JASSS
THE JOURNAL OF ARTIFICIAL SOCIETIES AND SOCIAL SIMULATION

Essential Workers at Risk: An Agent-Based Model (SAFE-ABM) with Uncertainty Quantification

Elizabeth B. Amona[1,2], **Indranil Sahoo**[2], **Ya Su**[2], **Edward L. Boone**[2], **Gwendoline Nelis**[3], **Ryad Ghanam**[4]

[1] *UofL Health Brown Cancer Center, University of Louisville, 529 South Jackson Street, Louisville, KY 40202, USA*
[2] *Department of Statistical Sciences and Operations Research, Virginia Commonwealth University, Harris Hall, 1015 Floyd Avenue, Box 843083, Richmond, VA 23284, USA*
[3] *Faculty of Public Health, Université catholique de Louvain, Clos Chapelle-aux-Champs 30, 1200 Woluwe-Saint-Lambert, Belgium*
[4] *Department of Liberal Arts and Sciences, Virginia Commonwealth University in Qatar, P.O. Box 8095, Doha, Qatar*
Correspondence should be addressed to *elizabeth.amona@louisville.edu*

Journal of Artificial Societies and Social Simulation 29(4) 7, 2026
Doi: 10.18564/jasss.6276 Url: http://jasss.soc.surrey.ac.uk/29/4/7.html

Received: 26-01-2026 Accepted: 31-08-2026 Published: 31-10-2026

Abstract: Essential workers face elevated infection risks due to their critical roles during pandemics, and protecting them remains a significant challenge for public health planning. This study develops SAFE-ABM, a simulation-based framework using Agent-Based Modeling (ABM), to evaluate targeted intervention strategies, explicitly capturing structured interactions across families, workplaces, and schools. We simulate key scenarios such as unrestricted movement, school closures, mobility restrictions specific to essential workers, and workforce rotation with and without quarantine, to assess their impact on disease transmission dynamics. To ensure robust uncertainty assessment, we integrate a novel Bayesian Uncertainty Quantification (UQ) framework, systematically capturing variability in transmission rates, recovery times, and mortality estimates. Our comparative analysis demonstrates that workforce rotation substantially reduces transmission among essential workers, even without quarantine, while workforce rotation combined with quarantine enforcement provides additional reduction in cumulative deaths and infection burden. Unlike less targeted interventions, the rotational workforce strategies preserve a larger portion of the susceptible population, resulting in a more controlled and sustainable epidemic trajectory. These findings offer critical insights for optimizing intervention strategies that mitigate disease spread while maintaining essential societal functions.

Keywords: Agent-Based Model, Uncertainty Quantification, Localized Interventions Analysis, Bayesian Statistics, Epidemiology

Introduction

1.1 Essential workers are individuals who maintain critical societal services despite widespread restrictions during pandemics. This group includes healthcare practitioners such as doctors, nurses, and hospital cleaning staff, as well as food vendors, delivery personnel, and ride-share drivers. During a pandemic, this group faces significantly elevated infection risks due to continued interactions in high-contact environments (Milligan et al. 2021). Infectious disease outbreaks frequently prompt policymakers to implement broad population-wide interventions, such as school closures, movement restrictions, and social interaction limitations, aimed at curbing general transmission (Chung et al. 2023; El Jaouhari et al. 2021). Although previous studies have shown these interventions, including lockdown and travel restrictions, effectively reduce infection surges when enacted early (El Jaouhari et al. 2021; Amona et al. 2022), they are not beneficial for essential workers. This oversight leaves

essential workers without specifically tailored protective measures. Despite their crucial societal role, there remains a significant gap in comprehensive modeling that explicitly assesses essential workers' contributions to disease transmission and evaluates targeted protective interventions (Milligan et al. 2021). Addressing this gap is essential for ensuring equitable and effective pandemic responses.

1.2 This situation raises several critical research questions: Does restricting the general population while allowing only essential workers to operate effectively reduce overall transmission, or does it instead concentrate risk within this crucial workforce? Additionally, how do specific intervention strategies such as school closures or mobility restrictions targeting non-essential populations, alter overall transmission dynamics? Could these policies unintentionally elevate transmission risks within workplaces and families? Clearly understanding these trade-offs is vital for designing interventions that reduce disease spread without disrupting essential services.

1.3 One proposed strategy for reducing workplace exposure involves splitting essential workers into rotating subgroups to minimize continuous workplace exposure (Milligan et al. 2021). While rotation may lower transmission within workplaces, it could inadvertently introduce unintended consequences, such as increased workloads on active workers or disruptions in critical service continuity (van Zoonen & Ter Hoeven 2022). For instance, healthcare settings have experienced workforce shortages from infection-driven absenteeism, exacerbating facility overcrowding, elevating transmission risks, and placing additional strain on healthcare systems. These complexities highlight the need for a systematic framework for evaluating intervention trade-offs, ensuring policies effectively balance infection control with operational feasibility (Phadke et al. 2021; McNeill 2022).

1.4 To systematically explore these trade-offs, this study employs a Structured Agent-based Framework for Essential workers (SAFE-ABM), a simulation-based platform using Agent-Based Modeling (ABM) (Bonabeau 2002; Hinch et al. 2021; Niazi & Hussain 2011). ABMs enable simulations of structured, localized interactions within workplaces, schools, and families, providing insights into how targeted interventions affect disease transmission at both individual and community levels. By explicitly modeling structured interactions, we evaluate various scenarios, including unrestricted movement, school closures, mobility restrictions tailored for essential workers, and workforce rotation schemes. Unlike traditional compartmental models that rely on aggregated population-level assumptions, ABMs capture detailed individual interactions within structured environments, enabling a granular assessment of intervention effectiveness.

1.5 A key innovation in our study is the integration of Bayesian Uncertainty Quantification (UQ) within the ABM. Although UQ has previously been explored in epidemiological modeling contexts, its explicit integration within ABMs, particularly utilizing Bayesian inference to quantify transmission uncertainties, remains underdeveloped. Our novel approach systematically evaluates variability in transmission rates, recovery periods, and mortality, ensuring that small fluctuations in parameter estimates do not disproportionately influence intervention conclusions. By employing Bayesian prior distributions informed by domain knowledge, we probabilistically explore epidemic outcomes, while accounting for stochastic agent interactions and epidemiological uncertainty.

1.6 ABMs have proven versatile in infectious disease modeling, enabling flexible exploration of transmission dynamics and intervention effectiveness (Kerr et al. 2021). Prominent ABM frameworks, including Covasim (Kerr et al. 2021), OpenABM-Covid19 (Hinch et al. 2021), and FRED (Grefenstette et al. 2013), have evaluated the effects of social distancing, vaccination, testing, and contact tracing on disease spread. Specifically, Covasim (COVID-19 Agent-based Simulator) is a flexible ABM explicitly designed for COVID-19, capturing detailed demographic structures, comprehensive transmission networks across households, schools, workplaces, and care facilities, and modeling age-specific disease outcomes. Covasim facilitates extensive evaluations of various policy scenarios including physical distancing measures, contact tracing, quarantine protocols, and vaccination strategies, and has been widely used to inform public health decisions globally. Similarly, OpenABM-Covid19 emphasizes computational performance, efficiently modeling structured interactions across households, schools, workplaces, and community settings. Its scalability enables timely large-scale scenario analyses, supporting rapid policy assessments. FRED (Framework for Reconstructing Epidemiological Dynamics), originally developed for influenza modeling and later adapted for COVID-19, integrates realistic synthetic populations with structured interactions across societal layers. It explicitly assesses the impacts of interventions such as vaccination, isolation, and quarantine at local and national scales, providing insights into epidemic dynamics. ABMs like EPINEST (Pinotti et al. 2024) also offer specialized frameworks for intervention assessment, while CityCOVID (Ozik et al. 2021; Hack & Papka 2020), explicitly models urban-scale transmission dynamics, employing spatially-resolved contact networks informed by local demographic and behavioral data. ABMs have been successfully applied beyond COVID-19 to influenza (Arduin et al. 2017), dengue (Miksch et al. 2019; Mahmood et al. 2020; Deng et al. 2008), and Ebola (Chen et al. 2021), demonstrating their adaptability across epidemiological contexts. These models incorporate age-stratified disease progression, spatial transmission networks, and heterogeneous in-

tervention effects, providing comprehensive insights into both pharmaceutical and non-pharmaceutical control measures (Lorig et al. 2021).

1.7 Despite these significant advances, most existing ABMs simplify workforce heterogeneity, assume uniform social interactions, or generalize intervention effects, thus limiting their effectiveness in evaluating strategies tailored specifically for essential workers. Our model directly addresses these limitations by explicitly incorporating structured movement patterns, targeted interventions, and workforce splitting. Our framework modifies the underlying model structure at each intervention point, reconfiguring agent interactions and movement dynamics to reflect the real-world implementation of policy changes over time. This dynamic restructuring allows us to capture the effects of interventions as they occur, rather than inferring them solely through parameter adjustments, as is typical in many existing ABMs (Dyer et al. 2023). Additionally, our framework integrates detailed family structures to accurately capture intra-family transmission dynamics which allows for precise assessments of how interventions affect transmission within and across structured social environments. Our explicit evaluation of workforce rotation strategies among essential workers further distinguishes this study from the previous ABMs. Essential workers are split into mutually exclusive groups such that while one group is active, the other remains at home, thus modeling a structure that has been proposed in policy discussions but rarely implemented within simulation frameworks.

1.8 By explicitly modeling heterogeneous social interactions and incorporating essential worker risk, workforce rotation, and family-level dynamics, our study significantly advances the realism and applicability of ABMs for pandemic response planning. Integrating Bayesian UQ within this essential-worker-focused framework ensures robust intervention evaluations that rigorously account for both stochastic variability and epidemiological uncertainty. These innovations provide crucial insights for optimizing targeted pandemic response strategies, safeguarding essential worker populations, and maintaining societal resilience. The study also explicitly focuses on public health implications rather than economic considerations.

1.9 The remainder of this paper is organized as follows. Sections 2, 3 and 4 present our the overall statistical methodology, including an overview of the ABM framework, mathematical formulation of the proposed model and key assumptions, simulation scenarios, model algorithm, and uncertainty quantification and validation methods. Section 5 discusses our simulation results, and Section 6 summarizes key findings and highlights future research directions.

SAFE-ABM Modeling Framework

2.1 The agent-based model developed in this study, SAFE-ABM (Structured Agent-Based Framework for Essential Workers), simulates disease transmission within structured social environments, explicitly capturing heterogeneous interactions across families, workplaces, and schools. Agents in SAFE-ABM are assigned demographic and behavioral attributes that determine their interactions, movement patterns, and contact structures. By explicitly incorporating individual-level variability, the model offers a detailed representation of the social interactions underlying infection spread.

Description of the agent-based model

2.2 Figure 1 presents the SAFE-ABM framework used in this study to simulate transmission dynamics across structured environments, including families, schools, and workplaces. The model represents a heterogeneous population, where agents interact in localized settings, mimicking real-world contact structures. Each agent is characterized by age, occupation (essential or non-essential worker), and movement behavior, all of which shape their interaction patterns and disease exposure.

2.3 At the core of the model is the family unit, where agents are grouped into 2,000 families of varying sizes, forming the primary transmission network. Within families, members engage in repeated, prolonged interactions, making this a high-contact setting for disease spread. Beyond families, schools and workplaces also form structured transmission networks. Schools are stratified into elementary and high school cohorts, facilitating age-dependent mixing patterns. The frequency and intensity of these structured peer interactions vary between elementary and high school students. Workplaces are differentiated by occupation, distinguishing between essential and non-essential workers. Essential workers, including healthcare professionals, delivery personnel, and public service employees, remain active during intervention periods, maintaining workforce participation even under movement restrictions. This sustained occupational exposure contributes to workplace-driven

transmission chains. In contrast, non-essential workers experience reduced external exposure, modifying overall disease propagation patterns.

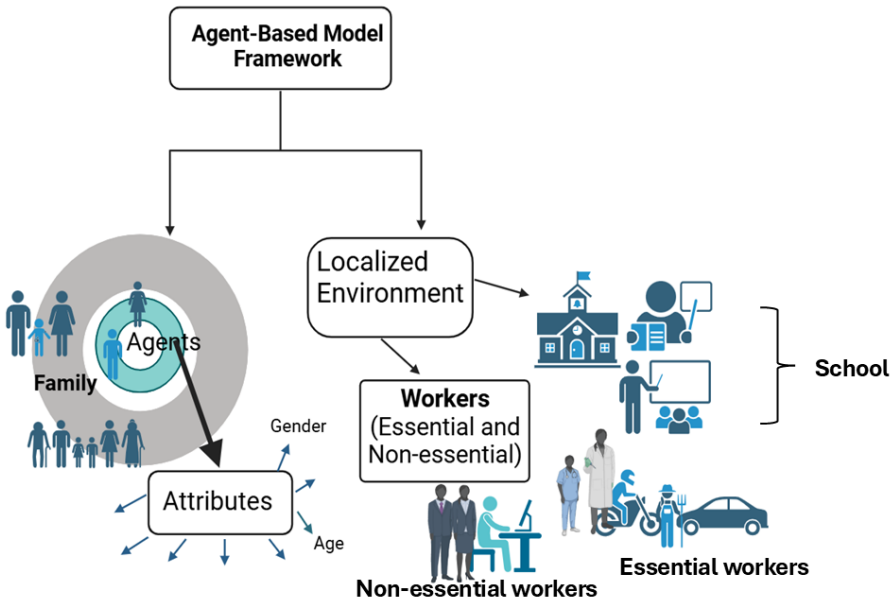


Figure 1: Agent-Based Modeling Framework illustrating agents, attributes, and localized environments such as families, schools (elementary and high school), and workplaces (essential and non-essential workers).

2.4 Each agent follows a set of predefined movement and contact rules which are used to calculate the probability of interactions within and across environments. Families represent high-contact networks with frequent interactions, schools facilitate structured peer interactions, and workplaces sustain consistent exposure among colleagues. Essential workers maintain continuous interactions even under restrictive measures, increasing exposure risks in both occupational and household settings. These interactions define the exposure risk landscape within the model, shaping infection pathways across the simulated population. Interventions in the model dynamically modify transmission pathways by adjusting movement dynamics and interaction frequencies. For instance, school closures shift student interactions from schools to families, potentially increasing intra-family transmission. Workforce modifications, such as split essential worker rotations, restructure workplace exposure by alternating work groups, thereby minimizing continuous high-risk interactions. Each intervention scenario captures the trade-offs between mobility restrictions and maintaining essential societal functions, offering insights into how public health measures influence both disease spread and economic stability.

