to represent naturally ventilated classrooms in the UK in wintertime in order to study the distribution of CO_2 and infected breath in the space. A number of assumptions had to be made to simplify the set-up ensuring however that estimates of CO_2 exposure remained within the expected order of magnitude and within the uncertainties of an infection risk model.

UK classrooms, whether in primary or secondary schools, vary significantly in their design and in order for the results to be applicable to a wide range of classrooms, it was decided not to base the simulations on a specific classroom. A generic classroom therefore had to be characterised, and this was achieved by ensuring it followed the framework defined by the Department for Education (2018) and Department for Education (2014) guidelines which all UK classrooms should meet.

It was decided to focus on a primary school classroom as younger children are more vulnerable to indoor contaminants and therefore a more at risk population. Unlike secondary schools, the same group of pupils can also generally be expected to remain in the same room during the day. This meant that longer periods of occupation could be considered in the simulations, increasing the likelihood of a steady state to be reached. Moreover, the infection risk model developed in Chapter 2 is applicable to schools in which pupils remain within the same classroom day-in-day-out — an operational procedure typical of UK primary schools, and less common for secondary schools. The occupants of the classroom were therefore assumed to be of primary school age (between the ages of 5 and 11 in the UK) when defining their CO_2 and heat outputs.

A single classroom was simulated with links to the rest of the building modelled only through the ventilation openings. Heat losses at the walls, air infiltration or thermal mass were not taken into account. Particular emphasis was directed at describing the ventilation system and how it impacted exposure in the classroom.

The ventilation design for UK classrooms varies widely but is predominantly natural ventilation. More than 90% of energy display certificates issued for school buildings in England and Wales from October 2008 to June 2021 describe their ventilation system as natural ventilation (Ministry of Housing, Communities & Local Government, 2021). A wide range of natural ventilation systems exist, and many of these have been deployed in schools (i.e. manual side windows, louvres, windcatchers). The goal of this work was not to look at a specific ventilation configuration, instead a generic buoyancy driven ventilation system was considered. For that reason, the effects of the openings were considered only through the ventilating flow they provide. These openings were modelled as flat open surfaces on the floor and ceiling of the room, therefore also avoiding the consideration of bi-directional flows through ventilation openings in the simulations. Using floor and ceiling openings is expected to be applicable in most buoyancy dominated scenarios with high and low-level vents and where wind doesn't directly oppose the flow (Hunt and Linden, 1999) or enhance mixing Carrilho da Graça (2018).

Two ventilation configurations were investigated. In each of these, the opening areas, heat input and therefore overall flow rate were kept constant but the position of the top vent was changed leading to:

(i) a buoyancy-driven cross-flow configuration with vents at opposite ends of the room, and

(ii) a buoyancy-driven single-sided configuration with vents on the same side of the room.

In both configurations, the vents will remain on the floor and ceiling, thus deviating from the traditional, often wind-driven, cross-flow and single-sided configurations where the opening will be found on walls. In the remainder of this thesis, each configuration will be referred to directly as either 'cross-flow' or 'single-sided' to refer to the changes in the relative horizontal position of low and high-level vents and their effect on the flow field. The heating season was investigated since, based on measurements in operational classrooms, winter was identified as the period of greatest risk for airborne infection (see Chapter 2) due to the lower inferred ventilation levels in the classrooms being observed (presumably, caused by the colder temperatures reducing the likelihood of windows being opened).

In the the remainder of this thesis, the effects of wind on ventilation are not taken into account illustrating the limiting case in which ventilation is provided only through temperature differences. Here, the ventilating flow is driven solely by buoyancy which is set by the classroom's internal heat sources (in this case the occupants and the heating provision). Convection was assumed to be the dominant heat transfer process and radiation and conduction were neglected. Chenvidyakarn (2013) argues that this is a reasonable assumption for naturally ventilated spaces where the building fabric is well insulated and the movements of air are strong, which is the case here as the room was considered independently and the focus was on wintertime when temperature differences between indoors and outdoors are largest in the UK. Including the effects of radiation in numerical simulations was shown to increase the solution accuracy when compared to displacement ventilated chambers experiments, but it also increased the complexity of the numerical set-up, notably in the choice of boundary condition and turbulence model (see, e.g. Deevy, 2006; Deevy and Gobeau, 2006; Chang, 2016). Since radiation was shown to mainly increase mixing, it was decided to neglect radiative effects with the knowledge that the results presented in this thesis represent a scenario where mixing within the classroom remains limited.

Classrooms are subject to a wide variety of heat sources, which constitute buoyancy sources to drive the flow. These heat sources during wintertime include: the designed heating provision, typically in the form of heating devices at low-level present over some fraction of the walls and floor; solar gains, covering patches of the floor and walls, typically at low-level; and occupants, i.e. the staff and pupils. It was assumed that all heat was released at low-level, including the heat generated by the pupils and staff. The effects of the human thermal plume, for instance, were not explicitly modelled. Although previous work (such as the numerical studies of Rim and Novoselac, 2010) showed that the human thermal plume can impact near-field exposure, using a floor heat source was deemed a sensible approximation to study ventilation at the room scale and far-field airborne exposure. Kaye and Hunt (2010) report that when the area over which low-level localised heat sources are input reaches approximately 15% of the total floor area, their interaction is sufficient to result in the heat within the room being relatively well-mixed and reasonably approximated by the equivalent distributed heat load over the whole floor area. In addition, classrooms are busy places with regular movement of staff and people within (and thus the heat sources they present) and this is especially true for the primary school classrooms which constitute our focus. As such, we take an approximation to the precise distribution of heat loads within any particular classroom by assuming the heat loads to be uniformly distributed on the floor of the classroom. This was designed to represent all internal heat loads in the classroom both from the heating system and occupants. Particular focus was drawn to the heating season in the UK and the heat input from radiators in the room was adjusted to enable the provision of a comfortable thermal environment to the pupils and staff.

Since our focus is on establishing the variations in steady states, the occupants' breathing was modelled as a steady source of CO_2, represented by the addition of a passive scalar. Additional tracers were introduced to represent infected breath arising from a specific individual (but not explicitly modelling viral particles) which is further described in Chapter 5. Due to our interests, the transient breathing pattern was not modelled, which simplified the choice of boundary conditions as the specific breathing rates did not have to be specified. This was deemed to be a suitable choice as the focus is on far-field exposure at the room level. In addition, the infected aerosols were not modelled as particles but as passive tracers, like the overall CO_2. Although this approximation might not capture the dynamics of infected aerosols, especially in the near-field, Ai et al. (2020) argued that, due to the complexity of approximating the size distribution and composition of viral droplets, the uncertainty of a tracer gas model is likely to be of the same order of magnitude as particle simulations. In addition, Pei et al. (2021) showed that passive tracers were representative of smaller aerosols, which are less likely to deposit and may be responsible for further-field airborne transmission. The use of passive tracers to represent infected breath implied that the natural decay of infected aerosols was not captured, although it was representative of a limiting case with increased likelihood of exposure.

3.1.1 Reference scenario

Following these assumptions the exact reference scenario to be simulated was selected following the available guidance: the classroom was assumed to have 32 occupants including 30 pupils and 2 staff following the BB101 guidelines (Department for Education, 2018). The outdoor temperature, internal heat gains and vent areas were set to result both in a desirable range of indoor temperatures and CO_2 concentrations that adhere to BB101 (Department for Education, 2018). This standard requires the average CO_2 concentration in a naturally ventilated classroom to be below 1,500 ppm and that the flow rate should not fall below 5 to 8 L/s/person. A temperature of 20°C during the heating season with a maximum of 25°C at full occupancy is also recommended by BB101 (Department for Education, 2018). The classroom was sized using the minimum requirements given by BB103 (Department for Education, 2014) and also corresponds to the reference case of Jones et al. (2021): the surface area was set to $55\,\mathrm{m}^2$, with a ceiling height of 2.7 m. An illustration of the set-up considered is shown in Figure 3.1. Ventilation is driven through high and low-level openings located on the floor and ceiling of the classroom. The bottom vent has an area A_B of $0.4\,\mathrm{m}^2$ ($0.8\,\mathrm{m} \times 0.5\,\mathrm{m}$) with its centre located at coordinates (1.4, 1.25, 0). Two ventilation configurations are investigated by changing the location of the top-level vent, while keeping its area A_T constant and equal to $0.2\,\mathrm{m}^2$ ($0.5\,\mathrm{m} \times 0.4\,\mathrm{m}$). In the cross-flow

(CF) configuration, it is positioned on opposite ends of the low-level opening centred around (8.75, 4.3, 2.7). In the single-sided (SS) configuration, the opening is positioned on the same side as the bottom opening at coordinates (1.25, 4.3, 2.7). Both configurations are illustrated in Figure 3.1.

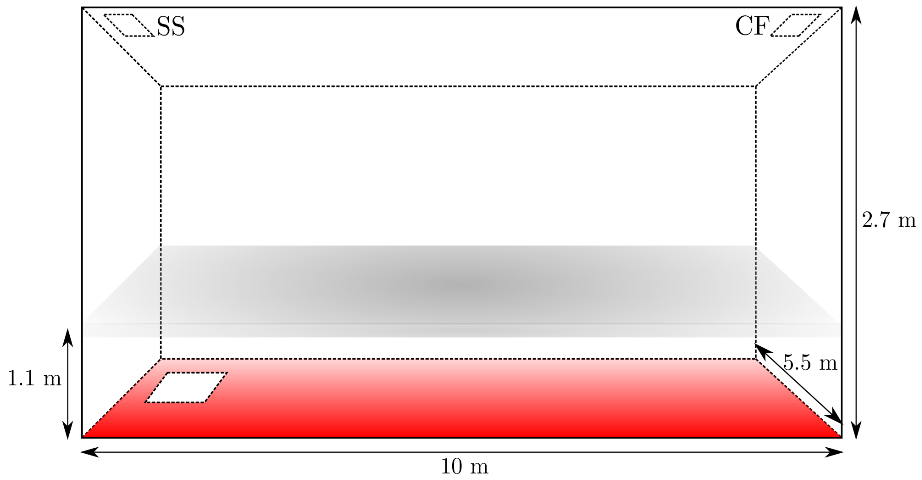

Figure 3.1: Illustration of the set-up used to represent a generic naturally ventilated UK classroom. Two high-level openings are considered to investigate both a buoyancy-driven cross-flow (CF) and single-sided (SS) ventilation system.

The 32 occupants in the room were each assumed to produce 60 W of heat (Department for Education, 2018) and generate $G = 3.35 \times 10^{-6}\,\mathrm{m^3/s}$ of CO_2, which is the average for primary aged school children found by Persily and de Jonge (2017), assuming an activity level of 1.4 met (which corresponds to an increase in the child's activity proportional to the increase in activity level of an office worker writing or standing, compared to the levels for rest). The heat, representing both the human presence and the input of the designed heating provision, was supplied at floor level within the classroom (as rationalised above). Tracers to represent CO_2 were input over the classroom between heights of 1.1 m and 1.2 m, using the typical breathing height of seated primary school pupils from the Department for Education (2018) BB101 guidance. The ambient outside temperature T_a was taken to be 5°C, assumed to be typical for the coldest months in the UK.

The overall heat input (W_c) and opening sizes (A_B and A_T) were defined using predictions of Gladstone and Woods (2001) and summarised by Chenvidyakarn (2013) for a distributed heat source, to provide a suitable ventilating flow and thermal environment. The flow rate was predicted using

$$Q = A^{*2/3} \left(\frac{g\alpha}{\rho_a c_p} \right)^{1/3} W_c^{1/3} H_c^{1/3} \, , \tag{3.1}$$

where g is the gravitational acceleration, α the thermal expansion coefficient, ρ_a the ambient density, c_p the specific heat capacity, W_c the heat input to the room and H_c is the classroom height. The different inputs and their respective values are summarised in Table 3.1. A^* is the effective vent area that characterises both vent areas and the effect of flow contraction at the openings. It is calculated from

$$A^* = \left(\frac{1}{2c_{d,B}^2 A_B^2} + \frac{1}{2c_{d,T}^2 A_T^2} \right)^{-1/2} \tag{3.2}$$

where A_B and A_T are respectively the bottom and top vent areas. The discharge coefficients $c_{d,B}$ and $c_{d,T}$ were taken to be the same for both openings and calculations were performed for several values of the discharge coefficients within the range $0.6 \leq c_d \leq 1.0$ (the typical value used in experiments Gladstone and Woods, 2001), for which a discharge coefficient of unity assumes no losses due to contraction at the opening.

The resulting temperature in the room was found from

$$T_c = T_a + W_c^{2/3} \left(\rho_a c_p A^*\right)^{-2/3} \left(g \alpha H_c\right)^{-1/3} \ , \tag{3.3}$$

where T_a is the ambient temperature, and the classroom CO_2 concentration at steady state from

$$C_c = \frac{NG}{Q} \times 10^6 + C_a \ , \tag{3.4}$$

with N the number of occupants, G their respective CO_2 generation rate and C_a the ambient CO_2 concentration, typically expressed in ppm.

Iterating over different values to achieve the best environment both in terms of ventilating flow and thermal comfort, the opening sizes and heat input were set to the values given in Table 3.1, with these parameters resulting in the predicted room averaged values shown in Table 3.2. For all three values of c_d considered, the ventilating flow allowed for sufficient values of ventilation as defined by Department for Education (2018), with flow rates comprised between 5 and 8 L/s/person. Specifically a steady state CO_2 value below 1,000 ppm was also achieved in all cases, which is deemed to be adequate for good indoor air quality in classrooms (Lowther et al., 2021).

Parameter	Symbol	Unit	Input
Floor surface area	S_c	m^2	55.0
Classroom volume	V_c	m^3	148.5
Classroom height	H_c	m	2.7
Number of occupants	N	-	32
Inlet area	A_B	m^2	0.40
Outlet area	A_T	m^2	0.20
Heat input	W_c	W	6200
CO_2 generation rate	G	m^3/s	3.35×10^{-6}
Outdoor temperature	T_a	°C	5
Ambient density	ρ_a	kg/m^3	1.268
Expansion coefficient	α	K^{-1}	0.00362
Heat capacity	c_p	J/kg/K	1005

Table 3.1: Inputs used to define the classroom simulation.

Well mixed prediction	Air Changes per Hour (ACH)	Q_p (L/s/person)	T_c (°C)	C_c (ppm)
$c_d = 0.6$	5.35	6.90	27.0	886
$c_d = 0.7$	5.93	7.64	24.9	838
$c_d = 1$	7.52	9.69	20.7	746

Table 3.2: Predicted steady state flow rate, room averaged temperature and CO_2 concentration in the classroom for different discharge coefficients c_d, assuming a distributed heat source and following Gladstone and Woods (2001).

3.2 Setting up a CFD simulation

The goal of this thesis is to explore how exposure within a naturally ventilated classroom might vary across a classroom and whether CO_2 can be used as a suitable proxy for far-field exposure. For this reason, CFD simulations were adopted as they provide access to the full three dimensional data fields and not just discrete measurement points, which could then be used to compare actual exposure to a potential single point measurement. CFD simulations also allow the specification of well defined and controlled conditions without relying on dynamic similarity, like water bath experiments do for instance. However, the various boundary conditions, numerical schemes and turbulence models that form the basis of a CFD simulation introduce uncertainties that need to be accounted for, as discussed in this section. The chosen simulation set-up needs to be able to simulate buoyancy driven ventilation flow and the transport of passive scalars representing infected breath while modelling the reference scenario described in §3.1.1. Only the cross-flow configuration is described here, with single-sided simulations run in the same way, with a different top vent position as described in §3.1.

The code employed in this work is OpenFOAMv2106 (OpenCFD, 2020). OpenFOAM is a free open source CFD code employed both in industry and research for a variety of applications, such as complex fluid flows involving chemical reactions, acoustics, solid mechanics and electromagnetics problems. It was selected here because it allows both buoyancy driven flows and passive scalar transport to be simulated. As such, the code is highly versatile, offering a range of options to model turbulence and is available on most supercomputer systems including ARCHER2, which was used for this work. The code is regularly updated with a release every 6 months, OpenFOAMv2106 (developed by OpenCFD) was used here because it included an option for including the buoyancy effects on turbulence production. A transient solver, `buoyantPimpleFoam`, was chosen as it is appropriate for buoyant and turbulent flows to model ventilation and heat transfer. The choice of a transient solver is discussed further in §3.3. OpenFOAM terminology is described in Appendix A.

3.2.1 Computational domain and mesh

The computational domain included the classroom along with two exterior boxes above and below the classroom linked to the classroom through inlets and outlets as shown in Figure 3.2. The need to model exterior boxes is discussed further in §3.4. The classroom was defined as a cube of dimensions $10\,\mathrm{m}\times\ 5.5\,\mathrm{m}\ \times 2.7\,\mathrm{m}$. The bottom exterior box was centred around the inlet to the classroom, with dimensions $2.8\,\mathrm{m}\times\ 2.5\,\mathrm{m}\times\ 1\,\mathrm{m}$. The top exterior covered the same surface as the classroom with dimensions $10\,\mathrm{m}\times\ 5.5\,\mathrm{m}\times\ 3.7\,\mathrm{m}$. The classroom was linked to the exterior boxes via an inlet and an outlet of the same size as the vents with respective areas A_B and A_T (as defined in §3.1) and height 0.2 m, which was the minimum necessary for meshing. The exact location of the openings along the diagonal of the room was investigated and was not found to impact the flow significantly. The computational grid was created directly with the OpenFOAM utility `blockMesh`. The mesh was defined as a perfect orthogonal mesh with $\Delta x = \Delta y = \Delta z = 0.05\,\mathrm{m}$ in the classroom, inlet and outlet and $\Delta x = \Delta y = \Delta z = 0.1\,\mathrm{m}$ in the exterior, leading to 1,399,356 hexahedral cells. The choice of computational grid is further discussed in §3.5

3.2.2 Turbulence modelling

Buoyancy driven flows found in naturally ventilated spaces are typically turbulent near vents and heat sources, but often exhibit low Reynolds number elsewhere. Hence, when simulating such flow

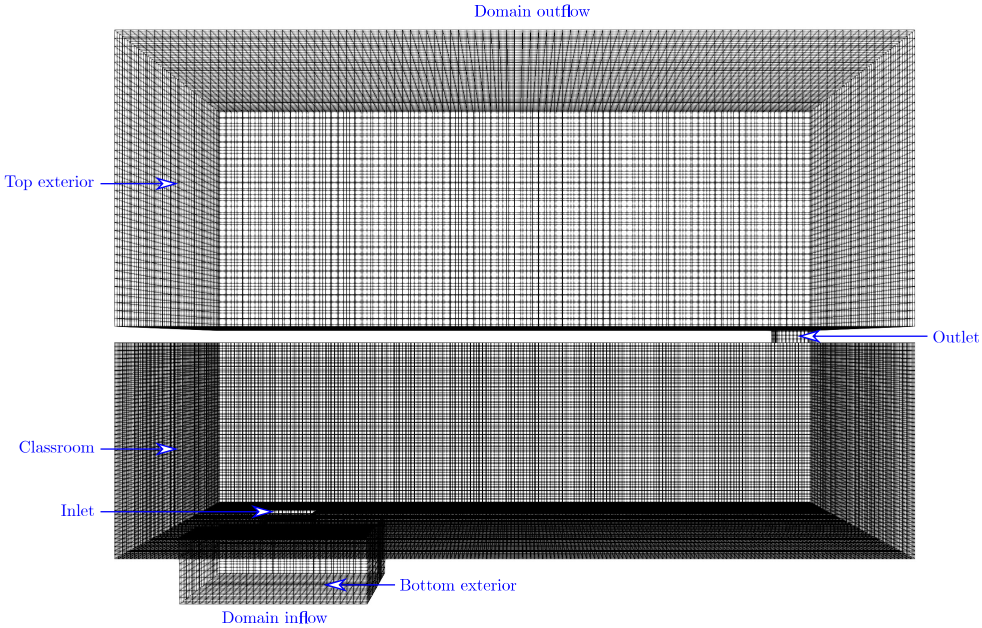

Figure 3.2: Computational grid used to simulate the classroom, including the modelled exterior, inlet and outlet.

the modelling must appropriately represent: turbulent flows driven by convection, low-Reynolds number flows, and both the transition to turbulence and relaminarisation. This presents a challenge which requires a balance of needs; highlighting why the choice of an appropriate turbulence model is of critical importance if the CFD simulations are to capture the main features of the flow in the room. Both Large Eddy Simulation (LES) and Reynolds-averaged Navier-Stokes (RANS) have been used for the simulations of flows in and around buildings as reviewed by Blocken (2018). LES can be more accurate than RANS simulations, as the larger eddies, often specific to the problem studied, are directly resolved and not modelled. This however also makes LES more computationally expensive and more complex to set up. Choice was made to exploit the speed of RANS simulations to explore the parameter space rather than investigate a single case in depth. With RANS simulations, turbulence is not directly simulated but is modelled which increases computational speed but also introduces inaccuracies.

Simulations were run using a RANS turbulence model: $k-\omega$ SST by selecting `kOmegaSST`. This model strikes a balance between the performance of $k-\epsilon$ models for shear flows and $k-\omega$ models at the walls (Pope, 2000), both features relevant to the capturing ventilating flows, by blending the two models. Large variations in the turbulence levels in the classroom can also be expected (Blocken, 2018), leading to potential relaminarisation of the flow which the $k-\omega$ SST model is expected to capture. The chosen OpenFOAM version also incorporated the effects of buoyancy on turbulence production by using the `buoyancyTurbSource` finite volume option (OpenCFD, 2021).

In order to limit the computational costs and the size of the computational grid, the boundary layer region layer was not fully resolved. Instead wall functions were used with standard parameters for: the turbulent thermal diffusivity using `alphatWallFunction`, the turbulent kinetic energy with `kqrWallFunction`, turbulent viscosity (`nutkWallFunction`) and specific dissipation rate with `omegaWallFunction`.

3.2.3 Fluid properties

The fluid properties were defined in the `thermophysicalProperties` input file which defines the constitutive equations between pressure and temperature from which other variables are calculated. Following typical examples of ventilation simulations, a density based model for a fixed composition was used and air in the room was set to a perfect mixture, with molecular weight set to 28.9 g/mol. The dynamic viscosity μ and Prandlt number were set to $1.74 \times 10^{-}5\,\mathrm{kg/m/s}$ and 0.7 respectively. The specific heat capacity was set to $c_p = 1005$ J/(gK). The Boussinesq equation of state was employed, where the density is found from $\rho = \rho_0\,(1 - \beta(T - T_a))$. This was deemed suitable in this context as relatively small density differences are expected within the classroom, and this approximation is often in experimental settings, e.g in Linden et al. (1990). The Boussinesq assumption was also employed when selecting the vent sizes and heat input in §3.1.1. The reference density ρ_a was set to $1.268\,\mathrm{kg/m^3}$ at ambient temperature $T_a = 278.15\,\mathrm{K} =$ 5°C with an expansion coefficient α equal to $0.00362\,\mathrm{K^{-1}}$. In addition, the reference pressure was set to 0 and the gravitational acceleration vector g set to $(0,0,-9.81)\,\mathrm{m/s^2}$.

3.2.4 Boundary conditions

A number of combination of boundary conditions were tested in the scenario presented in §3.1.1 with distributed heating, but also in cases without heating and ‘emptying’ cases, which were run to check several physical properties of the simulations were conserved, e.g. that there was no flow in the domain or that the heat was flushed out of the domain as expected. The chosen boundary conditions are described below.

At the domain inflow (the bottom face of the bottom exterior box), the temperature was set to the ambient value using a Dirichlet boundary condition `fixedValue`. On the floor of the classroom, `externalWallHeatFluxTemperature` was used to set the distributed heat source to 6200 W, with the accuracy of the input checked in post-processing. At the domain outflow (the top face of the top exterior box), the `inletOutlet` boundary condition was used, which uses a Neumann zero normal gradient (`zeroGradient`) boundary condition for flow out of the domain and sets any inflow with a Dirichlet boundary condition, here equal to the ambient temperature. The remaining boundary conditions were set to `zeroGradient` on the walls of the domain.

The velocity was set to `noSlip` at the classroom, inlet, outlet and exterior walls. `pressureInletOutletVelocity` was used at the domain inflow, which sets the inflow velocity based on the flux normal to a face boundary and any outflow is assigned a `zeroGradient` condition. `inletOutlet` was used at the outflow of the domain setting any potential inflow to (0, 0, 0) m/s.

The turbulent thermal diffusivity, kinetic energy, viscosity and dissipation used default wall functions at the walls as described above. The inflow and outflow to the domain were set to `zeroGradient`.

OpenFOAM buoyant solvers rely on the use of a modified pressure `prgh` to set the boundary conditions for pressure. This is an issue as it is calculated by subtracting the hydrostatic pressure contribution from the pressure field such that $prgh = p - \rho gh$ where ρ is the local density and h is an arbitrary height. This is problematic in the case of buoyancy driven flows where density varies in the domain and is not fixed at the domain inflow and outflow. This issue was avoided by using the `prghTotalPressure` condition which allows for the total pressure p_0 to be defined directly without using the modified pressure $prgh$. Defining the total pressure p_0 and not just the static pressure p was necessary to ensure a uniform flow across the opening. The total pressure across the total height of the domain H_D was found with $\Delta p_0 = -\rho_a g H_D$, and p_0 was defined as 14.93 at the domain inflow and -82.10 at the outflow, $z = 0$ taken to be on the classroom floor. The remaining wall boundary conditions were set to `fixedFluxPressure`, which is typically used.

The scalar representing the breath of occupants through the addition of CO_2 was introduced using a `scalarSemiImplicitSource` with a rate corresponding to $NG = 32 \times 3.35 \times 10^{-6}\,\text{m}^3/\text{s}$. It was added over a volume defined by the `topoSetDict` utility over the classroom surface between heights of 1.1 and 1.2 m. A `fixedValue` boundary condition set to 0 was used at the domain inflow, `inletOutlet` was used at the domain outflow and `zeroGradient` at the walls. Results were analysed by converting the concentration to ppm and adding a background CO_2 of 400 ppm so that it could directly be compared to classroom measurements.

3.2.5 Numerical discretisation

Numerical solver settings were set in the `fvSolution` and `fvSchemes` dictionaries. The `buoyantPimpleFoam` solver was used with a momentum predictor step, 3 outer corrector steps, 1 inner corrector step and 1 non orthogonal corrector step, following advice. The solver options used the `PCG` solver for the pressure field and `smoothSolver` for the remaining fields. Different solvers were also investigated but did not impact the final solution whilst also slowing down the simulations significantly. The absolute tolerance for the pressure field was set to 10^{-6}. The scalar fields used an absolute tolerance of 10^{-12} and remaining fields used a tolerance of 10^{-10}. All fields were solved using a relative tolerance of 0.01 for the inner loops and 0 otherwise. Different values were tested and the settings described above were found to provide both an accurate simulation and ensured the speed of the simulation. To increase efficiency all under relaxation factors were set to 1.

`backward` was used as the time discretisation scheme as it is both second order accurate and ensures stability. The gradient schemes were set to `cellLimited Gauss linear 1` to ensure the use of limiters and to avoid overshoots. For the pressure field, however, the scheme was changed to `Gauss linear` so that it satisfied the divergence criteria. The divergence schemes were set to `Gauss linear Upwind` which is second order and was found to improve the stability of the simulations. The passive scalar divergence scheme was set to `upwind`, which is recommended for the scalar transport utility to remain bounded. The Laplacian schemes were set to `Gauss linear orthogonal` since the mesh used was orthogonal, a `linear` interpolation scheme was selected for accuracy and the surface-normal gradient schemes schemes were also set to `orthogonal`. Finally the wall distance calculation necessary for the $k-\omega$ SST model used the default mesh wave method.

3.2.6 Running a simulation

The simulations were first initialised using a steady solver (`buoyantSimpleFoam`) for 50,000 iterations to establish stratification. This enabled the transient solver not to start from ambient conditions and reduced instabilities. The simulation was then run for 5,000 s with the transient solver (`buoyantPimpleFoam`) before averaging for 3,600 s. This process is further explained in §3.3. An adaptive time-step was used to ensure both stability and efficiency by setting a maximum CFL of 1. Overall, the simulation took 15.4 h to run on the ARCHER2 UK National Supercomputing Service (`https://www.archer2.ac.uk`) using an overall of 1,024 CPUs.

3.3 Using a transient solver

CFD simulations such as the ones described above could be expected to be run using a steady solver. Indeed, it can be assumed that since the turbulence is modelled and the input conditions are not varying, the system will reach a quasi-steady state such as the one observed by Gladstone and Woods (2001) in their distributed source experiments. This is advantageous numerically as once the simulation has reached a steady state no further averaging needs to be performed therefore reducing computational cost.

However, in the particular set-up studied herein, steady-state convergence proved challenging. Indeed, initial steady simulations were unable to converge due to the presence of large scale internal motions in the room, although the set-up was statistically steady. Several numerical parameters were investigated but convergence could not be achieved. The results shown in Figure 3.3 were obtained following the set-up described above, with a steady solver `buoyant-SimpleFoam` and upwind discretisation schemes, we note that 2^{nd} order schemes were also investigated and showed similar issues. Figure 3.3 shows that after an initial decrease, the initial residuals stabilise about a mean value and stop decreasing, never reaching convergence. This is accompanied by large fluctuations at points of the domain, as shown in Figures 3.3c) and d) for the vertical velocity and the CO_2 concentration for four probes in the domain. The two plots show an initial period where the ventilating flow and thermal stratification is established in the room followed by a period in which vertical velocity values keep oscillating about a mean value. This directly impacts the CO_2 distribution, as demonstrated by Figure 3.3d), in which an extreme variation of CO_2 concentrations can be observed at different iterations. Specifically, near the breathing zone at $z = 1.08\,\text{m}$, CO_2 values fluctuate by about 1,000 ppm across the simulation. This is of particular concern as understanding the resulting spatial distribution of CO_2 is one of the primary objectives of the CFD simulations used here. The use of a steady solver was therefore deemed inappropriate.

