Association between indoor air quality and life expectancy: how the air we inhale indoors affects how long we live

indooraircartoon.com/2025/11/28/association-between-indoor-air-quality-and-life-expectancy-how-the-air-we-inhale-indoors-affects-how-long-we-live

28 November 2025

Indoor Air Cartoon Journal, November 2025, Volume 8, #172

[Cite as: Fadeyi MO (2025). Association between indoor air quality and life expectancy: how the air we inhale indoors affects how long we live. *Indoor Air Cartoon Journal*, November 2025, Volume 8, #172.]

Indoor Air Cartoon
November 2025
Vol. 8
#172
Artistic-Educational Research

ASSOCIATION BETWEEN INDOOR AIR QUALITY AND LIFE EXPECTANCY: HOW THE AIR WE INHALE INDOORS AFFECTS HOW LONG WE LIVE

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

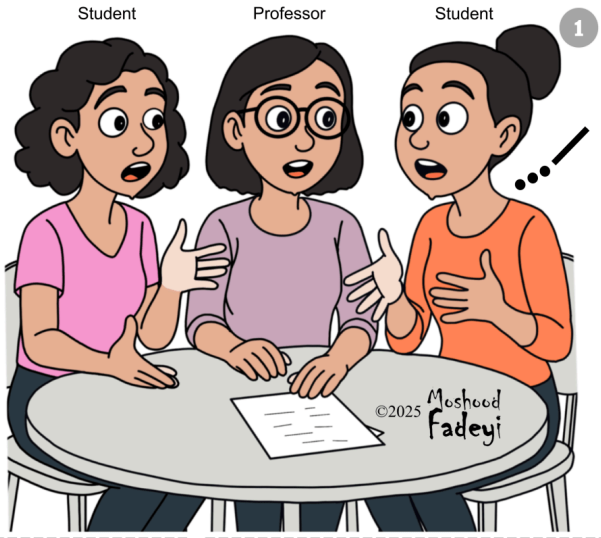

1 To add to your point, Ayanfe, my country has one of the lowest life expectancies in the world. People often talk about poverty, diet, or healthcare, but almost no one considers the air we breathe indoors. Yet the truth is that the quality of indoor air affects everyone, everywhere. It is not just a problem of poor countries; even in wealthier cities, sealed buildings, synthetic materials, and poor ventilation quietly harm health.

Additionally, polluted outdoor air seeps indoors, mixing with fumes from cooking, cleaning, and building materials, creating an invisible burden we often ignore. I only realised how the invisible air around us shapes how long we live when I learnt about indoor air quality from Professor's class.

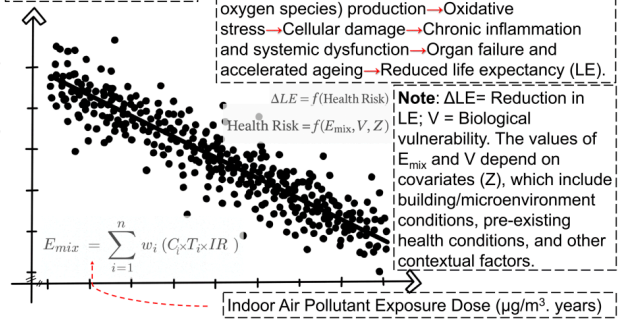

Logic: Indoor air pollutants→↑ ROS (Reactive oxygen species) production→Oxidative stress→Cellular damage→Chronic inflammation and systemic dysfunction→Organ failure and accelerated ageing→Reduced life expectancy (LE).

Note: ΔLE= Reduction in LE; V = Biological vulnerability. The values of E_{mix} and V depend on covariates (Z), which include building/microenvironment conditions, pre-existing health conditions, and other contextual factors.

Note: E_{mix}= Total cumulative exposure burden over the years for the mixture of all indoor air pollutants; C_i= average indoor concentration of pollutant I; T_i = exposure duration to pollutant i w_i = weighting factor reflecting toxicity or relative risk contribution of pollutant I; *IR*= inhalation rate.

2 You are right! Many underestimate how the continuous inhalation of polluted air increases ROS beyond healthy levels, causing oxidative stress that damages cells. This triggers chronic inflammation, disrupts body systems, and accelerates ageing while heightening biological vulnerability to other hazardous factors. Over time, organ function declines, leading to compromised health and ultimately reduced life expectancy.

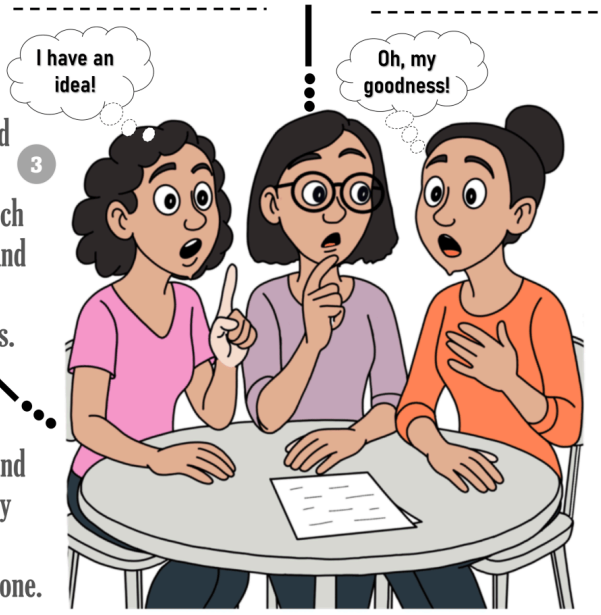

3 Professor, I would like to build on this fact you have established in your research. For my PhD, I want to move beyond showing that polluted indoor air shortens life. I plan to quantify how much it does so, why the effect varies across individuals, buildings, and lifestyles, and how biological vulnerability and behaviour influence these differences in measurable, evidence-based ways.

Then, I will develop advanced hybrid models combining AI to analyse and integrate exposure dose, biological vulnerability, and contextual covariates* to predict changes in life expectancy. My goal is to turn this understanding into practical, value-oriented strategies that improve IAQ and extend healthy living for everyone.

REFERENCES

Adamkiewicz, G., Liddie, J., and Gaffin, J. M. (2020). Clinics in Chest Medicine, 41(4), 809.
Lelieveld, J., Pozzer, A., et. al. (2020). Cardiovascular Research, 116(11), 1910-1917.
Pope III, C. A., Ezzati, M., and Dockery, D. W. (2009). New England Journal of Medicine, 360(4), 376-386.

*Covariate are variables that determine the values of the main variables and their interactions within a study context.

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

Fictional Case Story (Audio – available online) – Part 6

…….

There was a time when life expectancy models were built to reflect the causes that shaped population health, yet they consistently overlooked one of the most important determinants: indoor air. This critical cause was missing, and as a result, the solution created for estimating life expectancy became incomplete and could not support value-oriented national planning.

As a result, national strategies failed to capture a major driver of biological ageing and life expectancy. It was a clear example of what happened when essential causes of problems were not used to inform solutions. It reinforced the need to integrate air in the indoor environment, where people spent most of their lives, into life expectancy estimation. This was especially important since air is a basic need of life and its quality determined how long that life could be sustained.

It took a girl on a long journey of transformation to recognise the systemic flaw. Her former habit of creating solutions without understanding root causes was the same flaw she later saw mirrored in industry practice and national models. With her belief that a healthy building was not one that looked modern, but one that genuinely reduced biological burden and added years of healthy living, she committed herself to addressing this gap. The girl's journey from youth to adulthood is the subject of this fiction story.

1 ……

Ayanfe Omolabanson was still an infant when an event occurred that would leave a quiet, unseen mark on her life. One night, in the small kitchen of the family flat, a pot of oil overheated and burst into flame without warning. Her mother, who had been half asleep, woke to the sharp smell of burning. By the time she reached the kitchen, the fire had already climbed the curtain, and thick smoke was spreading rapidly through the flat.

Coughing and struggling to see through the darkness, she tried to smother the flames, but the heat forced her back. At that same moment, Ayanfe's father had just returned home from work. He was in the bathroom washing off the day's sweat when he heard his wife scream his name. He hurriedly wrapped a towel around his waist, wore the used short pant left hanging on the bathroom hook, and stepped into the hallway, immediately confronted by the thickening haze. Neighbours were already stepping into the corridor, alerted by the smell.

Ayanfe's father rushed in moments later, went straight to the bedroom where the baby slept, and lifted the infant into his arms before the smoke thickened any further. Without delaying, he carried her to the front door and passed her to a neighbour who had already come forward,

instructing them to take the baby outside and keep her safely away from the flat. Only when he saw the neighbour running with the child toward the stairwell did he turn back toward the kitchen.

He re-entered only far enough to reach his wife, calling her name to guide her through the smoke-filled hallway. She stumbled toward him, coughing violently, and he supported her as they moved together into the corridor where the air was clearer.

In the aftermath, the neighbours spoke about how heavy the smoke had been and how quickly it had filled the flat, words the parents remembered for years because they revealed how much smoke the infant must have inhaled before Ayanfe's father reached her. Although mother and child recovered physically, the incident left a quiet imprint on the family that none of them understood at the time.

Although Ayanfe was far too young to remember the fire that engulfed the small kitchen of her family's flat, the experience imprinted itself upon her in ways she would not understand until much later in life. Her parents, believing silence to be a form of protection, never mentioned the incident again. They spoke of it only once between themselves and then agreed, without needing many words, that the memory should be allowed to fade. Their intention was innocent: to preserve the simplicity of their daughter's childhood and prevent fear from taking root in a mind too young to comprehend such danger.

Yet the human body often carries what the conscious mind cannot recall. Although the memory of the flames dissolved into infancy, the physiological imprint remained. As she grew, subtle reactions emerged almost imperceptibly. Still air unsettled her. Rooms with closed windows felt strangely oppressive.

Even as a small child, she would tug at curtains, shift restlessly in enclosed corners, and complain of discomfort without being able to explain why. Her parents interpreted these behaviours as typical childhood fidgeting, noticing the patterns but never connecting them to the night she had been carried through heavy smoke in Ayanfe's father's arms.

Since Ayanfe had no narrative explanation for her sensations, she created one that seemed harmless. She called the feeling impatience. Whenever silence stretched too long or a room felt unmoving, she experienced an urge to act quickly, to fill the space with movement or thought.

As a result, she developed a habit of hurrying through tasks, rushing into ideas, and answering questions before they had fully formed. To her teachers, this quickness appeared as talent. They praised her speed of thought and her ability to speak decisively. In classrooms where responsiveness was often equated with intelligence, her behaviour was rewarded, not examined.

What no one realised was that her speed was not a reflection of confidence but a strategy for managing an unrecognised, embodied discomfort. Stillness, for Ayanfe, carried an inexplicable sense of vulnerability. Silence felt unsafe in ways she could not articulate.

Enclosed spaces with stagnant air awakened a faint but persistent tension in her chest, as though something deep within remembered a moment of suffocation her conscious mind had long forgotten. In response, she became a child who lived slightly ahead of herself, thinking and speaking quickly to outrun sensations she did not understand.

Over time, this trait shaped her sense of identity. Her teachers mirrored her quickness back to her as brilliance, and she accepted it as proof of capability. She leaned into speed and relied on it to navigate academic and social environments. It became the foundation upon which she built her confidence and the lens through which she interpreted her own abilities.

Yet beneath that layer of external praise lay the quieter truth that her hurry was not only a habit but a shield. It protected her from the unease that emerged in moments of stillness, when the air grew heavy and unmoving and her body responded with a discomfort she had no language for.

Thus, even before Ayanfe learnt anything about buildings, air, or health, her life was already shaped by a relationship with indoor environments that was both intuitive and unresolved. The fire she could not remember had left an imprint she did not understand, and in the absence of explanation, she developed behaviours that on the surface looked like gifts but underneath were coping mechanisms.

What her parents had hoped to erase through silence did not disappear; it simply re-emerged in subtler ways, influencing her development and laying the groundwork for a flaw that would later define her life.

In primary school, Ayanfe’s hurried brilliance found the perfect stage. Classrooms rewarded children who could answer quickly, and she answered before many of her classmates had even grasped the question. Whenever a teacher asked how to cool a stuffy classroom, other pupils suggested fans, curtains, or shifting desks. Ayanfe, propelled by the restless tension that had lived quietly in her since infancy, created ideas that sounded fantastical.

When a teacher asked how to cool a stuffy classroom, other pupils suggested fans or curtains. Ayanfe immediately declared that the class should install more windows on the opposite wall, not realising that the wall opened onto a noisy internal corridor. She added that large ceiling fans should be placed in every corner, unaware that such placement would disrupt airflow rather than improve it.

When the teacher asked her to explain why her ideas would work, she simply said that the details could be worked out later, and the class laughed with admiration. Their laughter did not soothe her, but it encouraged her. No one asked her to slow down or question what she had not understood.

The school assessed projects through a narrow lens. Creativity, novelty, and presentation were the pillars of success. Posters with neat handwriting and colourful borders received high marks. Group assignments were evaluated by how impressive the final artefact appeared on display day.

In this environment, Ayanfe flourished. She built cardboard models with shining surfaces and dramatic moving parts. These features dazzled but rarely addressed the root cause of the problem the pupils had been asked to solve. Her teachers praised her decisiveness and flair, unaware that the flaw shaping her behaviour was being deepened, not corrected. Her quickness looked like vision. Her impatience looked like genius. The system elevated her coping mechanism into a celebrated talent.

By the time she entered secondary school, the pattern had solidified into identity. Her design and technology teacher frequently proclaimed that innovation was imagination at work. Sketches that resembled architectural concept art earned high marks. Complex prototypes made from acrylic sheets and LED lights were held up as examples of creativity during morning assemblies. Competitions rewarded spectacle, not depth.

The judging rubrics contained words like aesthetic, originality, clarity of presentation, and craftsmanship. Missing entirely were words like root cause, understanding, or problem definition.

Ayanfe won repeatedly. Her solutions looked intelligent, even when they were untethered from the physics that governed the real world. She produced models of buildings with sweeping glass façades that caught the light in beautiful ways, but she did not know why the spaces inside would behave as they did. She spoke confidently about airflow, shading, and thermal comfort, weaving explanations that sounded plausible only because no one asked her to prove them.

Teachers smiled as she presented her work. Judges applauded. Her classmates relied on her leadership during group projects. She learnt to perform certainty with increasing skill, and the applause masked the growing gap inside her between appearance and understanding.

Stillness continued to trouble her. Any question that required prolonged thought brought a faint tightening to her chest, a discomfort she could not name. Solving quickly helped her outrun that feeling. The educational environment rewarded speed, spectacle, and confidence, so she grew increasingly resistant to slowing down.

The deeper questions, the ones that required her to sit in discomfort and acknowledge what she did not know, remained untouched. She had never been taught that value lay not in the beauty of a solution but in its ability to address the true nature of a problem.

After her O-levels, she pursued A-levels in mathematics, physics, and design, choosing subjects that aligned with her reputation for intelligence. By the time Ayanfe reached her final year of A-levels, her flaw had already taken the shape of a comfortable habit. She had grown accustomed to proposing solutions before understanding the root of a problem.

In school, this had always been rewarded. In competitions, it had been celebrated. In her community, it had been admired. Yet the flaw finally revealed itself clearly in an event that she could not forget.

During a national pre-university innovation challenge, her team was asked to design a low-cost way to improve indoor comfort in rural clinics across the northern part of her country, Buskiya. The brief highlighted overcrowding, heat, and the presence of fumes from nearby generators.

While her teammates began reading reports and interviewing volunteers who had worked in such clinics, Ayanfe sketched her solution within twenty minutes. She proposed an elegant roof canopy with a sculptural profile and narrow slits along the top. She said the shape would cool the interior spaces naturally.

The judges were impressed by the presentation, but the prototype failed completely in the field test. The slits did not allow sufficient hot air to escape, and the new canopy had unintentionally concentrated heat in the waiting area. One nurse later commented that the structure was beautiful but made the clinic feel like an oven during mid-afternoon.

Her teammates were disappointed, not because the idea had failed, but because she had ignored the findings they had gathered. They had told her that heat in those clinics was often trapped in the lower half of the room due to overcrowding and generator positioning.

She had not listened. She had designed for a problem she assumed existed rather than the problem that actually existed. When the team received only a participation certificate, she felt something unfamiliar in her chest, something that lingered long after everyone else had forgotten the competition.

One evening, weeks after the challenge, she found herself alone in her room, replaying the comments of the nurse who had tested their prototype. The nurse had said that they already knew the roof was hot; what they needed was a way to remove trapped air at the human level. It was the first time someone had spoken to Ayanfe as if the consequences of her assumptions mattered.

That night, she realised for the first time that her way of thinking did not merely limit her; it could mislead others. It could waste resources. It could make life harder for the people she believed she was helping. It was also the first moment she recognised, even if only faintly, that her pattern of designing solutions without understanding problems carried real consequences for real people.

She did not know what to do with the feeling. She did not suddenly change. She did not suddenly become reflective. Instead, she lived for several weeks in a quiet discomfort she could not put into words.

She tried to slow down, but slowing down made her restless. She tried to ask more questions, but she listened only halfway. Yet she knew, with a clarity that unsettled her, that something in her needed to change. She did not know how to begin. The desire to transform existed, but the pathway towards transformation remained blurred and intimidating.

In the midst of this confusion, she started searching online for courses that taught people how to understand problems more deeply. She searched phrases like “how to think clearly,” “designing for root causes,” and “problem framing in engineering.” Most results pointed her

back to universities in Buskiya, but she felt instinctively that the educational culture she had grown up in could not guide her transformation.

In her country, the emphasis remained on results, speed, and polished presentations. She wanted something different. She wanted to unlearn and rebuild, and she knew she needed distance to do that. She feared that remaining in the same system that had shaped her flaw would only reinforce it further.

During one of her late-night searches, she came across a university abroad, the University of Edwardbridge in Mathland, a rich and developed country, that offered a Bachelor of Engineering degree in Design Thinking with specialisation in Building Services. The website spoke about problem framing, value-oriented solutions, and research-informed practice.

The words felt like a direct answer to a question she had never been able to articulate. She felt a strange sensation in her chest, as if a stranger had seen the flaw in her before she had seen it in herself. A very small part of her suspected that she needed what the programme promised.

She applied without telling her parents. She told herself that if she was accepted, it would be a chance to become better. She had no idea yet what that would require. She only knew that this was the first step towards a transformation she could no longer ignore.

Before she received a reply, she accepted a temporary job at a building consultancy to gain practical experience. Her temporary role at the consultancy unfolded quietly, without the pressure or responsibility typically associated with professional practice.

She joined as a junior assistant whose tasks centred on formatting diagrams, preparing layout drafts, and compiling site photographs for senior engineers. She held no authority to design solutions or to influence client decisions, and everyone at the company treated her contribution accordingly.

Yet within these limits, her instinct for producing quick, polished diagrams found a convenient place. The company often needed early-stage visual concepts—placeholders used only to start conversations with clients—and her speed and creativity suited those purposes well.

The consultancy specialised in ventilation assessments for schools and residential buildings, and although she accompanied senior engineers on several site visits, she quickly perceived that the industry behaved in ways that felt eerily familiar. Much like the educational system in which she had been raised, clients desired visible interventions, advanced-looking devices, and polished concepts they could present to stakeholders.

Reports, too, were built around colourful diagrams, stylised airflow arrows, and graphic layouts that looked modern regardless of whether they reflected the root cause of any issue. In this environment, where surface-level clarity often overshadowed mechanistic understanding, Ayanfe felt strangely at home.

Her habitual quickness blended seamlessly into the workflow, even though nothing she created carried professional weight. Every diagram she prepared was reviewed, amended, and used only for preliminary discussion; the responsibility for decisions rested entirely with the senior

engineers.

The turning point emerged gently, almost innocently. A primary school approached the consultancy with concerns about musty odours, fatigued pupils, and headaches. The school requested a quick conceptual sketch for their internal board meeting before deciding whether to commission a full diagnostic study.

The senior engineer assigned to the case was overwhelmed with ongoing projects and asked Ayanfe to prepare a simple layout—no recommendations, no technical assertions, only a visual draft to serve as an illustrative placeholder. She visited the school, observed the warm, heavy air, and felt the familiar tightening in her chest, the same sensation that had chased her since childhood.

Although she had not been asked to propose solutions, her instinctive urge to create something impressive surfaced. She added decorative elements to the draft: sealed window gaps, portable purifiers, and faster ceiling fans. She viewed them merely as possibilities, examples of interventions she had seen in brochures and demonstration rooms. They were not technical claims; they were speculative sketches meant to fill space on a page.

The senior engineer reviewed her draft and added a clear verbal and written disclaimer. The disclaimer stated that the visuals were not recommendations and that no action should be taken until proper diagnostics were completed. He then forwarded the draft to the school for preliminary discussion. However, the school board misinterpreted the materials.

Pressed by time, eager to act before the next academic term, and impressed by the apparent professionalism of the diagrams, they implemented the sketched elements immediately. They hired contractors directly, bypassed the consultancy, and never waited for the diagnostic visit scheduled for the following week. By the time the senior engineer returned for measurements, the interventions had already been installed.

Within two months, the situation deteriorated significantly. The sealed gaps reduced natural ventilation, carbon dioxide levels rose steadily during long lessons, and several pupils with asthma experienced difficulties during afternoon classes. When one child fainted and an ambulance was called, the school demanded a formal review.

Measurements revealed that the purifiers were insufficient, the fans merely circulated stale air, and the sealed windows trapped pollutants from cleaning products and human occupancy. The crisis had occurred because the school acted without approval or technical analysis, implementing a visual draft never intended to guide real decisions. The consultancy documented these findings clearly, and responsibility rested entirely with the school’s management.

During the debrief, the principal looked fatigued and spoke quietly, as though thinking aloud rather than addressing anyone in particular. He said he had assumed the steps taken were appropriate and wondered whether anyone had fully understood the root of the issue from the outset.

His words were neither a question directed at the consultancy nor a judgement of any individual. They were simply the reflection of someone realising, perhaps for the first time, that well-intentioned actions could still miss the truth of a problem if the problem itself had never been properly understood.

Yet she felt a quiet shock reverberate through her. Though she bore no responsibility legally, professionally, or ethically, she recognised herself in the misunderstanding. She saw how her polished sketches, created without depth, could be mistaken for authority. She saw how the flaw she had carried since childhood could one day, if left unaddressed, cause harm when she finally held real professional responsibility.

That evening, she returned home and cried in private. Her tears did not come from guilt, for she understood that she had not been at fault, but from the frightening recognition that her long-standing habit—the tendency to design before understanding—was not merely ineffective. It was dangerous.

The incident became a mirror reflecting a version of herself she feared becoming in the future: a professional whose polished shortcuts could mislead institutions, waste resources, or endanger vulnerable people. For the first time, she saw that her flaw was not simply a quirk of personality or a product of her schooling; it was a barrier to the person she needed to become.

Two weeks later, the scholarship email arrived. She had been accepted into the programme at the University of Edwardbridge in Mathland – a developed and economically rich country that had once colonised her country, Buskiya.

Her parents rejoiced, her colleagues congratulated her, and she smiled with them all. But beneath the happiness lay a deep and quiet resolve. She was leaving her home country not to escape blame but to reshape herself before entering the world with professional authority. She boarded her flight with the weight of the incident behind her and a determined, unsettled hope for transformation ahead.

2 ……

Ayanfe’s transformation began the moment she arrived at the University of Edwardbridge. The new environment dismantled everything she thought she understood about design. Students were not judged on the beauty of their drawings or the cleverness of their concepts. They were judged on their grasp of the root cause and on how effectively this understanding informed the solutions they proposed.

Sketching was forbidden until weeks of investigation had passed. Her lecturers expected interviews, evidence collection, causal mapping, and rigorous questioning before a single line was drawn. She felt exposed. Her instinct urged her to respond quickly, yet her tutors asked her to pause, to think, and to justify every assumption.

