Durham University

Evaluating the Effectiveness of COVID-19 Non-Pharmaceutical Interventions in the UK: An Agent-Based Simulation using Wales

Daohui Liu

Laidlaw Scholars Program

August 2026

Project Supervisor: Dr Gavin Woolman

Institute for Particle Physics Phenomenology, Durham University

Abstract

This study used an agent-based simulation of Wales (MAY/JUNE2) (IDAS-Durham, 2026a, 2026b) to evaluate COVID-19 non-pharmaceutical interventions, motivated by the UK Covid-19 Inquiry's finding that earlier lockdown timing could have averted thousands of deaths. Each week of delay compounded outcomes exponentially, raising infections by 41% and deaths/hospitalisations by 57%. Sustained behavioural interventions, especially self-isolation on mild symptoms, outperformed lockdown timing alone, and combining interventions achieved the largest reduction (78–85% across outcomes). Mortality risk varied considerably by ethnicity, largely attributable to age structure, and interventions narrowed absolute disparities between groups far more than relative ones.

Introduction and background

As the world becomes increasingly interconnected through globalisation, international travel, and trade, the movement of people and goods has reached unprecedented levels (Findlater and Bogoch, 2018). Although this interconnectedness has delivered substantial economic and social benefits, it has also increased the speed and scale at which infectious diseases can spread (Gao, Gu and Yang, 2023).

In addition, as we continue to degrade the environment and encroach into undisturbed wilderness, not only does this impact the natural world and intensify existing environmental issues, e.g. wildfires from global warming — the UK recorded its highest-ever temperature of 40.3°C in July 2022 (Met Office, 2022) — but we also increase the risk of exposure to newly arising diseases originating from wildlife (IPBES, 2020).

COVID-19 has been one of the deadliest diseases in recent human history; having caused over 7 million confirmed deaths worldwide (World Health Organization, 2026), and with almost half of those affected reporting long-term effects such as extreme fatigue, cognitive impairment and shortness of breath (Aiyegbusi et al., 2021). Undoubtedly, in the UK, it exposed the many vulnerabilities of the national healthcare system and the challenges it faced with the overwhelming number of hospital admissions (NHS England, 2026). It also revealed the lack of preparation the UK had (UK Covid-19 Inquiry, 2024) and the policymakers' inability to make decisive decisions under the information given at the time (UK Covid-19 Inquiry, 2025).

Much of the grounding for this research is informed by the UK Covid-19 Inquiry (2024, 2025), which provided a rigorous, evidence-based account of the difficulties and shortcomings in the UK's pandemic response. Drawing on extensive documentary evidence, expert testimony, and — notably — epidemiological modelling and simulation data, the Inquiry has been valuable in identifying systemic weaknesses, from delays in decision-making to gaps in pandemic preparedness. Its findings offer a well-evidenced benchmark against which simulation-based research, such as this study, can test alternative intervention strategies and timings, providing a means of quantitatively exploring how different policy choices might have altered outcomes.

This study uses an agent-based simulation approach. Agent-based models simulate individual "agents" with distinct attributes and behaviours, from which transmission emerges through their interactions rather than population-level averages (Hunter, Mac Namee and Kelleher, 2017). This makes them well suited to this particular research: individual compliance with NPIs can be varied directly rather than assumed uniform, and the resulting heterogeneity and stochasticity — super-spreading events, uneven contact networks, demographic variation — better reflect real-world outbreaks. Crucially, this also allows counterfactual scenarios, such as an earlier lockdown, to be tested which a granularity aggregate model cannot offer.

Research Methodology

Goals and Aims

This study aims to quantitatively evaluate the effects of non-pharmaceutical interventions (NPIs) and their timings — including lockdowns, isolation strategies, social distancing (reduced contact with others), and improved hygiene practices such as regular handwashing — on infection and mortality rates, with particular attention to how these outcomes varied across different age groups and different ethnic and racial demographics. By testing a range of compliance rates across these groups, this research establishes an evidence-based account of which interventions, and at what level of adherence, most effectively reduce transmission and death — and for whom.

This is not a hypothetical exercise. The UK Covid-19 Inquiry's Module 2 report found that the government's initial response was marked by delay and a lack of urgency, and that had a full lockdown been imposed just one week earlier, existing modelling suggests approximately 23,000 deaths could have been prevented in the first wave alone (UK Covid-19 Inquiry, 2025). Delay was not a neutral choice; it carried a measurable and devastating cost — one that fell disproportionately on the elder population (Buffel et al., 2023) and Black, Asian, and minority ethnic communities, who experienced markedly higher rates of infection and mortality throughout the pandemic (Public Health England, 2020). These disparities did not end with the acute phase of the pandemic: years later, people from ethnic minority backgrounds continued to face compounded barriers in accessing long COVID care, including delayed diagnosis and restricted referral pathways, and remained underrepresented in both specialist clinics and the research evidence base itself (Smyth et al., 2024).

The aim of this research, therefore, is twofold: to equip policymakers with the quantitative evidence needed to act decisively — not retrospectively — when the next global health crisis arrives, and to highlight the unequal effects on specific demographics to ensure that future interventions are designed with the needs of disproportionately affected demographic groups in mind, rather than assuming a uniform impact across the population.

Structure of the Research

This research employs a three-stage computational pipeline for agent-based epidemiological modelling: MAY for synthetic population construction (IDAS-Durham, 2026a); JUNE2, for disease simulation and analysis of the data generated from the simulation through the built-in Jupyter notebook workflow (IDAS-Durham, 2026b), and MAY World Visualiser, which allows the user to inspect the synthetic world population and offers geographical information such as schools and hospitals on different geographical scales (gavdoubleu, 2026).

MAY

The synthetic population underlying each simulation was built using MAY (IDAS-Durham, 2026a), a configuration-driven population-generation framework developed by IDAS-Durham. MAY is geography-agnostic by design, using a hierarchical system of geographical units to distribute a synthetic population across residences, schools, workplaces, and other venues, meaning it can in principle be applied to any administrative hierarchy, era, or country. For this study, the modern-day UK configuration for Wales (based on 2021 census data) was used. By giving MAY the census data collected in 2021— geographical hierarchies, age/sex demographics, household composition counts, and venue inventories — MAY produces a single HDF5 world file containing every simulated individual, their place of residence, school, workplace or care setting, and their friendship and partnership networks. Critically, the entire pipeline is controlled through YAML and CSV configuration files rather than source code changes, allowing scenario parameters (e.g., geographic scope, household structure) to be adjusted without modifying the underlying Python engine. This is useful as building larger worlds require more processing power and memory, so testing different parameters in JUNE2 on smaller worlds before using on a larger model ensures the utility of the change.

MAY World Visualiser

To inspect the synthetic population underlying each simulation, this study used the MAY World Visualiser (gavdoubleu, 2026a), a companion toolset which offers an interactive geographic interface for exploring the spatial distribution of a MAY-generated population (WorldMap) and WorldExplorer, a file-browser-style interface for navigating the geographic-unit hierarchy and inspecting population and venue data in detail. These tools were used to visually verify the structure and spatial distribution of the synthetic populations generated by MAY prior to their use as inputs to JUNE2. Importantly, this allows the user to verify the synthetic world population and individual agents closely match with the data based on the 2021 census, in order to produce the most accurate data for analysis.