The remaining simulations were then run using unsteady RANS simulations by switching to the transient solver `buoyantPimpleFoam`. Each time-step was then resolved numerically and the solution was time-averaged, averaging out the effects of the large scale fluctuations. Unsteady RANS simulations have been used for several applications including indoor flows (Rossi and Iaccarino, 2013) and flows around buildings (Tominaga, 2015), for which they provided more accurate results than steady RANS. In the context of urban flows, Tominaga and Stathopoulos (2017) also showed that unsteady RANS simulations performed better than steady RANS to

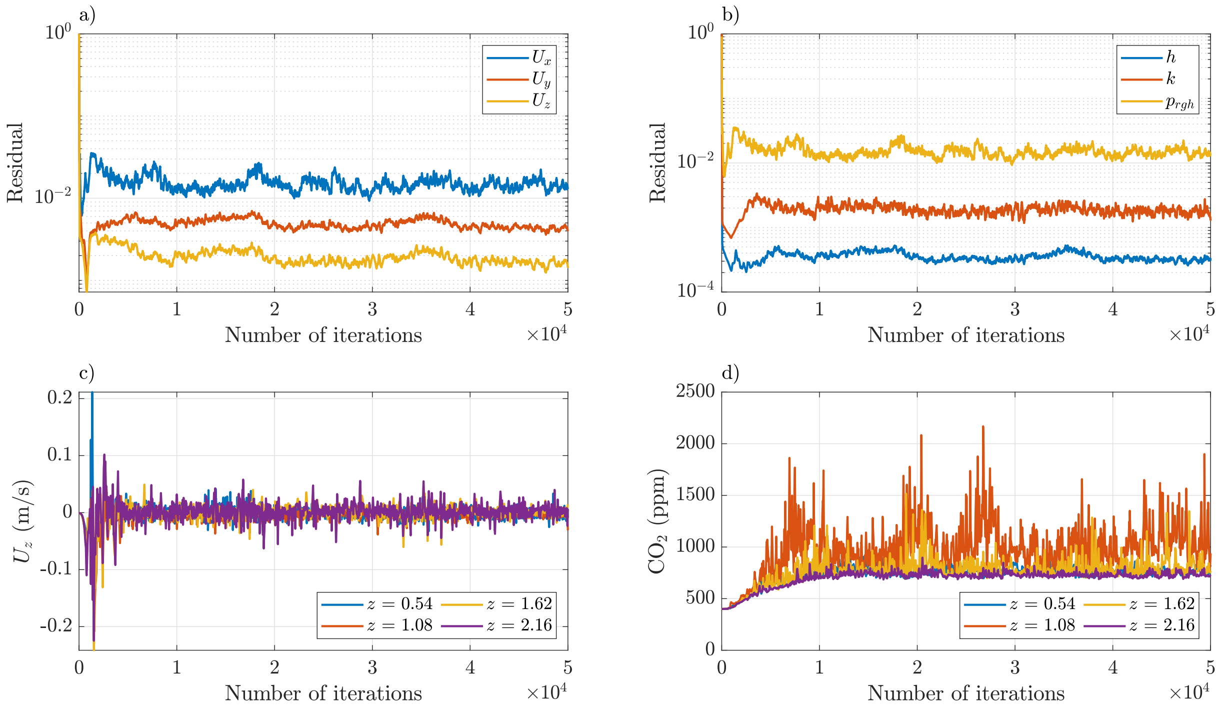

Figure 3.3: a) and b) show the initial residual values (for velocity and other fields respectively) of the solver with number of iterations. c) and d) show the probes values for velocity and CO_2 at 4 sample points in the room.

capture large scale fluctuations necessary to model pollution dispersion, which is something this work also aimed to achieved.

Simulations were then run with a transient solver after a precursor steady run to establish the flow and stratification in the room. The convergence of the room averaged temperature and CO_2 in this initial steady solver step is shown in Figure 3.4a) by plotting the room average value with number of iterations. This shows that the stratification has been established although the fluctuations discussed above are also visible in the room averaged CO_2 concentration. Stratification was deemed to be established enough to use 50,000 iterations as a precursor simulation from which to initialise transient simulations. Simulations were then run with a transient solver and averaged after 5,000 s. The resulting room averaged temperature and CO_2 are shown in Figure 3.4b). 5,000 s was deemed a suitable time after which the simulation had reached a statistically steady state and the time averaging could be performed. The final time-averaged values are shown with the horizontal dashed lines.

Figures 3.5a) and b) also show the instantaneous values of vertical velocity and CO_2 at four probes in the classroom for the first 5,000 s of the transient simulation. This can be directly compared to the results shown in Figures 3.3c) and d) obtained with the steady solver. The fluctuations observed for both the vertical velocity and CO_2 concentration are lower than those

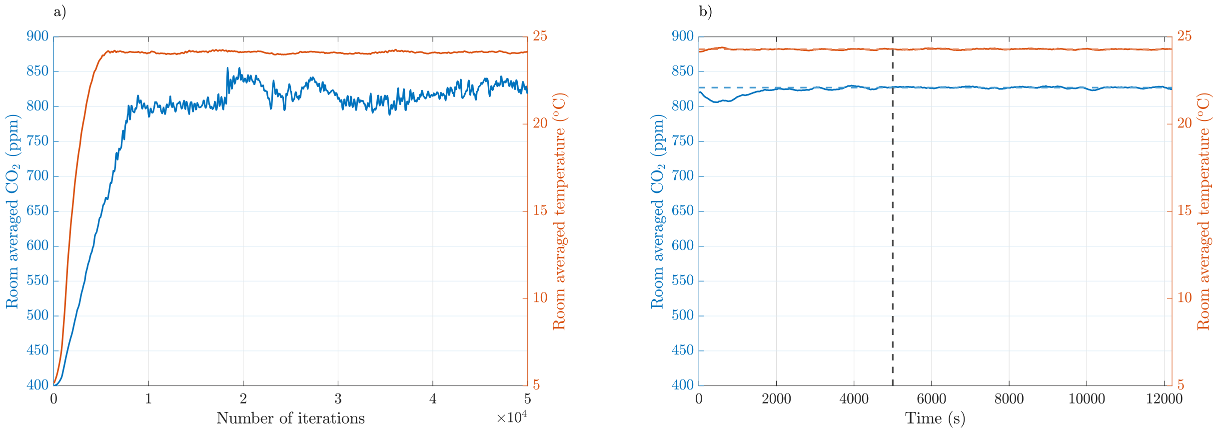

Figure 3.4: Variation of the room averaged CO_2 (blue) and temperature (red): a) with number of iterations for the precursor step using a steady solver and b) with time for the transient solver. Dashed horizontal lines show the final time averaged value and the vertical dashed line shows the start of the averaging period.

obtained previously. However, near the breathing zone (at $z = 1.05\,\mathrm{m}$ for instance), CO_2 values still vary within a range of 400 ppm (compared to 1,000 ppm previously) which shows the presence of large scale fluctuations and points out the need for time averaging.

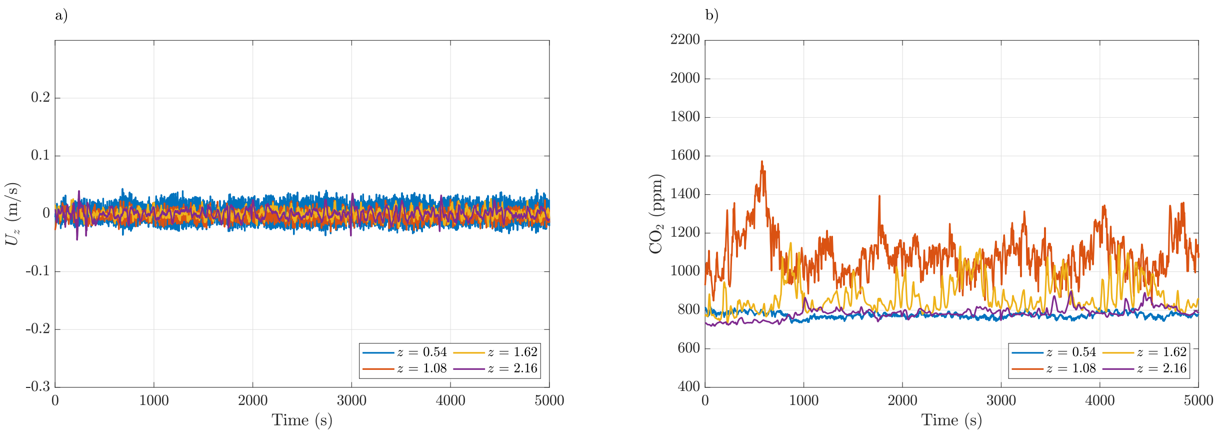

Figure 3.5: a) Vertical velocity and b) CO_2 concentration instantaneous values at 4 sample points in the room obtained with a transient solver over 5,000 s.

The required time to obtain a representative average was also investigated. Figure 3.6 shows the a) temperature and b) CO_2 concentration averaged at each height in the classroom, averaged over different durations after an initial 5,000 s transition period. Small discrepancies are visible between the 30 and 60 minutes profiles. These disappear between the 60 and 120 minutes time-averaged profiles. To limit computational cost and necessary simulation time it was decided to average simulations over 60 minutes.

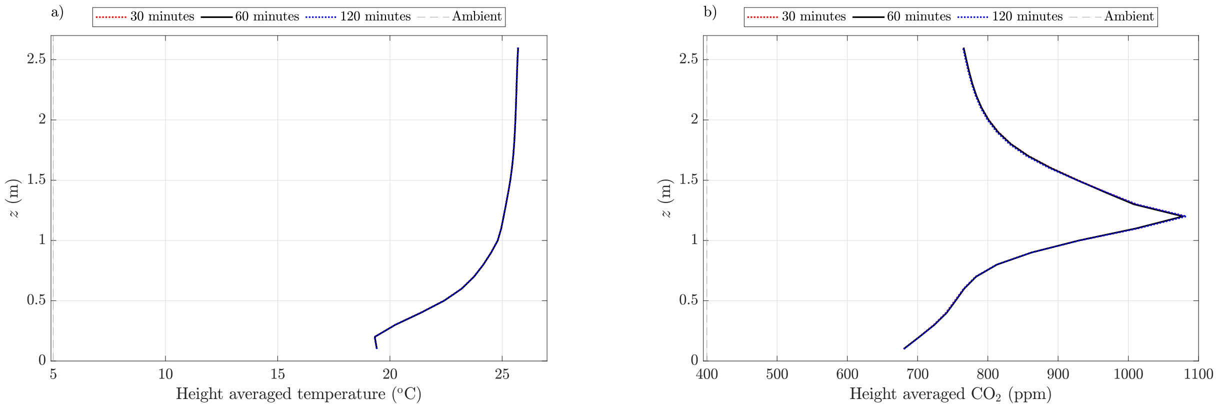

Figure 3.6: Horizontally averaged: a) temperature and b) CO_2 concentration in the classroom. After a 5,000 s transition period, simulations were averaged for 30, 60 or 120 minutes. Ambient temperature and CO_2 concentration are shown with dashed vertical lines.

3.4 Defining the computational domain

This section aims to describe how the computational domain was defined. In particular, the use of additional boxes to represent the exterior, as shown in Figure 3.2, was deemed necessary and this is detailed here.

3.4.1 Simulating the classroom only

In order to limit computational costs, a simulation modelling only the classroom was first investigated. This meant that the mesh covered only the area of direct interest therefore limiting the number of cells. The inflow and outflow of the domain were located directly at the bottom and top vents of the classroom. This was particularly challenging because the typical boundary conditions used in CFD, where the flow rate is directly imposed at the inflow, could not be used. Instead, the set-up had to be buoyancy-driven, and pressure boundary conditions were used at the inflow and outflow of the domain, directly on the classroom boundaries. The `pressureInlet-OutletVelocity` and `prghTotalPressure` coupling described in §3.2 was thought to be suitable to avoid issues created by the use of the modified pressure *prgh* introduced by OpenFOAM.

However, additional issues were encountered and could not be resolved despite the analysis of other numerical settings, such as different choices of boundary conditions or solver parameters. Following the process described in §3.3, the simulations were first initialised with a steady solver before running the transient solver and time-averaging. This revealed issues in which the final state of the transient simulation was dependent on the initial state provided by the precursor steady simulation. The exact same simulation (using the same computational grid), but initialised from a different steady simulation, for instance one that might use a coarse or medium mesh, led to very different time averaged final values, with a difference of up to 100 ppm or 0.5°C.

This was thought to be caused by a sensitivity to the values imposed at the openings by the chosen pressure boundary conditions, which couldn't be controlled. It was therefore decided to move the domain inflow and outflow boundaries further away from the classroom itself by modelling the exterior beyond the classroom's bottom and top vents. Although this added to the computational cost by requiring the simulation of a larger domain, it also allowed the proper resolution of the classroom inlet and outlet, thereby ensuring that the flow entering the classroom was realistic.

3.4.2 Simulating the exterior

Exterior boxes were then added below and above the classroom cube, connected via an inlet and outlet as shown in Figure 3.2. The inlet and outlet domains were kept to a minimum, with a height of 0.2 m to ensure the possibility of meshing even in the coarse case described in §3.5. At first, the top exterior box was designed to be smaller, covering only 6 m^2 centred around the 0.2 m^2 top opening. However this also led to convergence issues, as shown in Figure 3.7a) where the final temperature profile in the room was dependent on the height of the top exterior box, varying here between $3/4H_c$ and $2H_c$ (with H_c the height of the classroom) and leading to differences in temperature of up to 3°C.

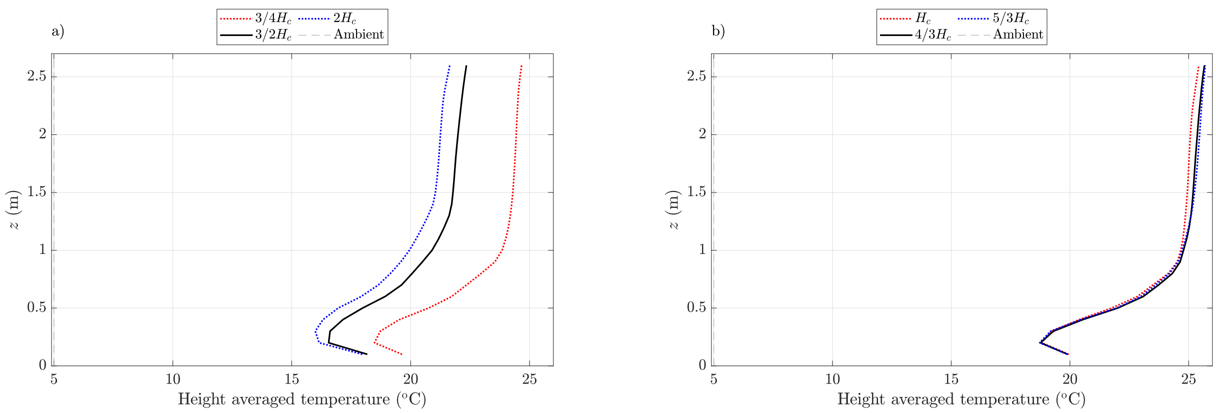

Figure 3.7: Horizontally averaged temperature in the classroom for different top exterior heights. a) shows simulations where the exterior domain covered a surface of 6 m^2 and b) results where the exterior has been extended to cover the entire classroom surface (55 m^2). Ambient temperature is shown with a vertical dashed line.

This issue was resolved by ensuring that the top exterior box covered the same horizontal area as the classroom, as shown in Figure 3.2 and that the height was at least 4/3 of the height of the classroom, as demonstrated by Figure 3.7b) for three exterior domains of different heights.

The scale of the bottom exterior box was also investigated but shown not to impact the results. In order to limit the mesh size, the bottom exterior domain was set to cover 1 m around the bottom opening vent leading to a bottom exterior box of dimensions 2.8 m $\times$2.5 m $\times$1 m.

For completeness, the effect of initialisation was also checked with this configuration. The steady initialisation was run either on a coarse or medium mesh, and the simulations were then continued with a transient solver on a medium mesh as described in §3.2. The results are shown in Figure 3.8a) for the room averaged temperature and Figure 3.8b) for the room averaged CO_2 concentration. Simulation A represents the case described in §3.2 where the initialisation is run on a coarse mesh, B is the case where the steady solver is run on a medium mesh. Despite a different initial state both simulations do tend to the same final value, unlike in the classroom only simulation described in §3.4.1. It was therefore decided to use a coarse steady simulation to initialise all further simulations as it was faster and led to comparable results to a medium mesh initialisation.

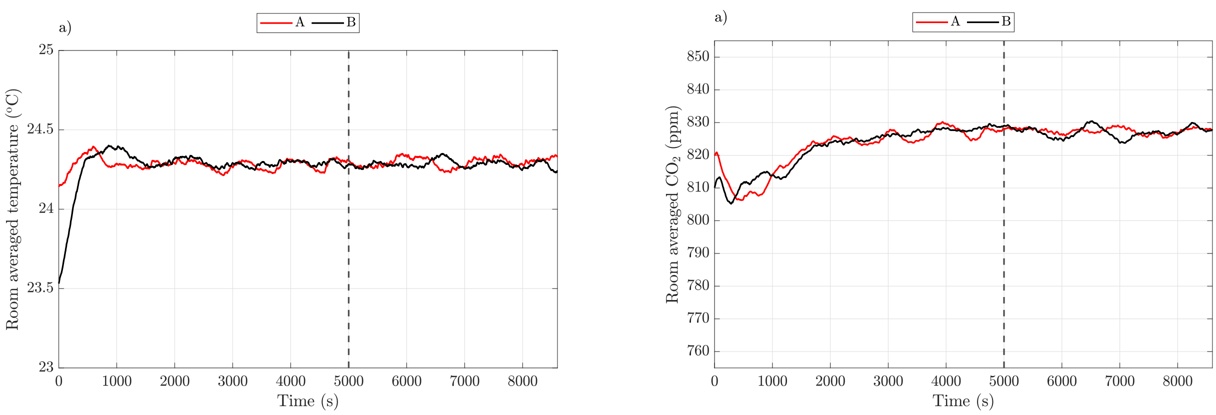

Figure 3.8: Room averaged: a) temperature and b) CO_2 concentration for a transient simulation. Simulation A was initialised from a precursor steady simulation run on a coarse mesh and simulation B from a run using a medium mesh.

3.5 Designing the computational grid

This section describes how the mesh used to run the simulations was chosen. When setting the computational grid and choosing the mesh resolution, it is essential to ensure the results are independent of the choices made. Three different meshes were investigated in which the exterior resolution was kept the same, to keep the computational cost down, but the mesh at the inlet, at the outlet and inside the classroom (the focus of this work) was refined to show mesh independence. All meshes were kept uniform and orthogonal to reduce the complexity and run time. All simulations were run using the method described in §3.2 (the simulation shown corresponds to the medium mesh used here).

The computational cost of running each simulation is summarised in Table 3.3. As expected, this shows the increasing computational demand with mesh refinement, a fine mesh simulation ran in 3.5 days on ARCHER2 with a significant computational cost. Such a long run time and cost would make the use of a fine mesh to investigate the parameter space challenging. In addition to the run time, the size of the data output also needs to be taken into account, any investigation of several time-steps would require significant data storage capabilities which would make the subsequent post-processing more complex. On the other hand, the medium simulation, remains at a lower cost with a run time of 15 hours, which is already significant. The coarse mesh, as can be expected, is the most advantageous in terms of cost and speed with the lowest run time, computational requirement and data output size.

Simulation	Number of cells	Classroom resolution (m)	Run time (h)	Nodes (CPUs)	Data output final time step (GB)
Coarse	359,120	0.1	3.4	2 (256)	0.4
Medium	1,399,356	0.05	15.4	8 (1024)	1.4
Fine	9,722,076	0.025	84.3	16 (2048)	9.1

Table 3.3: Summary of the meshes investigated along with the resulting computational cost of the simulations. The mesh described in §3.2 corresponds to the medium mesh.

The resulting room averaged values for temperature and CO_2 along with the measured flow rate both in terms of Air Changes per Hour (ACH) and L/s/person are shown in Table 3.4. The ACH results of the coarse and medium mesh are within 2% of the results calculated from using a fine mesh, with the coarse mesh, perhaps surprisingly, leading to a value closer to the fine results. This is not the case for the room averaged excess temperature where the coarse mesh underestimates the fine results by 5% and the medium by 3%. Agreement with the fine simulation is even closer for the averaged excess CO_2 concentration, where the coarse simulation underestimated the result by 5% and medium simulation by less than 1%. Overall the simulations were within the range of the values predicted by assuming a well-mixed environment for c_d values between 0.6 and 0.7 (Gladstone and Woods, 2001) as described in §3.1.1.

The independence of the mesh can also be investigated by looking at the temperature and CO_2 profiles across the classroom which show the implementation of stratification, plotted in Figure 3.9. Differences between profiles is visible for the temperature with an increase in resolution leading to an increase in the temperature across the height of the room. The shape of

Simulation		ACH	Q_p (L/s/person)	T (°C)	C (ppm)
Coarse		5.74	7.40	23.9	811
Medium		5.79	7.47	24.3	827
Fine		5.66	7.30	24.9	831
Well	$c_d = 0.6$	5.35	6.90	27.0	886
mixed	$c_d = 0.7$	5.93	7.64	24.9	838
prediction	$c_d = 1$	7.52	9.69	20.7	746

Table 3.4: Bulk flow measurements and room averaged results for the three resolutions considered. Well-mixed predictions are given for $c_d = 0.6$, 0.7 and 1 (calculations are detailed in §3.1.1).

the profile is, however, similar across simulations and the medium mesh does achieve closer results to the fine mesh near the floor and the ceiling, with differences of less than 1°C at every height between the medium and fine simulations. The profiles are a lot closer for the CO_2 concentration shown in Figure 3.9b). Once again, the agreement in the middle of the room is not as good as the agreement at the floor and ceiling for the fine and medium mesh, with a maximum difference of about 100 ppm observed at any given height.

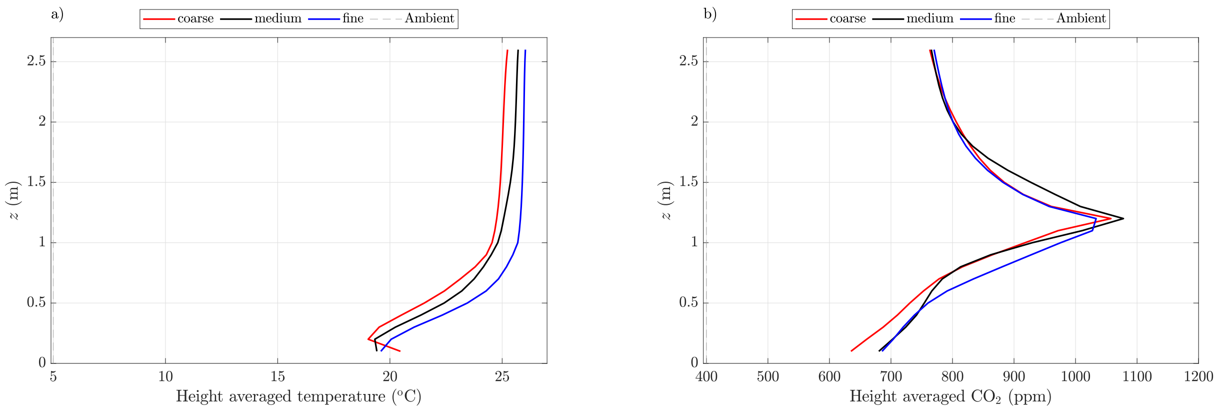

Figure 3.9: Horizontally averaged: a) temperature and b) CO_2 concentration in the classroom for simulations with different resolutions. Ambient temperature and CO_2 are shown with dashed vertical lines.

In addition to these profile plots, contour plots at the breathing zone, taken to be between 1 and 1.5 m to represent the seated and standing head heights of the pupils, were also used to investigate the results across the different resolutions. Figure 3.10 shows the vertical velocity averaged over those heights. It is critical to ensure that the velocity field is well resolved as it will directly impact the distribution of CO_2 and infected breath. In this configuration, the flow entering the room flows along the floor, is slowly heated by the heat input and then rises along the opposite wall from the inlet, as shown in Figure 3.10. The details of the ventilating flow are described in more detail in Chapter 4 and compared to the flow field obtained with a single-sided configuration. All three contour plots show this feature, with a considerable increase of vertical velocity on the right side of the plot, furthest away from where the bottom vent is located at (1.4, 1.25, 0). Here, the increase in resolution between a) the coarse and b) medium simulations does lead to better qualitative results, the region of high velocity being closer to what is observed for c) the fine mesh.

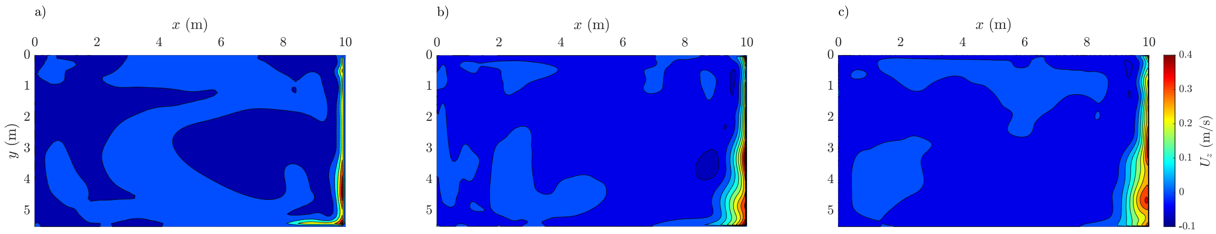

Figure 3.10: Horizontal cross section of time-averaged vertical velocity in the breathing zone ($1 \leq z \leq 1.5$ m) for: a) the coarse, b) medium and c) fine simulations.

In the case of the temperature plots shown in Figure 3.11, the general flow structure is the same. Once again the effect of the high velocity region can be seen with a region of higher temperature in the right side of all three plots which corresponds to the warmer air current flowing upwards. The rest of the domain shows a more uniform temperature with an increase in simulation resolution leading to an overall increase in temperature as observed in Figure 3.9a). b) the medium and c) fine simulations, however, do present a closer matched temperature structure in this region of high temperature.

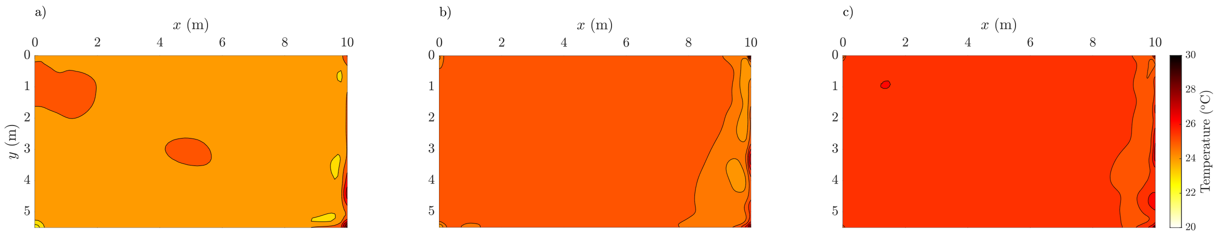

Figure 3.11: Horizontal cross section of time-averaged temperature in the breathing zone ($1 \leq z \leq 1.5$ m) for: a) the coarse, b) medium and c) fine simulations.

The CO_2 plots shown in Figure 3.12 show a slightly different structure than the velocity and temperature plots since the source of CO_2 is within the breathing zone plotted. The effect of the high velocity region is still visible however on the right with a low CO_2 concentration region visible for all three resolutions. The CO_2 distribution does differ slightly as the mesh resolution varies. However, the medium and fine mesh do show comparable regions of higher CO_2 levels (above 1,200 ppm) in the top side of the plots and away from the region of high velocity on the right.

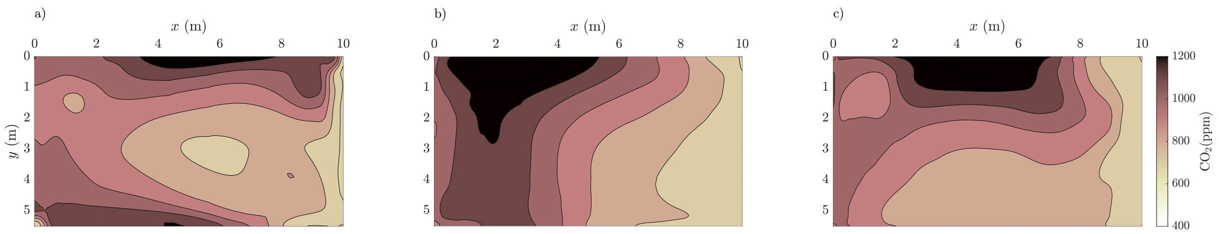

Figure 3.12: Horizontal cross section of time-averaged CO_2 in the breathing zone ($1 \leq z \leq 1.5\,\text{m}$) for: a) the coarse, b) medium and c) fine simulations.

The main focus of this thesis being on assessing the validity of a well-mixed assumption, the distribution of CO_2 was studied using histograms shown in Figure 3.13. This shows all three simulations considered together and, overall the agreement is acceptable for the range of values expected in the classroom. The medium mesh does, however, show a better prediction in regions around 800ppm with frequencies matching those observed with the fine simulation.

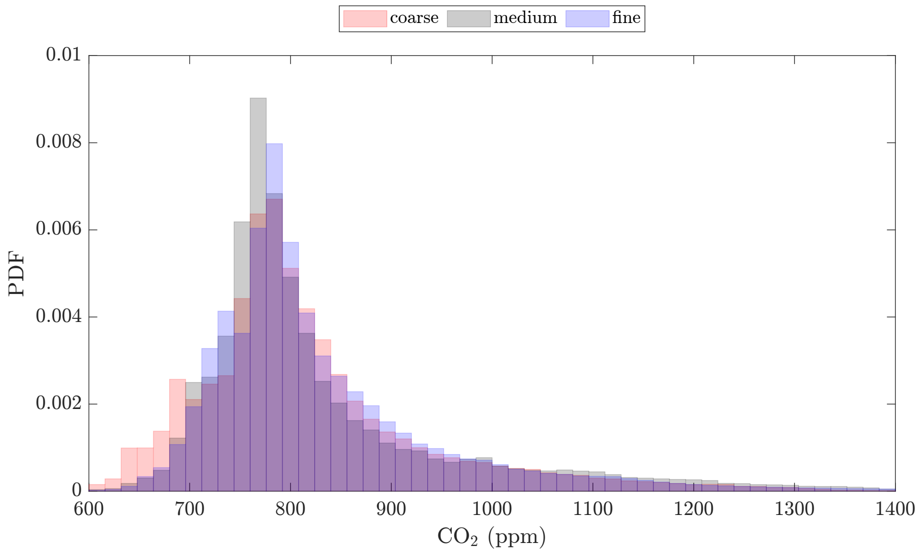

Figure 3.13: Histogram of the CO_2 concentration in the classroom using 50 bins for the coarse, medium and fine simulations.

The grid convergence study showed that changes in resolution between a coarse, medium and fine mesh led to comparable bulk flow measurements, although local differences are observed, particularly in the case of the coarse mesh, but also for a medium mesh. These were, however, deemed acceptable as the medium mesh is shown to capture the main flow features in the breathing zone as well as the overall resulting CO_2 distribution, which is the main focus of this work. The use of a medium computational grid also allows the investigation of a number of scenarios which would not possible with a fine mesh due to the significant run time and computational cost it incurs (both to run and analyse the simulation).

3.6 Conclusions

A generic naturally ventilated UK classroom was selected to run CFD simulations in which the distribution of CO_2 across the space can be investigated. A number of assumptions were made to ensure the set-up used was representative of the widest range of classrooms possible. It was decided to model a scenario in which the mixing in the space would be limited in order to better estimate the upper limit of errors introduced by the use of a well-mixed assumption. A wintertime scenario was considered and the buoyancy driven ventilation was chosen to be supplied through low and top-level vents in two configurations where the relative horizontal position of the vents was modified: cross-flow and single-sided. The occupants were solely modelled through the heat they release, which added to the heat supplied by the radiators, was imposed on the floor. The occupants' breath was modelled as a source of CO_2 across the breathing height (set here to be between 1.1 and 1.2 m). The opening size and remaining heat input due to radiators were defined to match government guidelines using predictions of Gladstone and Woods (2001) for a distributed heating source.

The CFD simulations were run using OpenFOAM and the solver `buoyantPimpleFoam`. The different set-up parameters were chosen in line with recommended best practice for the case studied here and following numerical investigations to obtain the optimal solver settings. The set-up contained 1.4 million cells and took approximately 15 hours to on ARCHER2. A transient solver had to be used due to the presence of large scale fluctuations which did not allow solutions generated by a steady solver to converge. The simulations were therefore initialised with a steady solver, so that stratification would be first established, and then run for 5,000 s with a transient solver before time averaging was performed for 3,600 s. These values were chosen to ensure a suitable convergence and accuracy of the results while limiting computational costs.

The exterior had to be simulated in order to properly model the effect of flow at the two openings and the resulting flow in the classroom. A grid convergence study was also performed to ensure the independence of the results from the resolution chosen. Although local variations were still observed, a medium mesh was shown to accurately represent both the bulk flow and CO_2 distribution when compared to a fine mesh.

