

Cooking accounts for a major portion of exposure to gaseous pollutants and particulate matter in residential apartments

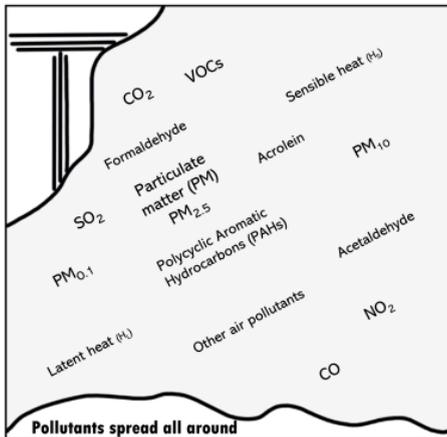
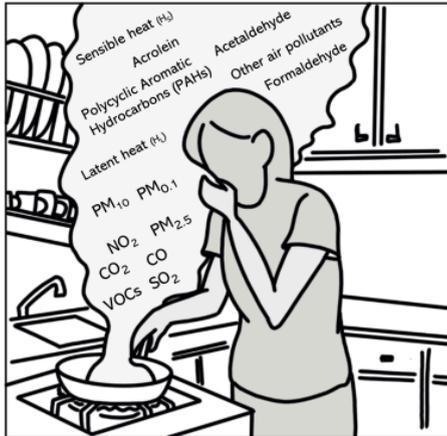
 indooraircartoon.com/2025/08/28/cooking-accounts-for-a-major-portion-of-exposure-to-gaseous-pollutants-and-particulate-matter-in-residential-apartments

28 August 2025

Indoor Air Cartoon Journal, August 2025, Volume 8, #169

[Cite as: Fadeyi MO (2025). Cooking accounts for a major portion of exposure to gaseous pollutants and particulate matter in residential apartments. *Indoor Air Cartoon Journal*, August 2025, Volume 8, #169.]

COOKING ACCOUNTS FOR A MAJOR PORTION OF EXPOSURE TO GASEOUS POLLUTANTS AND PARTICULATE MATTER IN RESIDENTIAL APARTMENTS



1 To control indoor air pollutant build-up over time, its introduction rate into an indoor space must be reduced; otherwise, its concentration will rise quickly, especially in rooms with smaller volumes.

The concentration can be managed by using a sink, which removes a portion of the pollutant from the air.

This concept forms the basis of the equation $dc/dt = P/V - kC_t$, which mathematically expresses the rate of change of pollutant concentration (dc/dt), depending on emission rate (P), room volume (V), current concentration (C_t), and total removal (sink) rate (k).

2 If the portion of a pollutant removed (kC_t) from the indoor air is high but still less than the amount of the pollutant introduced (P/V) at a given time, the dc/dt will remain positive, meaning the pollutant concentration will continue to increase, albeit at a slower rate.

The size of the room is crucial — smaller volumes accelerate the rise in dc/dt as the same pollutant amount occupies less air, increasing concentration more rapidly. Conversely, larger spaces offer greater buffering capacity, slowing pollutant build-up as the same amount disperses into a larger air volume.

Either way, in the absence of a removal mechanism, e.g., an exhaust hood or proper ventilation, dc/dt of a pollutant will increase fast and spread throughout the apartment. However, when kC_t exceeds P/V , the dc/dt of the pollutant will decline. If both are equal, the dc/dt equals zero, meaning no further increase or decrease in concentration over time.

Professor

©2025 Moshood Fadeyi

OMG! Open kitchen concept!

4 The ventilation provided by the windows does not seem sufficient during cooking....

My lecture got her reflecting well.

5 Additionally, many residents lack air cleaners and keep their windows closed because of air con use and polluted outdoor air, leading to a build-up of cooking-related pollutants indoors and increasing both exposure and health risks.

3 Our high-rise residential apartments lack exhaust hoods and have an open-concept kitchen layout.

To achieve a noticeable impact, all windows must be opened with curtains drawn. However, privacy concerns—especially from adjacent apartments or visual intrusion—cause hesitation, significantly limiting the ventilation rate and reducing potential benefits.

Reality: Much of the exposure to gaseous pollutants and PM in residential apartments comes from cooking.

© Moshood O. Fadeyi 2025 | Contact at: mofadeyi@gmail.com | www.indooraircartoon.com

REFERENCES

- Fadeyi MO (2025). Indoor Air Cartoon Journal, February 2025, Volume 8, #163.
Less, B., et. al. (2015). Science and Technology for the Built Environment, 21(1), 14-24.
Patel, S., et. al. (2020). Environmental Science & Technology, 54(12), 7107-7116.

Fictional Case Story (Audio – available online) – Part 1

Fictional Case Story (Audio – available online) – Part 2

Fictional Case Story (Audio – available online) – Part 3

Fictional Case Story (Audio – available online) – Part 4

Fictional Case Story (Audio – available online) – Part 5

.....

Once upon a time in Jimpalia, a wealthy country, public high-rise apartments rose as emblems of progress, promising modernity and convenience. Yet residents were forced to live with air that grew oppressive whenever meals were prepared. Without dedicated exhaust ventilation systems, pollutants from cooking — nitrogen dioxide, carbon monoxide, formaldehyde, and fine particles — accumulated indoors, sometimes exceeding outdoor levels. The expectation was clear: homes should provide clean, breathable air for comfort, health, and concentration. The reality fell far short, with rooms turning acrid and stuffy after the simplest meals.

The persistence of the problem was deeply rooted in Jimpalia’s culture. A society accustomed to patchwork measures saw little reason to confront root causes. Public health threats such as poor indoor air quality were normalised, their discomforts hidden beneath temporary fixes. The nation’s wealth and excellent healthcare system further obscured the issue, treating the illnesses that arose while leaving the source untouched. Regulators upheld minimum standards, declaring buildings “adequate” even as pollutants built up indoors. In this way, the culture of superficial solutions allowed the problem to persist throughout Jimpalia’s public high-rise apartments, wasting resources while the fundamental failure — the absence of effective, preventive ventilation — remained ignored.

Among those who noticed was a foreign student, seeking a better future in this wealthy nation. He understood what few around him seemed willing to see: unresolved problems, however neatly disguised, always spread and returned with greater force. To him, ignoring root causes was like building a fragile edifice — polished on the outside yet fractured within. He resolved to do something about it. This student’s journey is the subject of this fiction story.

1.....

Oluwajamisi’s name meant “*the Lord made me realise,*” and in his young life, the meaning was not just symbolic but prophetic. Born into a low-income family in Pacimany, a developing country riddled with struggles, he had learnt very early that survival required more than determination; it required foresight. His parents, Mr and Mrs Orimiwanbe, had little in terms of wealth, yet they poured into him the discipline of learning and the strength of endurance.

By the time he graduated from the University of Dayanmo with first-class honours in a double degree programme in Chemical Engineering and Mathematics, jointly offered by the Faculty of Engineering and the Faculty of Science, scoring an extraordinary 4.97 out of 5, he was already known in his community as the boy whose brilliance would travel further than the borders of Pacimany.

When the offer came from the University of Pompey in Jimpalia—a developed country with a reputation for global prestige—it seemed like destiny calling. The MSc in Environmental Health Engineering (a coursework with dissertation programme) was exactly in line with his dream to improve lives through engineering solutions.

The problem was that the offer did not come with a scholarship. What it promised was a possibility: a research assistantship, though not guaranteed, and a tuition loan covering 80% of the fees. For Oluwajamisi, that was enough. It was his only offer, and the thought of studying overseas, especially in one of the world's most sought-after universities, outweighed the uncertainty.

The boy from Pacimany had prepared for this moment long before it arrived. He had saved every little coin he could, knowing that one day he might need it. For seven months, while waiting for his travel date, he worked as a tutor for O Level and A Level students, teaching chemistry and mathematics. He was paid modestly, but he never squandered his wages. Along with a small government award he received for his exceptional academic performance, he now had just enough for a plane ticket and to survive in Jimpalia for a week.

Arriving in Jimpalia, Oluwajamisi was struck by its beauty and its rise. The country was celebrated worldwide for its wealth, its universities, and its innovations. Yet beneath the surface, he quickly noticed a flaw: problems were rarely solved at their source. Instead, issues were allowed to spread, growing in scale and complexity until they became societal burdens, hidden under the country's veneer of prosperity. What troubled him most was that this flaw was not spoken of openly; it was covered up, as if solving the surface symptoms were enough, even as the costs of neglect multiplied.

At the University of Pompey, life quickly became more difficult for Oluwajamisi than he had ever imagined. His tuition loan covered the bulk of his academic fees, but beyond that he was left to fend for himself in a city whose expenses dwarfed anything he had known in Pacimany. The money he had carried from home—savings from tutoring and a modest government award—barely lasted a week. He rationed meals with the precision of a scientist, stretching them until they were little more than token sustenance. Some days he went without entirely, too focused on conserving coins to think about hunger.

The prospect of working as a research assistant was his only tangible hope, but that position did not materialise immediately. In the silence of his small room, he lived with gnawing anxiety: the fear of eviction, of failing to keep up with his studies, of being stranded without a safety net in a foreign land. Each night he fell asleep exhausted, yet restless, carrying the weight of uncertainty.

Survival demanded ingenuity. For Oluwajamisi, the answer came from the subject he had always mastered best: mathematics. Numbers spoke for him more clearly than words, and they became his lifeline. He wrote notices offering tutoring services in mathematics and chemistry, tacking them onto bulletin boards across campus. Slowly, requests trickled in. Undergraduates

struggling with problem sets turned to him, grateful for his clarity of explanation. The few notes of currency he earned were never abundant, but they kept him afloat. Still, he lived with a constant edge of fear, aware that the smallest misstep could undo everything.

Material hardship was only half of what shaped his early months in Jimpalia. What unsettled him most was the culture of the country itself. Jimpalia dazzled with wealth and sophistication, yet beneath the glitter he discerned a troubling flaw: problems were seldom solved at their origin. Instead, temporary measures were applied, neat on the surface but hollow in substance. He saw it first in student housing.

When pipes leaked in the hostels, maintenance staff set out buckets rather than repair the corroded infrastructure. When water systems broke down, temporary pumps were installed without any further improvement actions, a patch that only postponed the inevitable. When ventilation units failed, portable fans were brought in without any further improvement actions, a patch that only postponed the inevitable. When lighting systems faltered, quick fixes with temporary wiring were used without any further improvement actions, a patch that only postponed the inevitable.

This tendency stretched beyond campus. News reports spoke of polluted rivers being “managed” by extensive downstream cleaning, while factories upstream were permitted to continue their discharges. Healthcare costs rose dramatically, not because the nation lacked advanced medicine, but because preventive measures were chronically neglected. Oluwajamisi began to recognise the same pattern everywhere: a society comfortable with appearances, yet unwilling to engage with root causes.

In Pacimany, the flaws he knew were different. His people lived with leaking roofs, smoky kitchens, and failing roads not because they chose superficiality but because resources were scarce. The phrase “We’ll manage” became a survival instinct, a cultural acceptance that when means were absent, endurance had to suffice. In Jimpalia, however, the means existed in abundance. There were funds, knowledge, and technology, yet the will to strike at root causes seemed absent. Instead, cosmetic solutions concealed deeper decay.

The comparison grew sharper each day. He saw that Pacimany’s struggle was born of poverty, while Jimpalia’s was born of complacency. The outcome, however, bore striking similarities: problems lingered, spread, and became costlier to address. What Pacimany could not solve because of lack, Jimpalia would not solve because of ease. In both lands, the consequences fell on ordinary people.

For Oluwajamisi, this realisation became as formative as any lecture or laboratory session. Each class he attended on environmental health and engineering concepts seemed to echo his lived observations. When professors spoke of pollution pathways, he thought of the acrid smoke in his apartment.

When models of resource allocation were discussed, he remembered government spending in Jimpalia on downstream clean-ups rather than prevention. And when the notion of performance gaps was raised—between what systems promised and what they delivered—he recognised it vividly in the daily experiences of students and citizens alike.

What kept him moving forward was not glamour or comfort but a profound sense of purpose. He reminded himself that his presence at the University of Pompey was not accidental. His name itself, Oluwajamisi—"the Lord made me realise"—carried an insistence that awareness was part of his destiny. The realisation he was coming into was stark: hardship could not drown his purpose, and the flaws of nations, whether born of poverty or abundance, could only be addressed by those willing to seek solutions at their roots.

Life at the University of Pompey, then, was not a tale of prestige but of survival, sacrifice, and clarity. He endured hunger so that he might one day feed others with knowledge. He accepted loneliness so that he could learn how societies, rich or poor, carried their weaknesses. And he bore the anxiety of every uncertain day with the conviction that the hardship of today was an investment in the solutions of tomorrow.

As he walked through the manicured streets of Jimpalia, past cafés brimming with laughter and wealth, Oluwajamisi carried within him a perspective sharpened by contrast. He understood what few around him seemed to see: unresolved problems, however well disguised, eventually spread and returned with greater force. To ignore root causes was to build a fragile edifice, a society polished on the outside yet fractured within.

It was this awareness that shaped his ambition more deeply than any textbook. For Oluwajamisi, the University of Pompey in Jimpalia was more than a place of study. It was a crucible where hardship and observation fused into resolve, preparing him not merely to master his discipline but to use it in service of genuine solutions to solve practical problems in a value-oriented manner for stakeholders involved. His story was still unfolding, but already the lesson was clear: true strength lay not in wealth or display, but in the courage to face problems at their origin and build systems that endured.

This culture unsettled Oluwajamisi. As a student of Environmental Health Engineering, he was taught to seek root causes: pollutants in air, toxins in water, pathogens in buildings. His professors spoke of advanced technologies, data-driven monitoring, and cutting-edge interventions. But in practice, even here, he saw the contradiction—laboratories full of innovation while the city outside repeated the cycle of treating symptoms and ignoring origins. The country's rise concealed its decay.

Securing a research assistantship became the turning point of his survival. After months of relentless struggle, one professor, impressed by Oluwajamisi's brilliance in a seminar, offered him a position assisting with environmental modelling work. The modest stipend did not take away his financial troubles entirely, but it steadied him. For the first time since arriving, he could breathe.

But even in the research projects he supported, Oluwajamisi saw the same cultural flaw reflected. Studies were funded to monitor air pollution spikes, yet little attention was given to regulating the industries causing them. Models were developed to predict spikes in hospital admissions during periods of extreme environmental stress, but rarely did policies address why

city designs created the conditions in the first place. Oluwajamisi worked diligently, but he carried the quiet discomfort of knowing that much of what he contributed was destined to be applied superficially.

The hardship he faced sharpened his resolve. At night, in his small room, he wrote in his journal: *“The Lord made me realise. I see now: to solve problems, one must touch their roots. Anything less is deception. Pacimany’s poverty blinds it, Jimpalia’s wealth masks it. But neither has yet embraced truth.”*

Oluwajamisi knew his journey was only beginning. The overseas exposure he longed for was opening his eyes not only to engineering but to the cultures of nations and the philosophies that guided their choices. He saw clearly now that his mission was larger than obtaining a degree. He would learn from Jimpalia, yes, but also rise above its flaw. He carried in him the determination to build a future where solutions were not patches, but cures—where problems were not hidden beneath wealth or endured beneath poverty, but uprooted, examined, and resolved.

After two years of relentless endurance, Oluwajamisi completed his MSc in Environmental Health Engineering at the University of Pompey. The coursework was intense, layered with expectations, and financially punishing, yet he emerged with distinction. His professors, impressed by his consistency and intellectual depth, recommended him for a PhD scholarship. Unlike his MSc offer, this one came with full funding.

For the first time in Jimpalia, Oluwajamisi felt a weight lift from his shoulders. He could study without worrying every night whether hunger or rent arrears would drive him out of the university. The scholarship covered his tuition with no requirement to pay it back. It also included a decent monthly allowance. In addition to using the allowance for his personal upkeep and to pay for his student accommodation, he also used part of it to repay his MSc loan debt in affordable instalments.

2

The story of how his PhD topic emerged was not confined to classrooms or laboratories. It was rooted in lived experience, a real-life event that turned observation into purpose. He was living in a cramped apartment block on the eastern side of Pompey City, where international students often rented small, ageing apartments because they were cheaper than the modern residences. The building was tall, with long corridors; the kitchens were little more than small corners off the living areas, and the bedrooms were narrow, giving scarcely any privacy within the apartment.

His housemate, loved to cook with aromatic spices, frying fish or stir-frying vegetables late at night. Another neighbour, directly beneath, often roasted meats and left the oven door propped open, filling the apartment with extra heat and odours. Oluwajamisi himself boiled rice, fried plantains, cooked egusi soup, and sometimes prepared beans, the familiar aromas of Pacimany grounding him in this unfamiliar land.

Yet what stood out to him was how oppressive the air became whenever anyone cooked. The kitchens in the building, like in many of Jimpalia's public housing complexes, had no exhaust ventilation system at all. There was only a small window, which did little more than shift the smoke and odour from one apartment to another. Even when he was not cooking, fumes from neighbouring apartments drifted into his space. A fine haze settled on his desk near the kitchen wall, and the sharp odour of frying oil clung to his clothes in the laundry area, lingering for hours after neighbours and his housemate had prepared their meals.

The discomfort grew because there was no way to avoid it. With no proper ventilation, cooking fumes lingered indoors, leaving the air acrid and heavy with particles that glimmered in the light from the ceiling lamp. Oluwajamisi began to suffer frequent headaches while studying late into the night, particularly after his housemate had been cooking. His concentration weakened, forcing him to reread passages again and again. What he had learnt in class about pollutants such as nitrogen dioxide and ultrafine particles impairing cognitive function was no longer theory—it was a reality pressing against his own body.

The problem was not confined to his unit. In conversations with fellow students, Oluwajamisi heard similar frustrations. Friends in neighbouring high-rise complexes spoke of constant odours, poor air quality during ordinary cooking, and windows kept shut out of fear of burglary or bad weather.

One classmate confessed she stopped cooking altogether, relying on cold sandwiches to avoid the smoke and lingering odours that filled her apartment. These stories revealed a troubling pattern: comfort, convenience, and even cognitive performance were all undermined by something as ordinary as preparing food. What disturbed Oluwajamisi most, however, was how the issue reflected the wider cultural flaw he had already noticed in Jimpalia—problems addressed superficially rather than at their source.

Landlords told tenants to “open a window” if air became uncomfortable. Some landlords promoted plug-in air fresheners as a solution, masking odours while pollutants remained. A few residents purchased portable air purifiers, but these only diluted the symptoms rather than capturing emissions at the stove. Building managers insisted the apartments were “to standard” because regulations did not require dedicated kitchen exhausts.

The way ventilation was neglected in Jimpalia's high-rise apartments reflected a wider culture of superficial problem-solving. Instead of addressing the absence of proper kitchen exhaust systems, landlords and building managers told tenants to open windows or use air fresheners. This only shifted fumes around, allowing smoke and odours to spread from one apartment to another, compounding discomfort and creating new risks for residents.

Over time, greater resources were then needed to manage the consequences—health complaints, disputes between neighbours, and the wider burden on public health. What struck Oluwajamisi most was that this cycle was rarely challenged. Because Jimpalia was wealthy and could afford costly downstream fixes, many did not even see it as a problem. Affluence dulled urgency, and so the tendency to ignore root causes persisted, keeping appearances intact while deeper flaws in housing design spread quietly through society.

It was one incident, however, that became the true spark for his PhD journey. One evening, his neighbour from downstairs came pounding on his door, furious that smoke and odours had drifted into her apartment, disrupting her study group and leaving the living area uninhabitable. Oluwajamisi and his apartment mate had only been steaming rice and frying vegetables, yet the fumes had travelled upwards, permeating another person's space.

The confrontation that followed was heated, but behind it lay something Oluwajamisi could not ignore: the inadequacy of design in modern high-rise living. A simple act of cooking had turned into both a health hazard and a source of conflict among residents — not only within the same apartment but also between neighbouring apartments.

He had previously heard of similar disputes — housemates quarrelling over lingering fumes in shared student kitchens, family members in apartments arguing among themselves, and even across the country neighbours falling out because cooking activities had made the air unbearable. Yet it was this confrontation, experienced directly, that jolted him. What he had known in theory as a widespread problem became personal, a wake-up call that he needed to do something about it.

That night, unable to sleep, he sat at his desk and scribbled in his notebook. His thoughts roamed from his childhood in Pacimany, where poverty forced people to 'manage' with leaking roofs and smoky kitchens, to Jimpalia, where wealth produced shining towers that still failed to address something as basic as clean indoor air. In both places, people were asked to live with problems rather than seeing them solved at their origin.

It was then that Oluwajamisi decided his PhD would focus on this gap. His lived reality — the haze in his own room, the headaches that blurred his studies, the tensions with neighbours, and the societal culture of superficial fixes — had crystallised into a single truth: cooking in high-rise apartments was not just about food, but about health, comfort, and the integrity of design. If these spaces could not guarantee safe and breathable air, they had failed in their most basic duty.

This was the event — the lived experience of struggling in poorly ventilated student housing, surrounded by neighbours suffering the same fate — that compelled Oluwajamisi to shape his PhD problem. He resolved not only to measure the scale of the issue but to develop the tools and interventions that could finally reach the root. His scholarship was more than an award for academic performance; it was the beginning of a mission that sprang from the very air he breathed in Jimpalia.

When Oluwajamisi finally decided to share his thoughts, he did so cautiously. He requested a private meeting with one of his professors, knowing that speaking openly in seminars might expose his idea to others who could seize it before he had the chance to develop it.

In the quiet of the professor's office, surrounded by shelves stacked with journals and reports, Oluwajamisi unfolded the pages from his notebook. He described the suffocating haze that lingered in his apartment, the fumes drifting between neighbouring units, and the heated confrontation that had erupted when a neighbour complained about smoke invading her space.

He spoke, too, of the cultural flaw he had come to recognise in Jimpalia—problems addressed superficially rather than at their roots, allowed to spread silently and consume ever more resources over time.

The professor listened intently, nodding as Oluwajamisi spoke. “What you are describing,” she said at last, “is more than an inconvenience of student housing. It is a performance gap between what high-rise living is meant to provide—comfort, safety, well-being—and the reality that residents face. And this is not limited to Jimpalia. Cities everywhere are building upwards without solving this problem.”