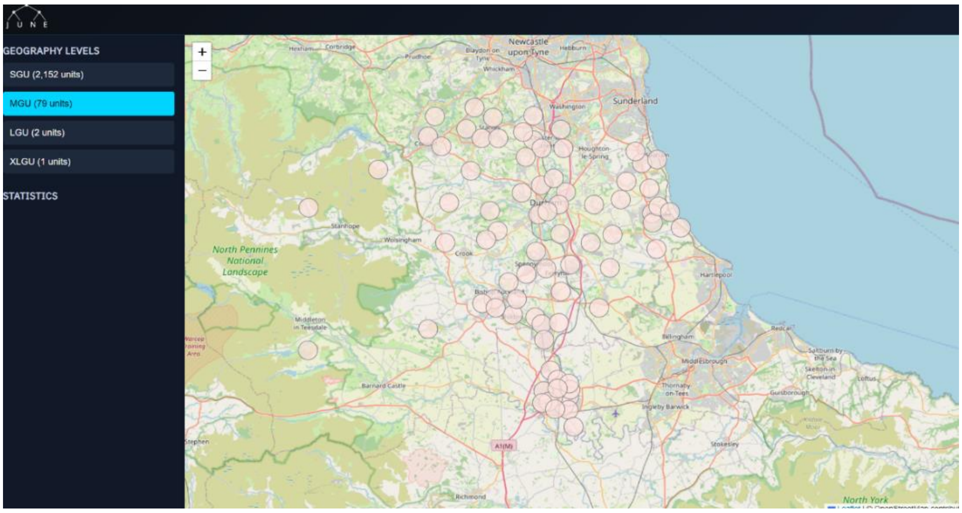

Figure 1. A view of the WorldMap of Durham and Darlington Region, generated using the MAY World Visualiser (gavdoubleu, 2026).

JUNE2

The simulation of the disease is performed through JUNE2 (IDAS-Durham, 2026b). Population files generated by MAY are used as direct inputs to JUNE2, a high-performance, event-driven, agent-based disease simulation framework, which is a C++ rewrite of the original JUNE (Aylett-Bullock et al., 2021), offering processing speed of up to 19x faster than the original, along with bug fixes. JUNE2 simulates every individual explicitly, tracking demographics, daily schedules, network ties, clinical state, and advances the simulation via discrete events (encounters, disease-state transitions, and policy triggers) rather than fixed time-steps. Each day of every individual is recorded and can be used as data for analysis. As with MAY, all scenario-defining elements — including the disease model, contact matrices, NPI's, and vaccination programmes — are specified through YAML/CSV configuration rather than source code, allowing the different NPI compliance rates central to this study to be tested without modifying the simulation engine itself. The framework also supports MPI-parallel execution, enabling a much faster time to conduct simulations at the population scale required for this research. For this study, the configuration files are adapted to model a 4-month epidemic window - closely modelled after the 2nd lockdown wave – but without vaccinations, allowing the interventions described above (lockdowns, isolation, social distancing, hygiene measures) to be varied and tested against simulated infection and mortality outcomes.

Analysis Tools (Part of JUNE2)

Output from each JUNE2 run is written to simulation_events.h5, an event log recording every infection, death, symptom change, and encounter, timestamped and linked to the individuals and venues involved. This output was processed using june_events and june_analysis, two Python libraries built specifically around this event log, which decode JUNE2's compact integer-coded fields into readable labels and join these against person- and venue-level metadata to produce analysis-ready tables. A Jupyter notebook workflow was built around these libraries, run in a dedicated WSL2/Ubuntu environment for reproducibility, to extract infection, mortality, and hospitalisation outcomes and compare results across the NPI scenarios described above.

Results

Time-dependent lockdown scenarios

To examine the effect of lockdown timing on infection, mortality, and hospitalisation outcomes, three counterfactual scenarios were compared against the baseline lockdown date: a lockdown one week earlier, one week later, and two weeks later. These are shown below:

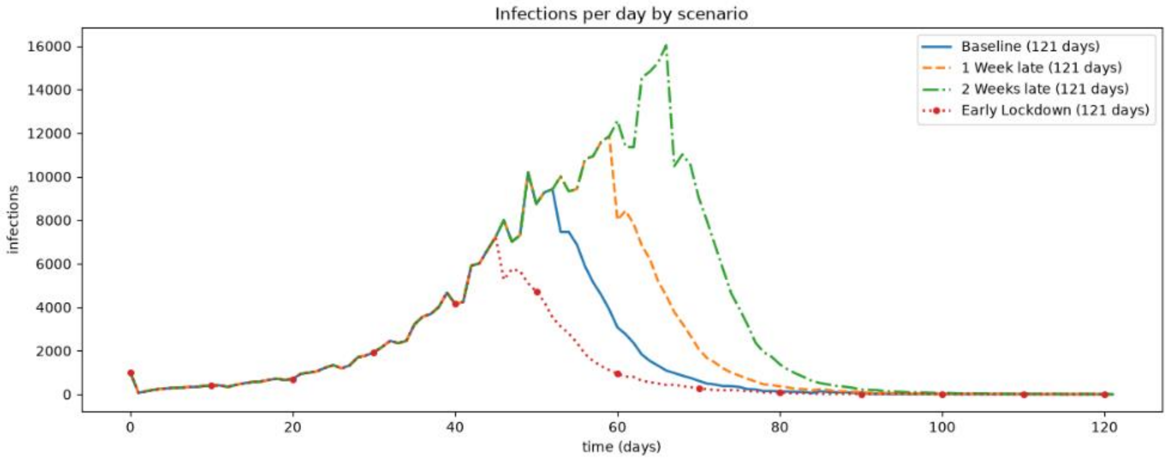

(a) A graph of infections per day for each timing scenario

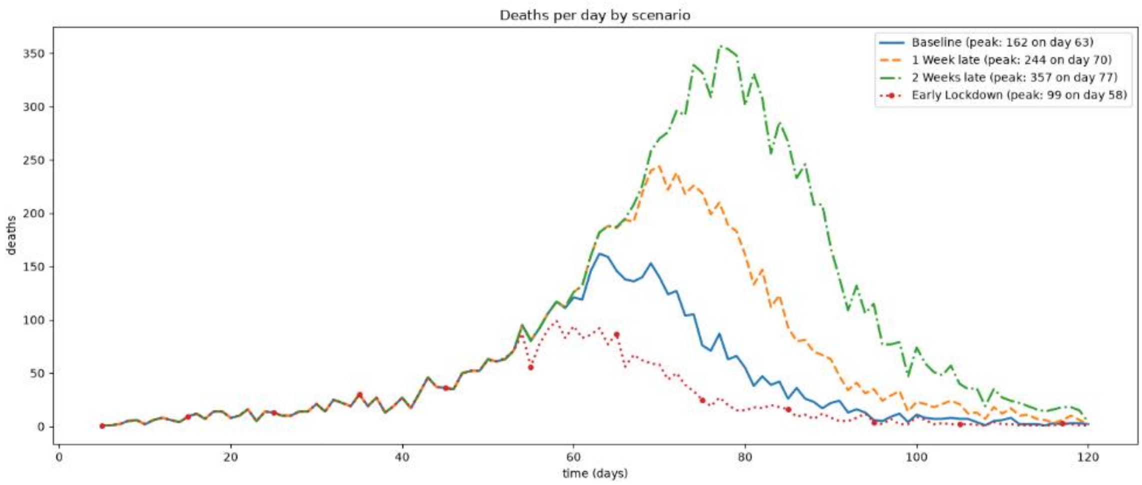

(b) A graph of deaths per day for each timing scenario

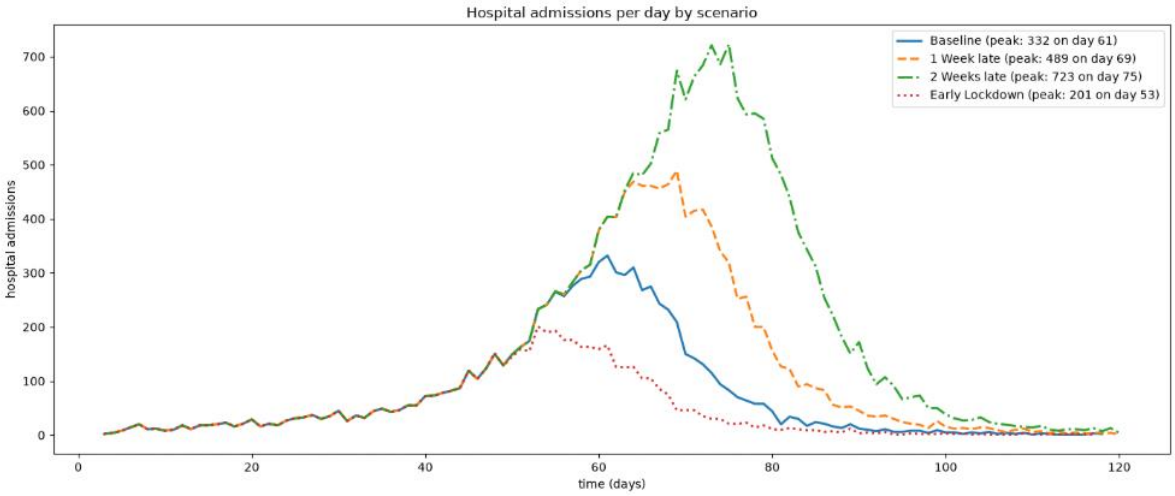

(c) A graph of the hospital admissions per day for each timing scenario

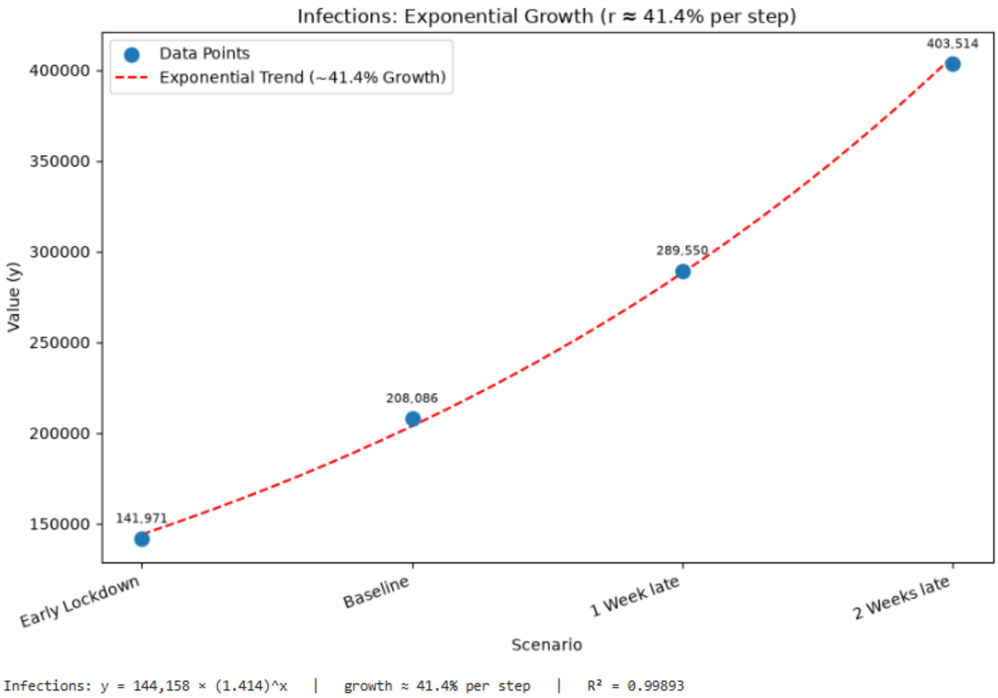

(d) A graph showing the relationship between lockdown time delay and number of infections

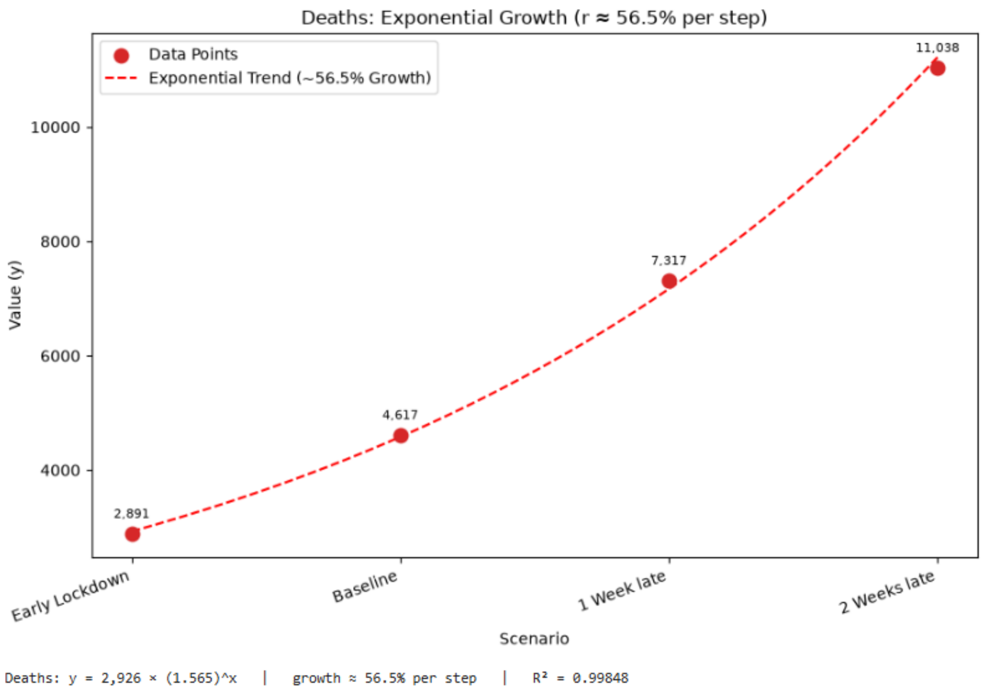

(e) A graph showing the relationship between lockdown time delay and number of deaths

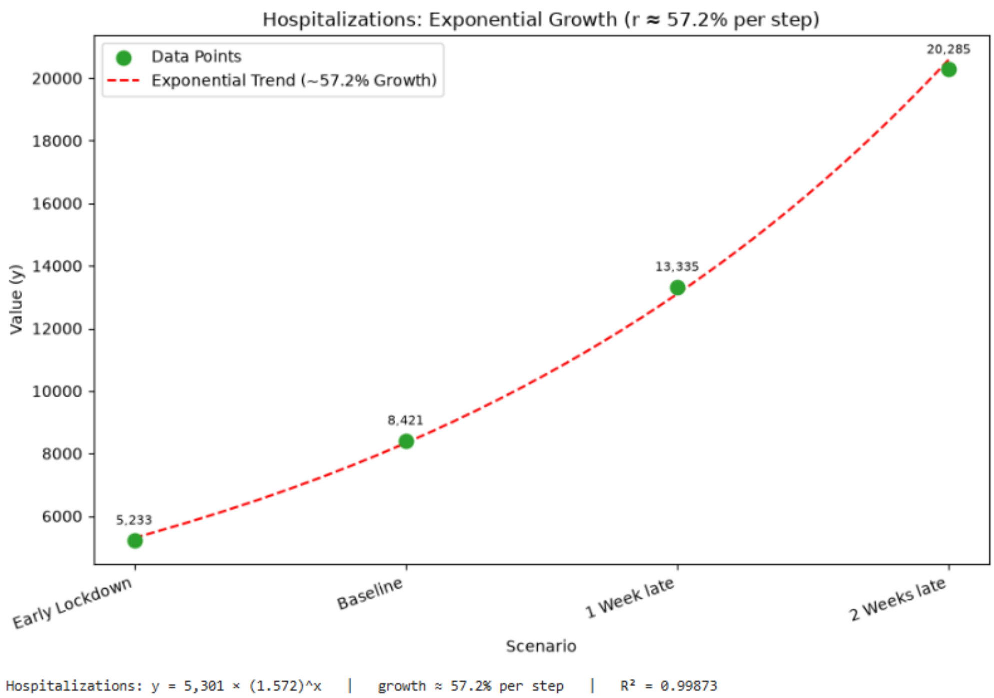

(f) A graph showing the relationship between lockdown time delay and number of hospital admissions

Figure 2: Infection, mortality, and hospitalisation outcomes for each lockdown timing scenario, and their relationship to delay length.