Chapter 4

CO_2 distribution within a buoyancy driven naturally ventilated classroom

In Chapter 3, the computational set-up used in the remainder of this thesis was described and the assumptions made, taken to characterise a generic naturally ventilated classroom, were described. Computational Fluid Dynamics (CFD) simulations were set up to model the transport of two passive scalars; one to mimic the general CO_2 concentration, and the other to represent potential sources of infected breath in the classroom. The set-up was designed to enable the distribution of these scalars to be compared and hence determine the extent to which CO_2 measurements are a suitable proxy to estimate far-field exposure to infected breath in classrooms — an already well established methodology (e.g. Rudnick and Milton, 2003) applied to UK classrooms in Chapter 2. In this chapter, the methodology presented in Chapter 3 is deployed to first analyse and describe the expected variations of CO_2 in a generic naturally ventilated classroom. Particular focus is drawn to whether a number of point measurements, viably obtainable from in-classroom sensors, can accurately represent the distribution of CO_2 and hence be relied upon to inform estimates of the bulk ventilation rate. Comparison will also be drawn to the scalar taken to mimic exposure to infected breath and released from a variety of local regions within the classroom in Chapter 5, in which the suitability of CO_2 as a proxy for far-field exposure will be investigated.

As such, the focus of the current chapter is the distribution of CO_2, which is inherently generated by all occupants within the classroom. The current chapter aims to assess whether for the reference scenario described in Chapter 3, in which the ventilation is buoyancy driven through low and high-level vents, the CO_2 concentration can be assumed to be well-mixed, and if not then: what is the degree of uncertainty (relative to other uncertainties common to infection modelling) that is introduced by assuming the air in the classroom is well-mixed? To inform the potential variability within a classroom, two limiting ventilation configurations were investigated. Between the two configurations the horizontal location of the top opening on the ceiling was changed: in one configuration, a flow from one end of the classroom to the other was promoted, herein the 'cross-flow' configuration; and in the other, both vents were positioned at one end of the classroom, herein the 'single-sided' configuration. The impact of the ventilating flow rate on the results was also assessed, both by varying the size of the openings to restrict or enhance the flow, and by varying the heat input in order to change the driving force underlying the ventilation of the classroom. Part of the aim of this chapter is to investigate the potential for assessing ventilation, and potentially far-field infection risk, by using point measurements from CO_2 sensors in operational classrooms. However, in practice, many CO_2 sensors are deployed by attaching them to walls (a viable choice from many perspectives) but the extent to which these measurements might be representative of CO_2 levels within the bulk of the space, and particularly in the breathing zone, is not known. As such, this chapter also investigates the degree of any additional uncertainties introduced when point measurements CO_2 are taken near the walls; in addition the sensitivity of the CO_2 measurements to height is investigated for both flow configurations considered.

This chapter is divided as follows: §4.1 compares the cross-flow and single-sided configurations for the reference scenario presented in Chapter 3. The resulting ventilating flow is described, the CO_2 concentration in the room and in the breathing zone are compared and the impact of using wall measurements is also discussed. In §4.2, the reference set-up described in Chapter 3 is modified to investigate the sensitivity of the previous results to changes in the ventilating flow rate, by varying the low and high-level vent areas or the heat input. Finally, conclusions are drawn in §4.3.

4.1 Comparing the cross-flow and single-sided configurations

Two ventilation configurations were investigated to study the impact of opening location on the ventilating flow and resulting CO_2 distribution. With all other parameters remaining the same, the location of the high-level vent, the classroom ventilation outlet, was modified leading to:

(i) a 'cross-flow' configuration in which the low and high-level vents are at opposite ends of the classroom (this was used for the simulations shown in Chapter 3),

(ii) a 'single-sided' configuration in which the low and high-level openings are both at one end of the classroom.

The precise position of the openings and their areas are detailed in Chapter 3 and presented in this chapter when relevant to the discussion. Using the limiting case that the air within the classroom is perfectly mixed and all properties of the air uniformly distributed throughout the volume, the so-called 'well-mixed' approximation, provided predictions which were used to define the set-up, including sizing the vents given the expected heat loads; full details are presented in Chapter 3. These prediction were based on the work of Gladstone and Woods (2001) who used a distributed heat source — such predictions, by their very nature, are unaffected by the horizontal locations of the ventilation openings. Thus, identical flow rates, room averaged temperature and steady state CO_2 concentrations were predicted for both configurations. The CFD simulations presented here aimed to investigate the extent to which this was realised, or whether the bulk ventilation flow through the classroom was in fact altered by the horizontal location of the vents. The buoyancy-driven 'cross-flow' and 'single-sided' configurations presented here are broadly representative of the limiting cases typically investigated throughout the literature for the flows in naturally ventilated rooms (e.g. Carrilho da Graça and Linden, 2016; Zhong et al., 2022), although in the cases studied here, the vents remain located on the floor and ceiling of the room. A room relying on single-sided ventilation is often assumed to be harder to ventilate as a whole as both vents are located on one side of the room (Carrilho da Graça and Linden, 2016). This can potentially lead to short-circuiting, which refers to circumstances in which the incoming fresh air exits the room through the outlet without passing through the bulk of the room (mixing and typically diluting as it travels). This often creates significant zones of relatively stagnating air, with elevated concentrations of indoor pollutants being indicative of the stagnant zones. However, the architecture of multi-zone spaces, like schools, results in the single-sided ventilation configuration being quite typical of UK classrooms, since it is practical to position windows on the same side of the building facade. The cross-flow configuration requires either openings on both sides of a room connected directly to the outdoor air, or can be achieved for classrooms with, typically the outlet, venting into a corridor (or other shared space) which then ventilates to the outdoors.

4.1.1 Bulk flow and ventilation pattern

The ventilating flow and mean temperature and CO_2 concentration in the classroom are given for both configurations in Table 4.1. The well-mixed predictions, as described in detail in Chapter 3, are presented for a range of discharge coefficients c_d.

		Q (m^3/s)	ACH (-)	Q_p (L/s/person)	ΔT (°C)	ΔT_e (°C)	ΔC (ppm)	ΔC_e (ppm)	η (−)
Cross-flow simulation		0.239	5.85	7.54	19.3	20.4	427	449	1.05
Single-sided simulation		0.241	5.85	7.54	18.9	20.2	357	444	1.24
Well mixed prediction	$c_d = 0.6$	0.221	5.35	6.90	22.0		486		1
	$c_d = 0.7$	0.245	5.93	7.64	19.9		438		1
	$c_d = 1.0$	0.310	7.52	9.69	15.7		346		1

Table 4.1: Results of the CFD simulations for the cross-flow and single-sided configurations (top two rows), and predictions assuming different values for the loss coefficient at the vents c_d (the three rows below), for: the ventilating flow rate, the room averaged excess temperature $\Delta T = T - T_a$ and excess CO_2 concentration $\Delta C = C - C_a$, with T_a and C_a the ambient temperature and CO_2 concentration respectively. In all five cases, T_e and C_e correspond to the predicted 'well-mixed' values of temperature and CO_2 calculated for the ventilating flow rate Q, the corresponding excess temperature and CO_2 concentration are calculated from $\Delta T_e = T_e - T_a$ and $\Delta C_e = C_e - C_a$ respectively. η gives a measure of the pollutant removal efficiency.

The bulk flow rates from the CFD simulations for both the cross-flow and single-sided configurations are very similar, as expected under the well-mixed assumption. The flow rates determined agree closely with predictions using the well-mixed assumption and a value of the loss efficient at the vents of $c_d = 0.7$ (a value which lies within the experimentally determined range of Heiselberg et al., 2001), predictions taking $c_d = 0.6$ (a widely used value at the lower end of loss coefficients in the literature and used by Gladstone and Woods, 2001) result in only a 10% change in the predicted flow rate. For both cases, Q_p, the flow rate per occupant, was close to 8 L/s/person, with the BB101 guidelines (Department for Education, 2018) advising ventilation rates between 8 and 9 L/s/person to achieve CO_2 levels of 1,000 ppm in a typical classroom.

Table 4.1 shows that there is a slight difference in the room averaged excess temperature ΔT, the single-sided simulation resulting in a lower excess temperature by approximately 2% when compared to the cross-flow configuration. Unsurprisingly, both values are quite close to the well-mixed predictions using $c_d = 0.7$, as expected given the similarities in flow rate and enforcing conservation of energy in the steady state. In addition, temperature is an active tracer such that variations in a horizontal plane result in buoyancy forces which act to promote motion towards horizontally uniform layers. Conversely, for the case of the room averaged excess CO_2 concentration ΔC, a passive tracer (within the simulations), the difference between the two simulations is around 16%, indicating a stronger spatial variation in CO_2 compared with temperature, i.e. within the CFD simulations, the passive tracer exhibits a stronger spatial variation compared to the active tracer, which is not unexpected. Interestingly, the average CO_2 concentration in single-sided configuration is *lower* than in the cross-flow configuration.

This indicates, perhaps counter-intuitively, that if one takes CO_2 concentration as an indoor air quality indicator, then the single-sided ventilation would be considered 'better' at supplying the classroom with fresh air, thus, reducing the overall CO_2 concentration. This finding is further investigated below, including using the metric of the 'age of air' within the classroom.

The differences between the two configurations and whether they represent a well-mixed environment was investigated by calculating the expected exhaust temperature T_e and CO_2 concentration C_e. They are found from the measured flow rate Q assuming a well-mixed classroom environment. Following Chenvidyakarn (2013), the exhaust temperature T_e was found from:

$$T_e = T_a + \frac{W_c}{\rho_a c_p Q} \; , \tag{4.1}$$

where T_a is the ambient temperature, W_c the heat input to the room, ρ_a the ambient density, c_p the specific heat capacity and Q the ventilating flow rate obtained in the simulation. The exhaust CO_2 concentration, C_e, was calculated using the method given in Chapter 3 with (3.4):

$$C_e = \frac{NG}{Q} \times 10^6 + C_a \; ,$$

with N the number of occupants, G their respective CO_2 generation rate and C_a the ambient CO_2 concentration expressed in ppm.

The performance of ventilation configurations has frequently been examined by analysing the contaminant (or sometimes 'pollutant') removal efficiency η (e.g. Novoselac and Srebric, 2003; Carrilho da Graça and Linden, 2016) defined as:

$$\eta = \frac{\Delta C_e}{\Delta C} = \frac{C_e - C_a}{C - C_a} \; , \tag{4.2}$$

with C the room averaged CO_2 concentration. Although similar to the ventilation efficiency or effectiveness defined by Sandberg (1981) and used for instance by ASHRAE (2002), contaminant removal efficiency doesn't rely on the use of a tracer gas supplied at the inlet and instead measures the efficiency in removing a contaminant, here CO_2, generated in the room. If η is equal to unity, the efficiency is equivalent to the one found in a perfectly mixed room. Values below unity indicate short-circuiting and values superior to unity are representative of displacement ventilation systems (Carrilho da Graça and Linden, 2016). For the single-sided simulation, the contaminant removal efficiency is found to be equal to 1.24 and indicative of the performance of displacement ventilation. In the cross-flow configuration, although still superior to unity, η

was calculated to be a lot closer to the efficiency expected in a perfectly well-mixed room, this once again showed that for the scenario considered here, the single-sided configuration could be considered to provide a more efficient ventilation by comparison of these metrics. For both configurations, the values calculated for η were higher than unity and showed that the classroom environment was not well-mixed. Interestingly, this was not the case in the numerical simulations performed by Tan et al. (2022), in which they considered mechanically ventilated rooms and found that the concentration estimated from a well-mixed mass balance provided an accurate estimate of the mean room concentration. This example shows the need to specifically consider buoyancy driven naturally ventilated environments to accurately represent the expected CO_2 variations.

To better understand the factors underlying the difference in performance of the two configurations, the flow pattern was investigated via the study of the streamlines within the classroom. Figure 4.1a shows the ventilating flow obtained in the cross-flow configuration and this is shown for the single-sided configuration in Figure 4.1b. In each case, the streamlines were coloured by the age of air — a statistic which represents, at a given point in the room, the amount of time that the air, at that point, has spent in the classroom since flowing through the inlet and hence gives a good measure of the 'freshness' of the air.

In both configurations, the incoming cold air can be seen flowing into the classroom through the low-level vent and falling back down to the floor as a fountain. This relatively cold air then flows along the floor across the room before hitting the opposite wall and flowing upwards. In the cross-flow configuration (Figure 4.1a), at least part of the flow directly leaves the classroom through the high-level vent. In the single-sided configuration however (Figure 4.1b), the air then flows along the ceiling before reaching the classroom outlet on the other side of the classroom. Figure 4.1b shows that the single-sided configuration did not lead to short-circuiting, and, on the contrary, created a ventilating flow that covered a larger proportion of the classroom than the cross-flow configuration. Indeed, Figure 4.1a shows a large area of the room away from the cross-flow outlet in which there is little ventilating flow. This area, covering more than the top left half of the classroom, contains relatively stale air, as shown by the high age of air, and likely led to an accumulation of CO_2 concentration, therefore increasing the overall room average. This is of particular concern as the region of stagnating air is located in the breathing zone, where the CO_2 is generated, therefore potentially increasing variations within the room. This investigation of the streamlines, and associated age of air, has shed light on the differences in the flow patterns between the two configurations which give rise to, and are evidenced by, the differences in the

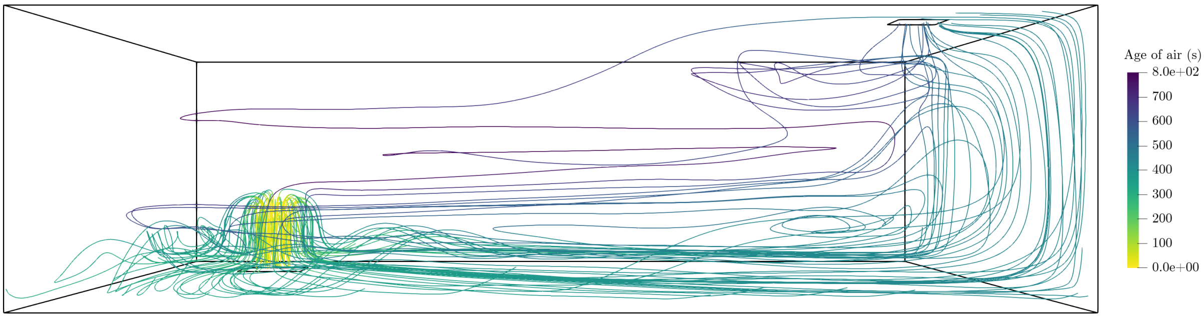

(a) Cross-flow configuration

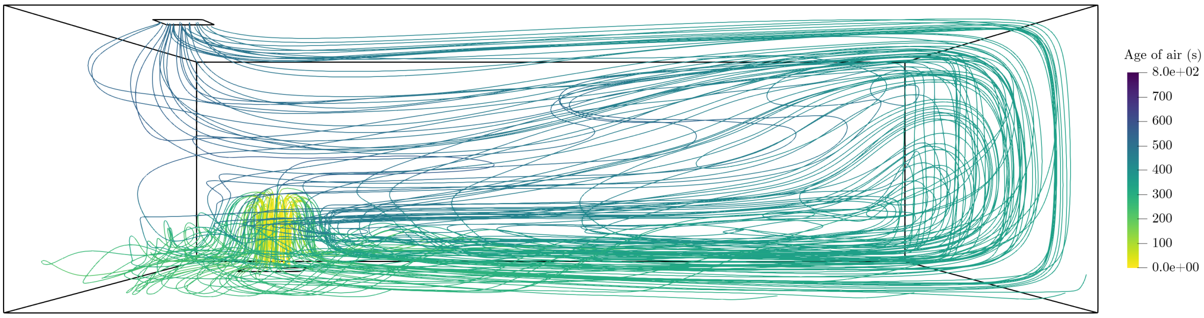

(b) Single-sided configuration

Figure 4.1: Streamlines originating from the low-level vent in both the cross-flow and single-sided configurations coloured by the age of air. In both cases, 100 streamlines were seeded at the classroom inlet.

bulk metrics for each configuration.

The differences in the ventilating pattern between the two configurations is also visible in the horizontally averaged temperature and CO_2 concentration (see Figure 4.2). The horizontally averaged temperature (Figure 4.2a) varies with height in the room from about 19°C near the floor to over 25°C near the ceiling. Up to 0.75 m high, the temperature profiles are similar for both ventilation configurations. Higher up in the classroom though, the effect of the outlet position can be observed and leads to an increase in temperature for the cross-flow configuration. In the single-sided configuration, the temperature at the top of the classroom is close to a well-mixed prediction T_e. However for the cross-flow configuration, the horizontally averaged temperature exceeds the corresponding T_e from a height of 1.5 m. However, the differences are relatively small, consistent with similarities in the room averaged temperatures reported earlier. If, similarly to η, a temperature removal efficiency is calculated from $(T_e - T_a)/(T - T_a)$, it results in very similar values for both configurations (1.06 for the cross-flow and 1.07 for the single-sided configurations respectively) which are close to a well-mixed environment. The difference in removal efficiency between CO_2 and temperature are likely due to one being a passive and the other an active scalar, as discussed above.

The similarity in the temperature profile is in contrast to Figure 4.2b) which shows the horizontally averaged CO_2 in the classroom. Both configurations exhibit a similar profile shape in which the CO_2 concentration peaks in the breathing zone between 1 and 1.5 m where CO_2 is introduced. This shows that for both configurations the CO_2 concentration was not well-mixed across the height of the room, varying by about 250 ppm and 400 ppm respectively for the single-sided and cross-flow configurations. However, in Figure 4.2b), the CO_2 concentration is consistently higher in the cross-flow scenario, exhibiting the impact of the stagnation zone highlighted in Figure 4.1a. The single-sided CO_2 concentration is below C_e over most of the classroom height, contrary to the cross-flow configuration in which the equivalent C_e was below the measured CO_2 between the heights of 0.75 m and 1.75 m. This highlights that whilst the cross-flow configuration contaminant removal efficiency was found to be close to a well-mixed value, $\eta = 1$, this occurred despite the fact that there are significant variations in the horizontally averaged CO_2 concentration from a well-mixed value, it just so happened that in this case the variations approximately sum to zero — there is no requirement that this is inherently the case, as illustrated by the findings for the single-sided configuration.

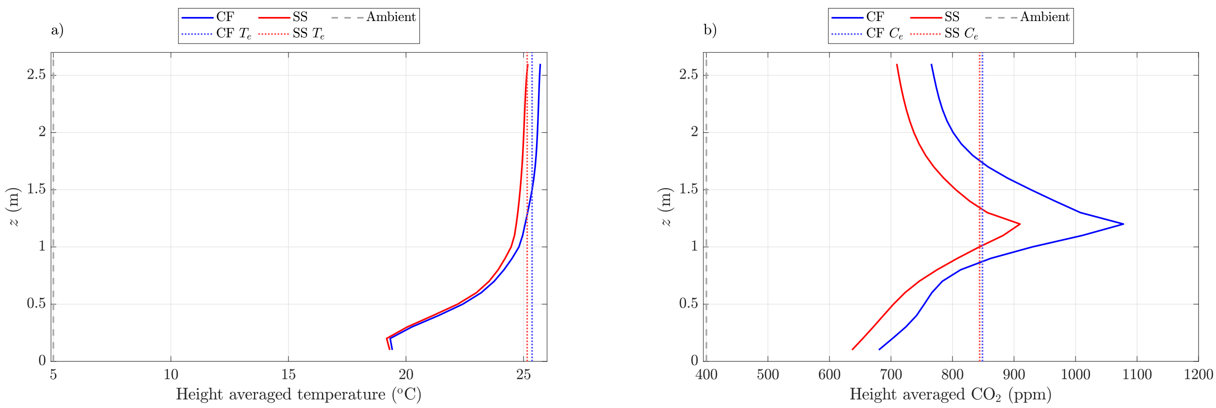

Figure 4.2: Horizontally averaged: a) temperature and b) CO_2 concentration in the classroom for the cross-flow and single-sided configurations. The dashed vertical lines show the ambient values T_a=5°C and C_a=400 ppm respectively. T_e and C_e are calculated from the flow rate assuming a well-mixed environment.

The flow field in the breathing zone was further studied as it is of direct concern and will determine the occupant's exposure. Figure 4.3 shows horizontal cross sections of the vertical velocity between the heights of 1 and 1.5 m, which accounts for the seated and standing head heights of the pupils, for both the cross-flow and single-sided configurations. Both plots show similar patterns, with small values of vertical velocity in most of the plane and a region of high vertical velocity on the right hand side, which corresponds to air flowing along the wall opposite the inlet, as described for Figure 4.1.

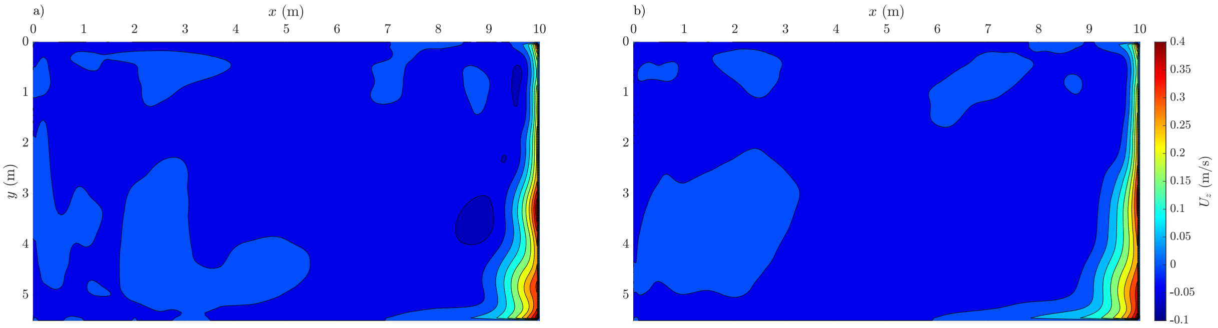

Figure 4.3: Horizontal cross section of the vertical velocity in the breathing zone ($1 \leq z \leq 1.5\,\mathrm{m}$) for: a) the cross-flow configuration, and b) the single-sided configuration.

The horizontal cross section of the temperature is shown in Figure 4.4. For both configurations, the temperature is uniform across the breathing zone apart from the right hand side of the cross-section where the high velocity region leads to a slight increase in temperature. As shown in Figure 4.2, the cross-flow configuration exhibits temperatures slightly higher than for the single-sided ventilation.

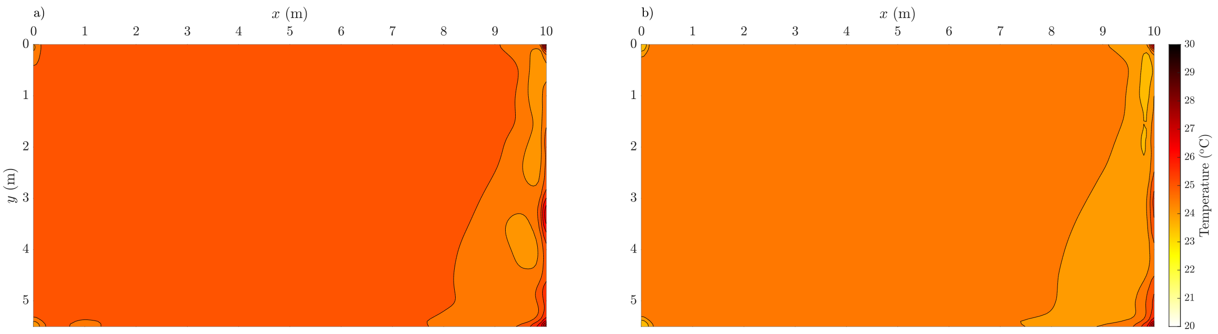

Figure 4.4: Horizontal cross section of the temperature in the breathing zone ($1 \leq z \leq 1.5\,\mathrm{m}$) for: a) the cross-flow configuration, and b) the single-sided configuration.

For the CO_2 concentration, shown in Figure 4.5, the horizontal cross section in the breathing zone differs between the two configurations. For both configurations however, there are significant variations in the CO_2 concentration in the breathing zones which would contradict a well-mixed assumption. In particular, for the cross- flow configuration, the concentration is shown to vary by at least 600 ppm, with the lowest concentration found on the right of the plane, in the region of high vertical velocity shown in Figure 4.3. This is also observed in Figure 4.5b), for the single-sided configuration, although the CO_2 variations are lower, spanning 300 ppm. This difference in the distribution of CO_2 in the breathing zone can be attributed to the different ventilating patterns observed for instance in Figure 4.1, where a larger stagnating region is present in the cross-flow configuration leading to an increase in CO_2 concentration.

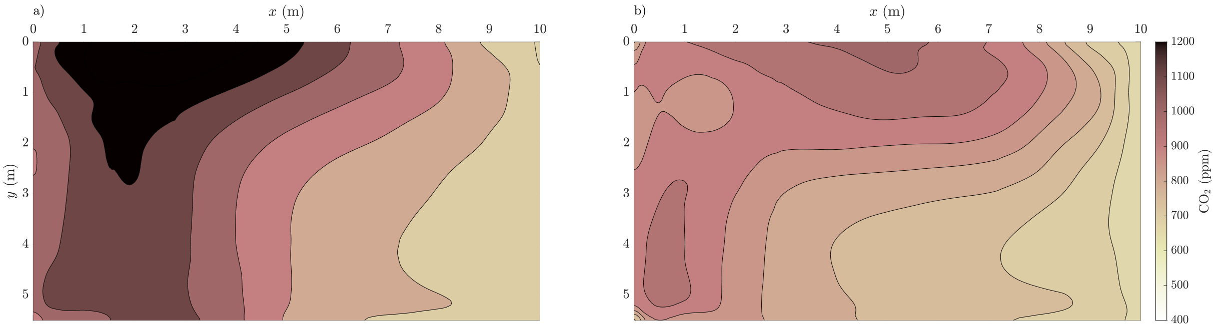

Figure 4.5: Horizontal cross section of the CO_2 concentration in the breathing zone ($1 \leq z \leq 1.5\,\mathrm{m}$) for: a) the cross-flow and b) the single-sided configuration.

4.1.2 CO_2 distribution

Given both the more significant differences between the two ventilation configurations and its prevalence as an Indoor Air Quality (IAQ) marker, the CO_2 distribution was examined more closely in each ventilation configuration by looking at a wider variety of statistics. Table 4.2 compares the CO_2 distribution in the breathing zone ($1 \leq z \leq 1.5\,\mathrm{m}$) and the overall classroom, by reporting values for the mean CO_2 concentration and the coefficient of variation in each region. The coefficient of variation of the CO_2 concentration, c_v, is calculated in each domain by dividing the unbiased estimate of the standard deviation by the mean. It gives a measure of the dispersion of the values of the CO_2 concentration for each region and configuration. In the

Configuration	Location	C (ppm)	c_v (%)
Cross-flow	Classroom	827	16.6
	Breathing zone	993	18.9
Single-sided	Classroom	757	11.6
	Breathing zone	852	11.7

Table 4.2: Statistics of the CO_2 concentration for the overall classroom and the breathing zone ($1 \leq z \leq 1.5\,\mathrm{m}$) in the two ventilation configurations. C is the mean CO_2 concentration in each domain and the coefficient of variation c_v is the ratio of the unbiased estimate of the standard deviation relative to the mean.

single-sided configuration, the mean CO_2 concentration in the breathing zone is higher than for the overall classroom by 13%. This increase is not unexpected given that the CO_2 is generated in this zone. The coefficient of variation c_v is similar for both regions under consideration. In the cross-flow configuration however, there is larger increase in the mean CO_2 in the breathing zone (20% compared to the overall classroom) and both values of c_v also increase. In this configuration, the coefficient of variation is also larger in the breathing zone, which implies a larger dispersion of the CO_2 concentration and can be explained by the stagnating region observed in the cross-flow configuration. In both regions, there is a significant difference in the room averaged CO_2 between the two ventilation configurations, with the highest values found for a cross-flow case. In the breathing zone for the cross-flow configuration, the mean CO_2 is 17% higher than in the single-sided configuration which is of the same order of magnitude as the coefficients of variation which is therefore not negligible.

Histograms of the distribution of CO_2 in both configurations are also shown in Figure 4.6: a) for the entire classroom, and b) in the breathing zone only. For reference the median values and interquartile ranges are also plotted. All four distributions are skewed with large tails towards higher CO_2 values and in both regions, higher values of CO_2 are present in the cross-flow configuration. The distributions in the classroom, plotted in Figure 4.6a), show a large number of values at low CO_2 levels, where the local value is likely to be below the mean in the room. Both the cross-flow and single-sided configurations exhibit a large tail of high concentration values, with CO_2 above 900 ppm. Figure 4.6b) shows the distribution in the breathing zone where larger concentrations are present. In this region, the cross-flow configuration displays a larger spread of values which is in agreement with the coefficient of variation given in Table 4.2.

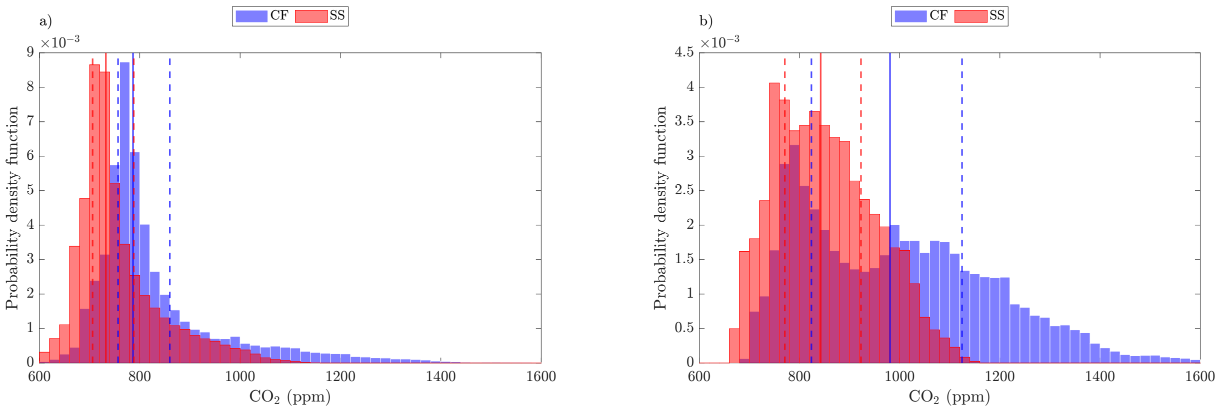

Figure 4.6: Histogram of the CO_2 concentration for the cross-flow and single-sided configuration for: a) the whole classroom and b) the breathing zone ($1 \leq z \leq 1.5$ m) using 50 bins. For each case, the median is shown with a full vertical line, the first and third quartiles are shown with dashed vertical lines.

This analysis demonstrates that a wide range of CO_2 values were present in the room, particularly in the breathing zone, for both ventilation configurations and that in both cases the classroom was far from well-mixed. In addition to larger room averaged values of CO_2, the cross-flow configuration was also shown to exhibit larger variations across the space, which are likely to be more challenging to represent accurately with a limited number of point measurements, typical of existing sensing technologies.

4.1.3 Sensor placement

The simulations presented here provide access to the full three dimensional fields in the classroom; we exploited this to assess where CO_2 sensors could be placed to take measurements of greater relevance to occupants. In particular, the CO_2 concentration across the whole room at different heights was compared to the concentration measured near walls only. Although the presence of walls might impact the transport of scalars in their near field, this was chosen because sensors are often placed out of the way and close to the walls of the room. Here, a measurement at the walls was taken to be the concentrations measured within a region 0.2 m away from the classroom walls.