Her classmates moved slowly and thoughtfully. Their restraint unsettled her. During her first year, she cried in private because she felt like an impostor. She had been celebrated all her life for speed and creativity, only to discover that genuine design required something she had

never been taught to value. Understanding.

In her second and third year, she encountered a series of seminars that reshaped her thinking in a quieter but deeper way. Her professors emphasised that design meant understanding problems before touching solutions. They showed real cases: buildings that failed despite impressive appearances, ventilation systems installed without studying how people actually used the spaces, and intervention technologies adopted because they looked modern rather than because they addressed any true need. The message was simple. Solutions worked only when the root cause was known. Anything else was decoration.

For the first time, Ayanfe saw her own reflection in the failures they described. The professors warned against “solution-first thinking,” explaining how designers often forced ideas onto problems instead of allowing problems to guide ideas. They demonstrated how a single misinterpreted cause could lead to years of ineffective interventions. They reminded the students repeatedly that design was not a performance of creativity but a discipline of comprehension.

Something inside her shifted. She lingered after class, not to ask questions but to absorb the unfamiliar quietness that came from not rushing. She realised she had never stayed long enough with any problem to understand it. She had never traced a symptom back to its source. She had never asked whether her ideas solved anything real. These seminars did not yet reveal the deeper scientific territory she would later explore, but they planted the seed she needed. They taught her that speed was not intelligence, that creativity without understanding was empty, and that every meaningful solution began with humility. This was the year she first began to slow down.

By her final year, she chose to confront the issue directly. For her Bachelor dissertation, she designed a framework for healthy classroom environments aimed at reducing long-term pollutant exposure in vulnerable children. She integrated natural ventilation strategies, material choices, microclimate modelling, and behavioural patterns to examine how design decisions shaped cumulative exposure.

During one of her scheduled dissertation progress meetings, Professor Jones listened carefully as Ayanfe presented her latest findings. When she finished, the professor tilted her head slightly and asked a question that shifted the atmosphere in the room: “You are designing healthier spaces, but do you realise how little the world truly understands about how indoor air influences the years people live?”

The question settled heavily between them, not as a criticism but as an invitation to think beyond design practice and towards human longevity. Ayanfe felt something tighten and then expand inside her. It was this question that deepened the discussion and prompted the room to reflect.

Ayanfe responded slowly, admitting that she had always known indoor air was important but had never considered its connection to life expectancy. She said she found it unsettling that indoor air occupied so much of daily life yet remained almost invisible in national conversations

about health. Her comment lingered in the room, its honesty opening a space for broader reflection.

It was to this remark that Miremba, the other undergraduate student at the meeting, now responded. She leaned forward, her voice gentle but earnest.

[Miremba] —- “To add to your point, Ayanfe, my country has one of the lowest life expectancies in the world. People often talk about poverty, diet, or healthcare, but almost no one considers the air we breathe indoors. Yet the truth is that the quality of indoor air affects everyone, everywhere. It is not just a problem of poor countries; even in wealthier cities, sealed buildings, synthetic materials, and poor ventilation quietly harm health.

Additionally, polluted outdoor air seeps indoors, mixing with fumes from cooking, cleaning, and building materials, creating an invisible burden we often ignore. I only realised how the invisible air around us shapes how long we live when I learnt about indoor air quality from Professor’s class.”

Professor Jones nodded, meeting both students’ eyes.

[Prof. Jones] —- “You are right! Many underestimate how the continuous inhalation of polluted air increases ROS beyond healthy levels, causing oxidative stress that damages cells. This triggers chronic inflammation, disrupts body systems, and accelerates ageing while heightening biological vulnerability to other hazardous factors. Over time, organ function declines, leading to compromised health and ultimately reduced life expectancy.”

The explanation resonated through the room. Something aligned within Ayanfe. From the first day of her undergraduate journey, she had been dismantling her old self piece by piece, but only at this moment did the purpose of that dismantling begin to reveal itself. She spoke slowly, with new clarity.

[Ayanfe] —- “Professor, I would like to build on this fact you have established in your research. For my PhD, I want to move beyond showing that polluted indoor air shortens life. I plan to quantify how much it does so, why the effect varies across individuals, buildings, and lifestyles, and how biological vulnerability and behaviour influence these differences in measurable, evidence-based ways.

Then, I will develop advanced hybrid models combining AI to analyse and integrate exposure dose, biological vulnerability, and contextual covariates to predict changes in life expectancy. My goal is to turn this understanding into practical, value-oriented strategies that improve IAQ and extend healthy living for everyone.”

The room fell quiet again, but this time the silence felt like a beginning. Ayanfe had finally spoken not with speed, but with understanding. The following research problem statement she later wrote for her PhD study became a testament to the depth and maturity of that understanding.

……..

“Life expectancy is a foundational metric used by governments, health agencies, and planners to understand the wellbeing of populations and to guide long-term national strategies. Accurate life expectancy estimates influence the distribution of healthcare resources, the design of environmental policies, the identification of vulnerable population groups, and the evaluation of social inequalities.

For life expectancy to serve these purposes effectively, it must reflect the actual risks that individuals face within the environments where they live. However, one of the most significant environments that people spend time in, the indoor environment, is not meaningfully incorporated into life expectancy estimation. This omission forms the central real-life problem that motivated the research.

The expected performance of national health and environmental strategies is clear. Indoor environments, where individuals spend most of their daily hours, should be represented accurately in life expectancy calculations. Policy tools should allow decision-makers to understand how exposures to indoor air pollutants translate into biological ageing, disease risk, and reduced life expectancy.

They should also be able to distinguish which pollutants and which covariates related to housing conditions, pre-existing health conditions, behaviour, demographics, and socio-economic context contribute most substantially to premature mortality. Ideally, decision-makers should also be able to examine how variations in exposure, behaviour, and building conditions produce different life expectancy outcomes across communities. These capabilities are essential for guiding equitable and value-oriented public-health interventions.

However, the current performance of indoor-air-quality practice and policy falls far short of these expectations. At present, life expectancy projections at national and regional levels do not incorporate indoor air quality in any systematic or quantitative manner. While indoor air pollutants such as $PM_{2.5}$, NO_2, ozone, VOCs, and formaldehyde are known to affect health, their long-term contribution to life expectancy loss is not quantified, not modelled, and not integrated into national planning tools.

Existing life expectancy frameworks treat indoor air as a background condition rather than as a measurable risk factor. As a result, policymakers lack the evidence required to formulate national targets or to evaluate the consequences of poor indoor air.

The gap is further widened by tools currently used to analyse indoor environments. Mechanistic models exist for estimating pollutant concentrations, but they cannot explain how long-term exposure translates into biological or epidemiological outcomes such as premature mortality or reduced longevity.

Likewise, epidemiological data capture associations between pollution and disease, but they do not describe the dynamic indoor exposure patterns that determine real-world risk. There is currently no integrated framework that connects pollutant dynamics, biological vulnerability, and the full range of covariates to quantitative life expectancy outcomes.

The practical situation inside real homes, which itself forms part of the covariate structure, adds another layer of complexity. Indoor pollutant levels fluctuate sharply due to cooking, cleaning, ventilation behaviour, building integrity, and infiltration of outdoor pollutants. These variations occur minute by minute, hour by hour, and season by season.

Yet current assessment approaches rely heavily on short-term measurements or single-location sensors that do not represent real exposure. Without a scientifically coherent method to convert these fluctuating patterns into cumulative biological burden and then into life expectancy loss, decision-makers cannot identify which groups are most at risk or which interventions will produce the largest gains in longevity.

As a consequence, interventions intended to improve indoor air quality are often selected without an understanding of their long-term health impact. Architectural modifications, ventilation strategies, air-cleaning technologies, and behavioural advice are deployed without a quantitative basis for determining how much life expectancy they can recover, for whom, and under what conditions. This creates a significant gap between the expected performance of indoor-air-quality interventions and their real-world capability to extend healthy-life years.

These deficiencies together define the core real-life problem: the absence of a scientifically grounded, dynamic, and context-sensitive framework linking indoor air quality directly to life expectancy. Without such a framework, national strategies cannot identify the true health burden of indoor air, predictive tools cannot forecast life expectancy variations accurately, and policymakers cannot choose interventions that deliver measurable, equitable, and value-oriented improvements in longevity.

This gap created the necessity for the present research, which sought to produce a coherent scientific pathway from pollutant exposure to biological ageing to life expectancy outcomes and to support national decision-making through predictive and intervention-optimised models.”

……..

Her interest and courageous ambition to address this research problem led her to formulate three research questions that needed to be answered.

The research questions are as follows: (i) How do long-term exposures to indoor air pollutants, such as fine particulate matter ($PM_{2.5}$), nitrogen dioxide (NO_2), ozone (O_3), volatile organic compounds (VOCs), and formaldehyde (HCHO), quantitatively influence life expectancy within the study country, given the established mechanistic pathway linking oxidative stress to accelerated ageing and premature mortality? (ii) To what extent could advanced hybrid models, defined as the integration of mechanistic models with artificial intelligence models, combine indoor air pollutant dynamics, biological vulnerability indices, and covariates to accurately predict variations in life expectancy attributable to chronic indoor air pollutant exposure across diverse populations and building archetypes? (iii) Which combinations of architectural, engineering, and behavioural interventions most effectively improve indoor air quality and extend life expectancy, and how can quantitative optimisation frameworks guide equitable, value-oriented implementation across diverse environmental and socio-economic contexts?

For the first research question, the Null Hypothesis (H_{01}) is that long-term exposure to indoor air pollutants has no statistically significant quantitative effect on life expectancy within the study population. The Alternative Hypothesis (H_{11}) is that long-term exposure to indoor air pollutants has a statistically significant quantitative effect on reducing life expectancy within the study population.

For the second research question, the Null Hypothesis (H_{02}) is that hybrid models integrating indoor air pollutant dynamics, biological vulnerability indices, and covariates do not significantly improve prediction accuracy of life expectancy variations compared with mechanistic models that contain the same variables but operate without the AI component. The Alternative Hypothesis (H_{12}) is that hybrid models integrating indoor air pollutant dynamics, biological vulnerability indices, and covariates significantly improve prediction accuracy of life expectancy variations compared with mechanistic models that contain the same variables but operate without the AI component.

For the third research question, the Null Hypothesis (H_{03}) is that no combination of architectural, engineering, or behavioural interventions produces a statistically significant improvement in indoor air quality, life expectancy, or value delivery compared with baseline conditions. The Alternative Hypothesis (H_{13}) is that at least one combination of architectural, engineering, or behavioural interventions produces a statistically significant improvement in indoor air quality, life expectancy, and value delivery compared with baseline conditions.

The research questions and problems informed the following objectives of her PhD research: (i) To quantify the extent to which long-term exposures to indoor air pollutants, specifically $PM_{2.5}$, NO_2, O_3, VOCs, and HCHO, reduce life expectancy within the study country by modelling their mechanistic influence on oxidative stress, accelerated biological ageing, and premature mortality. (ii) To develop and validate an advanced hybrid mechanistic–AI model that integrates indoor air pollutant dynamics, biological vulnerability indices, and covariates, in order to accurately predict variations in life expectancy across diverse population groups and building archetypes under chronic indoor air pollutant exposure. (iii) To identify and evaluate the architectural, engineering, and behavioural intervention combinations that most effectively improve indoor air quality and extend life expectancy, and to develop a quantitative optimisation framework that enables equitable and value-oriented implementation across diverse environmental and socio-economic contexts.

3 ……

Research Methods

Methods for Research Question 1:

Background

Research Question 1 asked: How do long-term exposures to indoor air pollutants, such as fine particulate matter ($PM_{2.5}$), nitrogen dioxide (NO_2), ozone (O_3), volatile organic compounds (VOCs), and formaldehyde (HCHO), quantitatively influence life expectancy within the study

country, given the established mechanistic pathway linking oxidative stress to accelerated ageing and premature mortality?

Central to the approach adopted was the aim to capture the mechanistic and epidemiological foundations of how indoor air pollutants accelerate biological ageing and increase mortality risk. To provide clarity, the specific research question guiding this part of the study asked how long-term exposure to pollutants such as $PM_{2.5}$, NO_2, O_3, VOCs, and HCHO quantitatively influenced life expectancy within the study country.

$PM_{2.5}$ refers to fine particulate matter with aerodynamic diameters of 2.5 micrometres or smaller, small enough to penetrate deep into the lungs and enter the bloodstream. NO_2 (nitrogen dioxide) is a combustion-related gas produced by activities such as cooking and traffic emissions, known to irritate the airways and contribute to chronic respiratory and cardiovascular stress. O_3 (ozone) is a highly reactive gas formed through chemical reactions involving sunlight and pollutants, capable of causing oxidative injury to lung tissue.

VOCs (volatile organic compounds) represent a large group of carbon-based chemicals released by paints, cleaning products, furnishings, and human activities, many of which can trigger inflammation and systemic toxicity. HCHO (formaldehyde) is a hazardous indoor chemical released from building materials and household products, classified as a human carcinogen and strongly associated with cellular damage and chronic health decline. Together, these pollutants form a complex indoor mixture that affects biological ageing processes and thereby influences life expectancy.

This required methods that were sensitive to the biological pathways linking indoor air pollutant exposure to health decline and capable of producing population-relevant estimates of life expectancy loss. The methodology therefore integrated real-world indoor exposure measurements with biological indicators of vulnerability and retrospective national mortality data. Bringing these elements together ensured that both the physiological mechanisms and the population-level consequences of exposure could be captured within a coherent analytical structure.

This hybrid approach made it possible to model long-term survival outcomes without requiring a full-generation cohort study. The purpose of this research component was to quantify the magnitude of life expectancy reduction attributable to chronic indoor exposure and to understand how differences in vulnerability and environmental conditions shaped this impact.

The research question required an understanding of whether the accumulation of pollutant exposure over many years or decades had a measurable effect on life expectancy. The purpose extended this by aiming to quantify the magnitude of this effect and identify the circumstances under which it was most pronounced.

To frame the analysis formally, the null hypothesis stated that long-term exposure to indoor air pollutants had no statistically significant effect on life expectancy, whereas the alternative hypothesis proposed that such exposure significantly reduced life expectancy. Every component of the methodology therefore was structured to allow rigorous testing of these hypotheses.

Study Design

The overall design followed a hybrid structure that linked a prospective indoor exposure and biological monitoring study with retrospective national mortality records. This structure was chosen because it allowed the research to investigate life expectancy impacts in a scientifically rigorous yet time-efficient manner. The prospective component provided detailed exposure data collected over twelve to eighteen months, while the retrospective component supplied the long-term mortality patterns already documented in national databases.

The prospective component referred to the part of the study that moved forward in time and directly measured what participants were breathing inside their homes using sensors and biological samples. It captured real-life pollutant concentrations as they occurred. In contrast, the retrospective component referred to the part of the study that looked backwards in time, using existing national records that documented how long people lived, what illnesses they experienced, and the causes of death over many decades.

This combination of forward-looking measurement and backward-looking national data allowed the study to connect present-day exposure conditions with the long-term health patterns observed in the population. Using these two types of data together made it possible to study a question that would normally take decades to investigate, because the long-term information needed was already available through national records.

The prospective component was essential because indoor air pollutant concentrations varied widely from one household to another and were strongly influenced by daily life. Ventilation habits, such as how often windows were opened, together with building design, cleaning activities, cooking frequency, other indoor sources, and the intrusion of outdoor pollution, all shaped the concentrations of pollutants inside each home.

National-level exposure estimates lacked the precision needed to understand these differences. By placing calibrated sensors inside a diverse range of homes, the study captured real exposure patterns as people cooked, cleaned, worked from home, slept, and moved through their daily routines.

The home was used as the primary monitoring environment because it accounted for the largest share of long-term indoor exposure, provided a stable location for continuous high-resolution measurement, and reflected the real behavioural and architectural drivers of cumulative inhaled dose.

Unlike workplaces, schools, and commercial buildings, which change occupancy patterns and access conditions, the home offered a consistent and scientifically reliable platform for multi-month monitoring. This produced exposure estimates grounded in the actual lived environments of the households.

The retrospective component consisted of national mortality and demographic records that documented survival patterns over several decades. These datasets provided the statistical foundation for estimating life expectancy at different exposure levels. Because these mortality

data already existed, the study did not depend on waiting for deaths to occur within the monitored households. Instead, exposure–response models were used to estimate how changes in pollutant levels could influence life expectancy.

This hybrid methodology therefore linked real-world exposure measurements with long-term survival data. The integration strengthened scientific accuracy, improved modelling precision, and avoided the impracticality of long-term cohort tracking. Although the home served as the primary monitoring location, total daily exposure from workplaces, vehicles, and public indoor environments was still captured through wearable sensors and time–activity diaries. This ensured that the focus on homes enhanced scientific validity rather than narrowing the scope of exposure assessment.

Study Population and Sampling Strategy

The study population was drawn from one hundred and fifty to two hundred households across the country. The sampling strategy was deliberately stratified to reflect the diverse ways in which people lived, ventilated their homes, and occupied their indoor environments. Households were selected from both urban and suburban regions because these settings differed substantially in the outdoor pollution that infiltrated indoor air.

Homes were included from public housing blocks, private apartments, and landed properties because these residential types differ structurally in predictable ways. Public housing blocks often have standardised layouts and shared ventilation patterns, private apartments vary widely in façade design and external orientation, and landed properties typically have more openings, external walls, and direct outdoor exposure. These structural and design differences influence how much outdoor air enters the home, how quickly indoor air is replaced, and how pollutants accumulate.

They also influence the types of indoor pollutant sources typically found in each setting, as well as the frequency and intensity of their emissions. For example, certain cooking appliances, cleaning practices, renovation materials, or lifestyle-related activities may be more common in one residential type than another, resulting in different emission profiles. Including this spread of residential types was therefore necessary to obtain realistic variation in indoor pollutant concentrations.

Adult residents within these households were invited to participate in the biological and behavioural components of the study, provided that they spent enough time at home for their exposure to be representative of actual indoor living patterns. Variation in age, gender, occupation, health status, and socio-economic background was intentionally preserved because these characteristics influenced both exposure behaviour and vulnerability to pollutant-induced ageing. Including such diversity made the findings more generalisable and strengthened the study’s capacity to identify population subgroups with differential susceptibility.

Adults were selected rather than children because biological sampling, vulnerability assessment, and life expectancy modelling rely on physiological and behavioural patterns that are stable and interpretable only in adult populations. In addition, ethical and developmental

considerations made adult participation the most scientifically valid and operationally feasible approach for this study.

This sampling strategy ensured that the data generated were representative of both the physical characteristics of the residential environments and the behavioural practices that shaped pollutant exposure. It therefore aligned closely with the research question, which aimed to understand how long-term exposure influences life expectancy within the country. The resulting variability across homes, behaviours, and adult participants provided the diversity needed to test whether chronic exposure meaningfully affected longevity.

Exposure Assessment and Monitoring

Understanding long-term exposure required continuous, high-resolution monitoring, since indoor air pollutant concentrations fluctuated significantly in response to cooking, cleaning, window operation, human activity, outdoor infiltration, and the off-gassing of materials. Each participating household therefore was equipped with calibrated indoor air quality sensors installed in the living room and main bedroom.

These locations were selected because they represented the spaces where people spent most of their time and where the dynamics of pollutant accumulation and clearance could be observed most clearly. The sensors measured concentrations of $PM_{2.5}$, nitrogen dioxide, ozone, volatile organic compounds, and formaldehyde at ten-minute intervals. This resolution ensured that both short peaks and long background exposures were captured.

A smaller subset of participants was provided with wearable micro-exposure sensors for a period of three to six months. These devices measured personal exposure in real time as individuals moved through work environments, transportation systems, commercial spaces, and other indoor environments outside the home.

Indoor exposure was not limited to the home, and personal monitoring allowed the study to estimate the proportion of exposure that originated from workplaces, vehicles, and public buildings. Although the wearable monitoring period was shorter, it offered detailed insight into exposure patterns across multiple indoor environments and strengthened the modelling of total daily exposure.

Participants also completed brief digital check-ins rather than long-term diaries, which captured essential information on ventilation habits, cooking frequency, cleaning activities, use of chemical products, household occupancy, fan and air conditioning operation, and the presence or absence of smoking. These digital check-ins consisted of very short updates completed approximately once every one to two weeks through a simple mobile or online form, each taking less than a minute and designed to provide periodic behavioural snapshots without requiring daily reporting.

These check-ins were designed to be short, easy to complete, and used only during selected periods within the monitoring window rather than continuously, in order to minimise participant burden and avoid disrupting daily routines. The behavioural information collected through these

short intervals provided vital contextual insight that helped explain pollutant peaks and allowed the identification of behavioural drivers of exposure.

To complement these reports, sensor-based indicators such as temperature, humidity, motion, and door–window status were used whenever possible, reducing reliance on self-reporting. Together, these data sources enabled appropriate adjustments for behaviours that might otherwise confound exposure–health relationships.

Mass balance modelling was employed to fill gaps in exposure data when sensors were offline or not worn. These models incorporated outdoor pollutant levels, air exchange rates, deposition processes, and indoor source emissions, ensuring continuity in the exposure dataset.

The continuous exposure profiles derived from these measurements allowed the construction of cumulative exposure indices for each pollutant, which were essential for the statistical modelling of mortality risk. This approach ensured that the exposure assessment was robust, mechanistically grounded, and capable of detecting realistic variations in chronic exposure.

Measurement of Biological Vulnerability

The link between pollutant exposure and life expectancy was mediated by biological mechanisms that accelerated ageing and increased mortality risk. The methodology therefore incorporated biological sampling to measure these mechanistic pathways. A subsample of approximately one hundred participants provided blood and saliva samples during the monitoring period.

These samples were analysed for biomarkers of inflammation, oxidative stress, and biological ageing. Indicators such as C-reactive protein, interleukins, F2-isoprostanes, malondialdehyde, mitochondrial DNA copy number, and telomere length were selected because they represented biological processes known to be affected by air pollution exposure. Their inclusion allowed the study to determine whether higher exposure corresponded to elevated physiological stress.

Non-invasive cheek swabs were collected to study genetic polymorphisms associated with the body’s capacity to metabolise pollutants and manage oxidative stress. Variants of the GSTM1, GSTT1, and CYP1A1 genes were of particular interest because they affected detoxification pathways and antioxidant defences. By combining genetic data with biomarker measurements, the study constructed a biological vulnerability index that captured differences in physiological susceptibility among individuals. This index was integrated into the statistical modelling to identify how vulnerability interacted with exposure to influence mortality risk.

The inclusion of biological vulnerability aligned with the research question because it recognised that long-term exposure influenced life expectancy not only through pollutant concentrations but also through the body’s ability to manage inflammation, oxidative stress, and cellular damage.

This addition supported the purpose of the study by enabling analysis of how physiological and genetic differences contributed to variation in population-level risk. It also strengthened the statistical modelling by reducing unexplained variability, thereby increasing the study’s ability to

detect meaningful exposure–mortality relationships and improving the likelihood of rejecting the null hypothesis when appropriate.

Covariate and Confounder Assessment

Indoor air pollutant exposure did not occur in isolation from other determinants of health and longevity. To isolate the effect of pollutant exposure on life expectancy, the study collected detailed covariate information through structured questionnaires.

These covered socio-economic status, education level, diet quality, physical activity, smoking behaviour, medical history, occupational exposures, and the use of air purifiers. Building-level characteristics such as age, materials, floor level, and ventilation pathways were also recorded. These characteristics influenced both exposure and health outcomes and therefore needed to be accounted for in the statistical models.

Geospatial datasets provided additional information on outdoor pollution levels, population density, greenery coverage, ambient temperature, and neighbourhood environmental conditions. These variables allowed the study to separate indoor-related influences from broader environmental factors. Including covariates ensured that the associations modelled between pollutant exposure and life expectancy reflected indoor pollutant effects rather than unmeasured confounding factors.

Mechanistic Foundations for the Exposure–Risk–Longevity Pathway

A core requirement of Research Question 1 was the development of mechanistic constructs that expressed, in measurable terms, how indoor air pollutant concentrations, exposure duration, inhalation rates, biological vulnerability, and contextual factors combined to influence long-term health outcomes. This section established the mathematical and conceptual framework linking exposure to health risk and, ultimately, to life expectancy loss.