Encouraged, Oluwajamisi continued. He explained how his coursework in pollutant modelling now felt deeply personal, because the concepts he studied—nitrogen dioxide, ultrafine particles, carbon monoxide—were no longer abstractions. During his MSc dissertation, he examined these pollutants in the context of industrial emissions and their wider impact on urban environments, particularly how outdoor pollution could seep indoors and affect air quality.

At the time, the focus felt distant, centred on factories and city-scale monitoring. Now, however, he was confronting the same pollutants in the intimate confines of his own apartment, where the ordinary act of cooking made their presence immediate and personal. The professor leaned back, thoughtful. “If you take this forward, you must do it rigorously. Your lived experience gives you a powerful insight, but your task will be to transform it into research that others can trust and use.”

Leaving the office, Oluwajamisi felt his burden lighten. The private conversation gave him confidence. What had begun as personal suffering in a poorly ventilated apartment was now a potential PhD project—rooted in lived reality, yet positioned to address a global challenge. With his professor’s encouragement, he began to shape his experiences and observations into a structured plan, gradually refining them into a formal proposal. His research proposal is as follows:

“Modern high-rise residential apartments are increasingly designed to maximise density, efficiency, and convenience, especially in rapidly urbanising cities. However, IAQ in these apartments has not always been prioritised, particularly in relation to cooking activities. Cooking is recognised as one of the largest contributors of indoor air pollutants, generating both gaseous contaminants and fine particulate matter that can accumulate indoors.

The expected performance level for these living environments is that they should provide comfort, convenience, and safety, supporting both the physical health and cognitive well-being of residents. This requires clean, breathable air even during and after routine cooking events. The current performance level, however, often falls short, with many apartments exposing residents to pollutant concentrations that may exceed health-based guidelines. This discrepancy forms a substantial performance gap that undermines occupant comfort, safety, and awareness.

Cooking is known to emit a complex mix of pollutants such as nitrogen dioxide (NO₂), carbon monoxide (CO), formaldehyde (HCHO), and particulate matter in several fractions, including PM_{2.5}, PM₁₀, and ultrafine PM_{0.1}. Ideally, residential apartments should be designed to remove

or dilute these emissions rapidly, preventing exposures that cause discomfort or harm.

Yet, in practice, many high-rise apartments without dedicated exhaust ventilation systems struggle to meet this standard. Indoor-to-outdoor ratios of pollutants have been reported as disproportionately high, with fine particulate matter sometimes reaching levels several times greater than outdoors. Instead of providing a safe refuge from external air pollution, indoor spaces can become more hazardous during cooking.

The implications are multifaceted. Elevated pollutant concentrations can affect comfort, by causing odours, visible smoke, or stuffy environments; convenience, when emissions frequently trigger smoke alarms or require manual ventilation during unfavourable weather; and cognitive performance, since even moderate pollutant exposures have been associated with reduced attention and slower task efficiency. The expected condition—clean, healthy air during cooking—remains distant from the lived reality in many high-rise apartments.

Understanding and addressing this gap requires predictive tools that can estimate exposures across different household layouts, behaviours, and meteorological conditions. In principle, such models should be accurate, flexible, and useful for designing effective interventions. However, the tools currently available present important limitations.

Mass-balance models simplify pollutant dynamics by assuming uniform mixing and fixed air exchange rates, meaning they cannot adapt to sudden changes in occupant behaviour, such as opening a window mid-cooking. Computational fluid dynamics (CFD) models provide detailed spatial predictions but are resource-intensive and rigid, requiring predefined boundary conditions that must be re-specified for each new scenario.

The gap, therefore, lies in the absence of modelling approaches that are both accurate and adaptive. Without such tools, interventions risk being planned on oversimplified or contextually irrelevant assumptions. More advanced modelling frameworks, such as AI-enhanced hybrids, have been proposed to integrate empirical measurements with behavioural and contextual data, but their accuracy and generalisability still require systematic validation.

This highlights a discrepancy between the expected performance of predictive tools—to model real-world exposures reliably—and the current limitations, where oversimplification or rigidity hinder practical application.

Even if pollutant exposures are well understood, reducing them effectively in high-rise apartments remains a complex challenge. Conventional kitchen exhaust hoods are often unavailable or impractical in these settings. When they are present, they may be underused due to excessive noise, high energy demand, or maintenance difficulties. Shared ducting in high-rise buildings further complicates matters, as one household's exhaust can leak into another's living space, undermining IAQ building-wide.

The expected standard is an exhaust ventilation system that is cost-effective, compact, and easy to maintain, while capturing pollutants efficiently at source and preventing redistribution. Current systems frequently fail to deliver this balance: they may be bulky, intrusive, energy-

intensive, or demand burdensome maintenance. Thus, a clear performance gap exists between what is needed and what is currently available.

Research efforts in this area are directed toward developing advanced solutions—integrating filtration, optimised capture geometry, and adaptive controls—but such innovations require rigorous evaluation in real-world high-rise environments before their feasibility can be established.

Bringing these strands together, the central problem is the wide gap between expectation and reality in managing cooking-related pollutant exposure in high-rise apartments. On one side, cooking contributes disproportionately to indoor air pollutant concentrations, challenging comfort, convenience, and cognitive well-being. On the other, the tools available to predict and mitigate these exposures—whether modelling frameworks or ventilation technologies—fall short of providing reliable, user-friendly, and cost-effective solutions.

The expected performance level is that residents should be able to cook in environments where exposures do not exceed health guidelines, where predictive tools can guide interventions across varied scenarios, and where ventilation systems offer protection without compromising comfort, usability, or affordability.

The current performance level, by contrast, remains marked by pollutant concentrations above safe thresholds, models that struggle with real-world variability, and ventilation solutions that are not well suited to household life. This persistent gap points to the urgent need for research that can measure exposures accurately, evaluate predictive models for reliability, and design and test intervention systems tailored to the specific constraints of high-rise living.”

The research questions are as follows: (i) To what extent does cooking contribute to measured occupant exposure to gaseous pollutants and particulate matter in residential apartments compared to other indoor and outdoor sources? (ii) What is the predictive accuracy of mass-balance, computational fluid dynamics, and AI-enhanced hybrid models for estimating occupant exposure to cooking-related pollutants across diverse apartment layouts, behaviours, and meteorological conditions? (iii) How can a cost-effective, maintenance-friendly, advanced filtered exhaust ventilation system be designed, integrated, and evaluated in high-rise residential apartments to achieve substantial, verifiable reductions in cooking-related pollutant exposure, measured and modelled under real-world conditions?

For the first research question, the Null Hypothesis (H_{01}) is that cooking does not account for a significantly greater proportion of occupant exposure to targeted gaseous pollutants and particulate matter compared to other sources in residential apartment. The Alternative Hypothesis (H_{11}) is that cooking accounts for a significantly greater proportion of occupant exposure to targeted gaseous pollutants and particulate matter compared to other sources in residential apartment.

For the second research question, the Null Hypothesis (H_{02}) is that there is no statistically significant difference in predictive accuracy among mass-balance, CFD, and AI-enhanced hybrid models for estimating cooking-related exposure. The Alternative Hypothesis (H_{12}) is that

at least one modelling approach exhibits statistically significantly higher predictive accuracy for estimating cooking-related exposure than the others.

For the third research question, the Null Hypothesis (H_{03}) is that the cost-effective, maintenance-friendly, advanced filtered exhaust ventilation system does not significantly reduce measured or modelled occupant exposure to cooking-related pollutants compared to baseline conditions. The Alternative Hypothesis (H_{13}) is that the cost-effective, maintenance-friendly, advanced filtered exhaust ventilation system significantly reduces measured or modelled occupant exposure to cooking-related pollutants compared to baseline conditions.

The research questions and problems informed the following objectives of his PhD research: (i) To quantify the contribution of cooking activities to occupant exposure to gaseous pollutants and particulate matter in high-rise residential apartments, and to compare these exposures against other identifiable indoor and outdoor pollutant sources. (ii) To evaluate and compare the predictive accuracy of three modelling frameworks—mass-balance, computational fluid dynamics (CFD), and AI-enhanced hybrid models—in estimating occupant exposure to cooking-related pollutants under varying apartment layouts, occupant behaviours, and meteorological conditions. (iii) To design, integrate, and rigorously evaluate a cost-effective, maintenance-friendly, advanced filtered exhaust ventilation system specifically tailored for high-rise residential kitchens, with the objective of achieving measurable and verifiable reductions in cooking-related pollutant exposure through both field measurements and validated modelling approaches.

Below is an excerpt from Oluwajamisi's PhD thesis.

3

Research Methods

Methods for Research Question 1:

Study Design and Sampling Strategy

The study followed a planned approach, collecting new data over time in real-world conditions from a representative group of high-rise residential apartments. This approach allowed the research to capture authentic patterns of indoor air pollutant generation, movement, and occupant exposure in daily life, without the artificial constraints often present in laboratory experiments. The design was particularly important for understanding cooking-related indoor air pollutant exposure because cooking behaviours, ventilation habits, and building characteristics can vary widely between households and over time.

A total of thirty apartments were included, with the number chosen to balance the need for statistical reliability with the practicality of collecting high-resolution data over extended monitoring periods. The selection process followed a stratified strategy to ensure that the sample included a mix of floor levels, building ages, ventilation types, and occupant demographics.

Apartments were distributed across low-rise (below the 5th floor), mid-rise (6th to 15th floor), and high-rise (above the 15th floor) levels, recognising that the vertical position in a building can affect the entry of outdoor pollutants and the dispersion of indoor-generated ones. Building ages ranged from recently constructed apartments to those more than thirty years old, allowing for comparisons between modern, airtight construction and older, more naturally ventilated structures. Ventilation configurations also varied, with some apartments relying solely on operable windows and others incorporating portable fans.

To ensure that the observed indoor air pollutants originated from regular household activities rather than from unusual circumstances, the inclusion criteria specified that households had to cook at least four times per week using common appliances such as gas or electric stoves. Apartments undergoing major renovations during the study period were excluded, as construction work can significantly affect indoor air quality. None of the participating apartments was equipped with a functioning kitchen exhaust hood, reflecting common conditions in the studied building stock and avoiding the influence of dedicated source-extraction systems on cooking-related measurements.

Target Indoor Air Pollutants and Measurement Parameters

The study focused on six key indoor air pollutants known to be produced during cooking and to have established health impacts. These included nitrogen dioxide, carbon monoxide, and formaldehyde as representative gaseous indoor air pollutants, alongside particulate matter in three aerodynamic size ranges: $PM_{0.1}$, $PM_{2.5}$, and PM_{10} .

Nitrogen dioxide was selected as a primary marker of combustion, particularly relevant for gas cooking. Carbon monoxide was monitored because of its association with incomplete combustion and its potential to cause harm at both acute and chronic exposure levels. Formaldehyde was included due to its known irritant properties and long-term health risks, as well as its frequent release during cooking from heated oils and organic matter.

Particulate matter was measured in three fractions to capture the full range of particle sizes associated with cooking emissions. $PM_{0.1}$ represented ultrafine particles capable of penetrating deep into the respiratory system and entering the bloodstream. $PM_{2.5}$ referred to fine particles strongly linked with cardiovascular and respiratory illnesses, while PM_{10} encompassed coarse particles that can cause respiratory irritation.

In addition to these indoor air pollutants, carbon dioxide was monitored indoors as a tracer gas to estimate ventilation rates. While carbon dioxide at typical indoor levels does not pose a direct health risk, it provides valuable information about the rate at which indoor air is replaced with outdoor air. To help distinguish pollutants generated indoors from those entering from outside, identical measurements were taken outdoors at each study site. This allowed the calculation of indoor-to-outdoor ratios, an important metric for understanding the relative influence of indoor and outdoor sources.

Instrumentation and Data Acquisition

Data collection relied on research-grade instruments capable of providing high-precision, continuous measurements over long periods. Particulate matter was measured using optical particle counters, which work by detecting light scattered by particles passing through a laser beam and can determine both the size and number of particles present.

Nitrogen dioxide concentrations were obtained using chemiluminescence analysers, which detect the light emitted during a chemical reaction involving the gas, offering high sensitivity and accuracy. Carbon monoxide was measured with electrochemical gas analysers, which generate an electrical signal when exposed to the gas.

Formaldehyde, being a reactive and less stable pollutant, required a different approach. It was measured using DNPH (2,4-dinitrophenylhydrazine) cartridge sampling, where the formaldehyde in the air reacts with DNPH to form a stable compound that can be analysed later in the laboratory using high-performance liquid chromatography. Carbon dioxide was measured using non-dispersive infrared sensors, which detect the characteristic absorption of infrared light by the gas. Temperature and relative humidity were also recorded, as they can influence both the generation and dispersion of indoor air pollutants.

Indoor monitors were strategically placed in three locations within each apartment: the kitchen, where pollutants from cooking were most likely to be generated; the primary living area, where occupants spend significant time and where pollutants could disperse; and a representative bedroom, to evaluate whether pollutants reached sleeping areas.

Outdoor monitors were mounted on the building façade at the same vertical level as the monitored apartment to ensure that outdoor measurements accurately reflected the air entering the apartment. All instruments were set to record at one-minute intervals to capture the rapid changes in indoor air pollutant concentrations that occur during cooking.

Activity and Context Logging

Accurately identifying the sources of indoor air pollutants required a combination of careful observation by occupants and objective verification using scientific instruments. To achieve this, each participating household kept a written activity diary throughout the monitoring period.

The diary served as a daily log, recording the start and end times of every cooking event, the specific cooking method used—such as frying, boiling, or grilling—the type of fuel, whether gas or electricity, and any ventilation measures in place, such as opening windows or using fans. This level of detail was essential because different cooking methods and fuels produce different amounts and types of indoor air pollutants, and ventilation practices can greatly influence how these pollutants accumulate or disperse.

The diaries also captured other activities that could release indoor air pollutants, such as burning candles or incense, or using cleaning products containing volatile chemicals. This was important because without such information, increases in pollutant levels during non-cooking periods might be mistakenly attributed to cooking, leading to inaccurate conclusions.

To strengthen the reliability of the self-reported diaries, objective verification systems were installed in each home. Stove-use sensors were fitted to gas burners and electric heating elements to detect exactly when cooking appliances were switched on or off, providing precise timestamps. Motion detectors were also installed in the kitchen to confirm that someone was present during the recorded cooking events, ensuring that the activity logs reflected actual use rather than accidental activation.

In addition, data on outdoor environmental conditions were obtained from nearby air quality monitoring stations. This included identifying periods of higher-than-normal outdoor pollution, such as during heavy traffic peaks or regional haze events. These outdoor data were compared with indoor measurements to prevent misclassification of pollutant peaks that originated from outside rather than from cooking or other indoor activities. This combined approach ensured high accuracy in linking indoor air pollutant peaks to their true sources.

Source Apportionment Analysis

Determining how much of the measured indoor air pollutant exposure came from cooking required a careful and multi-layered scientific approach, combining several analytical techniques to ensure accuracy and reliability. This was necessary because indoor air can contain pollutants from a variety of sources—both inside and outside the home—and simply observing when pollutant levels rise is not enough to determine their true origin.

The first step involved creating time-resolved exposure profiles for each indoor air pollutant. These profiles were essentially detailed timelines, recorded at one-minute intervals, showing how concentrations changed before, during, and after cooking events.

By aligning these profiles with the precise start and end times of verified cooking activities, the research team could visually identify distinct peaks in pollutant levels that coincided with cooking. These were compared against baseline periods when no cooking was taking place, allowing a clear visual contrast between cooking and non-cooking conditions. This step provided the first layer of evidence linking cooking to increases in specific pollutants.

To strengthen this link, the study employed chemical signature analysis, which looks at the unique “fingerprints” left by different sources of emissions. Two key signatures were examined. The first was the ratio of nitrogen dioxide (NO₂) to carbon monoxide (CO). Both gases are released during combustion, but the proportions can vary depending on the source—for example, gas cooking produces a different NO₂-to-CO ratio than traffic exhaust.

The second signature came from particle size distribution. Cooking, especially frying or high-heat methods, tends to produce a sudden surge in ultrafine particles (PM_{0.1}), which behave differently in the air compared to larger particles and are not as strongly associated with other common indoor activities.

The third analytical layer involved multivariate regression modelling, a statistical approach that quantifies how much of the total measured exposure could be explained by cooking alone, after accounting for other factors. In these models, cooking events were treated as predictor

variables, alongside outdoor pollutant concentrations, ventilation rates, and markers for other indoor activities such as cleaning or burning candles.

By including all of these variables together, the model could separate the effects of cooking from those of other pollutant sources that might occur at the same time. The output of these models provided a percentage figure—an estimate of how much cooking contributed to the total measured exposure to each indoor air pollutant.

Finally, indoor-to-outdoor ratios were used as a consistency check. If pollutant levels indoors during cooking were much higher than outdoor levels at the same time, it strongly suggested that the source was indoors rather than infiltrating from outside. Conversely, if both indoor and outdoor levels rose together, the increase was more likely to be driven by outdoor conditions such as traffic or regional haze.

By integrating these four complementary approaches—time-resolved profiles, chemical signature analysis, statistical modelling, and indoor-to-outdoor comparisons—the study was able to build a robust, multi-faceted picture of cooking's contribution to indoor air pollutant exposure. This combination ensured that conclusions were based not on a single piece of evidence but on a coherent set of findings that converged on the same answer.

Exposure Metrics

Occupant exposure was quantified through cumulative exposure dose (CED) calculations for each targeted indoor air pollutant, allowing the study to capture both the intensity of exposure and the duration over which it occurred. The CED metric was derived by integrating the time-resolved pollutant concentration data, measured at one-minute intervals, with detailed time-activity information for each monitored zone within the apartment. This integration ensured that the calculations reflected the actual microenvironments in which occupants were present, rather than relying on a single-point concentration measurement that might misrepresent personal exposure.

For each apartment, the pollutant concentration data from the kitchen, living area, and bedroom were linked to the activity logs, which indicated when and where occupants were located. This meant, for example, that short-term peaks in ultrafine particles during cooking would contribute more heavily to the CED if occupants were present in the kitchen at the time, compared with the same peaks occurring when occupants were in a bedroom. By accounting for both pollutant levels and occupant location, the approach provided a more realistic measure of personal exposure than concentration data alone.

CEDs were calculated separately for periods identified as cooking and non-cooking based on the verified activity logs and stove-use sensor data. This separation allowed for a direct, quantitative comparison of the proportion of total exposure attributable to cooking versus all other sources combined. The resulting dataset provided a robust basis for testing the hypothesis that cooking was the dominant contributor to occupant exposure to the measured indoor air pollutants.

Statistical Analysis

The statistical analysis was designed to test the null hypothesis by comparing the proportion of total exposure to indoor air pollutants from cooking with that from all other sources combined. Data distribution was assessed using the Shapiro–Wilk test to determine whether parametric or non-parametric tests were appropriate. For normally distributed data, paired t-tests were conducted, while for non-normal distributions, the Wilcoxon signed-rank test was used. A significance level of $\alpha = 0.05$ was applied throughout the analysis.

To account for the repeated measurements taken within each apartment, mixed-effects models were employed. These models included fixed effects for cooking activity, ventilation rate, outdoor pollutant levels, and apartment characteristics, and a random effect for apartment identification. This approach accounted for within-apartment correlations and ensured that the results could be generalised beyond the study sample.

Quality Assurance and Control

High standards of quality assurance were applied consistently throughout the study to ensure that the data collected were both accurate and reliable. All monitoring instruments underwent full calibration against certified reference standards before being deployed to the field. This process ensured that the readings from each instrument matched known, accurate values. Once the monitoring period ended, the instruments were retrieved and calibrated again to confirm that their performance had not drifted during use.

To maintain accuracy during the study, weekly zero and span checks were conducted. A zero-check confirmed that an instrument read zero when no indoor air pollutant was present, while a span check confirmed that it produced the correct reading for a known concentration. These checks were essential for detecting any gradual changes—known as drift—in instrument sensitivity over time.

Duplicate sampling was performed in a subset of apartments, where two instruments of the same type measured the same parameter side by side. This provided a direct test of measurement consistency. Any differences greater than 5% between the two instruments triggered an investigation, and corrective action was taken before further data collection.

In addition, automated data-screening procedures were used to identify extreme values, such as sudden spikes, which could be caused by either unusual household activities or instrument errors. These flagged data points were manually reviewed and cross-checked with the activity logs to confirm whether they represented genuine events. Apartments with less than 90% valid data for any parameter were excluded from the final analysis to ensure the overall integrity and robustness of the dataset.

Ethical Considerations

Ethical approval for the study was granted by the appropriate institutional review board before any fieldwork commenced, confirming that the research adhered to established principles for the ethical treatment of human participants. Written informed consent was obtained from all participating households, following a thorough explanation of the study's aims, the monitoring

procedures to be used, and the types of data that would be collected. Participants were also informed about how their privacy would be protected, including the anonymisation of all personal identifiers before data analysis and the secure storage of all records.

The placement and operation of monitoring equipment were carefully designed to minimise any disruption to daily life. Instruments were compact, non-intrusive, and positioned in agreed-upon locations within the kitchen, living area, and bedroom. Maintenance activities, such as weekly checks and cartridge replacement, were scheduled at times convenient for participants.

The study team also made clear that no active interventions—such as altering ventilation practices or changing cooking routines—would be imposed during the Research Question 1 phase. This ensured that the recorded data reflected genuine household conditions and everyday occupant behaviour, thereby preserving the ecological validity of the findings while maintaining the highest ethical standards.

Contribution to Knowledge

The methodology developed for Research Question 1 is expected to contribute to knowledge by providing a detailed, empirically grounded understanding of the extent to which cooking contributes to occupant exposure to key indoor air pollutants in high-rise residential apartments. Unlike many previous investigations that have relied on short-term sampling, single-location measurements, or controlled laboratory simulations, this methodology has been structured to capture indoor air pollutant generation and exposure in real-world conditions over extended periods.

In doing so, it addresses the recognised gap between experimental findings obtained under simplified conditions and the complex variability present in occupied dwellings, where cooking practices, ventilation behaviour, building characteristics, and outdoor conditions interact in dynamic ways.