Histograms of the CO_2 distribution at different heights, comparing the concentration over the entire plane with the concentration near the classroom walls, are plotted for the cross-flow (Figure 4.7) and single-sided (Figure 4.8) configurations. The greatest variation is observed in the breathing zone, particularly at $z = 1.1$ m. At lower heights and in the breathing zone, differences are visible between the overall CO_2 concentration at that height and the one measured at the walls. Depending on the ventilation configuration selected, this can be either an overestimate or an underestimate, with larger or smaller spreads observed near the walls. It is therefore difficult to reach any universal conclusions on the suitability of wall measurements in the bottom half of the classroom, shown here for heights between 0.3 and 1.5 m. This was due to the complexity of the flow in these regions and its effect on the CO_2 distribution. However, since in both configurations, the flow enters at low-level and exits at high-level, above the breathing zone the spread in the CO_2 distribution was shown to decrease — presumably as the flow mixes throughout its transport. At $z = 2.3$ m for instance, the distribution of the CO_2 measured near the classroom walls is close to the distribution across the whole room at that height, both in terms of the median and interquartile range. This was especially true in the single-sided configuration, which as discussed above, presented smaller variations in the CO_2 distribution compared to the cross-flow configuration.

If CO_2 measurements at the walls were to be used to estimate a ventilating flow rate, measurements taken above the breathing zone were found to be more likely to give accurate reading due to the lower spread of the results. This is consistent across 0.5 m between the distributions shown at heights 1.9 and 2.3 m, with distributions at 2.3 m showing the smallest variations and the closest agreement between the whole room and wall measurements for both the cross-flow and single-sided configurations. In this scenario, the high-level vent was located on the ceiling. It can be assumed that for ventilation designs promoting inflows at low-level and

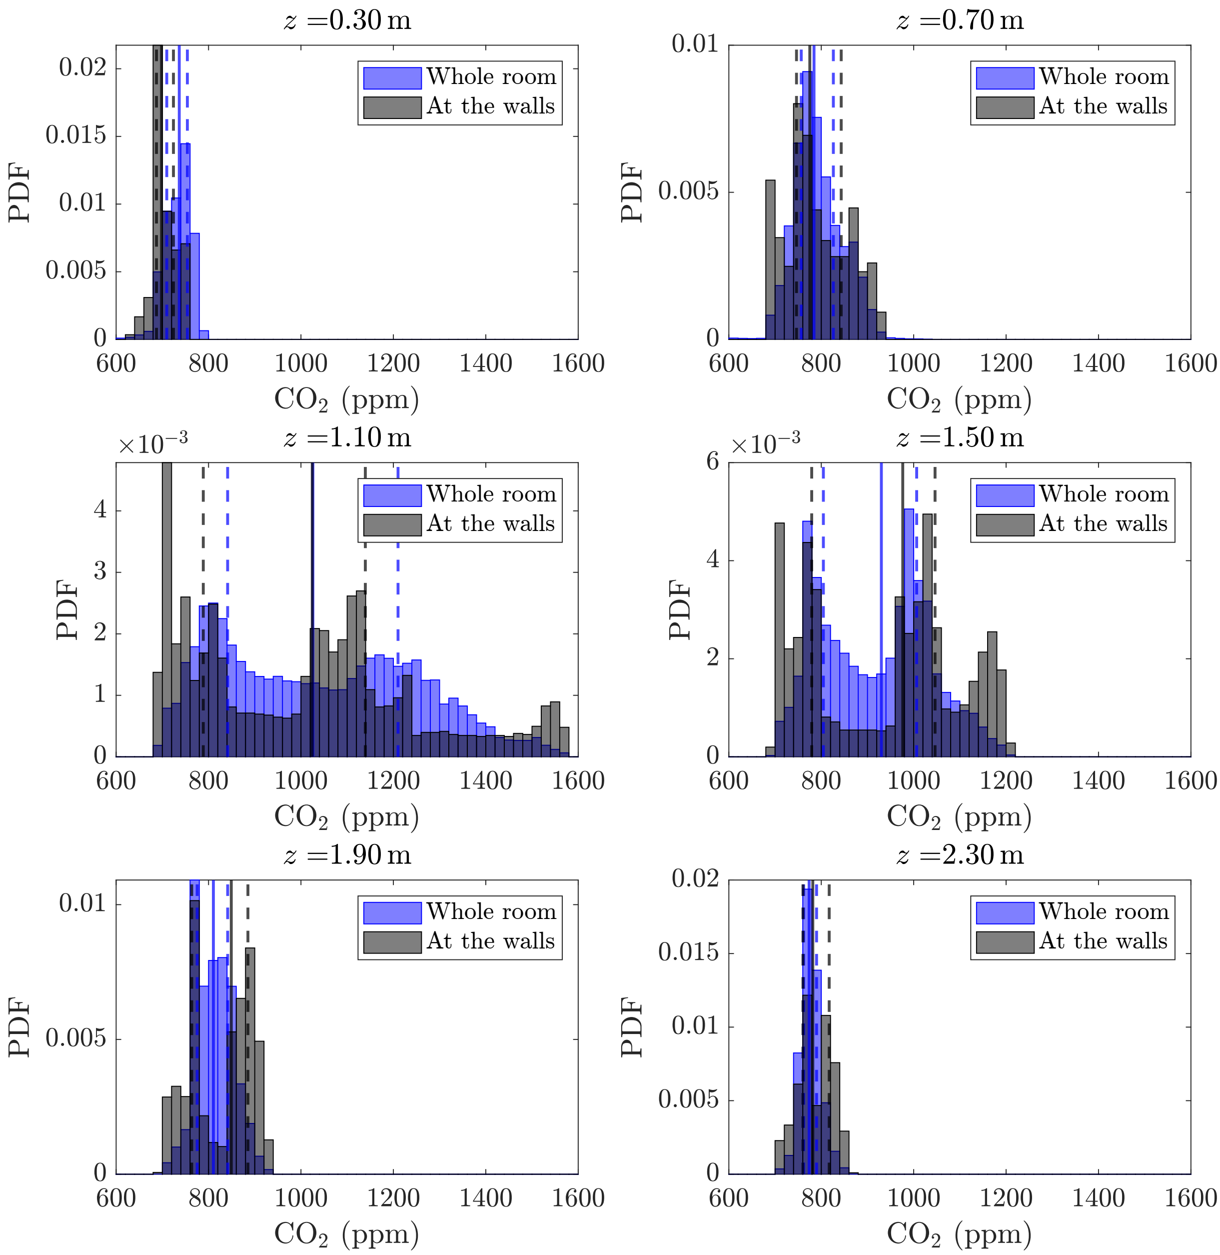

Figure 4.7: Histogram of the CO_2 concentration for the cross-flow configuration at different heights using 50 bins. The CO_2 concentration across the whole room at every height is compared to measurements at the walls (defined as within 0.2 m of a wall). For each case, the median is shown with a full vertical line, the first and third quartiles are shown with dashed vertical lines.

outflows at high-level, a case where the high-level vent is located on a wall below the ceiling, the best agreement will be achieved above the breathing zone and below the position of the vent, to ensure that measurements do not include concentrations above the vent where CO_2 might accumulate.

Table 4.3 summarises these results by comparing statistical results at heights 1.1 and 2.3 m which were respectively found to be the most and least dispersed distributions. The error ϵ was

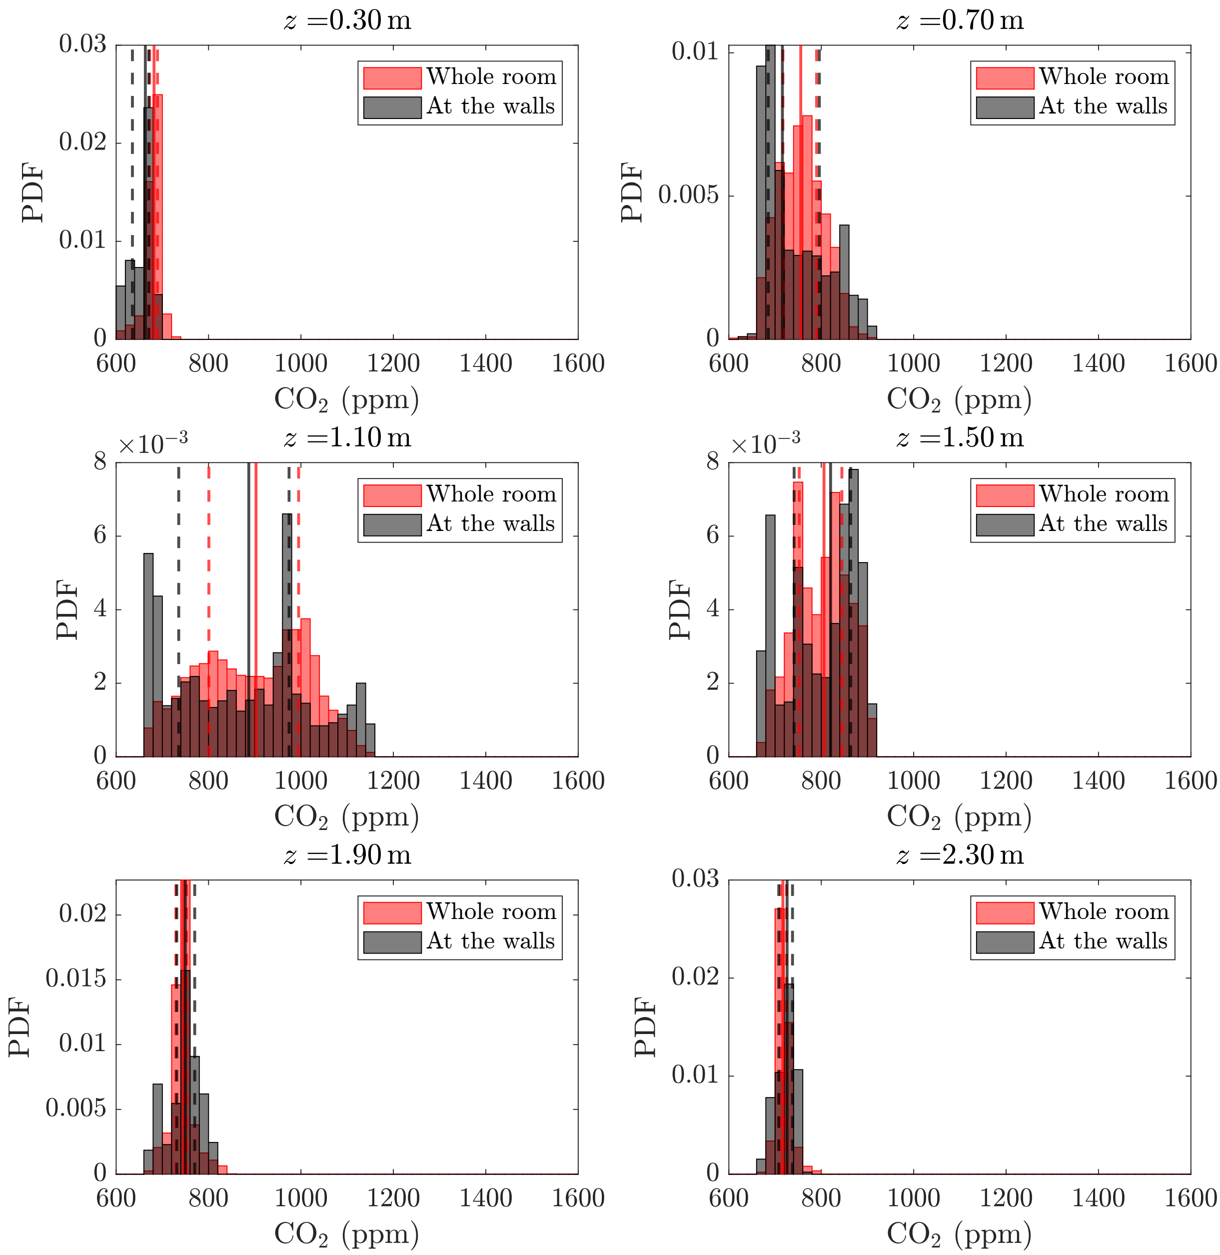

Figure 4.8: Histogram of the CO_2 concentration for the single-sided configuration at different heights using 50 bins. The CO_2 concentration across the whole room at every height is compared to measurements at the walls (defined as within 0.2 m of a wall). For each case, the median is shown with a full vertical line, the first and third quartiles are shown with dashed vertical lines.

first calculated to quantify the error in the predicted mean at every height by considering the wall region only against considering the overall cross section at that height. It was calculated from:

$$\epsilon = \frac{|\langle C_w \rangle - \langle C \rangle|}{\langle C \rangle} \,. \tag{4.3}$$

where $\langle C_w \rangle$ is the mean CO_2 concentration at a given height measured near the walls and $\langle C \rangle$ is the mean CO_2 concentration across the whole room at that height. This table shows that the use of wall measurements leads to an error exceeding 2% in the breathing zone ($z = 1.1$ m), with a larger error in the cross-flow configuration. Higher in the classroom, at $z = 2.3$ m, the error ϵ drops, it is more than halved in the cross-flow configuration and reaches 0.5% in the single-sided configuration.

Simulation	Height	ϵ (%)	c_v (%)	$c_{v(w)}$(%)
Cross-flow	$z = 1.1$ m	2.9	20.4	23.9
	$z = 2.3$ m	1.2	2.8	4.7
Single-sided	$z = 1.1$ m	2.5	13.0	16.7
	$z = 2.3$ m	0.5	2.1	3.0

Table 4.3: Accuracy of the CO_2 measurements at the walls at two heights for the two ventilation configurations. ϵ is the error in the predicted mean when considering the CO_2 concentration near the walls (defined as within 0.2 m of a wall). c_v is the coefficient of variation of the CO_2 concentration over the whole room at a given height and $c_{v(w)}$ quantifies the coefficient of variation at the same height considering measurements near the walls only.

The coefficients of variation were also used to assess whether wall measurements are capable of illustrating the spread of CO_2 distribution within the room. Overall, the coefficients of variation are higher in the cross-flow configuration, and, for both configurations and at both heights, the distributions exhibit a larger spread at the walls, as shown by larger values of $c_{v(w)}$. Both c_v and $c_{v(w)}$ drop higher in the classroom which matches the trends observed in Figure 4.7 and Figure 4.8. In the case of buoyancy driven ventilation simulated here, wall measurements were found to be able to capture accurately the CO_2 concentration if located high in the classroom above the breathing zone and below the ventilation outlet — these results are likely to remain valid for classrooms in which the ventilation design promotes low-level inflows and high-level outflows.

4.2 Sensitivity of the ventilation patterns and CO_2 distribution to changes in the flow rate

The results presented so far in this chapter have shown that the single-sided configuration provided more efficient ventilation, in terms of better dilution of IAQ markers, and led to lower CO_2 levels. The cross-flow configuration was also shown to result in higher variations of CO_2 in the classroom. In addition, for both configurations, the variation of CO_2 concentration might be deemed acceptably small (compared with the mean) at higher vertical location, i.e. above the breathing zone, which allows for the expectation of relatively accurate predictions of the mean CO_2 concentration based on a limited number of point measurements. In this section, a range of different vent areas are selected to investigate how robust these results might be.

Different simulations were used to perform this analysis and are summarised in Table 4.4. The reference set-up described in §4.1 and in Chapter 3 corresponds to the original vent set-up (OS). The incoming ventilating flow rate was modified in two ways:

(i) The areas of the classroom openings were varied, they were reduced in simulations SO1 and SO2 and enlarged in the LO1 and LO2 cases. In the SO1 set-up, both the inlet and outlet areas were halved and they were both doubled in the LO1 case. The SO2 and LO2 were designed by changing the area of the high-level vent only to achieve effective vent areas A^* equivalent to those calculated for the SO1 and LO1 cases.

(ii) The heat input was reduced in simulation LH, corresponding to a case where the sole heat input is from the occupants. This was designed so that the expected flow rate would correspond to the one predicted with small openings (SO1, SO2).

Simulation	Description	A_B (m^2)	A_T (m^2)	A^* (m^2)	W_c (W)
OS	Original set-up	0.40	0.20	0.18	6200
SO1	Small openings 1	0.20	0.10	0.09	6200
LO1	Large openings 1	0.80	0.40	0.35	6200
SO2	Small openings 2	0.40	0.10	0.10	6200
LO2	Large openings 2	0.40	0.80	0.35	6200
LH	Low Heat	0.40	0.20	0.18	1920

Table 4.4: List of simulations performed. The simulations described in Chapter 3 and §4.1 correspond to the original vent set-up (OS). A^* is calculated using $c_d = 0.7$ which previously gave the best agreement.

4.2.1 Bulk flow

The bulk flow rates, room averaged temperature and CO_2 concentration in the classroom are given in Table 4.5 for each simulation and ventilation configuration. Overall, the flow rates obtained are close to that predicted using the well-mixed predictions and $c_d = 0.7$ (Gladstone and Woods, 2001, further details see Chapter 3). These are summarised in square brackets in Table 4.5. The flow rate obtained in each simulation is lower than predicted with the well-mixed prediction (with the exception of the SO2 case), with single-sided configurations leading to larger flow rates. Cases with lower flow rates (LH, SO1 and SO2) correspond to those with relatively smaller differences between the flow rates in the single-sided and cross-flow configurations.

Name	Set -up	ACH		Q_p (L/s/person)		T (°C)		T_e (°C)	C (ppm)		C_e (ppm)	η
OS	CF	5.79	[5.93]	7.47	[7.64]	24.3	[24.9]	25.4	827	[838]	849	1.05
	SS	5.85		7.54		23.9		25.2	757		844	1.24
SO1	CF	3.73	[3.74]	4.81	[4.81]	35.4	[36.6]	36.6	1087	[1096]	1096	1.01
	SS	3.73		4.81		35.2		36.6	968		1096	1.23
LO1	CF	8.88	[9.41]	11.45	[12.13]	17.2	[17.5]	18.3	679	[676]	693	1.05
	SS	8.94		11.52		16.9		18.2	646		691	1.18
SO2	CF	3.99	[3.94]	5.15	[5.08]	33.1	[34.9]	34.5	1161	[1059]	1051	0.86
	SS	4.00		5.15		33.0		34.5	946		1050	1.19
LO2	CF	8.54	[9.41]	11.01	[12.13]	17.5	[17.5]	18.8	646	[676]	704	1.24
	SS	8.66		11.16		17.5		18.6	627		700	1.32
LH	CF	3.97	[4.01]	5.11	[5.17]	13.7	[14.1]	14.2	1017	[1048]	1055	1.06
	SS	4.00		5.16		13.5		14.1	902		1049	1.29

Table 4.5: Bulk flow and room averaged results for the different simulations. Predicted values with $c_d = 0.7$ are shown in square brackets. T_e and C_e are calculated based on the measured flow rate assuming a well-mixed environment as described in §4.1. η is a measure of the contaminant removal efficiency.

Similar trends are observed for the classroom averaged temperature, with smaller flow rates leading to higher temperatures in the room when the vent areas are modified (SO1 and SO2). In the case where the heat input was reduced (LH), the mean temperature is considerably lower, as can be expected. T_e was calculated assuming the room was well-mixed as detailed above using (4.1). In all cases, the actual room temperature was lower than T_e.

In general, a larger flow rate leads to lower mean CO_2 concentration, with the cross-flow configuration leading to higher values than the single-sided configurations in all cases. The contaminant removal efficiency η metric was calculated from (4.2) and allowed to ascertain whether the CO_2 concentration varied following the changes in flow rate. SO1, LO1 and LH gave similar results to those obtained in the original vent set-up (OS): a value close to 1, and therefore close to the well-mixed case, for the cross-flow configuration and a value close to 1.2 for the single-sided configuration, indicative of a displacement ventilation efficiency. For SO2 and LO2 this was not the case, despite respective effective areas A^* close to those used in SO1 and LO1 and similar flow rates. The single-sided configuration of the SO2 set-up had an efficiency similar to the other single-sided configurations. In the cross-flow configuration however, the SO2 set-up had a very low efficiency which indicates the presence of stagnation zones in the classroom.

For the LO2 set-up, on the other hand, the values of ventilation efficiency were higher in both configurations and superior to 1.2. The flow structure of both cases is detailed below but it highlighted that, in this context, the effective area A^* could not be used to predict the mean CO_2 and resulting contaminant removal efficiency, instead, the area of both openings needed to be accounted for.

The horizontally averaged temperature and CO_2 concentration are plotted in Figure 4.9 for all simulations. The left column, Figures 4.9a) and c) show the cross-flow configuration. The single-sided results are shown in the right column, in Figures 4.9b) and d). The temperature profiles follow the increase in room averaged temperature shown in Table 4.5. The variations of temperature observed in the room between the floor and the ceiling increase as the overall mean temperature also increases. For instance, in the LH cases (with the lowest mean temperature), the horizontally averaged temperature varies by about 4°C. On the other hand, in the SO1 simulations (with the highest mean temperatures), the temperature varies by about 8°C over the height of the classroom.

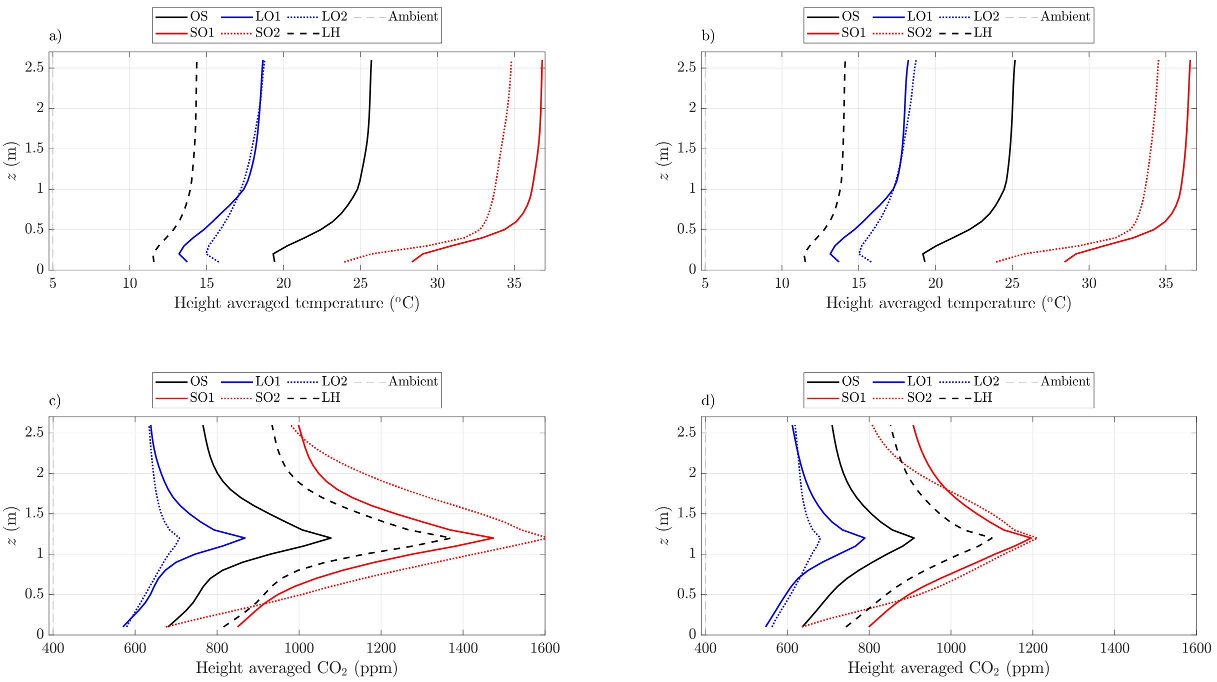

Figure 4.9: Comparison of the horizontally averaged temperature (top panes) and CO_2 concentration (bottom panes) in the classroom for different vent and heat input set-ups using the cross-flow (left panes) and single-sided (right panes) configurations.

Trends are similar to those observed in the original vent set-up (OS) for the CO_2 profiles shown in Figures 4.9c) and d), with a significant increase in CO_2 concentration observed for the different cross-flow configurations when compared to the respective single-sided set-ups. Two notable exceptions are observed: the CO_2 concentration in both LO2 simulations varies significantly less across the height of the classroom than what is observed for the LO1 simulations. Although the concentration is similar near the floor and ceiling, the CO_2 concentration in the breathing zone is considerably lower in the LO2 simulations when compared to LO1. Contrastingly, the concentration in both SO2 simulations largely peaks above all other simulations in the breathing zone but the CO_2 concentration near the floor is significantly smaller than what is observed for simulations with comparable flow rates (SO1 and LH).

The ventilation patterns for the SO2 and LO2 set-ups in the cross-flow configuration was studied further by looking at the streamlines in each case plotted in Figure 4.10. These can be directly compared to Figure 4.1a which showed the streamlines for the original vent set-up (OS), also in the cross-flow configuration. All three simulations shown have identical inlets and heat input, their only difference being the area of the high-level vent. By keeping the low-level vent area constant, the effect of the change in ventilating flow rate can be seen directly. Differences are most apparent in the fountain created at the inlet of the classroom by the relatively cold rising inflow. With a low flow rate, in the SO2 case, the fountain rises significantly less than the original vent set-up (OS) and this creates a much larger stagnating zone, with very stale air as indicated by the old age of air. In the LO2 case however, the increase in ventilating flow rate also increases the height of the fountain at the inlet which rises more than halfway across the height of the classroom. This encourages recirculation in the room and no clear stagnation zones are visible. Instead, the streamlines cover most of the classroom and the air is a lot fresher overall leading to age of air values close to 400 s overall, half of what could be observed in the OS and SO2 set-ups. Interestingly, although the general flow rate variations are comparable in the SO1 and LO1 cases, the impact on the height of the fountain is less striking which explains why these set-ups display similar patterns to the original vent set-up (OS). For both the SO2 and LO2 cases considered here, the contaminant removal efficiency was useful to identify cases which did not match previous patterns and which required further investigation, η being equal to 0.86 in the SO2 case (compared to 1.05 for the original vent set-up cross-flow configuration) and equal to 1.24 for the LO2 case (and therefore closer to the efficiency of the single-sided configuration).

4.2.2 CO_2 distribution

The impact of the changes of flow rate on the CO_2 concentration in the room is summarised by the box plots shown in Figure 4.11. This plot compares the classroom CO_2 distribution in each vent and heat input set-up for both the cross-flow and single-sided configurations.

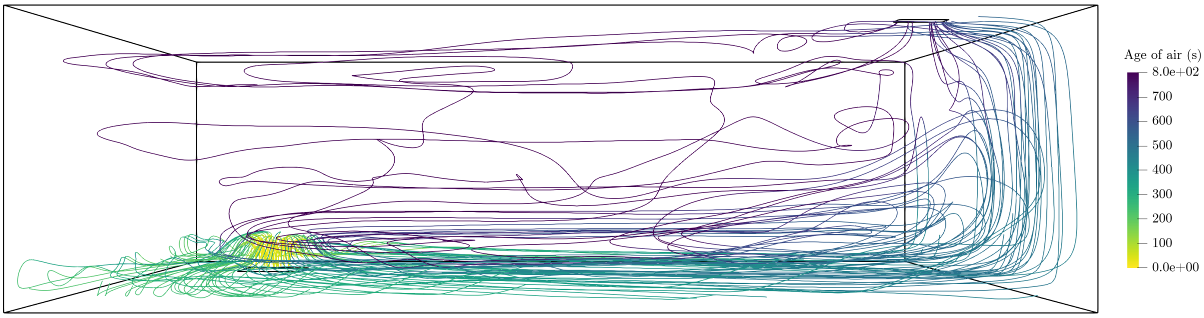

(a) SO2

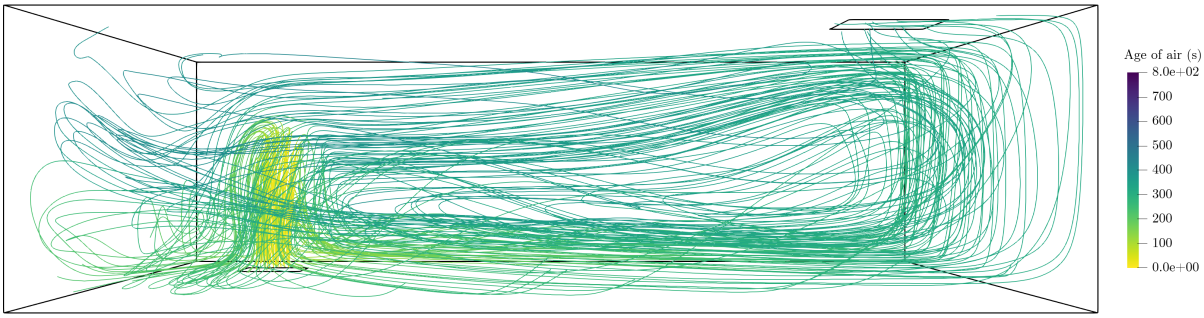

(b) LO2

Figure 4.10: Streamlines originating from the low-level vent in the cross-flow configuration for the SO2 and LO2 cases coloured by the age of air. In both cases, 100 streamlines are seeded at the classroom inlet.

Figure 4.11 shows that the values of CO_2 in the classroom are consistently lower in the single-sided configurations, often accompanied by lower spreads in the measured values. For similar geometries, the spread and overall CO_2 in the room increases as the flow rate decreases. This can be seen for instance by comparing the distributions of the LH case with the original vent set-up (OS). The differences between the cross-flow and single-sided distributions reduce as the ventilating flow rate increases, better agreement is seen for instance for the LO1 and LO2 set-ups. The SO2 case displays the highest spread in CO_2 concentrations, with values ranging from the ambient 400 ppm to over 2,000 ppm in the cross-flow configuration, although its median is similar to the one of the SO1 set-up. This is one example which shows the value of using CFD simulations to evidence patterns that would otherwise be challenging to observe and explain through a limited number of measurement points.

4.2.3 Sensor placement

The sensitivity of the results presented in §4.1.3 was studied to assess how they were affected by changes in the ventilating flow rate. Figure 4.12 summarises the errors in the predicted mean using wall measurements ϵ and the coefficient of variations found using the measurement over

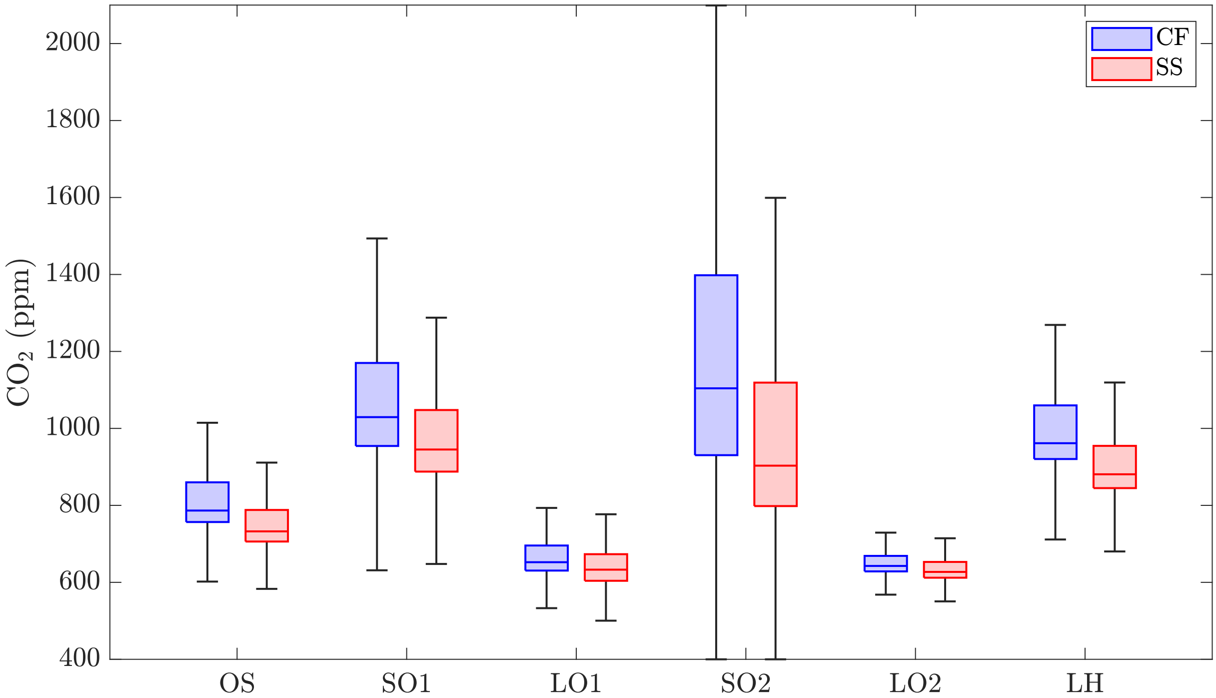

Figure 4.11: CO_2 distribution in the classroom for different vent and heat input set-ups in the cross-flow (blue) and single-sided (red) configurations. The bottom and top of the box show the 25^{th} and 75^{th} percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1^{st} and 3^{rd} quartiles.

the entire height c_v and at the walls only $c_{v(w)}$. This is shown for the two heights which showed the largest differences in §4.1.3: $z = 1.1$ m and $z = 2.3$ m. Wall measurement were defined as concentration taken within 0.2 m of a vertical wall and the results are shown both in the cross-flow and single-sided configurations.