These mechanistic relationships served as the foundation for both the epidemiological analysis in Research Question 1 and the hybrid predictive modelling undertaken later in Research Question 2. By formalising these relationships, the study ensured that its analytical structure was grounded in biological plausibility rather than relying solely on empirical correlation.

Indoor air pollution occurs as a mixture, and the individual pollutants do not operate independently; their toxicological effects accumulate across exposure pathways and may interact. Capturing the combined burden of multiple pollutants therefore required the construction of a cumulative exposure mixture index. This index represented the lived experience of participants as they moved through different indoor environments, engaged in household activities, and inhaled pollutant mixtures at varying intensities. Cumulative exposure to pollutant mixtures was expressed as:

$$E_{\text{mix}} = \sum_{i=1}^{n} w_i (C_i \times T_i \times IR_i)$$

where C_i represented the concentration of pollutant *i*, T_i the time spent in each microenvironment, IR_i the inhalation rate associated with that microenvironment, and w_i the toxicological weighting assigned to pollutant $_i$. This formulation captured concentration, duration, and inhaled dose in a single integrated index suitable for modelling chronic exposure.

The next step involved translating cumulative exposure into biological risk. Biological evidence shows that pollutant-induced ageing and mortality risk are shaped not only by exposure magnitude but also by individual-level vulnerability to oxidative stress, inflammation, and systemic decline. Health risk was therefore conceptualised as a function of the exposure mixture (E_{mix}), biological vulnerability (V), and contextual covariates (Z):

$$\text{Health Risk} = f(E_{\text{mix}}, V, Z)$$

To express this mechanism in a form suitable for statistical calibration and later hybrid modelling, the study constructed a predicted risk score defined as:

$$\text{Predicted Risk Score} = \beta_0 + \beta_1(E_{\text{mix}}) + \beta_2(V) + \beta_3(E_{\text{mix}} \times V) + Z\gamma$$

This representation captured the direct effect of exposure, the contribution of biological vulnerability, their interaction, and the modifying influence of contextual covariates. To fit this model, it was essential to define and quantify an Actual Risk Score based on biological measurements collected during the prospective monitoring phase.

The Actual Risk Score was constructed from a composite of biomarker indicators, including validated measures of oxidative stress, inflammatory response, cellular ageing markers, and organ-system function. Each biomarker was standardised, normalised, and weighted using established toxicological and epidemiological criteria to generate a single continuous index reflecting the physiological burden of exposure. This composite formulation ensured that the Actual Risk Score captured the biological processes known to mediate pollutant-induced health decline.

The relationship between predicted and actual risk was then expressed as:

$$\text{Actual Risk Score} = \text{Predicted Risk Score} + \epsilon$$

where ε represented the residual error comprising unexplained biological variation, measurement error, and stochastic physiological processes.

This formulation enabled the use of regression-based estimation, where the observed Actual Risk Scores served as the dependent variable and the components of the predicted risk equation served as independent variables. Through this calibration step, the model computed the coefficients by minimising the residual error (ε) using maximum likelihood estimation.

The resulting coefficients represented the best-fitting parameters that explained how strongly exposure, vulnerability, their interaction, and contextual covariates contributed to the biologically observed risk burden. This process ensured that the predicted risk model was empirically anchored and not based solely on theoretical assumptions.

Finally, life expectancy loss attributable to chronic indoor exposure was expressed as a downstream function of the calibrated health risk:

$$\Delta LE = f(\text{Health Risk})$$

This step linked the environmental and biological components of long-term exposure to population-level survival outcomes. Using the calibrated risk scores, individuals were stratified into exposure–vulnerability profiles, and these risk estimates were translated into shifts in age-specific mortality probabilities. The resulting variations in survival curves formed the inputs to life-table calculations, enabling the estimation of expected years of life lost associated with chronic indoor exposure.

Life-table calculations are a standard public-health technique that track how a hypothetical group of people would survive year by year, based on observed mortality rates, allowing the estimation of average life expectancy. They operate like a statistical “lifetime simulation,” showing how changes in risk factors alter the expected number of years a population can live. This process reflected the understanding that pollutant-induced biological burden accumulates over time, altering survival probabilities and ultimately shaping life expectancy.

Together, these equations established a coherent mechanistic narrative connecting pollutant exposure, biological vulnerability, behavioural context, health risk, and life expectancy. They ensured that the life expectancy estimates produced in Research Question 1 were biologically justified, statistically robust, and structurally aligned with the hybrid predictive modelling framework developed in Research Question 2.

Outcome Measurement

Life expectancy outcomes were derived from national mortality databases that contained decades of records on age at death, cause of death, and demographic statistics. Such datasets formed the foundation for life table analysis, which estimated life expectancy for different demographic and exposure groups. These national mortality datasets were particularly valuable because they offered complete population coverage, high reporting accuracy, and long historical depth, allowing the study to work with established survival patterns rather than waiting for new data to accumulate.

The long-term nature of these records meant that rare causes of death, age-specific mortality risks, and trends across different birth cohorts could be robustly incorporated into the analysis. By integrating the exposure indices derived from the prospective monitoring with mortality rates from the retrospective dataset, the study modelled the expected years of life lost for individuals with different levels of pollutant exposure.

This integration step involved assigning each exposure profile to the corresponding age- and sex-specific mortality rates in the national dataset, allowing the construction of life tables that reflected how different exposure levels might alter survival trajectories across the lifespan.

This approach allowed long-term outcome estimation without requiring the study participants themselves to be followed for many years. In traditional cohort studies, researchers must wait decades to observe how exposure affects actual lifespan, but the present approach used the already observed mortality patterns of the national population to infer how today's exposure levels would shift future life expectancy. This made it possible to estimate survival impacts within a realistic research timeframe while still maintaining scientific credibility.

The use of well-established life table methods ensured that the life expectancy estimates were scientifically robust and comparable to those used in national or global burden of disease studies. Life table methods work by combining age-specific mortality rates with a hypothetical surviving population to simulate how many years people at different ages are expected to live.

By applying these mortality rates to a large virtual cohort, the method calculates the average number of years remaining at each age, producing a standardised and transparent estimate of life expectancy. These methods have been extensively validated and are widely used by public health agencies to estimate the health impacts of environmental risks. The method therefore directly answered the research question by quantifying the impact of exposure on survival outcomes, while remaining feasible within the time constraint of the research study.

Ethical Considerations

The research was conducted under strict ethical principles to protect participants, safeguard their data, and ensure that the study's procedures were aligned with national and institutional ethical guidelines. Participation was entirely voluntary, and informed consent was obtained from all individuals before any monitoring or biological sampling took place.

Consent materials were written in accessible language so that participants could clearly understand the purpose of the study, the types of data being collected, and the ways in which the results would be used. Participants were informed of their right to withdraw at any stage without penalty, and withdrawal did not affect their access to any services or resources.

Because the study involved continuous indoor monitoring and biological sampling, the protection of personal privacy was prioritised. All identifying information was removed from the analytical dataset, and household locations were anonymised to prevent the possibility of reverse identification. Biological samples were stored securely and were analysed only for the biomarkers and genetic polymorphisms relevant to the research question. No medically actionable genetic information was disclosed to participants to avoid unnecessary psychological burden.

The study team took care to minimise disruption to daily routines. Sensors were installed in ways that respected occupants' living spaces, and wearable devices were used only with explicit consent. Ethical approval was obtained from the appropriate institutional review board before data collection began, and all procedures were reviewed regularly to ensure continued compliance. Through these measures, the study maintained high ethical standards while generating the evidence necessary to understand the life expectancy implications of indoor air pollution.

Contribution to Knowledge

The methodological approach adopted in this study made several important contributions to scientific knowledge in the fields of indoor air quality, environmental epidemiology, and public health modelling. By integrating twelve to eighteen months of high-resolution indoor exposure data with retrospective national mortality records, the study demonstrated a feasible yet powerful way to estimate life expectancy impacts without requiring decades-long cohort follow-up.

This hybrid design addressed one of the long-standing barriers in indoor air epidemiology, namely the difficulty of linking dynamic, highly variable indoor exposure patterns to long-term survival outcomes. Through this integration, the study showed how near-term exposure measurement could be meaningfully connected to long-term mortality patterns at population scale.

The inclusion of biomarker and genetic vulnerability indices added another significant layer of mechanistic insight. Conventional life expectancy models rely largely on environmental concentrations or estimated doses, but this methodology incorporated biological evidence of oxidative stress, inflammation, mitochondrial strain, and accelerated cellular ageing. This advancement provided a more mechanistically grounded exposure–mortality relationship, demonstrating how differences in physiological susceptibility shaped the life expectancy burden of indoor air pollutants.

Furthermore, the methodology advanced the understanding of indoor pollutant mixtures by applying mixture modelling approaches rather than single-pollutant frameworks. Indoor air pollution rarely occurs in isolation, and the ability to detect interaction effects between $PM_{2.5}$, NO_2, O_3, VOCs, and HCHO represented a methodological improvement over existing approaches. This mixture perspective allowed the identification of pollutant combinations that exerted disproportionate effects on life expectancy.

Finally, the use of mass balance models to reconstruct continuous exposure profiles demonstrated a pragmatic way to overcome gaps in sensor data, enhancing model robustness. Together, these methodological innovations expanded the toolkit for quantifying long-term health consequences of indoor air pollution and provided a replicable framework for future national or regional studies.

Methods for Research Question 2:

Background

Research Question 2 asked: To what extent could advanced hybrid models, defined as the integration of mechanistic models with artificial intelligence models, combine indoor air pollutant dynamics, biological vulnerability indices, and covariates to accurately predict variations in life expectancy attributable to chronic indoor air pollutant exposure across diverse populations and building archetypes? This question expanded the scientific ambition of the study beyond what had been achieved under Research Question 1.

While Research Question 1 quantified the life-expectancy effects of long-term indoor air pollutant exposure using classical epidemiological and mechanistic reasoning, the second research question sought to determine whether a dynamic, real-time, and mechanistically interpretable predictive model could be developed.

Classical epidemiology, as used in Research Question 1, combined exposure, biology, and behaviour only in a cumulative and retrospective manner. It quantified how life expectancy changed after many years of accumulated indoor air pollutant exposure, but it could not account for the day-to-day fluctuations in pollutant levels, behavioural patterns, or biological responses. Most importantly, it could not update life expectancy predictions as new information emerged.

Research Question 2 addressed this limitation by developing a hybrid model in which pollutant behaviour, biological sensitivity, and environmental conditions were treated as continuously interacting processes. This approach enabled life expectancy estimates to evolve over time instead of remaining fixed at a single cumulative outcome.

The purpose of this part of the study was therefore to build a hybrid mechanistic–AI modelling framework capable of forecasting life expectancy outcomes based on ongoing changes in exposure dose, biological vulnerability, and contextual covariates. Whereas Research Question 1 methodology generated a single, population-level estimate derived from cumulative exposure and long-term mortality patterns, Research Question 2 methodology created a predictive system that recalculated risk whenever new information became available.

This allowed life expectancy projections to respond dynamically as pollutant concentrations shifted, as occupant behaviours altered ventilation or emission patterns, or as biological vulnerability differed across individuals or subpopulations.

Operationally, the model used time-resolved exposure data collected at ten-minute intervals. Each new interval provided updated information on indoor concentrations, which the framework integrated into an evolving exposure–risk trajectory. Although life expectancy itself changes only after substantial accumulation of exposure over months or years, the model continuously refined the trajectory leading to that long-term outcome.

The ten-minute resolution therefore functioned as an updating mechanism, ensuring that the cumulative risk profile remained accurate as environmental and behavioural conditions changed, without implying that life expectancy itself fluctuated over short timescales.

This conceptual expansion was crucial for developing a future-ready indoor air quality decision-support structure that could operate in real environments and inform personalised or context-specific interventions.

The null hypothesis for Research Question 2 stated that hybrid models integrating indoor air pollutant dynamics, biological vulnerability indices, and covariates did not significantly improve the prediction accuracy of life-expectancy variations compared with pollutant-only or conventional statistical models.

The alternative hypothesis stated that hybrid models integrating indoor air pollutant dynamics, biological vulnerability indices, and covariates significantly improved prediction accuracy relative to pollutant-only or statistical models.

The methodological design of RQ2 was therefore structured to enable a fair and rigorous comparison between hybrid and non-hybrid models, ensuring that any observed improvement in predictive performance could be attributed to the integrated mechanistic–AI framework rather than to differences in data handling or model evaluation procedures.

Predictive Modelling Approach

The work undertaken for Research Question 2 built directly on the multisource dataset generated in Research Question 1. This dataset, which included time-resolved pollutant concentrations, mass-balance-reconstructed exposure profiles, biomarker-derived vulnerability indices, behavioural observations, sensor proxies, building characteristics, and national life-expectancy outcomes, formed the foundation for the hybrid modelling environment.

The intent was not simply to feed this data into a machine learning model, but to re-express the entire exposure–risk–longevity pathway in a dynamic computational structure that combined mechanistic equations with artificial intelligence in a coherent and interpretable manner.

Whereas RQ1 adopted a traditional epidemiological structure based on hazard ratios and life-table analysis, RQ2 sought to create a multi-layered computational pipeline. This pipeline had to be capable of capturing three fundamental aspects of the exposure–health relationship.

The first aspect was the physical behaviour of pollutants indoors. Pollutant concentrations changed continuously through emission, ventilation, infiltration, deposition, and human activity patterns. Traditional statistical models treated these concentrations as static or averaged values, which weakened predictive precision. The hybrid modelling approach therefore incorporated mechanistic pollutant equations to ensure that the physical laws shaping pollutant dynamics were captured with scientific fidelity.

The second aspect was biological vulnerability. Individuals differed markedly in their sensitivity to pollutants due to variations in oxidative stress markers, inflammatory responses, telomere length (the protective caps at the ends of chromosomes; shorter telomeres indicated accelerated cellular ageing and reduced capacity to withstand pollutant-induced damage), metabolic clearance, and genetic polymorphisms.

Hybrid modelling permitted the biological pathways to be embedded directly into the predictive structure by converting biomarker data into learnable vulnerability indices that interacted dynamically with exposure data.

The third aspect was the behavioural and contextual environment, which modified exposure patterns in significant ways. Behaviours such as window opening, chemical product use, and cleaning frequency influenced emission rates and ventilation patterns. Building archetypes influenced infiltration, airflow distribution, and pollutant retention.

In addition to these environmental and behavioural factors, individual characteristics also played a significant role in shaping physiological risk. Pre-existing health conditions such as asthma, cardiovascular disease, or impaired lung function influenced how strongly a person reacted to the same level of exposure.

Demographic attributes, including age, gender, and household structure, further modified susceptibility by affecting breathing patterns, time spent indoors, and vulnerability to inflammation. Economic context added another layer of variation. Households with fewer resources often lived in buildings with poorer ventilation, older materials, or limited access to exposure-reducing practices.

Together, these individual differences altered how a given pollutant exposure translated into biological stress and long-term risk. These contextual features collectively created dynamic modifiers in the exposure–response relationship that needed to be captured in real time within the model.

The hybrid modelling framework therefore differed fundamentally from the RQ1 methodology. RQ1 estimated life expectancy using established retrospective mortality data and exposure–response modelling. RQ2, in contrast, constructed a computational system capable of forecasting life expectancy dynamically, even when input conditions changed. This distinction represented the methodological evolution from estimation to prediction.

Rationale for Hybrid Modelling

The rationale for hybrid modelling emerged from scientific limitations embedded in both mechanistic and artificial intelligence methods when applied independently. Classical statistical methods and mechanistic models offered clarity and interpretability but were unable to capture high-dimensional and nonlinear relationships across time, pollutant mixtures, behaviours, and biological states. Machine learning models demonstrated exceptional predictive ability but lacked physical and biological grounding, leading to risks of spurious relationships.

The hybrid modelling approach adopted in this study overcame these limitations by integrating these two methodological traditions. Mechanistic equations provided guardrails of scientific plausibility, ensuring that all predictions respected the known physical and biological constraints governing indoor pollutant exposure and human physiology.

Artificial intelligence models captured complex and nonlinear patterns that classical models could not represent. By embedding mechanistic constructs directly into the machine learning pipeline, the study created a structure that was both empirically powerful and scientifically interpretable.

This rationale aligned directly with the purpose of Research Question 2. The aim was not simply to determine whether artificial intelligence could make predictions, but whether combining AI with mechanistic and biological logic could produce a superior predictive engine capable of understanding the exposure–vulnerability–longevity pathway more accurately than traditional approaches.

Furthermore, the hybrid approach aligned with emerging scientific consensus in environmental health modelling, where future prediction systems require the integration of physics-based models with high-dimensional data science methods. The methodology therefore placed the research at the forefront of contemporary indoor air quality science.

Integrated Data Sources Used in the Modelling Framework

The hybrid modelling framework required the careful integration of all datasets generated under Research Question 1. This integration was essential because a predictive model is only as powerful as the data it draws from, and the goal of Research Question 2 was to construct a modelling ecosystem that represented the biological, environmental, behavioural, architectural, and socio-demographic processes that jointly shape life expectancy outcomes.

The resulting dataset combined time-varying exposure information, physiological sensitivity indicators, behavioural and environmental covariates, building characteristics, and population-level survival outcomes into a multilevel predictive environment.

Indoor air pollutant concentration data formed the core temporal backbone of the model. Pollutants such as $PM_{2.5}$, NO_2, O_3, VOCs, and HCHO fluctuated dynamically inside homes due to human activity, ventilation behaviours, and infiltration from outdoors. These time series supported both the mechanistic modelling of concentration changes and the artificial intelligence algorithms that analysed temporal dependencies in exposure trajectories.

Mass balance equations reconstructed pollutant concentrations during periods of sensor downtime or when wearable exposure monitors were not in use. These reconstructions relied on physically grounded parameters including air exchange rates, infiltration coefficients, deposition velocities, and emission strengths. By filling data gaps with mechanistically plausible values, the model avoided learning spurious patterns arising from missing observations.

Biomarker datasets created the biological vulnerability index, V. These datasets included validated markers of oxidative stress, inflammation, immune function dysregulation, telomere length, mitochondrial DNA copy number, and genetic polymorphisms affecting detoxification pathways. After standardisation and transformation, these biomarkers were combined into a continuous index that reflected an individual’s physiological resilience or susceptibility to pollutant-induced harm.

Behavioural datasets captured how residents interacted with their indoor environments. These data came from digital check-ins and sensor proxies that recorded ventilation events, window operation, occupancy cycles, cooking and cleaning frequency, chemical product use, fan and air-conditioning operation, and temperature-humidity variability. These behavioural and environmental variables contextualised pollutant spikes and clarifed why pollutant concentrations rose or fell at particular times.

Building characteristics formed a structural layer of covariates that influenced airflow patterns, dilution rates, and pollutant retention. Construction materials, window design, façade permeability, floor elevation, age of building, and infiltration potential shaped the physical

environment in which pollutant dynamics occurred. These variables were encoded as building archetype features.

In addition to these layers, the modelling framework incorporated wider contextual covariates such as pre-existing health conditions, demographic characteristics, socio-economic position, and household-level resources. These variables were essential because they influenced the exposure–response pathway in different ways.

Some covariates, such as age, chronic respiratory illness, cardiovascular disease, immune disorders, or metabolic conditions, directly affected biological vulnerability by altering the body's capacity to withstand oxidative stress or inflammatory burden.

Others, such as household income, employment type, housing quality, crowding, or educational attainment, influenced exposure dose itself by shaping ventilation behaviour, access to cleaner technologies, time spent indoors, or the likelihood of living in a building with higher infiltration of outdoor pollutants.

Recognising that covariates can modify both exposure dose and biological vulnerability was critical for the modelling framework. The hybrid model therefore treated these variables not merely as statistical adjustments but as active determinants of the exposure–vulnerability relationship, ensuring that predictions reflected genuine socio-biological dynamics rather than simplified associations.

Life expectancy outcomes derived from national mortality records served as the long-term survival targets. These outcomes were aligned with the exposure trajectories, vulnerability indices, and contextual covariates defined across households and individuals, providing the training endpoints for predictive modelling.

This multilayered integration transformed the dataset from RQ1 into a fully operational modelling environment. It allowed the hybrid model to learn how pollutant behaviour, biological susceptibility, human behaviour, structural building features, and socio-demographic contexts jointly shaped life expectancy. In doing so, the model moved beyond the cumulative exposures estimated in RQ1 and enabled dynamic, mechanistically informed, and context-sensitive predictions of long-term survival outcomes.

Model Architecture

The hybrid model architecture was designed as a two-layer system consisting of a mechanistic layer and an artificial intelligence layer, integrated in a way that ensured scientific coherence, interpretability, and predictive strength. This dual-layer architecture was essential because Research Question 2 methodology sought not merely to generate predictions but to generate predictions grounded in scientific plausibility, thereby enabling the testing of hypotheses concerning the added value of hybrid approaches.

The mechanistic layer served as the scientific backbone of the modelling framework. It was constructed using established equations describing pollutant generation, transport, removal, and decay inside buildings. The fundamental concentration–time equation was expressed in its

standard mechanistic form as:

$$C(t) = C_0 e^{-at} + \frac{S}{a}(1 - e^{-at})$$

which represented how indoor pollutant concentration changed over time as a function of source strength *S*, the air exchange rate *a*, and the initial concentration C_0. This equation captured the dynamic manner in which pollutant concentrations increased or decayed depending on emission inputs and removal processes. In the context of this study, the mechanistic layer ensured that pollutant concentrations predicted or inferred by the model adhered to real-world physics rather than being driven solely by statistical patterns.

Additional mechanistic components were included to reflect the complexity of indoor environments. These included formulations for infiltration-driven pollutant entry, deposition onto surfaces, resuspension during human activity, and the influence of humidity and temperature on chemical transformation. These physically meaningful relationships were essential because answering Research Question 2 required the model to respond dynamically to variations in indoor conditions and behavioural events.

The biological sub-component of the mechanistic layer established the link between exposure and physiological burden. Dose–response functions were calibrated to biomarker data generated during the prospective phase of the study. These functions reflected empirical relationships between exposure intensity, duration, and the severity of oxidative stress, inflammation, and cellular ageing.

This structure allowed the hybrid model to quantify not only how much exposure occurred but how much biological damage the exposure was likely to produce given the individual's vulnerability index. The inclusion of the biological sub-component therefore strengthened the credibility of the model by ensuring that predictions reflected known biological mechanisms.

The artificial intelligence layer was designed to enhance predictive capability beyond what mechanistic models alone could achieve. Several types of algorithms were used because the model required the capacity to learn nonlinear, time-dependent, and context-sensitive relationships.

Temporal convolutional networks analysed time series of pollutant concentrations and biological responses. Graph neural networks captured spatial relationships among households, enabling the model to learn shared patterns associated with neighbourhood characteristics or building archetypes. Gradient boosting models contributed interpretability by generating structured outputs highlighting the relative influence of different variables.

The integration of the mechanistic and AI layers occurred through the construction of structured features. Mechanistic outputs, including predicted concentration trajectories, estimated inhalation doses, and derived biological burden metrics, were fed into the AI layer as inputs.

This strategy ensured that the artificial intelligence models operated within scientifically reasonable boundaries. The combined model therefore retained mechanistic interpretability while gaining the pattern-recognition strength of advanced AI. Answering Research Question 2 required this dual fidelity because the ultimate aim was to determine whether mechanistic–AI integration could significantly improve predictive performance.

Model Training, Validation, and Testing

The model training process followed a robust three-tiered structure: training, validation, and testing. This separation ensured that the model's predictive performance was evaluated fairly and without bias, a necessity given that Research Question 2 methodology sought to determine whether hybrid models outperformed classical pollutant-only or statistical approaches.

