engineering proceedings

MDPI

Proceeding Paper

Extending a Multilayer Agent-Based Model for COVID-19 Transmission with Public Transport and Church Attendance Layers: A Case Study for Patras, Greece †

Vasileios Thomopoulos * and Panagiotis Tzoudas

Department of Computer Engineering and Informatics, University of Patras, 26504 Patras, Greece; tzoudas@ceid.upatras.gr

* Correspondence: vthomopoulos@upatras.gr

† Presented at the 2nd International Conference on Engineering, Technology, and Applied Science Innovations (ICETASI 2026), Zakynthos, Greece, 24–26 September 2026.

Abstract

Epidemiological agent-based models (ABMs) offer a detailed, behavior-based alternative to aggregate compartmental approaches by representing heterogeneous populations and enabling the simulation of targeted non-pharmaceutical interventions. This study extends a previously developed synthetic-population ABM for COVID-19 transmission in the central district of Patras, Greece, by introducing two new interaction layers: a public transport layer and a church attendance layer. In addition, the school assignment mechanism is enhanced to reflect the age-appropriate structure of local educational institutions. The model follows a SEIRD compartmental structure and is implemented using the Mesa framework for Python 3.14. The public transport layer captures 13% of agents who use one of ten urban bus lines on weekdays and Saturdays, whereas the church attendance layer assigns 5% of the population to one of twenty parishes active exclusively on Sundays. Household transmission parameters are grounded in empirical household secondary attack rate estimates from the COVID-19 literature, yielding a simulated within-household secondary attack rate of approximately 22% and an estimated basic reproduction number $R_0 \approx 3.3$. Eight intervention scenarios are evaluated across ten stochastic replications each. Moderate and strict lockdown scenarios (Scenarios 7–8) reliably reduce the effective reproduction number below 1, limiting cumulative incidence to under 0.3%. Mild lockdown (Scenario 6), in contrast, exhibits near-critical stochastic behavior, with a mean incidence of 1.5%. Individual Scenario 6 replications range from near-extinction to a 7.2% attack rate. The outcome distribution is heavy-tailed, so the mean understates the risk of any single implementation near the epidemic threshold. Workplace closure and masks-and-distancing scenarios are found to operate near the epidemic threshold, producing bimodal outcomes: in 30% and 10% of replications, respectively, the epidemic goes extinct before becoming established. In the remaining runs, attack rates of approximately 40% and 58% are observed. The case-fatality rate across epidemic scenarios ranges from 0.85% to 1.18%. The age-targeted intervention selectively protects the highest-mortality subgroup, reducing cumulative mortality by 27.5%. As the model has not yet been validated against empirical Greek epidemiological time series, these findings should be read as an exploratory comparison of intervention scenarios rather than as calibrated predictions.

Keywords: agent-based modeling; COVID-19; SEIRD model; synthetic population; multilayer networks; epidemic simulation; Patras; reproduction number

Check for updates

Academic Editor: Konstantinos Poirazidis

Published: 16 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Infectious disease modeling has long been a core instrument for public health planning, and the COVID-19 pandemic renewed and intensified this role globally. In the years since the initial outbreak, a substantial body of literature has revisited foundational epidemiological concepts, extending classical compartmental models with behavioral, spatial, and demographic dimensions in order to improve fidelity to observed transmission dynamics. Despite this growing international effort, the application of agent-based epidemic simulation to Greek population contexts remains limited, particularly at the level of detailed multilayer interaction modeling.

Two broad paradigms dominate computational epidemiology. Ordinary differential equation (ODE) models represent population compartments as continuous quantities that evolve over time and are computationally efficient, but they treat the population as homogeneous and cannot capture individual-level behavioral heterogeneity. Agent-based models (ABMs), by contrast, represent each individual as a discrete entity with its own attributes and behavioral rules; population-level dynamics emerge from the aggregation of individual interactions rather than being imposed analytically [1,2]. This makes ABMs particularly suitable for simulating non-pharmaceutical interventions (NPIs) that operate through behavior change, such as school closures, mobility restrictions, or restrictions on public gatherings.