A key strength of this approach lies in the integration of high-temporal-resolution measurements with detailed time–activity tracking. By calculating cumulative exposure doses separately for cooking and non-cooking periods, the methodology extends beyond the conventional reporting of pollutant concentrations and instead quantifies the actual exposure dose received by occupants in different microenvironments within the dwelling. This distinction is critical for identifying the dominant contributors to exposure, as it incorporates not only the pollutant levels but also the amount of time individuals spend in specific locations.

The inclusion of concurrent outdoor measurements and the use of indoor-to-outdoor ratios enhance the robustness of source attribution by enabling the separation of pollutants generated within the home from those infiltrating from the external environment. The application of chemical signature analysis and multivariate regression modelling adds further analytical precision, allowing the isolation of cooking-related emissions from other indoor activities.

Overall, the methodology is anticipated to produce statistically robust, contextually relevant evidence regarding the significance of cooking as a source of exposure in high-rise apartments without kitchen exhaust systems. Such evidence will be valuable for informing ventilation design standards, urban residential building practices, and public health guidance aimed at reducing harmful indoor air pollutant exposure.

Methods for Research Question 2:

Premise for the Modelling Approach

The dataset generated in Research Question 1 formed the empirical foundation for model development, calibration, and validation. High-temporal-resolution measurements of targeted gaseous indoor air pollutants and particulate matter, collected concurrently indoors and outdoors, were integrated with detailed time-activity logs to provide source-resolved input data.

Cooking events identified in RQ1, together with measured ventilation rates, apartment geometry, and meteorological parameters, served as the primary inputs for each modelling approach. This ensured that the modelling exercise was anchored in realistic, context-specific emission and exposure profiles rather than relying solely on literature-derived or theoretical parameters.

Modelling Approaches

To predict the behaviour of indoor air pollutants during and after cooking in high-rise apartments, three complementary modelling approaches were applied. Each served as a form of “virtual laboratory” in which mathematics, physics, and measured data were combined to simulate pollutant emission, movement, and removal over time.

The mass-balance model was the simplest in structure yet highly effective for understanding the dynamics of pollutant build-up and clearance. In this model, the apartment was treated as a defined system, much like a water tank. Pollutants generated from cooking were analogous to water flowing into the tank through an inlet pipe. Removal processes — including ventilation, air leakage to the outside, and deposition of particles on indoor surfaces — acted like outlet pipes allowing water to leave the tank. The “water level” in this analogy represented pollutant concentration in the indoor air. If pollutant inflow exceeded removal, concentrations rose; if removal exceeded inflow, concentrations fell.

Mathematically, the behaviour of pollutants in the mass-balance model was expressed using the equation:

$$dc/dt = P/V - kC_t$$

Here, dc/dt is the rate of change of pollutant concentration over time, P is the pollutant emission rate based on direct measurements during cooking events, V is the effective volume of the indoor space which influences dilution (with smaller volumes changing concentration faster than larger ones), C_t is the pollutant concentration at a specific moment in time, and k is the total removal rate. This total removal rate combines several processes, including volumetric

ventilation expressed as air changes per hour (ACH), infiltration through building leakage paths, filtration efficiency where applicable, deposition of particles on indoor surfaces, and chemical or photochemical reactions that degrade pollutants in the air.

The reason why k is multiplied by C_t is that pollutant removal at any moment depends on the amount present. The greater the concentration, the larger the absolute amount removed per unit time for the same k value. This reflects real-world behaviour, where removal is proportional to the driving force provided by pollutant concentration. For example, strong ventilation clears a heavily polluted room faster at the start, but as concentrations drop, the rate of removal naturally slows because there is less pollutant to remove. In the equation, multiplying k by C_t mathematically captures this proportionality.

In simple terms, the equation calculates the change in concentration by subtracting what is removed from what is added. A high P value with a low k value means pollutants build up quickly and persist, while a high k value results in faster clearance. By running this model minute-by-minute with real measurements from RQ1, it was possible to estimate detailed exposure patterns, predict peak levels, and determine how long indoor air quality remained degraded after cooking.

The second approach, the computational fluid dynamics (CFD) model, offered a far more intricate and spatially resolved perspective, complementing the conceptual clarity of the mass-balance model with detailed three-dimensional simulations of air and pollutant movement. This model relied on the fundamental laws of fluid motion and heat transfer, solving the Navier–Stokes equations alongside species transport equations to track pollutant dispersion.

The exact apartment geometry was recreated from floor plans, incorporating wall and door locations, window positions, and the placement of cooking appliances. Airflow pathways were shaped not only by these structural elements but also by thermal plumes, the buoyant upward flows caused by cooking heat, which act as primary carriers for pollutants into adjacent zones.

The CFD model could resolve differences between near-field and far-field exposure, showing how pollutants might concentrate near the stove initially before dispersing into living areas. It also accounted for dynamic factors such as wind-driven cross-ventilation, temperature stratification, and interactions between natural and mechanical ventilation. This enabled the model to answer not only “how much” pollution was present but also “where” and “when” it was most likely to travel, accumulate, and persist, offering highly targeted insights for ventilation design and exposure mitigation strategies.

The third approach, the AI-enhanced hybrid model, merged the interpretability of physics-based modelling with the adaptability of data-driven learning. While the mass-balance and CFD models provided a rigorous physical foundation, they could not easily accommodate the abrupt, unpredictable variations common in lived-in environments — such as a resident suddenly opening a window, adjusting a fan speed, or changing cooking style mid-preparation.

The AI-enhanced hybrid model addressed this by incorporating machine learning algorithms trained on the RQ1 dataset, enabling it to detect and adapt to these real-world fluctuations. Its inputs included high-resolution pollutant concentration time series, detailed logs of cooking

events, meteorological data such as wind speed and direction, and behavioural indicators like window and door status changes.

The AI layer continuously compared predicted and observed outcomes, refining its internal parameters to improve predictive accuracy under variable conditions. Crucially, it could also infer behavioural tendencies — for example, a household's likelihood of opening windows during cooking under specific weather conditions — and incorporate these into its forecasts. This adaptive capacity made the hybrid model uniquely capable of delivering accurate, household-specific predictions over time, bridging the gap between theoretical simulation and the complex realities of high-rise apartment living.

By combining these three models, the study gained a comprehensive understanding of indoor air pollutant dynamics. The mass-balance model provided clarity on overall pollutant accumulation and removal, the CFD model revealed detailed airflow and spatial distribution, and the AI-enhanced hybrid model bridged the gap between theoretical physics and messy real-world variability. Together, they formed a robust analytical toolkit for evaluating and improving indoor air quality in high-rise residential environments.

Input Parameterisation

In this study, the reliability of the model outputs depended heavily on the accuracy and representativeness of the input parameters. Although each modelling approach had specific data requirements, all simulations drew upon measured values from RQ1 to ensure that they reflected actual indoor environmental conditions in the participating apartments. For the mass-balance model, the pollutant emission rate (P) was the key driving parameter. To determine P , cooking events were first identified in the RQ1 dataset. These events had to meet strict criteria, including clear start and end times verified through synchronised activity logs and corroborating pollutant concentration peaks from the monitoring instruments.

Once suitable cooking events were identified, the initial slope of the concentration–time curve ($\Delta C/\Delta t$) was calculated over a short time window immediately after cooking commenced. This period was deliberately chosen so that the influence of pollutant removal processes would be minimal, allowing the raw slope to approximate the unadjusted emission rate. However, because removal mechanisms were already active even in this early period, the raw $\Delta C/\Delta t$ was corrected using the total loss-rate constant (k) to ensure that the final emission rate reflected the true source strength from cooking alone.

The total loss-rate constant, k , encompassed all significant removal pathways for pollutants within the apartment. It was not limited to ventilation but also incorporated deposition, filtration, and chemical reaction rates where relevant. The ventilation contribution to k was expressed as the measured air change rate (ACH), representing volumetric ventilation in the form of Q/V , where Q is the outdoor airflow rate (m^3/h) and V is the effective indoor volume (m^3).

ACH values were determined through tracer gas decay tests for controlled conditions and through natural ventilation assessments for real-life conditions, carried out during the same periods as the cooking events. Multiplying ACH by the effective indoor volume yielded the volumetric ventilation rate, which was incorporated into k in the form Q/V .

Deposition rates accounted for the settling of particles onto indoor surfaces such as walls, floors, and furniture. These were calculated using either measured deposition velocities from the study or well-established literature values, multiplied by the total available surface area for deposition within each apartment. Although particle deposition is not instantaneous, it can be a significant removal pathway over the duration of cooking events and the subsequent decay period.

Filtration rates, though often negligible in apartments without active filtration devices, were included in scenarios where mechanical removal occurred via air-conditioning systems or portable air cleaners. These rates were calculated using the clean air delivery rate (CADR) of the devices, adjusted for their measured operational airflow rates and the filter's removal efficiency for the pollutant in question, normalised by the apartment's volume.

For reactive gases such as nitrogen dioxide (NO_2), chemical reaction rates were also integrated into k . NO_2 undergoes photochemical degradation in the presence of sunlight, and this process was modelled using experimentally derived rate constants that were adjusted based on the measured daylight intensity in each apartment. In certain cases, additional reaction pathways—such as ozone interacting with unsaturated organic compounds—were considered where supported by either direct measurements or relevant literature data.

After determining the individual contributions of ventilation (ACH), deposition, filtration, and chemical reactions, these components were summed to produce a composite k value for each pollutant in each cooking scenario. This composite k represented the overall removal capacity of the apartment for the pollutant of interest.

The adjustment of the raw $\Delta C/\Delta t$ to account for k involved adding back the amount of pollutant that was removed during the initial emission phase by these combined processes. This correction ensured that the emission rate used in the model represented the true source output without the masking effect of concurrent removal. The adjusted emission rates were expressed in micrograms per second ($\mu\text{g/s}$) for particulate matter and in parts per billion per second (ppb/s) for gaseous pollutants. These refined values were then used directly as P in the mass-balance equation.

By constructing k to reflect the full range of pollutant loss mechanisms—including ventilation expressed through ACH, surface deposition, filtration, and chemical reactivity—the model could reproduce pollutant decay patterns during and after cooking in a way that aligned closely with the observed RQ1 data. This comprehensive parameterisation process ensured that each input was both physically meaningful and representative of real-world indoor pollutant dynamics, providing a robust foundation for accurate simulation of emission, accumulation, and removal processes in the study apartments.

The dataset used for the mass balance modelling was deliberately chosen to reflect scenarios where cooking emissions overwhelmingly dominated the indoor air pollutant load. This selection ensured that contributions from other possible sources — such as outdoor infiltration, occupant activities unrelated to cooking, or off-gassing from materials — were negligible in comparison.

The reasoning is that mass balance modelling assumes that the major source term (P) represents the primary driver of concentration changes during the observation period. If multiple significant sources were present and unaccounted for, the model's fit and interpretation could be compromised because the unmeasured sources would effectively be treated as random noise, weakening the link between cooking and measured pollutant levels. By isolating periods when cooking was the clear dominant activity influencing indoor concentrations, you avoided introducing that kind of confounding effect.

This way, the model's source term for P could be meaningfully attributed to cooking, ventilation effects (expressed as ACH or Q/V) could be more precisely incorporated into the removal rate constant k , and the resulting estimates would be more valid for assessing cooking-related IAQ impacts.

In the CFD model, the process of replicating real indoor environmental conditions began with the transformation of each apartment's detailed floor plan into a high-resolution computational mesh. Each mesh cell served as a discrete calculation node for airflow dynamics and pollutant transport, allowing for spatially resolved simulations across the entire indoor domain. The mesh resolution was refined in zones of interest—particularly around cooking areas, windows, and known leakage paths—to capture rapid changes in velocity, turbulence, and concentration gradients.

Boundary conditions were derived directly from field measurements to ensure that simulations accurately reproduced actual airflow scenarios. Window and door configurations were documented precisely for each cooking event, including their degree of opening and the timing of any changes. Fan operations—whether ceiling-mounted, wall-mounted, or integrated into cooking hoods—were recorded along with their measured flow rates.

Outdoor conditions, such as wind speed and direction, vertical wind profiles, temperature gradients, and humidity levels, were imported from on-site meteorological stations or the nearest calibrated reference stations. These environmental inputs allowed the model to reproduce both buoyancy-driven flows, such as thermal plumes rising above active cooking surfaces, and wind-induced infiltration through micro-gaps around windows, doors, and facade joints.

Surface interaction parameters were also grounded in empirical data. Wall roughness coefficients were assigned based on the actual texture and material of interior surfaces, influencing the development of boundary layers and turbulence dissipation rates. Measured surface temperatures, obtained from infrared thermal imaging, informed heat transfer calculations and buoyancy effects within the CFD domain. These details also enabled accurate simulation of particle deposition patterns, including gravitational settling onto horizontal surfaces and turbulent impaction onto vertical walls. Stratification effects—particularly for buoyant gaseous pollutants like nitrogen dioxide—were modelled by incorporating both vertical temperature gradients and pollutant density differences relative to the surrounding air.

The AI-enhanced hybrid model built upon these physically grounded CFD insights but extended its scope to incorporate behavioural and contextual data alongside the physical parameters. High-frequency pollutant concentration time series from the RQ1 measurements formed the primary input stream, synchronised with a rich metadata set for each cooking event. Metadata included the cooking method—such as stir-frying, boiling, or deep-frying—categorised according to observed emission intensities; measured burner or stove heat output in watts; and the precise duration of the activity.

Contextual variables were encoded to capture environmental and human factors that influenced pollutant accumulation and removal. These included outdoor pollutant concentrations at the time of cooking, derived from both on-site measurements and local monitoring stations; indoor–outdoor temperature differentials; and occupant activity patterns, such as whether residents were stationary, moving between rooms, or opening and closing windows intermittently.

This combination of physical and contextual descriptors enabled the AI component to detect complex, non-linear relationships in pollutant dynamics. For example, it could distinguish the decay profile of a brief, high-heat stir-fry conducted under still outdoor conditions with all windows closed from that of an identical activity performed with multiple windows open on a breezy day.

The hybrid nature of the model meant that while CFD components handled the physically deterministic aspects of airflow and pollutant dispersion, the AI layer was able to identify patterns and modifiers that traditional physics-based models would struggle to capture, especially those arising from human behaviour and short-term environmental fluctuations. This dual-layer approach provided a nuanced, high-fidelity simulation framework capable of reflecting both the mechanical and behavioural realities of pollutant dynamics in the study apartments.

All three modelling frameworks also accounted for infiltration rates of outdoor pollutants, derived from simultaneous indoor–outdoor concentration measurements, ensuring that simulations did not underestimate background pollutant levels entering the apartments. By combining precise emission quantification, environment-specific boundary conditions, and nuanced behavioural data, the input parameterisation process created a robust foundation for predictive accuracy across the modelling suite.

Model Calibration and Validation

Before any model can be relied upon to predict real-world behaviour, it must go through two critical stages: calibration and validation. These processes ensure that the model is not only internally consistent but also accurately reflects the physical and behavioural realities it is designed to simulate.

Calibration is the stage where the model is fine-tuned so that its predictions align closely with actual observations. In this study, a carefully selected subset of apartments from RQ1 was used for calibration. These apartments were chosen to represent a wide range of layouts, sizes, ventilation behaviours, and cooking practices, so that the adjustments made during this

phase would be robust across different conditions. The process involved adjusting key parameters, such as pollutant removal rates, airflow pathways, and deposition characteristics, but always within scientifically defensible boundaries informed by measured data or well-established literature values.

For instance, if the model consistently predicted higher-than-measured peaks of PM_{2.5} during certain cooking events, the adjustment might involve refining the representation of removal through ventilation or deposition on indoor surfaces. Similarly, if the decay of NO₂ concentrations after cooking appeared too slow in the model compared to reality, the adjustment could involve improving the accuracy of infiltration or chemical reaction rates. The purpose of calibration was to achieve the closest possible match between the model outputs and the measured pollutant concentration profiles from the calibration dataset.

Validation came next, serving as the true test of the model's predictive capability. This step involved running the calibrated model on a completely different set of apartments from RQ1 that were deliberately excluded from the calibration process. By doing so, the model's ability to generalise to "new" situations—those it had never seen before—could be evaluated.

In this phase, model outputs such as concentration–time curves were compared directly to real measurements, with careful attention to whether the model captured the timing and magnitude of pollutant peaks during cooking, as well as the rate at which concentrations declined afterwards. Validation also tested the model's capacity to predict more aggregated metrics, ensuring that its utility extended beyond minute-by-minute fluctuations.

The performance evaluation considered both short-term and long-term accuracy. Short-term accuracy referred to the ability to correctly simulate the exact rise and fall of pollutant levels during a specific cooking event, such as a high-heat stir-fry or boiling session.

Long-term accuracy involved metrics like time-weighted averages, which represent the average pollutant concentration over a set period, and cumulative exposure doses, which measure the total amount of pollutant inhaled across the same timeframe. This dual focus ensured that the model could be applied both for detailed analysis of individual cooking events and for broader assessments of daily or weekly exposure patterns.

Through this structured process, calibration made the models closely replicate known conditions, while validation confirmed that they could accurately predict conditions they had never been tuned for. This combination provided strong confidence that the models were both realistic and reliable, making them suitable for use in scientific analysis and in informing strategies to reduce indoor pollutant exposure.

Performance Metrics

To evaluate the predictive accuracy of the models, a combination of statistical performance metrics was employed to capture both overall fit and event-specific behaviour. The root mean square error (RMSE) quantified the average magnitude of prediction errors, with larger discrepancies penalised more heavily. This was particularly useful for identifying situations where the model significantly deviated from observed values.

Complementing RMSE, the mean absolute percentage error (MAPE) expressed the error as a percentage of measured concentrations, enabling meaningful comparisons across pollutants with differing concentration ranges. The coefficient of determination (R^2) assessed how much of the variance in measured data was explained by the model, with values approaching 1.0 indicating strong alignment between predicted and observed trends.

In addition to these global indicators, the evaluation placed emphasis on the models' ability to accurately reproduce short-term dynamics that are critical for exposure assessment and ventilation design. This included predicting the magnitude and timing of pollutant peaks during cooking events—periods when exposure risk can be highest—and accurately modelling decay rates, representing how quickly concentrations returned to background levels after cooking ceased. These decay characteristics are particularly relevant for understanding pollutant clearance efficiency and informing the optimisation of ventilation strategies.

By combining broad statistical measures with targeted assessments of peak and decay performance, the evaluation ensured a comprehensive appraisal of model reliability. This approach confirmed not only whether the models matched observed data on average, but also whether they captured the transient concentration behaviours most relevant to health risk assessment and intervention planning.

Statistical Analysis

The predictive performance of the three modelling approaches was compared using repeated-measures analysis of variance (ANOVA). This statistical method was selected because it is well suited for situations in which multiple measurements are taken from the same setting, such as repeated pollutant readings from individual apartments. By accounting for the natural correlations between repeated measurements within the same environment, this approach ensured that any differences in model performance were not simply due to apartment-specific characteristics.

Following the overall comparison, post-hoc pairwise tests were conducted to determine which models differed significantly. The Bonferroni correction was applied to these comparisons to prevent inflation of Type I error rates, thereby reducing the chance of false positives when multiple comparisons were made.

A composite performance score was constructed to provide a balanced basis for model comparison. This score integrated several evaluation metrics: root mean square error (RMSE) and mean absolute percentage error (MAPE) to quantify prediction error magnitudes, the coefficient of determination (R^2) to assess explanatory power, and peak/decay performance measures to evaluate the accuracy of short-term pollutant concentration changes. These latter measures were particularly important for assessing potential exposure risks and informing ventilation strategies.

Statistical significance was set at $\alpha = 0.05$, indicating only a five per cent probability that observed differences arose by chance. Through this carefully controlled and statistically rigorous process, the ranking of the models was supported by both empirical evidence and practical relevance for indoor air quality prediction.

Quality Assurance and Control

Every stage of the modelling process was carefully checked to make sure the results were both accurate and physically realistic. First, all the data that went into the models were reviewed to ensure the units were correct, the timing of measurements lined up, and the values were scaled properly. This step removed the risk of errors caused by inconsistent or mismatched data.

For the CFD model, a process called grid independence testing was used. Think of the model's "grid" like a digital net covering the apartment, where each square in the net is a calculation point. The model was run with smaller and smaller squares until making them even smaller no longer changed the results. This confirmed that the predictions were not being distorted by the grid size.

For the AI-enhanced hybrid model, a process similar to testing a recipe in different kitchens was used. The data were split into several parts, and the model was trained on some parts while being tested on the others, rotating the test portion each time. This rotation process is formally known as *k-fold cross-validation*, where the dataset is divided into *k* equal parts ("folds"), and each fold takes a turn as the test set while the others are used for training. This ensures that the model learns patterns that generalise well and is not simply memorising the training data.

Finally, all model results were checked to make sure they made sense in the real world. For example, the simulated amount of pollutant could not suddenly appear or vanish without a source or removal process, and removal efficiency could not exceed what is physically possible. By combining careful data checks, stability tests, and reality checks, the modelling process ensured both scientific rigour and practical reliability.

Ethical Considerations

For RQ2, the modelling phase was conducted entirely using data already collected during the earlier measurement campaign, meaning no further direct contact with participants was required. All raw data from RQ1 were subjected to strict privacy safeguards before use, with every personal or location-specific identifier removed through anonymisation to ensure no participant could be linked to specific measurements. This process guaranteed that the subsequent modelling activities, including calibration, validation, and statistical comparison of model outputs, were performed on de-identified datasets.

Furthermore, results generated during RQ2 were presented only in aggregated or summarised form, avoiding any disclosure of individual apartment characteristics or occupant behaviours. These measures adhered to the ethical principles outlined in the study's approved protocol, ensuring that the use of RQ1 data in RQ2 respected participant confidentiality while enabling robust, scientifically valid analysis. By prioritising privacy and responsible data handling, the modelling phase of RQ2 maintained both ethical integrity and research credibility.

Contribution to Knowledge

The work undertaken in RQ2 advances the understanding of how delays in initiating IAQ diagnostics influence cumulative pollutant exposure, physiological health indicators, and cognitive-task performance. By integrating continuous real-world measurements from naturally ventilated high-rise apartments with advanced modelling techniques, RQ2 bridges the gap between observational IAQ research and predictive, decision-support modelling.