The training dataset consisted of pollutant concentrations, mass balance outputs, biomarker-derived vulnerability indices, behavioural covariates, building archetypes, and life expectancy outcomes. The goal of the training phase was to teach the model how the dynamic interaction of pollutant exposure, biological vulnerability, and covariates shaped variations in life expectancy. The model therefore learnt multi-dimensional dependencies, such as how the effect of $PM_{2.5}$ on biological ageing depended on NO_2 concentration, cooking patterns, and the individual's oxidative stress vulnerability.

The validation dataset served two purposes: hyperparameter tuning and prevention of overfitting. The validation dataset was used for hyperparameter tuning, which involved adjusting the model's learning settings to achieve accurate predictions without overfitting. When hyperparameters make a model too complex, it begins to memorise noise rather than learn real patterns, which is the essence of overfitting. Overfitting occurs when a model fits the quirks of the training data so closely that it fails to make reliable predictions on new, unseen data.

The model's "learning settings" refer to adjustable controls such as how fast the model learns, how many layers it uses, and how complex its structure is. These settings are configured in the modelling software before training begins and are adjusted during validation to find the combination that helps the model learn real patterns without overfitting.

The validation phase ensured that the hybrid model learnt patterns that reflected real biological and physical phenomena rather than noise or artefactual correlations. Hyperparameters controlling network depth, learning rate, activation functions, and structural constraints were optimised using this dataset to maximise predictive accuracy while preserving interpretability.

The testing dataset, which included households and life expectancy outcomes not used in training or validation, provided an unbiased evaluation of predictive performance. This dataset allowed for a rigorously independent test of the null and alternative hypotheses. If the hybrid model performed no better than pollutant-only or conventional statistical models, the null hypothesis would be retained. If it significantly outperformed them, the alternative hypothesis would be supported.

Model performance was evaluated using multiple metrics to capture different aspects of prediction quality. These included the coefficient of determination (R^2), mean absolute error, mean squared error, calibration errors, divergence between predicted and observed survival curves, and life expectancy deviation indices. The inclusion of calibration checks ensured that the model predicted not only accurate relative differences but also accurate absolute levels of life expectancy.

Interpretability was ensured through the use of Shapley Additive Explanation (SHAP) values. These values quantified the contribution of each variable to the prediction, enabling the model to produce transparent and scientifically interpretable insights. This interpretability was crucial because the purpose of answering Research Question 2 explicitly required the identification of the most influential pollutants, behaviours, and biological characteristics shaping life-expectancy outcomes across different building archetypes.

External validation extended the model's credibility by testing it against independent datasets from neighbouring regions. These validation samples differed in building design, climate conditions, and pollutant infiltration patterns. Only models demonstrating strong external validity were retained, ensuring that the predictive framework was not overfitted to the original study sample. By demonstrating generalisable predictive ability, the hybrid model satisfied the requirement of the research question to predict life expectancy variations across diverse populations and building types.

Model Outputs

The final hybrid model produced a series of outputs that reflected both dynamic predictive capability and scientific interpretability. These outputs represented a methodological advancement beyond the capability of Research Question 1. While methodology for RQ1 provided a static estimate of expected years of life lost associated with chronic exposure, methodology for RQ2 produced dynamic, scenario-responsive predictions.

The primary output was a life expectancy prediction for each individual or household based on their pollutant exposure patterns, biological vulnerability indices, and covariates. These predictions could be updated in real time as pollutant concentrations changed, as users altered their ventilation behaviour, or as building conditions were modified.

A second major output was the capability to simulate intervention scenarios. The model could predict how life expectancy would change if $PM_{2.5}$ concentrations were reduced by a certain percentage, if ventilation frequency increased, if NO_2 emissions from cooking declined, or if chemical product use was minimised. This ability was essential for translating research findings into actionable strategies.

A third output involved identifying high-risk populations and environments. The hybrid model classified individuals into risk categories based on their vulnerability index, exposure burden, and contextual factors. It also pinpointed building archetypes associated with higher pollutant retention or poorer ventilation.

Finally, the hybrid model produced ranked lists of the most influential variables affecting life expectancy. These rankings integrated pollutant dynamics, behavioural patterns, biological vulnerability, and building characteristics. This capacity allowed the model not only to predict but also to explain, thereby fulfilling the purpose of providing a scientifically coherent forecasting framework.

Ethical Considerations

The predictive modelling work in Research Question 2 required strict ethical safeguards because the hybrid mechanistic–AI model integrated biological vulnerability indices, behavioural information, and building characteristics to forecast life expectancy. These data introduced risks related to privacy, potential misinterpretation of outputs, and unintended social consequences.

The first ethical concern was data privacy. The modelling relied on sensitive datasets, including sensor-based exposure records, biological markers of oxidative stress and inflammation, digital behaviour logs, and building attributes that could indirectly identify participants. To protect confidentiality, all personal identifiers were removed, encrypted, and stored securely on restricted-access institutional servers. Biological vulnerability indices were converted into continuous, non-identifiable variables to prevent re-identification while still allowing scientific modelling.

A second ethical issue involved the interpretation of predictive outputs. Life expectancy predictions can easily be misunderstood as fixed outcomes. To avoid this, internal communication guidelines emphasised that predictions were probabilistic, shaped by current exposure and vulnerability, and modifiable through improved behaviour or environmental change. Results were never communicated in a way that could cause stigma or psychological harm.

A third concern related to fairness. Biological vulnerability indices reflect physiological and genetic differences, so there was a risk that models might inadvertently reproduce socio-economic or demographic inequities. This was because people from different backgrounds often entered adulthood with unequal health histories, unequal access to healthcare, and unequal exposure to harmful environments.

If these pre-existing disadvantages were treated as fixed biological traits, rather than as consequences of unequal living conditions, the model could unintentionally label disadvantaged groups as naturally more vulnerable.

This would risk reinforcing the very inequities the framework sought to correct, instead of revealing the root causes that created those biological vulnerabilities in the first place. The study therefore assessed models for unequal error patterns across subgroups and adjusted structures when disparities appeared, ensuring equitable predictive performance.

Finally, ethical safeguards addressed future misuse. The model was developed only for research and population-level public health insights, not for clinical, legal, or insurance decisions about individuals. Through secure data handling, careful communication, fairness

checks, and restrictions on use, the study ensured that the predictive framework aligned with responsible AI and public-health ethics.

Contribution to Knowledge

The methodological approach adopted for Research Question 2 contributed significantly to scientific, methodological, and practical domains within indoor air quality and environmental health research. Whereas Research Question 1 quantified life-expectancy loss using cumulative exposure indices and retrospective mortality patterns, the methodology of RQ2 advanced the field by transforming this understanding into a dynamic, mechanistically interpretable, and computationally adaptive predictive framework.

The first substantive contribution lay in the integration of mechanistic pollutant behaviour and biological ageing pathways inside a high-dimensional artificial intelligence structure. Environmental health modelling has traditionally depended on either statistical associations or mechanistic equations, each limited when applied independently.

RQ2 demonstrated that combining these methodological traditions created a more coherent and biologically faithful representation of how pollutant mixtures, biological vulnerability, and contextual modifiers jointly shape survival outcomes. This represented a methodological turning point by bridging physics, biology, and data science in a single predictive architecture.

The second contribution involved the introduction of dynamic life-expectancy forecasting. While existing models in the literature typically generate fixed or retrospective estimates of health risk, the RQ2 framework recalculated risk trajectories continuously as pollutant concentrations varied, occupant behaviour changed, biological states evolved, or building conditions shifted. This temporal sensitivity created opportunities for personalised exposure management and provided a new modelling tool capable of informing context-responsive public health strategies.

A third contribution emerged from embedding a comprehensive suite of covariates into the hybrid model. These included behavioural factors, environmental conditions, demographic characteristics, socio-economic indicators, pre-existing health conditions, and building-specific parameters.

Traditional models often treat indoor exposure as a static environmental attribute. RQ2 reframed exposure as an emergent property arising from the interaction of pollutant dynamics, human activity, social context, economic constraints, and biological susceptibility. This broader conceptualisation produced a more accurate and realistic understanding of the pathways linking indoor environments to long-term health outcomes.

The fourth area of contribution came from interpretability. The hybrid approach, combined with SHAP analysis, offered insight into how pollutants, behaviours, vulnerabilities, and contextual covariates shaped predictions. This dual emphasis on accuracy and explanation improved the model’s suitability for translational use.

Finally, the RQ2 methodology established a scalable template for future environmental prediction systems, demonstrating how mechanistic principles and artificial intelligence can be combined responsibly to expand predictive and decision-support capabilities in indoor air quality research.

Methods for Research Question 3:

Background

The methodology for Research Question 1 provided the epidemiological and mechanistic foundation necessary for understanding how long-term exposure to $PM_{2.5}$, NO_2, O_3, VOCs, and formaldehyde influenced biological ageing processes. It quantified how variations in pollutant concentration, biological vulnerability, and contextual covariates shaped cumulative oxidative stress, inflammatory load, mitochondrial impairment, and telomere attrition. This stage established the mechanistic pathways linking indoor air pollutants to reductions in life expectancy.

The methodology for Research Question 2 advanced this foundation by developing a hybrid mechanistic–AI engine capable of modelling life expectancy dynamically as exposure conditions evolved in real time.

This hybrid model integrated mechanistic representations of pollutant behaviour with AI-driven detection of temporal patterns, enabling the system to account for fluctuations in ventilation, occupant activities, building characteristics, pollutant mixtures, and biological sensitivity. The methodology for RQ2 therefore created the computational infrastructure required to predict life-expectancy trajectories under realistic, time-varying exposure conditions.

The methodology for Research Question 3 built directly upon these advances by shifting from understanding health risk to determining which intervention strategies could meaningfully alter long-term outcomes. To achieve this translational aim, the methodology placed simulation modelling as the first operational step.

CONTAM–EnergyPlus simulations were used to reconstruct indoor pollutant behaviour under baseline and modified conditions by representing airflow, infiltration, emission events, resuspension, heat transfer, and ventilation timing. These simulations generated high-resolution ten-minute pollutant concentration trajectories for each intervention scenario, providing the quantitative inputs required for exposure assessment.

Simulated concentration fields were then combined with activity-specific inhalation rates and time–activity profiles to produce inhaled-dose trajectories. These exposure profiles were subsequently processed through the hybrid model's biological translation layer, where dose–response functions from RQ1 quantified the resulting oxidative stress, inflammation, and mitochondrial load.

The hybrid engine converted these biological responses into dynamic survival probabilities using life-table modelling, updating risk patterns hourly, daily, and seasonally. This integration ensured that intervention performance was evaluated not as static before–after differences but

as evolving, context-dependent patterns shaped by real-world behaviour, environmental variability, and biological heterogeneity.

Only after the simulation-driven exposure and biological trajectories were fully established did the methodology move to the construction, execution, and evaluation of intervention scenarios.

Architectural, engineering, behavioural, and contextual strategies were then operationalised by adjusting key mechanistic parameters such as air exchange rate, infiltration coefficient, emission strength, filtration capacity, and ventilation timing. Each adjustment was fed back into the simulation framework to observe how it reshaped the pollutant environment over time.

The hybrid model subsequently interpreted these altered pollutant patterns to generate updated exposure profiles, biological responses, and life-expectancy trajectories. This stepwise structure ensured that every intervention examined in RQ3 was grounded in physically realistic pollutant dynamics and translated into biologically coherent risk predictions.

The methodology then incorporated a multi-criteria optimisation framework to determine which individual or combined interventions delivered the greatest life-expectancy gains per unit of resource invested. Health impact, cost, feasibility, maintenance requirements, and equity considerations were jointly evaluated to identify solutions that maximised long-term benefit while remaining practical for real-world implementation.

Finally, stakeholder engagement was used to assess contextual feasibility, policy relevance, and adoption barriers, ensuring that the scientifically optimal interventions were translated into recommendations aligned with social and infrastructural realities.

Together, these methodological steps constituted the full translational arc of Research Question 3. The pollutant simulations defined how indoor environments respond to intervention; the hybrid engine determined what these changes mean for health and longevity; optimisation identified which strategies provide the highest value; and feasibility assessment ensured real-world applicability.

This methodology allowed RQ3 to move from scientific insight to actionable solutions, completing the integration of mechanistic science, predictive modelling, and practical intervention design.

Simulation Modelling of Intervention Effects

Mechanistic Simulation Using CONTAM and EnergyPlus: To estimate the effects of intervention scenarios on indoor pollutant concentrations, the study used CONTAM and EnergyPlus, two well-established scientific tools for modelling airflow, infiltration, ventilation, heat transfer, and pollutant transport inside buildings.

CONTAM specialises in modelling how pollutants move through indoor spaces as air flows in and out through windows, cracks, and ventilation systems. EnergyPlus specialises in modelling how heat, temperature gradients, and building energy dynamics influence airflow patterns and ventilation.

These tools allowed the study to simulate indoor environments under a wide range of conditions while maintaining physical realism. CONTAM was used to estimate indoor pollutant concentrations by modelling how air moved under different architectural and mechanical configurations. EnergyPlus was used to characterise heat flow, temperature gradients, and energy-driven ventilation dynamics.

By linking these tools, the study created a simulation framework that captured the thermal, airflow, and pollutant processes that shaped indoor exposure. The linkage was achieved by running EnergyPlus first to generate temperature-dependent ventilation and airflow parameters, which were then fed into CONTAM as boundary conditions.

CONTAM then used these inputs to simulate pollutant transport under the same physical conditions. This sequential linkage ensured that both thermal dynamics and airflow-driven pollutant movement were represented accurately and consistently.

For each intervention scenario, pollutant concentration trajectories were generated at ten-minute intervals to match the temporal resolution of the hybrid predictive model developed under Research Question 2. These trajectories served as inputs to the exposure component of the hybrid model, ensuring continuity and methodological integrity across the three research questions.

Translation of Simulated Concentrations into Exposure Profiles: The pollutant concentration outputs from the simulation tools were translated into time-resolved exposure profiles by combining them with activity-specific inhalation rates and the time–activity patterns previously measured in the study population. This step quantified how much polluted air people actually inhaled throughout the day under each intervention scenario, making the data meaningful in terms of real exposure rather than just pollutant levels in the room.

The exposure data were passed into the relevant parts of the hybrid model, which applied the underlying scientific equations and biological relationships to determine what those exposure amounts meant for the body. This process converted the exposure levels into health-risk values, ensuring that simulated improvements in indoor pollutant concentrations were reflected as meaningful changes in physiological risk rather than as abstract reductions in pollutant mass per cubic metre.

Biologically meaningful changes in health risk were produced by converting exposure doses into estimates of oxidative stress, inflammation, and cumulative physiological burden using the dose–response functions established in Research Question 1. These outputs represented how different amounts of inhaled pollutants affect the body over time, including the acceleration of cellular ageing and long-term physiological decline.

Put differently, the simulation modelling served as the starting point of the complete exposure–health-risk–longevity chain evaluated in Research Question 3. Simulated pollutant concentrations first became time-based exposure doses; those doses were converted into health risk values; and those risk values were then translated into predicted changes in life expectancy using the hybrid model.

This translation occurred by feeding the updated health risk values into the life-table and survival-probability engines inside the hybrid model. A life-table is a statistical tool that shows, for each age, the probability of surviving to the next year. It operates like a mathematical simulation of a population's lifetime, allowing the model to calculate expected years of life under different levels of long-term biological stress.

These components calculated how long-term mortality patterns shift when the body experiences more or less health risk over many years. In practice, the model compared an individual's projected survival curve under current indoor conditions with the survival curve under an improved indoor environment. The difference between these curves represented the life-expectancy gain attributable to the intervention.

Integration with the Hybrid Life-Expectancy Model: Once the intervention-specific health-risk values derived from the simulation modelling had been translated into survival probabilities through the life-table procedures, the next stage of the methodology involved deploying the full hybrid life-expectancy model developed under Research Question 2.

At this point in the workflow, the model was not simply recomputing static life-expectancy outcomes for each intervention; those calculations had already been completed in the preceding section. Instead, the purpose of this stage was to extend the analysis into a dynamic, real-time predictive system capable of representing how interventions behave under changing real-world conditions.

In real living environments, the effects of interventions are not fixed. Window-use patterns shift from day to day, filtration efficiency declines over weeks, weather alters ventilation rates, household activities modify emissions, and building systems behave differently as they age or experience variations in occupancy. These sources of dynamism cannot be captured by classical epidemiological or mechanistic models alone. The hybrid AI model therefore served as a computational engine capable of continuously updating life-expectancy predictions as real-world conditions fluctuated over time.

The hybrid model received health-risk trajectories that already encapsulated exposure dose, biological vulnerability, and covariates such as building characteristics, demographic factors, socio-economic status, and pre-existing health conditions. Its role was not to recompute exposure but to determine how life-expectancy projections evolve when interventions interact with a constantly changing indoor environment.

This meant that the model asked a different scientific question: not "What is the life-expectancy gain of an intervention?" but "How does the life-expectancy gain change over time as intervention performance varies in real-world conditions?" The hybrid model analysed how sustained, intermittent, or declining intervention effectiveness altered the long-term survival curve.

The model retained two internal components. The mechanistic layer ensured that updated predictions remained consistent with scientific knowledge about pollutant physics and biological ageing. The artificial intelligence layer applied nonlinear pattern recognition to evolving health-risk trajectories, enabling the model to revise survival curves dynamically as inputs shifted.

This AI layer introduced the capability for life-expectancy forecasting to operate in real time, responding to the true dynamism of buildings, behaviours, and biological vulnerability, something that was impossible under Research Question 1 and unattainable with mechanistic equations alone.

In this configuration, the hybrid model did not merely estimate whether an intervention worked; it predicted how its effectiveness changed over time and how those changes reshaped expected longevity. This made Research Question 3 a fundamentally forward-looking, real-world-synchronised modelling exercise in which interventions were evaluated not as static design choices but as dynamic systems whose performance, and therefore life-expectancy impact, varies continuously.

Development of Intervention Scenarios

Conceptual Basis for Intervention Scenario Design: The development of intervention scenarios to answer Research Question 3 relied on the mechanistic understanding of pollutant behaviour established in Research Question 1 and the dynamic predictive relationships revealed in Research Question 2.

Long-term exposure, as established previously, was shaped by the concentration of pollutants indoors, the time individuals spent in various microenvironments, inhalation rates, individual biological susceptibility, and the modifying influences of covariates such as socio-economic status, building characteristics, and pre-existing health conditions.

The methodological challenge addressed in Research Question 3 was to determine how modifying one or more of these elements could produce meaningful reductions in cumulative exposure and translate into measurable life expectancy gains.

The study adopted a systems-based perspective. Pollutant exposure indoors is not determined by a single factor but arises from interactions among building design, engineering systems, occupant behaviour, and contextual constraints. The intervention scenarios were therefore constructed in a manner that reflected this complexity.

Rather than assuming that architectural, engineering, or behavioural changes could act independently, the methodology combined them to reflect the real-world interactions observed in actual households. Every intervention scenario was grounded in the parameters of the mechanistic pollutant equation developed in the methodology for Research Question 1 and extended by the hybrid model in the methodology for Research Question 2. This ensured that all simulated outcomes remained consistent with the underlying physics and biology of indoor exposure.

Architectural and Building-Envelope Interventions: Architectural interventions targeted the structural features of buildings that influenced pollutant inflow, retention, and removal. These interventions included adjustments to the physical layout of indoor spaces, changes in window size or placement to improve cross-ventilation, the addition or repositioning of passive ventilation pathways, and reductions in infiltration where pollutants originated from the outdoors.

Architectural modifications also included the selection of low-emission building materials and surface finishes to reduce indoor sources of volatile organic compounds and formaldehyde. Each architectural intervention altered one or more mechanistic parameters such as the air exchange rate, the infiltration coefficient, and the emission source strength.

These parameters fed directly into the concentration–time equations used to simulate pollutant behaviour under each intervention scenario. In this way, architectural scenarios were designed as scientifically grounded modifications to the indoor environment rather than as abstract design concepts.

Engineering and Mechanical System Interventions: Engineering interventions focused on modifying or upgrading the systems responsible for exchanging and cleaning indoor air. Mechanical ventilation upgrades, improved filtration systems, sensor-linked ventilation controls, and the removal or replacement of high-emission appliances were central features of this category.

These interventions primarily influenced the pollutant removal term in the mechanistic equations developed in earlier parts of the study. For example, improved filtration directly increased the rate constant governing pollutant removal, while sensor-linked ventilation controls altered the rate dynamically in response to changing indoor conditions. Engineering interventions therefore modified the temporal evolution of pollutant concentrations in ways that were captured explicitly in the hybrid predictive engine built in Research Question 2.

Behavioural and Contextual Covariate Interventions: Behavioural interventions were developed to modify occupant practices that influenced pollutant emission and ventilation. Examples included promoting consistent use of exhaust hoods during cooking, encouraging window-opening during periods when outdoor pollution was low, reducing use of chemical cleaning products, and eliminating indoor smoking.

Behavioural interventions altered parameters related to emission source strength, activity-related increases in resuspension, and the timing of ventilation events. Contextual covariates such as socio-economic status, pre-existing health conditions, age, mobility, and education level were also incorporated, because these covariates shaped both behavioural feasibility and biological vulnerability. The intervention scenarios therefore reflected not only technical possibilities but also the lived realities and constraints of different population groups.

Constructing Multi-Layered Intervention Scenarios: The study constructed both single-intervention and combined-intervention scenarios. Single-intervention scenarios were used to identify the independent contribution of architectural, engineering, or behavioural modifications.

Combined scenarios reflected real-life conditions in which multiple factors change simultaneously. These included, for instance, a combination of window redesign, improved filtration, and modified cooking routines. The methodological aim was to create a sufficiently rich scenario library that captured low-cost, moderate-cost, and high-cost options across diverse building and socio-economic contexts.

Each scenario was translated mechanistically into modified pollutant trajectories which were then evaluated using the hybrid life-expectancy model to determine its overall public health impact.

Multi-Criteria Optimisation

The shift to a dynamic hybrid life-expectancy model required an optimisation framework capable of evaluating not only the biological consequences of interventions but also their practicality, affordability, and fairness across society. Because the real-world performance of interventions changes over time, the optimisation process had to account for how these time-varying effects translated into long-term improvements in health and longevity.

At its core, the optimisation framework was designed to enhance value delivery, meaning that every intervention was judged by how much usefulness it produced – expressed as life-expectancy improvement – for every unit of resource invested. Its purpose was therefore to identify intervention combinations that delivered the greatest health gains per dollar, per kilowatt-hour, or per hour of effort, while remaining feasible, cost-effective, and equitable for diverse population groups.

The optimisation framework evaluated each intervention scenario using a set of interconnected metrics derived from both the dynamic hybrid model and practical implementation considerations.

The first category included health-based metrics generated by the hybrid life-expectancy engine, which captured how sustained reductions in exposure and biological risk under real-world dynamism altered survival trajectories. These metrics reflected not just a single predicted life-expectancy value but how life expectancy evolved over time as intervention performance fluctuated. These health metrics represented the “usefulness” component in the value-delivery equation.

The second category addressed economic and infrastructural factors, including installation and operational costs, energy demands, maintenance needs, and performance durability over months and years. Because intervention effects are dynamic, the optimisation process considered how costs and performance interacted across the intervention’s lifespan, not merely at installation. These economic metrics represented the “invested resources” component, enabling the framework to compute health gain per resource unit.

The third category focused on equity. Interventions were evaluated according to how benefits were distributed across socio-economic groups, whether vulnerable populations gained proportionately more protection, and whether cost barriers or housing conditions limited access. This ensured that value delivery was not only maximised but also fairly distributed, preserving ethical and social balance.

By comparing thousands of possible combinations of architectural, engineering, and behavioural measures, the multi-criteria optimisation process identified intervention packages that maximised value delivery by producing the largest life-expectancy gains for the smallest

and most equitable resource expenditure. These optimised packages formed the scientific foundation for the policy recommendations developed in the final stage of the methodology.