As shown previously, the error and coefficients of variations are worse near the breathing zone, which can be expected. Agreement between wall measurements and those over the entire plane is a lot closer higher in the classroom for all cases, with errors in the predicted mean ranging from less than 1% to 4%, coefficient of variations varying between 1% and 8% in the domain and 1% and 12% at the walls. All parameters show distinct peaks at 2.3 m in the cross-flow configuration for the smaller opening scenarios (SO1, SO2) which were found to have a larger spreads in the CO_2 distributions therefore leading to higher errors and coefficient of variations. In the breathing zone ($z = 1.1$ m), trends are not as clear, the error ϵ ranges from 1% to 12% over the different vent and heat input set-ups, with the single-sided configuration leading to better results (and a smaller error) in all but 2 cases (SO1 and LO1). In both ventilation configurations, the maximum error is reached for the SO2 case. At $z = 1.1$ m, the coefficients of variation vary between 9% and 25% in the domain and 9% and 29% at the walls, with consistently higher

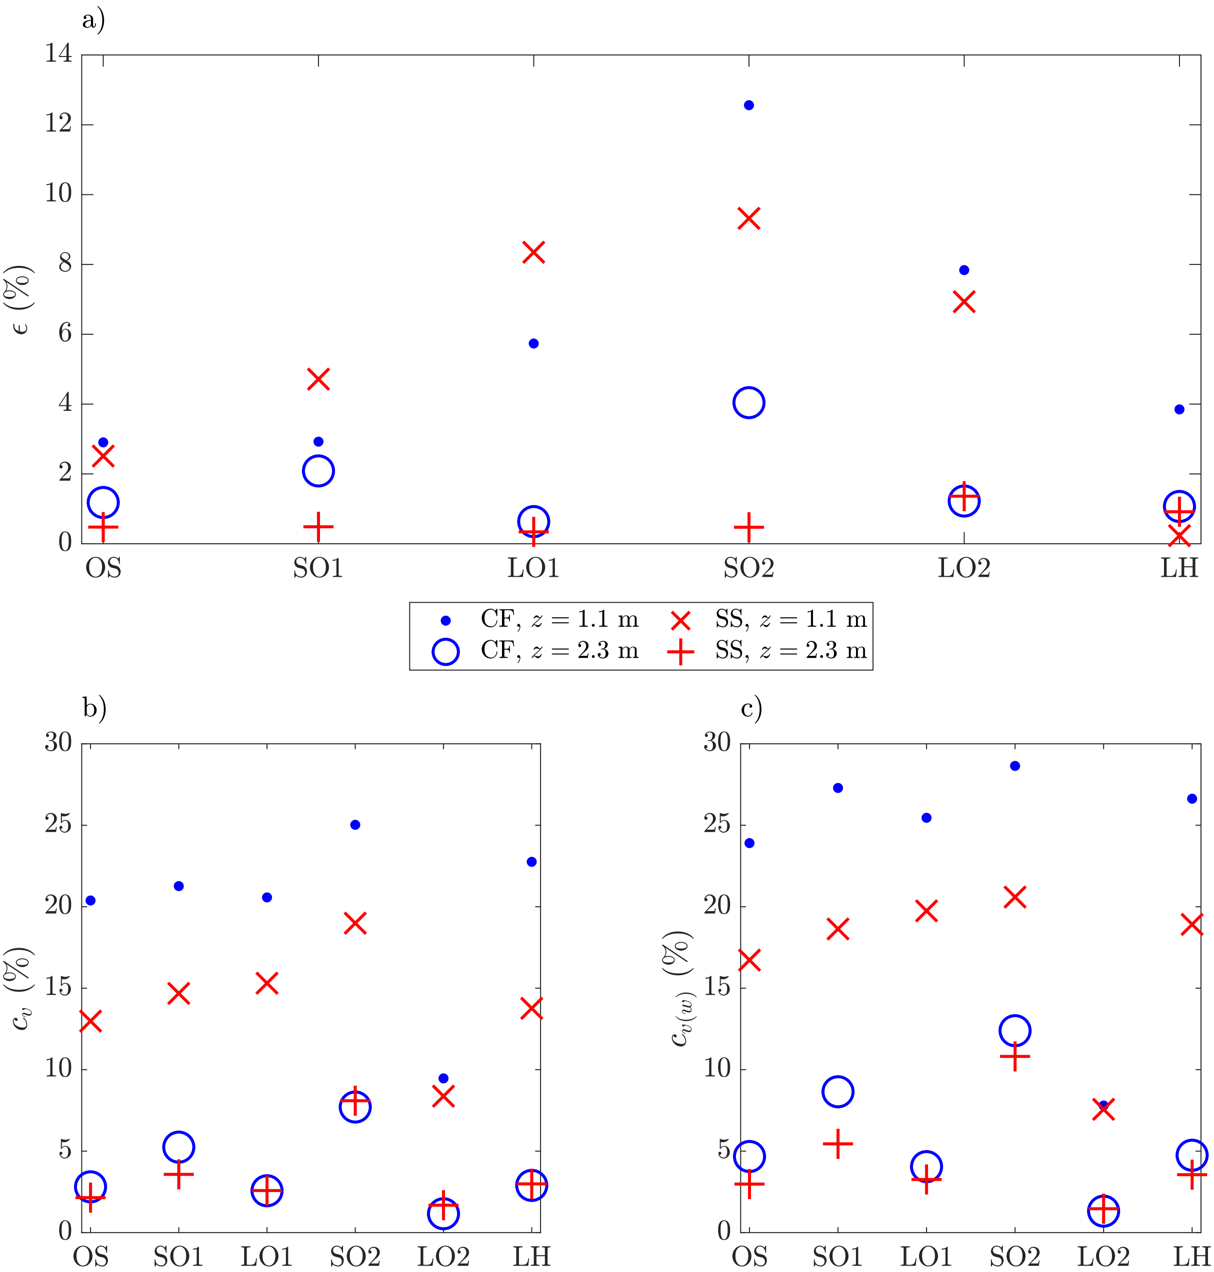

Figure 4.12: Accuracy of the CO_2 measurements at the walls at $z = 1.1$ m and $z = 2.3$ m for the different vent and heat input set-ups. Results for both the cross-flow and single-sided configurations are shown. ϵ is the error in the predicted mean when considering the CO_2 concentration near the walls (defined as within 0.2 m of a wall). c_v is coefficient of variation of the CO_2 concentration over the whole room at a given height and $c_{v(w)}$ quantifies the coefficient of variation at the same height considering measurements near the walls only

values for the cross-flow configuration. Here, a significant decrease in the coefficient of variation is visible for the LO2 simulation, which is explained by the increased recirculation in the room observed in Figure 4.10b.

Across the different simulations considered here, wall measurements were found to be more accurate at points higher in the room. The CO_2 distribution in the breathing zone displayed larger variations for all vent and heat input set-ups and would be hard to measure accurately with a limited number of sensors. Despite significant changes in the ventilation pattern created by changes in the geometry or the heat input, conclusions drawn for the original vent set-up held up for the different vent sizes and heat inputs considered. Larger variations of CO_2 were observed using the cross-flow configuration and wall measurements taken above the breathing zone and below the high-level vent were found to be the most accurate to represent the distribution at that height.

4.3 Conclusions

Significant spatial variations in the CO_2 concentration were observed in simulations intended to represent a broadly generic naturally ventilated UK classroom. These variations are apparent even when averaged either horizontally or vertically, showing that the environment was far from well-mixed despite the distributed input of the heat loads. Somewhat surprisingly, when the low-level inlet and high-level outlet vents were positioned at one end of the classroom (single-sided configuration) the ventilation strategy was shown to be most efficient, as determined by a number of standard measures, consistently leading to lower CO_2 concentrations in the classroom. This highlights that in addition to the areas of the vents, the relative position of these vents needs to be carefully considered to determine the effectiveness of the ventilating flow established, and the resulting CO_2 concentration in spaces fitted with buoyancy driven ventilation.

The variations of CO_2 were shown to be highest within the breathing zone ($1 \leq z \leq 1.5\,\mathrm{m}$) and the spread in the observed concentration decreased at greater heights. Measurements near the wall (within $0.2\,\mathrm{m}$) were compared to the distribution found over the whole room at each height. This is pertinent in classrooms, as sensors are often placed out of the way and fixed to walls. This comparison showed that, for ventilation strategies which promote low-level inflows and high-level outflows, sensors placed near walls could be useful to predict the mean CO_2 at that height, provided they were above the breathing zone and below the classroom ventilation outlet. CO_2 measurements should however be used with caution when used to infer ventilation rates, even if they are broadly representative of the mean CO_2. In the single-sided configuration for instance, where the pollutant removal efficiency was superior to unity (the well-mixed efficiency), calculating the ventilation flow rate from the room averaged CO_2 would lead to overestimating the actual ventilation supply by 25%.

The robustness of these results was investigated by considering a range of scenarios in which the ventilating flow rate was varied either by changing the opening areas or the heat input. Overall, a lower flow rate led to higher CO_2 levels and higher variations in the concentration within the classroom. Changes in the relative sizes of the vents was shown to significantly impact the CO_2 distribution even if the overall effective area A^* was kept constant. This was shown to be a result of the fountain like flow at the classroom inlet. The fountain height specifically, was shown to have a direct impact on recirculation in the classroom and the contaminant removal efficiency in the cross-flow configuration. CFD simulations proved to be essential to understand the different phenomena involved which could have not be identified through room averaged measurements.

This chapter has showed that considerable variations in CO_2 can be expected in generic naturally ventilated classrooms, even in the limiting case that the heat loads are well distributed, and that the well-mixed assumption should be used with caution, ideally alongside account of the expected variations. Whether CO_2 measurements are suitable to assess far-field exposure to infected breath remains an open question — this forms the focus of the subsequent chapter.

Chapter 5

Variations in far-field exposure and infection risk

In Chapter 4, Computational Fluid Dynamics (CFD) simulations of a generic naturally ventilated classroom showed significant spatial variations of CO_2 concentration, with the scalar being far from 'well-mixed'. This was the case for both the cross-flow and single-sided configurations. CO_2 measurements taken above the breathing zone were shown to be broadly representative of the overall room average. This held for a range of scenarios with different vent areas and heat inputs, despite demonstrable changes in the ventilation pattern. In Chapter 2, CO_2 measurements were used to calculate the rebreathed fraction and thus estimate the risk of airborne infection within the classroom. In that chapter, CO_2 measurements were assumed to represent exposure to infected breath. The aim of this chapter is, given the findings of Chapter 4, to determine how accurate the assumptions underpinning the analysis of Chapter 2, and other works (e.g. Rudnick and Milton, 2003), might be within a generic naturally ventilated classrooms studied within this thesis.

This chapter first introduces how the infected breath is modelled in §5.1. §5.2 shows the distribution of infected breath depending on the infector location and ventilation configuration. The resulting exposure is compared to the one calculated using CO_2 measurements in §5.3 and §5.4 addresses the question of the optimal sensor placement in the context of far-field risk assessment. §5.5 assesses the robustness of the results to changes in the ventilating flow rate through varying the vent areas and heat inputs. Conclusions are drawn in §5.6.

5.1 Methodology

In Chapter 4, the distribution of CO_2, produced by all occupants, was analysed. CO_2 was introduced as a passive scalar over the entire cross section of the classroom between the heights of $1.1 \leq z \leq 1.2$ m, as described in Chapter 3. Here the same simulations are analysed but, in addition, nine other passive scalars were introduced in the simulations to represent the breath of a potential infector. These were all introduced at once in the simulation for efficiency but will be described consecutively, which corresponds to a single infector being present at once, as considered in Chapter 2.

The focus of this thesis is far-field airborne transmission and not the close range. A choice was therefore made to introduce a passive tracer not at a point source, which would require careful validation to ensure appropriate dispersion to the far-field, but via a volume source so that the 'breath' of an occupant was already dispersed over the individual's fraction of the breathing zone within the room. Each 32 occupants of the classroom were assumed to occupy an equal cross-sectional area of the classroom. The scalars representative of potentially infected breath were introduced between the heights of 1.1 and 1.2 m over an area of $S/N = 1.7\,\mathrm{m}^2 \approx (5.5\times 10)/32$ m^2. These were distributed across the room at the nine locations shown in Figure 5.1 so that the influence of the potential infector location could be investigated. Their generation rate was set to the average CO_2 generation rate of 0.00335 L/s per person, as detailed in Chapter 3.

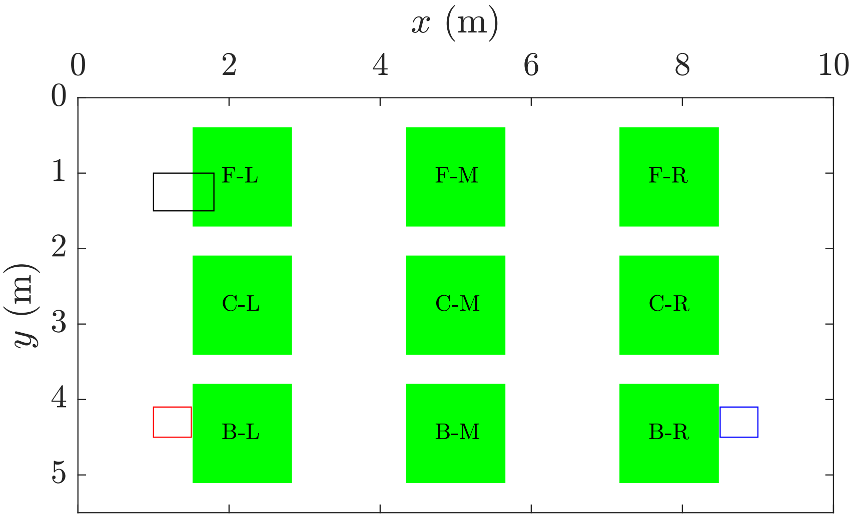

Figure 5.1: Classroom horizontal cross section. A scalar representing the infected breath was introduced at 9 different locations shown by the green areas between the heights of 1.1 and 1.2 m. The black rectangle (towards the top left) shows the horizontal position of the inlet. The blue rectangle (towards the bottom right) shows the position of the cross-flow high-level vent and the red rectangle (towards the bottom left) shows its position in the single-sided configuration.

5.2 Exposure to infected breath within the classroom

The distribution of potentially infected breath B was analysed for each potential infector location and for both the cross-flow and single-sided configurations. The tracer concentration B is defined as the scalar concentration, released at one of the nine locations (shown in Figure 5.1) with equivalent fluxes to that of CO_2 production within an individual's breath. This enables the fraction of the CO_2 concentration that would be measured at any given location, originating from a particular individual's breath, to be directly evaluated. This aims to determine how the concentration of B varies with source location, the position of the infector, ventilation configuration and resulting ventilation pattern, as described in Chapter 4. The results shown here correspond to the original vent set-up (OS) and the results for the other set-ups with varying vent sizes and heat inputs (as described in Table 4.4) are detailed in §5.5.

Figure 5.2 shows an exposure matrix which summarises the resulting normalised concentration in the areas shown in Figure 5.1 for different scalar release locations. The columns show the location of the infector (i.e. the scalar source) and the rows the location where the concentration and exposure is assessed. Here, the breathing zone only ($1 \leq z \leq 1.5\,\mathrm{m}$) is considered as the focus is on far-field transmission and this region is where occupants will be exposed to infected breath. Figure 5.2 shows the normalised concentration of the scalar $NB/(\langle C \rangle - C_a)$ where B is the mean concentration of the scalar in a specific exposure location, N is the number of occupants (here 32), $\langle C \rangle$ is the mean CO_2 in the breathing zone and C_a is the ambient CO_2 concentration.

The focus of this work is on far-field exposure, the near-field, defined as the diagonal row, is not considered here. High values of $NB/(\langle C \rangle - C_a)$, shown in the orange spectrum when superior to unity, indicate areas of the classroom where the scalar concentration is higher than the individual fraction, $(\langle C \rangle - C_a)/N$, estimated from the averaged CO_2 concentration in the breathing zone which would be expected if the room was well-mixed. The highest values of normalised concentration are located close to the diagonal, demonstrating that high scalar concentrations were found in proximity to the corresponding scalar source. This is detailed further below in Figure 5.3.

In the cross-flow configuration, shown in Figure 5.2a, the highest normalised concentrations are found at exposure locations (shown by the rows of the matrix) F-L and C-L. This corresponds to the area of the classroom where the flow is stagnant described in §4.1 and shown in Figure 4.1a. The lowest normalised concentrations (shown in the purple spectrum) are calculated at exposure locations F-R, C-R and B-R, where $NB/(\langle C \rangle - C_a)$ is below unity regardless of the scalar release

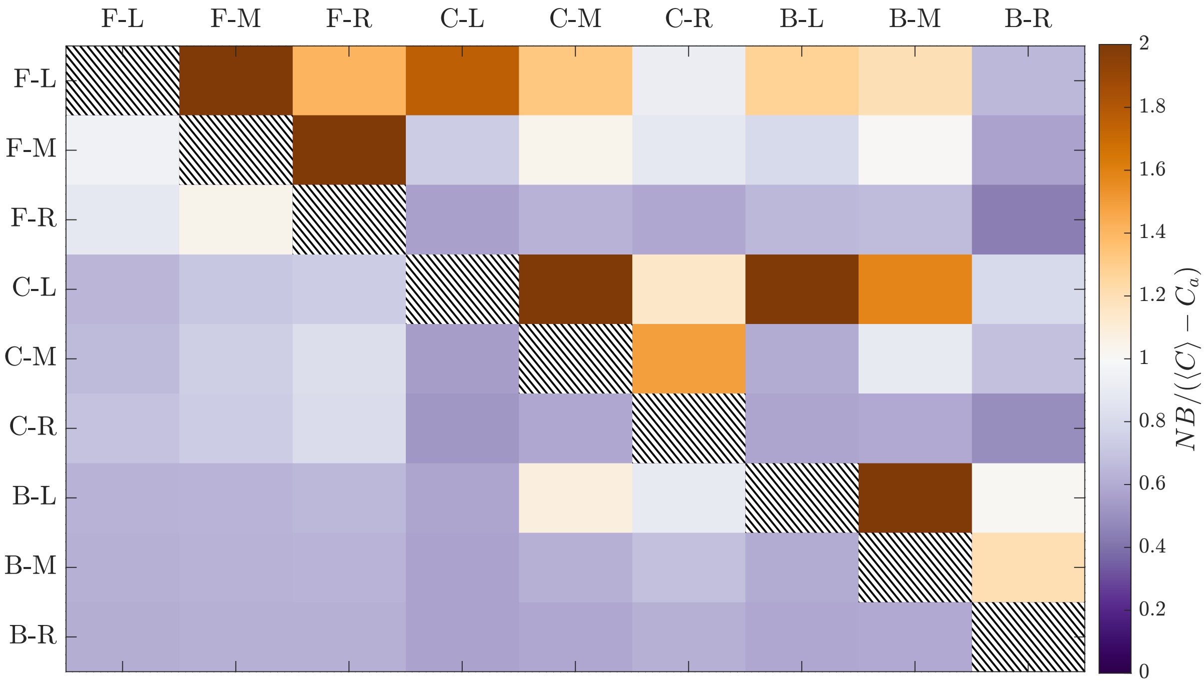

(a) Cross-flow

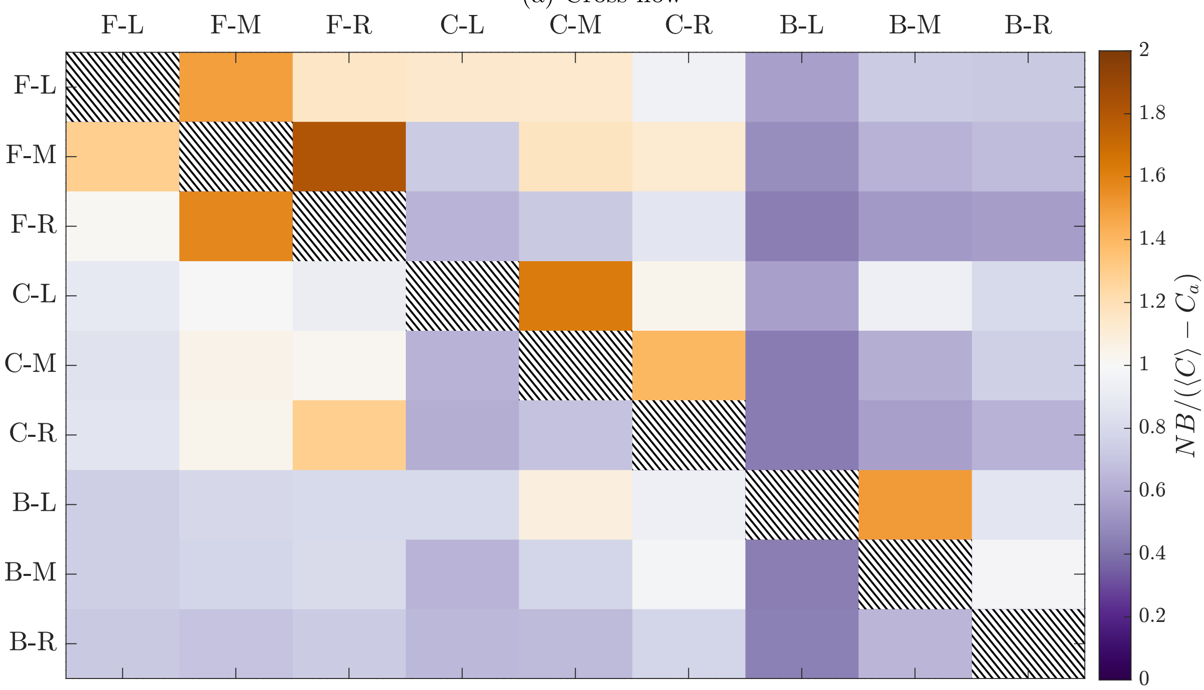

(b) Single-sided

Figure 5.2: Normalised scalar concentration $NB/(\langle C \rangle - C_a)$ in the zones shown in Figure 5.1 for: a) the cross-flow configuration, and b) the single-sided configuration. The columns show the source of the scalar (the potential infector location) and the rows the location where the concentration is measured (the exposure location). B is the mean concentration of the scalar in a specific exposure location, $\langle C \rangle$ is the mean CO_2 in the breathing zone, C_a is the ambient CO_2 concentration and N is the number of occupants. The hatched diagonal represents the near-field exposure and is not considered.

location. These exposure locations are located on the right hand side of the classroom near the region of high vertical velocity (c.f. Figure 4.3a). If the release is located in areas F-L, C-L and B-R (shown by the columns of the matrix) the spread of the scalar in the classroom is limited, most exposure locations showing a normalised concentration below unity. F-L and C-L are both situated within the stagnating zone which might not encourage the spread of a scalar originating from a location. B-R on the other hand corresponds to the release zone closest to the outlet of the cross-flow configuration.

Overall in the single-sided configuration, shown in Figure 5.2b, the values of normalised concentration are closer to unity, the number of areas in lighter shades of blue and orange is larger than in the cross-flow configuration. In Chapter 4, the single-sided configuration was shown to lead to a lower spread in the CO_2 concentration, and this seems to also apply to the distribution of scalar B originating from various locations.

In Figure 5.2b, high normalised concentrations are also found near the diagonal as well as at exposure locations F-L and F-M. Low values of $NB/(\langle C\rangle - C_a)$ are calculated at location B-M and B-R across several release locations. The exposure locations highlighted are different to those observed for the cross-flow configuration and show the impact of the different ventilation patterns on the distribution of scalar in the breathing zone. A scalar released at location C-M is found in high normalised concentration in the largest number of locations, which might be promoted by the source being in the centre of the classroom. On the contrary, the normalised concentration calculated for B-L as a release concentration was below unity in all exposure areas considered, which could indicate that this scalar did not spread significantly around the room which might be expected as, in the single-sided configuration, B-L is the zone closest to the outlet.

Overall for both configurations, the normalised concentration is highest closest to the location of the infector, with high normalised concentrations often found near the diagonal of Figure 5.2. This was investigated further by plotting the average normalised concentration with respect to the distance from the infector location. This was was averaged for all potential infector location considered and is shown in Figure 5.3 for both ventilation configurations. This does not show distances contained within the source, taken to be representative of the near-field and therefore not studied here. Figure 5.3 demonstrates that, for both ventilation configurations, the scalar concentration depends on the distance from its source, the infector location. As highlighted in Figure 5.2, however, it also depends on the ventilating flow within the room and it is critical to understand how this might affect far-field exposure at different points of the classroom.

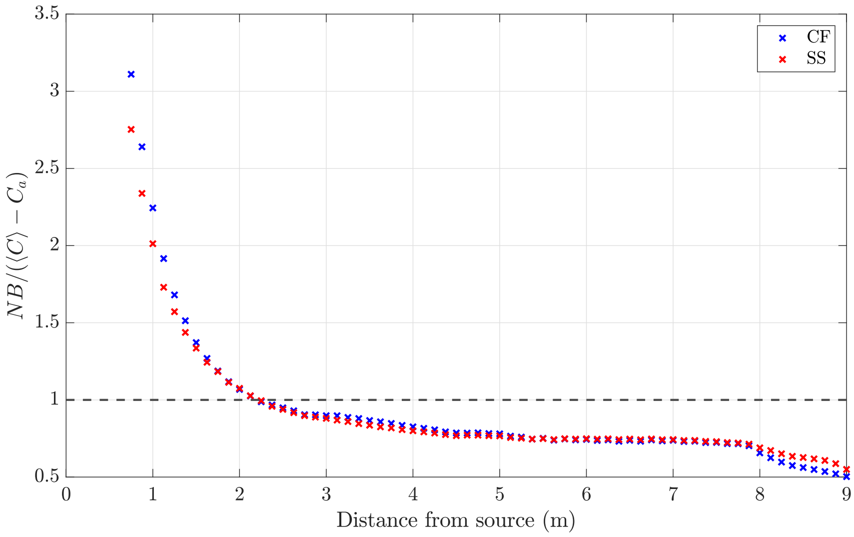

Figure 5.3: Normalised scalar concentration $NB/(\langle C\rangle - C_a)$ with respect to the distance from the scalar source for the cross-flow and single-sided configurations. B is the scalar concentration, N the number of occupants, C_a the ambient CO_2 and $\langle C\rangle$ is the mean CO_2 in the breathing zone. This is averaged for all 9 release locations. The horizontal dashed line shows $NB/(\langle C\rangle - C_a) = 1$.

5.3 Comparing actual and proxy far-field exposure

In addition to studying the changes in the distribution of the scalar B, within the breathing zone of the classroom as the release position varies, this chapter aims to address the extent to which CO_2 can used as a proxy to determine far-field exposure to infected breath. That is, whether CO_2 can represent the exposure to exhaled breath B regardless of the location of its source, the infector, within the bounds of the uncertainties associated with other factors in predicting exposures.

In Chapter 2, the probability of infection when the space is occupied was calculated from (2.2), reproduced here:

$$P_A = 1 - \exp\left(-\int_0^{T_A} \frac{I}{n}\, f\, q\, \mathrm{d}t\right),$$

with the rebreathed fraction calculated from $f = (C - C_a)/C_p$, where C is the monitored CO_2 within the space, C_a indicates the outdoor ambient level and C_p is the concentration of CO_2 added through exhaled breath. This in effect means that the proxy exposure E_p can be calculated by

$$E_p = \frac{I(C - C_a)}{N\, C_p}\,, \tag{5.1}$$

where in the case considered herein, the number of infector I is taken to be one and the classroom is assumed to be fully occupied with n taken to be the total number of occupants $N = 32$.

In this simulation and depending on the infector location considered, the actual exposure can be found from

$$E_a = \frac{B}{C_p}\,, \tag{5.2}$$

with B the concentration of the chosen scalar in the scenario studied. The ratio of both, the actual to proxy exposure ratio

$$R = \frac{E_a}{E_p} = \frac{B\,N}{I(C - C_a)}\,, \tag{5.3}$$

gives a measure of whether CO_2 is a suitable proxy for the exposure to B. In the context of the CFD simulations shown here, the proxy exposure calculated from CO_2 can be measured in several ways since the entire three dimensional field is available. Two options are discussed here:

(i) The proxy exposure was calculated from the room averaged CO_2 concentration $\overline{C}$ with

$$R_c = \frac{B\,N}{I(\overline{C} - C_a)}\,. \tag{5.4}$$

(ii) The proxy exposure was calculated from the local CO_2 concentration C at a point with

$$R_l = \frac{B\,N}{I(C - C_a)}\,. \tag{5.5}$$

If this ratio is above unity, the exposure is likely to be underestimated when using the CO_2 measurements and this might then lead to an underestimation of the risk of airborne infection predicted in Chapter 2. On the other hand, if the actual to proxy exposure ratio is below unity, the far-field exposure to infected breath will be overestimated by the use of a proxy CO_2 concentration. Although underestimating the risk of infection is the principal worry, due to its potential negative heath impacts, overestimating the risk should not be discounted. Indeed, overestimating the risk might lead to the implementation of stringent countermeasures that might not be necessary and negatively impact the classroom's occupants.

In §5.3.1, the room averaged CO_2 concentration is used to determine the proxy far-field exposure to calculate R_c and in §5.3.2, the local concentration of CO_2 at a point is used to determine R_l. Here, the cross-flow and single-sided configurations are compared for the original vent set-up (OS), the robustness of the results to changes in the ventilating flow rate is explored in §5.5.

5.3.1 Room averaged CO_2 as proxy

The proxy exposure was first calculated using the room averaged CO_2 concentration $\overline{C}$, calculated from the CFD simulations. This is useful as it allows the study of the spatial variations of the scalar B only. However, to be applicable in reality, we are typically required to assume that the point measurements of CO_2 from sensors are capable of capturing the room averaged concentration appropriately; Chapter 4 has shown how challenging it might be to obtain highly accurate estimates of the room average based on multiple point measurements. As such in §5.3.2, we investigate the scale of uncertainties introduced when using a single point measurement of CO_2 so that errors within these (typical) approaches might be better understood.