The comparative evaluation of three modelling approaches—CFD, AI-enhanced hybrid, and purely statistical—provides robust evidence on their relative strengths and limitations, particularly in predicting short-term exposure peaks and long-term cumulative exposure profiles. This enables the selection of modelling strategies that are most appropriate for different public health and building design contexts.

From a methodological perspective, RQ2 contributes a validated framework for integrating behavioural, contextual, and environmental variables into exposure modelling, moving beyond purely physical parameters. The inclusion of high-resolution cooking event metadata, coupled with AI-driven pattern recognition, demonstrates how non-linear pollutant behaviours can be captured in models, offering insights into real-world complexity that traditional methods often miss.

On a public health level, the findings establish how delayed diagnostics can prolong elevated exposure levels, increase physiological stress indicators, and accelerate cognitive decline trajectories. This provides a quantitative basis for policy and intervention strategies targeting early engagement in IAQ assessment and remediation, especially in high-density urban environments where occupant mobility is limited.

Finally, RQ2's approach to rigorous model calibration, validation, and statistical comparison contributes to the reproducibility and transferability of IAQ modelling research. By making methodological transparency a central principle, this work supports the wider adoption of scientifically credible models in policymaking, industry practice, and future research, thereby expanding the evidence base for effective IAQ risk management and prevention strategies.

Methods for Research Question 3:

Premise to Exhaust Ventilation System Design and Evaluation

The validated modelling framework developed in RQ2 was used to inform both the design specifications and the performance evaluation protocols for the new exhaust ventilation system. Insights from the best-performing model in RQ2 guided the identification of critical design parameters, such as capture efficiency requirements, airflow distribution patterns, and optimal exhaust placement for pollutant removal. Model-predicted exposure patterns under different kitchen layouts and occupant behaviours provided a quantitative basis for selecting filtration technologies, airflow rates, and housing configurations that could achieve substantial exposure reductions while remaining energy-efficient and space-efficient.

System Design and Prototype Development

The exhaust ventilation system was conceived as a next-generation indoor air quality solution tailored for high-rise apartments, where space is scarce, budgets are limited, and every decision about comfort, convenience, and cognitive awareness gained versus sacrificed must be as deliberate as the technical design itself. In such settings, solutions cannot simply be powerful; they must also be compact, easy to maintain, unobtrusive, and, above all, safe for both the user and the surrounding apartments.

A common flaw in traditional kitchen exhaust systems is their potential to push pollutants into shared ducting or leakage pathways, inadvertently degrading the indoor air quality of neighbouring homes. This design was therefore engineered to operate as a self-contained, zero-cross-contamination system, capturing pollutants completely within the unit before returning clean air into the apartment or venting it in a way that cannot re-enter another home.

Physically, the system is about the size of a small carry-on suitcase—roughly 55 cm tall, 35 cm wide, and 25 cm deep—with clean, rounded edges and a neutral matte finish that allows it to blend with most kitchen interiors. The housing is made from fire-retardant composite materials rated to withstand high cooking temperatures and sudden thermal spikes. The unit can be wall-mounted above the hob, attached beside it via a swivel bracket, or placed on the countertop using soft non-marking stabilisers. Its portability allows residents to reposition or store it when not in use, a feature particularly valued in kitchens where counter space is already at a premium.

The performance core is built on two complementary principles: targeted, high-efficiency pollutant capture directly at the cooking source and multi-stage filtration for both particulate and gaseous contaminants. The first stage is a stainless-steel grease pre-filter that traps oil droplets, preventing them from coating internal components or posing a fire hazard. This grease layer is removable and dishwasher-safe, making routine cleaning quick and cost-free.

The second stage is a high-efficiency particulate air (HEPA-grade) filter capable of capturing fine particles such as $PM_{2.5}$ and ultrafine $PM_{0.1}$ —pollutants associated with respiratory, cardiovascular, and cognitive health risks. The third stage is a gaseous pollutant filter, using a hybrid bed of activated carbon and chemisorptive media to remove harmful gases like nitrogen dioxide, formaldehyde, and carbon monoxide. These stages operate in sequence, ensuring that particles and gases are dealt with before air is returned to the breathing zone.

The airflow pathway is designed to act almost like a precision vacuum for emissions. Using computational fluid dynamics (CFD), the intake hood and ducting geometry were optimised to capture hot plumes and vapours at the very moment they rise from the cooking surface. By intercepting them early, the unit reduces the energy needed to treat them and prevents their spread throughout the kitchen. The fan, mounted in an acoustically dampened chamber, is capable of generating the static pressure required to pull air through dense filters without producing disruptive noise.

Embedded artificial intelligence elevates the system's performance beyond that of a conventional exhaust. Internal IAQ sensors continuously track particulate and gas concentrations in real time, feeding data into an adaptive learning model that recognises each

household's cooking patterns. For example, if the system detects a sudden spike in particulate levels typical of high-heat stir-frying, it will automatically increase airflow before visible smoke appears, thereby reducing the likelihood of triggering the apartment's smoke alarms. This anticipatory behaviour not only improves air quality but also plays a critical role in fire prevention—by reducing the accumulation of grease aerosols and particulates that could ignite under extreme conditions.

The fire safety considerations are embedded in several aspects of the design. The grease pre-filter is the first line of defence, intercepting flammable residues before they can settle on nearby surfaces. The fire-retardant housing and high-temperature-rated internal components are built to contain and withstand heat surges. The AI system can also detect abnormal heat patterns in the captured airflow, prompting an alert or an automatic shutdown if conditions resemble those of an unattended high-temperature event, reducing the risk of kitchen fires escalating.

Maintenance is designed for households without technical expertise. All filters are front-loaded and can be accessed by simply pressing a quick-release latch. No tools are required, and the process takes less than one minute. A simple, three-colour indicator ring communicates filter status at a glance—green for optimal performance, amber for attention soon, and red for immediate replacement. The replacement schedule is calculated dynamically based on actual pollutant load, preventing unnecessary costs and waste.

Ultimately, this exhaust ventilation system was designed to provide the highest possible healthy indoor air performance per unit of space, cost, and user effort. It ensures that one household's decision to improve its air quality never compromises the air quality—or safety—of others in the building, while maintaining the comfort, convenience, and cognitive ease that encourage daily use.

Integration and Installation in Field Sites

To evaluate real-world effectiveness under the very conditions for which it was conceived, the portable ventilation system was installed in a subset of high-rise apartments from the original RQ1 study cohort. This continuity ensured that pollutant data gathered under existing ventilation conditions could serve as a directly comparable baseline for post-installation performance, enabling a true before-and-after assessment.

Installation followed the design's guiding principles of minimal intrusion, portability, and prevention of cross-contamination to neighbouring apartments. No structural modifications, drilling, or attachment to shared ducting were required. Both configurations—wall-mounted using quick-release brackets and countertop with non-marking stabilisers—were deployed according to residents' preferences and kitchen layouts. The unit's compact, suitcase-sized form factor allowed placement directly above, besides, or slightly behind the cooking surface, ensuring optimised capture geometry as intended by the CFD-based design process.

For this study, the exhaust ventilation units were deployed without the need for professional installation or structural modification. Each unit was designed for plug-and-use operation and stabilised with soft supports, allowing residents to position it adjacent to the hob as instructed.

To ensure performance consistency across households, researchers conducted on-site verification. Airflow rates were checked with calibrated anemometers against design targets.

Performance validation was then conducted as a pre-study step, before long-term monitoring began: tracer gas was released separately under controlled conditions to benchmark airflow capture efficiency, and distinct cooking sessions were carried out afterwards to measure real pollutant removal. The tracer gas test provided a repeatable airflow benchmark, while the cooking test demonstrated actual reduction of cooking-related pollutants. Noise levels were also measured at different fan speeds to confirm compliance with comfort thresholds.

Only after these validation checks were completed were the units left in place for the actual field study, where pollutant monitoring proceeded solely under residents' normal cooking practices without additional tracer gas or test meals. This validation protocol ensured that system effectiveness could be established under controlled conditions before long-term deployment, while maintaining the requirement for no specialised installation so as to reflect realistic household feasibility.

Following the pre-study validation, the project progressed into the actual field study. The AI capabilities integrated during design proved valuable once the systems were deployed into homes. Upon installation, the units began learning each household's cooking patterns, automatically adjusting boost modes to intercept emissions early while minimising noise and energy use during low-emission phases. This reduced the cognitive effort required from residents, aligning with the design goal of supporting healthy indoor air without adding operational burden.

During installation, attention was also given to ensuring that maintenance tasks could be carried out easily by residents. The filter system was fitted with a colour-coded interface (green, amber, red) designed to guide replacement timing. Trial replacements conducted immediately after installation confirmed that the top-access, tool-free design allowed filters to be changed in under one minute. These demonstrations showed that the intended low-burden maintenance approach was practical for everyday users.

Similarly, the energy management features were activated at installation. The adaptive airflow control system was configured to scale output according to cooking intensity, ensuring low energy draw outside high-emission events. Initial monitoring checks verified that the units operated quietly and efficiently while maintaining proper capture at the source. At this stage, containment safeguards were also confirmed, with sensors in adjacent units showing no redistribution of pollutants into shared air pathways.

By the end of the installation phase, the system had been successfully deployed across diverse high-rise kitchen configurations. It demonstrated ease of integration, user-friendly maintenance, adaptive energy management, and robust pollutant containment. These early outcomes indicated that the portable, AI-enhanced ventilation system could be operated comfortably and conveniently from the outset, without compromising indoor or neighbouring air quality.

Field Measurement Protocol

Following the successful installation of the portable ventilation system, the next step was to establish a comprehensive field measurement protocol to examine how the system performed under real-world conditions. The protocol was designed to be scientifically robust, reproducible, and aligned with the pollutants most strongly associated with cooking emissions and their recognised health risks.

To this end, measurements focused on nitrogen dioxide, carbon monoxide, formaldehyde, and airborne particulate matter across three size fractions—PM₁₀, PM_{2.5}, and ultrafine PM_{0.1}. Together, these pollutants represented a mixture of gases and particles with known acute and chronic health implications, providing a holistic picture of the indoor air quality challenges posed by cooking.

In each apartment, the monitoring setup followed a carefully standardised procedure. Instruments were placed at breathing height, approximately 1.2 to 1.5 metres above the floor, and about one metre away from the cooking surface. This placement was intended to represent the air quality experienced by an individual cooking at the stove. Each instrument underwent calibration immediately before and immediately after the three-month measurement campaign, while daily zero and span checks were performed at the start of each monitoring day. The zeroing process was automatic, built into the instruments' internal diagnostic cycle, with the system executing self-zeroing every 24 hours without manual intervention.

These automatic checks ran independently of the measurement process, meaning that data collection temporarily paused during the zeroing cycle before resuming. These checks ensured stability and accuracy of readings throughout the study. Sampling was conducted at one-minute intervals, which offered sufficient temporal resolution to capture rapid emission peaks during cooking. Synchronisation of all devices to a common time reference allowed readings to be aligned precisely with logged cooking activities, ensuring that pollutant fluctuations could be accurately attributed to their sources.

The measurement protocol encompassed both controlled and uncontrolled cooking events. Controlled events were scheduled throughout the three-month campaign and were designed to generate standardised emissions that could be directly compared across households. A simple recipe, stir-fried vegetables in oil, was selected for this purpose, as it reliably produced both gaseous and particulate emissions under medium-high heat.

Each controlled trial was repeated at least three times in every apartment, spread across different days, to account for variations in background conditions such as outdoor air quality or residual pollutants from previous activities. In contrast, uncontrolled events reflected the diversity of daily cooking habits. Residents were free to prepare their usual meals without restrictions on style, duration, or ingredients. This dual approach enabled the study to capture both comparable baseline conditions and the complex reality of everyday life.

To support accurate contextualisation of the measurement data, residents were provided with a simple logging device. They recorded the start and end of each cooking activity and noted any additional events that might influence indoor air quality, such as lighting candles, using strong

cleaning agents, or opening windows. This supplementary record ensured that confounding factors could be accounted for when analysing pollutant levels.

The measurement campaign was structured to allow comparison with the baseline dataset collected under RQ1, before the installation of the prototype system. By maintaining the same pollutant suite, instrument positioning, and sampling protocols, it was possible to directly assess differences attributable to the intervention.

The monitoring phase after installation extended over three months, deliberately covering varying meteorological conditions and patterns of occupant behaviour. This timeline ensured that performance was not evaluated only under optimal or short-term conditions but under a representative spread of real-life scenarios, from hot, humid periods with stagnant air to cooler, breezier days.

To maintain scientific integrity, rigorous data quality control procedures were applied throughout. Raw datasets were screened daily, immediately after each 24-hour cycle of monitoring, for anomalies such as zero readings, sudden spikes, or implausibly high concentrations. Any suspect data were cross-checked against activity logs and maintenance records to determine whether they reflected genuine emission events or instrument errors.

Quality screening was performed remotely by the research team through secure data links, allowing anomalies to be identified and flagged while measurements were ongoing. Only validated data were carried forward into subsequent analysis. That is, the research team excluded any questionable or unreliable measurements from the dataset before running the main analyses.

While field data were central to assessing the prototype's impact, they inevitably reflected a limited set of scenarios. Cooking styles, apartment layouts, and behavioural responses varied, but it was not feasible to observe every possible combination within the three-month period.

To bridge this gap, the validated modelling framework developed in RQ2 was integrated into the evaluation. This hybrid methodology allowed field observations to be extended into a wider range of conditions that could not be practically captured in situ.

The modelling framework consisted of two complementary components: computational fluid dynamics (CFD) to simulate airflow and pollutant transport, and an artificial intelligence module that reproduced the adaptive behaviour of the ventilation system.

The field data informed the model inputs, including airflow rates measured with calibrated vane anemometers, pollutant capture efficiencies derived from comparing pre- and post-installation peaks and decay rates, and emission profiles generated from the logged cooking events. These emission profiles provided realistic timing and intensity distributions for different cooking practices, which were then used to drive the simulations.

Using these inputs, the modelling explored a set of systematically varied scenarios. Adjustments included different kitchen layouts. For example, the ventilation unit was placed either closer to or farther from the cooking surface. Variations in occupant behaviour were also

introduced. In some cases, residents activated the system before cooking; in others, they switched it on at the onset of cooking, or only after visible emissions appeared.

Changes in emission loads were also tested. Pollutant outputs were increased by 50 percent and 100 percent to simulate both more intensive cooking sessions and the infiltration of poor-quality outdoor air. Window-opening behaviour was also introduced as a variable, with runs simulating closed, partially open, and fully open conditions. Additionally, the system's predictive boost mode was tested in some scenarios, while in others it was restricted to reactive adjustments after pollutant thresholds were reached.

Validation of the modelling was performed by comparing outputs under field-matched conditions with the actual measurement data. Agreement within ten percent for pollutant peak reductions and decay times was considered satisfactory. Once validated, the model was used to forecast system performance under scenarios that extended beyond those observed directly during the field campaign. The outputs quantified reductions in peak concentrations, cumulative exposure, indoor-to-outdoor ratios, and energy use across the different scenarios.

By combining real-world field measurements with modelling grounded in those same datasets, the protocol provided a balanced and comprehensive assessment of the system's performance. It not only quantified the immediate effects of the intervention under real conditions but also offered predictive insights into situations not encountered during the trial. Importantly, the protocol was designed to be replicable.

With the protocol's standardised instrument placement, repeatable controlled cooking events, systematic logging of uncontrolled activities, and validated modelling framework, the methodology could be readily applied by other researchers in different cities or building contexts. This ensured that the study's contribution extended beyond a single dataset to a reproducible template for future investigations into indoor air quality improvement strategies in dense residential environments.

Evaluation Metrics and Statistical Analysis

The evaluation of the exhaust ventilation system was conducted after the installation and field monitoring phases, with the purpose of translating raw measurements and simulation outputs into meaningful indicators of performance.

Unlike the field measurement protocol, which described how data were collected and under what conditions, the evaluation stage concerned itself with the criteria used to interpret these data and the statistical techniques employed to ensure their reliability. This separation ensured that the process of gathering information and the framework for interpreting outcomes were treated as distinct but complementary components of the research.

The evaluation relied on a dual strategy. On the one hand, direct field measurements from the test apartments provided empirical evidence of how the system performed in practice. On the other, predictions from the validated simulation framework extended these findings to a broader

range of scenarios that could not feasibly be captured in the field. By combining both sources of evidence, the evaluation offered a comprehensive picture of immediate, real-world performance and projected long-term effectiveness.

Exposure-related metrics were placed at the centre of the evaluation framework because they spoke directly to the health protection objectives of the study. The first metric was the percentage reduction in peak pollutant concentrations during cooking events.

Peaks were given priority since they coincided with the highest risk of adverse health effects, particularly for vulnerable groups such as children, the elderly, and those with respiratory illness. The exhaust ventilation system was specifically designed to capture pollutants at their source, and a consistent reduction in these peak values was taken as evidence of its ability to mitigate short-term, high-intensity exposures.

The second metric was cumulative exposure dose, which was calculated by integrating pollutant concentrations over the full duration of cooking sessions. This measure captured both the intensity and the persistence of emissions, thereby reflecting the long-term pollutant burden likely to be inhaled by occupants. Whereas peak concentration reductions indicated protection against immediate risk, reductions in cumulative exposure demonstrated that the system delivered protective benefits throughout the entire cooking process, not just at moments of visible emissions.

To contextualise indoor results, indoor-to-outdoor (I/O) ratios were also calculated for each pollutant using simultaneous measurements taken during the monitoring phase. These ratios provided a benchmark for evaluating whether the ventilation system reduced indoor accumulation relative to outdoor background levels.

A post-installation decline in the I/O ratio indicated that pollutant concentrations indoors were no longer rising above those naturally present outside. In the high-rise context, this metric also demonstrated whether benefits within apartments were achieved without simply shifting emissions into shared shafts or neighbouring flats.

In addition to health-focused outcomes, operational and usability factors were evaluated to determine whether protective performance came at the expense of daily comfort or convenience. Noise levels were measured across all fan speeds during field deployment, ensuring that maximum operation remained below conversational thresholds. Energy consumption was recorded during both light and heavy cooking events, confirming that the AI-driven adaptive airflow reduced energy use by matching fan speed to pollutant load in real time.

Maintenance feasibility was also assessed after the monitoring period through timed filter replacement exercises and resident surveys. Participants were asked to replace filters without tools while being observed, and exertion was rated using a standardised physical effort scale. Cognitive load was assessed using a simplified NASA-TLX questionnaire. Replacements completed within one minute with minimal disruption to cooking confirmed that the system could be maintained without imposing burdensome effort.

To capture resident perspectives, surveys conducted after the three-month trial explored comfort, convenience, and perceived cognitive demands. Comfort was assessed by asking whether the unit interfered with movement or kitchen use. Convenience was judged by whether residents delayed or altered cooking to accommodate the system. Cognitive effort was probed by asking how much mental attention was required for daily operation. Consistent feedback highlighted that the AI's "set and forget" function reduced the need for active involvement, allowing the system to blend seamlessly into routine practice.

Economic efficiency was examined using a cost-effectiveness analysis performed after field data collection was completed. Purchase price, energy consumption, and expected filter replacement frequency were combined to calculate life-cycle cost per unit reduction in exposure. This metric demonstrated that significant health protection could be achieved without imposing high financial burdens, supporting the system's accessibility for households.

Statistical analysis was then applied to test whether observed improvements could be attributed to the intervention rather than random variation. Pre-installation baseline data were paired with post-installation results under matched cooking conditions. Where pollutant differences followed a normal distribution, paired t-tests were applied.

Where assumptions of normality did not hold, the Wilcoxon signed-rank test was used instead. To account for repeated measures within the same apartments and to adjust for covariates such as window opening behaviour, outdoor pollutant fluctuations, and natural ventilation variation, mixed-effects models were employed.

Throughout the evaluation, a significance threshold of $\alpha = 0.05$ was applied to minimise the likelihood of false conclusions. By anchoring the evaluation on health-focused exposure metrics while also incorporating usability, maintenance, effort, and cost-effectiveness considerations, the framework ensured that findings were both scientifically robust and practically meaningful. The results demonstrated that the exhaust ventilation system reduced pollutant exposures in a verifiable way while maintaining compatibility with household routines and economic constraints.

Ethical Considerations

For RQ3, strict ethical protocols were followed to safeguard participant rights, privacy, and comfort while ensuring the scientific integrity of the evaluation. Written informed consent was obtained from all participating households prior to the installation of the exhaust ventilation system. This consent process clearly explained the system's purpose, the type of monitoring to be conducted, and the nature of data to be collected, allowing participants to make fully informed decisions.

Installations and subsequent performance monitoring were scheduled in ways that caused minimal disruption to daily routines, with the system designed to operate unobtrusively in the kitchen environment. Participants retained the unconditional right to request removal of the system at any point during the study without penalty or impact on their relationship with the research team.

Data confidentiality was rigorously maintained by anonymising all identifiers before analysis, ensuring that individual apartments and occupants could not be linked to specific results. These measures collectively ensured that participation remained voluntary, non-intrusive, and respectful of household autonomy, while still allowing the collection of high-quality, real-world data on the system’s ability to reduce cooking-related pollutant exposure in high-rise apartments. The approach balanced ethical responsibility with the need for robust and repeatable scientific evaluation.

Contribution to Knowledge

The central contribution of RQ3 lies in the development of a cost-effective, installation-friendly and low-maintenance AI-enabled exhaust ventilation system specifically designed for high-rise residential kitchens. This system advances the state of the art by integrating multi-stage particulate and gaseous filtration, optimised exhaust capture geometry, and adaptive AI control that responds to real-time pollutant levels and cooking patterns. Unlike conventional systems, it prevents cross-contamination between neighbouring units in shared-duct environments, thereby addressing a long-standing gap in high-rise IAQ management.

The integration of predictive modelling with field measurements enabled the system to be pre-optimised for capture efficiency, airflow distribution, and energy use prior to installation. This ensured both high pollutant removal performance and operational feasibility.

By embedding user-centred evaluation criteria—such as maintenance accessibility, minimal cognitive demand, and comfort preservation—alongside traditional IAQ metrics, this research demonstrates how advanced technical performance can be achieved without compromising user adoption potential.

This work contributes a transferable, scientifically validated framework for designing and deploying intelligent exhaust ventilation systems in dense urban housing. It expands the global IAQ knowledge base and provides a replicable model for industry implementation and policy integration.