Model structure and agent interactions

2.5 Agents in our model interact dynamically within predefined social structures, with movement patterns and contact networks evolving over time. The first 14 days of the simulation establish baseline interaction dynamics, during which individuals maintain normal mobility and interaction patterns without intervention. During this period, families serve as primary contact units, workplaces allow for occupational interactions, and schools facilitate peer-based mixing. Essential workers remain active throughout the simulation, continuously interacting in both workplace and household environments, while non-essential workers follow routine schedules until movement restrictions are applied.

2.6 The implemented interaction rules are summarized in Table 1. During the first 14 days, all scenarios follow the same baseline contact structure. After day 14, scenario-specific interventions modify the eligible contact environments while preserving the same stochastic disease-transition rules. Under the baseline dynamics, infected agents initiate general random contacts drawn from a Poisson distribution with mean 10, whereas exposed agents initiate general random contacts drawn from a Poisson distribution with mean 15. In addition to these general contacts, transmission can occur through structured household, school, and workplace networks.

Table 1: Implemented contact structure across intervention scenarios. All scenarios follow the baseline contact structure during days 1–14, while scenario-specific restrictions are applied after day 14.

Contact component	Baseline rule (days 1–14)	Scenario-specific rule (after day 14)
General random contacts	Infected agents initiate random contacts drawn from Poisson(10), while exposed agents initiate random contacts drawn from Poisson(15).	Maintained in Scenario 1. In intervention scenarios, random contacts are restricted to agents who remain eligible for external interaction under the scenario rules.
Household contacts	Family members are eligible contacts within each household.	Maintained in all scenarios because household transmission remains possible.
Elementary school contacts	Elementary students interact within assigned elementary school networks.	Maintained in Scenario 1 and removed in Scenarios 2–5 after day 14.
High school contacts	High school students interact within assigned high school networks.	Maintained in Scenario 1 and removed in Scenarios 2–5 after day 14.
Non-essential worker contacts	Non-essential workers interact within assigned workplace networks.	Maintained in Scenarios 1 and 2 and restricted in Scenarios 3–5 after day 14.
Essential worker contacts	Essential workers interact within assigned workplace networks.	Maintained in Scenarios 1–3. In Scenarios 4 and 5, contacts are restricted to the active rotating essential-worker subgroup.
Rotation and quarantine	Not applied during the baseline period.	Rotation is applied in Scenarios 4 and 5, where essential workers rotate on a 7-day schedule. Quarantine of symptomatic active workers is applied only in Scenario 4.

2.7 As interventions are introduced, interaction structures adapt accordingly. School closures shift student interactions from classroom settings to family units, altering exposure risks within households. Workforce modifications, such as the split essential worker rotation scheme, reduce occupational exposure by alternating worker schedules, thereby limiting prolonged, high-risk interactions. At each time step, disease transmission is governed by contact-based stochastic processes, where susceptible agents face infection risk based on interaction frequency and the infectious status of their contacts. Upon exposure, individuals enter an incubation period before transitioning to either asymptomatic recovery or symptomatic infection. Symptomatic cases progress to recovery or mortality based on predefined health transition probabilities. The model continuously tracks these transitions, ensuring a dynamic representation of how movement and contact patterns shape epidemic trajectories.

Mathematical framework and model assumptions

2.8 The mathematical framework of SAFE-ABM defines the rules governing disease progression and agent movement dynamics within structured environments. Agents transition between epidemiological states through stochastic processes, with infection, recovery, and mortality determined by predefined probability distributions. Movement patterns establish contact structures within families, workplaces, and schools, thereby shaping exposure risks. Figure 2 illustrates the epidemiological states and transition pathways implemented within SAFE-ABM, structured by a stochastic Susceptible, Exposed, Infected, Recovered, Quarantined, and Death (SEIRQD) framework.

2.9 The model assumes a closed population, with the total number of agents remaining unchanged throughout the simulation. Natural births, deaths, emigration, or immigration are not introduced during the simulation period. Agents begin in the susceptible (S) state, meaning they have not yet been exposed to the virus. Upon exposure, they transition to the exposed (asymptomatic) (E) state, undergoing an incubation period before either recovering without symptoms or developing symptomatic infection. Symptomatic infected (I) individuals may recover, succumb to the disease, or, in scenarios incorporating interventions, be quarantined (Q) to reduce further transmission. Recovered (R) agents are assumed immune, while deceased (D) agents remain in the simulation, ensuring the total population count remains constant.

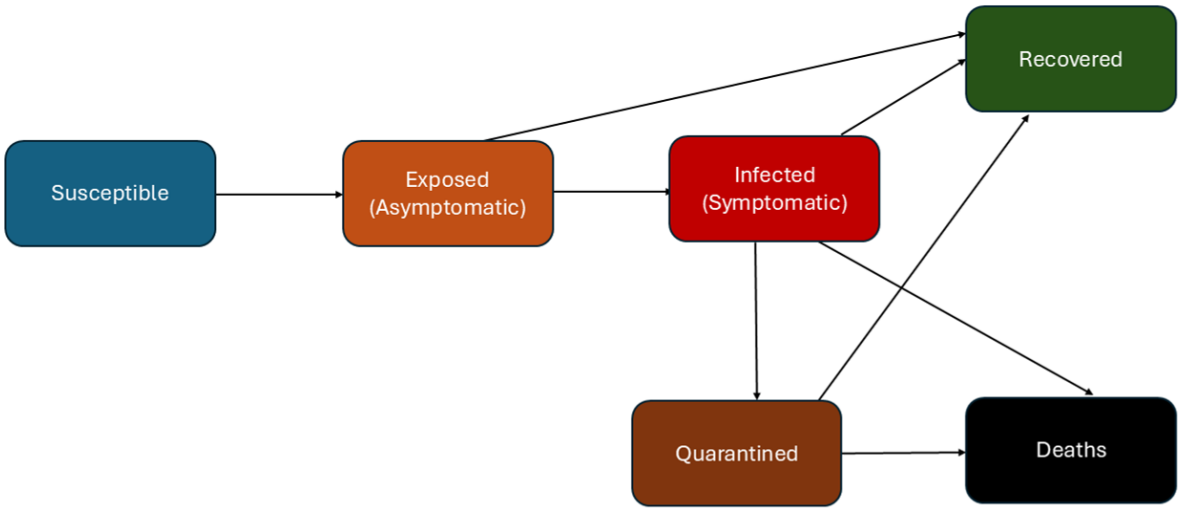


Figure 2: Epidemiological states and transition pathways in the agent-based model. The model follows a stochastic SEIRQD framework, with agents transitioning between states based on probabilistic rules.

Infection dynamics

2.10 Susceptible (S) agents contract the infection through contact-based transmission upon interaction with an infected individual. The probability of infection is defined as:

$$P_{\text{infect}} = 1 - e^{-\beta C I}$$

where $\beta \sim \mathcal{U}(0.05, 0.1)$ is the per-contact transmission probability, $C \sim \text{Poisson}(\lambda_{\text{env}})$ denotes the daily contact rate within structured environments (schools, workplaces), and I is the number of infected (symptomatic) contacts in the agent's network. Infection occurs if a random draw $U \sim \mathcal{U}(0, 1)$ is less than P_{infect}.

2.11 Upon infection, agents transition into the Exposed (Asymptomatic) (E) state, entering an incubation period given by:

$$\tau_{\text{incubation}} \sim \mathcal{U}(2, 15) \text{ days} \tag{1}$$

2.12 This incubation-period prior was selected to reflect reported incubation periods across acute epidemic infections relevant to respiratory and emerging viral disease modeling. Table 2 summarizes reported incubation periods for selected infections. The prior is interpreted as a broad uncertainty range rather than a disease-specific average, allowing the model to represent both short-incubation respiratory infections and moderately longer incubation periods observed in emerging viral infections.

Table 2: Reported incubation periods for selected acute epidemic infections used to justify the incubation-period prior.

Disease	Reported incubation period	Source
Influenza	Usually 1–4 days after exposure	Centers for Disease Control and Prevention (2025)
COVID-19	Average incubation approximately 5–7 days, with a possible range of 1–14 days	Cimolai (2021)
MERS-CoV	Symptoms usually develop approximately 5 days after exposure, with a reported range of 2–14 days	Centers for Disease Control and Prevention (2024a)
Ebola virus disease	Symptoms may appear 2–21 days after exposure, with average onset approximately 8–10 days	Centers for Disease Control and Prevention (2026)
Bundibugyo virus disease outbreak	Incubation period reported as 2–21 days	World Health Organization (2026)

2.13 At the end of incubation, exposed agents either recover without symptoms ($E \to R$) with probability:

$$P_{E \to R} \sim \mathcal{U}(0.02, 0.03) \tag{2}$$

or progress to symptomatic infection ($E \to I$) with probability:

$$P_{E \to I} \sim \mathcal{U}(0.05, 0.1) \tag{3}$$

2.14 Recovered asymptomatic individuals (R) do not contribute to further transmission. Symptomatic infected (I) agents subsequently face two possible outcomes: recovery ($I \to R$) with probability:

$$P_{I \to R} \sim \mathcal{U}(0.03, 0.07) \tag{4}$$

or mortality ($I \to D$) with probability:

$$P_{I \to D} \sim \mathcal{U}(0.01, 0.02) \tag{5}$$

2.15 Recovered symptomatic individuals gain immunity and no longer contribute to transmission. Agents remain in the Exposed state until they reach the minimum required duration, specifically a sampled incubation period. Once this condition is met, they transition based on the defined probabilities. In contrast, Infected agents do not require a minimum infectious duration; instead, they recover or die based solely on daily transition probabilities. Consequently, the transition probabilities do not sum to 1, as transitions from the Exposed state are only considered after the minimum duration, while transitions from the Infected state may occur at any time.

2.16 In the split essential worker rotation with quarantine setting, symptomatic infected agents (I) in the active rotating subgroup are immediately quarantined (Q) to limit transmission within workplaces. The number of quarantined agents follows:

$$Q \sim \text{Poisson}(\lambda_{\text{active}}) \tag{6}$$

where λ_{active} represents the expected number of actively infectious essential workers. Quarantined individuals either recover or remain isolated according to predefined transition probabilities. In this setting, λ_{active} dynamically changes throughout the simulation based on the current subgroup of essential workers assigned to on-site work during each rotation cycle. At any given time, only one subgroup is active, and therefore λ_{active} reflects the number of symptomatic infections within that group. This dynamic tracking ensures that only those currently participating in workplace interactions contribute to potential quarantine events, thereby capturing the temporally structured nature of rotational workforce strategies.

Agent movement rules

2.17 In our SAFE-ABM model, agents interact within structured environments, moving among *families, workplaces, and schools*, where their interactions shape their exposure risk. Each agent is assigned attributes based on *age and occupation*, determining their movement patterns and contact frequency within and across these environments.

2.18 Agents belong to family units, forming the primary social structure of the model. A total of $N_f = 2000$ families are generated, with family sizes (S_f) following a Poisson distribution:

$$S_f \sim \text{Poisson}(\lambda_f) + 2 \tag{7}$$

where $\lambda_f + 2$ defines the base family size, ensuring a minimum household size of two individuals. Families include adults and children, stratified by age to reflect realistic household compositions. Within families, agents engage in frequent and prolonged interactions, making this setting a primary transmission environment.

2.19 Beyond family interactions, workplaces and schools structure additional transmission pathways. A total of $N_c = 100$ workplaces accommodate adult agents, who are randomly assigned to companies. Interaction frequency within workplaces varies by occupation, with essential workers (including healthcare professionals, delivery personnel, and public service employees) maintaining continuous interactions and higher occupational exposure.

2.20 School-aged children are assigned to either elementary or high school environments, with interactions structured by education level. The model includes $N_e = 20$ elementary schools and $N_{hs} = 50$ high schools, reflecting realistic age-specific mixing patterns. Elementary students experience higher contact rates, whereas high school students have fewer interactions.

2.21 At each time step, corresponding to one day in the simulation, agents move between their assigned family, workplace, or school environments, accumulating exposure based on their contact frequency. The daily number of contacts per agent in each environment follows a truncated Poisson distribution, upper-truncated at the number of other agents, k_{env} present in that environment.

$$C_{\text{env}} | C_{\text{env}} \leq k_{\text{env}} \sim \text{Poisson}_T(\lambda_{\text{env}}) \tag{8}$$

where Poisson_T denotes a truncated Poisson distribution and λ_{env} is the expected number of daily interactions within a given environment (env $= \{$family, workplace, or school$\}$).

Simulation Scenarios

3.1 To evaluate how movement restrictions and workplace policies influence disease transmission, we simulate multiple scenarios representing varying intervention intensities. These scenarios structure agent interactions across environments, altering exposure risks and transmission pathways. The *No Restriction scenario* serves as the baseline, reflecting unrestricted movement, where agents interact freely without intervention. We refer to the no-restriction baseline as Scenario 1, the school-from-home intervention as Scenario 2, the essential-workers-only intervention as Scenario 3, the rotational workforce strategy with quarantine as Scenario 4, and the rotational workforce strategy without quarantine as Scenario 5. Subsequent scenarios progressively introduce various levels of mobility restrictions and workplace adjustments, enabling comparative analyses of how public health interventions modify movement patterns and shape epidemic trajectories.

No Restriction scenario

3.2 The No Restriction scenario establishes a baseline condition where individuals move freely across environments without interventions. It reflects typical daily activities in the absence of control measures, allowing agents to interact naturally without mobility constraints. This scenario is considered as a baseline for evaluating the effectiveness of subsequent interventions.

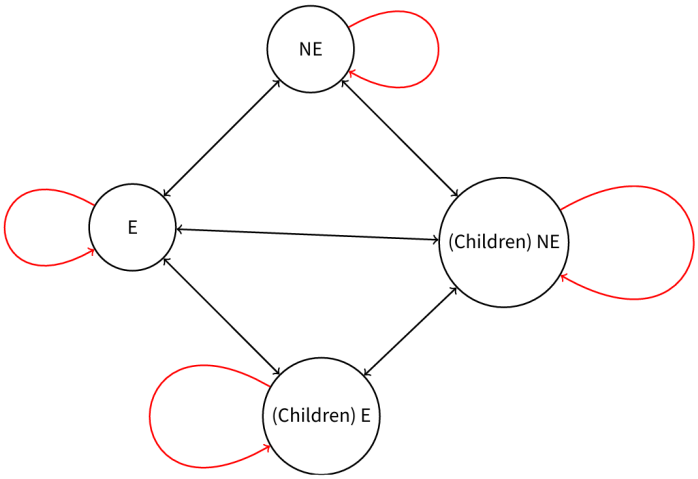


Figure 3: Transmission networks for Scenario 1 (No Restriction). Essential workers (E), non-essential workers (NE), and their children interact freely, reflecting an unrestricted movement scenario.