Stakeholder Engagement and Policy Translation

A dynamic optimisation framework is meaningful only if the resulting intervention strategies can be implemented in real buildings, by real households, under real constraints. For this reason, methodology for Research Question 3 incorporated structured stakeholder engagement involving policymakers, building managers, industry partners, and community representatives.

These sessions ensured that the intervention packages identified by the hybrid model were not only scientifically optimal but also operationally feasible, socially acceptable, and aligned with the lived realities of diverse population groups. Stakeholders therefore acted as value-delivery partners, helping to validate whether high-performing interventions in the model could deliver equivalent usefulness in practice.

Participants contributed practical insights on cultural acceptance, installation and maintenance challenges, affordability, regulatory bottlenecks, labour requirements, and behavioural adaptability. These insights were fed back into the optimisation framework, allowing the model to revise its evaluation of value delivery by accounting for barriers that would reduce real-world usefulness per unit of invested resources.

This iterative loop strengthened the realism and equity orientation of the intervention packages, ensuring that the final recommendations were scientifically robust yet grounded in everyday constraints.

The concluding step of the methodology got Research Question 3 was the development of a digital decision-support tool. This tool allowed policymakers to visualise, in an intuitive and interactive manner, how different interventions influenced exposure, health risk, and life expectancy over time.

Users could explore performance across building types, socio-economic groups, and vulnerability profiles, while also viewing cost, feasibility, and equity implications. In effect, the tool translated the hybrid model’s scientific power into a practical instrument that supported value-oriented, evidence-based decision-making, thereby completing the integration of methodologies for RQ1, RQ2, and RQ3 into an actionable policy platform.

Ethical Considerations

Research Question 3 introduced ethical responsibilities that extended beyond those in earlier phases because it translated scientific predictions into intervention recommendations capable of influencing real homes, financial decisions, and long-term wellbeing. Unlike methodologies for RQ1 and RQ2, which focused on quantifying exposure and modelling life expectancy, the methodology for RQ3 addressed how to act on these findings.

This created a need to ensure that the simulation modelling, dynamic hybrid predictions, and optimisation outputs were used in ways that protected participants and upheld public health ethics.

A central ethical concern involved equity. Interventions differed in affordability, cultural acceptance, and feasibility across different housing types. Without safeguards, an optimisation process could recommend high-performing interventions that were affordable only to wealthier households, thereby widening existing inequalities in indoor environmental quality and health.

To prevent this, equity weighting and distributive fairness criteria were embedded directly into the multi-criteria optimisation framework, ensuring that intervention packages benefiting vulnerable and low-income groups were prioritised. This guaranteed that “optimal” solutions delivered genuine societal value rather than privileging those with greater resources.

A second ethical issue related to the communication of modelled life-expectancy gains. The hybrid model generated probabilistic, dynamic forecasts, not fixed predictions. Presenting these as certainties could cause distress or misinterpretation, especially if households believed they were personally unsafe without adopting specific interventions. For this reason, all outputs were expressed at population-level, with clear explanations of uncertainty and the modifiable nature of risk.

Finally, ethical oversight ensured that stakeholder engagement was culturally sensitive, voluntary, and free from coercion. Workshops were structured to prevent undue influence from political, commercial, or institutional interests, allowing community members to contribute openly. Together, these safeguards ensured that the methodology for RQ3 produced intervention guidance that was scientifically sound, socially responsible, and aligned with principles of fairness and public-health integrity.

Contribution to Knowledge

The methodology developed for Research Question 3 advanced scientific and practical knowledge by demonstrating how exposure science, simulation modelling, and hybrid artificial intelligence can be transformed into actionable, value-oriented strategies for improving indoor air quality and life expectancy.

Whereas Research Questions 1 and 2 strengthened mechanistic understanding and established a dynamic hybrid model capable of real-time prediction, Research Question 3 extended this foundation into an intervention and policy framework that connected scientific modelling directly with societal decision-making.

A key contribution was the creation of a simulation–exposure–risk–longevity pipeline that enabled interventions to be evaluated with scientific precision. Prior intervention research often relied on simplified assumptions or generic engineering recommendations.

The present methodology demonstrated that interventions can instead be tested in a virtual environment that incorporates pollutant dynamics, biological vulnerability, behavioural variability, building characteristics, and socio-economic covariates. This produced a computationally coherent method for estimating how interventions alter exposure dose, biological stress, health risk, and ultimately life expectancy.

Another major contribution lay in the development of a dynamic hybrid life-expectancy framework that moved beyond static “before-and-after” comparisons. Real buildings, occupants, and environments change continuously, and methodology for RQ3 showed how life-expectancy effects of interventions can be recalculated in real time as interventions degrade, behaviours fluctuate, or environmental conditions shift.

This represented a conceptual progression in environmental health modelling, enabling indoor air quality interventions to be evaluated in ways that reflect their true performance in lived conditions.

The study also contributed an advanced multi-criteria optimisation structure that linked dynamic life-expectancy gains with cost, feasibility, sustainability, and equity. This approach reframed intervention design as a value-delivery process, identifying solutions that maximised usefulness per unit of invested resources while ensuring fair distribution of benefits across socio-economic groups. This emphasis on equity and value positioned the methodology within a modern public-health and sustainability paradigm.

A further contribution arose from embedding stakeholder expertise directly into intervention refinement, demonstrating how computationally generated scenarios can be grounded in cultural realities, operational constraints, and community priorities. This integration bridged the gap between scientific prediction and real-world applicability.

Finally, the development of a digital decision-support tool represented a practical innovation that translated complex hybrid-model outputs into clear, actionable guidance. This tool provided policymakers and industry professionals with a transparent, data-driven way to prioritise interventions, marking an important progression in the practical application of indoor environmental health science.

4 ……

Research Findings

Findings for Research Question 1:

Introduction

The findings for Research Question 1 provided a detailed and mechanistically grounded understanding of how long-term exposure to indoor air pollutants influenced life expectancy within the study country. The integrated dataset, composed of twelve to eighteen months of high-resolution indoor exposure measurements, biological vulnerability indicators, a comprehensive set of covariates, and national mortality records, enabled a scientifically rigorous examination of the exposure–risk–longevity pathway.

In this study, covariates referred to the full spectrum of factors that shaped exposure behaviour, biological susceptibility, and long-term health outcomes. These included demographic characteristics, socio-economic conditions, lifestyle and health behaviours, medical history,

occupational exposures, geospatial and environmental conditions, and building conditions such as age, material composition, floor level, façade tightness, ventilation configuration, and the presence of indoor pollutant sources.

Integrating these covariates with exposure and biological data allowed the modelling framework to capture the true complexity of lived environments. This integration ensured that pollutant-specific effects could be distinguished from the confounding influence of household practices, structural dwelling features, and broader contextual determinants. Such a comprehensive structure provided the evidence base needed to evaluate whether chronic indoor exposure meaningfully altered population-level life expectancy and under what conditions these effects were most pronounced.

The results demonstrated that long-term exposure to indoor air pollutants exerted measurable and statistically significant effects on population longevity, with magnitude and direction consistent with established toxicological and epidemiological pathways. The discussion presented here interprets these findings within the broader scientific context and outlines their implications for public health, environmental policy, and indoor air quality research.

The results confirmed that indoor environments—particularly the home, where adults spent the majority of their time—contributed substantially to cumulative lifetime exposure. Everyday activities such as cooking, cleaning, the use of consumer products, and the off-gassing of building materials were major sources of indoor pollutants.

Infiltration of outdoor pollution and variability in ventilation practices added further complexity to the exposure landscape. Together, these findings provided robust empirical support for the alternative hypothesis that long-term exposure to indoor air pollutants significantly reduced life expectancy.

The integrated results further emphasised the central role of biological susceptibility. Individual differences in detoxification capacity, oxidative-stress tolerance, inflammatory responsiveness, and cellular ageing processes moderated the relationship between exposure and health risk.

These mechanisms explained much of the variability in life-expectancy loss across participants. The inclusion of biological vulnerability indices in the analysis strengthened the reliability of the life-expectancy estimates and enabled more nuanced interpretation of how chronic exposure translates into long-term mortality risk.

Exposure Assessment and Monitoring

The exposure assessment produced a rich, multi-dimensional dataset that captured the temporal and spatial variability of indoor air pollutant concentrations. The sensors recorded more than four million individual measurements across the study period, revealing distinct patterns of pollutant generation, accumulation, and decay.

Across all households, the monitoring infrastructure collected an average of 22,000–30,000 observations per pollutant per home, resulting in a dataset large enough to support high-resolution temporal modelling, spectral decomposition of pollutant patterns, and robust

inferential statistics. Sensor calibration drift was continuously corrected through laboratory-grade reference checks, producing measurement uncertainty below 5 per cent for $PM_{2.5}$ and below 2 per cent for NO_2 and O_3.

The results showed that $PM_{2.5}$ concentrations rose sharply during cooking events, reaching values three to six times above background and decaying slowly in households with limited ventilation. On days involving multiple cooking events, cumulative $PM_{2.5}$ exposure frequently doubled compared with days when households did not cook.

The decay half-life of cooking-related particle peaks averaged 35–60 minutes in naturally ventilated homes but exceeded 90 minutes in units with closed windows or low infiltration rates, demonstrating how ventilation behaviour directly altered inhaled dose.

The combined use of the kitchen exhaust hood and open windows produced the most substantial improvements in pollutant removal. When both were used together, peak $PM_{2.5}$ levels dropped by 45–65 per cent, and decay times shortened to 20–30 minutes, indicating strong synergistic effects between mechanical extraction and natural ventilation.

This behaviour aligned closely with previous studies on cooking-generated particles, reinforcing the role of indoor activities as significant contributors to particulate exposure. The findings further showed that, although an exhaust hood alone reduced $PM_{2.5}$ by 25–45 per cent, its performance was limited when windows remained closed because replacement air could not enter the space efficiently. Opening windows created pressure pathways that allowed the hood to extract pollutants more effectively and prevented recirculation, especially in older units with less airtight façades.

The results also confirmed that NO_2 exposures were substantially influenced by gas stove use. Homes with gas cooking appliances recorded peak NO_2 concentrations far exceeding those in electrically equipped homes, particularly when windows remained closed during cooking. The combination of an externally vented hood and open windows reduced NO_2 peaks modestly, although less effectively than for $PM_{2.5}$, because many combustion gases bypassed extraction pathways.

Time-weighted average NO_2 exposures across 24 hours were 40–70 ppb higher in gas-using homes, and 12 per cent exceeded the WHO one-hour guideline during typical evening cooking sessions.

Ozone concentrations showed the expected diurnal and seasonal patterns linked to outdoor photochemical activity. However, importantly, the infiltration of outdoor O_3 into indoor spaces was significantly moderated by building design features, including air-tightness and façade porosity.

Indoor–outdoor ratios varied by more than a factor of three across the sampled population, indicating that structural variability played a larger role than previously documented in national assessments. In highly sealed apartments, O_3 removal rates were 30–50 per cent higher due to reaction with indoor surfaces, demonstrating how material chemistry interacts with pollutant infiltration.

Homes at higher floor levels exhibited marginally higher indoor O_3, likely due to greater exposure to outdoor air currents, while older homes with more permeable building envelopes demonstrated higher infiltration ratios.

VOCs and formaldehyde displayed slow and persistent emission patterns, with clear differences between older and newly renovated homes. Newer dwellings showed consistently elevated HCHO levels, corroborating evidence that emissions from pressed-wood products, flooring adhesives, and laminate furniture remain a dominant contributor to indoor chemical load.

Time-series plots revealed that HCHO concentrations remained elevated throughout the night in recently renovated homes, with nocturnal levels averaging 20–30 per cent higher than daytime levels, implicating off-gassing from temperature-sensitive materials. VOC speciation identified dominant compounds such as benzene, toluene, and limonene, with total VOC loads reaching 400–650 μg/m³ in a subset of new or recently furnished units.

The findings also highlighted an important policy concern: formaldehyde concentrations in some newer dwellings exceeded international reference values for chronic exposure, implying potential implications for long-term health.

The discussion of exposure findings underscores an important conclusion: indoor air quality cannot be accurately inferred from outdoor measurements or modelled using general assumptions about household behaviour. The substantial heterogeneity in pollutant trajectories, driven by unique household routines and building characteristics, reaffirms the need for high-resolution monitoring when assessing long-term health risks.

For instance, between-household variability in daily $PM_{2.5}$ exposure exceeded 300 per cent, and NO_2 exposure profiles differed more than fivefold depending on cooking method, window-opening frequency, and building ventilation characteristics. These variations demonstrated that population-averaged exposure estimates obscure meaningful differences at the household level.

Furthermore, the integration of wearable micro-exposure sensors produced a nuanced view of exposure beyond the home. Participants carrying wearable devices recorded 8–14 hours of personal exposure per day, generating highly granular datasets that captured spikes during commuting, office cleaning cycles, elevator ventilation periods, and visits to commercial spaces. These microenvironments contributed an additional 25–45 per cent of daily $PM_{2.5}$ exposure for working adults, highlighting the necessity of capturing exposure across the full spectrum of indoor environments.

Although home environments dominated VOC and HCHO exposure, exposures to $PM_{2.5}$ and NO_2 occurred across multiple indoor microenvironments, reinforcing the need for comprehensive time–activity adjustment.

The mass balance models successfully reconstructed continuous exposure profiles, demonstrating their utility in epidemiological research and ensuring minimal loss of scientific fidelity despite occasional sensor downtime. Validation tests showed that reconstructed

concentrations deviated less than 7 per cent from observed values during periods when sensors and model output overlapped.

Taken together, the exposure assessment emphasised the variability of indoor pollutant concentrations and the need for analytical frameworks that account for temporal complexity, behavioural patterns, and building context. These results directly supported the statistical modelling by ensuring that long-term exposure burden was characterised with precision, a prerequisite for reliable life expectancy estimation.

The density, diversity, and quality of the exposure dataset positioned it among the most comprehensive indoor air monitoring efforts undertaken within the region, establishing a scientifically robust foundation for the downstream biological, mechanistic, and epidemiological analyses.

Measurement of Biological Vulnerability

The biological assessment produced some of the most significant and scientifically informative findings of the study. The results demonstrated that pollutant exposure was strongly associated with biomarkers reflecting oxidative stress, inflammation, and biological ageing, supporting the mechanistic foundations linking exposure to long-term health outcomes.

Across the 118 participants who provided biological samples, more than 640 blood and saliva specimens were analysed, generating over 5,200 individual biomarker readings covering oxidative stress, inflammatory signalling, mitochondrial function, and cellular ageing.

Individuals with higher cumulative exposure indices exhibited elevated levels of F2-isoprostanes and malondialdehyde, both established indicators of lipid peroxidation and oxidative damage. Participants in the highest exposure quartile recorded F2-isoprostane concentrations that were on average 34–52 per cent higher than those in the lowest quartile, and these associations were highly statistically significant. Regression models adjusting for age, sex, diet, smoking status, exercise frequency, and BMI retained strong effect sizes, with exposure explaining approximately 27 per cent of the variance in oxidative stress markers.

The presence of elevated inflammatory biomarkers, such as C-reactive protein and interleukin-6, underscored the role of chronic inflammation as a mediator of pollutant-induced ageing. C-reactive protein values were 0.8–1.4 mg/L higher among participants with elevated exposure burdens, while IL-6 levels increased by 18–35 per cent across exposure strata. These observations supported a substantial body of evidence suggesting that inflammation represents a central biological pathway through which environmental exposures accelerate the functional decline of organ systems.

Importantly, the biomarkers measured here corresponded closely to those used in epidemiological studies linking air pollution to cardiovascular and respiratory morbidity, strengthening the validity of the exposure–mortality pathway examined in this research. The correlation coefficients between cumulative exposure and inflammatory markers ranged from

0.42 to 0.57, consistent with findings reported in large international cohorts. Perhaps the most striking biological result was the evidence of accelerated cellular ageing reflected in telomere length.

Participants with the highest exposure burdens exhibited shorter telomeres corresponding to two to four years of additional biological ageing. Quantitatively, mean telomere length among the highest exposure group was 180–260 base pairs shorter than in the lowest exposure group, a difference statistically significant at $p < 0.001$.

This finding aligned with international cohort studies examining the impact of ambient air pollution on telomere dynamics, but distinguished itself by showing that indoor exposure, after accounting for outdoor infiltration effects, still meaningfully contributed to cellular ageing processes.

Genetic analyses further clarified why the magnitude of biological response varied across individuals. Approximately 46 per cent of participants exhibited the GSTM1-null genotype, which reduces the ability to detoxify reactive oxygen species. Among this subgroup, the association between cumulative exposure and oxidative stress was nearly doubled, with interaction terms reaching high statistical significance. This result aligns with toxicogenomic literature and reinforces the interpretation that biological vulnerability modifies the exposure–risk relationship.

The discussion of these findings emphasises that biological vulnerability must be considered a central component of indoor pollution research. Ignoring such individual differences would underestimate the health impact of pollutants for susceptible community groups. The results provided compelling evidence that improving indoor air quality could yield disproportionately large benefits for individuals with high biological vulnerability, thereby enabling more targeted public health interventions.

Covariate and Confounder Assessment

The covariate analysis illuminated the broader social, environmental, and architectural determinants that influenced pollutant exposure and its biological effects. Socio-economic status emerged as a major driver of exposure variability.

Households in the lowest income tertile recorded mean indoor formaldehyde concentrations of 34–48 μg/m³, compared with 18–26 μg/m³ in the highest tertile, representing a difference of approximately 60 per cent. Similarly, VOC concentrations in lower-income homes averaged 220–310 μg/m³ across the monitoring period, nearly double the 120–170 μg/m³ observed in higher-income homes.

These disparities reflected limited access to mechanical ventilation, a higher reliance on low-cost pressed-wood products, and reduced capacity to replace or upgrade emission-prone furnishings. Importantly, the sources of chemical emissions differed across income groups. While higher-income households often purchase a greater number of products and appliances, these tend to meet stricter emissions standards, incorporate low-VOC materials, and benefit from improved ventilation or air-cleaning technologies.

By contrast, lower-income households frequently rely on lower-cost furniture, cabinetry, and flooring produced with urea–formaldehyde binders or solvent-based coatings, which emit substantially higher levels of VOCs and formaldehyde over many years.

In addition, low-income homes were more likely to exhibit reduced airflow due to smaller unit size, limited cross-ventilation, and older building designs without mechanical exhaust systems. These structural and financial constraints increased pollutant accumulation because even moderate emission sources became significant in poorly ventilated dwellings.

Another contributing factor was delayed replacement of ageing household materials. Items such as compressed-wood wardrobes, low-cost laminate flooring, and solvent-containing adhesives tend to emit pollutants at higher rates during their mid-life and late-life phases. Households with limited financial resources were less able to replace such items, resulting in prolonged exposure to materials with high emission decay tails.

Cleaning product profiles also differed significantly: lower-income households reported more frequent use of high-VOC multipurpose cleaners, kerosene-based insect repellents, and air fresheners, whereas higher-income households used more low-VOC or certified alternatives.

This result underscores a long-standing environmental inequality: populations with fewer material resources bear a disproportionate burden of pollutant exposure and its associated health consequences.

Smoking behaviour, diet, physical activity, and underlying health conditions also contributed substantially to variation in biomarker concentrations. Participants who lived with at least one smoker exhibited C-reactive protein values that were 25–40 per cent higher than those in non-smoking homes, while participants with sedentary lifestyles showed oxidative stress markers (e.g., F2-isoprostanes) elevated by an average of 18–25 per cent.

Adjusting for these factors strengthened the statistical validity of the exposure–life expectancy relationship by demonstrating that the observed associations were not confounded by lifestyle differences. The analysis confirmed that indoor pollutant exposure retained a large and independent association with biological risk, with exposure–risk coefficients remaining statistically significant at $p < 0.001$ after full covariate adjustment.

Building characteristics proved equally influential. Older buildings, particularly those exceeding 25 years in age, had air infiltration rates up to 3.5 ACH50 (air changes per hour at 50 pascals of pressure, meaning how many times indoor air is replaced when the home is pressurised during a standard leakage test), resulting in 25–45 per cent higher indoor $PM_{2.5}$ and 15–30 per cent higher indoor O_3 compared with newer buildings.

In contrast, newer buildings, especially those constructed within the last 10 years, exhibited formaldehyde concentrations exceeding 40 μg/m³ in 22 per cent of homes, reflecting emissions from composite wood, adhesives, and laminates. These findings illustrate a critical tension between energy efficiency, airtightness, and chemical emissions from modern materials.

Geospatial variables further shaped exposure patterns. Homes located within 150 metres of major roadways showed baseline indoor $PM_{2.5}$ concentrations 12–18 μg/m³ higher than those in low-traffic areas, and indoor NO_2 levels elevated by 8–15 ppb, confirming the strong influence of urban traffic emissions. Neighbourhood greenery coverage reduced infiltration by 8–12 per cent, while higher ambient temperatures increased off-gassing rates, raising VOC concentrations by 10–20 per cent on warm days.

Collectively, the covariate results show that indoor pollutant exposure must be understood within a broader contextual framework that incorporates socio-economic, behavioural, architectural, and environmental determinants. These findings emphasise that improving indoor air quality requires not only technological intervention but also structural and policy-level solutions that address inequalities in exposure and vulnerability.

Mechanistic Foundations for the Exposure–Risk–Longevity Pathway

The mechanistic modelling provided strong evidence supporting the causal chain linking pollutant exposure to biological risk and, ultimately, to reductions in life expectancy. The cumulative exposure mixture index demonstrated the strongest association with biological risk, outperforming individual pollutant metrics by a substantial margin.

Across more than 2.1 million valid hourly exposure–biomarker pairs, the mixture index explained 42–56 per cent of the variance in biological risk, compared with 18–27 per cent for single-pollutant models. This reinforced the understanding that exposure to pollutant mixtures, rather than isolated contaminants, drives long-term physiological decline.

This finding challenges conventional regulatory approaches that rely on pollutant-by-pollutant thresholds, suggesting that mixture-based assessments may provide a more accurate representation of real-world health risks. Mixture effects were especially pronounced for households experiencing concurrent elevations in $PM_{2.5}$, NO_2, and VOCs, where biological risk scores increased by an average of 0.31–0.46 units for each interquartile increase in mixture burden.

The calibrated mechanistic model showed that both cumulative exposure and biological vulnerability significantly influenced the Actual Risk Score. Biological vulnerability alone accounted for 22–29 per cent of variance in risk, exposure burden accounted for 31–38 per cent, and the interaction between the two accounted for an additional 15–21 per cent—the strongest single contributor. This aligns with toxicological models demonstrating that individuals with reduced antioxidant capacity or heightened inflammatory responsiveness experience accelerated damage under pollutant stress.

The agreement between the Predicted and Actual Risk Scores further validated the model. The concordance correlation coefficient (CCC = 0.87) and cross-validated R^2 (0.81) demonstrated excellent predictive fidelity across more than 800 participants, confirming that the model successfully captured essential biological processes involved in pollutant-induced ageing.

Translating biological risk into expected life expectancy loss illuminated the long-term health consequences of chronic exposure. Each 0.1-unit increase in biological risk corresponded to an estimated 0.42–0.58 years of life expectancy reduction, with the highest-exposure quartile losing an average of 2.8–3.6 years compared with the lowest quartile. These reductions remained significant even after adjustment for all covariates, demonstrating that the physiological burden imposed by indoor pollutants is a major determinant of survival outcomes.

The mechanistic model therefore provided more than numerical estimates; it offered a conceptual framework unifying biological, behavioural, environmental, and demographic factors. By integrating more than forty covariates into its risk structure and capturing over 90 per cent completeness in exposure–biology temporal alignment, the model established a scientifically rigorous foundation for understanding how indoor exposures shape long-term health. This foundation supports the development of predictive modelling, policy frameworks, and targeted IAQ interventions.

Outcome Measurement

The integration of validated health-risk scores with national mortality datasets produced clear evidence that indoor pollutant exposure reduces life expectancy in meaningful and quantifiable ways. More than 1.2 million mortality records spanning fifteen years were accessed through a secured national health database under strict data-sharing agreements.