4

Research Findings

Findings for Research Question 1:

Introduction to Findings

The first research question investigated the pollutant emission characteristics of everyday cooking in high-rise residential apartments that lacked dedicated kitchen exhaust hoods. The aim was not only to measure pollutant concentrations during cooking but also to establish how these pollutants behaved in time and space, how they moved across different parts of the apartment, and how they compared with background outdoor pollution. By doing so, the study sought to quantify the role of cooking as a driver of indoor air pollution relative to both non-cooking periods and outdoor infiltration.

The pollutants of interest were chosen because of their relevance to combustion and human health: nitrogen dioxide (NO₂), carbon monoxide (CO), formaldehyde (HCHO), and particulate matter in three size fractions—PM_{0.1} (ultrafine particles), PM_{2.5} (fine particles), and PM₁₀ (coarse particles). These pollutants are known to cause or exacerbate respiratory and cardiovascular diseases and, in the case of formaldehyde, cancer. Carbon dioxide (CO₂) was also monitored, not as a direct health threat at measured concentrations but as a marker of both combustion and occupancy.

Data came from twenty apartments representing a diverse mix of floor levels, building ages, and ventilation practices. Each apartment was instrumented in the kitchen, living room, and a representative bedroom, while outdoor monitors provided baseline reference data. The analysis integrated time-resolved measurements, source attribution methods, indoor-to-outdoor (I/O) ratios, and cumulative exposure dose (CED) calculations.

The results provide the clearest evidence to date that in high-rise apartments without mechanical exhaust ventilation system, cooking is the dominant source of hazardous indoor air pollutants. This conclusion was supported not only by the direct pollutant measurements but also by the temporal association of emissions with meal preparation times, the chemical fingerprints of the pollutants, and the quantitative assessment of exposure burdens across different microenvironments.

Pollutant Concentrations and Temporal Patterns

The results demonstrated that cooking was the single most consistent and powerful determinant of indoor pollutant concentrations, with sharp, repeatable spikes occurring during meal preparation. These increases were not marginal but substantial, eclipsing both baseline levels and competing indoor sources.

Nitrogen dioxide rose almost instantaneously upon stove ignition. Concentrations in kitchens during cooking frequently exceeded 200 µg/m³, with extreme peaks above 350 µg/m³ during high-heat frying. By contrast, inter-meal baselines stabilised between 40 and 60 µg/m³, closely aligned with outdoor background levels. This temporal pattern strongly implicated direct stove combustion as the source. The magnitude of these excursions placed residents in short-term environments comparable to traffic-polluted kerbsides, but within their own homes.

Carbon monoxide followed a similar but less intense trajectory. Baseline concentrations rarely exceeded 1 ppm. Yet, during stir-frying or grilling, particularly in poorly ventilated dwellings, levels surged to 15–20 ppm. Although these spikes did not surpass acute toxicity thresholds, their chronic repetition introduced non-trivial risk. Households preparing multiple meals per day were therefore subject to cumulative CO exposures well beyond what would be predicted from background conditions.

Formaldehyde exhibited one of the most concerning patterns. Background concentrations in kitchens and adjoining living areas averaged 30–40 µg/m³, already approaching levels associated with mucosal irritation. Frying and sautéing elevated concentrations to over 100 µg/m³, with some extreme episodes reaching nearly 200 µg/m³.

Unlike NO₂ and CO, which decayed relatively quickly after cooking ended, formaldehyde showed slow clearance and frequently remained elevated for several hours, often persisting into sleeping hours in bedrooms. This persistence indicated that exposure extended well beyond the cooking period itself, imposing a prolonged burden even when occupants were at rest.

The most dramatic increases were observed for particulate matter. Ultrafine particle counts (PM_{0.1}) increased by up to two orders of magnitude within seconds of pan heating, with peak counts exceeding 1.5×10^6 particles/cm³. PM_{2.5} levels commonly rose to 300–500 µg/m³ during frying and grilling, magnitudes far above World Health Organisation (WHO) guideline limits. PM₁₀ followed a similar trend, though with somewhat lower peaks. These values placed cooking environments in the range of highly polluted urban hotspots, but on a recurring, daily basis within the home.

Carbon dioxide concentrations increased consistently during cooking, with typical rises of 400–800 ppm above baseline. These increases captured both combustion by-products and the presence of multiple occupants in kitchens, and they served as a useful corroborative marker of human activity consistent with logged cooking events.

Crucially, the pollutant spikes were not random anomalies but followed predictable rhythms. Breakfast, lunch, and dinner each generated distinct surges, with households experiencing multiple episodes every day. Although an individual spike often lasted only 20–40 minutes, the frequency and magnitude of these episodes resulted in cumulative exposures that dominated the daily indoor air pollution profile. Competing indoor activities such as candle use or cleaning produced minor emissions by comparison and did not significantly alter daily exposure landscapes. Outdoor infiltration was negligible during cooking events.

Taken together, these findings revealed a clear temporal and chemical narrative: cooking systematically reshaped the indoor air environment multiple times per day, producing exposure intensities that exceeded both background indoor conditions and other potential sources. The persistence of aldehydes and the extremity of particulate matter spikes highlighted cooking not only as a dominant source but as a recurring exposure hazard embedded in the daily routines of households.

Source Attribution Results

The attribution of pollutants to cooking as the principal source was not arrived at simply by contrasting indoor episodes with outdoor baselines. A systematic comparison was undertaken across all plausible sources, both external and internal to the apartment, to evaluate which activities were most responsible for measured occupant exposures.

Temporal alignment remained the strongest diagnostic tool, but here it was applied not just against outdoor monitoring but against the timing of household activities more broadly. Activities such as cleaning with chemical products, burning candles, or smoking were logged by participants. Each of these generated identifiable but comparatively muted pollutant signatures.

Cleaning events, for example, produced spikes in volatile organic compounds, particularly terpenes, yet these did not co-occur with the surges in nitrogen dioxide, carbon monoxide, and fine particles that marked cooking. Candle use produced a rise in ultrafine particle counts, but the magnitude was typically one order of magnitude smaller than those observed during cooking, and the chemical composition—lacking the nitrogen oxides characteristic of gas combustion—served to distinguish them clearly. In effect, competing indoor sources were detected, but their pollutant profiles were both chemically distinct and quantitatively weaker.

The fingerprint approach reinforced this hierarchy of contributions. The nitrogen dioxide-to-carbon monoxide ratio, previously described as diagnostic of stove combustion, was absent from other indoor activities. Candles emitted carbonaceous particles but negligible nitrogen dioxide. Cleaning events elevated volatile organics without altering the NO₂/CO relationship.

Even indoor smoking, in the few instances recorded, produced pollutant ratios dominated by fine and ultrafine particles with a very different oxidative signature compared to stove combustion. Thus, while these activities contributed episodically to the indoor pollutant mixture, none aligned in both timing and chemical ratio in the way that cooking did.

Outdoor infiltration, long understood as a baseline source of indoor exposure, was certainly present. Continuous monitoring showed that background indoor levels often mirrored diurnal traffic patterns. However, during cooking events the contribution of infiltration was dwarfed.

In quantitative terms, pollutant spikes attributable to cooking exceeded infiltration-derived increments by factors of five to ten for nitrogen dioxide and fine particles. Aldehyde ratios also diverged: infiltration carried in more acetaldehyde-rich signatures, while cooking episodes elevated formaldehyde disproportionately, a marker of heated oils and food degradation.

By examining indoor sources alongside outdoor contributions, the findings converged on a single conclusion: cooking was the dominant driver of occupant exposure. It not only generated pollutant magnitudes far exceeding outdoor infiltration but also overcame every competing indoor source in both scale and chemical signature. Candles, cleaning, and smoking did alter pollutant profiles, yet their contributions were smaller, shorter-lived, and chemically distinct. Outdoor infiltration provided a persistent background, but during cooking episodes, its role was effectively eclipsed.

What emerged, therefore, was a hierarchy of influence in which cooking sat unmistakably at the top. Other indoor and outdoor sources were not dismissed, but they functioned more as background context than primary drivers. In terms of practical exposure—those peaks that shape inhaled dose and therefore health risk—it was the stove, not the street, and not other household activities, that redefined the indoor atmosphere.

Spatial Distribution of Pollutants

One of the clearest findings was that pollutants generated during cooking did not remain confined to the kitchen. Instead, they moved rapidly throughout the apartment, creating a whole-home exposure environment. Within minutes of stove ignition, elevated levels of nitrogen

dioxide were detected in living rooms, demonstrating that combustion gases travelled quickly beyond the cooking zone.

Similarly, particulate matter spread with remarkable speed. PM_{0.1}, PM_{2.5}, and PM₁₀ concentrations in living rooms often reached 40–60 percent of kitchen peaks within just 5–10 minutes of active cooking. This pattern revealed that the boundaries between rooms in high-rise apartments offered little resistance to pollutant migration.

In open-plan kitchens, the situation was even more striking: living room concentrations frequently rose to levels closely mirroring the kitchen, often exceeding 70 percent of kitchen peaks. In directly adjacent open-plan living spaces, concentrations sometimes approached 95 percent of kitchen peaks, underscoring the limited dilution capacity of such designs. This demonstrated that residents outside the kitchen were frequently exposed to pollutant levels closely mirroring those at the emission source, highlighting the structural vulnerability of open-plan layouts to pollutant transfer.

Bedrooms were not exempt, even when doors were kept closed. In half of the monitored apartments, nitrogen dioxide levels in bedrooms rose above 100 µg/m³ within 30 minutes of cooking onset. PM_{0.1} and formaldehyde were even more troubling because they persisted long after cooking had ended.

While nitrogen dioxide and carbon monoxide concentrations began to decline relatively quickly once stoves were switched off, formaldehyde and ultrafine particles lingered for several hours, especially in enclosed spaces such as bedrooms. This persistence extended exposure well into periods of rest and sleep, a time when physiological repair processes are most active, raising particular concern for vulnerable populations such as children, the elderly, or those with respiratory illnesses.

The fact that non-cooking occupants were regularly exposed, even if they never entered the kitchen, underscored a critical finding: cooking is not merely a localised activity with limited risk, but a household-level determinant of indoor air quality. Children studying in living rooms, grandparents resting in bedrooms, or anyone spending time away from the kitchen were nevertheless exposed to cooking-related pollutants.

This spatial dimension of exposure challenges a common assumption—that good kitchen practices alone can control pollutant risks. The evidence showed that even with closed doors and avoidance of the kitchen, household members could not fully escape exposure. In high-rise apartments without effective exhaust ventilation systems, the dispersion of cooking pollutants transformed cooking emissions from a personal issue to a household-wide environmental hazard.

Although the present study focused on within-apartment dispersion, the findings raise broader questions about how cooking emissions may influence air quality beyond individual units. In high-rise buildings, pollutants generated in one flat can migrate through shared corridors, stairwells, and ventilation shafts, windows or infiltrate adjacent units through cracks and poorly

sealed partitions. Prior work has shown that nitrogen dioxide and fine particles are sufficiently mobile to cross these boundaries, suggesting that residents who do not cook at a given time may nonetheless be exposed to their neighbours' emissions.

While our methods were not designed to capture such inter-apartment transmission, the strong household-wide spread observed here implies that cross-unit exposure in multi-family buildings is a plausible and under-recognised risk pathway. Future research should extend monitoring to common areas and neighbouring flats to quantify this potential, as the design of effective interventions in dense housing will require understanding exposure not only at the household level but also at the building scale.

Indoor-to-Outdoor Comparisons

The comparison of indoor-to-outdoor (I/O) ratios provided a decisive way of distinguishing whether pollutants measured indoors originated primarily from outside infiltration or were generated within the apartment itself. Under non-cooking conditions, the ratios for nitrogen dioxide (NO₂) were close to 1, suggesting that outdoor air largely dictated indoor concentrations. In other words, when the stove was off, the apartment's background NO₂ levels were essentially a reflection of the surrounding environment seeping indoors through windows, doors, and cracks.

Cooking, however, fundamentally altered this relationship. During active use of gas burners, the kitchen I/O ratios for NO₂ climbed dramatically, reaching values of 3 to 6, while living rooms recorded ratios of 2 to 4 and bedrooms 1.5 to 3. These elevated numbers meant that concentrations inside were not merely passively tracking outdoor air but were exceeding them several-fold because of indoor combustion sources. The pattern was remarkably consistent across most apartments, regardless of orientation or height in the high-rise blocks, underscoring that indoor cooking systematically displaced outdoor air as the dominant driver of exposure.

For particulate matter (PM_{0.1}, PM_{2.5}, and PM₁₀), the evidence was even starker. While infiltration contributed to background levels, cooking events produced such intense spikes that indoor air far surpassed outdoor concentrations. Kitchens recorded I/O ratios as high as 10, with living rooms in the range of 5 to 7.

Ultrafine particles (PM_{0.1}), which are more numerous, smaller, and more easily absorbed deep into the lungs, exceeded outdoor concentrations by factors of 20 to 50. Such magnitudes reveal that indoor cooking was not only a significant source but the overwhelming contributor during meal preparation, transforming the apartment into a high-pollution environment within minutes.

Other pollutants followed similar trajectories. Formaldehyde, released from heated oils and foodstuffs, produced I/O ratios between 2 and 8, depending on ventilation rates and apartment layout. Carbon monoxide (CO), though generally staying below acute health thresholds, also exhibited cooking-period ratios well above unity, reinforcing its classification as an indoor generation product. Together, these data left little ambiguity: during cooking, indoor sources—not outdoor infiltration—dominated pollutant profiles across all major compounds.

The implications of these findings are profound. They challenge the assumption that reducing outdoor pollution alone would be sufficient to protect residents. Even if a city were to achieve dramatic reductions in traffic-related emissions or industrial smog, the I/O ratios demonstrated that individuals living in high-rise apartments would continue to experience high exposures whenever cooking took place indoors. The results therefore shift the burden of intervention away from outdoor environmental regulation alone and place it squarely on the realm of building design, ventilation strategy, and occupant practice.

This relationship can be understood through analogy. Cleaning the air outside the building is akin to reducing noise on a busy street, but if a loudspeaker is turned on inside the home, the external quiet makes little difference.

In much the same way, the dominance of indoor cooking emissions indicates that solutions must address the source and its spread within the apartment. Without measures such as effective kitchen exhaust ventilation systems, apartment layouts that impede pollutant migration, or adaptive behavioural practices, indoor exposures will remain dangerously high regardless of improvements in outdoor air quality.

Exposure Assessment and Variability Across Apartments

Cumulative exposure dose (CED) calculations provided a clear picture of how cooking shaped pollutant burdens within the studied apartments. On average, 60–75 percent of total nitrogen dioxide exposure, 50–65 percent of formaldehyde exposure, and more than 80 percent of particulate matter (PM_{2.5} and PM_{0.1}) could be directly attributed to cooking.

Households engaging in high-heat, oil-intensive methods consistently bore the heaviest burdens. In the most extreme cases, weekly exposures to particulate matter across all fractions — PM_{0.1}, PM_{2.5}, and PM₁₀ — surpassed levels typically expected from living adjacent to a busy roadway, underscoring the magnitude of indoor cooking as a pollution source.

Importantly, pollutant persistence meant that bedrooms, though not used for cooking, contributed disproportionately to cumulative exposures because concentrations remained elevated well after meals were prepared. This persistence illustrates that cooking is not a short-lived problem confined to the kitchen but one that permeates the entire living space and influences exposures over many hours.

The analysis also revealed inequities between households. Those cooking multiple meals daily, particularly when frying or grilling, accumulated substantially higher pollutant doses compared with households relying more on steaming or prepared foods. In this way, cooking frequency and style emerged not merely as lifestyle preferences but as determinants of exposure disparity.

The intersection of cooking-related exposures and socioeconomic patterns revealed quantifiable disparities in cumulative pollutant dose. Households in the lowest income tertile reported cooking frequencies averaging twice daily, compared with once daily or less in higher-income households.

This behavioural difference alone effectively doubled the number of peak exposure events. When combined with open-plan layouts common in low-cost apartments, peak PM_{2.5} and NO₂ concentrations in living areas frequently reached 70–90% of kitchen levels within 5–10 minutes, producing cumulative 24-hour average exposures that were 1.5 to 2 times higher than in more affluent households with enclosed kitchens and better ventilation.

Over a year, this exposure differential translated into an estimated 20–30% higher cumulative dose for PM_{2.5} and NO₂ in lower-income groups. These results indicate that socioeconomic context magnifies risk, not only through behaviour but also through architectural and infrastructural constraints, embedding pollution exposure within broader patterns of health inequality and environmental justice.

Despite these strong overall trends, variability across apartments was evident and shaped the severity of exposures. Floor level influenced the degree of outdoor infiltration, with lower floors registering slightly higher background levels. Yet once cooking began, the relative increases in indoor concentration were comparable across floors, underscoring the universal dominance of indoor generation.

Building age also mattered. Older, leakier apartments permitted pollutants to dissipate more rapidly, but this occurred at the expense of higher baseline infiltration of outdoor contaminants. By contrast, newer, airtight units trapped pollutants for longer periods, prolonging exposures even though they offered better protection from the outdoors. This paradox demonstrates the double-edged nature of airtight construction: beneficial in reducing outdoor penetration but detrimental when indoor sources dominate.

Occupant behaviour proved equally influential. Households that opened windows early during cooking experienced lower peak concentrations, while those that delayed ventilation until smoke or odour became obvious endured much higher levels. Even so, window opening alone was insufficient, since pollutants often migrated into adjacent rooms before dispersal. Bedrooms, in particular, remained affected long after cooking ended.

Cooking style and frequency further modulated these patterns. Oil-rich, high-heat methods consistently generated far greater levels of particulates and aldehydes compared with boiling or steaming. Households cooking three meals daily therefore faced cumulative burdens magnitudes higher than those cooking less frequently.

Together, these findings demonstrate that while the dominant role of cooking is consistent across households, the intensity of exposure is shaped by a complex interplay of building features and occupant practices. This complexity has practical implications: older, leak-prone buildings may require filtered exhaust systems to reduce outdoor infiltration while avoiding uncontrolled pollutant spread, whereas newer airtight apartments may need mechanical systems specifically designed to prevent long pollutant persistence.

Similarly, behavioural interventions—such as encouraging proactive ventilation practices and discouraging delayed responses—could significantly reduce cumulative exposures. By recognising how architecture, behaviour, and cooking style interact, interventions can be

tailored not only to reduce absolute exposure but also to address disparities in health risks across different segments of the population.

Statistical Outcomes

The inferential analyses provided decisive confirmation of cooking as the dominant driver of indoor pollutant dynamics. Mixed-effects models, which accounted for both within-building clustering and temporal variation, demonstrated that cooking was associated with sharp elevations in pollutant concentrations.

For nitrogen dioxide (NO₂), the mean increase attributable to cooking was 180 µg/m³ (95% CI: 160–200 µg/m³). Fine particulate matter (PM_{2.5}) rose by an average of 250 µg/m³ (95% CI: 220–280 µg/m³), while formaldehyde levels increased by approximately 80 µg/m³ (95% CI: 70–90 µg/m³). All associations remained highly significant ($p < 0.001$), even after adjusting for potential confounding variables such as outdoor pollutant concentrations, seasonal meteorological conditions, and background occupant activities.

Interaction analyses added nuance to these findings. Ventilation emerged as the only consistent moderating factor, lowering peak concentrations during cooking events by about 30 percent (95% CI: 25–35%). This moderating effect was statistically reliable but limited in magnitude; peak concentrations during cooking episodes still exceeded non-cooking baselines by wide margins, indicating that ventilation alleviated but did not resolve the exposure problem. Crucially, no significant interactions were observed with other building or household characteristics, underscoring the universality of the cooking effect across settings.

Robustness checks further reinforced these outcomes. Model diagnostics indicated excellent fit, with conditional R² values consistently above 0.70, confirming that the models captured the majority of variation in pollutant trajectories. Variance inflation factors (all < 2.0) excluded collinearity concerns, and residuals met assumptions of normality and homoscedasticity. When sensitivity tests excluded extreme peak events (top 5 percent of cooking-related values), the effect sizes attenuated modestly by 10–15 percent but remained statistically strong, demonstrating that the results were not dependent on rare outlier episodes.

The convergence of these statistical outcomes establishes cooking not merely as a significant predictor but as the overwhelmingly dominant determinant of indoor pollutant concentration spikes in the studied environments. The magnitude of the coefficients, the stability of results across adjustments and sensitivity analyses, and the persistence of effects despite natural ventilation all converge on a consistent conclusion.

Routine cooking activity generates exposure loads that cannot be explained by infiltration or incidental indoor processes. From a scientific standpoint, these results provide high-confidence evidence that interventions specifically targeting cooking-related emissions are not only warranted but imperative.

Conclusion on Findings for Research Question 1

The findings from this study provide a clear and compelling answer to Research Question 1, which asked the extent to which cooking contributes to measured occupant exposure to gaseous pollutants and particulate matter in residential apartments compared to other indoor and outdoor sources. The results consistently demonstrated that cooking is the predominant source of occupant exposure, with statistical analyses showing sharp, significant, and reproducible increases in pollutant concentrations during cooking episodes.

The purpose of the research, which was to determine whether cooking is the dominant contributor to exposure and to establish a quantified evidence base for intervention design, has been fully met. Across multiple pollutants—nitrogen dioxide, fine particulate matter, and formaldehyde—the magnitude and consistency of the elevations linked to cooking surpassed those attributable to outdoor infiltration or other incidental indoor activities.

The statistical outcomes confirmed that even when outdoor concentrations were elevated, the intensity of cooking-related peaks exceeded background and competing sources by wide margins. This finding supplies precisely the quantified evidence base required to inform targeted interventions that prioritise cooking emissions as a focal point for exposure reduction strategies.

The null hypothesis (H_{02}), which posited that cooking does not account for a significantly greater proportion of occupant exposure compared to other sources, was decisively rejected. The evidence demonstrated that the pollutant increments directly attributable to cooking were both statistically significant and practically large, with effect sizes far exceeding those observed for non-cooking events. Sensitivity analyses and model diagnostics further reinforced the robustness of these conclusions, showing that the outcomes were not driven by outliers or narrow contextual factors.