Non-pharmaceutical intervention scenarios

To examine the effect of different non-pharmaceutical interventions on infection, mortality, and hospitalisation outcomes, two additional scenarios were compared against the baseline: a scenario with increased hygiene and reduced contact, and a scenario with greater public awareness of isolation and symptoms. A combined intervention scenario is also included.

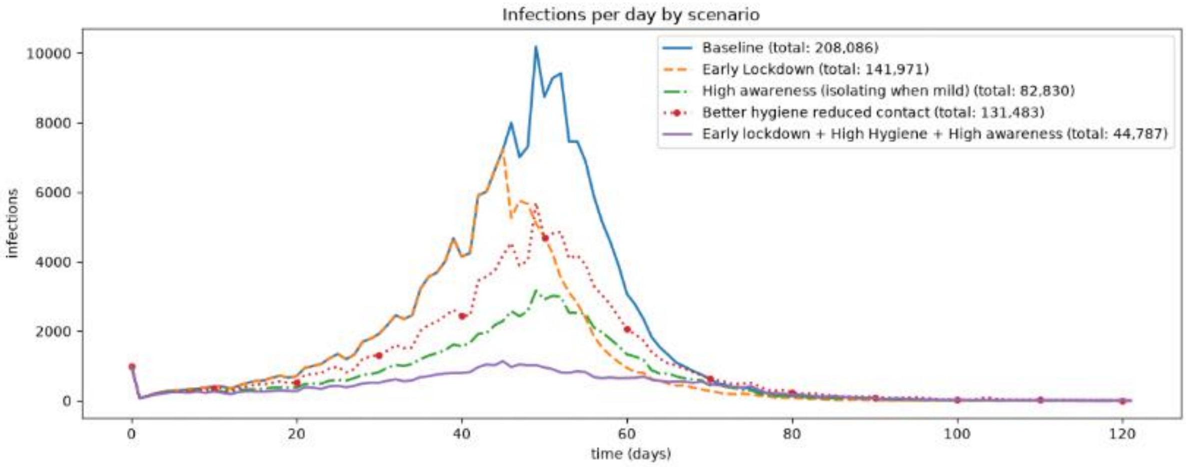

(a) A graph of infections per day (with total infections) for each NPI scenario

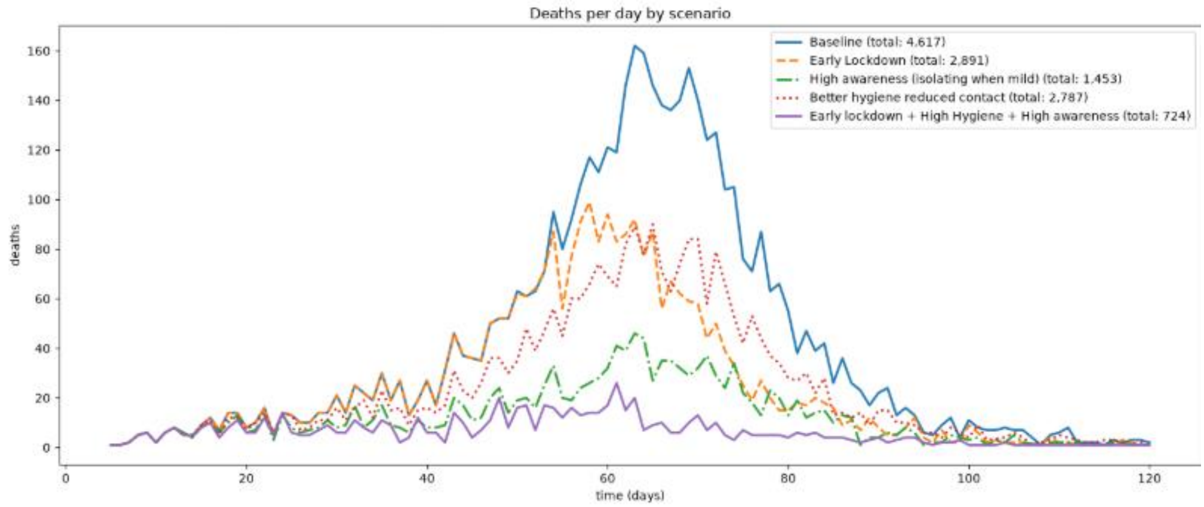

(b) A graph of deaths per day (with total deaths) for each NPI scenario

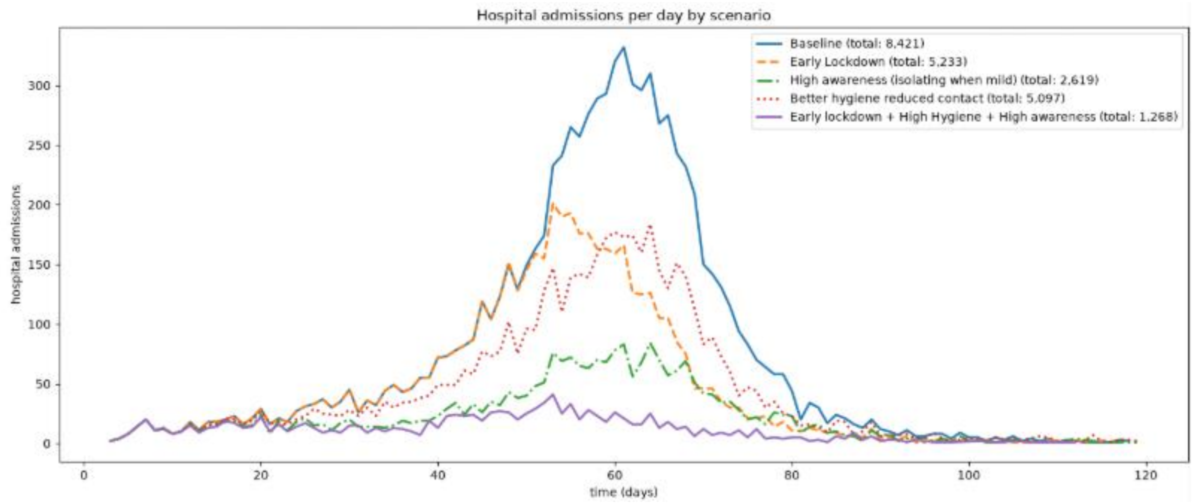

(c) A graph of hospital admission per day (with total hospital admissions) for each NPI scenario

Figure 3: Infections, deaths, and hospital admissions per day, and their totals, for each NPI scenario.