The two interaction settings added in this study are motivated by evidence from the pandemic period. Public transport was identified as a setting of elevated transmission risk: the concentration of individuals in enclosed vehicles with limited ventilation created conditions favorable for airborne spread [3]. Religious gatherings presented a distinct challenge: systematic review evidence confirms that places of worship were associated with COVID-19 transmission events across multiple countries and denominations [4]. In Greece, restrictions on Orthodox Sunday services were socially contentious and created conditions of partial non-compliance with public health guidance.

A prior study [2] developed an ABM for epidemic dynamics in Greece, grounded in a synthetic population of 19,600 agents calibrated to census data from ELSTAT (the Hellenic Statistical Authority) for the central district of the Municipality of Patras. That model incorporated five daily interaction layers: household, workplace, school, peer, and random community contacts. Both public transport and religious attendance were absent. This study addresses those gaps by extending [2] with a public transport layer and a church attendance layer, both constructed from locally grounded data for Patras. The school assignment mechanism is also refined to produce an age-appropriate distribution of students across educational institution types. To the best of our knowledge, no prior ABM study of COVID-19 transmission in Greece has incorporated both interaction settings within a multilayer synthetic-population framework.

2. Related Work

Agent-based modeling has been applied extensively to COVID-19 epidemic simulation at city and regional scales. Singh et al. [5] developed an ABM for the Madrid metropolitan area that incorporated age-stratified hospitalization and fatality probability structures, providing a methodological reference for demographic mortality calibration. Nitzsche and Simm [6] evaluated lockdown scenarios of varying severity using an abstract city ABM, demonstrating a monotonic relationship between restriction intensity and epidemic suppression. Wan et al. [7] showed that higher-order interactions in multilayer contact networks shape epidemic trajectories in ways that single-layer representations cannot capture, motivating the treatment of each contact setting as a structurally distinct network layer.

https://doi.org/10.3390/engproc2026167021

Luo et al. [3] review COVID-19 transmission risk and control in land public transport, identifying crowding, ventilation, and contact duration as primary determinants of infection probability on buses and trains. Goumballa et al. [4] provide a systematic review of SARS-CoV-2 transmission at religious mass gatherings, confirming transmission events across multiple countries and identifying gathering density and preventive-measure compliance as risk factors. Von Hoene et al. [8] present an approach to incorporating health-related behavioral characteristics into spatially explicit synthetic populations for use in ABMs. Madewell et al. [9] conducted a meta-analysis of household secondary attack rates across 54 COVID-19 studies, reporting a pooled estimate of 16.6% (95% CI: 13.4–19.6%), which provides an empirical anchor for household transmission parameter calibration. Jing et al. [10] reported a household secondary attack rate of 17.1% (95% CI: 13.3–21.8%) from a contact-tracing study in Guangzhou, China. The model of Thomopoulos and Tsichlas [2] is the direct baseline for the present study.

3. Model Description

Figure 1 summarizes the model pipeline described in detail in this section: synthetic population generation (Section 3.1), the seven interaction layers and their day-of-week scheduling (Sections 3.2–3.4 and 3.6), transmission and mortality parameter calibration (Section 3.5), and the daily state-transition loop that is executed once per simulated day for each of the eight scenarios × ten stochastic replications (Section 4).

3.1. Baseline Architecture

The model is implemented in Mesa for Python [11] (Mesa version 3.5.1, Project Mesa, open source; Python version 3.14, Python Software Foundation, Wilmington, DE, USA) and operates on the 19,600-agent synthetic population representing the central district of the Municipality of Patras [2]. Each agent carries demographic attributes including age, household membership, and employment or school status. Disease dynamics follow a SEIRD compartmental structure with five states: Susceptible (S), Exposed (E), Infected (I), Recovered (R), and Deceased (D). The Deceased state extends the prior SEIR formulation [2] using an age-stratified mortality mechanism adapted from [5] and scaled to Greek conditions as described in Section 3.5.