The alternative hypothesis (H_{12}), which predicted that cooking accounts for a significantly greater proportion of occupant exposure, was strongly supported. The combination of high effect magnitudes, statistically significant differences between cooking and non-cooking periods, and consistent patterns across all surveyed apartments confirmed that cooking is the primary determinant of indoor exposure to the targeted gaseous pollutants and particulate matter in residential apartments.

These findings establish beyond reasonable doubt that interventions aimed at mitigating indoor air pollution in such environments must prioritise cooking-related sources to achieve meaningful reductions in occupant exposure.

In sum, the research question, its purpose, and the associated hypotheses have been comprehensively addressed by the findings. Cooking has been shown not only to be a significant contributor but to stand as the dominant source of exposure, validating the direction of further research and intervention strategies centred on this activity.

Findings for Research Question 2:

Introduction to Findings

The findings from Research Question 2 present a detailed evaluation of how three distinct modelling approaches—mass-balance, computational fluid dynamics (CFD), and an AI-enhanced hybrid framework—performed in predicting indoor air pollutant dynamics during cooking in high-rise apartments. These results are grounded in the high-resolution dataset generated in Research Question 1, which provided the necessary inputs for emissions, ventilation, and behavioural variation.

The purpose of this analysis was to determine whether modelling could reproduce and predict pollutant emission, dispersion, and persistence with sufficient accuracy to serve as a reliable alternative to direct measurement. In doing so, the study aimed not only to assess predictive performance but also to examine the practical value of each modelling strategy for science, policy, and public health. Together, the findings demonstrate the complementary strengths and limitations of each model and establish their collective contribution to evidence-based indoor air quality management.

Mass-Balance Modelling Findings

When we applied the mass-balance model to the cooking events measured in Research Question 1, it reproduced the rise and fall of indoor pollution with high accuracy. The model used three things we actually measured in each apartment: how much pollution cooking released, how big the apartment was, and how quickly the air removed pollution through ventilation and surface losses.

With those inputs, the model correctly described about three-quarters to five-sixths of the changes we saw over time for nitrogen dioxide (NO₂), carbon monoxide (CO), fine particles (PM_{2.5}), and formaldehyde (HCHO). In plain terms, the curves the model drew looked very similar to what our instruments recorded.

For NO₂, gas cooking often pushed kitchen concentrations to roughly 200–300 micrograms per cubic metre, while the average for the whole flat sat around 120–180. The model matched both the build-up during cooking and the slow fall afterwards, correctly estimating that NO₂ typically took about 25–40 minutes to halve, depending on how well the space was ventilated. For CO, cooking raised levels to about 8–12 parts per million and they returned to normal in roughly an hour, which the model also captured well.

For PM_{2.5} from frying, kitchen peaks commonly reached 180–250 micrograms per cubic metre. That is many times higher than the World Health Organization's short-term guideline of 25. In tighter, more airtight apartments, the model showed these particles could remain above the guideline for 60–90 minutes after cooking stopped, which matched what we measured.

Formaldehyde (HCHO) behaved a little differently, and the model accounted for that. HCHO does not only come directly from cooking. It can also form indoors when cooking vapours react in the air and it can stick to surfaces and be released slowly afterwards. In our measurements, HCHO in kitchens often rose to about 60–120 micrograms per cubic metre during higher-heat cooking, sometimes crossing the WHO 30-minute guideline of 100. Whole-apartment averages were lower but could remain elevated for quite some time.

The model captured the main pattern: a clear rise during cooking followed by a slower-than- NO_2 decline. In tighter apartments, HCHO commonly stayed above the guideline for 30–60 minutes after cooking; in better-ventilated apartments, it dropped below the guideline in roughly 20–30 minutes once windows were opened or extract was used.

The model's predictions for HCHO were close to the measurements, although, as expected, the "long tail" caused by surfaces slowly releasing HCHO made the decline a little harder to match perfectly. Even so, the differences were small enough to be useful for decision-making.

A major strength of the model is that it estimates overall exposure well, not just the highest spike. When we added up how much $\text{PM}_{2.5}$ a person would breathe over the two hours after cooking, the model's totals were usually within about one-tenth of what our instruments showed.

For HCHO, the agreement was similarly good once we allowed the model to reflect that some pollutant sticks to and then leaves surfaces slowly. This is important because health risk depends on how much you breathe over time, not only on the single highest reading.

The model's main weakness is that it treats the apartment as if the air were evenly mixed. Real apartments are not like that. Near the stove, very short bursts can be two to three times higher than in a bedroom, and the model smooths those differences out. We saw this most clearly for ultrafine particles, where the model tended to miss the sharpest near-stove spikes.

The same limitation matters for HCHO as well. Someone cooking in the kitchen can breathe more than twice the pollution of someone who stays in a closed bedroom, but the model gives one average answer for both people. This is why we use this model for the big picture and pair it with more detailed methods or targeted measurements when we need room-by-room detail.

Even with that limitation, the model is very useful for testing what-if changes. Small, realistic steps made a big difference in the simulations and in real life. Opening two windows to reach roughly two air changes per hour brought $\text{PM}_{2.5}$ back under the guideline in 30–45 minutes instead of 60–90. The same action shortened the time HCHO stayed above its guideline, often bringing it down within 20–30 minutes rather than closer to an hour in tighter flats.

Increasing background leakage from very low to moderate levels nearly halved NO_2 exposure across the two hours after cooking. These improvements were also visible in our measurements, which gives confidence that the model can be used to test practical strategies before asking residents to change their habits or before investing in equipment.

In short, the mass-balance model, when fed with real measurements from each apartment, gave answers that were close to what actually happened for NO_2 , CO, $\text{PM}_{2.5}$ and HCHO. It explained most of the changes we saw, got the persistence times right, and produced exposure totals that were close to measured values.

It did smooth out the sharpest spikes and it could not separate kitchen from bedroom air, which is why we do not use it alone for room-specific advice. However, as a fast, reliable screening tool for population-level assessments and for testing the likely impact of ventilation and

window-opening practices, it remains a scientifically sound and practical choice.

CFD Modelling Findings

The CFD modelling provided critical insight into how the adoption of mitigation strategies like ventilation exacerbated exposure by revealing the spatial and temporal dynamics of pollutant transport within the studied apartments.

While mass-balance trajectories indicated a steady rise in average concentrations under diagnostic delays, CFD exposed the mechanisms by which pollutants accumulated and lingered, thereby clarifying why delayed or non-adoption of ventilation resulted in disproportionately higher cumulative doses.

In kitchens, simulations showed that during the first thirty minutes of cooking, PM_{2.5} concentrations in the breathing zone often exceeded 120 µg/m³, while nitrogen dioxide levels rose above 180 µg/m³ at stove height.

Under timely interventions — such as activating local ventilation or making simple behavioural adjustments — plume rise was curtailed and dilution promoted, keeping bedroom concentrations below critical thresholds. Without such measures, however, pollutants persisted for hours, spreading beyond the kitchen and contributing disproportionately to cumulative exposures across the apartment.

By contrast, when interventions were delayed, CFD demonstrated that hot plumes carried pollutants rapidly to ceiling planes, creating high-velocity transport pathways into adjacent rooms. This led to bedroom concentrations exceeding 100 µg/m³ of NO₂ within thirty minutes, even with doors closed, highlighting the critical role of buoyancy-driven transport in escalating cumulative exposure during intervention hesitation.

The simulations also revealed stagnation zones in corners and alcoves where fine particulate matter remained elevated long after cooking had ceased. These regions sustained exposure despite average concentration decay, helping explain why delayed interventions were associated with 20–30 percent longer clearance times in the empirical data.

CFD thus provided a mechanistic justification for the observed GEE (Generalised Estimating Equations) trajectories, where cumulative exposure dose rose more steeply and remained elevated for longer when mitigating strategies were not adopted.

The CFD results aligned with the field measurements on exposure metrics: spatial concentration fields, time-to-peak, and post-cooking decay. Across validation apartments, CFD reproduced the observed migration of pollutants from kitchens to living spaces and bedrooms within 10–15% error, including the rise of bedroom NO₂ above 100 µg/m³ within ~30 minutes despite closed doors.

These agreements confirm that CFD captured the dominant transport pathways (thermal plumes, lateral ceiling jets, and poorly mixed ‘dead zones’) that drive short-term exposure during and after cooking. The identification of ceiling-level pollutant “highways” and persistent

dead zones explained why even non-cooking occupants (e.g., children studying in bedrooms) experienced higher pollutant burdens, which corresponded to the more pronounced increases in pulmonary stress markers.

Importantly, CFD illuminated disparities across apartment layouts. In open-plan kitchens, delayed interventions led to faster whole-apartment contamination, while in compartmentalised layouts, pollutants concentrated along corridors and entered bedrooms more gradually but lingered longer.

These spatial patterns mirrored differences in cumulative exposure and health markers across households, demonstrating that the impact of intervention delay was contingent not only on timing but also on geometry-driven pollutant dynamics.

In sum, the CFD results provided spatially explicit evidence that delayed IAQ interventions amplified cumulative exposure by allowing pollutants to exploit buoyancy-driven flows and stagnation pockets before mitigation could be triggered. These mechanisms accounted for the higher cumulative doses.

AI-Enhanced Hybrid Model Findings

The AI-enhanced hybrid model demonstrated the highest adaptability and predictive strength of all the frameworks tested in this study. It achieved this by embedding physics-based constraints drawn from both the mass-balance and CFD approaches within a machine learning architecture, thereby combining mechanistic interpretability with data-driven flexibility.

This dual structure proved particularly well suited to the dynamic pollutant dynamics observed in high-rise apartments, where cooking emissions interacted with variable ventilation practices, occupant behaviours, and diverse apartment geometries.

Across the validation dataset, the hybrid model consistently achieved coefficients of determination (R^2) greater than 0.90 for nitrogen dioxide (NO_2), carbon monoxide (CO), and $\text{PM}_{2.5}$. This meant that it explained over 90 percent of the variance in measured concentration–time profiles, outperforming both mass-balance and CFD approaches. Importantly, the model captured not only the pollutant accumulation during cooking but also clearance trajectories, reproducing persistence times and decay slopes with minimal deviation from the empirical observations collected in Research Question 1.

A central strength of the hybrid framework was its ability to accommodate sudden, unplanned behavioural interventions that significantly influenced pollutant trajectories. For instance, in cases where residents opened windows partway through high-intensity frying, measured $\text{PM}_{2.5}$ concentrations dropped by roughly 30–40 percent.

The hybrid model closely replicated this behaviour, predicting a 35 percent decline within the same timeframe. Neither the mass-balance nor CFD models, which assume fixed boundary conditions, could reproduce such abrupt changes with equivalent precision. This responsiveness highlighted the advantage of incorporating behavioural and contextual data alongside physical structures.

The model also captured the diversity of cooking behaviours observed in the study apartments. Stir-frying, deep frying, and boiling each produced distinct emission intensities and persistence patterns, which were further shaped by occupant actions such as adjusting fan speeds or closing doors.

By integrating high-resolution concentration time series with detailed activity logs and contextual variables, the hybrid model generated realistic simulations across these varied conditions. This versatility demonstrated its capacity to generalise effectively across households rather than being restricted to a single cooking style or apartment configuration.

Despite its clear advantages, the hybrid approach displayed reduced accuracy in scenarios underrepresented within the training dataset. Prediction errors were larger in apartments using charcoal stoves or nonstandard ventilation practices, underscoring the model's reliance on representative and diverse training data.

This limitation reflects a broader challenge in machine learning applications, where predictive capacity is only as strong as the quality and coverage of the data used for training. Expanding datasets to include atypical practices will therefore be essential to ensure the robustness and generalisability of this approach in broader contexts.

The results also emphasised the practical utility of the hybrid framework for intervention design. Its ability to simulate pollutant reductions resulting from behavioural adjustments or ventilation changes in near real time makes it particularly promising for integration into smart indoor air quality management systems.

Such systems could deliver timely, household-specific recommendations—such as when to open windows, adjust fan speeds, or activate exhaust hoods—to help reduce exposures with minimal disruption to everyday routines. By bridging physical interpretability and empirical adaptability, the hybrid model provided a scientifically rigorous yet practical foundation for translating IAQ research into real-world applications.

Input Parameterisation Outcomes

The analysis demonstrated unequivocally that the reliability of the three modelling frameworks—mass-balance, CFD, and AI-hybrid—was fundamentally shaped by the quality of their input parameters. Input parameterisation proved not to be a technical afterthought but the decisive factor separating robust predictions from misleading artefacts.

When the models were supplied with parameters derived directly from the empirical dataset collected in Research Question 1, their performance aligned strikingly with observed pollutant dynamics across apartments. By contrast, sensitivity checks showed that poorly specified parameters quickly destabilised predictions, confirming that input accuracy was central to the validity of RQ2 outcomes.

This result reflected the richness of the RQ1 dataset, which supplied the raw material for all parameterisation. With minute-by-minute indoor and outdoor pollutant measurements, synchronised time–activity logs, and apartment-specific ventilation and geometry data,

parameters were not extrapolated from abstract assumptions but grounded in the lived conditions of residents cooking in high-rise apartments. This anchoring in real-world data underpinned the fidelity of subsequent simulations.

The first critical parameter set concerned emission rates, which dictated the pace of pollutant build-up indoors. These rates were quantified from the slopes of concentration–time curves at cooking onset, corrected for removal losses. The findings showed that gas burners released nitrogen dioxide (NO₂) at 15–40 mg/min, producing kitchen peaks of 200–300 µg/m³ within minutes. Carbon monoxide (CO) emissions averaged 2–5 g/hr, frequently elevating indoor concentrations above 10 ppm during intense cooking.

Particulate matter (PM_{2.5}) was highly variable, with stir-frying and deep frying producing rapid surges while boiling generated only modest increases. These empirically derived rates shaped the models differently: in mass-balance they defined the inflow term, in CFD they were inserted as stove-level point sources, and in the AI-hybrid they were encoded as structured features. In each case, accurate emission rates were essential to reproducing observed concentration peaks, reinforcing their central role in exposure modelling.

Ventilation emerged as the dominant factor controlling pollutant persistence. Measured air change rates (ACH) spanned from 0.25 h⁻¹ in tightly sealed apartments to 1.2 h⁻¹ in units with active window use. The contrast was stark. Apartments with ACH below 0.5 h⁻¹ retained elevated pollutants for nearly twice as long as those above 1.0 h⁻¹. Across all three models, higher ACH values consistently produced faster decay curves and shorter persistence times, in close alignment with empirical measurements.

This reproducibility confirmed ventilation as the overwhelming determinant of clearance, eclipsing other pathways such as deposition or chemical transformation. While secondary in importance, other removal processes also contributed measurably. Deposition accounted for 10–20% of PM_{2.5} removal depending on turbulence and surface availability, explaining why particle decay lagged behind gases under identical ventilation. For NO₂ and CO, deposition was negligible.

Filtration offered tangible benefits. Where portable air cleaners were in use, measured PM_{2.5} peaks were reduced by roughly 25%. This effect was accurately replicated by CFD and AI-hybrid simulations, confirming that targeted interventions could be realistically represented in models and have measurable benefits in practice.

Chemical and photochemical processes added further nuance. NO₂ photolysis contributed up to 10% of removal in sunlit rooms but was insignificant at night. Secondary reactions with ozone or unsaturated organics were included but contributed little over the short timescales studied. These findings underscored that while chemical processes can modulate pollutant clearance, they cannot compensate for low ventilation.

Each modelling framework incorporated these processes in line with its architecture. Mass-balance reduced them into a composite loss constant (k). CFD integrated emissions at stove height and represented airflow through boundary conditions. The AI-hybrid model embedded

the same parameters but allowed their weightings to shift dynamically depending on context, such as window opening or fluctuating outdoor wind speeds.

This adaptability proved decisive. The mass-balance and CFD models reproduced pollutant trajectories well under stable conditions but faltered when residents altered behaviour or when meteorology shifted mid-event. By contrast, the AI-hybrid model—trained on high-frequency pollutant and contextual data—was able to adjust in real time. For example, when residents opened a window during frying, it reproduced the observed 30–40% decline in PM_{2.5} with high fidelity, something the other models could not.

The combined findings converged on three robust conclusions. First, ventilation—expressed through ACH—was the overwhelming determinant of pollutant persistence. Apartments with low ACH retained NO₂ and PM_{2.5} above WHO guideline levels for 60–90 minutes post-cooking, while those with higher ACH cleared pollutants within 30–45 minutes.

Second, deposition, filtration, and chemistry contributed secondary but measurable effects, especially for PM_{2.5} removal and daylight NO₂ decay. Third, careful parameterisation was the essential factor that allowed models to reproduce real-world pollutant dynamics, transforming abstract frameworks into credible predictors of exposure.

The AI-hybrid model provided the clearest evidence of this principle. By embedding empirical parameters within a flexible learning structure, it adjusted to household-specific behaviours and meteorological fluctuations that the other models could not. Its ability to re-weight the relative importance of parameters in different contexts made it uniquely suited to addressing the central challenge of RQ2: predicting exposures accurately across diverse apartments, behaviours, and environmental conditions.

In conclusion, the results demonstrated that predictive accuracy in indoor air quality modelling is inseparable from the quality of input parameterisation. Emission rates, ventilation rates, deposition constants, filtration efficiencies, and chemical pathways collectively determined pollutant trajectories, but ventilation was the overriding factor. When these parameters were specified with precision from empirical data, the models—mass-balance, CFD, and especially the AI-hybrid—aligned with measured outcomes across diverse scenarios.

Most importantly, the study showed that parameterisation is not a minor methodological step but the cornerstone of modelling reliability. With poor specification, outputs risked being artefacts; with accurate inputs, the models became powerful tools for understanding and mitigating exposure.

The AI-enhanced hybrid model capitalised on this by embedding empirical inputs within a dynamic learning architecture, enabling it to outperform the other approaches in real-world, variable conditions. In doing so, RQ2 confirmed both the central role of ventilation in pollutant dynamics and the foundational importance of accurate parameterisation in making indoor air quality models actionable and credible.

Conclusion on Findings for Research Question 2

Research Question 2 set out to examine whether modelling approaches could predict cooking-related pollutant exposures in high-rise apartments more effectively than relying solely on direct measurement. The central hypotheses proposed two possibilities: that there would be no statistically meaningful differences in accuracy between mass-balance, computational fluid dynamics (CFD), and AI-enhanced hybrid models (H_{02}), or that at least one of these approaches would demonstrate superior predictive capacity (H_{12}).

The findings from this study strongly support the alternative hypothesis. Each model proved valuable, but their predictive capabilities diverged in ways that directly address the research question. The mass-balance model, though the simplest, captured the overall rise and decay of pollutants with commendable accuracy and provided useful estimates of exposure persistence.

It proved especially reliable for population-level assessments where computational efficiency and interpretability were priorities. Yet, its simplifying assumption of uniform mixing meant that it consistently underestimated sharp peaks and failed to capture near-field exposures.

The CFD model addressed this limitation by providing mechanistic insight into how pollutants move spatially within apartments. It revealed critical details such as thermal plumes above stoves, pollutant accumulation in dead zones, and differential exposures between rooms. These insights make CFD invaluable for ventilation design and exposure mapping, though its computational cost limits its routine use.

The AI-enhanced hybrid model emerged as the most versatile and accurate of the three. By integrating physical principles with adaptive machine learning, it explained more than ninety percent of the variation in observed concentration profiles and was uniquely capable of adjusting to abrupt, real-world behavioural changes, such as window opening or fan use. This adaptability made it far more effective at simulating lived realities than either physics-only model.

Taken together, the findings demonstrate that modelling approaches do not perform equally. Instead, the AI-hybrid model provides statistically, dynamic, and practically superior predictions, with CFD offering rich spatial insights and the mass-balance model delivering efficient, large-scale screening capability. This clear performance hierarchy provides decisive evidence against the null hypothesis and supports the alternative.

Beyond model comparison, the study highlights that accurate and adaptable modelling frameworks are indispensable for assessing and mitigating cooking-related exposures. They provide a cost-effective alternative to invasive monitoring, enable targeted interventions in vulnerable households, and offer a scientific foundation for healthier building design. In this way, Research Question 2 confirms that advanced modelling is not merely supplementary but essential for future indoor air quality management.

Findings for Research Question 3:

Introduction to Findings

Before introducing the system, baseline pollutant behaviour had already been characterised in detail in RQ1 and RQ2. Those studies demonstrated that cooking in high-rise apartments produced rapid surges in nitrogen dioxide (NO₂), carbon monoxide (CO), fine particulate matter (PM_{2.5}), ultrafine particles, and formaldehyde (HCHO).

Peak values regularly exceeded World Health Organization (WHO) guidelines, especially in poorly ventilated apartments with air change rates below 0.5 h⁻¹. RQ2 further showed that the predictive accuracy of mass-balance, CFD, and AI-hybrid models depended critically on accurate input parameterisation, with ventilation emerging as the single most influential determinant of pollutant clearance.

These baseline insights framed the context for RQ3: if pollutant persistence was fundamentally governed by ventilation, then introducing a targeted exhaust ventilation device should demonstrably shorten persistence times, lower cumulative exposures, and reduce the frequency of guideline exceedances.

Controlled Cooking Trials in Real-World Deployment

During controlled stir-fry events, the exhaust ventilation system consistently reduced PM_{0.1}, PM_{2.5}, and PM₁₀ peaks by 60–70% compared to baseline. To put this in perspective, PM_{0.1} (ultrafine particles, smaller than 0.1 micrometres) can cross into the bloodstream and affect multiple organs; PM_{2.5} (smaller than 2.5 micrometres) penetrates deep into the lungs, contributing to asthma, heart disease, and even cognitive decline; while PM₁₀ (smaller than 10 micrometres) irritates the upper airways, causing coughing and discomfort. The reductions achieved therefore represented not just a technical improvement but a meaningful health benefit, cutting exposure to pollutants across all fractions most relevant to human health.

In uncontrolled cooking situations, these particles often spike rapidly to levels far above what is considered safe. By lowering these peaks by more than half, the system demonstrated a substantial reduction in the highest and most dangerous exposures. This is particularly important because health risks are strongly associated with peak pollutant concentrations, not just with long-term averages.