3.3 As illustrated in Figure 3, agents are categorized as essential (E) and non-essential (NE) workers, and their children are integrated into interaction patterns that mirror real-world social mixing. Essential workers regularly interact within professional environments, primarily engaging with colleagues but also interacting with non-essential workers. This setup represents workplace settings where individuals frequently share spaces and communal facilities. Non-essential workers maintain occupational interactions but typically experience greater flexibility and variability in their external contacts. Children's interaction patterns depend on parental workforce classification. They attend either elementary or high school, following established age-stratified mixing

patterns. Schools thus serve as significant transmission hubs due to unrestricted peer interactions contributing to higher overall contact frequency. Additionally, children from essential and non-essential worker families interact freely, simulating natural social behaviors common in educational and recreational settings. The interaction patterns include intra-group contacts (depicted by red self-loops in Figure 3), highlighting that agents primarily engage with others within their own category. However, given the absence of movement restrictions, interactions across groups also remain unrestricted. This baseline scenario quantifies the natural progression of disease spread prior to applying interventions and provides a comparative foundation for evaluating subsequent movement restrictions and targeted workplace policies.

School from home intervention

3.4 We now introduce a targeted intervention in our model under which children remain at home for online schooling, thereby reducing their external contacts. This School-from-Home intervention models a widely implemented pandemic measure intended to minimize children's exposure by limiting their physical presence in schools and communal spaces.

3.5 In this scenario (Figure 4), the population continues to consist of essential workers (E) and non-essential workers (NE), along with their children. However, unlike the No Restriction scenario, children no longer participate in peer interactions, substantially reducing school-based transmission. Instead, their contacts become restricted primarily to immediate family members, removing the cross-group interactions commonly observed within classrooms. The interaction dynamics depicted in Figure 4 clearly illustrate these structural adjustments. Essential workers maintain their occupational contacts, primarily interacting with other essential workers, as indicated by the self-loops. Similarly, non-essential workers continue interacting predominantly within their own group. Cross-group interactions between essential and non-essential workers persist through shared environments such as workplaces and community settings, highlighting the ongoing role of occupational interactions as potential channels for family-level transmission.

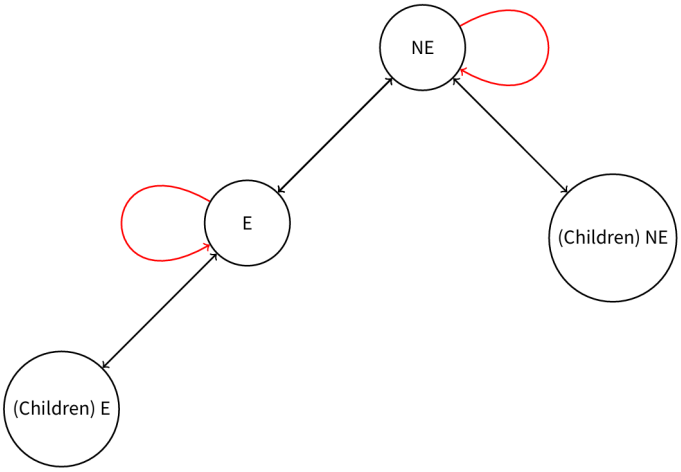


Figure 4: Transmission networks for Scenario 2 (School from Home). Essential workers (E), non-essential workers (NE), and their children interact as shown. Children are confined to family units and do not interact across groups, while adults maintain limited interactions within and between groups.

3.6 Although children remain isolated from their peers, secondary transmission via parental exposure remains possible. This aligns closely with real-world pandemic scenarios, where children indirectly acquire infections when parents interact externally and subsequently introduce pathogens into family settings. By eliminating direct interactions among children, this intervention strategically disrupts a significant route of virus spread, reinforcing public health recommendations aimed at minimizing transmission risks among school-aged populations (Vaivada et al. 2022; Centers for Disease Control and Prevention 2024b, 2022). The School-from-Home scenario enables the model to assess trade-offs between protecting children from external exposure and maintaining necessary workforce interactions. Despite restricted mobility for children, essential and non-essential workers continue their external activities, emphasizing the role occupational interactions play in shaping family-level transmission patterns.

3.7 In parallel, this approach resembles a natural protective mechanism observed among eusocial species. In bee colonies, younger bees typically remain protected within the hive, while older foragers leave to collect

resources, thereby reducing exposure risks for vulnerable colony members. Analogously, keeping children home during a pandemic provides protective isolation while still permitting essential workforce participation. This structured strategy effectively balances disease mitigation and societal functionality, mirroring adaptive division-of-labor behaviors documented in nature (Seeley 1982; Wilson-Rich et al. 2009; Johnson 2010; Naug 2008).

Essential workers only intervention

3.8 Here, only essential workers continue working outside home, while all other individuals remain at home with mobility restrictions. This scenario mimics real-world pandemic policies, where critical services such as healthcare, emergency response, and food distribution must remain operational despite widespread lockdown measures. Essential workers include healthcare professionals, first responders, delivery personnel, and public utility employees, who sustain critical societal functions during health crises (Centers for Disease Control and Prevention 2020b; Bielicki et al. 2020).

3.9 Essential workers face increased exposure risks due to persistent workplace interactions. As shown in Figure 5, these workers primarily engage in occupational settings, represented by self-loops illustrating frequent and repeated workplace contacts. In environments such as hospitals, transportation networks, and service hubs, regular interpersonal interactions significantly elevate the potential for disease transmission. For example, healthcare workers operating in close proximity may inadvertently spread infections even with protective measures, particularly if safety protocols become compromised (Centers for Disease Control and Prevention 2020a; World Health Organization 2020; Khunti et al. 2021; Sim 2020; Wei et al. 2022). Beyond their occupational settings, essential workers also maintain interactions with external contacts, though typically at reduced frequencies. Service-oriented roles such as food delivery and public transportation involve brief but frequent interactions, providing additional pathways for potential transmission. Although precautionary measures (e.g., contactless delivery) aim to reduce risks, brief customer exchanges still pose a residual risk of disease spread (Rahimi et al. 2021; Nguyen et al. 2021; Romero-Michel et al. 2021; Karlsson & Fraenkel 2020). A critical dimension of this scenario involves household transmission risks associated with essential workers. Although children remain at home participating in online schooling, Figure 5 illustrates how parental occupational exposure can introduce infections into family settings. The directed link between essential workers (E) and their children highlights the elevated secondary transmission risk. Children do not engage in external social activities; however, if parents contract infections through workplace exposure, the family environment becomes a significant route for disease spread (Lorenzo & Carrisi 2020; Bi et al. 2021; Karlsson & Fraenkel 2020).

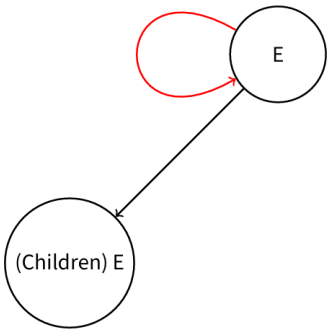


Figure 5: Transmission networks for Scenario 3 (Only Essential Workers' Movement). Essential workers (E) interact within their group, as shown by the self-loop, and maintain connections with their children. Children are isolated within their family units, with no direct interactions outside these units.

3.10 By isolating non-essential workers and enforcing online schooling, this intervention substantially reduces community-level transmission through limited cross-group interactions. Nonetheless, it emphasizes the occupational burden on essential workers, who face elevated direct exposure risks and an increased likelihood of transmitting infections to their families. Consequently, the scenario underscores the critical importance of workplace protections – including enhanced safety protocols and effective personal protective equipment – to minimize secondary household transmission risks (Wang et al. 2020; Madewell et al. 2020; Wu et al. 2020;

Madewell et al. 2021). Since external mobility is concentrated within this specific group, essential workers represent a key transmission link between occupational and family environments, thereby highlighting workplace safety measures as central components for reducing both direct and indirect disease spread.

Split essential workers (rotational workforce with and without quarantine)

3.11 The final intervention category, represented by Scenarios 4 and 5 (Figure 6), implements a rotation system where essential workers are divided into two mutually exclusive subgroups, E1 and E2, alternating weekly between active duty and staying home. These scenarios reduce workplace density while maintaining essential services during a pandemic. Each subgroup operates on a 7-day alternating schedule: one subgroup (say, E1) actively works for one week while the other subgroup (E2) remains at home, with no interaction between groups during this period. In Scenario 4, if workers in the actively working subgroup become symptomatic, they are promptly quarantined to prevent interaction with healthy workers, further reducing transmission risks. Scenario 5 applies the same rotational structure without quarantine. After seven days, roles reverse; E2 becomes active, and E1 stays home. This structured, mutually exclusive rotation explicitly tests whether reducing the number of concurrently active workers, with or without rapid quarantine of symptomatic individuals, effectively reduces disease transmission, or if it unintentionally increases risks due to concentrated interactions within smaller groups.

3.12 As illustrated in Figure 6, when subgroup E1 is active, members interact among themselves, depicted by the self-loop around E1, and subsequently return home to their families (indicated by the arrow from E1 to 'Children E1'). During the same period, the inactive subgroup (E2) stays home, interacting solely with household members. After one week, the subgroups switch roles, and this alternating pattern continues throughout the simulation. Although rotation schemes are intended to lower workplace density, they can inadvertently increase interaction frequency within the actively working subgroup. Unlike situations where the entire workforce is active simultaneously, splitting the workers results in more frequent interactions among a smaller group of colleagues during each rotation. Real-world studies indicate that workforce rotations might unintentionally amplify infection rates if interactions among active subgroup members intensify (Sah et al. 2018; Sim 2020).

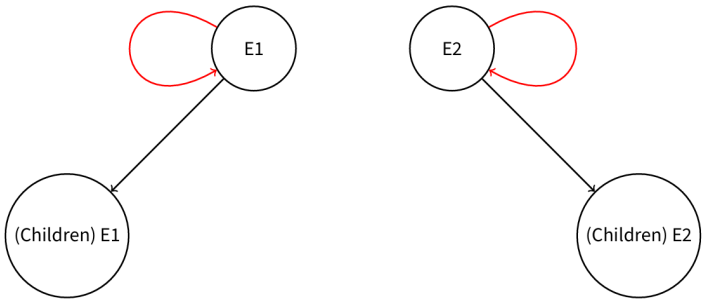


Figure 6: Transmission networks for Scenarios 4 and 5 (Split Essential Workers: Rotation Scheme). Essential workers are divided into two subgroups (E1 and E2) to alternate their movement and duties. Within each subgroup, essential workers interact among themselves (depicted by self-loops), and they also interact with their children. Scenario 4 applies this rotational structure with prompt quarantine of symptomatic active workers, whereas Scenario 5 applies the same rotational structure without quarantine. This setup minimizes sustained high-risk interactions across the entire essential workforce while maintaining household exposure dynamics.

3.13 We hypothesize that while rotational workforce strategies effectively reduce overall workplace exposure, increased interactions within active subgroups could raise within-group transmission risks. However, we also hypothesize that enforcing quarantine for symptomatic individuals could effectively mitigate these heightened risks, resulting in an overall reduction in workplace and secondary household transmission.

3.14 Given that rotation alone might inadequately control transmission risks and could potentially intensify exposure among active workers, we incorporate an additional measure: quarantining symptomatic workers. Quarantine ensures symptomatic workers are quickly removed from circulation, limiting further disease spread within the subgroup. By promptly isolating infected individuals before transmission can occur extensively, quarantine effectively interrupts transmission chains. This practice aligns closely with historical and contemporary public health strategies, where quarantine remains fundamental for controlling infectious disease spread

(Tognotti 2013). Combining workforce rotations with targeted quarantine thus allows evaluation of whether a hybrid intervention is more effective at mitigating workplace and household transmission compared to strategies that solely involve essential workers working continuously.

3.15 This intervention introduces two significant changes in exposure dynamics. First, dividing essential workers into rotating subgroups might increase intra-group transmission risks due to more frequent interactions among a smaller, actively working group. Because only half of the essential workforce is active at any given time, interactions within each subgroup become more frequent, potentially intensifying transmission risks. Second, introducing quarantine as a complementary measure provides a counterbalance, promptly removing symptomatic workers and preventing subgroup-level outbreaks. The effectiveness of rotation in reducing overall exposure therefore heavily depends on how efficiently quarantine measures limit transmission among active subgroup members.

3.16 Contrasting Scenario 3, where all essential workers remain continuously active, Scenario 4 attempts to disrupt prolonged workplace transmission by implementing alternating work schedules combined with quarantine measures, while Scenario 5 evaluates the same rotational structure without quarantine. By explicitly modeling these competing dynamics, Scenarios 4 and 5 allow us to evaluate whether rotating essential workers, with or without quarantine, can reduce both workplace outbreaks and secondary household infections. Results from these scenarios can directly inform public health policies designed to optimize workforce safety while maintaining essential societal functions during pandemics. By assessing structured workforce rotations with and without quarantine practices, the model provides key insights into effectively managing workplace outbreaks, limiting secondary family transmission, and balancing infection control with continued economic and societal stability. These insights emphasize the value of targeted public health measures aimed specifically at protecting high-risk workers while ensuring critical infrastructure remains operational, thereby providing valuable guidance for developing effective pandemic response frameworks.

Algorithm

3.17 Algorithm 1 outlines our SAFE-ABM model by explicitly structuring agent movements, infection dynamics, and intervention mechanisms within a computational framework. The algorithm implementation aligns precisely with the mathematical model described in Section 2 and incorporates the distinct intervention strategies detailed in Subsections 3.1–3.16. By systematically integrating these components, the algorithm ensures that each simulated scenario accurately represents the dynamic effects of policy changes on disease transmission within heterogeneous populations. Specifically, it captures the dynamic restructuring of agent interactions under various intervention conditions, enabling clear comparisons of epidemic outcomes across scenarios.

Uncertainty Quantification and Model Validation

4.1 Ensuring the reliability of simulation outcomes requires a rigorous approach to uncertainty quantification, parameter-range selection, and simulation-based validation. This section introduces the Bayesian Uncertainty Quantification (UQ) framework implemented in our SAFE-ABM model to systematically account for parameter variability and stochastic transmission dynamics. We also outline the parameter-range selection process, simulation-based internal validation approach, and sensitivity analysis used to support robust statistical inference in intervention assessments.