The horizontal cross section of the actual to proxy exposure ratio R_c is shown in Figure 5.4 for the cross-flow configuration. This shows how R_c varies in the horizontal cross-section for different release locations of potentially infected breath, the location of the source in each pane is shown by the white square. The highest values of the actual to proxy exposure ratio R_c are typically found near the scalar source, which shows that when estimated from CO_2 measurements, the exposure is likely to be underestimated in these regions. This is to be expected as this is also where the scalar concentration was the highest in Figure 5.2. The distribution in the rest of the cross section is affected by the ventilation pattern and the results will vary depending on the source location. For instance, if the scalar originates from C-L (central row, left column), the spread is limited and exposure is likely to be underestimated only near the source. In the rest of the cross section, almost three quarters of the classroom, the use of CO_2 will lead to an overestimation of exposure (with an actual to proxy exposure ratio R_c below unity). In the case of a source located at C-M however, where B is input in the centre of the classroom (central row, middle column), the region where exposure is underestimated is larger, showing that the scalar B originating from that location has spread more significantly which might lead to underestimation over a larger proportion of the breathing zone. This is also the case if the scalar is introduced at B-L (bottom row, left column) where R_c reaches high values, equal to or above three for almost a quarter of the room, which could lead to significant underestimates of the exposure and resulting infection risk.

A similar figure is shown for the single-sided configuration in Appendix B.1.1 (Figure B.1). The statistics for both configurations and all nine potential infector locations are shown in Figure 5.5. Overall a larger spread for the values of the actual to proxy exposure ratio R_c is found in the cross-flow configuration, specifically for scalars introduced at locations C-M, C-R, B-L and B-M, which was also evident from the contours shown in Figure 5.4. In both configurations,

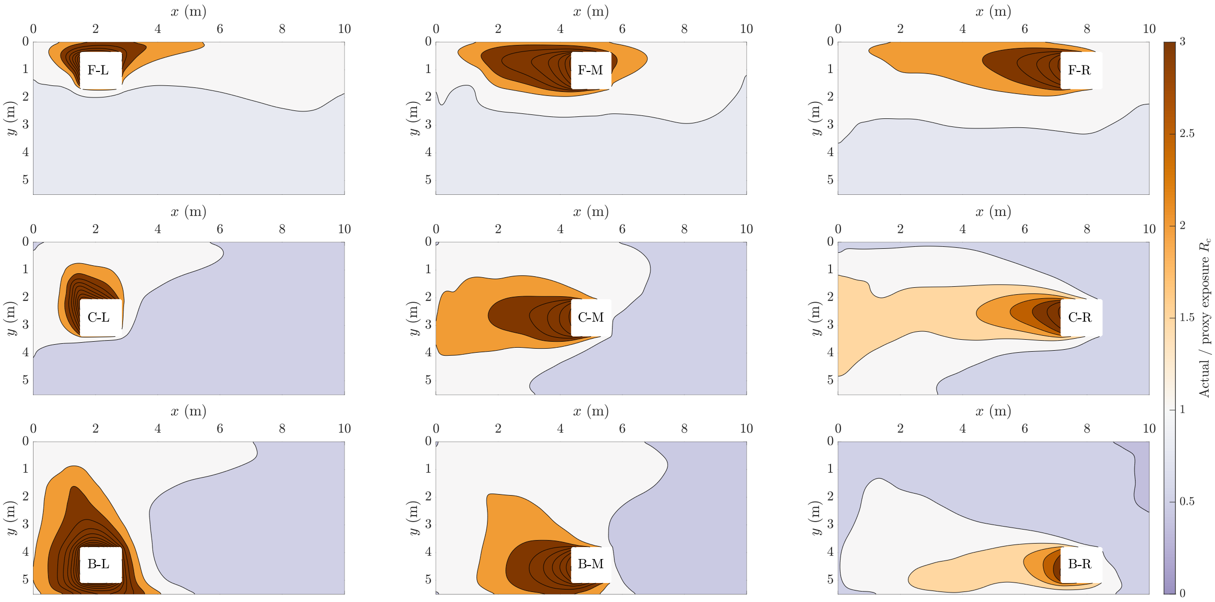

Figure 5.4: Horizontal cross section of the actual to proxy far-field exposure ratio R_c for the cross-flow configuration in the breathing zone shown for all nine infector locations. The position of the infector is shown by the white square. The proxy exposure is calculated from the room averaged CO_2 concentration $\overline{C}$. Colours in the purple spectrum indicate potential for overestimation of infection risk, colours in the orange spectrum indicate potential for underestimation of infection risk.

one of the smallest spreads in the ratio calculated is found when B originates from C-L. For this release location, Figure 5.5 also shows that exposure is likely to be overestimated: in almost 75% of the breathing zone the ratio R_c is below unity. This corresponds to the case highlighted in Figure 5.4 where the scalar was shown not to spread significantly from its source.

Notably, a significant difference is found between the cross-flow and single-sided configurations when the scalar is introduced at location B-L. In the single-sided configuration, unlike the cross-flow, the actual to proxy exposure ratio calculated is below unity in the entire breathing zone, with a median value close to 0.5, which shows that the use of CO_2 will overestimate the far-field exposure. In the single-sided configuration, the horizontal position of the classroom outlet is very close to B-L and the difference in the actual to proxy exposure ratio R_c found in this case highlights the impact of the top vent position on the flow field and resulting scalar distribution. This difference is unusual, for instance if scalars are introduced at the front (F-L, F-M, F-R) or at C-L the distributions between the cross-flow and single-sided configurations are relatively similar, with comparable medians and spreads.

For all cases shown in Figure 5.5, the interquartile range was comprised between 0.5 and 2 and the proxy exposure calculated from the room averaged CO_2 was within a factor of two of the actual far-field exposure in the majority of the breathing zone, despite variations due to the

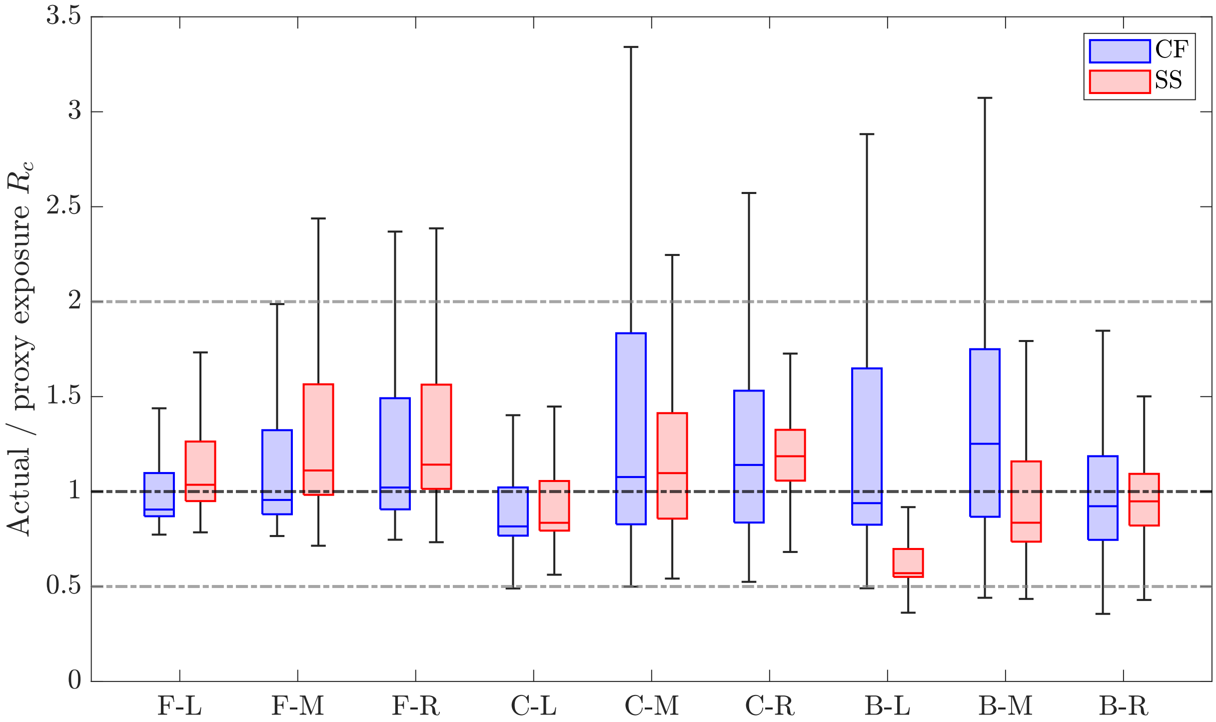

Figure 5.5: Actual to proxy exposure ratio R_c distribution for all nine infector locations (horizontal axis) for both the cross-flow (blue) and single-sided (red) configurations. The proxy exposure is determined from the room averaged CO_2 concentration $\overline{C}$. The bottom and top of the boxes show the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1$^{\text{st}}$ and 3$^{\text{rd}}$ quartiles. Horizontal dashed lines show a ratio equal to 0.5, 1 and 2 respectively.

source location and chosen ventilation configuration. This analysis however assumed that an accurate estimate of the room averaged CO_2 could be obtained, the impact of the CO_2 variations within the breathing zone on the calculation of a proxy exposure are discussed in the following section.

5.3.2 Local CO_2 as proxy

In Chapter 4, the CO_2 concentration was shown to vary significantly in the breathing zone. However, if these CO_2 measurements made within the breathing zone are to be used to inform the risk of airborne infection (as in Chapter 2), it is pertinent to determine how they compare to the potentially infected breath at that point. This is examined by calculating the proxy exposure from the local CO_2 concentration at that point and comparing it to the actual exposure calculated from B. The resulting far-field actual to proxy exposure ratio R_l accounts for both the variations of the CO_2 scalar and the infected breath scalar, and by examining the ratio R_l within the breathing zone we directly assess the variation in the potential to infer the exposure to infected breath that might arise via a CO_2 sensor placed at a given location within the breathing zone.

As previously, the ratio of the actual to proxy exposure R_l for the cross-flow configuration is shown in Figure 5.6. For reference, the contour of the CO_2 concentration in the breathing zone was shown in Figure 4.5 for both ventilation configurations. Overall, this shows similar qualitative trends to those obtained for R_c when using the room averaged CO_2 concentration $\overline{C}$ to determine the proxy exposure. The areas where underestimation is likely (with a ratio superior to unity) are scaled down and closer to the source of the scalar, compared to what was shown in §5.3.1. This might be because in regions where the flow is stagnating, both the scalar and CO_2 are likely to accumulate, which improves the predictions of the proxy. Overall, exposure is more likely to be overestimated when using local CO_2 compared to using room averaged CO_2 as a proxy, and this is exacerbated in regions of high CO_2 concentrations. This is visible for instance in the contours shown for a scalar input at B-R (bottom row, right column), where high levels of CO_2 at the top of the classroom lead to a larger overestimation of exposure in this region.

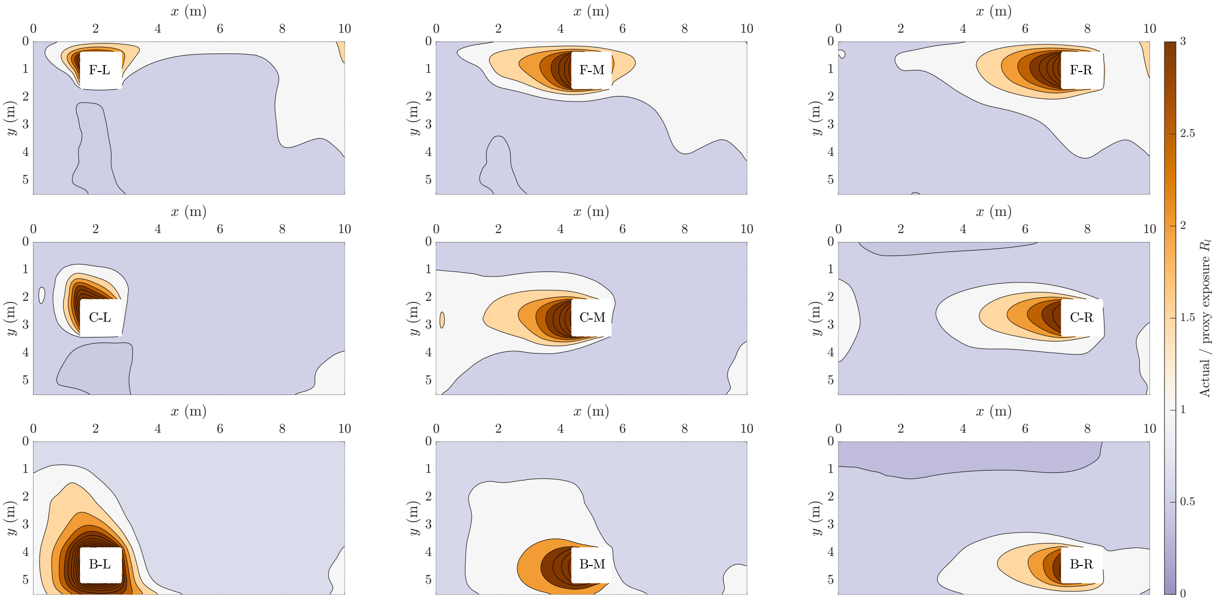

Figure 5.6: Horizontal cross section of the actual to proxy far-field exposure ratio R_l for the cross-flow configuration in the breathing zone shown for all nine infector locations. The position of the infector is shown by the white square. The proxy exposure is calculated from the local concentration of CO_2 C at a point. Colours in the purple spectrum indicate potential for overestimation of infection risk, colours in the orange spectrum indicate potential for underestimation of infection risk.

The actual to proxy exposure ration R_l is also shown for the single-sided configuration in Appendix B.1.2 (Figure 5.6). The statistics are summarised for both configurations and shown in Figure 5.7. Comparing this figure, showing R_l, to the earlier example, i.e. Figure 5.5 (in which the room averaged CO_2 was used as a proxy to calculate R_c), the spread of the values calculated for the actual to proxy exposure ratio is significantly reduced in all cases — almost all distributions are entirely contained between ratios of 0.5 and 2, i.e. the actual to proxy exposure ratio R_l can, for both configurations, be considered accurate to within a factor of two. Both the infected breath and CO_2 are modelled as passive scalars and are therefore transported by the same mechanisms around the classroom. Regions where CO_2 accumulates are likely to also lead to an accumulation of scalar concentration, leading to a comparable increase in proxy and actual exposure and ratios R_l closer to unity. Actual exposure is likely to remain high near the source and this can lead to underestimations of the proxy exposure and conversely, local increases in CO_2 concentration might lead to overestimates, as described previously. Generally, similar trends to those described for Figure 5.5 for R_c are observed, but the value of the ratio R_l is lowered, all medians are now below unity and the interquartile range is comprised between 0.5 and 1.25.

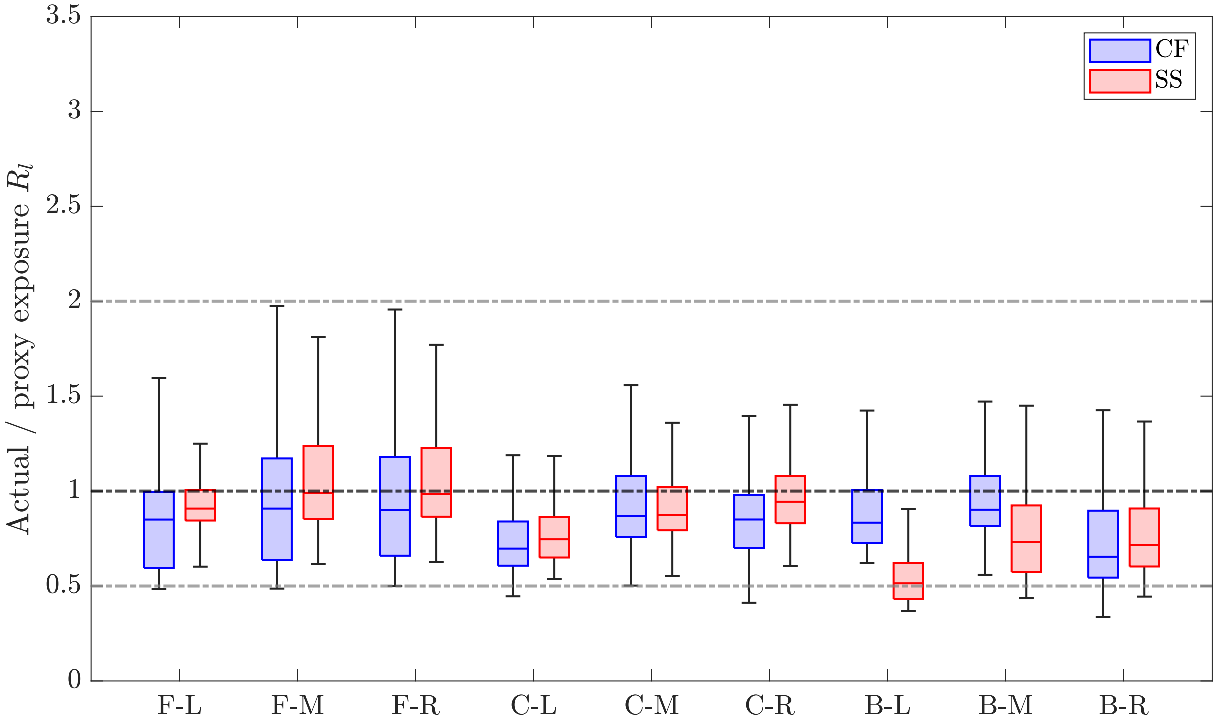

Figure 5.7: Actual to proxy exposure ratio R_l distribution for all nine infector locations for both the cross-flow (blue) and single-sided (red) configurations. The proxy exposure is found from the local CO_2 concentration C at a point. The bottom and top of the boxes show the 25^{th} and 75^{th} percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1^{st} and 3^{rd} quartiles. Horizontal dashed lines show a ratio equal to 0.5, 1 and 2 respectively.

In summary, the use of a local CO_2 measurement to estimate potential exposures to infected breath was found to be less likely to underestimate the far-field exposure to infected breath than when a room averaged CO_2 concentration was used (c.f §5.3.1). Across the breathing zone the actual to proxy exposure ratio, R_l, was found to be within the range 0.5 to 2, indicating that for almost all infector locations infection exposure estimates based on local point CO_2 measurements would lie within a factor of two of the actual exposure risk for those points — this held true for both ventilation configurations examined. This method was however more likely to lead to overestimates in exposure, where depending on the infector location, higher CO_2 concentrations might not have been indicative of a higher exposure to infected breath; this may, or may not, be significant.

5.4 Sensor placement

The use of CFD simulations was also exploited in this context to determine where CO_2 sensors should be placed in the classroom to best assess far-field exposure. A similar analysis was performed in Chapter 4, but in that chapter, the focus was to achieve reliable CO_2 measurements for estimating the mean CO_2 concentration in the classroom, here the focus is subtly distinct. As previously, the ratio of actual to proxy exposure, R_l, was calculated. All nine potential scalar sources were considered to estimate the variations in the exposure calculated in different regions of the classroom. The nine locations shown in Figure 5.1 were considered, disregarding the exposure resulting from the scalar originating from that location (i.e. when assessing exposure in the F-L area, the scalar input at F-L was not taken into account since this would constitute near-field, not far-field, exposure). For practical reasons, measurements taken at the walls were also analysed, following the method described in Chapter 4 and considering concentrations calculated within 0.2 m of a classroom wall. The actual to proxy exposure ratio R_l was calculated using the local CO_2 concentration C so that the performance of a sensor at a given point could be directly assessed. In this section, only the breathing zone was examined and the results are described for original vent set-up (OS) for both the cross-flow and single-sided configurations.

Figure 5.8 gives an estimate of the actual to proxy exposure ratio R_l (using local CO_2 to assess proxy exposure) for different measurement regions in the classroom for all potential scalar release locations. For all cases, the median value was between 0.6 and 0.9 which is indicative of a slight overprediction of the exposure to infected breath. The spread in the results was lowest when exposure was assessed in the B-R region. This is located closest to the high vertical velocity region near the far wall that was described in §4.1, and where both the concentration of CO_2 and scalar are likely to transported upwards by this current, leading to similar exposures and a

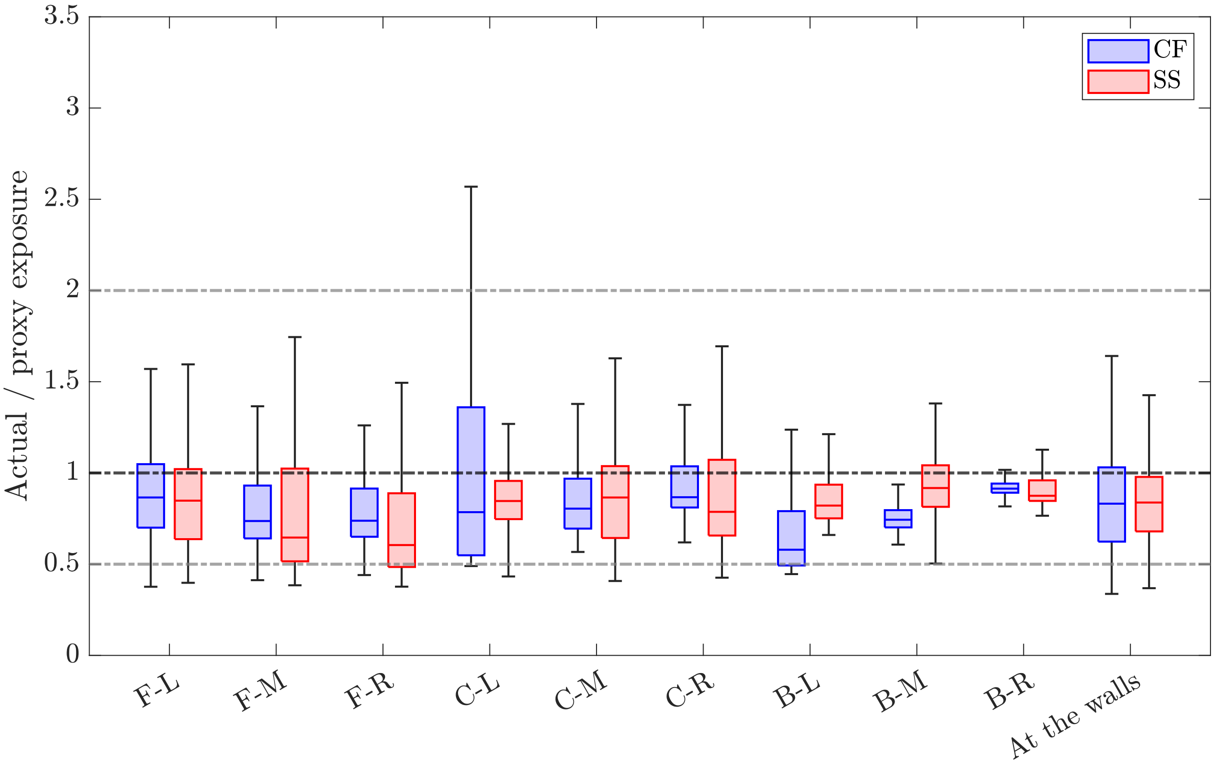

Figure 5.8: Actual to proxy exposure ratio R_l distribution at 10 measurement locations in the breathing zone for the cross-flow (blue) and single-sided (red) configurations. The proxy exposure is found from the local CO_2 concentration C at a point. The bottom and top of the boxes show the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1$^{\text{st}}$ and 3$^{\text{rd}}$ quartiles. Horizontal dashed lines show a ratio equal to 0.5, 1 and 2 respectively.

ratio closer to unity as shown in Figure 5.8. Here, the differences between the single-sided and cross-flow configurations are most significant when the assessment of exposure is made in the C-L region. In the cross-flow configuration, this corresponds to a region of stagnating flow where high concentrations can be expected and highlights that this can be challenging to estimate with a proxy, as shown with the large spread in the actual to proxy exposure ratios R_l calculated. This region is not present in the single-sided configuration where recirculation around the room is stronger (c.f. Figure 4.1b), which results in a smaller spread in the values for R_l.

The impact of measuring CO_2 at the walls to assess far-field exposure to infected breath was investigated in a similar manner. The distributions shown in Figure 5.8 give an estimate of the actual to proxy exposure ratio R_l if the exposure is assessed within 0.2 m of a wall, using local CO_2 as a proxy. In both ventilation configurations, the median of the distribution is close to 0.8 which suggests that exposure is likely to be slightly overestimated via assessment of CO_2 concentrations near the walls. This is not unexpected, the different scalars have been introduced some distance away from the walls and the resulting exposure was shown to depend on the distance from the scalar source. The overall accuracy of the wall measurements is however within the range observed at other locations in the breathing zone.

The accuracy of using local CO_2 measurements as a proxy for potentially infected breath varied across the breathing zone. In regions where air is likely to stagnate, the greater dependencies of the ratio of actual to proxy exposure on the location of the infected breath source were observed. Despite this, for measurement locations throughout the classroom the proxy exposure calculated based on the local CO_2 concentration was within a factor of two from the actual exposure at that point for all but the most extreme locations.

5.5 Sensitivity of far-field exposure to changes in the flow rate

The robustness of the previous results to changes in the ventilating flow rate was investigated using various vent sizes and heat inputs, introduced in §4.2 and summarised in Table 4.4. The distribution of actual to proxy exposure ratio R_c for all potential scalar release locations sources and for the different vent sizes and heat inputs, is shown in Figure 5.9 (the distributions of R_c and R_l are shown for each individual release location in Appendix B.2). In this case, the proxy exposure was determined from the room averaged CO_2 concentration $\overline{C}$ (as was the case in §5.3.1).

Overall, the distributions of R_c obtained with different vent sizes and heat inputs showed similar results to those obtained in the original vent set-up (OS), described in previous sections and summarised in Figure 5.9. Across the breathing zone for all potential release locations of scalars considered, the distributions of R_c have a median close to unity and interquartile ranges comprised between 0.8 and 1.5. Over all potential scalar source locations considered, the proxy exposure was within a factor of 1.5 of the actual exposure for 50% of the classroom and the differences between the cross-flow and single-sided configurations were small.

For the various vent sizes and heat inputs, the same trends described for the CO_2 distribution in Chapter 4 were also visible for the actual to proxy exposure ratio R_c. Notably a higher spread of this ratio was observed for the SO2 set-up which had also displayed large variations in the CO_2 distribution. On the other hand, the LO2 set-ups showed the smallest spread, with almost identical distributions for the cross-flow and single-sided configurations. This was investigated in §4.2 where the increase in flow rate in the LO2 set-up led to an increase in the pollutant removal efficiency of the cross-flow configuration, approaching what was observed for the single-sided configuration. CO_2 and infected breath were both modelled as passive scalars, higher variations

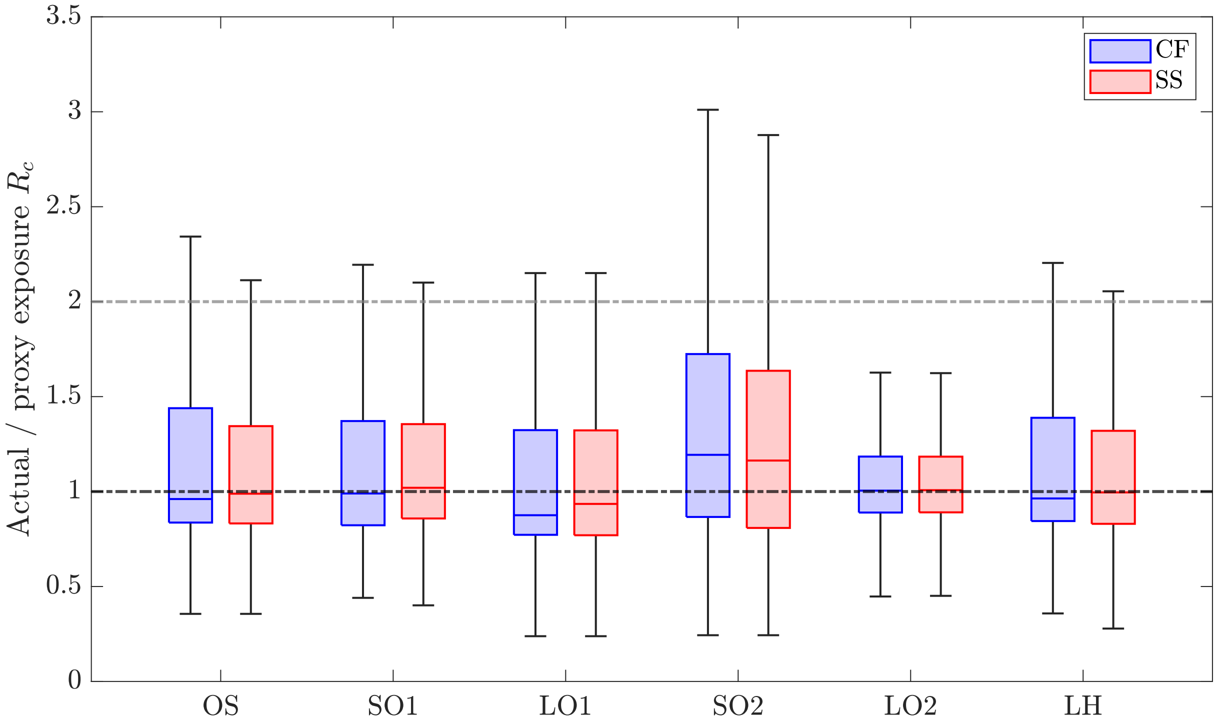

Figure 5.9: Actual to proxy exposure ratio R_c distribution, based on results from all nine infector locations, summarised for the different vent sizes and heat inputs tested in both the cross-flow (blue) and single-sided (red) configurations. The proxy exposure is determined from the room averaged CO_2 concentration $\overline{C}$. The bottom and top of the boxes show the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1$^{\text{st}}$ and 3$^{\text{rd}}$ quartiles. Horizontal dashed lines show a ratio equal to 0.5, 1 and 2 respectively.

in the distribution of CO_2 due to different ventilation patterns was likely to indicate a larger spread in scalar concentration and resulting exposure.

Despite the significant variations in flow pattern and CO_2 distributions described in Chapter 4, the different vent sizes and heat inputs considered showed similar accuracy in assessing exposure from room averaged CO_2 concentration than what was observed for the original vent set-up (OS). In the scenario with the highest exposure (SO2), the interquartile range doubled when compared to OS and in the case with the lowest exposure considered here (LO2) was halved. In all cases however, the proxy exposure was within a factor of two of the actual exposure in the majority of the breathing zone.

5.6 Conclusions

Like the CO_2 distribution described in Chapter 4, a scalar representing potentially infected breath originating from different locations displayed large spatial variations in the breathing zone and could not be assumed to be well-mixed within the classroom. Accounting for the far-field distribution only, and not taking into account the area where the scalar was introduced, showed that the absolute concentration of scalar (infected breath) was indeed dependent on the distance from the release location. Regions closest to a specific scalar source often showed higher concentrations than the rest of the domain. The overall distribution of the scalar did however depend on where it was introduced and the ventilation pattern in this region. Scalars introduced in stagnating areas for instance had a tendency to spread less than those originating in areas with significant ventilating flows.