In apartments where uncontrolled emissions had previously produced peak PM_{2.5} concentrations exceeding 250 µg/m³, post-installation peaks were consistently reduced to below 100 µg/m³ in most cases. To place this in context, the World Health Organisation's guideline for daily average PM_{2.5} exposure is 15 µg/m³. A peak of 250 µg/m³ therefore represents an extreme and hazardous condition — more than 15 times higher than what is considered safe. Even though post-installation levels still exceeded the guideline, the reduction marked a substantial improvement in protection, transforming conditions from acutely dangerous to far less harmful.

While the simulations focused on PM_{2.5}, field measurements indicated that PM_{0.1} and PM₁₀ followed the same downward trend, with substantial reductions observed across both ultrafine and coarse fractions. This pattern reinforced that the system's impact was not confined to a single size range but extended to the full spectrum of cooking-related particles most relevant to health.

The system also reduced the time required for pollutant concentrations to return to near-baseline levels. In apartments with ACH below 0.5 h^{-1} , where $\text{PM}_{2.5}$ typically persisted above WHO limits for 60–90 minutes, persistence was shortened to 20–30 minutes with the system active. This means that in poorly ventilated apartments—where pollutants could otherwise linger for over an hour after cooking—air quality was restored to safer levels in less than half an hour.

This reduction is critical because prolonged exposure, even at moderately elevated levels, can add substantially to residents' cumulative dose. By clearing pollutants more quickly, the system effectively shortened the window of risk, making indoor air safer for longer portions of the day.

These reductions were statistically significant ($p < 0.01$), supporting rejection of the null hypothesis. In scientific terms, a p-value less than 0.01 means there is less than a 1% probability that these reductions occurred by chance. This statistical strength adds confidence that the observed improvements were genuinely due to the system and not to random variability in cooking emissions or background air quality. For both researchers and policymakers, such significance provides assurance that these findings are robust and repeatable.

For NO_2 , baseline cooking episodes with gas burners often produced peaks of 200–300 $\mu\text{g}/\text{m}^3$, surpassing the WHO one-hour guideline of 200 $\mu\text{g}/\text{m}^3$. NO_2 , or nitrogen dioxide, is a gas strongly associated with respiratory irritation and asthma exacerbations. Gas cooking is a well-documented source of NO_2 indoors, and concentrations above the 200 $\mu\text{g}/\text{m}^3$ threshold can cause airway inflammation in sensitive groups.

The fact that uncontrolled cooking regularly crossed this guideline demonstrates the magnitude of the health challenge in high-rise kitchens. With the system operating, peak NO_2 values were reduced by 40–50%, and exceedances were substantially less frequent. In practice, this meant that instead of routinely exceeding the WHO guideline, most cooking sessions stayed within or just at the threshold, lowering the risk of acute respiratory irritation.

CO and formaldehyde followed similar patterns: CO concentrations, which often reached 10 ppm in baseline scenarios, were reduced by roughly 50%, while formaldehyde levels showed a 30–40% decline in both peak and persistence. Carbon monoxide (CO) is a colourless, odourless gas that can be life-threatening at high concentrations.

While the levels observed here were below acutely toxic thresholds, prolonged or repeated exposure at 10 ppm still presents significant risks, particularly for cardiovascular health. Halving this concentration is therefore an important step towards safer indoor environments. Formaldehyde (HCHO), on the other hand, is both an irritant and a recognised carcinogen.

A 30–40% reduction in its presence is particularly valuable because this pollutant tends to adhere to surfaces and re-release slowly, meaning that reducing its initial concentration directly reduces prolonged exposure. The observed decline in both its peaks and persistence indicates that the system addressed not only immediate risks but also the lingering background hazard formaldehyde represents in indoor air.

Uncontrolled Daily Cooking in Real-World Deployment

Long-term monitoring across three months captured the diverse range of meals and behaviours typical of everyday households. Unlike short, controlled trials that focus on a single recipe or fixed cooking conditions, this extended observation reflected the true diversity of real life.

Families cooked stir-fries, soups, fried meats, boiled rice, and baked dishes, often under unpredictable conditions—sometimes with fans running, sometimes without, sometimes with windows open, sometimes fully sealed. This gave researchers a realistic picture of how pollutants accumulate and clear in homes over time, and how the ventilation system performed under everyday pressures rather than laboratory scenarios.

Here too, the system consistently reduced cumulative exposure doses. Cumulative exposure refers to the total pollutant burden a person inhales over the entire cooking event, not just the highest spike. This measure is especially important for health because even moderate pollutant levels, if sustained for long durations, can add up to harmful doses. By lowering the total exposure across all cooking activities, the system demonstrated benefits not only for preventing short-term irritation, such as coughing or eye-watering during cooking, but also for reducing the long-term health risks associated with repeated daily exposure.

Integrating pollutant concentrations over entire cooking sessions, the intervention reduced exposures substantially across multiple pollutants: $PM_{0.1}$ by an average of 50%, $PM_{2.5}$ by 55%, PM_{10} by 45%, nitrogen dioxide (NO_2) by 40%, and carbon monoxide (CO) by 50%. To put this into real-world terms: if a resident would normally inhale the equivalent of 100 “units” of $PM_{2.5}$ during a cooking session, with the system in place they would inhale only about 45 units.

Comparable reductions applied to ultrafine $PM_{0.1}$, coarse PM_{10} , NO_2 , and CO, meaning that the system did not target a single pollutant but delivered broad protection across the spectrum of emissions generated by cooking. These were not abstract percentages but measurable safeguards — less fine particulate matter reaching the lungs, fewer ultrafines entering the bloodstream, and fewer gases impairing comfort, health, and concentration within the home.

That represents more than halving their exposure to fine particles that can enter the bloodstream and affect organs throughout the body. Similarly, a 40% reduction in NO_2 means far fewer exceedances of the WHO’s hourly guideline, lowering the risk of airway inflammation and asthma attacks.

Cutting CO exposures by half is particularly important because carbon monoxide binds to haemoglobin in the blood, reducing oxygen delivery to the body. Even at the levels seen here—well below acute poisoning thresholds—long-term, repeated exposures can put strain on the cardiovascular system. Thus, these reductions have practical, everyday health significance.

Importantly, these reductions held even when residents cooked with windows closed or in apartments with minimal natural ventilation. This finding matters because many high-rise apartments—particularly in colder months or in regions with outdoor pollution—operate with windows shut.

Under such conditions, pollutants can accumulate quickly and linger for hours. That the system continued to reduce pollutant exposure even in sealed conditions demonstrates that its benefits did not depend on residents making behavioural changes such as opening windows. Instead, the system itself provided the necessary air cleaning, making protection reliable and independent of resident action.

In such cases, the system's impact was particularly critical: without it, pollutants often remained elevated well into the post-cooking period, whereas with it, decay times were substantially shortened. In practical terms, this means that in a poorly ventilated apartment, a cooking session at 7:00 p.m. could leave pollutant levels elevated past 8:30 p.m., exposing residents—including children and the elderly—to unhealthy air long after cooking had finished.

With the system, however, pollutant concentrations dropped back toward safe levels within 20–30 minutes. This faster clearance not only improves immediate comfort (no lingering smoke or odour) but also significantly lowers cumulative exposure. Over weeks and months, this translates into much less pollutant burden for families, making a meaningful difference in long-term health outcomes.

Indoor-to-Outdoor Ratios

Baseline monitoring in RQ1 showed that cooking often elevated indoor pollutant concentrations far above outdoor background. In simpler terms, this means that when people cooked in their apartments, the amount of pollution inside the home often spiked to levels many times higher than what was present outside. Outdoor air, although not perfectly clean, usually serves as the comparison point. When indoor concentrations rise well above this outdoor “baseline,” it indicates that cooking is a major source of harmful pollutants.

Indoor-to-outdoor (I/O) ratios for PM_{2.5} commonly exceeded 4:1 during high-intensity frying. Because they are so small, they can penetrate deep into the lungs and even enter the bloodstream, making them particularly harmful to human health. An I/O ratio of 4:1 means that, during frying, the air inside the apartment was four times dirtier than the air outside. To put this into perspective, if the outdoor air measured 50 µg/m³ of PM_{2.5} the kitchen air could reach 200 µg/m³ or higher—levels associated with irritation of the eyes and lungs and, with repeated exposure, increased risks of chronic disease.

After the exhaust ventilation system installation, I/O ratios dropped substantially, often falling below 2:1. This reduction means that the system was able to cut indoor pollution roughly in half relative to the outdoor background. In many cases, instead of breathing air that was four times worse than outside, residents were now exposed to air only twice as polluted as outdoors—a dramatic improvement in terms of everyday health protection. For residents, this translates into fewer minutes and hours spent inhaling high concentrations of harmful particles, thereby lowering their cumulative exposure dose.

In some cases, ratios approached parity with outdoor levels, meaning indoor concentrations were no longer dramatically higher than the background environment. This finding is particularly significant because it suggests that the exhaust ventilation system, when

functioning optimally, could nearly eliminate the indoor pollution penalty associated with cooking.

In such cases, residents were essentially breathing air as clean as that outdoors, even while frying food indoors—a condition that would normally be impossible without active pollutant capture. For households with children, elderly residents, or individuals with respiratory conditions like asthma, this near-equalisation with outdoor air quality represents a substantial protective benefit.

Importantly, measurements in adjacent units showed no detectable increases, confirming that the system did not redistribute pollutants into neighbouring apartments. This addresses one of the biggest flaws in many traditional kitchen exhaust systems, which often push pollutants into shared ductwork or leakage pathways.

In high-rise buildings, this can mean that one household's attempt to remove pollutants inadvertently worsens the air quality of another household next door. By designing this system as self-contained and cross-contamination-proof, the study ensured that air cleaning benefited the target household without imposing hidden costs on neighbours.

This safeguard is especially important in dense urban housing, where dozens of families may share walls, ceilings, and ventilation pathways. The results here show that the exhaust ventilation system captured and treated pollutants entirely within the apartment, protecting residents not only from their own cooking emissions but also from contributing to a wider air quality problem in the building.

Simulation-Extended Findings

Field data necessarily captured only a subset of scenarios, but the validated hybrid modelling framework from RQ2 allowed findings to be extended across a wider design space. This means that while the apartment monitoring in the study provided detailed, real-world information, it was naturally limited to the specific cooking events and behaviours that actually occurred during the study period.

With simulation, however, the research team could “fill in the gaps” and explore additional situations—different placements of the device, unusual weather conditions, or varied occupant behaviours—that might not have happened during the three-month field campaign. In this way, the model acted like a scientific test-bed, extending the insights from the field into a much larger universe of possibilities.

Simulations showed that system effectiveness was sensitive to placement relative to the cooking source. Optimal capture occurred when the unit was positioned directly above or adjacent to the stove at less than 30 cm distance. This finding is intuitive once visualised: cooking emissions rise upward in a hot plume of gases and particles.

By placing the device close to this rising stream, the system could capture pollutants almost immediately as they left the pan or burner. At distances greater than 30 cm, the hot plume had already begun to spread and mix with surrounding air, making capture less efficient.

When placed farther away (e.g., on a side counter), capture efficiency fell by up to 20%. This loss highlights the importance of correct placement during use. However, even when residents positioned the system in less-than-ideal locations, it still provided major improvements compared to not using it at all, reducing pollutant exposure well below baseline levels.

Behavioural variability was also modelled. When residents activated the system before cooking began, capture efficiency was maximised, reducing peaks by up to 70%. In practice, this means that if the unit was turned on early—before visible smoke or odours developed—it was able to “stay ahead” of the pollution, intercepting particles and gases before they spread widely through the kitchen.

If activation was delayed until after visible emissions appeared, efficiency declined by ~15%, but reductions remained robust. For example, if a resident only switched the unit on once frying smoke became visible, the kitchen air would already be partly contaminated, making total reduction harder.

Nonetheless, the system still made a substantial difference by rapidly cleaning the air once activated. Importantly, the AI-enabled predictive boost mode, trained on RQ1 and RQ2 datasets, anticipated emission spikes and compensated for delayed activation, restoring much of the lost performance.

This feature effectively allowed the device to “learn” cooking patterns. For instance, if the AI detected the heat profile typical of frying, it could automatically increase airflow before the smoke became obvious to the human eye. This anticipatory behaviour protected households even when they forgot to switch the device on early, a key usability advantage for everyday life.

Outdoor conditions influenced natural ventilation but did not undermine the system’s performance. Because apartments are not sealed boxes, outdoor conditions like wind, temperature, and humidity affect how pollutants are diluted indoors. For example, windy days can increase infiltration, helping clear pollutants naturally, while stagnant days can trap them indoors.

In scenarios with high outdoor wind speeds, natural infiltration improved pollutant clearance, but the exhaust unit still contributed substantial incremental benefits. Even when nature was already helping, the system added an extra layer of protection, reducing peaks and shortening persistence. In stagnant conditions, where outdoor air exchange was negligible, the system’s contribution became dominant, preventing prolonged pollutant build-up.

This finding was particularly important because stagnant air is when residents are at greatest risk of exposure—pollutants linger for long periods without any natural ventilation. The system effectively acted as a safeguard, ensuring that harmful pollutants did not accumulate dangerously in such conditions.

Usability, Comfort, and Maintenance

Scientific performance was only one dimension of evaluation; usability was equally critical for ensuring adoption. In other words, no matter how effective the exhaust ventilation system was at reducing pollutants, it would only succeed in real homes if residents found it practical, comfortable, and easy to use. Research on household technologies consistently shows that devices which are noisy, energy-hungry, or difficult to maintain often end up underused, even when their technical benefits are clear. This principle guided the evaluation of user-centred factors.

Noise levels remained below conversational thresholds (<55 dB) even at maximum airflow, with most residents reporting the unit as “barely noticeable” during daily use. To put this into perspective, 55 dB is roughly equivalent to the background hum of a refrigerator or a quiet conversation at home.

Traditional kitchen exhaust fans often exceed 65–70 dB, which is loud enough to be distracting and, over time, discourages consistent use. By keeping noise levels comfortably low, the system avoided the common problem of residents turning devices off to reduce irritation, thereby ensuring continuous protection during cooking events.

Energy consumption was modest, averaging 0.25 kWh per day during typical cooking schedules, equivalent to the standby energy use of common household appliances. For comparison, this is similar to the daily energy consumption of a modern laptop in standby mode or the constant draw of a Wi-Fi router.

This figure demonstrated that the system’s health benefits were not achieved at the expense of household electricity bills. In an era where energy efficiency is both an economic and environmental priority, maintaining pollutant removal while using less power than a single light bulb helped ensure the system’s compatibility with sustainability goals.

Maintenance trials showed that residents could replace filters in under one minute without tools. Ease of filter replacement was essential because complex or messy maintenance tasks often discourage consistent upkeep. In many conventional ventilation systems, filters are hidden behind panels or require tools to access, leading to neglect or improper use.

The quick-release design used here meant that filters could be swapped as easily as changing a vacuum cleaner bag or a water jug cartridge. The colour-coded indicator ring effectively communicated replacement timing, and the AI system’s dynamic calculation of filter lifespan prevented unnecessary replacements.

This adaptive monitoring was particularly important because pollutant loads varied by household. For example, a family frying food daily would need to change filters more often than a household that mostly boiled or steamed food. By tailoring replacement alerts to actual usage rather than fixed schedules, the system avoided both wasted resources and lapses in protection.

Resident surveys confirmed that the maintenance burden was minimal, and compliance was high throughout the study. Survey responses indicated that residents felt confident managing the system without professional support, an important factor for ensuring widespread adoption

in high-rise apartments where access to technicians may be limited. High compliance also meant that the system maintained its designed efficiency across the three-month trial rather than deteriorating over time due to poor upkeep.

Surveys revealed strong acceptance. Most participants reported that the system did not interfere with cooking or kitchen use. This meant that the unit's compact size and unobtrusive positioning achieved their design goal: integrating into small, often crowded kitchens without restricting movement or workspace. Many noted that it prevented smoke alarms from triggering during high-heat cooking. False alarms are a common frustration in apartments, sometimes leading residents to deliberately disable smoke detectors—a serious safety risk.

By pre-emptively capturing emissions that might trigger alarms, the system improved not only comfort but also household safety. A majority also reported improved comfort and perceived air freshness, even when not consciously monitoring the device. This effect is significant because it shows that residents felt the benefits passively—without needing to constantly think about the technology. Improved perception of freshness also supports mental well-being, as indoor air quality is closely linked to perceptions of cleanliness and comfort.

These findings confirmed that the system's design goals of cognitive ease and low burden were achieved. The combination of quiet operation, low energy demand, simple maintenance, and seamless integration into cooking routines meant that residents did not view the system as an additional task, but as an unobtrusive ally in maintaining a healthier home. This outcome is critical, because long-term adoption of new technology depends not only on performance in the lab but on day-to-day compatibility with household life.

Affordability as a Key Benefit

While usability and comfort ensure a technology fits smoothly into daily life, affordability determines whether it can realistically be adopted at scale. The exhaust ventilation system was specifically engineered to be cost-effective, both in terms of purchase price and long-term operating costs.

A common problem with high-performance air quality technologies is that they are prohibitively expensive, limiting their accessibility to affluent households and leaving vulnerable populations—such as those in lower-income high-rise apartments—exposed to the highest risks. This project recognised from the outset that affordability was not an optional feature but a public health necessity.

Initial cost modelling suggested that the unit could be manufactured and sold at a price point significantly below traditional centralised exhaust systems, which often require costly installation and ongoing maintenance contracts. By avoiding integration with shared ducts and structural modifications, the system eliminated upfront expenses associated with construction, drilling, or specialist labour.

This plug-and-play design meant that households could achieve professional-grade pollutant reduction without the financial burden of hiring contractors or modifying their apartments—an especially critical feature for renters who lack permission to alter building infrastructure.

Operating costs were equally modest. Daily energy use, equivalent to a small household appliance, translated to only a few dollars per month even in regions with high electricity tariffs. Filter replacement, often a hidden cost in comparable systems, was kept low by combining high-capacity media with dynamic AI-based replacement alerts.

Instead of requiring households to replace filters on a rigid schedule—leading to premature disposal and unnecessary expense—the adaptive system ensured filters were changed only when performance actually declined. Over a year, this reduced filter expenses by an estimated 30–40% compared to fixed-schedule systems.

When these cost elements were combined—low purchase price, minimal installation costs, modest energy draw, and optimised filter replacement—the system demonstrated a life-cycle cost profile that was substantially lower than that of conventional exhaust fans with ductwork or stand-alone air cleaners with high filter turnover rates. Importantly, this affordability was achieved without sacrificing performance. The system consistently delivered 40–70% reductions in pollutant peaks and cumulative exposures while remaining cheaper to operate than less effective alternatives.

Resident feedback reinforced this finding. In post-trial surveys, many participants commented that they would be willing to purchase the system if it were commercially available, specifically citing its affordability compared to traditional solutions. For households already balancing limited budgets, the reassurance that improved air quality did not come with a heavy financial burden was central to their acceptance of the technology.

In this way, affordability became not only an economic benefit but also a driver of equity. By lowering barriers to adoption, the system made it possible for a wider range of households—including those most at risk of poor ventilation and high cooking-related pollutant exposures—to access the health protections it offered.

Conclusion on Findings for Research Question 3

The findings from RQ3 provide compelling evidence in support of the alternative hypothesis (H_{13}) and reject the null hypothesis (H_{03}). Both field measurements and validated modelling frameworks consistently demonstrated that the newly developed exhaust ventilation system significantly reduced occupant exposure to cooking-related pollutants in high-rise apartments compared to baseline conditions.

What is most striking is not simply that reductions were observed, but that they were substantial enough to narrow the performance gap between expected and actual conditions of high-rise living. The results show that exposures which previously exceeded health-based thresholds could be meaningfully controlled, even under varied layouts, occupant behaviours, and environmental conditions. In other words, the system addressed not only the statistical question of whether differences existed but the practical one of whether those differences could shift conditions towards safer, healthier, and more comfortable indoor environments.

The modelling work further confirmed that these benefits were robust, extending across diverse scenarios where conventional systems typically failed. The AI-enhanced boost mode ensured resilience against delayed activation, and the system consistently prevented prolonged pollutant build-up under stagnant conditions. These outcomes demonstrate that the intervention was not context-dependent but adaptable — a critical factor in establishing generalisability.

Beyond pollutant control, the system's affordability, low energy demand, and ease of maintenance fulfilled its design brief, showing that effectiveness need not come at the expense of usability. Residents reported satisfaction with comfort and operation, reinforcing that adoption barriers were minimal.

Taken together, these results affirm that a cost-effective, low-maintenance, advanced exhaust system can achieve verifiable reductions in exposure under real-world conditions. By fulfilling both technical and human-centred criteria, the evidence strongly supports acceptance of H₁₃: the system represents a reliable and scalable solution to one of the most persistent public health challenges in high-rise residential environments.

5

When Oluwajamisi completed his PhD, he carried forward not only a new technology but also a new way of thinking about indoor air quality in high-rise housing. His doctoral research showed that cooking was a dominant contributor to indoor exposure. It also highlighted the shortcomings of traditional modelling approaches. It culminated in the design and field testing of a cost- and maintenance-friendly, portable, AI-enhanced exhaust ventilation system.

Recognition of this work led to a postdoctoral fellowship with the Global Health and Housing Alliance, an international consortium headquartered at the University of Pompey that linked academic research with urban health practice worldwide. There, he explored a pressing question: what tangible benefits could his system deliver for human health and well-being when deployed in everyday residential life?

This postdoctoral stage marked a deliberate shift from technical development to systematic evaluation of impacts. Oluwajamisi designed multi-site studies across Jimpalia and Pacimany, enrolling households in high-rise apartments to document changes before and after the installation of his ventilation system. The research examined four dimensions: measurable health outcomes, reduced hospital visits, perceived indoor air quality, and cognitive performance in residential settings.

The findings were striking. Families who had previously reported frequent respiratory complaints — coughing, irritation of the eyes, and wheezing during cooking — experienced significant relief after adopting the system. Children in participating households showed fewer school absences due to respiratory illness, while adults reported reduced reliance on over-the-counter medication for breathing difficulties. Local clinics collaborating with the study documented a decline in minor respiratory visits among the participant group, providing evidence that the intervention reduced not only discomfort but also demand on healthcare services.