Bayesian Uncertainty Quantification (UQ) framework

4.2 A few studies have explored the use of Bayesian methods for uncertainty quantification in agent-based modeling. Liu et al. (2022) developed a Bayesian framework to quantify uncertainty in ABMs of networked group anagram games, which focuses on clustering diverse player behaviors rather than epidemic dynamics. Their approach integrates Bayesian nonparametric clustering and multinomial regression to model behavioral transitions, making it well-suited for decision-making dynamics rather than structured epidemiological processes. McCulloch et al. (2022) addressed ABM calibration by combining History Matching and Approximate Bayesian Computation to refine parameter spaces and provide credible intervals. While their approach enhances parameter estimation and model fitting, it does not focus on uncertainty propagation in epidemic modeling.

Algorithm 1 Agent-Based Model (ABM) for Epidemic Spread

Step 1: Initialization
Input: Population size $N = 10,000$, Families $n_{\text{fam}} = 2000$
Output: Agent population with attributes
Assign agent attributes: $\rightarrow$ Unique ID, Age, Gender ({M, F}), Cohort ({E, NE, HS, e})
$\rightarrow$ NE = Non-Essential Worker, E = Essential Worker, HS = High School, e = Elementary Student
Assign Family, School, and Work Networks
Assign meeting frequencies for family, work, and school
Save the initialized population
Step 2: Disease Seeding
Input: Initialized population
Output: Initial infected and exposed individuals $\rightarrow$ Exposed: $E_0 = 10$ agents $\rightarrow$ Infected: $I_0 = 5$ agents
$\rightarrow$ Compute initial susceptible population: $S_0 \leftarrow N - E_0 - I_0$
Step 3: Epidemic Dynamics
Input: Initial conditions (S_0, E_0, I_0), time steps $t = 1, \ldots, T$
Output: Time-series $(S_t, E_t, I_t, R_t, D_t)$
Define transmission parameters: $\rightarrow$ Transmission Probability: $\beta \sim \text{Uniform}(0.05, 0.1)$
$\rightarrow$ Exposure-to-Infection Probability: $P_{E\rightarrow I} \sim \text{Uniform}(0.05, 0.1)$
$\rightarrow$ Exposure-to-Recovery Probability: $P_{E\rightarrow R} \sim \text{Uniform}(0.02, 0.03)$
$\rightarrow$ Recovery Rate: $P_{I\rightarrow R} \sim \text{Uniform}(0.03, 0.07)$
$\rightarrow$ Death Rate: $P_{I\rightarrow D} \sim \text{Uniform}(0.01, 0.02)$
for $t = 1$ to T **do**
    Transmission Dynamics: Each infected agent (I) contacts a number of others drawn from a truncated Poisson distribution, $C_{\text{env}}|C_{\text{env}} \leq k_{\text{env}} \sim \text{Poisson}_T(\lambda_{\text{env}})$, upper-bounded by the number of other agents, k_{env} in the environment.
    Infection is governed by the probabilistic function $P_{\text{infect}} = 1 - e^{-\beta CI}$, where β is the per-contact transmission probability and C is drawn as described under Infection Dynamics.
    Update exposed population:

$$E_{t+1} \leftarrow E_t + \sum_{i \in I} \mathbb{I}(U_i < P_T)$$

    State Transitions:
    Exposed $\rightarrow$ Infected with probability P_E
    Infected $\rightarrow$ Recovered with probability γ
    Infected $\rightarrow$ Deaths with probability μ
    Update susceptible population:

$$S_{t+1} \leftarrow S_t - \left(E_{t+1} - E_t\right)$$

end for
Step 4: Intervention Phase (After $t > 14$**)** $\rightarrow$ **No Restriction:** Continue baseline dynamics
$\rightarrow$ **School-From-Home:** Remove school-based contacts
$\rightarrow$ **Essential Workers Bubble:** Restrict non-essential workers' contacts
$\rightarrow$ **Split Essential Workers with Quarantine:** Rotate essential workers in subgroups and quarantine symptomatic active workers
$\rightarrow$ **Split Essential Workers without Quarantine:** Rotate essential workers in subgroups without applying quarantine
Return: Time-series $(S_t, E_t, I_t, R_t, D_t)$

4.3 In contrast, our study introduces a Bayesian uncertainty quantification framework specifically tailored to *epidemiological agent-based models*. Unlike clustering-based and calibration-focused approaches, our framework directly propagates uncertainty through stochastic simulations, systematically quantifying intervention impacts on epidemic trajectories. By simulating 100 independent runs with distinct population samples, our framework captures the variability introduced by heterogeneous population structures and stochastic transmission dynamics. This structured approach enables a rigorous assessment of uncertainty in disease outcomes, offering novel insights into intervention effectiveness within agent-based epidemiological models.

4.4 Our Bayesian UQ framework assigns prior distributions to key epidemiological parameters (Martinez & Achcar 2014; MacLehose & Hamra 2014), explicitly capturing variability in disease transmission, recovery, and mortality

rates, thus yielding probabilistic epidemic trajectories rather than deterministic outcomes. To reflect plausible disease dynamics, we define prior distributions over core transmission parameters:

$$\begin{aligned} \text{Transmission Probability} &\sim U(0.05, 0.1) \\ \text{Recovery Rate} &\sim U(0.03, 0.07) \\ \text{Mortality Rate} &\sim U(0.01, 0.02) \end{aligned} \tag{9}$$

4.5 For each simulation run, values of the epidemiological parameters in Equation 9 were sampled from their specified uniform prior distributions, with the same prior ranges applied across all intervention scenarios to ensure comparable scenario evaluation.

4.6 Each simulation explicitly captures stochastic interactions among agents (intra-run variability) and differences in epidemiological parameters and population structures across independent runs (inter-run variability), providing a comprehensive uncertainty assessment. This methodology enhances model robustness by incorporating prior knowledge and uncertainty directly into the agent-based framework.

4.7 To fully capture uncertainty, we conduct 100 independent simulation runs, each initialized with a unique synthetic population realization. At initialization, 15 agents were randomly selected from the full synthetic population, with 10 assigned to the exposed state and 5 assigned to the infected state, using the same outbreak-seeding procedure across all scenarios. No additional outbreak takeoff criterion or reseeding procedure was imposed. This design ensures that observed variability in outcomes is not solely due to stochastic agent interactions but also accounts for differences in underlying population structures and parameter variability. The Force of Infection (FoI), along with state variables (Susceptible, Exposed, Infected, Recovered, Deaths), is recorded across runs, facilitating post-simulation statistical analysis of uncertainty bounds.

4.8 To efficiently handle large-scale uncertainty propagation, the simulation workflow was optimized and parallelized using Python's `ProcessPoolExecutor` from the `concurrent.futures` module. For each intervention scenario, we generated 100 distinct synthetic populations, each consisting of 10,000 agents and simulated over 120 days. This design allowed uncertainty to reflect both sampled epidemiological parameters and variation across population realizations. The optimized implementation was executed on a multi-core workstation using 37 processing cores. The no-restriction scenario required the longest computational time, approximately 8 hours, while the school-from-home scenario required approximately 5 hours; the remaining optimized scenarios completed more quickly under the same parallel-processing setup.

4.9 The computational implementation of our Bayesian uncertainty quantification framework is summarized in Algorithm 2, which outlines the structured approach used to quantify parameter variability and stochastic uncertainty across multiple simulation runs.

Parameter-range selection and plausibility assessment

4.10 To ensure the robustness of our SAFE-ABM model and the reliability of simulated outcomes, we implement parameter-range selection, simulation-based validation, and sensitivity analysis within a Bayesian uncertainty quantification (UQ) framework. Since our study uses synthetic populations rather than empirical epidemiological data, a structured simulation-based validation approach is crucial to maintain statistical rigor and avoid biases from single-run stochastic variability. The simulation framework, including the generation of synthetic populations and model initialization, was implemented in Python version 3.12.7.

4.11 This study did not perform formal empirical calibration to an external ground-truth epidemic dataset. Instead, parameter ranges were selected through exploratory simulation-based plausibility assessment. Candidate ranges were evaluated by repeated simulations to determine whether they produced qualitatively reasonable epidemic dynamics, including an initial growth phase, a discernible epidemic peak, and a subsequent decline. Parameter ranges that produced immediate extinction, highly unstable trajectories, or unrealistically rapid saturation across most simulations were not retained. Thus, the term "plausible" refers to qualitative consistency with expected epidemic progression rather than agreement with a specific observed outbreak dataset.

4.12 To ensure fair comparisons across intervention strategies, all scenarios employ identical prior distributions for the epidemiological parameters. Parameters are drawn from uniform distributions within the retained ranges so that differences in outcomes across scenarios reflect differences in intervention structure under the same epidemiological assumptions, rather than scenario-specific parameter choices.

4.13 Model validation in this study refers to simulation-based internal validation rather than empirical validation against an external observed outbreak dataset. Specifically, we assessed whether the SAFE-ABM framework

Algorithm 2 Uncertainty Quantification (UQ) in Agent-Based Model (ABM)

Step 5: Initialization
Input: Number of simulation runs per seed n_{sim}, Number of seeds n_{seeds}, Time steps T
Output: Quantified uncertainty estimates for epidemic dynamics
Define Bayesian priors for epidemic parameters:
→ Per-contact transmission probability: $\beta \sim \mathcal{U}(0.05, 0.1)$
→ Recovery rate: $\gamma \sim \mathcal{U}(0.03, 0.07)$
→ Death rate: $\mu \sim \mathcal{U}(0.01, 0.02)$
Sample distinct agent populations for each simulation run ($N = 10,000$ per run)
Step 6: Scenario-Based Simulation
Input: Initialized agent populations, priors
Output: Time-series epidemic data for all scenarios
for seed $= 1$ to n_{seeds} **do**
  for sim $= 1$ to n_{sim} **do**
    Sample (β, γ, μ) from prior distributions
    Load agent population for current seed
    Run scenario-based simulation:
    - **No Restriction:** Baseline transmission dynamics
    - **School From Home:** Remove school-based contacts
    - **Essential Workers Only:** Restrict non-essential workers' interactions
    - **Split Essential Workers with Quarantine:** Alternate essential workers in rotations of 7-day shifts and quarantine symptomatic active workers
    - **Split Essential Workers without Quarantine:** Alternate essential workers in rotations of 7-day shifts without applying quarantine
    Record epidemic trajectories $(S_t, E_t, I_t, R_t, D_t)$ for each run
  end for
end for
Step 7: Aggregation of Results
Input: Time-series data from multiple runs
Output: Uncertainty estimates (quantiles)
Estimate quantiles $(q_{2.5}, q_{50}, q_{97.5})$ for:
→ Exposed (E_t)
→ Infected (I_t)
→ Recovered (R_t)
→ Deaths (D_t)
→ Force of Infection
Step 8: Save and Analyze Results
Store quantified uncertainty estimates for further Bayesian inference
Return: Uncertainty estimates for epidemic states

produced coherent and consistent system-level epidemic behavior across repeated stochastic simulations and distinct synthetic populations. For each scenario, 100 independent simulation runs were performed, with each run initialized using a distinct synthetic population and sampled epidemiological parameters.

4.14 The validation assessment focused on trajectory-level consistency across susceptible, exposed, infected, recovered, and deceased states. Median trajectories and 2.5%–97.5% simulation-based uncertainty intervals were computed to evaluate whether scenario outcomes were stable across repeated runs rather than being driven by a single stochastic realization. Thus, the validation approach was based on internal simulation consistency, uncertainty quantification, and robustness across population seeds, not on calibration or validation against a specific empirical ground-truth dataset.

4.15 We analyzed cumulative recoveries and cumulative deaths as downstream resolved-outcome measures of epidemic burden. Although cumulative infection burden is central to transmission dynamics, cumulative recover-

ies and deaths provide complementary information about how infections resolve over time. Cumulative deaths represent the severe public-health consequence of transmission, while cumulative recoveries reflect the number of agents who passed through infection and survived. Because recovery and death probabilities were drawn from the same prior ranges across scenarios, differences in these outcomes arise from scenario-specific transmission structure and timing rather than scenario-specific outcome probabilities. We therefore interpret cumulative recoveries and deaths alongside the infected trajectories rather than as independent measures of transmission alone.

4.16 Sensitivity analysis further evaluated the stability of model conclusions under parameter variability. Key epidemiological parameters – including transmission probability, minimum exposure duration, and recovery rate – were systematically varied within their prior ranges to assess impacts across different scenarios. In addition, the use of distinct synthetic populations across simulations naturally captures heterogeneity in population structures and stochastic transmission pathways. This approach supports robust comparison of intervention strategies across a broad range of plausible simulation conditions. By combining parameter-range selection, simulation-based internal validation, uncertainty quantification, and sensitivity analysis, we strengthen the credibility and reliability of SAFE-ABM for evaluating intervention strategies under uncertainty.

4.17 To support numerical comparison across intervention scenarios, we computed inflection points, final cumulative recovery and death summaries, pairwise Welch's t-tests (Welch 1947), and mutual information sensitivity rankings (Shannon 1948; Kraskov et al. 2004). Inflection points were computed numerically from the cumulative recovery and death trajectories as the time point at which the rate of change began to decline. For a cumulative trajectory $f(t)$, the daily increment was defined as $\Delta f(t) = f(t) - f(t-1)$. The inflection day was identified when the daily increment at the next time point became smaller than the current daily increment, that is, when $\Delta f(t+1) < \Delta f(t)$, equivalently $f(t+1) - f(t) < f(t) - f(t-1)$. This finite-difference criterion provides a discrete-time analogue of the second-derivative criterion and directly captures the point at which the rate of change starts to decline. The criterion was applied to the scenario-specific median cumulative trajectories shown in Figures 8 and 9 to obtain the inflection day reported for each scenario and outcome. The same criterion was applied separately across the 100 run-level trajectories to summarize the mean, standard deviation, and 95% interval of the inflection days.

4.18 Final cumulative recoveries and deaths were computed at the final simulation time point and summarized using the mean and standard deviation across the 100 simulation runs. Pairwise two-sided Welch's t-tests were used to compare scenario-specific outcome distributions, allowing for unequal variances. Statistical significance was assessed at $\alpha = 0.05$.