The use of CO_2 as a proxy for far-field exposure was assessed by comparing CO_2 concentration to the actual exposure to potentially infected breath, B. The proxy exposure was calculated in two ways: 1) from the room averaged CO_2 concentration, to determine the actual to proxy exposure ratio R_c, and 2) from the local CO_2 concentration at a given point to assess the ratio R_l. The first method led to estimates of exposure to within a factor of two of the actual exposure in the majority of the classroom for the different sources considered. In practice this gives an estimate of the accuracy of CO_2 measurements as a proxy, provided a suitable estimate of the room average can be obtained. The effect of the CO_2 distribution within the classroom on this result was evaluated using the second method, where the local CO_2 concentration was used to determine proxy exposure and calculate R_l. This showed better agreement between actual and proxy exposure than the previous method with a reduction of 30% in the spread of the results. This was because high levels of CO_2 indicated regions where the flow was likely to stagnate and where scalars could also accumulate. This however led to an overestimation of the exposure in cases where increases in the CO_2 concentration were not correlated to increases in the specific scalar. Comparing different measurement locations, the proxy exposure measured at or near the walls, showed a similar accuracy to exposure determined elsewhere in the breathing zone, despite tending to an overestimation of the exposure.

These findings regarding the accuracy of CO_2 as a proxy for potentially infected breath were investigated for five other sets of simulations where both the rate, and the pattern, of the ventilating flow was varied by modifying the vent areas and/or the heat input. In these cases, the room averaged CO_2 concentration was used to estimate the proxy exposure so that the results could be easily related to the method used in Chapter 2 (where the exact location of the CO_2 sensor was not known but was assumed to be suitable to manage the ventilation and estimate the number of secondary infections). This analysis showed that, despite significant variations in the classroom ventilating flow, and both the resulting distributions of CO_2 and potentially infected breath, the proxy exposure obtained from CO_2 was still within a factor of two in the majority of the breathing zone for all scalar sources considered. Although significant, these variations need to be seen in the context of airborne infection risk modelling where another parameter, the quanta generation rate q is predicted to vary by at least three orders of magnitude for a given activity level and strain of the virus (Mikszewski et al., 2021). In situations where the location of the infector is unknown, as is typical in classrooms, CO_2 provides a useful tool with which to assess far-field exposure to infected breath with areas of high concentration often indicating low ventilating flow and the likely accumulation of infected breath — this exposure to infected breath being an important input into infection risk models.

Chapter 6

Conclusions

The provision of good Indoor Air Quality (IAQ) is critical for the health, well-being, and in some cases productivity, of building occupants, as an excessive exposure to contaminants indoors can cause serious chronic and acute health issues, or a degradation in cognitive performance. It remains, however, a complex challenge to address, with regulatory guidelines often missing (Morawska and Huang, 2022). The public's awareness of the health impact of IAQ was heightened during the COVID-19 pandemic since infections were found to primarily spread indoors, with airborne transmission playing a key role in badly ventilated spaces. Taking advantage of the increased public awareness, there is now a significant opportunity to build upon IAQ research findings and improve the indoor environments we all live in. During the pandemic, a primary location of concern for transmission was schools, and, in particular, classrooms (Corsi et al., 2021). In the UK, almost 11 million pupils and staff spend significant amounts of time in these spaces (roughly 6 hours a day, 5 days a week), and were therefore at risk of increased infections if mitigating measures were not introduced. COVID-19 infection is one of many issues related to IAQ in schools, with pollution in classrooms also shown to impact pupils' health and learning more widely (Salthammer et al., 2016; Fisk, 2017). This thesis set out to determine how exposure to contaminants and infection risk were likely to vary in naturally ventilated classrooms intended to be broadly typical of those found in the UK, with a focus on COVID-19 due to the pressing need for concrete IAQ research driven by the pandemic response. Although naturally ventilated spaces can be more challenging to investigate than mechanically ventilated rooms, they are prevalent, with around 90% of classrooms in the UK being naturally ventilated, and were therefore the focus of this thesis.

6.1 Summary of research and key findings

In this thesis, the behaviour of contaminants within naturally ventilated classrooms was investigated using:

(i) CO_2 measurements in 45 operational UK classrooms, and,

(ii) CFD simulations of a generic classroom in wintertime with low and high-level openings.

These two methods led to the following findings:

- The airborne infection risk predicted using CO_2 measurements exhibited seasonal variations and was considerably higher in the winter months than in summer.
- The calculated likelihood of infection also varied between classrooms of the same school despite similar ventilation systems.
- The ventilation pattern in the simulated classroom was dependent on the position of the top vent, with a single-sided configuration leading to a better contaminant removal efficiency in the scenario considered herein.
- The distribution of passive scalars representing either CO_2 or potentially infected breath varied considerably both in the vertical and horizontal directions and were not well-mixed within the classroom.
- Proxy far-field exposure calculated from CO_2 concentration was within a factor of two of the actual exposure to infected breath, whereas estimates for quanta generation rates used in airborne infection risk modelling vary by up to three orders of magnitude.

The following section further details the research conducted in each chapter.

6.2 Overview of results and future work

6.2.1 CO_2 monitoring to estimate airborne infection risk

In Chapter 2, CO_2 measurements were used to estimate the risk of airborne infection in 45 classrooms in the UK. The method was adapted from the works of Riley et al. (1978) and Rudnick and Milton (2003), typically used to study the transmission of airborne diseases such as COVID-19. Overall, and perhaps reassuringly, the absolute risk of far-field infection via the airborne route from one original infector was shown to be low. It must however be noted that the

CO_2 measurements were obtained in classrooms with a hybrid ventilation system which is likely to be more sophisticated than that of an average UK classroom ventilation. One of the main assumptions of the model was a quanta generation rate which was chosen for a relatively quiet activity and for the early strain of COVID-19. Changes in both (e.g. from loud classroom activity levels and to account for the more highly transmissible variants such as Omicron) are likely to lead to increases in the quanta generation rate and the resulting risk of transmission, as described for different values in Chapter 2. The lengthy datasets used, spanning 5 years, allowed for seasonal variations to be exhibited with higher risk expected in winter months. The increased risk in winter was due to higher levels of CO_2 in the winter period, and, assuming similar occupancy levels throughout the year, indicated lower ventilation rates. Significant variations were also visible between classrooms of the same school fitted with the same ventilation system; a result which highlights just how complex it is to assess the factors that can impact CO_2 concentration, notably the ventilation, occupancy, activity levels and occupants' behaviour. The presented risk analysis could be deployed more widely as more data become available and applied to assess the risk posed by other airborne diseases, beyond COVID-19.

In the future, CO_2 measurements could be extended by including a larger set of classrooms with varying ventilation systems to get a better representation of UK schools, including classrooms with no controlled ventilation provision. As the accuracy of the risk analysis of Chapter 2 relies upon the data, generating more extensive and richer datasets would lead to a significantly more powerful tool in which details on the occupants and their activities could be related to potential changes in the monitored CO_2. The risk analysis could also be enhanced by improving the evaluation of the quanta generation rate q. The model used in Chapter 2 can be easily extended to incorporate time varying q, provided the time dependence is known, or determine the impact of interventions on risk predictions, such as face masks. The methodology presented in Chapter 2 was based upon a number of assumptions, chiefly the quanta generation rate and the implicit assumption of a well-mixed space when calculating the number of secondary infections (which is discussed below). The resulting uncertainty can be addressed by considering the relative risk, as was done in the context of seasonal variation. Despite these assumptions, one of the conclusions of the chapter was that CO_2 measurements remain useful to calculate the absolute risk of infection provided the scenario is well defined and explained. This could allow the comparison of the risk resulting from different transmission routes, and thus identify spaces where airborne infection is most likely and needs to be addressed as a priority.

6.2.2 CFD to examine CO_2 distribution and exposure

In Chapter 3, a generic naturally ventilated UK classroom was characterised and modelled via Computational Fluid Dynamics (CFD) simulations. Since the focus of this thesis was on far-field exposure and its relationship to the natural ventilation design, a number of assumptions were made to construct these CFD simulations, including the simulation of buoyancy driven ventilation in winter, which was identified as the riskiest season in Chapter 2. The set-up, although simplified, was characterised to be representative of any ventilation strategies that promote low-level inflow and high-level outflow to a classroom. The occupants and heat provision to the room were modelled as a distributed heat source, and the input of CO_2 and potentially infected breath to the classroom was modelled with passive scalars. The use of a buoyancy-driven set-up proved challenging to simulate. In particular, obtaining a suitable inflow and outflow to the classroom with the use of boundary conditions directly at the classroom vents was not feasible and instead the exterior had to be modelled. The resulting simulations matched the bulk flow rate and room averaged temperature predicted by Gladstone and Woods (2001) using distributed heating and a well-mixed assumption.

The analysis of Chapter 4 showed that CO_2 can vary within a generic classroom both in the horizontal and vertical directions. Despite a uniform introduction of CO_2, the CFD simulations revealed that the classroom's environment could not be considered to be well-mixed. As such, these simulations suggest that the 'well-mixed' assumption of Chapter 2, which was used to calculate the number of secondary infections, may not hold for typical UK classrooms. The relative horizontal position of the vents on the floor and ceiling was shown to impact the flow, with a single-sided configuration leading to a different flow in the room than a cross-flow configuration. In the specific scenario considered here, the cross-flow configuration led to a stagnation region away from the outflow and higher CO_2 levels. The simulations also showed that the single-sided configuration had a higher pollutant removal efficiency. The largest variation of CO_2 was found in the breathing zone, where the CO_2 was introduced, which might increase the inaccuracies of measurements made at that height. CO_2 distribution was found to be more uniform higher up in the classroom, which meant that a suitable estimate of the room averaged CO_2 could be achieved by placing sensors near the walls. Changes in vent areas led to differing flow rates and dynamics, with the height of the incoming fountain affecting recirculation and mixing of CO_2.

In Chapter 5, an additional scalar was introduced to represent a source of potentially infected breath at different locations in the classroom and was compared to the overall CO_2 concentration to assess its use as a proxy for far-field exposure. Similarly to the CO_2 distribution, the scalar was also shown to not be well-mixed in the breathing zone of the classroom. The highest concentrations were often found near the source and the spread of the scalar across the classroom was affected by the location of its source. Actual exposure to potentially infected breath was compared to that predicted from CO_2, either from local measurements or the room-averaged concentration. Typically, using the local CO_2 concentration showed better agreement than the room averaged CO_2, with a correlation observed between regions of high CO_2 and stagnation regions where the scalar concentration could accumulate. This correlation could also be a result of both the CO_2 and scalars being modelled as passive tracers originating from the same height. However, even when calculating the proxy exposure from the room averaged CO_2, the proxy exposure was within a factor of two of the actual far-field exposure in most of the breathing zone for all potential scalar sources. This held across set-ups with different vent sizes and heat inputs, although a lower flow rate was found to lead to larger discrepancies between the actual and proxy exposure. Overall, the CO_2 concentration was not a perfect proxy, especially near the source but it did give good indications of the regions where flow was likely to stagnate. When combined, the results of this chapter suggest that CO_2 measurements can be a valuable tool for identifying regions with insufficient ventilation where the risk of infection might be higher. In the context of risk of infection modelling, the uncertainty introduced by using a well-mixed assumption quantified in this thesis has to be contrasted to the one introduced by the use of a quanta generation rate, which can vary by up to three orders of magnitude for COVID-19 (Mikszewski et al., 2021) and therefore represents a much higher uncertainty.

In the future, some of the assumptions made in the modelling of a generic classroom could be challenged to assess whether they impact significantly the results presented here and are therefore necessary to best represent a naturally ventilated classroom. These could be related to the simulation of: the heat input, the ventilation system, the occupants' breath and the classroom design. For instance, the results obtained here with a distributed heat source could be compared to those with a combination of horizontal and vertical heat patches to better represent a non uniform distribution of occupants in the room or heat arising from other sources. The low and high-level vents could also be moved to the classroom walls to assess how their vertical position, in addition to their horizontal position, would change the overall ventilating flow. The addition of wind could also be investigated, since it can play a strong role in driving naturally ventilating flows, and the overall heat input varied further to determine whether this leads to the typical short-circuiting expected of the single-sided configuration. The simulation of infected breath could also be improved by including a decay term to account for deposition. In order to approach what is observed in reality, it would be interesting to account for the presence of furniture within the classroom and assess whether it significantly alters the ventilation patterns described herein. Finally, future work might also look to quantify ventilation efficiency in several other ways (notably by studying the age of air) and compare those to the results presented herein using CO_2 measurements.

6.3 Role of CFD

Chapter 4 and Chapter 5 have shown that CFD simulations can be a useful tool to investigate problems that can be difficult to measure in reality, for example CO_2 or scalar distribution across a space and the well-mixed assumption, as considered here. Even with a large number of sensors, only a finite number of discrete points in space can be measured experimentally, from which it might be challenging to obtain a general view of the ventilating flow in the room. In particular, the use of RANS simulations offers the opportunity to investigate a large number of set-ups, which would not be feasible with other methods such as LES or DNS. Simulations of naturally ventilated flows are, however, quite challenging to set-up and solve, as described in Chapter 3, and often require a thorough knowledge of the physics involved. In particular, the definition of boundary conditions at the domain openings is of critical importance to the simulations' accuracy and, ideally, in the case of OpenFOAM used here, the exterior domain would not have to simulated so that the simulations remain as simple and efficient as possible. If the use of CFD is to become widespread for building simulations, further work is required to better understand the potential issues that might arise, of which examples specific to OpenFOAM were shown in this work, and this should include extensive validation and documentation.

6.4 CO_2 as a proxy of indoor air quality

With the development of cheap CO_2 sensors, vast amounts of data are likely to be generated by all buildings, including schools. This anticipated influx of CO_2 data for occupied buildings is expected to represent a great opportunity to gain a better understanding of our indoor environment and how it is ventilated. However, in order for these data to be useful and representative, the sensors have to be carefully calibrated and placed at sensible locations within the room. As the ideal locations within the room are likely to vary depending on the circumstances, clear guidelines need to be provided by ventilation professionals to support the wider public and effectively deploy these sensors in buildings.

Display CO_2 monitors will also be critical in providing the information needed to manage ventilation systems in real time. In addition to the location and calibration issues mentioned previously, guidelines are also required so that the sensors can be used appropriately. These guidelines should include a thorough overview of the benefits of using such monitors, including what actions are required at different CO_2 levels to improve ventilation, and also of their limitations, highlighting what can go wrong and encouraging a measured interpretation of the results. An improved understanding of how the use of these sensors can impact the occupants' behaviour and ventilation management is required, especially concerning whether any changes are sustainable over long periods of time.

Overall CO_2 is a useful indicator for IAQ especially if the pollutant of concern originates from indoors and from the building occupants, which is the case for airborne risk. However, it is reiterated that CO_2 sensors are not a perfect proxy for exposure to infected breath. Care must be taken in using CO_2 measurements as quantitative evaluations of risk, and in particular, care must also be taken in extrapolating some conclusions from one room to another since their effectiveness strongly depends upon the particular scenario being investigated. For instance, if air cleaners are to become widespread, CO_2 sensors will not be appropriate to determine the risk of airborne infection and different approaches will be required.

6.5 Indoor air quality in UK schools

When the contaminant source cannot be removed, ventilation is essential to manage the level of pollutants in classrooms. Increasing the ventilating flow rate might, however, lead to other issues, such as the introduction of outdoor pollutants or a decrease in thermal comfort. The use of display CO_2 monitors provides a solution to mitigate some (but not all) of these issues and allow better ventilation management. Despite the variations shown in the naturally ventilated classrooms simulated here, CO_2 sensors remain useful, in particular to identify spaces where ventilation might be lacking.

Historically, research has often focused on identifying the best ventilation system design to satisfy different objectives. In actual classrooms however, maintenance and operation cannot be disregarded. In addition to designing the ventilation of future schools, research should also aim to improve the ventilation currently found in schools and ensure that it can respond to new challenges as they arise, whether they are caused by a new pandemic or a changing climate.

This requires understanding the current state of IAQ in classrooms around the country and working with all stakeholders to understand what issues might arise so that improvements can be widespread and not limited to schools with the largest resources. Working with schools to improve their indoor environment should also be seen as an opportunity to educate pupils, and by extension their families and social circles, so that the importance of air quality for health can become as familiar as that of water quality.

6.6 Final remarks

In this thesis, COVID-19 was the primary focus since it represented an immediate threat to the health of pupils and staff in schools. The pandemic was terrible and it cannot be stressed enough how difficult it was for many people around the world, it did however present glimmers of hope that we can use this experience to improve our lives and prepare for future events, including a compelling opportunity to improve IAQ, notably in classrooms. The public awareness on the topic improved and it provided the motivation for excellent research, primarily COVID-19 specific, but that could be extended to IAQ as a whole, and this momentum should be harnessed. The research question that prompted the work shown in this thesis was to determine how exposure in naturally ventilated classrooms could be assessed. Two methods were used, one using CO_2 measurements in classrooms and another simulating a generic naturally ventilated classroom. Both proved to be useful although each had their own drawbacks and uncertainties. However, these drawbacks were somewhat complementary, with best performance achieved when both methods were used in combination and it is the author's opinion that such approaches, which combine detailed models and real world data, can be increasingly developed and deployed in the future. The conclusions of this thesis highlight the importance of a collaborative and interdisciplinary approach to IAQ research, both to challenge some of the common assumptions and to provide the public the tools needed to improve their own IAQ – making them healthier, happier people.

Bibliography

I. E. Abdalla, M. J. Cook, S. J. Rees, and Z. Yang. Large-eddy simulation of buoyancy-driven natural ventilation in an enclosure with a point heat source. *International Journal of Computational Fluid Dynamics*, 21(5-6):231–245, 2007.

I. E. Abdalla, M. J. Cook, and G. R. Hunt. Numerical study of thermal plume characteristics and entrainment in an enclosure with a point heat source. *Engineering Applications of Computational Fluid Mechanics*, 3(4):608–630, 2009.

D. C. Adam, P. Wu, J. Y. Wong, E. H. Y. Lau, T. K. Tsang, S. Cauchemez, G. M. Leung, and B. J. Cowling. Clustering and superspreading potential of SARS-CoV-2 infections in Hong Kong. *Nature Medicine*, 26(11):1714–1719, 2020.

Z. Ai, C. M. Mak, N. Gao, and J. Niu. How suitable is tracer gas as a surrogate of exhaled droplet nuclei for studying airborne transmission in the built environment. *16th Conference of the International Society of Indoor Air Quality and Climate: Creative and Smart Solutions for Better Built Environments, Indoor Air 2020*, pages 489–496, 2020.

ASHRAE. *Measuring air-change effectiveness.* American Society of Heating, Refrigerating and Air-Conditioning Engineers Standard 129-1997 (RA 2002), 2002. ISSN 1041-2336.

ASHRAE. *ASHRAE position document on indoor carbon dioxide.* American Society of Heating, Refrigerating and Air-Conditioning Engineers, 2022.

J. Axley. Multizone airflow modeling in buildings: history and theory. *HVAC&R Research*, 13 (6):907–928, 2007.

Z. Bakó-Biró, D. Clements-Croome, N. Kochhar, H. Awbi, and M. Williams. Ventilation rates in schools and pupils' performance. *Building and Environment*, 48(1):215–223, 2012.

R. M. Baloch, C. N. Maesano, J. Christoffersen, S. Banerjee, M. Gabriel, É. Csobod, E. de Oliveira Fernandes, I. Annesi-Maesano, et al. Indoor air pollution, physical and comfort parameters related to schoolchildren's health: Data from the European SINPHONIE study. *Science of The Total Environment*, 739:139870, 2020.

BESA. Key UK education statistics, British Educational Suppliers Association, 2021. https://www.besa.org.uk/key-uk-education-statistics/ [accessed 26/05/2022].

R. K. Bhagat, M. S. Davies Wykes, S. B. Dalziel, and P. F. Linden. Effects of ventilation on the indoor spread of COVID-19. *Journal of Fluid Mechanics*, 903:F1, 2020.

B. Blocken. LES over RANS in building simulation for outdoor and indoor applications: A foregone conclusion? *Building Simulation*, 11(5):821–870, 2018.

D. T. Bolster and P. F. Linden. Contaminants in ventilated filling boxes. *Journal of Fluid Mechanics*, 591:97–116, 2007.

D. T. Bolster and P. F. Linden. Particle transport in low-energy ventilation systems. Part 1: theory of steady states. *Indoor Air*, 19(2):122–129, 2009a.

D. T. Bolster and P. F. Linden. Particle transport in low-energy ventilation systems. Part 2: Transients and experiments. *Indoor Air*, 19(2):130–144, 2009b.

D. T. Bolster and D. M. Tartakovsky. Probabilistic risk analysis of building contamination. *Indoor Air*, 18(5):351–364, 2008.

British Standard Institution. *BS ISO 16814:2008: Building environment design – Indoor air quality – Methods of expressing the quality of indoor air for human occupancy.* BSI Standards Limited, 2008. ISBN 978 0 580 65961 4.

P. J. Bueno de Mesquita, C. J. Noakes, and D. K. Milton. Quantitative aerobiologic analysis of an influenza human challenge-transmission trial. *Indoor Air*, 30(6):1189–1198, 2020.

G. Buonanno, L. Morawska, and L. Stabile. Quantitative assessment of the risk of airborne transmission of SARS-CoV-2 infection: Prospective and retrospective applications. *Environment International*, 145(August):106112, 2020.

H. C. Burridge, R. K. Bhagat, M. E. J. Stettler, P. Kumar, I. De Mel, P. Demis, A. Hart, Y. Johnson-Llambias, et al. The ventilation of buildings and other mitigating measures for COVID-19: a focus on wintertime. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 477(2247):rspa.2020.0855, 2021a.

H. C. Burridge, S. Fan, R. L. Jones, C. J. Noakes, and P. F. Linden. Predictive and retrospective modelling of airborne infection risk using monitored carbon dioxide. *Indoor and Built Environment*, page 1420326X2110435, 2021b.

N. Canha, C. Mandin, O. Ramalho, G. Wyart, J. Ribéron, C. Dassonville, O. Hänninen, S. M. Almeida, and M. Derbez. Assessment of ventilation and indoor air pollutants in nursery and elementary schools in France. *Indoor Air*, 26(3):350–365, 2016.

G. Carrilho da Graça. A technical note on simplified modeling of turbulent mixing in wind-driven single sided ventilation. *Building and Environment*, 131(December 2017):12–15, 2018.

G. Carrilho da Graça and P. Linden. Ten questions about natural ventilation of non-domestic buildings. *Building and Environment*, 107:263–273, 2016.

C. C. Chang. *Computational modelling of buoyancy-driven displacement ventilation flows.* PhD thesis, University of Manchester, 2016.

P. M. Chapman. Determining when contamination is pollution — Weight of evidence determinations for sediments and effluents. *Environment International*, 33(4):492–501, 2007.

L. Chatzidiakou, D. Mumovic, and A. J. Summerfield. What do we know about indoor air quality in school classrooms? A critical review of the literature. *Intelligent Buildings International*, 4 (4):228–259, 2012.

L. Chatzidiakou, D. Mumovic, and A. Summerfield. Is CO_2 a good proxy for indoor air quality in classrooms? Part 1: The interrelationships between thermal conditions, CO_2 levels, ventilation rates and selected indoor pollutants. *Building Services Engineering Research and Technology*, 36(2):129–161, 2015a.

L. Chatzidiakou, D. Mumovic, and A. Summerfield. Is CO_2 a good proxy for indoor air quality in classrooms? Part 2: Health outcomes and perceived indoor air quality in relation to classroom exposure and building characteristics. *Building Services Engineering Research and Technology*, 36(2):162–181, 2015b.

T. Chenvidyakarn. *Buoyancy Effects on Natural Ventilation*, volume 9781107015. Cambridge University Press, Cambridge, 2013. ISBN 9781139058407.

D. Clements-Croome, H. Awbi, Z. Bakó-Biró, N. Kochhar, and M. Williams. Ventilation rates in schools. *Building and Environment*, 43(3):362–367, 2008.

S. Coldrick, A. Kelsey, M. J. Ivings, T. G. Foat, S. T. Parker, C. J. Noakes, A. Bennett, H. Rickard, and G. Moore. Modeling and experimental study of dispersion and deposition of respiratory emissions with implications for disease transmission. *Indoor Air*, 32(2):1–16, 2022.

M. J. Cook, Y. Ji, and G. R. Hunt. CFD modelling of natural ventilation: combined wind and buoyancy forces. *International Journal of Ventilation*, 1(3):169–179, 2003.

R. Corsi, S. Miller, M. Van Ry, L. Marr, L. Cadet, N. Pollock, D. Michaels, J. Jones, E.R., et al. Designing infectious disease resilience into school buildings through improvements to ventilation and air cleaning. *Report of the Lancet COVID-19 Commission Task Force on Safe Work, Safe School, and Safe Travel*, April, 2021.

M. Deevy. CFD modelling of the benchmark displacement ventilation test case HSL/2006/86. Technical report, UK Health & Safety Laboratory, 2006.

M. Deevy and N. Gobeau. CFD modelling of benchmark test cases for flow around a computer simulated person HSL/2006/51. Technical report, UK Health & Safety Laboratory, 2006.

Department for Education. Building Bulletin 103: Area guidelines for mainstream schools, 2014. https://www.gov.uk/government/publications/area-guidelines-and-net-capacity [accessed 23/09/2021].

Department for Education. Building Bulletin 101: Guidelines on ventilation, thermal comfort and indoor air quality in schools, 2018. https://www.gov.uk/government/publications/building-bulletin-101-ventilation-for-school-buildings [accessed 23/09/2021].

Department for Education. All schools to receive carbon dioxide monitors, 2021. https://www.gov.uk/government/news/all-schools-to-receive-carbon-dioxide-monitors [accessed 03/02/2022].

X. C. Dopico, M. Evangelou, R. C. Ferreira, H. Guo, M. L. Pekalski, D. J. Smyth, N. Cooper, O. S. Burren, et al. Widespread seasonal gene expression reveals annual differences in human immunity and physiology. *Nature Communications*, 6(1):7000, 2015.

F. Durrani, M. J. Cook, and J. J. McGuirk. Modelling buoyant thermal plumes in naturally ventilated buildings. *Proceedings of the First Building Simulation and Optimization Conference*, September:117–123, 2012.

F. Durrani, M. J. Cook, and J. J. McGuirk. Evaluation of LES and RANS CFD modelling of multiple steady states in natural ventilation. *Building and Environment*, 92:167–181, 2015.

P. Engzell, A. Frey, and M. D. Verhagen. Learning loss due to school closures during the COVID-19 pandemic. *Proceedings of the National Academy of Sciences*, 118(17):e2022376118, 2021.

European Commission, Joint Research Centre, Directorate-General for Health and Consumers, Institute for Health and Consumer Protection, M. Varró, A. Hyvärinen, P. Rudnai, et al. SINPHONIE : Schools indoor pollution & health observatory network in Europe : final report. Technical report, The Regional Environmental Centre for Central and Eastern Europe, 2014.

W. J. Fisk. The ventilation problem in schools: literature review. *Indoor Air*, 27(6):1039–1051, 2017.

S. D. Fitzgerald and A. W. Woods. Natural ventilation of a room with vents at multiple levels. *Building and Environment*, 39(5):505–521, 2004.

C. Gladstone and A. W. Woods. On buoyancy-driven natural ventilation of a room with a heated floor. *Journal of Fluid Mechanics*, 441:293–314, 2001.

U. Haverinen-Shaughnessy, D. J. Moschandreas, and R. J. Shaughnessy. Association between substandard classroom ventilation rates and students' academic achievement. *Indoor Air*, 21 (2):121–131, 2011.

X. He, E. H. Y. Lau, P. Wu, X. Deng, J. Wang, X. Hao, Y. C. Lau, J. Y. Wong, et al. Temporal dynamics in viral shedding and transmissibility of COVID-19. *Nature Medicine*, 26(5):672–675, 2020.

Health Effects Institute. State of global air 2020. Technical report, Health Effects Institute, Boston, MA, 2020.

P. Heiselberg, K. Svidt, and P. V. Nielsen. Characteristics of airflow from open windows. *Building and Environment*, 36(7):859–869, 2001.

T. D. Higton, H. C. Burridge, and G. O. Hughes. Natural ventilation flows established by a localised heat source in a room with a doorway and a high-level vent. *Building and Environment*, 203(June):108093, 2021.

G. Hunt and P. Linden. The fluid mechanics of natural ventilation—displacement ventilation by buoyancy-driven flows assisted by wind. *Building and Environment*, 34(6):707–720, 1999.

G. R. Hunt and N. B. Kaye. Pollutant flushing with natural displacement ventilation. *Building and Environment*, 41(9):1190–1197, 2006.

Y. Ji, M. J. Cook, and V. Hanby. CFD modelling of natural displacement ventilation in an enclosure connected to an atrium. *Building and Environment*, 42(3):1158–1172, 2007.

B. Jones, P. Sharpe, C. Iddon, E. A. Hathway, C. J. Noakes, and S. Fitzgerald. Modelling uncertainty in the relative risk of exposure to the SARS-CoV-2 virus by airborne aerosol transmission in well mixed indoor air. *Building and Environment*, 191(January):107617, 2021.

N. Kaye and G. Hunt. The effect of floor heat source area on the induced airflow in a room. *Building and Environment*, 45(4):839–847, 2010.

N. Kaye, Y. Ji, and M. J. Cook. Numerical simulation of transient flow development in a naturally ventilated room. *Building and Environment*, 44(5):889–897, 2009.

N. B. Kaye and G. R. Hunt. Time-dependent flows in an emptying filling box. *Journal of Fluid Mechanics*, 520:135–156, 2004.

W. Kierat, M. Bivolarova, E. Zavrl, Z. Popiolek, and A. Melikov. Accurate assessment of exposure using tracer gas measurements. *Building and Environment*, 131(October 2017):163–173, 2018.

N. E. Klepeis, W. C. Nelson, W. R. Ott, J. P. Robinson, A. M. Tsang, P. Switzer, J. V. Behar, S. C. Hern, and W. H. Engelmann. The National Human Activity Pattern Survey (NHAPS): a resource for assessing exposure to environmental pollutants. *Journal of Exposure Science & Environmental Epidemiology*, 11(3):231–252, 2001.

R. Kumar and A. Dewan. URANS computations with buoyancy corrected turbulence models for turbulent thermal plume. *International Journal of Heat and Mass Transfer*, 72:680–689, 2014.

P. S. Lakey, Y. Won, D. Shaw, F. F. Østerstrøm, J. Mattila, E. Reidy, B. Bottorff, C. Rosales, et al. Spatial and temporal scales of variability for indoor air constituents. *Communications Chemistry*, 4(1):1–7, 2021.

P. J. Landrigan, C. A. Kimmel, A. Correa, and B. Eskenazi. Children's health and the environment: public health issues and challenges for risk assessment. *Environmental Health Perspectives*, 112(2):257–265, 2004.

Y. Li, G. M. Leung, J. W. Tang, X. Yang, C. Y. H. Chao, J. Z. Lin, J. W. Lu, P. V. Nielsen, et al. Role of ventilation in airborne transmission of infectious agents in the built environment - a multidisciplinary systematic review. *Indoor Air*, 17(1):2–18, 2007.

P. F. Linden. The fluid mechanics of natural ventilation. *Annual Review of Fluid Mechanics*, 31 (1):201–238, 1999.

P. F. Linden, G. F. Lane-Serff, and D. A. Smeed. Emptying filling boxes: the fluid mechanics of natural ventilation. *Journal of Fluid Mechanics*, 212(-1):309, 1990.

S. Liu and A. Novoselac. Lagrangian particle modeling in the indoor environment: A comparison of RANS and LES turbulence methods (RP-1512). *HVAC&R Research*, 20(4):480–495, 2014.

S. D. Lowther, S. Dimitroulopoulou, K. Foxall, C. Shrubsole, E. Cheek, B. Gadeberg, and O. Sepai. Low level carbon dioxide indoors—A pollution indicator or a pollutant? A health-based perspective. *Environments*, 8(11):125, 2021.