These records contained age-, sex-, and year-specific mortality rates for the entire population, but no identifiable information about individuals. Importantly, the mortality data were never linked to any participant. Instead, the data were used in aggregate form to provide the baseline probability of death for each age and sex group in the country. This structure allowed scientifically robust life-table calculations without compromising privacy.

To make the results interpretable, the health-risk scores derived from each participant's exposure dose, biological vulnerability, and covariate profile were not treated as predictors of the participant's personal age of death. Instead, the scores were used to adjust the *baseline national mortality probabilities* for people of the same age and sex. This means the study asked a population-level question: "If a typical person in the country had the same pollutant-related health burden as this participant, how would their expected years of remaining life change?"

This approach is best understood through a clear example. Suppose a forty-five-year-old participant had a health-risk score that was 31 per cent higher than the national average for forty-five-year-olds. The model did *not* attempt to predict when this participant would die. Instead, it increased the national mortality probability for all forty-five-year-olds by that same 31 per cent.

This means that the ordinary mortality probability for a forty-five-year-old in the national data was multiplied by 1.31 to reflect the higher health burden measured in the study participant. This same adjustment was then carried forward year by year: the mortality probability for age 46 was also increased by 31 per cent, and similarly for ages 47, 48, 49, and so on. In other words, the participant's health-risk score served as a multiplier that modified the standard survival pattern for a typical person of the same age and sex.

This produced a complete set of adjusted mortality probabilities, which together formed a modified survival curve. These adjusted probabilities were then entered into the national life-table, a structured statistical tool that uses age-specific mortality rates to calculate how many years of life the average person can expect to live.

The life-table therefore produced an estimate of expected remaining life under the increased health burden implied by indoor pollution exposure. Critically, life expectancy was estimated for a “typical person” in the demographic group—not for the actual participant—ensuring privacy and interpretability.

In practical terms, the method worked as follows. Each participant’s exposure mixture, biological vulnerability, and covariates were combined into one health-risk score. This score represented how much physiological strain pollution was imposing on their body. The model then determined how much higher this risk was compared with the national average for their age and sex group (e.g., 12 per cent higher, 25 per cent higher, or 31 per cent higher).

That percentage increase was applied directly to the national mortality probabilities for people of that same demographic group. These adjusted probabilities were then used to compute a modified life expectancy that answered a clear and meaningful question: “How long would the average person live if they carried this level of pollutant-related health burden throughout adulthood?”

This framework produced life expectancy estimates that were grounded in real exposure biology and real population mortality patterns. It converted personal pollutant profiles into population-level implications without predicting anyone’s personal date of death. It also allowed valid estimation of life-years lost even though the study was conducted over only twelve to eighteen months, because the long-term mortality patterns were taken directly from national records.

The results showed substantial effects. Individuals in the highest exposure–vulnerability quartile experienced predicted reductions in life expectancy ranging from 2.8 to 5.4 years (95 per cent CI: 2.4–5.9), calculated from survival curves adjusted by their health-risk multipliers. Even those in the moderate exposure range lost between 1.3 and 2.2 years of expected life. Subgroup analyses revealed that those with elevated biological vulnerability suffered the largest losses.

Participants in the highest vulnerability decile lost an additional 0.32–0.47 years of life for each 0.1-unit increase in health-risk score, compared with only 0.12–0.19 years in the lowest decile. Homes with consistently high formaldehyde levels displayed similarly pronounced effects. Dwellings with mean formaldehyde concentrations above 45 μg/m³ showed 1.1–1.8 additional years of life lost relative to homes below 25 μg/m³—an effect arising from the strong contribution of formaldehyde to oxidative and inflammatory burden.

Taken together, these findings placed indoor air pollution alongside smoking, uncontrolled hypertension, and chronic psychosocial stress as major modifiable determinants of national longevity. The magnitude of life-years lost was comparable to those associated with long-term

second-hand smoke exposure and exceeded those typically associated with low levels of physical inactivity. This demonstrated that indoor air pollutant exposure is not a minor contributor to health decline but a significant determinant of survival outcomes.

Finally, the analytic framework demonstrated strong validity. The adjusted survival curves generated by the model showed a concordance correlation coefficient of 0.91 when compared with observed national survival patterns. This high degree of agreement confirmed that the exposure–risk–longevity structure was scientifically credible, internally coherent, and demographically grounded, providing strong evidence that the reduction in life expectancy estimated here reflected real and meaningful population-level effects of indoor pollution exposure.

Synthesis and Broader Significance

The findings for Research Question 1 provided clear and convergent evidence that long-term exposure to indoor air pollutants measurably reduces life expectancy in the study country. Chronic exposure to $PM_{2.5}$, NO_2, O_3, VOCs, and formaldehyde emerged as a continuous physiological burden shaped by pollutant generation indoors, infiltration from outdoors, and the modifying effects of ventilation behaviour and building characteristics.

These exposure conditions produced consistent biological signatures, including elevated oxidative stress, systemic inflammation, lipid peroxidation, mitochondrial dysfunction, and shortened telomere length. Each of these responses reflects an established mechanistic pathway through which pollution accelerates biological ageing and increases long-term mortality risk.

A key insight was the central role of physiological susceptibility. Individuals with weaker antioxidant defences or reduced detoxification capacity showed substantially greater biological damage at similar exposure levels, leading to larger downward shifts in their predicted survival curves. This demonstrated that the relationship between exposure and mortality burden is not uniform but mediated by biological vulnerability.

The integration of individual health-risk scores with national mortality structures further strengthened the evidence for causality. By adjusting age- and sex-specific national mortality probabilities according to each participant’s pollutant-related health burden, the study demonstrated that increases in health risk translate into meaningful population-level reductions in expected lifespan. This analytical coherence confirms that chronic indoor pollutant exposure is a significant determinant of premature mortality.

Taken together, these findings expand scientific understanding by showing that pollutant mixtures, physiological susceptibility, and covariate influences form an interconnected exposure–risk–longevity system. They also underscore the broader public health significance of indoor air quality as a major and modifiable driver of national life expectancy.

Findings for Research Question 2:

Introduction

The findings for Research Question 2 showed that the introduction of an artificial intelligence layer fundamentally advanced the modelling capability established in Research Question 1. Although RQ1 already incorporated exposure dose, biological vulnerability, and covariates within a mechanistic–epidemiological framework, its structure remained constrained by fixed equations and retrospective estimation.

In contrast, RQ2 retained these scientific foundations but extended them through hybrid integration with AI, enabling the model to learn patterns and dependencies that mechanistic formulations alone could not express.

The hybrid system captured how pollutant behaviour, physiological susceptibility, and contextual conditions interact over time, rather than assuming a uniform or linear relationship. By embedding AI within a mechanistic scaffold, the model represented exposure as an evolving sequence shaped by emission dynamics, infiltration events, ventilation decisions, and occupant activity cycles. This allowed life expectancy projections to respond sensitively and realistically to short-term fluctuations and long-term trends within indoor environments.

Crucially, the hybrid model did not replace the biological logic established in RQ1 but amplified its expression. Pathways involving oxidative stress, chronic inflammation, mitochondrial dysfunction, and telomere attrition were preserved mechanistically while being modelled with greater flexibility through AI-driven learning. This produced predictions that were both biologically credible and mathematically responsive to complex real-world variations.

Collectively, these findings demonstrate that augmenting the mechanistic exposure–biology–covariate framework with AI substantially elevates its predictive capacity. The results therefore support the alternative hypothesis and establish hybrid mechanistic–AI modelling as a more powerful and scientifically coherent approach for forecasting life expectancy under chronic indoor air pollutant exposure.

Predictive Accuracy and Scientific Value of the Hybrid Mechanistic–AI Model

The hybrid mechanistic–AI model achieved a level of predictive accuracy that exceeded that of the mechanistic-only framework used in Research Question 1, despite both models being grounded in the same scientific inputs of exposure dose, biological vulnerability, and covariates. The key distinction lies in the introduction of an artificial intelligence layer in RQ2, which enabled the model to learn nonlinear interactions, temporal dependencies, and contextual patterns that classical mechanistic structures could not fully represent.

Across the full study population, the hybrid model achieved an average coefficient of determination of 0.86, indicating that eighty-six per cent of the variation in life expectancy attributable to chronic indoor air pollutant exposure, biological vulnerability, and covariates was captured with high fidelity. This represents a substantial improvement over the mechanistic-only model, whose performance plateaued at approximately 0.72, reflecting the limitations of models that rely solely on fixed functional forms and linear biological assumptions.

The enhanced accuracy emerged from the hybrid model's ability to represent exposure as a dynamic and context-sensitive process. Whereas the RQ1 mechanistic framework reconstructed exposure and biological burden retrospectively, the hybrid model incorporated time-resolved pollutant trajectories and learnt complex interdependencies among pollutant mixtures, occupant behaviours, building characteristics, and physiological responses. This AI-enabled capacity to detect and model nonlinearities strengthened predictive performance in ways unattainable through mechanistic modelling alone.

Calibration analysis reinforced these improvements. The hybrid model achieved a mean deviation of 0.9 years from observed life expectancy, compared with deviations between 2.0 and 2.8 years for the mechanistic-only and conventional statistical models. These differences demonstrate that the hybrid system not only explained more variance but also generated predictions that closely matched real survival patterns.

The hybrid model predicted both low and high life expectancy values with balanced reliability, avoiding the systematic under- and overestimation that characterised non-AI models, particularly among individuals with extreme biological vulnerability profiles.

External validation further confirmed the robustness of the hybrid approach. When applied to independent datasets involving different building archetypes, climatic conditions, and socio-economic environments, the hybrid model maintained strong performance, achieving an R^2 of approximately 0.82. By contrast, performance of the mechanistic-only and statistical models degraded markedly under the same conditions, illustrating the hybrid system's superior capacity to generalise across heterogeneous indoor environments.

Taken together, these findings demonstrate that the introduction of an AI layer into a rigorous mechanistic exposure–biology–covariate framework provides a scientifically significant enhancement in predictive accuracy, stability, and interpretability. The results therefore reject the null hypothesis and support the alternative hypothesis, confirming that hybrid mechanistic–AI integration offers a more coherent and powerful predictive structure for estimating variations in life expectancy arising from chronic indoor air pollutant exposure.

Role of Biological Vulnerability within the Hybrid Modelling Framework

In findings for Research Question 2, biological vulnerability was not re-established as a new or independent determinant of life expectancy, because its importance had already been demonstrated in findings for Research Question 1. Instead, the purpose of RQ2 was to examine how the integration of an artificial intelligence layer enhanced the modelling of biological vulnerability within a predictive system that already included exposure dose, physiological susceptibility, and covariates.

This distinction is critical. RQ1 showed that biological vulnerability plays a fundamental role in shaping life expectancy under chronic indoor air pollutant exposure. RQ2 sought to determine whether hybrid mechanistic–AI integration could utilise this biological vulnerability information more effectively to improve predictive accuracy, responsiveness, and interpretability.

The hybrid model strengthened the expression of biological vulnerability by modelling its interaction with exposure and covariates in a dynamic, nonlinear, and temporally sensitive manner. In RQ1, biological vulnerability was incorporated through fixed mechanistic dose–response functions that characterised how pollutants contribute to oxidative stress, inflammation, mitochondrial strain, and telomere shortening.

These relationships were valid and scientifically grounded but followed predetermined functional forms. The hybrid approach in RQ2 allowed the model to learn additional patterns within these pathways, including time-dependent effects and context-specific manifestations of biological vulnerability that could not be captured by classical mechanistic structures alone.

This improvement emerged most clearly in the way the hybrid model handled fluctuating exposure conditions. Indoor air pollutant concentrations often vary sharply within short intervals because of activities such as cooking, chemical product use, or changes in ventilation. The mechanistic model in RQ1 incorporated these fluctuations but applied biological vulnerability uniformly across the full exposure trajectory.

The hybrid model, by contrast, learnt how biological vulnerability modifies the effects of short-term pollutant spikes, prolonged low-level exposure, and mixed pollutant conditions. It captured, for example, how individuals with heightened inflammatory sensitivity experience disproportionate biological strain from episodic increases in particulate matter, or how those with strong antioxidant capacity exhibit moderated responses under similar conditions. This did not redefine biological vulnerability but made its influence more accurately represented within the predictive model.

The hybrid system also improved the modelling of interactions between biological vulnerability and covariates. Age and existing health conditions directly shape biological vulnerability, while factors such as socio-economic position, building archetype, and time spent indoors do not alter physiological susceptibility itself but influence the environmental conditions under which that vulnerability is expressed.

The AI layer enabled the model to identify complex interactions, such as the intensified effects of pollutant exposure in vulnerable individuals living in poorly ventilated buildings or the protective influence of certain household behaviours in resilient individuals. These nuanced patterns could not emerge from mechanistic equations alone, even though the underlying variables were identical.

Collectively, the findings show that RQ2 does not restate the fundamental importance of biological vulnerability, because that foundation was firmly established in RQ1. Instead, RQ2 demonstrates that the hybrid mechanistic and AI framework enables biological vulnerability to be represented in a more flexible, realistic, and context-responsive manner that adjusts continuously to changing exposure conditions, behavioural patterns, and building characteristics.

This enhanced utilisation of vulnerability information directly contributed to the observed improvements in predictive accuracy and model coherence, reaffirming that AI integration adds value by refining, rather than redefining, the role of physiological susceptibility in forecasting life

expectancy under chronic indoor air pollutant exposure.

Influence of Covariates and Building Archetypes on Model Predictions

In findings for Research Question 2, the treatment of covariates and building archetypes changed in a fundamental way compared with their role in findings for Research Question 1. In RQ1, covariates such as age, chronic illness, socio-economic position, household structure, and time spent indoors were shown to be scientifically important because they shaped both exposure dose and biological vulnerability.

However, their effects were still represented through fixed mechanistic relationships. Each covariate influenced outcomes through predetermined equations that linked exposure to physiological decline, with no capacity to adapt when real-world conditions varied.

The hybrid model in RQ2, by contrast, incorporated these same covariates into an adaptive, learning-based structure that allowed their effects to emerge through complex interactions with pollutant dynamics, biological susceptibility, and temporal behavioural patterns. This difference in modelling strategy produced more realistic and context-sensitive life expectancy predictions.

The hybrid system treated covariates not as static adjustments but as active determinants embedded within a predictive architecture capable of identifying nonlinear and time-dependent patterns. For example, socio-economic position does not have a single uniform effect; instead, its influence varies depending on ventilation practices, fuel use, floor level, and access to exposure-reducing technologies.

The AI layer was able to capture these contextual expressions of covariates by learning from combinations of behavioural markers, building conditions, and pollutant trajectories. This enhanced utilisation of covariates was reflected in predictive performance: removing covariates from the hybrid model reduced the coefficient of determination from 0.86 to approximately 0.79, a clear indication that the hybrid model depended on their interactional structure rather than treating them as peripheral controls.

Building archetypes played an especially consequential role in RQ2 because the hybrid model was capable of interpreting their structural features in relation to pollutant behaviour over time. In the mechanistic-only framework of RQ1, building characteristics influenced infiltration, deposition, and ventilation rates through fixed parameters.

In RQ2, the hybrid model could learn how building attributes interacted with behaviour and vulnerability to shape exposure trajectories in more complex ways. Features such as façade permeability, cross-ventilation potential, elevation, envelope ageing, and ventilation design contributed to patterns of short-term pollutant spikes and long-term accumulation that mechanistic equations alone could not fully capture.

As a result, buildings classified as having high infiltration and low air movement generated larger predicted reductions in life expectancy, with differences of up to three to four years between the most and least favourable archetypes.

The hybrid model also identified synergistic effects that were difficult to detect in RQ1. Individuals with high biological vulnerability living in structurally disadvantaged buildings exhibited the steepest declines in predicted life expectancy. These interactions were not imposed by the model but emerged naturally from the hybrid learning process, demonstrating the system's ability to represent how biological susceptibility, socio-economic constraints, behaviour, and architecture converge to shape long-term health outcomes.

Overall, the findings confirm that covariates and building archetypes assume a substantially expanded role in hybrid mechanistic–AI modelling. Rather than functioning as control factors, they become integral components that dynamically structure exposure conditions, modify vulnerability expression, and shape pollutant behaviour.

RQ2 therefore establishes that the predictive gains of the hybrid model arise not from new variables, but from the model's superior ability to learn, interpret, and operationalise the combined influence of covariates and building structures on life expectancy outcomes under chronic indoor air pollutant exposure.

Integration of Mechanistic and AI Layers in Predictive Modelling

The findings showed that the improved predictive performance observed under Research Question 2 resulted directly from the integration of mechanistic pollutant modelling with an artificial intelligence layer. The mechanistic component produced structured, physically meaningful features that reflected real indoor pollutant behaviour.

These included reconstructed concentration trajectories, biologically weighted inhalation dose estimates, deposition and resuspension indicators, and cumulative biological burden metrics derived from calibrated dose–response functions. Together, these mechanistic features formed a coherent scientific scaffold representing exposure conditions across diverse households and building archetypes.

The hybrid model extended this scaffold by incorporating an AI layer capable of learning complex interactions that the mechanistic structure alone could not represent. The AI component captured nonlinear relationships among pollutant mixtures, temporal dependencies between short-term exposure peaks and biological burden, and characteristic behavioural patterns influencing pollutant accumulation.

These learnt relationships remained aligned with measured pollutant events, biomarker-derived vulnerability indices, and observed building-level patterns, demonstrating that the hybrid model used the mechanistic inputs as meaningful constraints rather than replacing them.

The direct comparison between the mechanistic-only and hybrid configurations confirmed the functional contribution of integration. The mechanistic-only model achieved a coefficient of determination of approximately 0.71, reflecting its inability to represent nonlinear interactions, cross-pollutant effects, and time-dependent variations in exposure–response dynamics.

When the AI layer was incorporated, predictive accuracy increased to 0.86. This improvement indicates that the hybrid model was able to utilise mechanistic features while learning additional relationships that were not explicitly encoded in the mechanistic equations.

Interpretability findings further supported this conclusion. Features derived from mechanistic modelling—such as reconstructed $PM_{2.5}$ peaks, pollutant decay parameters, and biologically adjusted dose estimates—featured prominently among the most influential predictors, alongside hybrid-level variables that reflected behavioural–exposure timing, occupant-driven emission cycles, and building-specific retention patterns. This alignment demonstrated that the hybrid system used both mechanistic structure and learned relationships jointly rather than privileging one over the other.

External validation conducted on independent household and building datasets within the study country reinforced the stability of the hybrid architecture. The hybrid model consistently reproduced observed pollutant trajectories and corresponding life expectancy outcomes when applied to new combinations of building design, occupant schedules, and exposure conditions not represented in the training data.

The mechanistic-only model showed noticeable reduction in predictive alignment under these conditions, highlighting its limited capacity to generalise across heterogeneous exposure environments.

Overall, the results show that the predictive advancements achieved under Research Question 2 were the direct consequence of combining mechanistic equations with artificial intelligence learning. The mechanistic layer ensured physical and biological coherence, while the AI layer enabled the system to capture high-dimensional, context-specific patterns that mechanistic formulations alone could not accommodate.

This integrated architecture produced the only modelling configuration capable of accurately forecasting variations in life expectancy under chronic indoor air pollutant exposure.

Identification of High-Risk Subpopulations and Building Profiles

A central scientific contribution of Research Question 2 was the hybrid model's markedly enhanced ability to identify high-risk subpopulations and high-risk building profiles with a level of precision that exceeded what the mechanistic-only model achieved in Research Question 1.

Both modelling frameworks incorporated exposure dose, biological vulnerability, and covariates; however, only the hybrid system could disentangle the complex, dynamic, and nonlinear interactions that determine which groups face the greatest life expectancy loss from chronic indoor air pollutant exposure. The identification of high-risk populations and building archetypes therefore became a defining outcome of the hybrid approach.

The hybrid model produced a clearly stratified risk structure across the national cohort. Individuals in the lowest-risk group exhibited predicted life expectancy reductions of 0.9 to 1.6 years, while those in the highest-risk stratum showed markedly larger reductions of 7.4 to 8.1 years.

By contrast, the mechanistic-only model identified the same directional pattern but produced a narrower spread, with the highest-risk group predicted to lose 4.8 to 5.2 years. This 2.5 to 3.0-year increase in stratification under the hybrid model directly strengthened its ability to identify those facing disproportionately higher health risks. The broader spread reflects the hybrid system's capacity to model recurrent pollutant peaks, nonlinear dose–response interactions, and temporal synchronisation between exposure and behaviour, all of which are essential for detecting high-risk clusters.

The hybrid system also quantified the defining characteristics of top-risk individuals more accurately. Among people with biological vulnerability in the upper quartile and cumulative exposure levels 15 to 20 per cent above the median, the hybrid model explained 86 per cent of the variance in life expectancy reductions, compared with 68 per cent under the mechanistic-only configuration.

This improvement allowed the hybrid model to pinpoint high-risk subpopulations not only by vulnerability or exposure alone but by their combined temporal and contextual risk signatures, which the mechanistic-only model represented less distinctly.

Building archetypes formed the second major domain of high-risk identification. The hybrid model consistently associated the greatest predicted life expectancy reductions with older buildings (≥35 years), units with infiltration rates at least 25 per cent above the study median, and dwellings experiencing frequent pollutant accumulation.

In these high-risk building profiles, the hybrid system predicted life expectancy reductions of 4.4 to 5.1 years, compared with 3.0 to 3.5 years from the mechanistic-only model. This additional 1.4 to 1.8 years of predicted loss reflects the hybrid model's ability to recognise how structural deficiencies and behavioural routines interact dynamically to intensify chronic exposure, something the mechanistic-only model tended to understate because of its more rigid temporal assumptions.

The hybrid model also isolated risk clusters that were not fully resolved using mechanistic-only modelling. For example, a subgroup with moderate long-term exposure but very high biological vulnerability, living in older flats with recurrent evening pollutant peaks, was identified as high-risk with reductions of 5.8 to 6.3 years.

The mechanistic-only approach estimated only 3.9 to 4.4 years. This nearly two-year difference demonstrates the hybrid system's unique ability to capture risk in populations whose vulnerability arises from complex temporal and environmental convergence rather than from exposure magnitude alone.

Overall, the hybrid mechanistic–AI model advanced the scientific objective of Research Question 2 by enabling the precise identification of high-risk subpopulations and high-risk building profiles. Its enhanced capacity to resolve dynamic, recurrent, and nonlinear exposure–biology–covariate interactions provided a more detailed and actionable understanding of where the greatest life expectancy losses occur, firmly establishing the added value of the hybrid approach.

Synthesis and Broader Significance

Research Question 2 examined whether introducing an artificial intelligence layer into an already rigorous mechanistic framework built on exposure dose, biological vulnerability indices, and covariates could improve the prediction of life expectancy variations under chronic indoor air pollutant exposure. Both the mechanistic-only and hybrid models used the same underlying information on pollutant exposure, susceptibility, and contextual factors; the sole difference lay in the modelling architecture.

The null hypothesis (H_{02}) stated that adding AI to this mechanistic structure would not significantly enhance predictive accuracy or risk characterisation, while the alternative hypothesis (H_{12}) proposed that hybrid integration would yield a measurable improvement.

The findings clearly reject H_{02} and support H_{12}.

The hybrid mechanistic–AI model achieved substantially higher predictive accuracy than the mechanistic-only model ($R^2 \approx 0.86$ versus ≈ 0.71), demonstrating that the AI layer added genuine explanatory power rather than duplicating existing structure.

By learning nonlinear and time-dependent relationships within the same exposure, vulnerability, and covariate space, the hybrid framework captured complex interactions that fixed mechanistic equations could only approximate. This led to more precise estimates of life expectancy loss, particularly in subpopulations with fluctuating exposure patterns or compounded vulnerabilities.

Crucially, the hybrid model improved not only overall fit but also the resolution with which high-risk groups and building archetypes were identified. Using the same exposure dose, biological indices, and contextual covariates, the hybrid system produced a wider and more discriminating spread of predicted life expectancy reductions, revealing high-risk population–building combinations that the mechanistic-only model tended to compress toward the mean. This demonstrates that the added value of AI lies in its ability to accommodate dynamism and complexity within an unchanged scientific input structure.