Additional Findings:

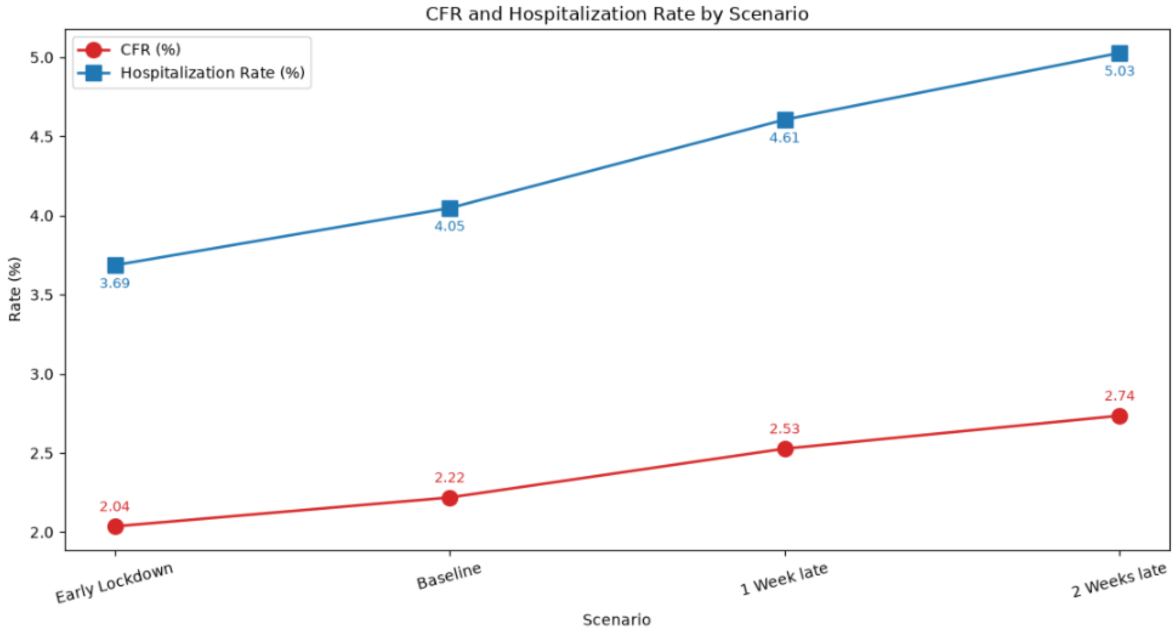

(a) A graph which shows the mortality rate and hospitalization rate of people who were infected from each timing of the lockdown

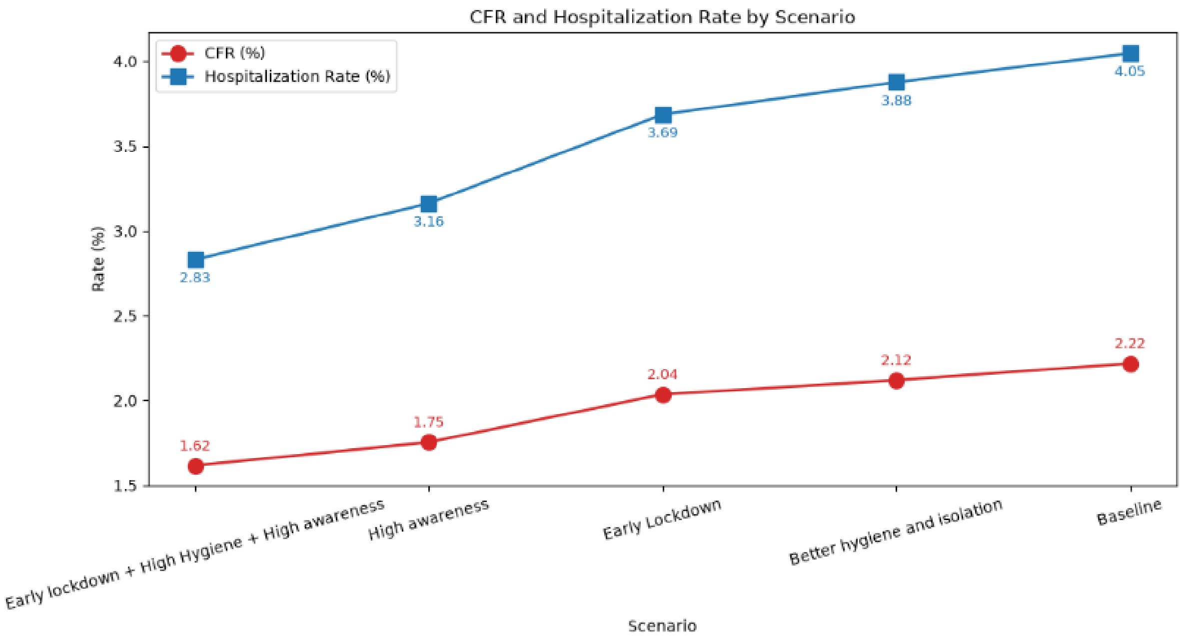

(b) A graph which shows the mortality rate and hospitalization rate of people who were infected from each NPI

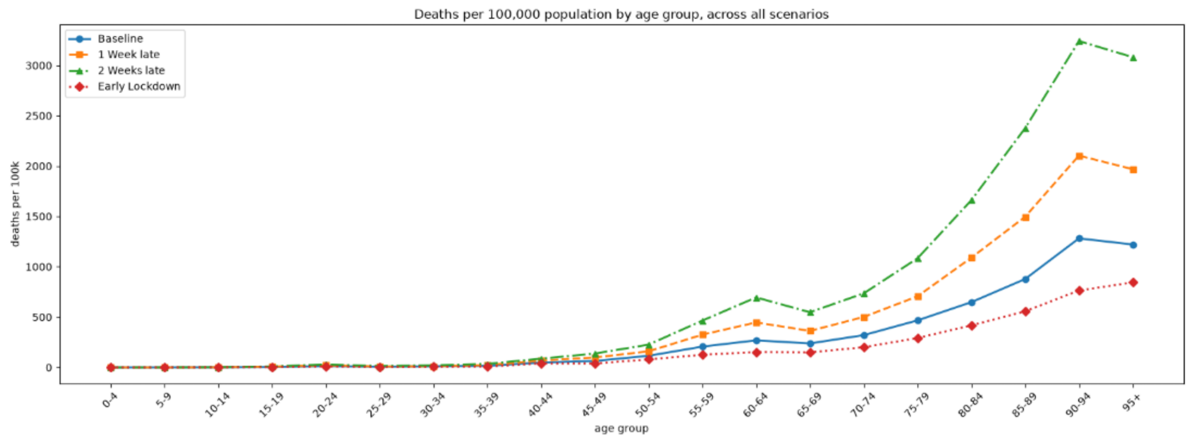

(c) A graph which shows the total death by age group per 100,000 people in each timing scenario

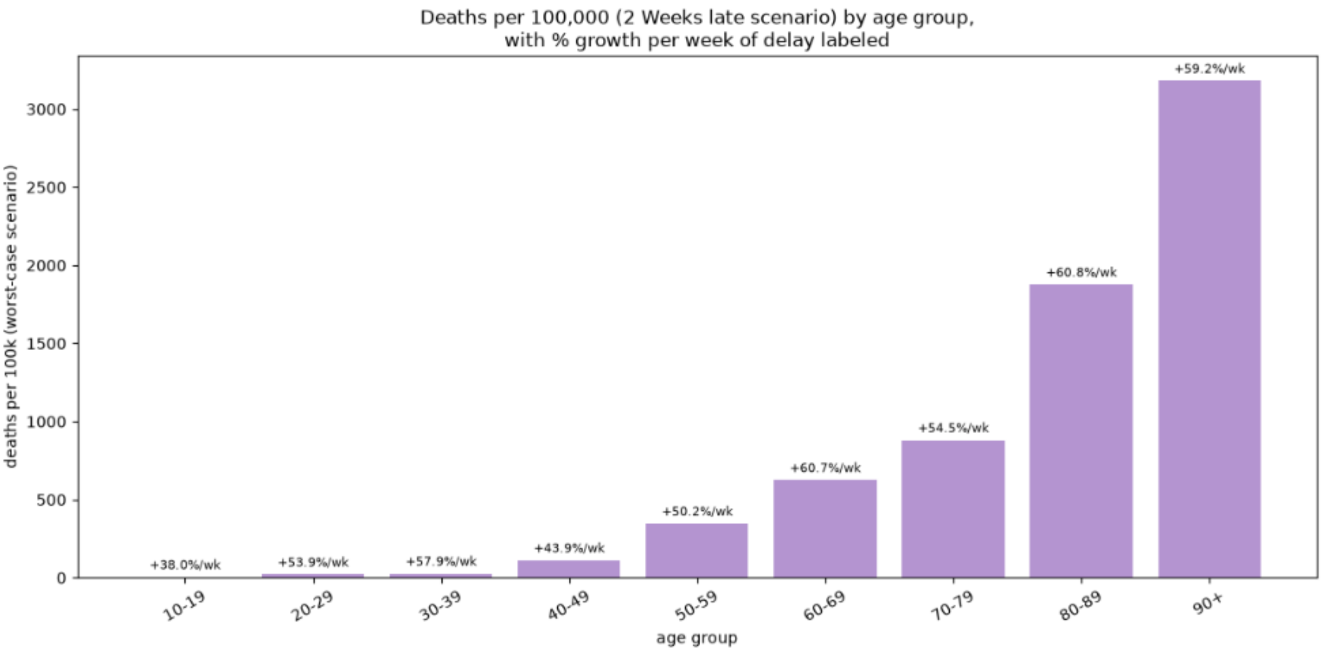

(d) A bar chart isolating the worst-case lockdown timing scenario from (d) for analysis for correlation between age group and deaths per 100,000 people