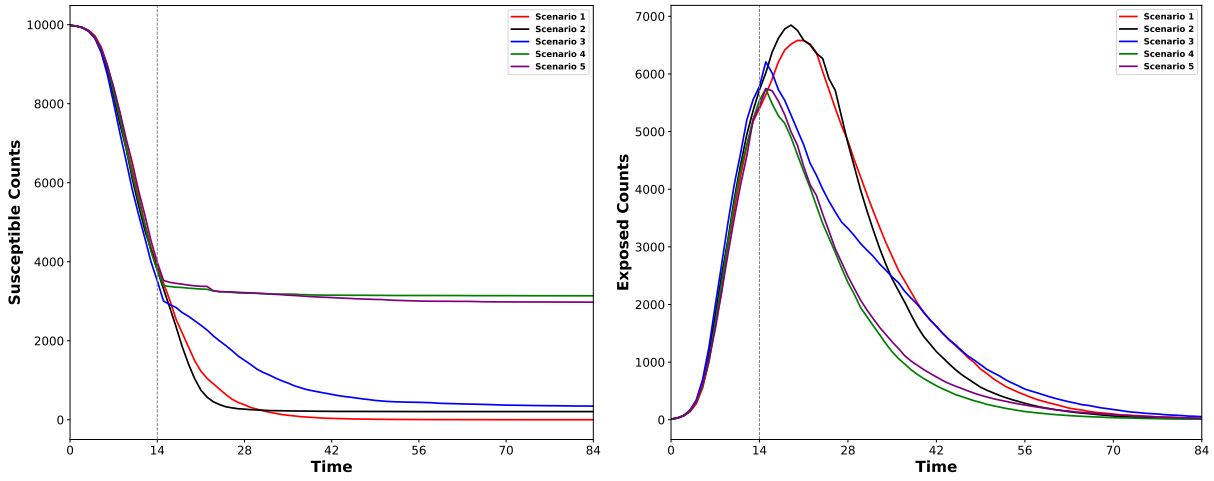
4.19 To assess input-output sensitivity, mutual information was computed between each model input and selected simulation outputs using `mutual_info_regression` from `scikit-learn` (Pedregosa et al. 2011). The input features included the sampled epidemiological parameters and the seed index, while the outputs included final cumulative recoveries, final cumulative deaths, recoveries at day 30, and deaths at day 30. Larger mutual information values indicate stronger association between an input feature and variation in the corresponding output.

4.20 While our simulations explore plausible parameter ranges and epidemic behaviors through exploratory analyses, formal calibration to empirical data remains outside the scope of this study. This decision reflects our immediate focus: to develop, implement, and validate a flexible agent-based framework capable of realistically simulating structured interventions. Future work will extend this framework by incorporating empirical epidemiological data for full parameter calibration, enabling more precise, data-driven policy assessments.

Results

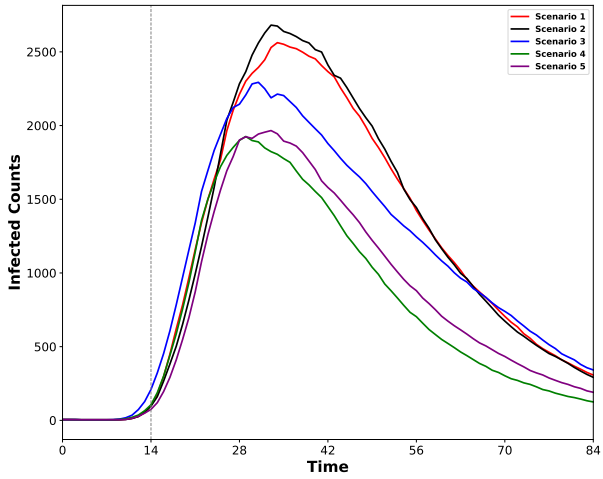
Epidemic progression across all scenarios

5.1 Figure 7 illustrates the epidemic trajectories for susceptible, exposed, infected (symptomatic), recovered, and deaths populations across five distinct intervention scenarios, with a vertical reference line at day 14 indicating the start of scenario-specific interventions. Scenario 1 (no restrictions, red) serves as the baseline, while Scenario 2 (school-from-home, black), Scenario 3 (essential-workers-only, blue), Scenario 4 (rotational shifts among essential workers with quarantine, green), and Scenario 5 (rotational shifts among essential workers without quarantine, purple), represent distinct structural intervention strategies.

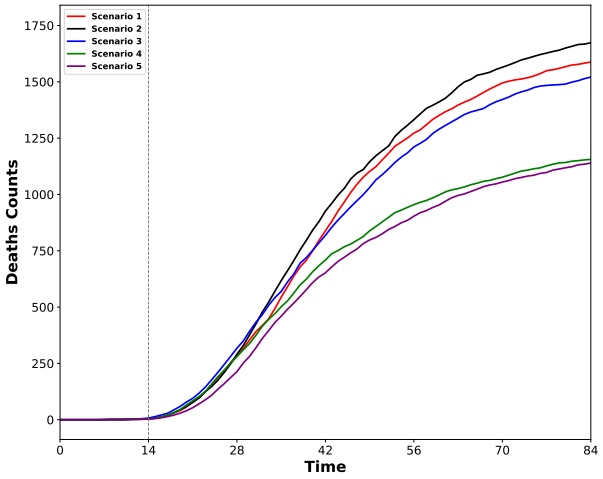


(a) Susceptible Counts

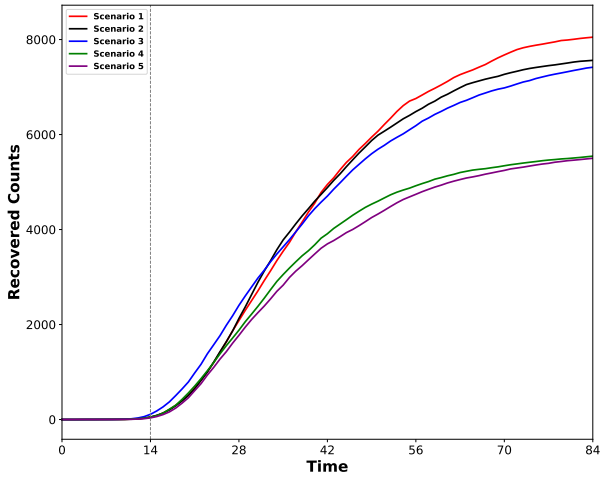
(b) Exposed (Asymptomatic) Counts



(c) Infected (Symptomatic) Counts



(d) Deaths



(e) Recovered

Figure 7: Median epidemic trajectories for the susceptible, exposed, infected, recovered, and deceased populations under five intervention scenarios: Scenario 1 (no restrictions), Scenario 2 (school-from-home), Scenario 3 (essential workers only), Scenario 4 (rotational shifts among essential workers with quarantine), and Scenario 5 (rotational shifts among essential workers without quarantine). Results reflect the median across 100 simulation runs, each conducted with a distinct synthetic population of 10,000 agents. The vertical reference line at day 14 indicates the start of the intervention period; all scenarios follow the same baseline contact structure before day 14, after which scenario-specific intervention rules are applied.

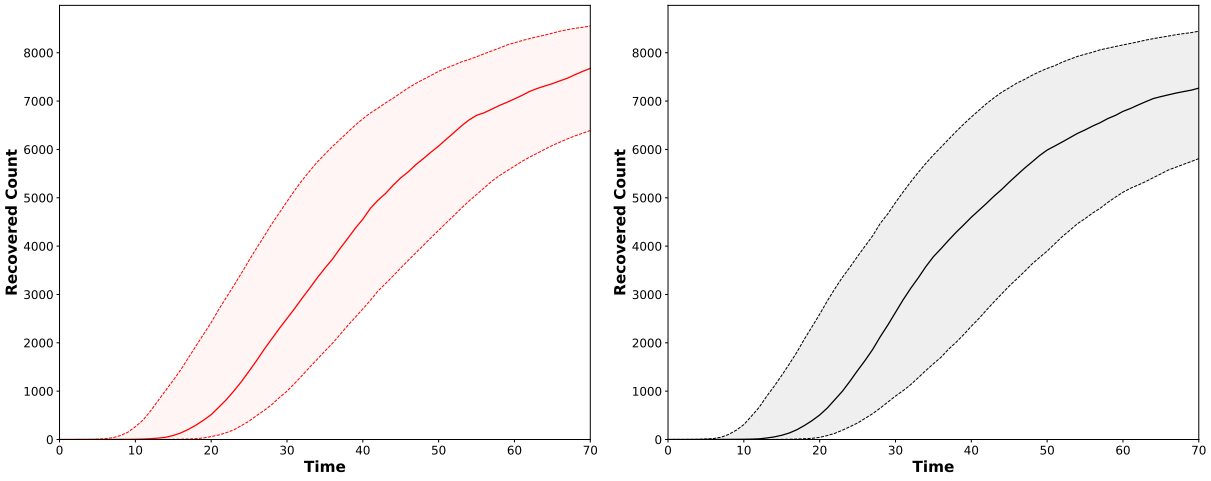
5.2 Panels (a) and (b) present the susceptible and exposed (asymptomatic) populations. In Scenario 2 (school-from-home), the susceptible population drops significantly starting on day 16, soon after the intervention begins, indicating substantial movement of agents from the susceptible state into the exposed state during the early intervention period. This early decline should not be interpreted as intervention success; rather, it reflects unintended family transmission dynamics. Children confined at home, while parents continue external interactions through work and other necessary activities, may experience increased within-family contacts and infection risk. Although the susceptible trajectory in Scenario 2 stabilizes after approximately day 31, the initial rapid decline suggests that school-from-home alone may temporarily shift transmission into household settings when adult external contacts are not also reduced. This demonstrates a crucial insight: closing schools without reducing parental interactions outside the home may inadvertently increase family-level transmission for some period of time. Scenario 3 (essential-workers-only, blue) shows a slower reduction in susceptible individuals compared with Scenarios 1 and 2, reflecting the benefit of restricting external mobility to essential workers. However, panel (b) shows that the exposed population under Scenario 3 remains elevated for a longer period. By around day 40, the exposed count in Scenario 3 becomes comparable to Scenario 1, and after approximately day 48, it exceeds Scenario 1. This indicates prolonged vulnerability among essential workers and their families when external mobility is concentrated within this group. In contrast, the rotational shift strategies provide stronger protection. Scenario 5, which applies rotational shifts without quarantine, shows substantial improvement compared with Scenarios 1–3. By the end of the simulation, approximately 2,976 agents remain susceptible, indicating that splitting essential workers into mutually exclusive groups and alternating their workdays every 7 days substantially reduces transmission even without quarantine. However, the best strategy remains Scenario 4, where rotational shifts are combined with quarantine. Scenario 4 consistently preserves the largest susceptible population and produces the lowest exposed counts over time. By the end of the simulation, approximately 3,136 agents remain susceptible, showing additional protection beyond rotation alone. These results indicate that workforce rotation reduces sustained exposure among essential workers, while quarantining symptomatic active workers further strengthens transmission control and provides more sustained protection for this critical population.

5.3 Panel (c) aligns logically with expected epidemiological patterns: higher exposure leads directly to more symptomatic infections. Scenario 2 experiences the greatest infection counts from approximately day 28 to day 52; after day 52, Scenarios 1 and 2 show similarly high infection counts, directly reflecting their elevated exposure levels. Scenario 4 achieves the lowest infection counts over time, distinctly separating from Scenario 5 around day 28 and maintaining consistently lower counts thereafter. It is also worth noting that Scenario 5 has a lower infection burden compared with Scenarios 1–3, further showing the effectiveness of splitting essential workers into mutually exclusive groups. This sustained improvement emphasizes the benefit of structured workforce rotation, with the strongest reduction observed when rotation is combined with targeted quarantine practices. Panels (d) and (e) illustrate cumulative deaths and recoveries, consistent with the infection trajectories. Scenarios 1 and 2 show higher recoveries due to elevated infections over time, confirming that recovery counts alone should not be interpreted as intervention success. Scenario 5 shows the lower median cumulative recovery trajectory, whereas Scenario 4 produces the lowest cumulative deaths and preserves the largest susceptible population, aligning with the effectiveness of rotational workforce strategies in reducing initial exposures and symptomatic infections. Clearly, the strategy of splitting essential workers into mutually exclusive groups and promptly quarantining symptomatic individuals appears to be the most protective approach for this high-risk population, particularly during pandemics when critical societal functions must continue.

Uncertainty quantification across all scenarios

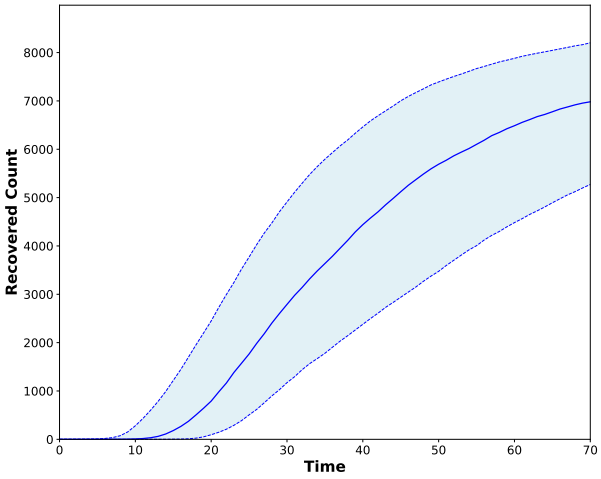
Predictive uncertainty in cumulative recoveries

5.4 Figure 8 shows predictive uncertainty in cumulative recoveries for each of the five intervention scenarios. For consistency, Scenario 1 (no restrictions) is shown in red, Scenario 2 (school-from-home) in black, Scenario 3 (essential-workers-only) in blue, Scenario 4 (rotational shifts among essential workers with quarantine) in green, and Scenario 5 (rotational shifts among essential workers without quarantine) in purple.