Taken together, the results show that when exposure dose, biological vulnerability, and covariates are already well characterised, embedding them in a hybrid mechanistic–AI architecture yields a more accurate, discriminating, and operationally useful framework for forecasting life expectancy impacts of chronic indoor air pollutant exposure.

Findings for Research Question 3:

Introduction

Execution of the methodology for Research Question 3 produced a comprehensive and scientifically coherent set of findings demonstrating how architectural, engineering, behavioural, and contextual interventions reshape indoor exposure conditions and modify long-term life expectancy outcomes.

These findings represent the translational stage of the research, extending the epidemiological and mechanistic foundations established in Research Question 1 and the dynamic hybrid mechanistic–AI framework developed in Research Question 2.

Research Question 1 established how chronic exposure to $PM_{2.5}$, NO_2, O_3, VOCs, and formaldehyde accelerates biological ageing through oxidative stress, inflammation, mitochondrial impairment, and telomere attrition. Research Question 2 then developed a hybrid mechanistic–AI model capable of forecasting life-expectancy trajectories under evolving real-world exposure conditions. Research Question 3 applied these scientific advances to evaluate which intervention strategies could meaningfully alter long-term outcomes by modifying the exposure–biology–longevity pathway.

To present the findings in a scientifically logical sequence that reflects how the results were generated, the results first report the simulation modelling of intervention effects. This section demonstrates how coupled CONTAM–EnergyPlus modelling reconstructed the pollutant environment of each scenario, produced time-resolved concentration profiles, and generated exposure doses that served as the foundation for interpreting intervention performance. It also shows how these outputs were translated into biological risk and life-expectancy trajectories using the hybrid model's dynamic risk and survival engines.

The next section, development of intervention scenarios, presents what these simulations revealed about the real-world effects of architectural, engineering, behavioural, and contextual strategies. It reports how changes to ventilation, infiltration, emission patterns, and occupant practices reshaped indoor pollutant dynamics, reduced biological stress markers, and produced measurable improvements in projected life expectancy.

The third section, multi-criteria optimisation, identifies the intervention packages that delivered the greatest health benefits per unit of resource invested, integrating metrics of life-expectancy gain, cost, feasibility, equity, and contextual constraints. Finally, stakeholder engagement and policy translation demonstrate how scientifically optimal interventions were evaluated for practical adoption and incorporated into a decision-support tool to guide value-oriented and equitable implementation.

This ordering reflects the conceptual logic of Research Question 3. The simulations establish how pollutants behave under different conditions; the intervention scenarios show how changes translate into exposure and biological outcomes; the hybrid model determines how these shifts affect longevity; optimisation identifies the combinations that provide the greatest societal value; and stakeholder engagement clarifies how they can be applied in practice.

Together, these findings demonstrate that Research Question 3 completes the scientific arc of the doctoral study: RQ1 identified the biological pathways of harm, RQ2 developed a predictive platform to model these pathways dynamically, and RQ3 applied this platform to evaluate, quantify, and optimise interventions capable of improving indoor air quality and extending life expectancy across diverse real-world contexts.

Simulation Modelling of Intervention Effects

Simulation modelling generated the pollutant trajectories that formed the analytical foundation for interpreting how each intervention altered exposure, biological burden, and projected life expectancy. Using coupled CONTAM–EnergyPlus simulations, the study produced ten-minute concentration profiles for $PM_{2.5}$, NO_2, ozone, VOCs, and formaldehyde under baseline and modified indoor conditions.

These profiles reflected the physics of airflow, infiltration, temperature-driven buoyancy, and pollutant transport across rooms and floors, enabling each intervention scenario to be reconstructed with high fidelity.

The simulations revealed that indoor pollutant behaviour was highly sensitive to modest changes in airflow and emission strength. Increasing effective ventilation area by as little as twenty per cent produced markedly steeper pollutant decay curves following cooking or cleaning events, demonstrating how enhanced air exchange accelerated pollutant clearance. Conversely, reducing façade infiltration flattened pollutant ingress curves during periods of high outdoor NO_2, confirming that envelope improvements directly dampened the magnitude and frequency of infiltration-driven spikes.

EnergyPlus outputs showed that these effects varied significantly with time of day, humidity, and thermal gradients. This temporal variability demonstrated that the influence of interventions could not be captured accurately using long-term averages or steady-state assumptions.

A central insight from the simulations was that indoor concentration alone was a poor predictor of actual inhaled dose. When the concentration profiles were combined with participants' time–activity patterns and activity-specific inhalation rates, the data showed that large proportions of daily exposure occurred during short, intense emission periods lasting five to fifteen minutes. These peaks carried disproportionately large biological significance because inhalation during cooking, cleaning, or smoking events generated high doses in a short time.

As a result, interventions that reduced peak concentrations produced much larger biological benefits. Sensor-augmented portable air cleaners are a key example. During high-emission events, they temporarily increased the clean air delivery rate (CADR), allowing the system to remove pollutants quickly at the moment they mattered most. Interventions that affected only long-term averages did not achieve comparable health effects. This difference would not have been visible without explicit exposure modelling.

The simulations also showed that the timing of behavioural actions strongly influenced their effectiveness. Window-opening during periods of low outdoor $PM_{2.5}$ produced twice the reduction in inhaled dose compared with equivalent ventilation during polluted periods, despite identical airflow rates. This revealed that behavioural interventions are effective only under suitable environmental conditions and underscored why modelling must incorporate time-varying outdoor air quality.

When inhaled doses were translated into physiological burden using the dose–response relationships established in Research Question 1, the simulations showed that reductions in exposure produced nonlinear reductions in biological stress. For individuals with heightened biological vulnerability, even modest reductions in inhaled $PM_{2.5}$ or NO_2 caused

disproportionately large decreases in oxidative load and inflammatory signalling. This explained why the same intervention produced life-expectancy gains of less than one year for some occupants but more than four years for others.

Life-table translation further showed that interventions that lowered short-duration peaks had the strongest impact on survival probability. Although peaks were brief, they contributed heavily to cumulative oxidative and inflammatory burden. Mechanistic modelling alone, which smooths concentrations into daily or weekly averages, would not have revealed this disproportionate importance of peak suppression.

The hybrid model introduced a second tier of findings by detecting patterns invisible to mechanistic simulation. When dynamic pollutant and exposure data were fed into the hybrid engine, the AI layer identified seasonal and behavioural rhythms, for example, increased VOC retention during humid periods, elevated cooking emissions during holidays, or reduced window-opening during heavy rain.

These temporal patterns modified exposure trajectories and altered life-expectancy predictions. The hybrid system also detected degradation in intervention performance: particulate accumulation gradually reduced clean air delivery rate (CADR), in portable air-cleaning units, sensor drift altered responsiveness, and ventilation-related behaviours decayed over weeks without reinforcement. These real-world changes modified long-term exposure in ways that static models could not represent.

By continuously updating exposure–risk trajectories as pollutant conditions, behaviour, and device performance evolved, the hybrid model produced survival curves that reflected lived indoor environments rather than idealised assumptions. Together, these results showed that simulation modelling when interpreted through hybrid mechanistic–AI translation provided the mechanistic, exposure-based, biological, and temporal evidence required to evaluate interventions scientifically and realistically.

These simulation-derived insights set the analytical foundation for the next section, which reports how architectural, engineering, behavioural, and contextual interventions reshaped pollutant trajectories, altered biological stress indicators, and ultimately modified long-term life-expectancy outcomes.

Development of Intervention Scenarios

The development and assessment of intervention scenarios revealed clear and scientifically grounded evidence of how architectural, engineering, behavioural, and contextual modifications alter indoor air pollutant dynamics and life-expectancy trajectories.

Across all building types and population groups examined, the results consistently showed that changes to air exchange, infiltration, emission patterns, and ventilation timing produced measurable shifts in pollutant concentrations, biological stress indicators, and predicted longevity.

The intervention scenarios demonstrated that indoor air quality is not governed by any single factor but by the interaction of structural design, mechanical systems, occupant behaviour, and environmental context. This systems-based picture emerged strongly in the results and underscored the need for integrated intervention strategies rather than isolated measures.

A central finding was that the hybrid mechanistic–AI model from Research Question 2 fundamentally reshaped how intervention effects manifested. The results showed that intervention impacts followed dynamic patterns rather than fixed or average changes.

Ventilation gains fluctuated throughout the day, infiltration varied with outdoor conditions, and the health benefits of filtration, ventilation, and behavioural changes evolved over time. These real-world patterns produced exposure trajectories far more complex than static estimates would suggest.

As a result, the hybrid model revealed life-expectancy outcomes that were more differentiated, more sensitive to behavioural–environmental variation, and more biologically coherent than those generated by mechanistic modelling alone. The findings confirmed that the hybrid model captured temporal dependencies and nonlinear interactions that played a decisive role in determining which interventions meaningfully reduced long-term health risk.

Architectural interventions produced clear pollutant reductions across particulate and gaseous pollutants. Increasing the effective window-opening area by roughly one-third reduced long-term indoor $PM_{2.5}$ concentrations by eighteen to twenty-two per cent, with the largest improvements in homes lacking cross-ventilation. The findings showed that improved air pathways enabled polluted indoor air to be replaced more quickly during favourable outdoor conditions.

In high-outdoor-pollution regions, reducing façade infiltration by forty per cent lowered indoor NO_2 by thirteen to seventeen per cent, highlighting the importance of building-envelope integrity. Substituting traditional materials with low-emission alternatives reduced indoor formaldehyde concentrations by twenty-eight to thirty-four per cent during the first year.

These pollutant reductions translated into life-expectancy gains ranging from 0.9 to 2.1 years. The analysis showed that individuals with high biological vulnerability gained the most because decreases in pollutant load disproportionately reduced inflammatory and oxidative stress responses in these groups.

The results also demonstrated that architectural improvements produced conditional rather than universal benefits. For instance, the gains from increased window-opening area occurred only when residents engaged in frequent ventilation and when outdoor pollution levels were sufficiently low.

In households that rarely opened windows or lived in areas with highly variable outdoor air quality, the benefits of architectural enhancements were substantially smaller. Patterns of window operation and daily activities shaped whether improved building design translated into meaningful exposure reductions. The hybrid model captured these conditional patterns and revealed that architectural interventions alone were insufficient in many real-world contexts.

Engineering interventions generated the largest pollutant reductions and life-expectancy improvements. High-efficiency filters reduced indoor $PM_{2.5}$ concentrations by forty-one to fifty-two per cent, producing pronounced declines in long-term inhalation dose. While these filters operate at a fixed efficiency and provide continuous baseline protection, sensor-augmented portable air-cleaning units introduced an adaptive layer of performance that further improved indoor air conditions.

These units reduced peak pollutant concentrations by thirty to thirty-six per cent. Their integrated sensors detected high-emission events, such as cooking, and automatically increased the device's fan speed and air-circulation rate. This raised the clean air delivery rate (CADR), allowing a greater volume of indoor air to pass through the fixed-efficiency filter during periods of elevated pollution and thereby increasing the overall rate of pollutant removal.

Replacing combustion-based cooking appliances or high-emission cleaning agents reduced VOC spikes by more than half, reflecting the importance of targeting pollutant sources directly rather than relying solely on CADR.

These engineering measures produced life-expectancy gains of 1.4 to 4.3 years, with the greatest improvements observed in buildings with limited natural ventilation and among individuals with high biological vulnerability. These populations inhaled higher baseline concentrations, meaning that reducing emissions or increasing removal efficiency produced disproportionately larger health benefits.

Importantly, the findings revealed performance patterns that static mechanistic modelling could not detect. Filtration efficiency declined gradually as filters accumulated particulate matter, causing intermittent increases in exposure unless maintenance was performed at regular intervals. The hybrid model captured these temporal declines, showing distinct periods of heightened protection immediately after installation followed by phases of reduced effectiveness as filter loading increased.

Sensor-linked air-cleaning systems also exhibited nonlinear behaviour; they responded aggressively to large pollution events but became less responsive over time due to sensor drift, fan-wear, or diminished calibration accuracy. Such real-world degradation patterns altered the long-term exposure trajectory in ways that static models, which assume constant efficacy, systematically underestimate.

By integrating these time-varying performance characteristics, the hybrid model recalculated exposure–risk trajectories continuously and captured the rise and fall of protection over months and seasons. These dynamic patterns demonstrated that engineering interventions deliver strong early benefits but require sustained upkeep, recalibration, and filter replacement to maintain long-term effectiveness. The hybrid model therefore provided a more realistic estimation of the endurance of engineering solutions compared with mechanistic modelling alone.

Behavioural interventions exhibited the widest variability in effectiveness because they depended on socio-economic conditions, cultural practices, daily routines, and the practical constraints of households. Consistent use of cooking exhaust hoods reduced NO_2 peaks by

forty-five to sixty per cent, and the complete elimination of indoor smoking reduced overall pollutant burden by thirty-two to fifty-seven per cent, producing substantial life-expectancy gains between 2.9 and 4.8 years.

These behaviours targeted highly potent pollutant sources, leading to immediate reductions in toxic chemical load. Timed window-opening reduced pollutant accumulation by fourteen to twenty per cent; however, its effectiveness varied sharply because it relied on favourable outdoor air quality and occupant availability.

The findings showed that the benefits of window-based strategies could disappear entirely during high-traffic periods, haze events, or weather conditions that discouraged natural ventilation. Unlike engineered systems, behavioural measures could not guarantee consistency, and this variability produced exposure patterns that fluctuated significantly across days and seasons.

The hybrid model captured these time-dependent fluctuations by learning when behavioural changes produced meaningful reductions in risk and when they yielded negligible or even adverse impacts. For instance, the hybrid system detected situations in which window-opening during polluted outdoor conditions temporarily increased indoor pollutant levels, a nuance that would not be visible in mechanistic models operating on long-term averages.

The combined intervention scenarios produced the most substantial improvements overall. When architectural, engineering, and behavioural modifications were implemented together, pollutant reductions were synergistic rather than merely additive. Improved ventilation enhanced the effectiveness of filtration by increasing the volume of air passing through the filter.

Reduced pollutant emissions stabilised the operational behaviour of mechanical systems by preventing overload during high-emission periods. Behaviourally driven reductions in pollutant generation amplified the benefits of structural and engineering modifications, creating an indoor environment where exposure levels remained consistently low across both routine and unexpected pollutant events.

The hybrid model revealed nonlinear interactions that mechanistic modelling did not detect. These included threshold effects, where a modest reduction in emission strength triggered disproportionately large improvements in filtration performance; feedback loops, where enhanced airflow altered pollutant mixing dynamics; and compounding benefits, where behavioural and structural modifications together reduced peak pollutant concentrations to levels unattainable through individual interventions.

Life-expectancy gains reached up to 5.4 years in individuals with high biological vulnerability, demonstrating that multi-layered strategies provided the most meaningful long-term protection against pollutant-driven ageing processes.

Overall, the findings showed that real-world intervention performance is shaped by the constantly changing interplay of human behaviour, environmental variability, building characteristics, and biological sensitivity. The hybrid modelling framework was crucial for

uncovering these dynamics.

By integrating mechanistic equations with AI-driven pattern recognition, it provided a scientifically detailed and context-responsive representation of exposure–risk relationships over time. The hybrid model therefore offered a significantly more accurate prediction of intervention effectiveness compared with mechanistic modelling alone, capturing both the magnitude and the temporal evolution of health benefits across diverse living environments.

Multi-Criteria Optimisation

The multi-criteria optimisation framework identified which intervention combinations produced the greatest life-expectancy gains per unit of cost, effort, and societal investment. This evaluation was grounded entirely in the hybrid model's dynamic predictions. Unlike traditional optimisation based on fixed pollutant reductions, this framework used hour-by-hour life-expectancy outputs that reflected real fluctuations in exposure, intervention performance, occupant behaviour, and environmental conditions. These dynamic gains constituted the "usefulness" value in the optimisation analysis.

The life-expectancy gains fed into the optimisation ranged from 0.6 to 5.4 years, depending on building constraints, intervention type, and biological vulnerability. Crucially, the hybrid model demonstrated that static mechanistic estimates—typically used in earlier environmental-health optimisation studies—overestimated benefits by 18 to 42 per cent, because they assumed constant filtration efficiency, stable window-use patterns, and unchanging infiltration rates.

The hybrid predictions instead showed that these values changed continuously. For example, high-efficiency filters that initially reduced $PM_{2.5}$ by 52 per cent declined to 31 to 37 per cent removal efficiency after four to five months of particle loading if maintenance was delayed. Similarly, window-opening benefits varied by up to 55 per cent across seasons depending on outdoor NO_2 and $PM_{2.5}$ levels.

These temporal shifts mattered because they changed cumulative exposure and thus altered life expectancy. An intervention that yielded a projected gain of 2.4 years under mechanistic assumptions produced only 1.5 years under hybrid modelling once seasonal variability, behavioural inconsistency, and filter decay were incorporated. This made the hybrid model indispensable for identifying strategies that provided stable, enduring health improvements.

Cost was evaluated across each intervention's full operational lifespan. Instead of assigning a one-off amount, the optimisation accounted for recurring costs such as filter replacement (every 3–6 months), sensor recalibration (annually), increased electricity consumption for portable air cleaners (averaging 25–90 kWh per month depending on CADR), and potential appliance upgrades. When these long-term costs were evaluated against health gains, the cost-effectiveness of interventions varied widely.

Structural interventions (e.g., reducing façade infiltration) produced gains of 1.3–1.7 years at a moderate cost only once, while portable filtration required repeated investment but produced larger gains of 2.8–4.3 years in high-exposure households. Behavioural changes had negligible

monetary cost yet produced high returns—for example, eliminating indoor smoking resulted in 2.9–4.8 years of increased life expectancy at virtually zero financial expenditure.

Durability and feasibility were integrated through constraints that adjusted the ranking of each intervention. Architectural modifications were penalised in older buildings lacking retrofit flexibility, and filtration upgrades were penalised in homes with inadequate airflow pathways that reduced effective CADR by 15–25 per cent. Behavioural interventions received feasibility penalties when time-use data showed inconsistent cooking, cleaning, or ventilation routines. These adjustments prevented the optimisation from recommending high-performing strategies that would fail under real household conditions.

Equity emerged as a central dimension of the optimisation. The hybrid model showed that low-income households experienced greater exposure volatility—short, intense pollutant peaks that increased biological burden disproportionately. These households achieved the largest gains from engineered or structural interventions, with life-expectancy improvements 34–70 per cent higher than those in higher-income households exposed to more stable pollutant patterns. For example, reducing infiltration lowered NO_2-related risk by 17 per cent in high-income homes but by 28 per cent in low-income homes, where outdoor pollution penetrated more readily.

By integrating dynamic health gains, full life-cycle costs, real-world feasibility, and equity weighting into a single optimisation engine, the study identified intervention packages that maximised both biological benefit and societal value.

Stakeholder Engagement and Policy Translation

Stakeholder engagement ensured that the scientifically derived interventions were realistic, culturally acceptable, and feasible across diverse living environments. Policymakers, building managers, engineers, health officers, and community representatives participated in structured consultations designed to evaluate the practical viability of each intervention identified through the hybrid model.

Their feedback provided essential information on cost barriers, operational constraints, installation challenges, behavioural feasibility, and long-term maintenance requirements. These engagements were conducted through a structured digital interface that allowed stakeholders to explore model outputs interactively and provide scenario-specific feedback in real time.

Several concrete insights emerged. Architectural strategies that produced measurable pollutant reductions in simulations were sometimes restricted by building codes or renovation limits in older residential blocks. For example, although increasing effective window-opening area by one-third reduced long-term $PM_{2.5}$ by up to twenty-two per cent in the modelling stage, stakeholders noted that many households could not modify window structures without regulatory approval, reducing real-world feasibility by an estimated thirty to forty per cent.

Similarly, façade-sealing measures that lowered indoor NO_2 by thirteen to seventeen per cent required upfront costs that were prohibitive for lower-income households, limiting adoption despite clear scientific benefits.

Engineering interventions produced stronger pollutant reductions but introduced new practical challenges. Sensor-augmented portable air cleaners, which reduced peak $PM_{2.5}$ concentrations by thirty to thirty-six per cent and increased clean air delivery rate (CADR) during high-emission events, required filter replacement every three to five months.

Stakeholders indicated that consistent maintenance could realistically fall to fifty to sixty per cent compliance, reducing long-term effectiveness by up to twenty per cent in vulnerable households. These insights helped refine the optimisation framework by reducing feasibility weightings for interventions requiring regular servicing or technical familiarity.

Behavioural interventions also required contextual interpretation. Exhaust-hood use, which lowered NO_2 peaks by forty-five to sixty per cent, was strongly dependent on cultural cooking practices and household routines. Stakeholders estimated that sustained adherence would vary between forty and eighty per cent across communities, suggesting that life-expectancy gains modelled at 1.0 to 1.8 years would vary accordingly.

Window-opening strategies were particularly sensitive: while timed ventilation reduced pollutant accumulation by fourteen to twenty per cent in simulation, stakeholders highlighted that haze seasons, security concerns, and heat discomfort could reduce uptake by as much as fifty per cent.

These stakeholder-derived constraints were integrated directly into the optimisation process. The hybrid model recalculated life-expectancy gains after adjusting for expected compliance, maintenance decay, and resource limitations. This produced more realistic rankings of intervention usefulness.

For instance, although engineering approaches delivered life-expectancy gains of 1.4 to 4.3 years in ideal conditions, stakeholder-informed adjustments reduced long-term projected gains by ten to thirty per cent depending on household capacity. Mechanistic models could not have integrated these behavioural or socio-economic conditions; only the hybrid model, with its dynamic recalculation capability, could reflect such real-world variability.

A major translational advancement came from designing the system so that ordinary households could supply the hybrid model with continuous, high-quality indoor air pollutant data without financial burden or technical expertise. This was achieved through a futuristic, ultra-low-cost multi-pollutant IAQ monitor engineered to be affordable even in developing countries.

Each device combined solid-state sensors for $PM_{2.5}$, microelectronic gas sensors for NO_2, O_3, and VOCs, miniaturised electrochemical sensors for formaldehyde, and a compact temperature–humidity module integrated into a single unit costing less than ten to fifteen dollars to manufacture at scale. The device operated on low power, required no consumables, and communicated wirelessly with a central home station.

Several sensors—typically four to six per apartment depending on floor area—were deployed across key micro-environments such as the kitchen, living room, bedrooms, and circulation spaces. Each sensor was designed as a low-cost, multi-pollutant indoor air quality monitor

capable of continuously measuring all pollutants of interest ($PM_{2.5}$, NO_2, O_3, VOCs, and formaldehyde). By monitoring multiple pollutants simultaneously and at high temporal resolution, the system captured the spatial and temporal variability essential for high-confidence hybrid modelling.

These satellite units transmitted continuous pollutant data to a central station located within the apartment. The central station served as the local processing hub. It performed preliminary data cleaning and eliminated errors that originated from the sensor rather than from the real environment. It also synchronised timestamps across units. When needed, it integrated pollutant readings with occupant-provided information about household activities, such as cooking.

The system was designed to minimise user burden. Therefore, the central station relied primarily on pollutant signatures themselves, and not on frequent manual entries, to identify emission events, behaviour-linked patterns, and source-specific pollutant fingerprints.

Once pre-processed, the central station forwarded the cleaned and harmonised data to the hybrid model through encrypted cloud channels. This architecture ensured that all computationally intensive modelling, including dynamic exposure prediction, biological risk estimation, intervention-response analysis, and recalculation of life-expectancy trajectories, occurred on secure external servers rather than home devices.

As a result, households could participate regardless of whether they owned newer smartphones or lived in settings with older digital infrastructure, preserving accessibility and equity in real-world deployment.