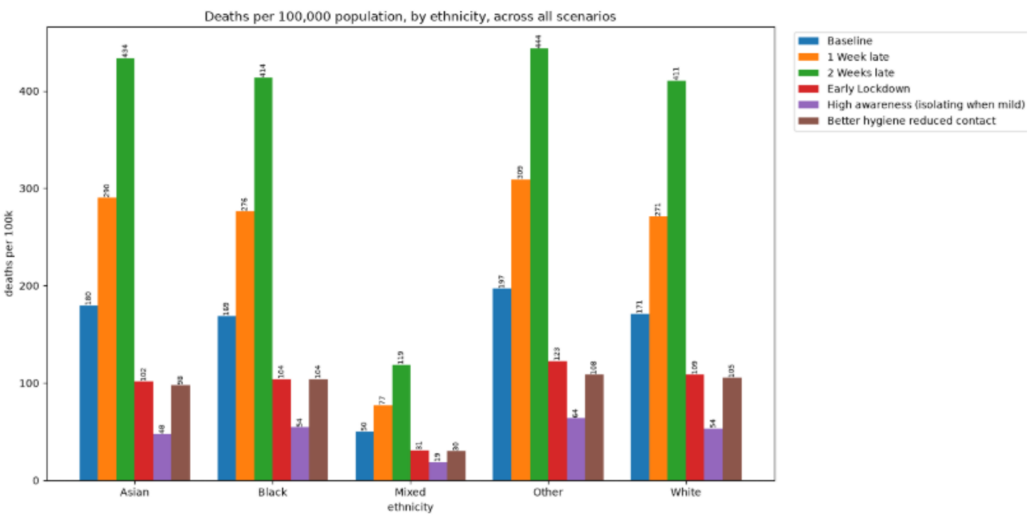

(e) A set of bar charts which show the deaths per 100,000 people of each ethnicity

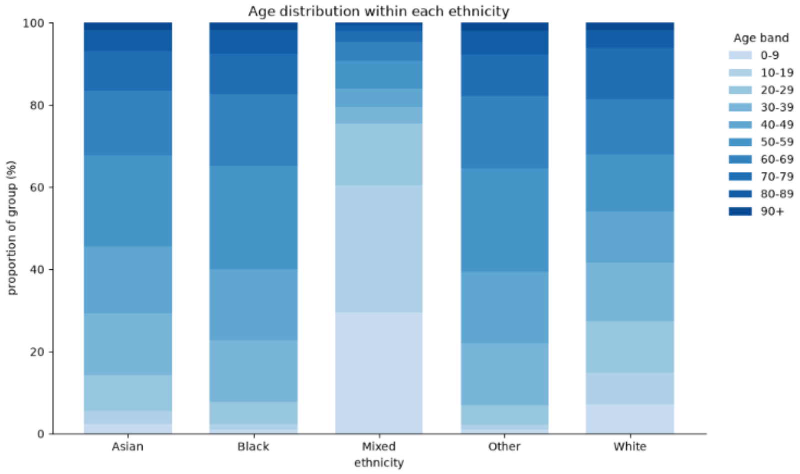

(f) A stacked- bar chart showing the age distribution within each ethnicity group

Figure 4: Severity-adjusted, age-stratified, and ethnicity-stratified outcomes across scenarios.

Discussion

In Figure 2(a), for each increasing delay in the introduction of the lockdown intervention, the peak of the infection curve is both higher and occurs 6-10 days later. This reflects a straightforward mechanism: with a delayed lockdown, more of the population is already infected by the time restrictions take effect, so the intervention has a smaller pool of susceptible transmission chains left to interrupt. The epidemic is consequently allowed to run further before being suppressed, producing the higher and later peaks seen in Figure 2(a) and the correspondingly larger total infection counts in Figure 2(c). This increase in total infections propagates directly into higher total deaths [Figure 2(b)] and hospital admissions [Figure 2(e)], since both scale with the number of people infected.

In Figure 2, the infection data 2(a), mortality data 2(b), and hospitalisation data 2(c) were used to produce 2(d)(e)(f), which plots total infections, deaths, and hospital admissions respectively against lockdown delay. In each case, an exponential model fit the data closely ($R^2 > 0.998$), indicating that outcomes do not simply scale linearly with delay but compound: each additional week of delay was associated with an approximate 41.4% increase in total infections, a 56.5% increase in total deaths, and a 57.2% increase in hospital admissions.

The disparity between these growth rates is itself a notable finding: deaths and hospitalisations grew faster than infections did, implying that delay does not only increase the *number* of people infected but also increases the *severity* of the epidemic's outcome per infection, an effect consistent with rising case fatality rates and hospitalisation rates under later-lockdown scenarios [Figure 4(a)]. Practically, this means a two-week delay increased total deaths by more than 150% (over double) the baseline figure — a substantial cost attributable to timing alone, holding disease severity otherwise constant. Given that mortality risk in this model is strongly concentrated in older age groups [4(c)(d)], a delayed intervention of this kind would be expected to disproportionately affect elderly populations, consistent with the demographic pattern reported in Figure 4(c). While the percentage increase in mortality per week of lockdown delay was broadly similar across age groups (approximately 38–61%), the absolute risk level from which this growth compounds differs dramatically by age: deaths per 100,000 population in the oldest age band (90+) exceeded 3,100 under the most delayed scenario in Figure 4(d), compared to near-zero in the youngest working-age bands. Because delay's proportional cost is applied to a substantially larger baseline risk in older populations, the absolute burden of delayed intervention falls disproportionately on elderly people — not because they are differentially more sensitive to delay itself, but because they were already at far greater risk before any delay occurred.

Across all three outcome measures [Figure 3(a)(b)(c)], the introduction of a non-pharmaceutical intervention produced a substantial reduction relative to the baseline scenario, though the magnitude of this reduction varied considerably by intervention type. Shifting the lockdown earlier (Early Lockdown) reduced total infections by 31.8% (208,086 to 141,971), total deaths by 37.4% (4,617 to 2,891), and total hospital admissions by 37.9% (8,421 to 5,233) relative to baseline. Reducing contact rates and improving hygiene compliance (washing hands more regularly and mask wearing) produced a broadly comparable reduction in deaths (39.6%, to 2,787) and hospital admissions (39.5%, to 5,097), but achieved a larger reduction in total infections (36.8%, to 131,483) than the timing-based intervention. The most substantial effect was observed under the high-awareness scenario, in which 90% of individuals self-isolated upon experiencing even mild symptoms, compared to the baseline 70%. This reduced total infections by 60.2% (to 82,830), total deaths by 68.5% (to 1,453), and total hospital admissions by 68.9% (to 2,619), roughly double the reduction achieved by either of the other two interventions on every outcome measure.