In this discrete-time Markovian formulation, transition parameters σ and γ are daily transition probabilities: $\sigma = 0.2$ means each Exposed agent has a 20% daily probability of becoming Infectious; $\gamma = 0.1$ means each Infected agent has a 10% daily probability of recovering. The ten-day mean infectious period is consistent with reported temporal dynamics of SARS-CoV-2 viral shedding [12]. Both parameters are held constant across all scenarios; scenario variation is expressed exclusively through adjustments to per-layer transmission rates (beta). Table 1 lists the baseline transmission and transition parameters.

Table 1. Baseline transmission and transition parameters.

Parameter	Value	Interpretation
FAMILY (β)	0.025	Household daily transmission probability (SAR ≈ 22%)
SCHOOL (β)	0.025	School daily transmission probability
WORK (β)	0.012	Workplace daily transmission probability
CHURCH (β)	0.005	Church attendance daily transmission probability
PUBTPT (β)	0.0005	Public transport daily transmission probability
PEERS (β)	0.001	Peer contact daily transmission probability
RANDOM (β)	0.0003	Random community daily transmission probability
σ (Sigma)	0.200	Daily probability: Exposed → Infectious (mean incubation 5 days)
γ (Gamma)	0.100	Daily probability: Infectious → Recovered (mean infectious period 10 days)
HOSP_FACTOR	0.75	Scale factor on Singh et al. [5] hospitalization probabilities
DEATH_FACTOR	0.75	Scale factor on Singh et al. [5] fatality probabilities

Note: σ and γ are daily transition probabilities under the Markovian assumption (see Section 3.1).