J. Lu and Z. Yang. COVID-19 outbreak associated with air conditioning in restaurant, Guangzhou, China, 2020. *Emerging Infectious Diseases*, 26(11):2789–2791, 2020.

N. Mahyuddin and H. Awbi. The spatial distribution of carbon dioxide in an environmental test chamber. *Building and Environment*, 45(9):1993–2001, 2010.

N. Mahyuddin, H. B. Awbi, and M. Alshitawi. The spatial distribution of carbon dioxide in rooms with particular application to classrooms. *Indoor and Built Environment*, 23(3):433–448, 2014.

Mariam, A. Magar, M. Joshi, P. S. Rajagopal, A. Khan, M. M. Rao, and B. K. Sapra. CFD simulation of the airborne transmission of COVID-19 vectors emitted during respiratory mechanisms: Revisiting the concept of safe distance. *ACS Omega*, 6(26):16876–16889, 2021.

L. C. Marr, J. W. Tang, J. Van Mullekom, and S. S. Lakdawala. Mechanistic insights into the effect of humidity on airborne influenza virus survival, transmission and incidence. *Journal of The Royal Society Interface*, 16(150):20180298, 2019.

A. Melikov and J. Kaczmarczyk. Measurement and prediction of indoor air quality using a breathing thermal manikin. *Indoor Air*, 17(1):50–59, 2007.

F. R. Menter, M. Kuntz, and R. Langtry. Ten years of industrial experience with the sst turbulence model. *Turbulence, heat and mass transfer*, 4(1):625–632, 2003.

A. Mikszewski, L. Stabile, G. Buonanno, and L. Morawska. The airborne contagiousness of respiratory viruses: A comparative analysis and implications for mitigation. *Geoscience Frontiers*, page 101285, 2021.

S. L. Miller, W. W. Nazaroff, J. L. Jimenez, A. Boerstra, G. Buonanno, S. J. Dancer, J. Kurnitski, L. C. Marr, L. Morawska, and C. Noakes. Transmission of SARS-CoV-2 by inhalation of respiratory aerosol in the Skagit Valley Chorale superspreading event. *Indoor Air*, 31(2): 314–323, 2021.

Ministry of Housing, Communities & Local Government. Energy performance of buildings data: England and Wales, 2021. data retrieved from Energy Performance Certificates, https://epc.opendatacommunities.org/ [23/09/2021].

L. Morawska and W. Huang. WHO Health Guidelines for Indoor Air Quality and National Recommendations/Standards. In *Handbook of Indoor Air Quality*, pages 1–20. Springer Nature Singapore, Singapore, 2022. ISBN 9789811051555.

L. Morawska, A. Afshari, G. N. Bae, G. Buonanno, C. Y. H. Chao, O. Hänninen, W. Hofmann, C. Isaxon, et al. Indoor aerosols: from personal exposure to risk assessment. *Indoor Air*, 23 (6):462–487, 2013.

L. Morawska, J. W. Tang, W. Bahnfleth, P. M. Bluyssen, A. Boerstra, G. Buonanno, J. Cao, S. Dancer, et al. How can airborne transmission of COVID-19 indoors be minimised? *Environment International*, 142(April):105832, 2020.

L. Morawska, J. Allen, W. Bahnfleth, P. M. Bluyssen, A. Boerstra, G. Buonanno, J. Cao, S. J. Dancer, et al. A paradigm shift to combat indoor respiratory infection. *Science*, 372(6543): 689–691, 2021.

E. Mundt. Non-buoyant pollutant sources and particles in displacement ventilation. *Building and Environment*, 36(7):829–836, 2001.

M. Nicas. Estimating exposure intensity in an imperfectly mixed room. *American Industrial Hygiene Association Journal*, 57(6):542–550, 1996.

C. J. Noakes and P. A. Sleigh. Mathematical models for assessing the role of airflow on the risk of airborne infection in hospital wards. *Journal of The Royal Society Interface*, 6(suppl 6), 2009.

A. Novoselac and J. Srebric. Comparison of air exchange efficiency and contaminant removal effectiveness as IAQ indices. *ASHRAE Transactions*, 109 PART 2(4663):339–349, 2003.

Office of Research and Development's National Center for Environmental Assessments (NCEA). Exposure factors handbook: 2011 Edition (EPA/600/R-09/052F). Technical report, United States Environmental Protection Agency, 2011.

P. Ole Fanger. What is IAQ? *Indoor Air*, 16(5):328–334, 2006.

OpenCFD. `buoyancyTurbSource` finite-volume option, 2020. https://www.openfoam.com/news/main-news/openfoam-v20-12/solver-and-physics-solver-and-physics-buoyancy-fvoption [accessed 23/05/2022].

OpenCFD. ESI OpenCFD Release OpenFOAM®v2106, 2021. https://www.openfoam.com/news/main-news/openfoam-v2106 [accessed 23/05/2022].

OpenCFD. OpenFOAM®Documentation, 2022. https://www.openfoam.com/documentation/overview [accessed 05/10/2022].

J. S. Park, N.-Y. Jee, and J.-W. Jeong. Effects of types of ventilation system on indoor particle concentrations in residential buildings. *Indoor Air*, 24(6):629–638, 2014.

B. Pavilonis, A. M. Ierardi, L. Levine, F. Mirer, and E. A. Kelvin. Estimating aerosol transmission risk of SARS-CoV-2 in New York City public schools during reopening. *Environmental Research*, 195(January):110805, 2021.

G. Pei, M. Taylor, and D. Rim. Human exposure to respiratory aerosols in a ventilated room: Effects of ventilation condition, emission mode, and social distancing. *Sustainable Cities and Society*, 73(June):103090, 2021.

Z. Peng, A. P. Rojas, E. Kropff, W. Bahnfleth, G. Buonanno, S. Dancer, J. Kurnitski, Y. Li, et al. Practical indicators for risk of airborne transmission in shared indoor environments and their application to COVID-19 outbreaks. *Environmental Science & Technology*, 56(2): 1125–1137, 2022.

L. Pérez-Lombard, J. Ortiz, and C. Pout. A review on buildings energy consumption information. *Energy and Buildings*, 40(3):394–398, 2008.

A. Persily and L. de Jonge. Carbon dioxide generation rates for building occupants. *Indoor Air*, 27(5):868–879, 2017.

S. B. Pope. *Turbulent Flows*. Cambridge University Press, Cambridge, UK, 2000. ISBN 9780521591256.

F. Pourfattah, L.-P. Wang, W. Deng, Y.-F. Ma, L. Hu, and B. Yang. Challenges in simulating and modeling the airborne virus transmission: A state-of-the-art review. *Physics of Fluids*, 33 (10):101302, 2021.

H. Qian, T. Miao, L. Liu, X. Zheng, D. Luo, and Y. Li. Indoor transmission of SARS-CoV-2. *Indoor Air*, 31(3):639–645, 2021.

O. Ramalho, C. Mandin, J. Ribéron, and G. Wyart. Air stuffiness and air exchange rate in French schools and day-care centres. *International Journal of Ventilation*, 12(2):175–180, 2013.

E. C. Riley, G. Murphy, and R. L. Riley. Airborne spread of measles in a suburban elementary school. *American Journal of Epidemiology*, 107(5):421–432, 1978.

D. Rim and A. Novoselac. Occupational exposure to hazardous airborne pollutants: effects of air mixing and source location. *Journal of Occupational and Environmental Hygiene*, 7(12): 683–692, 2010.

R. Rossi and G. Iaccarino. Numerical analysis and modeling of plume meandering in passive scalar dispersion downstream of a wall-mounted cube. *International Journal of Heat and Fluid Flow*, 43:137–148, 2013.

S. N. Rudnick and D. K. Milton. Risk of indoor airborne infection transmission estimated from carbon dioxide concentration. *Indoor Air*, 13(3):237–245, 2003.

T. Salthammer, E. Uhde, T. Schripp, A. Schieweck, L. Morawska, M. Mazaheri, S. Clifford, C. He, et al. Children's well-being at schools: Impact of climatic conditions and air pollution. *Environment International*, 94:196–210, 2016.

M. Sandberg. What is ventilation efficiency? *Building and Environment*, 16(2):123–135, 1981.

J. Smagorinsky. General circulation experiments with the primitive equations. *Monthly Weather Review*, 91(3):99–164, 1963.

L. Stabile, A. Pacitto, A. Mikszewski, L. Morawska, and G. Buonanno. Ventilation procedures to minimize the airborne transmission of viruses in classrooms. *Building and Environment*, 202(March):108042, 2021.

W. Su, B. Yang, A. Melikov, C. Liang, Y. Lu, F. Wang, A. Li, Z. Lin, X. Li, G. Cao, and R. Kosonen. Infection probability under different air distribution patterns. *Building and Environment*, 207(October):108555, 2021.

G. N. Sze To and C. Y. H. Chao. Review and comparison between the Wells-Riley and dose-response approaches to risk assessment of infectious respiratory diseases. *Indoor Air*, 20(1): 2–16, 2010.

S. Tan, Z. Zhang, K. Maki, K. J. Fidkowski, and J. Capecelatro. Beyond well-mixed: A simple probabilistic model of airborne disease transmission in indoor spaces. *Indoor Air*, 32(3):1–12, 2022.

K. W. Tham. Indoor air quality and its effects on humans—A review of challenges and developments in the last 30 years. *Energy and Buildings*, 130:637–650, 2016.

Y. Tominaga. Flow around a high-rise building using steady and unsteady RANS CFD: Effect of large-scale fluctuations on the velocity statistics. *Journal of Wind Engineering and Industrial Aerodynamics*, 142:93–103, 2015.

Y. Tominaga and T. Stathopoulos. Steady and unsteady RANS simulations of pollutant dispersion around isolated cubical buildings: Effect of large-scale fluctuations on the concentration field. *Journal of Wind Engineering and Industrial Aerodynamics*, 165(November 2016):23–33, 2017.

UNESCO. Covid-19 school closures around the world will hit girls hardest, https://en.unesco.org/news/covid-19-school-closures-around-world-will-hit-girls-hardest, 2020.

D. Vernez, S. Schwarz, J. Sauvain, C. Petignat, and G. Suarez. Probable aerosol transmission of SARS-CoV-2 in a poorly ventilated courtroom. *Indoor Air*, April:ina.12866, 2021.

C. V. M. Vouriot, H. C. Burridge, C. J. Noakes, and P. F. Linden. Seasonal variation in airborne infection risk in schools due to changes in ventilation inferred from monitored carbon dioxide. *Indoor Air*, 31(4):1154–1163, 2021.

W. F. Wells. *Airborne contagion and air hygiene. An ecological study of droplet infections.* Cambridge, MA, Harvard University Press, 1955.

WHO. Q&A: Ventilation and air conditioning in public spaces and buildings and COVID-19, https://www.who.int/news-room/q-a-detail/q-a-ventilation-and-air-conditioning-in-public-spaces-and-buildings-and-covid-19, 2020.

R. Yang, C. S. Ng, K. L. Chong, R. Verzicco, and D. Lohse. Do increased flow rates in displacement ventilation always lead to better results? *Journal of Fluid Mechanics*, 932:A3, 2022.

X. Yang, K. Zhong, Y. Kang, and T. Tao. Numerical investigation on the airflow characteristics and thermal comfort in buoyancy-driven natural ventilation rooms. *Energy and Buildings*, 109: 255–266, 2015.

S. Zhang and Z. Lin. Dilution-based evaluation of airborne infection risk - Thorough expansion of Wells-Riley model. *Building and Environment*, 194(October 2020):107674, 2021.

Z. Zhang, W. Zhang, Z. J. Zhai, and Q. Y. Chen. Evaluation of various turbulence models in predicting airflow and turbulence in enclosed environments by CFD: Part 2—comparison with experimental data from literature. *HVAC&R Research*, 13(6):871–886, 2007.

H.-Y. Zhong, Y. Sun, J. Shang, F.-P. Qian, F.-Y. Zhao, H. Kikumoto, C. Jimenez-Bescos, and X. Liu. Single-sided natural ventilation in buildings: a critical literature review. *Building and Environment*, 212(September 2021):108797, 2022.

K. Zhong, X. Yang, W. Feng, and Y. Kang. Pollutant dilution in displacement natural ventilation rooms with inner sources. *Building and Environment*, 56:108–117, 2012.

N. Ziauddeen, K. Woods-Townsend, S. Saxena, R. Gilbert, and N. A. Alwan. Schools and COVID-19: Reopening Pandora's box? *Public Health in Practice*, 1(July):100039, 2020.

Appendix A

OpenFOAM terminology

This appendix aims to detail the OpenFOAM terminology used in Chapter 3, further details on the code can be found online (OpenCFD, 2022).

A.1 Boundary conditions

- `externalWallHeatFluxTemperature`: this temperature boundary condition sets the heat flux normal to the boundary. In this thesis, a fixed power input (in W) is used.
- `fixedFluxPressure`: this pressure boundary condition sets a fixed gradient at the boundary to match the flux calculated from the velocity boundary condition. This needs to be used at boundaries where the velocity is set.
- `fixedValue`: this Dirichlet boundary condition sets the value at a boundary to one specified by the user.
- `inletOutlet`: this is an outflow boundary condition that applies a `zeroGradient` boundary condition if the flux is positive (flow out of the domain) and a `fixedValue` condition if the flux is negative (flow into the domain or return flow).
- `noSlip`: this boundary condition sets the velocity to zero at the boundary.
- `pressureInletOutletVelocity`: like the `inletOutlet` boundary condition, this is an outflow boundary condition for the velocity field. In the case of positive flux (flow out of the domain), the gradient of the velocity is set to zero (`zeroGradient`). When the flux is negative (flow into the domain or return flow) the velocity is calculated from the flux. This needs to be applied to boundaries where the pressure is set.

- `prghTotalPressure`: this pressure boundary condition sets the total pressure p_0 at a boundary.

- `zeroGradient`: this is a Neumann boundary condition that sets the gradient of a specific field to zero at the boundary.

A.2 Numerical discretisation

- Temporal discretisation (`ddtSchemes`): the `backwardEuler` scheme is used to discretise the time derivative such that:

$$\frac{\partial}{\partial t}\int_V \rho\phi dV = \frac{3(\rho_P\phi_P V)^n - 4(\rho_P\phi_P V)^o + (\rho_P\phi_P V)^{oo}}{2\Delta t}$$

  where ϕ is a given field of interest, P is the centre of a given cell, ρ is the density, V is the control volume and t is time. The time step being solved for is $\phi^n = \phi(t+\Delta t)$, the previous time step is $\phi^o = \phi(t)$ and $\phi^{oo} = \phi(t-\Delta t)$.

- Gradient discretisation (`gradSchemes`): `cellLimited Gauss linear 1` applies Gauss' theorem to discretise the cell gradient $\int_V \nabla\phi dV$. The interpolation scheme used is `linear` which corresponds to a central differencing scheme. For all fields apart from pressure, a `cellLimited` scheme is used with a coefficient of 1, ensuring that the face values are bounded by the neighbouring cells' maximum and minimum values. For `snGradSchemes`, the component of gradient normal to a cell face is found using the `orthogonal` scheme since the mesh is orthogonal.

- Divergence discretisation (`divSchemes`): `Gauss linearUpwind` applies Gauss' theorem to discretise $\nabla \cdot \phi$. For most fields, a `linearUpwind` interpolation scheme is used. For the passive scalars an `upwind` scheme is used.

- Laplacian schemes (`laplacian`): `Gauss linear orthogonal` is used for terms such as $\nabla^2\phi$, which are discretised using Gauss' theorem and interpolated using a central differencing scheme (`linear`). `orthogonal` refers to the interpolation scheme used for the surface normal gradient.

- Interpolation scheme (`interpolationSchemes`): a `linear` scheme (central differencing) is used to interpolate values from cell centres to face centres.

A.3 Others

- `blockMesh`: utility responsible for mesh generation.
- `buoyancyTurbSource`: option that takes into account the effects of buoyancy on turbulence production.
- `buoyantPimpleFoam`: transient OpenFOAM application for buoyant and turbulent flows.
- `buoyantSimpleFoam`: steady OpenFOAM application for buoyant and turbulent flows.
- `fvSchemes`: dictionary controlling the numerical schemes used by the application.
- `fvSolution`: dictionary defining the equation solvers, tolerances and algorithms.
- `kOmegaSST`: $k - \omega$ shear stress transport turbulence model, based on the work of Menter et al. (2003).
- `PCG`: preconditioned conjugate gradient solver.
- `scalarSemiImplicitSource`: option that adds a source in the transport equation.
- `smoothSolver`: solver using a smoother.
- `thermophysicalProperties`: dictionary where the thermophysical models are defined.

Appendix B

Far-field actual to proxy exposure ratio

This appendix shows the remaining figures of the analysis performed in Chapter 5.

B.1 Original vent set-up (single-sided configuration)

B.1.1 Room averaged CO_2 as proxy

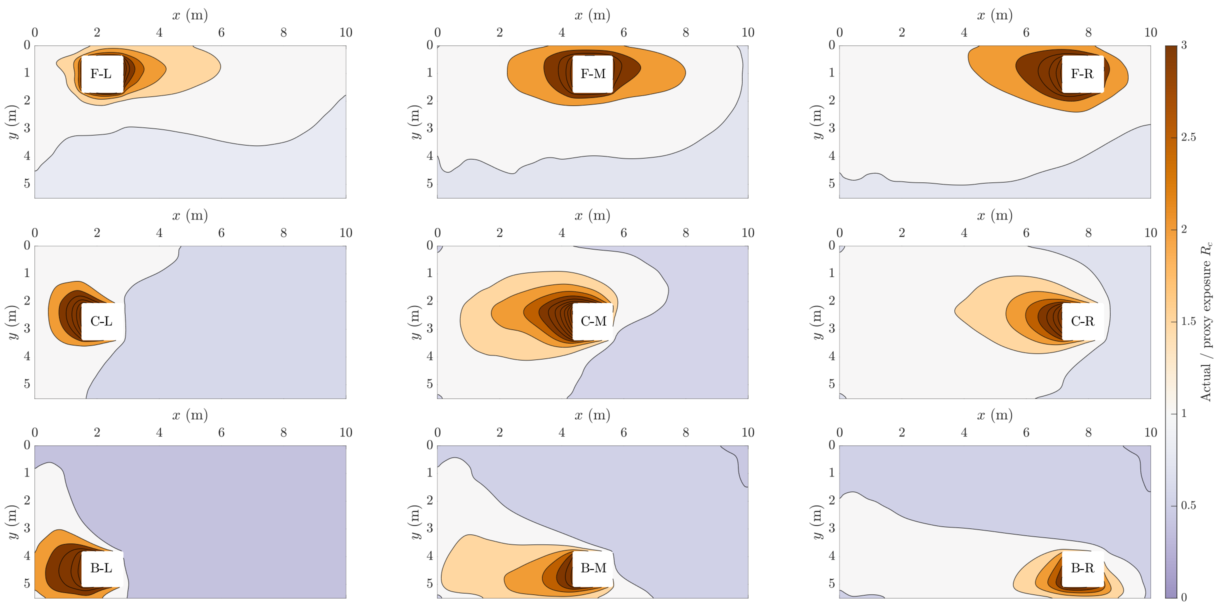

Figure B.1: Horizontal cross section of the actual to proxy far-field exposure ratio R_c for the single-sided configuration in the breathing zone shown for all nine infector locations. The position of the infector is shown by the white square. The proxy exposure is calculated from the room averaged CO_2 concentration $\overline{C}$.

B.1.2 Local CO_2 as proxy

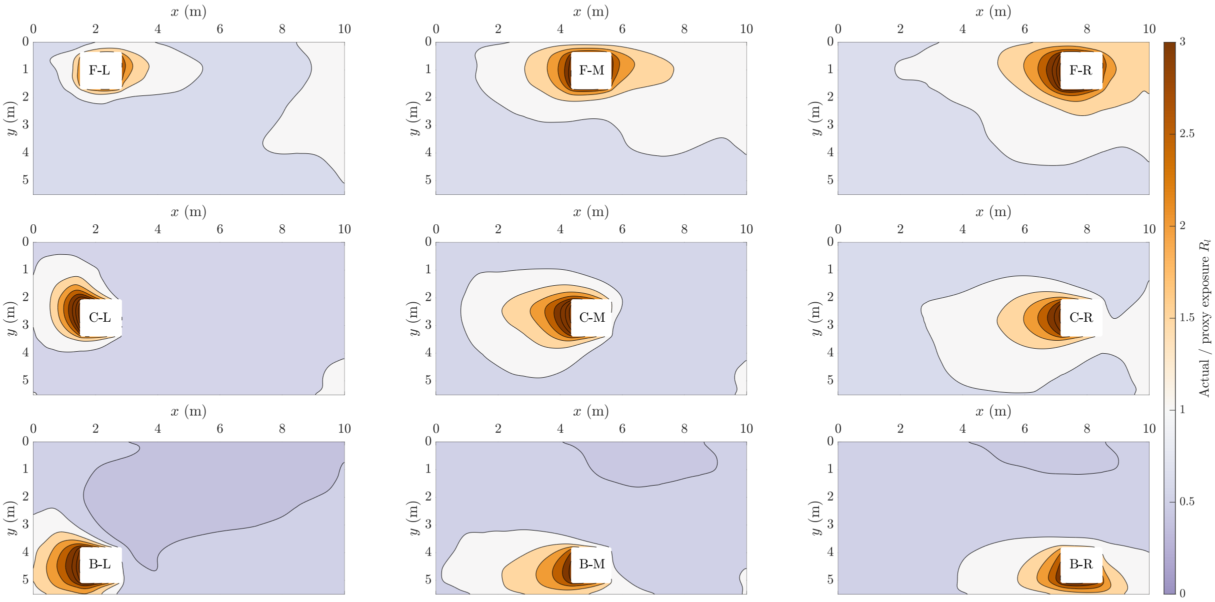

Figure B.2: Horizontal cross section of the actual to proxy far-field exposure ratio R_l for the single-sided configuration in the breathing zone shown for nine infector locations. The position of the infector is shown by the white square. The proxy exposure is calculated from the local CO_2 concentration C at a point.

B.2 Sensitivity to changes in the flow rate

Figure B.3 and Figure B.4 show the distribution of the actual to proxy exposure ratio for all potential infector locations and set-ups considered by varying the vent sizes and heat input. In Figure B.3, the proxy exposure is calculated from the room averaged CO_2 concentration $\overline{C}$ and this is summarised across all infector locations in Figure 5.9. For reference, the distributions obtained with a proxy exposure calculated from the local CO_2 concentration C is also shown in Figure B.4.

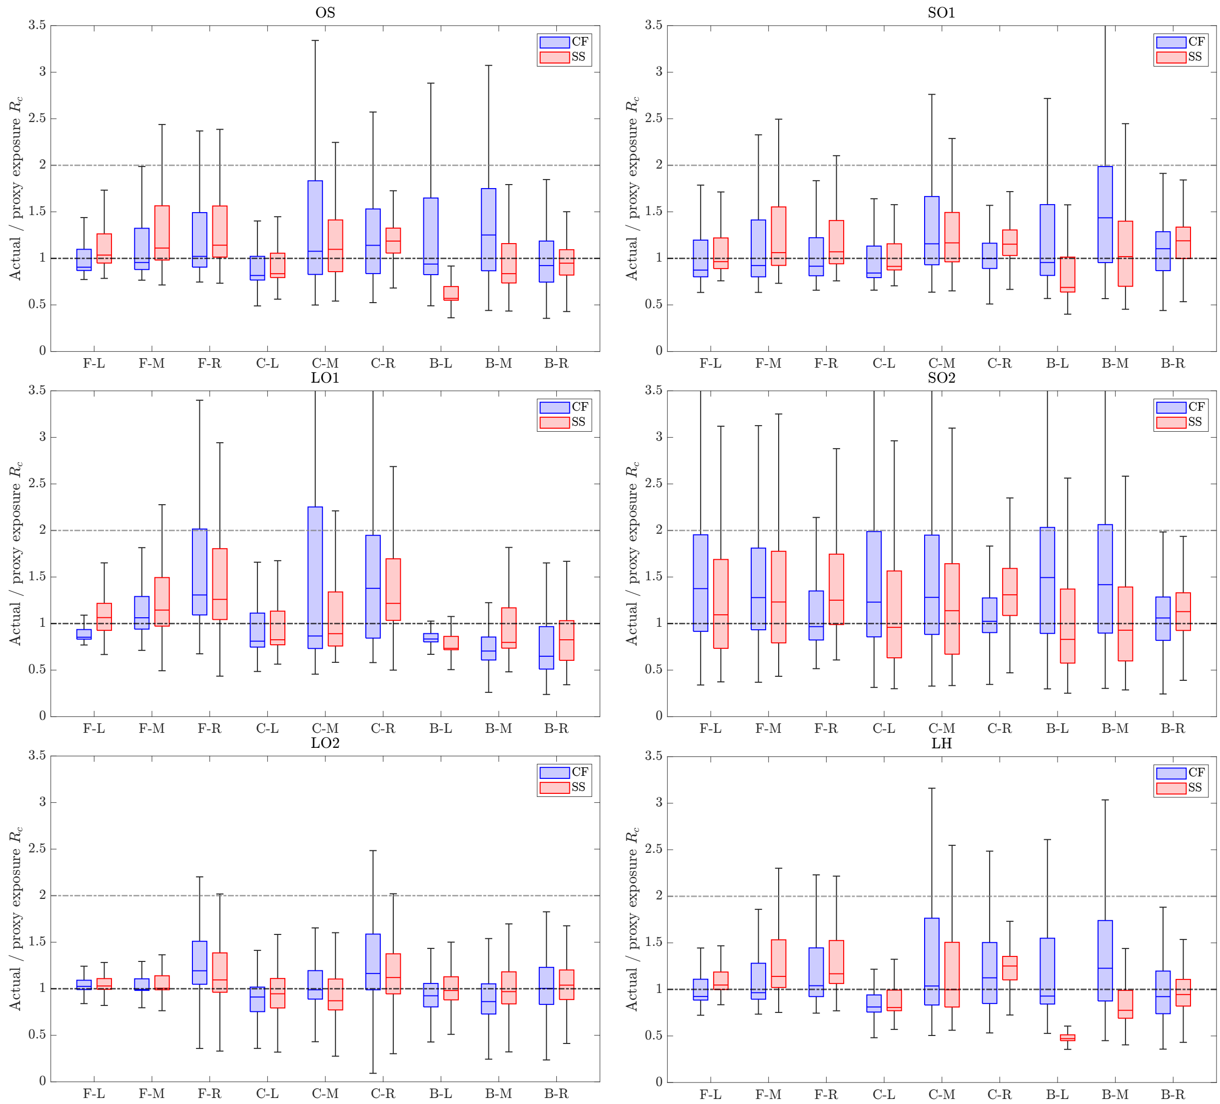

Figure B.3: Actual to proxy exposure ratio R_c distribution for all nine infector locations for both the cross-flow (blue) and single-sided (red) configurations, for the 6 set-ups described in Table 4.4. The proxy exposure is found from the room averaged CO_2 concentration $\overline{C}$. The bottom and top of the boxes show the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1$^{\text{st}}$ and 3$^{\text{rd}}$ quartiles. Horizontal dashed lines show a ratio equal to 0.5, 1 and 2 respectively.

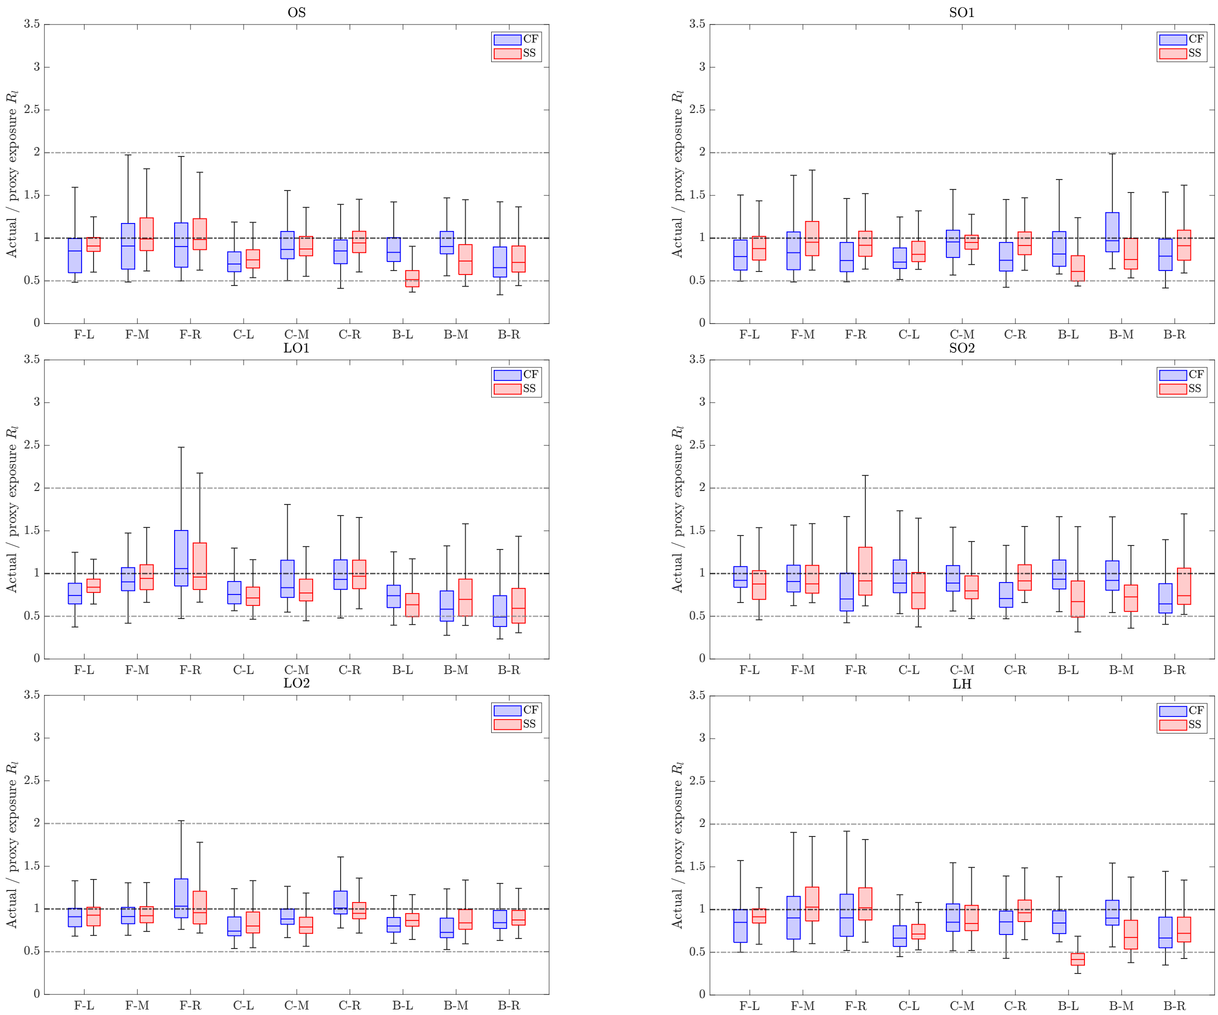

Figure B.4: Actual to proxy exposure ratio R_l distribution for all nine infector locations for both the cross-flow (blue) and single-sided (red) configurations, for the 6 set-ups described in Table 4.4. The proxy exposure is found from the local CO_2 concentration C at a point. The bottom and top of the boxes show the 25th and 75th percentiles, the median is shown by the central line. The whiskers include data within one and a half of the interquartile range from the 1st and 3rd quartiles. Horizontal dashed lines show a ratio equal to 0.5, 1 and 2 respectively.