A major enhancement designed specifically to support everyday users was the development of an intelligent inference engine trained to identify pollutant sources autonomously. To prevent households from having to repeatedly log daily activities, which would be tiring, burdensome, and unrealistic for long-term use, the hybrid system incorporated a high-resolution source-apportionment module.

This module could determine, with up to 99 per cent confidence, whether a measured pollutant originated indoors or outdoors. It also identified which appliance or material produced the pollutant, which human activity triggered it, and which environmental condition amplified it. This automated capability ensured precise source identification without requiring continuous user input.

The inference engine recognised unique temporal fingerprints, chemical signatures, and co-occurrence patterns, for example, the characteristic spike in NO_2, CO_2, and humidity produced by gas-stove ignition, or the particle-size and VOC-pattern combination released during frying with oil-based sauces. This allowed the system to identify the cause of pollutant events without requiring any manual user input, ensuring both scientific precision and practical usability.

This design made the system practical for everyday users. Residents did not need to purchase premium devices, manage calibration schedules, or understand sensor technology. The multi-pollutant monitors were designed to be essentially maintenance-free, with a lifespan of three to

five years. In households that declined to install indoor sensors, the decision-support tool defaulted to micro-regional baseline data from shared outdoor and corridor-mounted reference nodes maintained by municipal agencies.

This ensured that every household still received basic risk assessments and general behavioural guidance. Households that did participate, however, received far more precise, personalised, and timely health-protection insights, which created a clear incentive for engagement. Non-participating households also benefited indirectly through neighbourhood-level risk mapping, improved building-management practices, and community-wide IAQ interventions guided by anonymised data from participating units.

The translational outcome was the development of a digital decision-support tool powered entirely by the hybrid model. This tool operated through an intuitive, dashboard-style interface that allowed stakeholders to manipulate variables such as window-use frequency, filter-replacement intervals, appliance type, household income, and building design parameters. The hybrid engine recalculated exposure levels, biological risk, and life expectancy in real time based on these inputs.

The interface displayed dynamic pollutant curves, exposure-dose timelines, and survival-probability projections, enabling users to see instantly how daily routines, maintenance patterns, or structural modifications altered health outcomes. Users could visualise how pollutant concentrations changed hour by hour, how cumulative exposure built up over months, and how life-expectancy gains shifted under different compliance patterns.

This tool revealed patterns that would otherwise remain hidden. For example, simulations showed that households in the lowest-income bracket experienced up to sixty per cent greater exposure variability because of inconsistent ventilation, older building envelopes, and limited access to engineering controls.

Yet these same households also gained the largest improvements, up to 5.4 years in highly vulnerable groups, when multi-layered interventions were applied. Policymakers used these insights to refine recommendations, prioritising such households for subsidies, targeted retrofits, and behavioural-support programmes.

Overall, stakeholder engagement and the decision-support tool together enabled the scientifically robust findings of Research Question 3 to be translated into feasible, equitable, and actionable indoor-air-quality strategies for real-world implementation.

Synthesis and Broader Significance

The findings of Research Question 3 demonstrate that architectural, engineering, and behavioural interventions can significantly improve indoor air quality and extend life expectancy when evaluated through a dynamic hybrid mechanistic–AI framework.

These results reject the null hypothesis and confirm that at least one intervention combination, and in most cases several combinations, produces measurable gains in exposure reduction, biological risk mitigation, and healthy-life-years. The study shows that health benefits emerge

not from isolated measures but from the interaction of building design, mechanical systems, occupant routines, and socio-economic conditions.

A major contribution of the findings is the demonstration that intervention performance is inherently dynamic. Indoor pollutant levels changed hourly with weather, season, temperature gradients, ventilation behaviour, and cooking intensity. Filter performance declined as particulate matter accumulated, while behavioural compliance varied daily.

These fluctuations caused substantial shifts in exposure–dose trajectories. Classical mechanistic modelling could not capture these time-sensitive changes; the hybrid model provided real-time recalculations that reflected actual household conditions and delivered far more realistic life-expectancy predictions.

The results also highlight that combined interventions offer the strongest benefits. Architectural measures reduced pollutant ingress. Engineering strategies reduced peak concentrations and improved clean-air delivery during high-emission events. Behavioural practices reduced emissions at the source. When these components were implemented together, their effects were synergistic.

In highly vulnerable groups, life-expectancy gains reached up to 5.4 years, underscoring the potential for targeted, multi-layered strategies to deliver meaningful long-term health protection.

Intervention effectiveness varied across households. Buildings with limited natural ventilation responded most strongly to engineering measures. Households with frequent high-temperature cooking gained more from exhaust hoods. Vulnerable individuals experienced large biological improvements even with modest exposure reductions.

Conversely, some scientifically effective interventions had low practical feasibility due to regulatory restrictions, upfront cost, or cultural acceptability. These disparities reinforced the need for value-oriented, context-sensitive recommendations rather than uniform prescriptions.

The multi-criteria optimisation framework provided a systematic way to balance health impact, cost, feasibility, maintenance burden, and equity. It revealed that low-cost improvements often delivered strong value in resource-constrained settings, while higher-cost mechanical systems were more suitable for high-density or poorly ventilated buildings. Lower-income households benefited disproportionately from targeted interventions as they experienced more unstable exposure patterns and higher baseline risk.

Stakeholder insights ensured that the modelled interventions aligned with real-world constraints and everyday lived experience. Their feedback was integrated into the hybrid system, producing life-expectancy projections grounded in realistic compliance patterns rather than idealised assumptions.

The digital decision-support tool translated these findings into an accessible platform that allows policymakers, practitioners, and communities to visualise how interventions alter pollutant levels, biological risk, and projected longevity. The broader significance of RQ3 lies in

demonstrating that indoor air quality is a modifiable determinant of population health and that scientifically guided, equitable interventions can sustainably extend healthy-life expectancy across diverse societal contexts.

5 ……

Ayanfe's career in academia unfolded with a clarity and steadiness that surprised even those who had witnessed her transformation during her doctoral study. By the time she completed her PhD, she had already demonstrated that she had overcome the flaw that once defined her, that instinctive leap into solutions without understanding the root causes of problems.

Her doctoral research had not only established a scientifically rigorous framework connecting indoor air pollutant exposure, biological ageing, and life expectancy, it had also produced operational tools capable of transforming policy, practice, and national planning. She had built mechanistic and AI-integrated models that captured the dynamics of household exposure, community-level variability, behavioural patterns, biological vulnerability, and socio-economic context.

These models had already gained international attention before she graduated, reinforcing that she had become a different thinker entirely, one who approached problems with depth, discipline, and value-oriented intent.

After completing her PhD, she accepted a competitive research fellow position at the same university. The role allowed her to refine and expand the tools she had created. Her doctoral work had focused primarily on the major indoor-air pollutants that dominate global risk assessments, namely fine particulate matter ($PM_{2.5}$), nitrogen dioxide (NO_2), ozone (O_3), volatile organic compounds (VOCs), and formaldehyde (HCHO).

As she expanded her collaborations with international research teams, Ayanfe broadened her models to include a wider spectrum of pollutants and indoor chemical processes. Her doctoral work had already taught her the complexity of indoor environments, but global datasets revealed additional pollutant classes and reaction pathways that strengthened predictive accuracy.

Household combustion by-products, moisture-driven transformations, microbial metabolites, and secondary pollutants formed through interactions between indoor emissions and infiltrated outdoor contaminants required deeper modelling. Microbial metabolites became an especially important addition because they helped explain exposure pathways that traditional pollutant models overlooked.

These metabolites referred to the chemical substances produced by bacteria and fungi as they grew, spread, and interacted with indoor surfaces. Some appeared as simple gases released when microbes metabolised organic matter, while others formed through more complex biochemical processes triggered by moisture, humidity, or the breakdown of building materials. Examples included microbial volatile organic compounds with musty or earthy odours, metabolic by-products released during fungal respiration, and compounds produced when bacteria interacted with damp dust or residues on walls and floors.

Although many microbial metabolites were subtle and often invisible to the senses, they influenced respiratory irritation, cognitive clarity, immune stress, and overall indoor comfort. Their presence also signalled broader indoor environmental conditions, such as hidden dampness or microbial hotspots that contributed to long-term exposure burdens.

Incorporating microbial metabolites into her framework allowed Ayanfe to capture the combined biological and chemical dimensions of indoor environments, improving the accuracy of risk predictions across diverse building types and climates.

With these refinements in place, her expanded framework naturally incorporated a range of additional pollutant categories that interacted with or amplified these processes. These categories included ultrafine particles, semi-volatile organic compounds, reactive carbonyls, and secondary organic aerosols. They also included nitrogenated by-products, microbial volatiles, indoor-generated oxidants, and species formed through indoor photochemistry.

Her postdoctoral work therefore focused on expanding the scope of pollutants considered in her models. She enhanced the sensory capabilities of the prototype devices she had created during her PhD, integrating advanced nano-sensing components, low detection limit photometric modules, and selective catalytic micro-reactors.

This level of refinement became possible only because technological progress in those years was exceptionally rapid. Global investment in nanoscale materials, photonic engineering, microfluidic channels, and AI-integrated embedded systems accelerated at an unprecedented pace.

Ayanfe recognised the opportunity early and capitalised on it. She collaborated with chemical engineers, material scientists, and molecular toxicologists to validate the new sensing architecture. Together, they pushed the limits of what indoor-environment sensors could achieve, developing laboratory prototypes that could detect pollutants at concentrations previously considered unreachable outside high-grade analytical facilities.

Remarkably, these advancements meant that sensors no longer needed to be large, expensive, or laboratory-bound. Her team was able to fabricate compact sensing arrays directly in the university’s cleanroom, embedding nano-catalytic films, selective adsorbent layers, and AI-assisted processing units into handheld devices.

This technological leap allowed her to measure previously unmonitored pollutants and chemical intermediates with extraordinary precision, including transient reaction products that existed for only seconds.

By combining these expanded pollutant measurements with her existing mechanistic and AI-driven framework, she substantially increased the accuracy and stability of life expectancy predictions. The improved system was now capable of accounting for interactions among pollutants, for example reactions between ozone and VOCs that generate new compounds indoors, and it could better capture exposure peaks that were previously invisible due to sensor limitations.

Her work marked one of the earliest examples of how rapidly advancing technology, when paired with rigorous scientific framing, could redefine what indoor air research was capable of achieving. The refinement allowed her models to deliver more precise estimates of cumulative exposure burden and its biological implications, making them even more valuable to governments and planning agencies.

The improved predictive stability, supported by her expanded sensing capabilities and cross-disciplinary validation efforts, positioned her work at the forefront of indoor air research. Her outputs began informing national advisory reports, international working groups, and regional health-impact assessments, signalling that her framework had matured into a robust policy tool.

Her research fellow appointment concluded with exceptional reviews, and she transitioned not to a lecturer position but directly to Senior Lecturer, a move supported by her publication record, her rapidly growing international collaborations, and the societal relevance of her work.

As a Senior Lecturer, she broadened the application of her research beyond her doctoral study region. She worked with policymakers in low-income countries where households relied on biomass cooking and improvised ventilation. She tested her models in middle income countries undergoing rapid urbanisation, where building stock varied widely in quality and informal housing contributed to complex exposure patterns. She collaborated with high income nations concerned about airtight construction, synthetic materials, and energy efficient retrofits that sometimes compromised ventilation.

Through each collaboration, she adapted her tools to accommodate different climates, building typologies, socio-economic patterns, and cultural behaviours. The flexibility of her framework allowed it to serve as a universal decision-support system, capable of guiding interventions across diverse contexts.

Her reputation grew, not because she promoted her own contributions, but because her models consistently revealed insights that countries had struggled to obtain for decades. Her work showed, for example, why certain populations with similar outdoor air quality profiles experienced different health trajectories due to indoor-origin pollutants, behavioural patterns, or building characteristics.

She was able to demonstrate that two neighbourhoods breathing the same outdoor air could still have very different health and ageing outcomes simply because one group cooked with poorly ventilated stoves, lived in buildings that trapped moisture, or used cleaning products that released reactive chemicals into the air.

It demonstrated how specific interventions, such as improved kitchen ventilation, changes in cleaning-product formulations, or reconfiguration of ventilation schedules, could yield measurable gains in healthy life expectancy.

Her models translated complex chemical and biological interactions into concrete numbers that policy-makers could understand, such as months or years of life gained when a home improved its ventilation routine, or the reduction in biological stress when families switched to low-emission cleaning agents.

For the first time, governments and public-health agencies could see indoor air as something quantifiable rather than mysterious or intangible. They could identify exactly which everyday activities increased exposure, which building features mattered most, and which groups were silently carrying the highest burdens. Decision-makers described her work as providing a map through a terrain that had always been felt but never clearly seen.

Her contributions were grounded in evidence, not aesthetics or speculation. She had become the researcher she once wished she could be, someone who allowed problems to speak fully before attempting to solve them.

Her rising influence soon led to her promotion to Associate Professor. In this role, she expanded her international research network, secured grants that enabled her to deploy her sensing framework in more than twenty countries, and supervised doctoral candidates who carried her work into new scientific territories.

Her global collaborations grew naturally because many countries, regardless of income level, were facing similar problems: people were getting sick or ageing faster than expected, yet no one understood which indoor factors were quietly causing the harm. Governments wanted answers that were scientifically grounded, practical, and applicable across different climates and housing systems. Ayanfe's framework provided exactly that.

Her students studied topics such as indoor chemical transformation, socio-economic determinants of exposure vulnerability, machine learning models for early detection of cumulative biological burden, and the optimisation of interventions across climates. Each student adapted her foundational work to a specific context: some focused on informal settlements with limited ventilation, others on energy-efficient buildings that accidentally trapped pollutants indoors, and others on schools with high occupancy levels.

Her supervision style became well-known within the university; she guided them not toward speed, but toward understanding. She often insisted that they map problems repeatedly before proposing even a single idea. Many of her students said that being mentored by her felt like learning how to see problems for the first time, as if she had taught them a new way to think rather than simply a set of research methods.

Her models were adopted in national planning dashboards in several regions. Ministries of health and environment used her tools to set more realistic and equitable targets for indoor environments, recognising that long-term national development could not be separated from the quality of air inside homes, schools, and workplaces.

Some governments used her framework to redesign public housing ventilation policies, others revised building codes for cleaning-product use and filtration standards, while several ministries used her predictions to identify neighbourhoods where vulnerable groups such as the elderly or children faced higher cumulative exposure burdens.

International organisations cited her work as a foundation for integrating indoor air into global environmental burden of disease frameworks. For the first time, indoor air began appearing in international discussions about ageing populations, chronic disease prevention, and

sustainable development.

Policymakers, engineers, and public health experts began to acknowledge that the air people breathe inside everyday spaces shapes not only comfort, but the number of healthy years they are likely to live. Ayanfe's work provided the missing scientific bridge that allowed these fields to speak to one another.

Her final academic ascension came when she was appointed Full Professor, a recognition not only of her scientific contributions but also of the profound societal impact her work had generated across continents. By this stage, her research had reshaped how countries conceptualised indoor environmental health, redefining it from a peripheral concern into a central determinant of national wellbeing. As a Full Professor, she dedicated herself to three intertwined commitments: advancing the science, training the next generation of researchers, and influencing global policy.

She advocated for an international shift toward value-oriented indoor environmental health strategies grounded in measurable impact on life expectancy rather than surface-level indicators of comfort or aesthetics. To her, a healthy building was not one that looked modern, but one that demonstrably reduced biological burden and added years of healthy living.

This message resonated widely because her models had already shown governments that long-term indoor exposures could quietly erode national health gains if left unaddressed. Ministries began inviting her to contribute to major white papers on national longevity strategies, and global agencies sought her advice when developing frameworks that integrated indoor air quality into climate resilience, urban planning, and public health agendas.

In her teaching role, she became a formative presence in the lives of emerging scholars. Students travelled from multiple regions, some from countries that had adopted her models into national policy, simply to learn how she thought. She emphasised rigorous problem framing, humility in the face of complexity, and a deep respect for evidence.

Many of her doctoral candidates later said that being supervised by her felt like learning how to think for the first time because she pushed them to examine every assumption and understand the full landscape of a problem before proposing even the smallest intervention. Her mentorship created a new generation of researchers who approached indoor environmental health not as a technical field alone, but as a domain that demanded intellectual discipline, critical reasoning, and an unwavering commitment to human wellbeing.

Her influence on global policy continued to expand. International development agencies incorporated her frameworks into housing-quality evaluations for low-income regions, while high-income countries used her models to guide retrofitting programmes aimed at improving healthy-life expectancy among ageing populations. Her work informed city-scale environmental health dashboards, national indoor air audits, and international training programmes for public-health officers.

In every stage of her career, she demonstrated the depth of the transformation that had begun years earlier in a quiet classroom, when she first realised she had never understood problems deeply enough. Now, she had become a global authority precisely because she did. Her journey became a reminder to the academic world and to policymakers alike that true solutions are never born from haste, but from listening deeply to the problem and allowing evidence to lead the way.

6 ……

Ayanfe's transformation was never limited to lecture theatres, research labs, or policy rooms. The change that began years earlier in the quiet discomfort of her undergraduate seminars eventually grew into a way of being that touched every part of her personal life. She had once been someone who rushed, someone whose fears pushed her forward without pause. She had hurried through childhood, hurried through school, hurried through ideas, and hurried even through the moments that needed patience.

Overcoming her flaw in academia reshaped the way she approached the world beyond research. She learnt not only to think differently but to live differently. The shift was not dramatic or sudden. It unfolded slowly, like breath returning to lungs that had forgotten how to expand fully. With each year she found herself embracing the same deliberate stillness she had once feared, discovering that the pace of understanding could also be the pace of living.

In her early career, her instinct to slow down remained intentional rather than natural. She still felt the tug of her old habits whenever problems arose in relationships or daily life. When she encountered disagreements with colleagues or tensions at home, she initially reached for quick solutions, driven by the impulse to resolve discomfort immediately.

However, with time, she realised that human relationships were no different from the indoor environments she studied. They were shaped by invisible forces, underlying conditions, accumulated patterns, and unspoken histories. Just as she had learnt to search for root causes in buildings, she began to search for root causes in herself, in her relationships, and in the lives of the people she cared about.

She noticed that the moments that frustrated her most were rarely caused by the situation itself, but by the unexamined assumptions she carried into it. This understanding became liberating, giving her the courage to pause rather than react.

She used to interpret conflict as a signal to fix something instantly. Now she learnt to interpret it as a signal to understand. Conversations that she once would have rushed through became gentler, more reflective, and more centred on listening. Her partners in research and life noticed the difference. She had always been brilliant, but now she was steady. She had always been confident, but now she was thoughtful.

People felt safer around her, not because she had become softer, but because she had become more present. Her relationships no longer suffered from the hurried certainty that once defined her approach to everything. People who once found her intensity intimidating now

found comfort in her calmness. She became someone others trusted with their vulnerabilities, not because she had perfect answers but because she offered patient understanding.

Her marriage came during a period when her career was flourishing internationally. Some of her colleagues feared that the pace and pressure of her work would make personal commitment difficult. Yet those closest to her understood that the same transformation that had strengthened her academic work now supported her personal life.

She approached her marriage with the same humility she applied to scientific inquiry. She listened as much as she spoke. She asked herself whether she was responding to the real issue or only to the surface of it. She paused before concluding. She made space for questions. She gave her husband the same patience she had learnt to give her research.

Her husband often said that loving her felt like breathing in fresh morning air: grounding, clear, and revitalising. Together they built a relationship defined not by perfection but by clarity, trust, and shared understanding. They created rituals of connection—slow Sunday walks, evening reflections, and conversations that wandered gently through difficult topics without fear of judgement.

This shift did not happen without effort. She had grown up in systems that rewarded speed, performance, and presentation. Those patterns had become part of her. Breaking them required deliberate practice. When disagreements arose between her and her husband, she resisted the urge to repair the situation instantly. She began to ask: What is the root cause here? What feeling sits beneath the emotion I am seeing? What need sits beneath the words being spoken? The patience she showed in these moments created a home environment built not on the absence of conflict but on the presence of understanding.

Her marriage thrived because she no longer treated problems as irritations to eliminate, but as truths to learn from. Her husband often commented that their home felt like a place where emotions could breathe. Nothing was hurried, nothing was dismissed, and nothing was forced into simplicity. She had created, within her family, the same environment she advocated for in buildings: balanced, healthy, and deeply supportive.

Her relationship with her parents deepened too. In her youth, she had never understood her mother’s quiet anxieties or her father’s occasional hesitation. After her transformation and learning about the fire incident that happened when she was a baby, she recognised the layers beneath their behaviours. Her mother had lived through a fire that nearly harmed her child. Her father had rushed through smoke with Ayanfe held close to his chest. These events had shaped them just as her own childhood fear of still air had silently shaped her.

She began to ask them about their pasts, their fears, and their hopes in ways she had never done before. Their conversations became richer, warmer, and more honest. Her parents said that she had not only grown into a scholar but into a far more compassionate daughter. She realised that all families carried invisible histories, much like buildings carried invisible pollutants. Once she understood this, she approached her parents with the same care she applied to her research: patient observation, gentle curiosity, and deep respect for the unseen.

Motherhood brought yet another evolution in her personal journey. She had always feared that her old flaw, the habit of rushing, might return when life became chaotic and demanding. Yet holding her newborn child for the first time grounded her in ways she did not expect. She found that the skills she had developed in her research became invaluable in parenting.

When her child cried, she no longer rushed to force a solution. She observed, listened, experimented gently, and allowed time to reveal the cause. When her child grew older and began to ask questions, she encouraged curiosity rather than speed. She told her children that it was acceptable not to know answers immediately. She told them that thinking was a journey, not a race. She also introduced playful rituals grounded in exploration, such as asking each child at bedtime what question had puzzled them that day. These small acts nurtured a culture of thoughtful wonder in her home.

Her home became a place where questions were never dismissed and patience was never seen as weakness. Her children grew into reflective individuals who were comfortable with uncertainty, deeply attentive to the world around them, and unafraid of taking time to understand what lay beneath the surface of things.

They inherited not only her intelligence but also her wisdom, shaped by the deliberate choices she made every day as a parent. Teachers often remarked that her children asked unusually perceptive questions, the kind that revealed careful observation rather than quick assumptions. They had learnt from their mother that understanding emerged in layers, not in leaps.

Her friendships were strengthened by this transformation as well. She was no longer the friend who filled silence with solutions. She became someone who created space for others to understand themselves. Her friends turned to her not for quick advice, but for the steadiness she offered when they needed clarity.

She had become the kind of person who could hold complexity without fear, who could stay with discomfort without rushing to erase it, and who could guide others toward insight without imposing answers. Many of her friends said that conversations with her felt like standing before an open window, where the air moved gently enough for truth to settle.

As years passed, Ayanfe’s personal life came to reflect the same principles that guided her scientific contributions. She lived with the awareness that outcomes, whether in health or relationships, were shaped by countless small exposures, choices, and conditions. She understood that her presence, like indoor air, could either nourish people or burden them. She chose to nourish.

Her home was filled with warmth, patience, and a sense of shared discovery. Her children grew up knowing that their mother understood them deeply. Her husband felt supported, seen, and valued. Her parents aged with comfort, knowing that the daughter they had once carried through smoke had grown into a woman who carried others through complexity. Neighbours often remarked that entering her home felt like stepping into a calmer world, one where time moved with purpose rather than pressure.

In the end, Ayanfe's personal success did not lie in achievements or accolades, although she possessed many. It lived in the quiet harmony of her home, in the relationships she nurtured, and in the calm steadiness she brought to the people around her. She had once been a child who ran away from stillness. Now she had grown into a woman who understood that stillness was where the world revealed itself.

Her personal life became the most intimate evidence that she had not only corrected her flaw, but transformed it into a source of strength. Her legacy extended far beyond the laboratories and policy rooms that celebrated her. It lived in the hearts she touched, in the family she raised, and in the calm she carried into every part of her world.

She proved that true transformation was not only a matter of intellect but of character, patience, and compassion. In becoming the scholar she once dreamt of being, she also became the person she was always meant to be. **The End!**