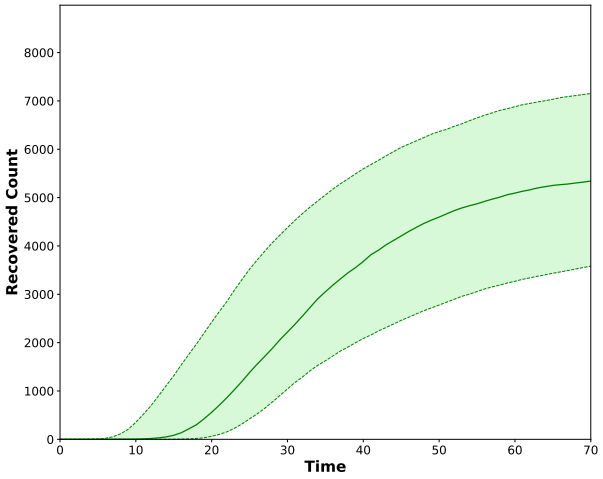


(a) Scenario 1

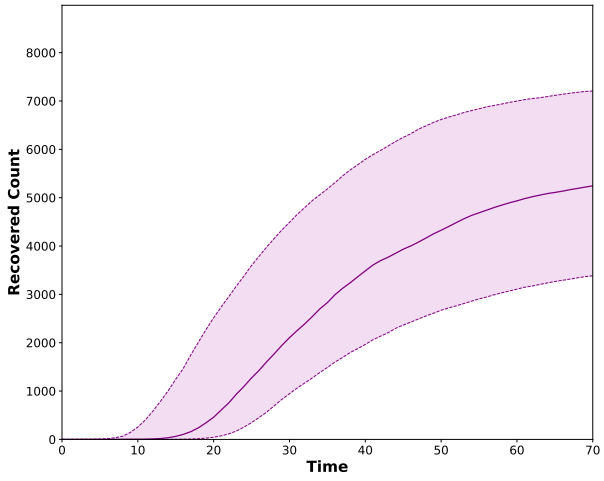
(b) Scenario 2



(c) Scenario 3



(d) Scenario 4



(e) Scenario 5

Figure 8: Predictive uncertainty in recovered population across five intervention scenarios. Each panel shows the median trajectory (solid line) with the 2.5th and 97.5th percentiles (shaded region) based on 100 simulation runs. Each run used a distinct synthetic population of 10,000 agents. The quantile bands reflect uncertainty from both stochastic dynamics and parameter variation. Colors are consistent across scenarios, with each scenario represented by the same color used in earlier figures for ease of comparison.

5.5 In Scenarios 1 and 2, cumulative recoveries rise quickly, with Scenario 1 reaching the highest overall recovery count by the end of the simulation and Scenario 2 remaining slightly lower cumulatively. However, these higher recovery counts should not be interpreted as evidence of better epidemic management. Rather, they reflect widespread infections under no restriction and incomplete control under the school-from-home intervention, where transmission continues through household and adult external contacts. The broad predictive interval (2.5th–97.5th percentiles) in Scenario 1 further highlights substantial variability and unpredictability in epidemic outcomes across simulation runs under minimal restrictions. Scenario 3 displays slower growth in recoveries, reflecting fewer infections due to structural interventions that restrict contact and mobility.

5.6 Scenario 3 shows slower cumulative recovery growth than Scenarios 1 and 2, reflecting fewer infection-driven recoveries under the essential-workers-only intervention. However, its predictive interval is not consistently narrower than those of Scenarios 1 and 2 across the full simulation period, indicating that reduced cumulative recovery growth should be interpreted separately from recovery-trajectory uncertainty. Scenarios 4 and 5 produce lower cumulative recovery trajectories than Scenarios 1–3, reflecting the effectiveness of limiting overall infections through strategic workforce splitting. Among the two rotational workforce strategies, Scenario 5 shows the lower median cumulative recovery trajectory, whereas Scenario 4 shows a narrower predictive interval, indicating greater consistency in recovery outcomes under the quarantine-based rotation strategy. Thus, by dividing essential workers into mutually exclusive groups alternating every 7 days and promptly quarantining symptomatic individuals, the quarantine-based rotational intervention reduces transmission opportunities and improves the stability of recovery outcomes within this high-risk population.

5.7 Moreover, Scenarios 4 and 5 show more stable recovery trajectories compared with Scenarios 1–3, indicating more controlled epidemic dynamics across simulation runs. This is important because Scenario 5 demonstrates that workforce rotation alone provides meaningful benefit even without quarantine, while Scenario 4 shows that adding quarantine further strengthens this protection. These findings demonstrate the value of explicitly modeling interventions tailored specifically to essential workers, a critical population during pandemics when societal functions must continue. By structurally modifying interactions rather than relying solely on generalized parameter adjustments, this targeted approach substantially reduces epidemic uncertainty and stabilizes long-term recovery outcomes, offering critical insights for high-risk population management.

5.8 To further illustrate how the interventions alter epidemic timing, we examine the inflection points in the recovery trajectories. The inflection point is defined as the time at which the rate of cumulative recoveries begins to decline, indicating a slowdown in the number of new recoveries per time step. Because recoveries occur after infection, this transition reflects changes in the preceding infection process as well as the effect of intervention strategies on epidemic progression. In Scenario 1, the inflection point in the median recovery trajectory occurs around day 27, with a mean of 30.71 ± 6.62 days across simulation runs, reflecting an uncontrolled surge in infections followed by widespread recovery. Scenario 2 shows a similar but slightly later inflection point in the median recovery trajectory around day 28, with a mean of 29.75 ± 6.79 days across simulation runs, suggesting that school-from-home alone does not substantially shift the recovery trajectory into an earlier slowing phase.

5.9 By contrast, Scenarios 3, 4, and 5 show earlier recovery inflection patterns. For Scenario 3, the inflection point in the median recovery trajectory occurs around day 23, with a mean of 25.37 ± 5.16 days, indicating that restricting mobility to essential workers reduces the pace of infection-driven recoveries compared with Scenarios 1 and 2. For Scenario 5, the inflection point in the median recovery trajectory also occurs around day 23, with a mean of 25.19 ± 4.21 days, showing that rotational splitting of essential workers remains effective even without quarantine. For Scenario 4, the inflection point in the median recovery trajectory occurs around day 25, with the lowest mean recovery inflection day, 24.36 ± 4.57 days, indicating a similarly early and stable transition under rotational shifts with quarantine. These results show that stronger structural interventions move the recovery trajectory into a slower accumulation phase earlier than Scenarios 1 and 2, with Scenarios 4 and 5 further supporting the benefit of rotational splitting among essential workers.

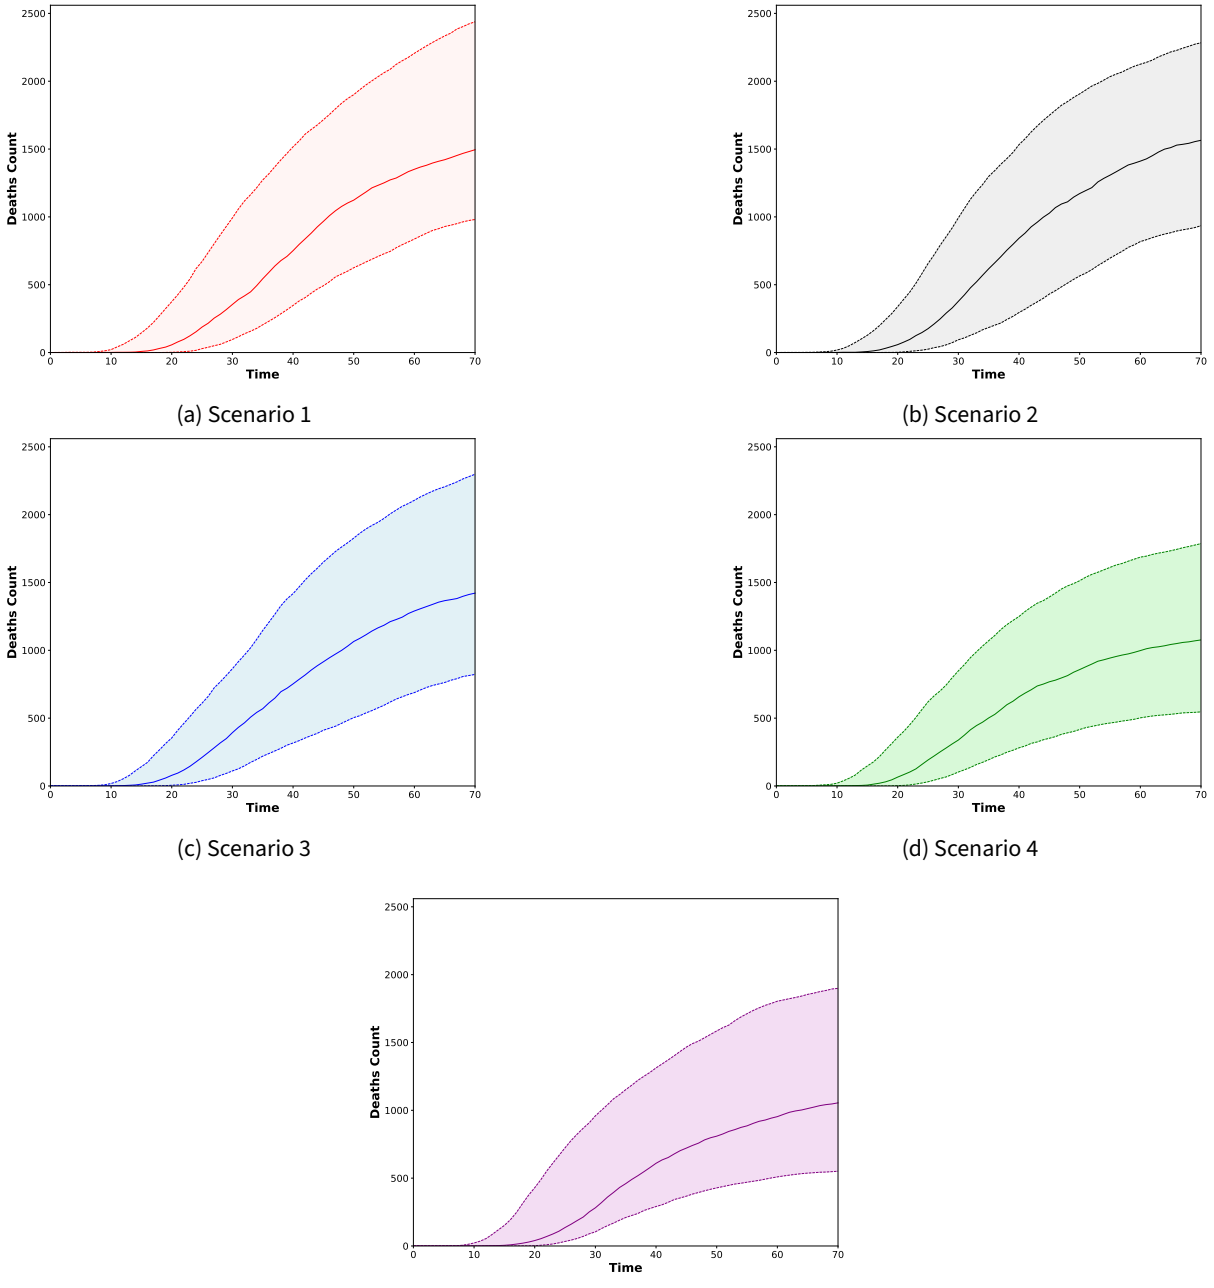
Predictive uncertainty in cumulative deaths

5.10 Similarly, Figure 9 examines predictive uncertainty in cumulative deaths under each intervention scenario, using the same consistent scenario color scheme as before: Scenario 1 (no restrictions, red), Scenario 2 (school-from-home, black), Scenario 3 (essential-workers-only, blue), Scenario 4 (rotational shifts among essential workers with quarantine, green), and Scenario 5 (rotational shifts among essential workers without quarantine, purple).

Table 3: Inflection points for cumulative recovery and death curves across simulation scenarios.

	Recovery inflection day			Death inflection day		
Scenario	Inflection point	Mean ± SD	95% interval	Inflection point	Mean ± SD	95% interval
Scenario 1: No Restriction	27	30.71 ± 6.62	20.00–42.00	35	35.31 ± 7.54	22.00–47.52
Scenario 2: School From Home	28	29.75 ± 6.79	18.00–43.52	32	34.41 ± 6.79	23.00–46.00
Scenario 3: Essential Workers Only	23	25.37 ± 5.16	17.00–34.00	38	34.56 ± 8.42	21.00–53.57
Scenario 4: Split Essential Workers + Quarantine	25	24.36 ± 4.57	17.00–32.00	32	30.48 ± 5.48	21.00–40.00
Scenario 5: Split Essential Workers without Quarantine	23	25.19 ± 4.21	17.48–32.52	32	30.94 ± 5.67	21.48–41.52

Note: Inflection days were computed numerically from the cumulative recovery and death trajectories using the finite-difference rate-decline criterion described in Section 4.17. The inflection point column reports the value obtained by applying this criterion to the median trajectory shown in Figures 8 and 9, while the mean ± SD and 95% interval summarize the inflection days obtained by applying the same criterion separately across 100 simulation runs.



(a) Scenario 1

(b) Scenario 2

(c) Scenario 3

(d) Scenario 4

(e) Scenario 5

Figure 9: Predictive uncertainty in cumulative deaths under five intervention scenarios. Each panel presents the median trajectory (solid line) and the 2.5th and 97.5th percentiles (shaded region), computed from 100 simulation runs using distinct synthetic populations of 10,000 agents. The uncertainty bands incorporate variability due to both stochastic interactions and parameter uncertainty. Scenario colors are consistent with earlier figures for ease of comparison.

5.11 In Scenario 1, cumulative deaths increase rapidly, resulting in the highest death burden and one of the broadest predictive intervals among all scenarios. The extensive variability in this scenario reflects the inherent unpredictability of epidemics under no restriction, where unchecked transmission can lead to substantially different mortality trajectories across simulation runs. Compared with the recovery trajectories, the death trajectories generally show greater uncertainty, which is expected because deaths occur downstream of infection and depend on stochastic progression from exposure to symptomatic infection and from infection to mortality. Introducing structural interventions reduces cumulative deaths and improves the overall stability of mortality outcomes. Scenario 3 produces lower cumulative deaths compared with Scenario 1, while Scenario 2 shows more limited improvement, indicating that school-from-home alone does not fully control death burden when household transmission pathways remain active.

5.12 Scenarios 4 and 5 show the greatest reduction in cumulative deaths, with Scenario 4 remaining the most effective. Scenario 5 demonstrates that rotational workforce splitting alone reduces mortality compared with Scenarios 1–3, while Scenario 4 shows the additional benefit of combining rotation with timely quarantine of symptomatic active workers. Through strategic rotation of essential worker groups and prompt quarantine, Scenario 4 reduces sustained workplace exposure and limits secondary transmission among this critical high-risk group. Although uncertainty remains visible in the death trajectories, the lower cumulative deaths in Scenarios 4 and 5 indicate more controlled mortality outcomes under targeted structural interventions. These findings demonstrate that explicit modifications of interaction patterns can significantly improve outcomes for essential workers and their families, even under varying population compositions. They also highlight that interventions designed specifically for essential workers, who are critical to maintaining societal functions during pandemics, can substantially reduce epidemic severity, minimize mortality, and improve the stability of epidemic outcomes.