This ordering — high awareness outperforming improved hygiene, which in turn outperformed early lockdown timing — held consistently across all three Mioutcome metrics, indicating that the relative ranking of intervention effectiveness in this model is robust to the choice of outcome measure rather than being an artefact particular to any single metric. Notably, the margin between interventions was not uniform across outcomes: early lockdown and improved hygiene differed by only 2.1–2.2% in their effect on deaths and hospital admissions, yet differed by 5.0% in their effect on total infections, suggesting that sustained behavioural interventions confer a disproportionate benefit on transmission reduction specifically, relative to their benefit on severe outcomes, when compared with a purely timing-based intervention of similar overall stringency.

The severity-adjusted outcome measures [Figure 4(b)] complicate this ranking rather than simply reinforcing it. High awareness produced the lowest hospitalisation rate (0.0316) and lowest case fatality ratio (0.0175) of all scenarios, consistent with its dominance on every other outcome measure. However, Early Lockdown and Better hygiene reverse their relative positions on these severity-conditional measures compared with their positions on total burden: Better hygiene reduced total infections, deaths, and hospitalisations by a larger margin than Early Lockdown, yet Early Lockdown produced a marginally lower hospitalisation rate (0.0369 vs. 0.0388) and case fatality ratio (0.0204 vs. 0.0212) than better hygiene. This suggests that, conditional on infection occurring, timing-based suppression and sustained behavioural suppression may affect the probability of progression to severe outcomes somewhat differently than they affect the total number of infections prevented. However, with only a single simulation run per scenario this difference is modest enough that it should be treated as a preliminary observation rather than a robust finding without repeated runs to confirm it.

The mortality per 100,000 people by ethnicity [Figure 4(e)] provides a more accurate picture from raw death counts, which are dominated by the White population's larger share of the total simulated population rather than by relative risk. In the baseline scenario, Other had the highest rate (196.9 per 100,000), followed closely by Asian (179.7), White (171.2), and Black (168.9); The mixed ethnicity sat much lower at 50.4, roughly a quarter to a third of every other group. This ordering was stable across all three interventions — the mixed group remained the lowest group throughout (30.9, 18.9, and 29.9 under early lockdown, high awareness, and improved hygiene respectively), and Other remained the highest or near-highest — indicating a structural rather than scenario-specific effect.

The age distribution data [Figure 4(f)] substantiate this directly rather than leaving it as speculation. Mixed has a markedly younger population profile than every other group: approximately 29% of the Mixed population falls in the 0–9 age band alone, and roughly 60% is under 30, compared with a much flatter, older-skewing distribution in Asian, Black, Other, and White, each of which has only 3–6% in the 0–9 band. Since COVID-19 mortality risk is strongly age-graded, this age skew is a highly plausible primary driver of mixed ethnicities' substantially lower mortality rate, rather than reflecting anything mechanistically distinct about ethnicity itself. This should temper how causally the ethnicity-mortality relationship is framed in the dissertation: the appropriate claim is that mortality varies by ethnicity group in this model, with age structure as a likely major contributor, not that ethnicity independently drives risk.

Intervention effects were also not proportionally uniform: Asian consistently saw the largest relative risk reduction (43.5% under early lockdown, 73.5% under high awareness, 45.5% under improved hygiene), moving it from the second-highest baseline rate to the second-lowest under every intervention — a reordering no other group underwent. Absolute disparity between groups also narrowed substantially with intervention effectiveness (a range of 146.5 per 100,000 at baseline, falling to 45.0 under high awareness), though the relative gap between the other ethnicities and the mixed ethnicity narrowed far less (roughly 3.9x to 3.4x).

As expected, stacking all three interventions together produces by far the largest effect in the whole comparison [Figure 3(a)(b)(c)]. Total infections reduced by 78.5%, deaths by 84.3%, and hospital admissions by 84.9% — roughly 1.2 to 1.4 times better than even the best single intervention (High awareness alone, which achieved 60.2%/68.5%/68.9%

respectively). Comparing the observed combined effect against a simple multiplicative-independence benchmark, however, suggests the three interventions were not fully additive: the actual reduction in total case counts was somewhat smaller than would be predicted if each intervention reduced the remaining risk pool independently of the others, indicating a degree of overlap between the transmission pathways each intervention addresses.

Beyond the outcomes captured directly in this model, it is worth noting that a lower death toll does not equate to less suffering, nor does it capture the full distribution of harm a policy produces. COVID-19 itself generates morbidity independent of mortality: an infection avoided is not merely a death avoided, but also potential long COVID, hospitalisation, and lasting organ or cognitive impairment in survivors (Aiyegbusi et al., 2021) — outcomes this simulation's framework, which focuses on infection and death events, does not directly quantify. Equally, NPIs are not costless interventions applied against an otherwise-neutral baseline. Lockdown-style measures generate their own morbidity: delayed diagnosis and treatment for non-COVID conditions (Maringe et al., 2020), deteriorating mental health from isolation and economic insecurity (Pierce et al., 2020), and disrupted education — the Education Endowment Foundation (2022) documents substantial learning loss and widening educational inequality among affected children specifically, rather than developmental harm more broadly. An earlier or longer lockdown that averts COVID deaths may simultaneously extend the duration of these secondary harms — a trade-off no single mortality figure can represent.

These ethnicity-stratified findings also sit within a wider pattern of unevenly distributed pandemic harm. Neither the pandemic's harms nor the interventions' costs were distributed evenly across the population, and an aggregate outcome measure such as mortality systematically obscures that unevenness. UK-wide surveillance data showed substantially elevated COVID-19 death rates among Black and South Asian populations (up to 50% higher than White ethnic groups) and among those with existing comorbidities (Public Health England, 2020); mortality was separately shown to be significantly elevated in more deprived local areas of England — a national rather than Wales-specific pattern, and evidence of deprivation-linked rather than ethnicity-specific risk (Office for National Statistics, 2020a). Lockdown measures carried their own unequal burden. Key and frontline workers — disproportionately lower-paid, and disproportionately from minority ethnic backgrounds — could not shelter at home in the way office-based workers could, and so bore higher occupational exposure risk while simultaneously facing greater economic precarity (Office for National Statistics, 2020b). School closures fell hardest on children in low-income households lacking private space, devices, or reliable internet for remote learning, widening pre-existing educational inequality (Bayrakdar and Guveli, 2023). Care home residents faced some of the strictest and most prolonged isolation of any group, with attendant harms to mental health and wellbeing that a mortality-only lens is poorly suited to register (Sims et al., 2024). Women and children in unsafe households faced elevated risk from prolonged confinement with an abuser, a harm invisible to any COVID-specific outcome measure entirely (Office for National Statistics, 2020c).

Conclusion and limitations

This study uses deaths averted as its primary outcome measure of success. Mortality is quantitative and objectively countable, comparatively well-recorded even amid pandemic reporting strain, and carries an unambiguous moral weight which makes it a natural headline metric. The finding explored here — that an earlier first lockdown could plausibly have averted a substantial number of deaths, especially in the elder population— sits within this tradition (UK Covid-19 Inquiry, 2025). However, treating aggregate mortality reduction as a sufficient measure of policy success risks obscuring both the genuine complexity of what NPIs cost and achieve, and who actually bears those costs and benefits. These costs, and their uneven distribution across the population, are discussed in detail in the Results and Analysis section above.

This suggests that a more complete evaluation framework would move toward composite and disaggregated measures — quality-adjusted life years (QALYs) or disability-adjusted life years (DALYs) (National Institute for Health and Care Excellence, 2013) weighed against distributional analysis by age, ethnicity, deprivation, and occupation — rather than a single aggregate figure that implicitly treats a death or a lost year of healthy life as equivalent regardless of whose it was. Such approaches are not without their own contested assumptions (Arnesen and Nord, 1999), but they at least make trade-offs and inequalities explicit rather than erasing them by omission.