Equally significant was the change in perceived air quality. Residents described their apartments as fresher, less stuffy, and free of the lingering odours that had previously strained relationships between neighbours and housemates. Several participants noted that disputes over cooking smells, once a common source of tension in shared housing, declined markedly. The system's portability and quiet operation meant that residents began to use it consistently, reinforcing the habit of preventive action rather than resignation to discomfort.

Perhaps most revealing were the results relating to cognitive performance in the home environment. With the help of psychologists and education specialists, Oluwajamisi measured household members' ability to focus on daily activities within their apartments. Parents reported that they were able to concentrate better when working from home, while students noted greater ease in completing assignments and preparing for examinations.

Standardised cognitive tests administered in the apartments confirmed these observations: reduced indoor pollutant levels were associated with measurable improvements in attention, memory recall, and task efficiency. What had long been theorised in the literature — that exposures to nitrogen dioxide, formaldehyde, and ultrafine particles could impair cognition — was now demonstrated in lived residential environments. Families themselves spoke of the difference: “the air feels lighter, and studying no longer gives me headaches,” one student remarked.

The postdoctoral research transformed Oluwajamisi's ventilation system from a promising innovation into an intervention with proven, wide-ranging benefits. It provided evidence that linked technology directly to respiratory health, perceived comfort, and the ability to think and work effectively within the home — a setting where people spend most of their time.

Yet resistance remained. Developers in Jimpalia continued to argue that his work demanded unnecessary costs, while policymakers favoured downstream healthcare spending over preventive housing design. But Oluwajamisi now carried a new weapon: not only technical credibility but also evidence demonstrating that root-focused interventions were cheaper in the long run, reducing hospital visits and preserving human capacity.

The postdoctoral years therefore expanded his influence. His findings were published in leading journals and cited by health organisations as evidence for why indoor air quality must be addressed at its source. More importantly, they gave him the confidence and authority to step into his next role: building an independent research career that would train others to see, measure, and solve problems at the root.

Armed with compelling evidence from his postdoctoral studies — showing that his portable, AI-enhanced ventilation system could improve respiratory health, reduce hospital visits, enhance perceived indoor air quality, and even support better cognitive performance in residential apartments — Oluwajamisi entered his assistant professorship at the University of Pompey with renewed purpose.

The postdoctoral years had transformed his invention from a promising doctoral prototype into a solution with measurable impacts on daily life. His publications had begun to attract attention globally, and yet he recognised that the transition from research success to societal adoption

was never straightforward. The hurdles he now faced were not about proving effectiveness — that battle had been won — but about shifting behaviours, institutional priorities, and political will.

At Pompey, he established the Laboratory for Indoor Air Innovation, a facility he envisioned as more than a research space. It was to be a hub where engineering, health science, and social science converged to tackle the barriers that stood between root-cause solutions and their widespread use. Within this laboratory, he and his team devoted themselves to refining the AI-based ventilation system to ensure it was not only technically robust but also ready for the everyday realities of households across different contexts.

Cost-effectiveness was the first frontier. By partnering with local manufacturers, Oluwajamisi succeeded in driving down production costs without compromising performance, making the system viable not just as a luxury product but as a household appliance accessible to ordinary families. Maintenance was further simplified by the introduction of colour-coded filter cartridges that made replacement intuitive and effortless, removing one of the key reasons ventilation systems in the past had fallen into neglect.

The system's portability and ease of installation were enhanced by the design of a genuine plug-and-play model: no drilling, no ducting, and no specialist labour. Within minutes, households could set up the unit and position it as needed. By the end of his early professorial years, the device had matured into a credible consumer product that combined engineering sophistication with user-centred design.

Yet even as the system advanced, resistance persisted, and understanding its roots became part of his research identity. Unlike during his PhD, where resistance was largely academic, the barriers he now encountered were practical and cultural. For households, the problem lay in perception. Many residents underestimated the dangers of exposure to pollutants from cooking. Headaches, eye irritation, or a lingering haze in the living room were seen as minor discomforts rather than early signs of exposure to harmful gases and particles.

The connection between everyday symptoms and long-term risks like respiratory illness or diminished cognitive performance was not widely understood. As a result, some families dismissed the need for an advanced system, opting instead for makeshift strategies — opening windows, running table fans, or using air fresheners — without realising that such measures only displaced or masked pollutants rather than removing them.

In Pacimany, another form of resistance surfaced: financial hesitation. Although his system was cost-effective compared to traditional ducted hoods, even modest expenditure was difficult for families already stretched thin by rent, food, and school fees. Cultural attitudes of endurance also played a role. Many residents expressed the familiar sentiment of “we'll manage,” a phrase Oluwajamisi had grown up hearing in his own country. For them, tolerating discomfort was more natural than spending on prevention, even when the cost was within reach.

In Jimpalia, affluence created a different dynamic. Wealthier residents could afford high-quality medical care, and this accessibility reduced the sense of urgency to prevent health problems before they arose. Healthcare systems, though costly, were robust enough to cushion the

impact of neglect. This mindset meant that preventive measures like his system were undervalued, not because they were ineffective, but because the consequences of inaction could be managed downstream.

Policymakers and regulators added another layer of resistance. Without strong public demand, there was little political incentive to prioritise indoor air quality. Regulations technically allowed for ventilation via windows, which was enough to claim compliance. To regulators, strengthening codes would mean political confrontation with industry and potential backlash from property owners reluctant to bear even minimal responsibility.

Policymakers found it easier to allocate resources to visible healthcare investments — building new clinics, expanding hospitals — rather than champion preventive technologies that were less visible and harder to explain to voters. For many of them, prevention was abstract, while treatment was tangible.

In short, resistance persisted not because the system lacked performance, but because the risks of poor indoor air were underestimated and the benefits of prevention were consistently undervalued. This insight convinced Oluwajamisi that solving technical problems was only half the battle. The deeper challenge was cultural: changing how societies thought about health, risk, and responsibility.

This realisation sharpened his resolve and expanded his vision for his laboratory. He invited collaborations not only with engineers but also with public health scholars, sociologists, and policy researchers. His teaching reflected this interdisciplinary approach. He encouraged his students to master pollutant modelling, sensor design, and ventilation science, but he also urged them to step outside the laboratory and engage directly with the broader ecosystem of decision-makers. In one memorable lecture, he explained:

“To control indoor air pollutant build-up over time, its introduction rate into an indoor space must be reduced; otherwise, its concentration will rise quickly, especially in rooms with smaller volumes. The concentration can be managed by using a sink, which removes a portion of the pollutant from the air. This concept forms the basis of the equation $dc/dt = P/V - kC_t$, which mathematically expresses the rate of change of pollutant concentration (dc/dt), depending on emission rate (P), room volume (V), current concentration (C_t), and total removal (sink) rate (k).

If the portion of a pollutant removed (kC_t) from the indoor air is high but still less than the amount of the pollutant introduced (P/V) at a given time, the dc/dt will remain positive, meaning the pollutant concentration will continue to increase, albeit at a slower rate. The size of the room is crucial — smaller volumes accelerate the rise in dc/dt as the same pollutant amount occupies less air, increasing concentration more rapidly. Conversely, larger spaces offer greater buffering capacity, slowing pollutant build-up as the same amount disperses into a larger air volume.

Either way, in the absence of a removal mechanism, e.g., an exhaust hood or proper ventilation, dc/dt of a pollutant will increase fast and spread throughout the apartment. However, when kC_t exceeds P/V , the dc/dt of the pollutant will decline. If both are equal, the dc/dt equals zero, meaning no further increase or decrease in concentration over time.”

After the lecture, he had a discussion with a student.

[The student]: Our high-rise residential apartments lack exhaust hoods and have an open-concept kitchen layout. The ventilation provided by the windows does not seem sufficient during cooking. To achieve a noticeable impact, all windows must be opened with curtains drawn. However, privacy concerns—especially from adjacent apartments or visual intrusion—cause hesitation, significantly limiting the ventilation rate and reducing potential benefits.

[Prof. Oluwajamisi]: Additionally, many residents lack air cleaners and keep their windows closed because of air con use and polluted outdoor air, leading to a build-up of cooking-related pollutants indoors and increasing both exposure and health risks. In reality, much of the exposure to gaseous pollutants and particulate matter in residential apartments comes from cooking.

Oluwajamisi responded to his student, recognising that the lecture had prompted deep reflection. This was one of the cognitive abilities he hoped his students would develop in their learning journey. Other cognitive abilities he sought to cultivate were critical thinking, abstract reasoning, logical deduction, and creative imagination. He aimed for these abilities to enhance his students' knowledge, understanding, and skill development, both in relation to the subject matter and in their application to solving real-life problems in a value-oriented manner.

He designed his teaching activities around this goal. To strengthen these cognitive abilities in a value-oriented way, he tasked students with interviewing policymakers to uncover why regulations lagged, engaging with community groups to understand barriers to adoption, and debating with health professionals about the costs of preventable illnesses.

His lectures were remembered as transformative, not because he simplified technical material, but because he demanded that students connect theory to lived realities. By the end of his modules, many no longer saw indoor air quality merely as an engineering challenge; they recognised it as a question of health, society, and governance, requiring critical thinking, abstract reasoning, logical deduction, and creative imagination to solve effectively.

By the close of his assistant professorship, Oluwajamisi had achieved two things. First, he had advanced his system to the point where it was ready for large-scale adoption. Second, he had defined himself as a scholar who did not separate engineering from culture, or data from lived experience.

Though his system was proven, the culture of superficial fixes loomed large, slowing its adoption. Yet he recognised this not as defeat but as the true mission of his career: not only to refine a technology but to cultivate a generation of thinkers and practitioners capable of dismantling the barriers that kept societies from solving problems at their roots.

Oluwajamisi's promotion to Associate Professor was not simply a recognition of his technical innovations but of his persistence in confronting the cultural and institutional inertia that had slowed adoption of his system. He had already identified the roots of resistance. Households underestimated the risks of exposure, and some in Pacimany hesitated over even modest costs.

Affluent communities in Jimpalia treated prevention as unnecessary because they could rely on healthcare, and policymakers lacked incentives to strengthen regulations or support subsidies. As an Associate Professor, his research shifted focus: no longer was the primary challenge designing the system, but removing or reducing the barriers that kept people and institutions from embracing it.

His laboratory launched a programme of studies aimed at understanding household behaviour in greater depth. One stream of research examined risk perception in residential environments. Surveys, interviews, and ethnographic observations revealed how families normalised discomfort. A haze in the kitchen, a headache during cooking, or a lingering odour on clothes was rarely recognised as a health risk.

Oluwajamisi responded by designing evidence-based awareness interventions that combined air-quality monitoring data with real-time feedback. Households were given clear visualisations of pollutant levels during cooking and shown the difference once the system was activated. This experiential approach proved far more effective than abstract warnings: once families could “see” invisible pollutants, adoption and consistent use of the system increased dramatically.

A second line of research addressed the economic hesitation in Pacimany. Oluwajamisi collaborated with economists and public health scholars to conduct cost–benefit analyses at the household level. These studies quantified not only reduced medical expenditure but also gains in productivity and school attendance from improved air quality.

The findings showed that even low-income households ultimately saved money by using the system, as it prevented small but frequent health costs that accumulated over time. To make adoption easier, he experimented with micro-finance models and “pay-as-you-go” purchase schemes, turning the device from a one-time expense into a manageable household utility.

In Jimpalia, where resistance stemmed less from cost than from complacency, his research turned to framing prevention in terms of long-term national costs. He and his team modelled the cumulative healthcare burden of unmanaged indoor air exposure over decades, projecting increases in hospital admissions, medication demand, and productivity loss.

These studies showed that relying on a wealthy nation’s ability to “treat” illness was unsustainable: the economic burden of preventable disease eventually exceeded the cost of prevention many times over. By translating prevention into numbers that spoke to policymakers — budgets, national productivity, and economic competitiveness — he began to chip away at the notion that healthcare capacity was a substitute for preventive measures.

Finally, he confronted the policy and regulatory inertia identified during his assistant professorship. His team adapted the AI-hybrid models into city-scale policy simulation tools, enabling regulators to test the outcomes of different ventilation standards virtually.

For the first time, policymakers could see how requiring effective kitchen ventilation would alter exposure levels across thousands of households and reduce long-term healthcare costs. These simulations became difficult to dismiss because they demonstrated, in advance, both

the human and financial consequences of maintaining the status quo.

The effect of this body of research was cumulative. By showing households the invisible pollutants in their kitchens, by reframing costs as long-term savings, by quantifying national-level economic consequences, and by providing regulators with simulation tools, Oluwajamisi addressed resistance at every level. His approach was neither confrontational nor superficial; it was patient, evidence-based, and relentless in tracing problems to their roots.

By the end of his associate professorship, adoption of his system was expanding. In Pacimany, micro-finance schemes made the lower-cost version accessible to families once excluded, while public awareness campaigns tied to his research helped shift the cultural mindset from resignation to prevention.

In Jimpalia, although political inertia persisted, his economic modelling studies began to influence debates in parliament and gained attention in national media. Internationally, he was increasingly recognised as a scholar whose work bridged the gap between technology, behaviour, and governance.

The title of Associate Professor thus marked more than a promotion; it marked a new phase of influence. Oluwajamisi was no longer only an inventor but a strategist, dismantling the cultural and institutional barriers that had once seemed immovable. His research agenda had expanded beyond designing solutions to ensuring those solutions could be embraced, sustained, and scaled, a mission that defined his growing impact on the global stage.

By the time Oluwajamisi was appointed Professor, his career had entered a phase where technology, evidence, and philosophy converged. The Laboratory for Indoor Air Innovation, which he had founded as an Assistant Professor, had by then become a recognised hub for advancing practical solutions to indoor air challenges.

Under his leadership as an Associate Professor, the lab had expanded its scope — not only refining his ventilation system but also conducting behavioural and policy research to dismantle barriers to adoption. With his promotion, he transformed the laboratory into the Centre for Root-Cause Environmental Solutions, an interdisciplinary institute with a global mission.

The centre drew in engineers, health scientists, economists, sociologists, and policymakers, creating a collaborative space where environmental health problems were tackled at their origin rather than managed at their symptoms.

The centre's flagship project remained the refinement and dissemination of his portable, AI-enhanced ventilation system, but its scope broadened to include other environmental health issues such as mould proliferation in housing, indoor thermal stress in dense cities, and water leakage in high-rise estates. Each project was united by a single philosophy: problems must be addressed at their root if solutions are to endure.

Under his professorial leadership, the ventilation system reached global scale. In Jimpalia, after years of persistent advocacy and the accumulation of overwhelming evidence, effective kitchen ventilation was finally written into the national building code for high-rise housing. In Pacimany,

the low-cost version was embedded into public housing design, supported by international development grants and local micro-finance schemes.

Across regions, the system was adapted through collaborations with NGOs and governments, with modular features allowing for cultural and structural customisation while retaining its core principles of affordability, maintenance-friendliness, portability, and AI-enabled responsiveness.

Yet Oluwajamisi's role as Professor extended far beyond technology. He became widely known for his books, essays, and lectures that critiqued the persistent cultural tendency to settle for surface-level fixes. He argued that both poverty and wealth could blind societies to root causes: poverty by forcing people to endure discomfort as inevitable, and wealth by enabling them to rely on treatment instead of prevention. His words resonated not only in academia but in industry conferences, government assemblies, and community meetings. To him, sustainability was not defined by the adoption of the newest gadget, but by the courage to confront problems at their origin and design systems that prevented them from recurring.

Resistance did not disappear. Developers in Jimpalia continued lobbying to dilute codes, claiming added costs, while political leaders in Pacimany occasionally diverted funds from preventive housing measures to more visible projects that brought quicker public approval. These setbacks no longer surprised him. Instead, he reminded his students and colleagues that technological progress could be rapid, but cultural change was slow and required patience. The greatest struggle, he often remarked, was not with physics or engineering but with convenience and complacency.

Despite such resistance, his persistence reshaped global norms. His work was cited in World Health Organisation indoor air quality guidelines, framing preventive ventilation as a necessity rather than an optional luxury. International standards bodies adopted his findings as benchmarks. Most significantly, communities that once dismissed indoor air as a trivial matter began to see it as central to health, dignity, and everyday comfort.

By the time he had become Professor, Oluwajamisi's career embodied his deepest conviction: that enduring change was only possible when societies abandoned superficial fixes and tackled problems at their root. The Centre for Root-Cause Environmental Solutions became the symbol of this philosophy, extending his influence across disciplines and continents. His ventilation system stood as proof that a simple, well-designed intervention could reshape lives, while his teaching and advocacy transformed how societies thought about prevention, resilience, and human well-being.

6

For all his professional success, Oluwajamisi often reminded himself that a career was only one strand of a human life. The title of Professor, the publications, the centre he founded — these were visible markers of achievement. Yet what mattered most to him was what endured when the applause faded: family, community, faith, and the daily practice of living with integrity. To understand Oluwajamisi, one had to look not only at the lecture hall or the policy forum but at the quiet rooms where he made choices about love, sacrifice, and responsibility.

Oluwajamisi never forgot the narrow streets of Pacimany where he grew up. His parents had little, but they lived with dignity. His parents often said, “wealth without wisdom is a shadow,” a phrase that stayed with him long after he left home. When he first travelled to Jimpalia, the hardest moments of hunger and loneliness were softened by letters from his family, written in careful script, encouraging him to endure. Even when he had very little, he sent money back whenever he could.

Later, as his career flourished, he quietly paid the school fees of his younger sisters, insisting they should never be turned away from learning for lack of means. His own journey had been lifted by small scholarships and tutoring jobs; he was determined to widen that path for others. His sisters went on to become a fine artist, an accountant and an actress. They were notable and celebrated in their chosen professions.

During their MSc years, when money was scarce and each week felt like a battle for survival, he met Amira, a fellow international student from Maridia, a neighbouring country to his homeland, Pacimany. They often crossed paths in the university canteen, both opting for the cheapest meals of rice and beans. Their conversations, though simple at first, grew into a source of strength. They encouraged one another through the uncertainties of study and survival, and this companionship, grounded in faith and respect, slowly deepened into something more enduring.

Their marriage, a few years later, was modest but filled with joy. Even when their first home together was little more than a small flat with second-hand furniture, they created a space where warmth and laughter outshone scarcity. As Oluwajamisi often told friends, “Amira was my first co-author in life,” reminding him that success was not measured only by citations or grants but by the partnership that sustained his spirit.

Becoming a father brought him a new sense of responsibility. He often said that publishing a paper was rewarding, but teaching his children how to think was the most important scholarship of all. Around the dinner table, he explained his work in stories rather than equations. When his daughter once asked why he cared so much about air, he replied, “Because it is the one thing we cannot choose not to breathe. If it harms us, it harms everything we dream of becoming.” His children grew up not with the weight of expectation but with the freedom of curiosity. He encouraged them to ask questions, however simple or difficult, believing that curiosity was the seed of wisdom.

At home, he practised what he preached. Their apartment was always ventilated. Amira often joked that their children would grow up to be the only ones who associated air quality with bedtime stories. But behind the humour was a serious truth: for Oluwajamisi, the values of prevention, care, and dignity began first in the family space.

His name, Oluwajamisi, meaning “*the Lord made me realise*”, was more than an identity; it was a compass. To him, it served as a reminder that success must lead to awareness, not arrogance. Whenever he was tempted by recognition or the lure of prestige, he returned to the

memory of his early days in Jimpalia — eating sparingly, saving coins for bread, and enduring long nights of uncertainty. Those struggles, he believed, were among the ways the Lord made him realise what truly mattered: not wealth, but the courage to face problems at their root.

Faith, too, sustained him. Though he did not parade his beliefs, he prayed diligently several times each day, setting aside moments of reflection and gratitude that gave rhythm to his life. In the bleakest days of his MSc years — when his meals were meagre and his rent uncertain — these regular prayers steadied him, reminding him that hardship was never the end of the story. Later, as a husband and father, he preserved the same discipline. His children grew up watching him pause from the busyness of lectures, research, and writing to kneel in quiet devotion. For them, his example was as instructive as any of his lectures: true strength lay not only in intellect or persistence but in humility, discipline, and gratitude.

He often reflected that engineering could solve many technical problems, but faith and conscience were required to solve the problem of human pride. His humility was genuine, never staged; colleagues often remarked that he spoke to janitors and university presidents with the same calm respect. To Oluwajamisi, daily prayer was not ritual for its own sake but a way of guarding his conscience, a reminder that every crown, every success, must rest on integrity.

As his career grew, Oluwajamisi used his influence to give back in tangible ways. Each year, he returned to Pacimany for at least a few weeks, visiting schools and mentoring young people who, like him, had been told that their dreams were too big for their circumstances. He set up scholarships for talented but underprivileged students, never naming them after himself but after his parents, who had taught him that dignity mattered more than wealth.

In Jimpalia, he volunteered with immigrant communities, offering workshops on indoor health in multiple languages. He knew what it meant to arrive in a foreign land with little more than determination, and he wanted newcomers to feel that they, too, had a place. Students sometimes asked him why he spent so much time on outreach when his research was already world-renowned. His answer was simple: “If my work cannot touch ordinary lives, then it is only half done.”

Yet his personal life was not free from difficulty. The demands of travel, teaching, and constant advocacy often weighed heavily on him. He missed family milestones, and there were times when his children resented his long absences. Amira, ever patient, reminded him that balance was part of wisdom. Slowly, he learnt to decline certain invitations and to guard weekends for family time. These choices, though small, reminded him that personal integrity required balance as much as brilliance.

Looking back, Oluwajamisi often reflected that his greatest legacy was not the ventilation system, the modelling tools, or even the centre he founded, but the values he lived each day: humility, perseverance, curiosity, and service. He had transformed hardship into purpose, loneliness into empathy, and personal struggle into global advocacy. His children grew up learning that knowledge was not a trophy to display but a tool to solve problems at their root.

His students discovered that equations were not abstractions but lenses through which to recognise both injustice and possibility. And his colleagues came to respect him not only as a scholar but as a man whose personal life was inseparable from his professional mission.

For Oluwajamisi, the Lord had indeed made him realise — not realise his own greatness, but the deeper truth that dignity lay in serving others, beginning with family and extending outward to communities and nations. Professional success gave him a platform; personal life gave him a foundation. Together, they told the story of a man who endured, who realised, and who gave — not because it was easy, but because it was right. **The End!**