5.13 To better evaluate intervention impacts on mortality dynamics, we also examine the inflection points of cumulative death trajectories. The inflection point represents the time when the rate of cumulative deaths begins to slow down, marking a crucial turning point where new deaths per time step start to decrease either due to intervention effects or the natural progression of the epidemic. In Scenario 1, the death inflection point occurs around day 35, with a mean of 35.31 ± 7.54 days across simulation runs, reflecting a prolonged period of mortality accumulation under no restriction. Scenario 3 shows the latest inflection point from the median trajectory, around day 38, and also has the widest 95% interval, suggesting greater variability in the timing of mortality slowing under the essential-workers-only intervention. Scenario 2 has an inflection point around day 32, but its mean inflection timing remains closer to Scenarios 1 and 3 than to the rotational shift strategies, indicating that school-from-home modifies transmission patterns but does not stabilize mortality dynamics as strongly as the rotational approaches.

5.14 Scenarios 4 and 5 both show death inflection points around day 32 from the median trajectories, but the summaries across the 100 simulation runs indicate earlier and more stable mortality transitions than Scenarios 1–3. Scenario 4 has the lowest mean death inflection time, 30.48 ± 5.48 days, while Scenario 5 is very close at 30.94 $\pm$ 5.67 days. This suggests that splitting essential workers into mutually exclusive rotating groups helps move the death trajectory into a slower accumulation phase earlier in the epidemic course. Importantly, Scenario 5 shows that this improvement remains evident even without quarantine, supporting the effectiveness of the rotational shift structure itself. However, the slightly stronger performance of Scenario 4 indicates that adding prompt quarantine of symptomatic active workers provides additional protection and further stabilizes mortality outcomes.

Statistical validation and sensitivity analysis

5.15 To provide a formal numerical evaluation of the cumulative recovery and death outcomes shown in Figures 8 and 9, we summarized the final cumulative values across the 100 simulation runs. As shown in Table 4, Scenarios 4 and 5 produced substantially lower final cumulative recoveries and deaths than Scenarios 1–3. Since cumulative recoveries reflect individuals who were previously infected, the lower final recovery counts under Scenarios 4 and 5 should be interpreted as evidence of reduced infection spread, rather than poorer recovery performance. Scenario 4 produced the lowest final cumulative deaths, followed closely by Scenario 5, indicating that rotational splitting of essential workers strongly reduces mortality outcomes, with additional descriptive benefit when quarantine is included.

5.16 To determine whether scenario-specific outcome distributions differed statistically, we conducted pairwise two-sided Welch t-tests, allowing for unequal variances, with statistical significance assessed at $\alpha = 0.05$. The results in Table 5 show that Scenarios 4 and 5 differed significantly from Scenarios 1–3 for both final cumulative recoveries and final cumulative deaths. In contrast, the comparison between Scenarios 4 and 5 was not

Table 4: Mean and standard deviation of final cumulative recovery and death outcomes across simulation scenarios.

Scenario	Final cumulative recovered	Final cumulative deaths
Scenario 1: No Restriction	8202.88 ± 465.64	1745.26 ± 445.24
Scenario 2: School From Home	7820.77 ± 506.65	1763.69 ± 413.81
Scenario 3: Essential Workers Only	7681.28 ± 570.50	1661.42 ± 427.83
Scenario 4: Split Essential Workers + Quarantine	5624.76 ± 991.37	1214.79 ± 371.25
Scenario 5: Split Essential Workers without Quarantine	5650.54 ± 957.98	1259.52 ± 428.35

Note: Values are reported as mean $\pm$ SD across 100 simulation runs. Final cumulative recovered and death counts were computed at the final simulation time point (day 120).

statistically significant for either final cumulative recovered or final cumulative deaths, suggesting that the rotational shift strategy without quarantine performs similarly to the rotational shift strategy with quarantine in terms of final cumulative outcomes. The Welch tests in Table 5 for inflection days further show that several scenario differences in curve timing are statistically supported, particularly for comparisons involving Scenarios 4 and 5. However, these inflection results are best interpreted together with the cumulative trajectories and final outcome summaries, since inflection timing alone does not fully characterize intervention effectiveness.

Table 5: Pairwise Welch t-test results for final outcomes and inflection days across simulation scenarios.

	Final recovered	Final deaths	Recovery inflection day	Death inflection day
Comparison	Welch p	Welch p	Welch p	Welch p
S1 vs S2	< 0.001	0.762	0.313	0.376
S1 vs S3	< 0.001	0.176	< 0.001	0.508
S1 vs S4	< 0.001	< 0.001	< 0.001	< 0.001
S1 vs S5	< 0.001	< 0.001	< 0.001	< 0.001
S2 vs S3	0.069	0.087	< 0.001	0.890
S2 vs S4	< 0.001	< 0.001	< 0.001	< 0.001
S2 vs S5	< 0.001	< 0.001	< 0.001	< 0.001
S3 vs S4	< 0.001	< 0.001	0.144	< 0.001
S3 vs S5	< 0.001	< 0.001	0.787	< 0.001
S4 vs S5	0.852	0.431	0.183	0.560

Note: Welch p values are from two-sided Welch t-tests comparing scenario distributions while allowing for unequal variances. S1 = Scenario 1, S2 = Scenario 2, S3 = Scenario 3, S4 = Scenario 4, and S5 = Scenario 5.

5.17 Finally, Table 6 summarizes the mutual information ranking of input effects on selected model outputs. The results show that $\tau_{\text{incubation}}$ is consistently among the strongest contributors to day-30 recovery and death outcomes, indicating that early epidemic dynamics are strongly influenced by the incubation period. For final cumulative outcomes, $P_{I \to R}$ and $P_{I \to D}$ also appear among the leading contributors across several scenarios, reflecting their direct roles in determining recovery and mortality accumulation over time. Seed index appears among the top-ranked inputs only in limited cases, suggesting that variation in the sampled epidemiological parameters generally contributes more strongly to output variability than population-seed effects alone. Overall, these statistical comparisons and sensitivity rankings provide additional quantitative support for the scenario-based findings observed in the cumulative recovery and death trajectories.

Table 6: Mutual information ranking of input effects on selected model outputs.

Scenario	Output	Rank 1	Rank 2	Rank 3
S1	Final recovered	$P_{I\to R}$ (0.567)	$P_{I\to D}$ (0.264)	$\tau_{\text{incubation}}$ (0.063)
S1	Final deaths	$P_{I\to R}$ (0.454)	$P_{I\to D}$ (0.331)	$\tau_{\text{incubation}}$ (0.036)
S1	Recovered at day 30	$\tau_{\text{incubation}}$ (0.854)	β (0.044)	$P_{I\to D}$ (0.031)
S1	Deaths at day 30	$\tau_{\text{incubation}}$ (0.906)	$P_{I\to D}$ (0.141)	β (0.055)
S2	Final recovered	$P_{I\to R}$ (0.207)	β (0.183)	$P_{E\to R}$ (0.165)
S2	Final deaths	$P_{I\to D}$ (0.263)	$P_{I\to R}$ (0.258)	$\tau_{\text{incubation}}$ (0.115)
S2	Recovered at day 30	$\tau_{\text{incubation}}$ (0.883)	β (0.175)	$P_{E\to I}$ (0.087)
S2	Deaths at day 30	$\tau_{\text{incubation}}$ (0.781)	$P_{I\to D}$ (0.249)	$P_{E\to R}$ (0.093)
S3	Final recovered	$\tau_{\text{incubation}}$ (0.364)	$P_{E\to R}$ (0.091)	$P_{I\to R}$ (0.085)
S3	Final deaths	$P_{I\to R}$ (0.387)	$P_{I\to D}$ (0.237)	$\tau_{\text{incubation}}$ (0.107)
S3	Recovered at day 30	$\tau_{\text{incubation}}$ (0.701)	β (0.098)	$P_{I\to D}$ (0.090)
S3	Deaths at day 30	$\tau_{\text{incubation}}$ (0.590)	$P_{I\to D}$ (0.107)	$P_{E\to I}$ (0.106)
S4	Final recovered	$\tau_{\text{incubation}}$ (0.297)	β (0.144)	$P_{I\to D}$ (0.003)
S4	Final deaths	$P_{I\to R}$ (0.130)	$P_{E\to I}$ (0.091)	$\tau_{\text{incubation}}$ (0.078)
S4	Recovered at day 30	$\tau_{\text{incubation}}$ (0.673)	β (0.083)	$P_{E\to I}$ (0.000)
S4	Deaths at day 30	$\tau_{\text{incubation}}$ (0.706)	β (0.102)	$P_{E\to R}$ (0.054)
S5	Final recovered	$\tau_{\text{incubation}}$ (0.294)	β (0.286)	Seed index (0.157)
S5	Final deaths	$P_{I\to R}$ (0.447)	$P_{E\to R}$ (0.217)	$\tau_{\text{incubation}}$ (0.170)
S5	Recovered at day 30	$\tau_{\text{incubation}}$ (0.815)	$P_{I\to R}$ (0.091)	β (0.058)
S5	Deaths at day 30	$\tau_{\text{incubation}}$ (0.515)	$P_{I\to D}$ (0.131)	β (0.072)

Note: Values in parentheses are mutual information scores. Larger values indicate stronger association between the input feature and the model output. Only the top three ranked inputs are shown for each scenario-output combination. Day 30 outcomes were included to assess input effects at an intermediate stage of the epidemic, while final recovered and final deaths summarize end-of-simulation outcomes. S1 = Scenario 1, S2 = Scenario 2, S3 = Scenario 3, S4 = Scenario 4, and S5 = Scenario 5.

Conclusion

6.1 Essential workers play a critical societal role but remain disproportionately vulnerable during pandemics due to sustained high-risk interactions in occupational and community environments. To address this significant challenge, this study presents SAFE-ABM (Structured Agent-Based Framework for Essential Workers), a novel stochastic agent-based modeling framework explicitly designed to evaluate targeted intervention strategies, with an emphasis on protecting essential workers. Unlike conventional ABM approaches that often assume homogeneous mixing or simplify workforce structures, our model captures structured social interactions across families, schools, and workplaces. This detailed representation allows precise evaluation of policies such as school closures, mobility restrictions, essential-worker-only restrictions, and targeted workforce rotation among essential workers, both with and without quarantine.

6.2 A significant innovation of our framework is its integration of prior-informed uncertainty quantification, enabling systematic exploration of stochastic variability and parameter uncertainty inherent in disease transmission dynamics. Through these prior-informed simulations, our analysis quantifies intervention effectiveness probabilistically while explicitly accounting for uncertainty across simulation runs. The scenario analysis demonstrates that rotational splitting of essential workers substantially reduces transmission, even in the absence of quarantine. This is shown by Scenario 5, where splitting essential workers into mutually exclusive rotating groups lowers infection burden, cumulative recoveries, and cumulative deaths compared with Scenarios 1–3, as shown in Figure 7 and Table 4. When rotational splitting is combined with quarantine enforcement for symptomatic active workers, as in Scenario 4, the model produces the strongest improvement, with the lowest cumulative deaths and the largest remaining susceptible population. These findings indicate that the rotational structure itself provides meaningful protection, while quarantine adds further benefit by reducing continued transmission from symptomatic active workers. In addition, the inflection-point analysis of cumulative recoveries and deaths in Table 3 shows that structured interventions alter the timing and stability of epidemic slowing, supporting a more controlled epidemic trajectory under the rotational shift strategies.

6.3 The formal statistical comparisons further support the scenario-based findings. The two-sided Welch tests in Table 5 show that Scenarios 4 and 5 differ significantly from Scenarios 1–3 for final cumulative recoveries and final cumulative deaths, confirming that the rotational shift strategies produce measurable reductions in epidemic outcomes. In contrast, Scenario 4 and Scenario 5 are not statistically different from each other for final cumulative recoveries or deaths, indicating that rotational splitting without quarantine performs comparably to rotational splitting with quarantine in the final outcome distributions, although Scenario 4 remains descriptively lower for deaths. The inflection-day comparisons also show statistically supported differences in curve timing for several scenario pairs, but these timing results are best interpreted alongside the cumulative trajectories and final outcome summaries rather than as a standalone ranking of intervention effectiveness. The mutual information results in Table 6 provide additional sensitivity evidence by identifying the inputs most strongly associated with variation in model outputs. $\tau_{\text{incubation}}$ is consistently influential for day-30 recovery and death outcomes, while $P_{I\to R}$ and $P_{I\to D}$ contribute strongly to several final cumulative outcomes. Seed index appears among the leading inputs only in limited cases, suggesting that parameter uncertainty contributes more consistently to output variability than population-seed effects alone.

6.4 While exploratory simulations confirm epidemiological plausibility, our current framework has not yet undergone calibration to empirical data. Our immediate focus was methodological rigor, structural validation, scenario analysis, statistical comparison, and sensitivity ranking rather than predictive accuracy. Formal model calibration to real-world epidemiological data thus constitutes an essential next step. Future research will integrate empirical datasets, refining parameter estimates to enhance predictive reliability and facilitate real-time evaluation of intervention strategies. In applying SAFE-ABM to a specific pandemic, posterior parameter distributions from calibrated epidemic models may also be incorporated so that parameter draws reflect the target disease and population more directly. However, this transfer requires careful consideration because infection-rate parameters estimated from population-level models may depend on population size, contact structure, and scaling assumptions. Appropriate transformations may therefore be needed before such posterior distributions can be used within an agent-based framework. Future work may also consider Bayesian parameter estimation directly within SAFE-ABM when disease parameters can reasonably be treated as stable over the study period. In practice, however, effective transmission patterns often change over time as governments and public health agencies introduce interventions that alter contact behavior and infection risk. Capturing these time-varying effects would require explicit information on the timing and nature of policy changes. Ultimately, our modeling framework, combining detailed agent interactions with prior-informed uncertainty quantification, formal scenario comparison, and input-output sensitivity ranking, provides a foundation for data-informed intervention planning. Our approach offers valuable guidance for equitable and effective public health responses in current and future pandemics by explicitly emphasizing the critical yet often overlooked role of essential workers.

Model Documentation

The SAFE-ABM model code and reproducibility materials are available through a private CoMSES Net review archive: `https://www.comses.net/codebase-release/3200bb24-c060-47c5-a5b5-0d45415c4279/`. The archive contains the source code, scenario-specific simulation scripts, comparison-analysis scripts, saved derived outputs, README files, software requirements, and reproducibility notes needed to review and reproduce the main results.