The ethnicity-stratified findings are further limited by the synthetic nature of the modelled population: because ethnicity, age, household structure, and geographic or occupational placement are jointly generated within MAY's population-generation process, the extent to which the observed mortality disparities reflect ethnicity-linked risk specifically, as opposed to structural factors that happen to correlate with ethnicity in this synthetic population, cannot be established from the aggregate outputs used here and would require either stratified modelling that holds age constant across ethnicity groups, or documentation of how MAY assigns these attributes relative to one another. Finally, as this study was only able to model the second national lockdown rather than the first, any extrapolations of these findings to claims about the earlier lockdown's effectiveness — including the counterfactual figures cited in the UK Covid-19 Inquiry — should be made with caution, given that the two epidemic waves differed in shape and the population's baseline behaviour and immunity differed substantially between them.

An extension of this work could include repeating this analysis at the scale of England and Wales combined – given the 2021 census data – which would allow for a greater population number and therefore a more robust set of results. Additionally, each scenario in this analysis was generated from a single stochastic simulation run rather than an average across repeated realisations; repeating each scenario multiple times would be a natural extension of this work.

Acknowledgments

I would like to thank my supervisor, Dr Gavin Woolman, for being incredibly helpful and for giving up a lot of his time, both within the research period and before submission of my application. I would also like to thank him for his guidance on using the computation pipeline of MAY, JUNE2 and May Visualiser, and helping me learn and understand how to use them, without which this research wouldn't have been possible. Thank you to the Laidlaw Foundation for funding this research project.

References

Aiyegbusi, O.L. et al. (2021) 'Symptoms, complications and management of long COVID: a review', Journal of the Royal Society of Medicine, 114(9), pp. 428–442. Available at: https://doi.org/10.1177/01410768211032850 (Accessed: 23 August 2026).

Arnesen, T. and Nord, E. (1999) 'The value of DALY life: problems with ethics and validity of disability adjusted life years', BMJ, 319(7222), pp. 1423–1425. Available at: https://doi.org/10.1136/bmj.319.7222.1423 (Accessed: 23 August 2026).

Aylett-Bullock, J. et al. (2021) 'June: open-source individual-based epidemiology simulation', Royal Society Open Science, 8(7), 210506. Available at: https://doi.org/10.1098/rsos.210506 (Accessed: 23 August 2026)

Buffel, T., Yarker, S., Phillipson, C., Lang, L., Lewis, C., Doran, P. and Goff, M. (2023) 'Locked down by inequality: older people and the COVID-19 pandemic', Urban Studies, 60(8), pp. 1465–1482. Available at: https://doi.org/10.1177/00420980211041018 (Accessed: 30 August 2026).

Education Endowment Foundation (2022) The impact of COVID-19 on learning: a review of the evidence. Available at: https://educationendowmentfoundation.org.uk/guidance-for-teachers/covid-19-resources/best-evidence-on-impact-of-covid-19-on-pupil-attainment (Accessed: 23 August 2026).

Findlater, A. and Bogoch, I.I. (2018) 'Human mobility and the global spread of infectious diseases: a focus on air travel', Trends in Parasitology, 34(9), pp. 772–783. Available at: https://doi.org/10.1016/j.pt.2018.07.004 (Accessed: 23 August 2026).

Gao, Z., Gu, Z. and Yang, L. (2023) 'Effects of community connectivity on the spreading process of epidemics', Entropy, 25(6), 849. Available at: https://doi.org/10.3390/e25060849 (Accessed: 23 August 2026).

gavdoubleu (2026a) MAY World Visualiser. GitHub repository. Available at: https://github.com/gavdoubleu/may_world_visualiser (Accessed: 23 August 2026).

Hunter, E., Mac Namee, B. and Kelleher, J.D. (2017) 'A taxonomy for agent-based models in human infectious disease epidemiology', Journal of Artificial Societies and Social Simulation, 20(3), 2. Available at: https://doi.org/10.18564/jasss.3414 (Accessed: 23 August 2026).

IDAS-Durham (2026a) MAY. GitHub repository. Available at: https://github.com/IDAS-Durham/MAY (Accessed: 23 August 2026).

IDAS-Durham (2026b) JUNE2. GitHub repository. Available at: https://github.com/IDAS-Durham/JUNE2 (Accessed: 23 August 2026).

IPBES (2020) Workshop report on biodiversity and pandemics. Bonn: IPBES secretariat. Available at: https://doi.org/10.5281/zenodo.4147317 (Accessed: 23 August 2026).

Maringe, C. et al. (2020) 'The impact of the COVID-19 pandemic on cancer deaths due to delays in diagnosis in England, UK: a national, population-based, modelling study', The Lancet Oncology, 21(8), pp. 1023–1034. Available at: https://doi.org/10.1016/S1470-2045(20)30388-0 (Accessed: 23 August 2026).

Met Office (2022) A milestone in UK climate history. Available at: https://www.metoffice.gov.uk/about-us/news-and-media/media-centre/weather-and-climate-news/2022/july-heat-review (Accessed: 23 August 2026).

National Institute for Health and Care Excellence (2013) Guide to the methods of technology appraisal 2013. London: NICE. Available at: https://www.nice.org.uk/process/pmg9/resources/guide-to-the-methods-of-technology-appraisal-2013-pdf-2007975843781 (Accessed: 23 August 2026).

NHS England (2026) COVID-19 hospital activity. Available at: https://www.england.nhs.uk/statistics/statistical-work-areas/covid-19-hospital-activity/ (Accessed: 23 August 2026).

Office for National Statistics (2020a) Deaths involving COVID-19 by local area and socioeconomic deprivation: deaths occurring between 1 March and 31 July 2020. Available at: https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/deaths/bulletins/deathsinvolvingcovid19bylocalareasanddeprivation/deathsoccurringbetween1marchand31july2020/pdf (Accessed: 23 August 2026).

Office for National Statistics (2020b) Coronavirus and key workers in the UK. Available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/articles/coronavirusandkeyworkersintheuk/2020-05-15 (Accessed: 23 August 2026).

Office for National Statistics (2020c) Domestic abuse during the coronavirus (COVID-19) pandemic, England and Wales: November 2020. Available at: https://www.ons.gov.uk/peoplepopulationandcommunity/crimeandjustice/articles/domesticabuseduringthecoronaviruscovid19pandemicenglandandwales/november2020 (Accessed: 23 August 2026).

Pierce, M. et al. (2020) 'Mental health before and during the COVID-19 pandemic: a longitudinal probability sample survey of the UK population', The Lancet Psychiatry, 7(10), pp. 883–892. Available at: https://doi.org/10.1016/S2215-0366(20)30308-4 (Accessed: 23 August 2026).

Public Health England (2020) Disparities in the risk and outcomes of COVID-19. Available at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/908434/Disparities_in_the_risk_and_outcomes_of_COVID_August_2020_update.pdf (Accessed: 23 August 2026).

Smyth, N. et al. (2024) 'People from ethnic minorities seeking help for long COVID: a qualitative study', British Journal of General Practice, 74(749), pp. e814–e822. Available at: https://doi.org/10.3399/BJGP.2023.0631 (Accessed: 23 August 2026).

UK Covid-19 Inquiry (2024) Module 1 report: the resilience and preparedness of the United Kingdom. Available at: https://covid19.public-inquiry.uk/reports/module-1-report-the-resilience-and-preparedness-of-the-united-kingdom/ (Accessed: 23 August 2026).

UK Covid-19 Inquiry (2025) Modules 2, 2A, 2B, 2C report: core decision-making and political governance. Available at: https://covid19.public-inquiry.uk/reports/modules-2-2a-2b-2c-core-decision-making-and-political-governance-volume-i/ (Accessed: 23 August 2026).

World Health Organization (2026) WHO COVID-19 dashboard. Available at: https://data.who.int/dashboards/covid19/deaths (Accessed: 23 August 2026).