The archive is organized by intervention scenario, including Scenario 1 (no restrictions), Scenario 2 (school-from-home), Scenario 3 (essential-workers-only), Scenario 4 (rotational shifts among essential workers with quarantine), and Scenario 5 (rotational shifts among essential workers without quarantine). It also includes scripts used to reproduce the median epidemic trajectories, predictive-uncertainty plots, cumulative recovery and death summaries, inflection-point summaries, Welch-test comparisons, and mutual-information sensitivity rankings.

To preserve the anonymised review process, the uploaded model files and documentation have been anonymised, and the archive remains private during review. The archive will be made public and updated with full citation and author information upon article publication.

Acknowledgements

The authors gratefully acknowledge Virginia Commonwealth University in Qatar and Qatar Foundation for supporting this work through the VCU Qatar Mathematical Data Sciences Lab.

References

Amona, E. B., Ghanam, R. A., Boone, E. L., Sahoo, I. & Abu-Raddad, L. J. (2022). Incorporating interventions to an extended SEIRD model with vaccination: Application to COVID-19 in Qatar. arXiv preprint. arXiv:2204.10961

Arduin, H., Domenech de Cellès, M., Guillemot, D., Watier, L. & Opatowski, L. (2017). An agent-based model simulation of influenza interactions at the host level: Insight into the influenza-related burden of pneumococcal infections. *BMC infectious diseases*, *17*, 1–12

Bi, Q., Lessler, J., Eckerle, I., Lauer, S. A., Kaiser, L., Vuilleumier, N., Cummings, D. A., Flahault, A., Petrovic, D., Guessous, I. et al. (2021). Insights into household transmission of sars-cov-2 from a population-based serological survey. *Nature communications*, *12*(1), 3643

Bielicki, J. A., Duval, X., Gobat, N., Goossens, H., Koopmans, M., Tacconelli, E. & van der Werf, S. (2020). Monitoring approaches for health-care workers during the COVID-19 pandemic. *The Lancet Infectious Diseases*, *20*(10), e261–e267

Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, *99*(suppl$_3$), $7280 - -7287$

Centers for Disease Control and Prevention (2020a). Guidance for healthcare workers during the covid-19 pandemic. Accessed: 2023-11-12. `https://www.cdc.gov/coronavirus/2019-ncov/hcp/guidance-risk-assesment-hcp.html`

Centers for Disease Control and Prevention (2020b). Interim guidance for businesses and employers responding to coronavirus disease 2019 (COVID-19), may 2020. `https://www.cdc.gov/coronavirus/2019-ncov/community/guidance-business-response.html`

Centers for Disease Control and Prevention (2022). A child's health is the public's health. Https://www.cdc.gov/childrenindisasters/features/children-public-health.html

Centers for Disease Control and Prevention (2024a). Clinical overview of MERS. Https://www.cdc.gov/mers/hcp/clinical-overview/index.html

Centers for Disease Control and Prevention (2024b). Everyday actions for schools to prevent and control the spread of infections. Https://www.cdc.gov/orr/school-preparedness/infection-prevention/actions.html

Centers for Disease Control and Prevention (2025). Influenza. CDC Yellow Book 2026. Available at: `https://www.cdc.gov/yellow-book/hcp/travel-associated-infections-diseases/influenza.html`

Centers for Disease Control and Prevention (2026). Ebola disease basics. Https://www.cdc.gov/ebola/about/index.html

Chen, P., Fan, W. & Guo, X. (2021). A hybrid simulation model to study the impact of combined interventions on ebola epidemic. *PLoS One*, *16*(7), e0254044

Chung, P.-C., Chen, K. J., Chang, H.-M., Chan, T.-C. et al. (2023). Evaluating the effectiveness of school closure in COVID-19-related syndromes from community-based syndromic surveillance: Longitudinal observational study. *Interactive Journal of Medical Research*, *12*(1), e44606

Cimolai, N. (2021). In pursuit of the right tail for the COVID-19 incubation period. *Public Health*, *194*, 149–155

Deng, C., Tao, H. & Ye, Z. (2008). Agent-based modeling to simulate the dengue spread. Geoinformatics 2008 and Joint Conference on GIS and Built Environment: Geo-Simulation and Virtual GIS Environments

Dyer, J., Bishop, N., Felekis, Y., Zennaro, F. M., Calinescu, A., Damoulas, T. & Wooldridge, M. (2023). Interventionally consistent surrogates for agent-based simulators. arXiv preprint. arXiv:2312.11158

El Jaouhari, M., Edjoc, R., Waddell, L., Houston, P., Atchessi, N., Striha, M. & Bonti-Ankomah, S. (2021). Impact of school closures and re-openings on COVID-19 transmission. *Canada Communicable Disease Report*, *47*(12), 515

Grefenstette, J. J., Brown, S. T., Rosenfeld, R., DePasse, J., Stone, N. T., Cooley, P. C., Wheaton, W. D., Fyshe, A., Galloway, D. D., Sriram, A., Guclu, H., Abraham, T. & Burke, D. S. (2013). FRED (A Framework for Reconstructing Epidemic Dynamics): An open-source software system for modeling infectious diseases and control strategies using census-based populations. *BMC public health*, *13*, 1–14

Hack, J. J. & Papka, M. E. (2020). The US high-performance computing consortium in the fight against COVID-19. *Computing in Science & Engineering*, *22*(6), 75–80

Hinch, R., Probert, W. J., Nurtay, A., Kendall, M., Wymant, C., Hall, M., Lythgoe, K., Bulas Cruz, A., Zhao, L., Stewart, A., Ferretti, L., Montero, D., Warren, J., Mather, N., Abueg, M., Wu, N., Legat, O., Bentley, K., Mead, T., Van-Vuuren, K., Feldner-Busztin, D., Ristori, T., Finkelstein, A., Bonsall, D. G., Abeler-Dörner, L. & Fraser, C. (2021). OpenABM-Covid19 — An agent-based model for non-pharmaceutical interventions against covid-19 including contact tracing. *PLoS Computational Biology*, *17*(7), e1009146

Johnson, B. R. (2010). Division of labor in honeybees: Form, function, and proximate mechanisms. *Behavioral Ecology and Sociobiology*, *64*, 305–316

Karlsson, U. & Fraenkel, C.-J. (2020). Covid-19: Risks to healthcare workers and their families

Kerr, C. C., Stuart, R. M., Mistry, D., Abeysuriya, R. G., Rosenfeld, K., Hart, G. R., Núñez, R. C., Cohen, J. A., Selvaraj, P., Hagedorn, B., George, L., Jastrzębski, M., Izzo, A. S., Fowler, S., Palmer, A., Delport, D., Scott, N., Kelly, S. L., Bennette, C. S., Wagner, B. G., Chang, S. T., Oron, A. P., Wenger, E. A., Panovska-Griffiths, J., Famulare, M. & Klein, D. J. (2021). Covasim: An agent-based model of COVID-19 dynamics and interventions. *PLOS Computational Biology*, *17*(7), e1009149

Khunti, K., Griffiths, A., Majeed, A., Nagpaul, C. & Rao, M. (2021). Assessing risk for healthcare workers during the Covid-19 pandemic. *BMJ*, *372*

Kraskov, A., Stögbauer, H. & Grassberger, P. (2004). Estimating mutual information. *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics*, *69*(6), 066138

Liu, X., Hu, Z., Deng, X. & Kuhlman, C. J. (2022). A Bayesian uncertainty quantification approach for agent-based modeling of networked anagram games. 2022 Winter Simulation Conference (WSC)

Lorenzo, D. & Carrisi, C. (2020). COVID-19 exposure risk for family members of healthcare workers: An observational study. *International Journal of Infectious Diseases*, *98*, 287–289

Lorig, F., Johansson, E. & Davidsson, P. (2021). Agent-based social simulation of the COVID-19 pandemic: A systematic review. *Journal of Artificial Societies and Social Simulation*, *24*(3), 5

MacLehose, R. F. & Hamra, G. B. (2014). Applications of Bayesian methods to epidemiologic research. *Current Epidemiology Reports*, *1*, 103–109

Madewell, Z. J., Yang, Y., Longini, I. M., Halloran, M. E. & Dean, N. E. (2020). Household transmission of SARS-CoV-2: A systematic review and meta-analysis. *JAMA Network Open*, *3*(12), e2031756–e2031756

Madewell, Z. J., Yang, Y., Longini, I. M., Halloran, M. E. & Dean, N. E. (2021). Factors associated with household transmission of SARS-CoV-2: An updated systematic review and meta-analysis. *JAMA Network Open*, *4*(8), e2122240–e2122240

Mahmood, I., Jahan, M., Groen, D., Javed, A. & Shafait, F. (2020). An agent-based simulation of the spread of dengue fever. Computational Science - ICCS 2020: 20th International Conference, Amsterdam, The Netherlands, June 3-5, 2020, Proceedings, Part III 20

Martinez, E. Z. & Achcar, J. A. (2014). Trends in epidemiology in the 21st century: Time to adopt Bayesian methods. *Cadernos de Saúde Pública*, *30*(4), 703–714

McCulloch, J., Ge, J., Ward, J. A., Heppenstall, A., Polhill, J. G. & Malleson, N. (2022). Calibrating agent-based models using uncertainty quantification methods. *Journal of Artificial Societies and Social Simulation*, *25*(2), 1

McNeill, M. (2022). Extraordinary impacts on the healthcare workforce: Covid-19 and aging. *Delaware Journal of Public Health*, *8*(5), 164

Miksch, F., Jahn, B., Espinosa, K. J., Chhatwal, J., Siebert, U. & Popper, N. (2019). Why should we apply ABM for decision analysis for infectious diseases? An example for dengue interventions. *PLoS One*, *14*(8), e0221564

Milligan, W. R., Fuller, Z. L., Agarwal, I., Eisen, M. B., Przeworski, M. & Sella, G. (2021). Impact of essential workers in the context of social distancing for epidemic control. *PLoS One*, *16*(8), e0255680

Naug, D. (2008). Structure of the social network and its influence on transmission dynamics in a honeybee colony. *Behavioral Ecology and Sociobiology*, *62*, 1719–1725

Nguyen, C., Tran, D., Nguyen, A. & Nguyen, N. (2021). The effects of perceived risks on food purchase intention: The case study of online shopping channels during COVID-19 pandemic in Vietnam. *Journal of Distribution Science*, (pp. 19–9)

Niazi, M. & Hussain, A. (2011). Agent-based computing from multi-agent systems to agent-based models: A visual survey. *Scientometrics*, *89*(2), 479–499

Ozik, J., Wozniak, J. M., Collier, N., Macal, C. M. & Binois, M. (2021). A population data-driven workflow for COVID-19 modeling and learning. *The International Journal of High Performance Computing Applications*, *35*(5), 483–499

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Müller, A., Nothman, J., Louppe, G., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830

Phadke, I., McKee, A., Conway, J. & Shea, K. (2021). Analysing how changes in the health status of healthcare workers affects epidemic outcomes. *Epidemiology & Infection*, *149*, e42

Pinotti, F., Lourenço, J., Gupta, S., Das Gupta, S., Henning, J., Blake, D., Tomley, F., Barnett, T., Pfeiffer, D., Hoque, M. A. & Fournié, G. (2024). EPINEST, an agent-based model to simulate epidemic dynamics in large-scale poultry production and distribution networks. *PLoS Computational Biology*, *20*(2), e1011375

Rahimi, E., Shabanpour, R., Shamshiripour, A. & Mohammadian, A. K. (2021). Perceived risk of using shared mobility services during the COVID-19 pandemic. *Transportation Research Part F: Traffic Psychology and Behaviour*, *81*, 271–281

Romero-Michel, J. C., Mokay-Ramírez, K. A., Delgado-Machuca, M., Delgado-Enciso, J., Aurelien-Cabezas, N. S., Tiburcio-Jimenez, D., Meza-Robles, C., Delgado-Enciso, O. G., Esquivel, J. G., Zaizar-Fregoso, S. A. et al. (2021). Health and economic measures in response to the COVID-19 pandemic-effect on street vendors. *The Journal of Infection in Developing Countries*, *15*(02), 198–203

Sah, P., Mann, J. & Bansal, S. (2018). Disease implications of animal social network structure: A synthesis across social systems. *Journal of Animal Ecology*, *87*(3), 546–558

Seeley, T. D. (1982). Adaptive significance of the age polyethism schedule in honeybee colonies. *Behavioral Ecology and Sociobiology*, *11*, 287–293

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(3), 379–423

Sim, M. R. (2020). The COVID-19 pandemic: Major risks to healthcare and other workers on the front line. *Occupational and Environmental Medicine*, *77*(5), 281–282

Tognotti, E. (2013). Lessons from the history of quarantine, from plague to influenza A. *Emerging Infectious Diseases*, *19*(2), 254

Vaivada, T., Sharma, N., Das, J. K., Salam, R. A., Lassi, Z. S. & Bhutta, Z. A. (2022). Interventions for health and well-being in school-aged children and adolescents: A way forward. *Pediatrics*, *149*(Supplement 6)

van Zoonen, W. & Ter Hoeven, C. L. (2022). Disruptions and general distress for essential and nonessential employees during the COVID-19 pandemic. *Journal of Business and Psychology*, *37*(2), 443–458

Wang, Z., Ma, W., Zheng, X., Wu, G. & Zhang, R. (2020). Household transmission of SARS-CoV-2. *Journal of Infection*, *81*(1), 179–182

Wei, C.-F., Lan, F.-Y., Hsu, Y.-T., Lowery, N., Dibona, L., Akkeh, R., Kales, S. N. & Yang, J. (2022). Risk of SARS-CoV-2 infection among essential workers in a community-based cohort in the United States. *Frontiers in Public Health*, *10*, 878208

Welch, B. L. (1947). The generalization of 'STUDENT'S' problem when several different population varlances are involved. *Biometrika*, *34*(1–2), 28–35

Wilson-Rich, N., Spivak, M., Fefferman, N. H. & Starks, P. T. (2009). Genetic, individual, and group facilitation of disease resistance in insect societies. *Annual Review of Entomology*, *54*(1), 405–423

World Health Organization (2020). Infection prevention and control during health care when COVID-19 is suspected: Interim guidance, 19 March 2020. `https://apps.who.int/iris/handle/10665/331495`

World Health Organization (2026). Ebola disease caused by Bundibugyo virus - Democratic Republic of the Congo. Https://www.who.int/emergencies/disease-outbreak-news/item/2026-DON603

Wu, J., Huang, Y., Tu, C., Bi, C., Chen, Z., Luo, L., Huang, M., Chen, M., Tan, C., Wang, Z. et al. (2020). Household transmission of Sars-Cov-2, Zhuhai, China, 2020. *Clinical Infectious Diseases*, *71*(16), 2099–2108