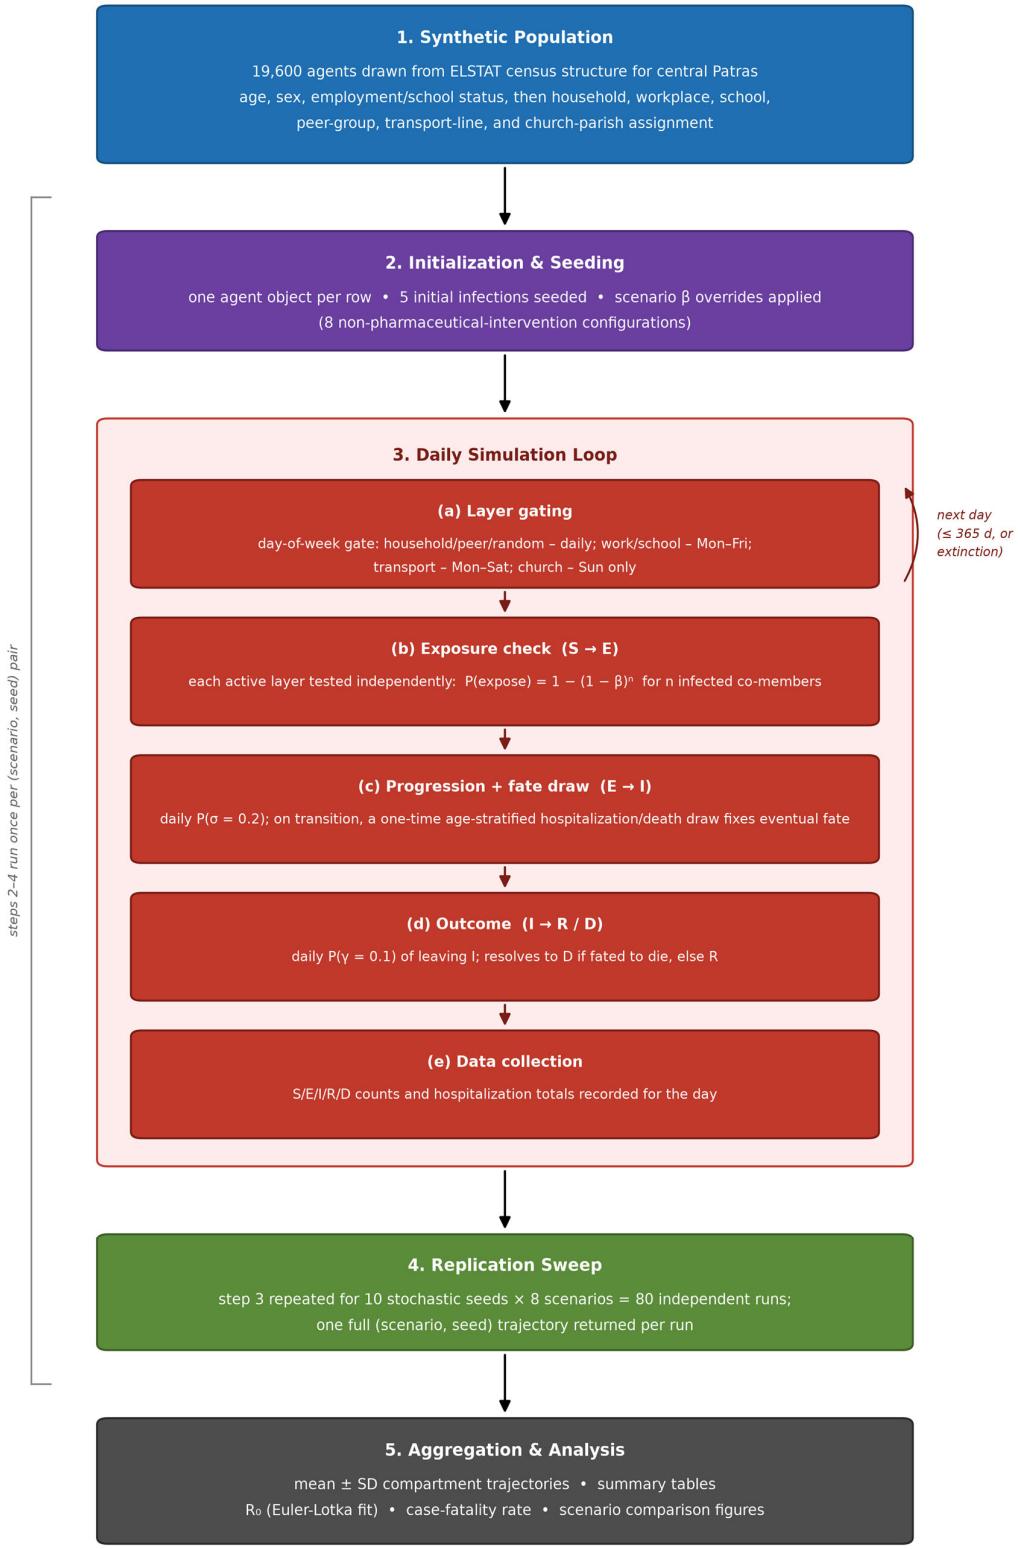

Figure 1. Model pipeline, from synthetic population generation through the daily simulation loop to results aggregation. The three stages numbered 2 to 4 in the figure (initialization, the daily loop, and the replication sweep) run once for every (scenario, seed) pair. The full sweep spans 8 scenarios × 10 seeds, for 80 runs total.

3.2. Public Transport Layer

Urban public transport is a documented setting for airborne disease transmission, particularly under conditions of high passenger density and limited ventilation [3]. The layer

is constructed from data obtained through direct contact with the Patras Urban Transport Authority. Three parameters ground its design: (i) urban buses are the predominant form of public transport in the city; (ii) 13% of the simulated population uses public transport as its primary commuting mode; and (iii) ten distinct bus lines operate across the central district, each with a documented ridership share. Line assignment is drawn probabilistically from a ridership-weighted distribution.

Agents aged 12 and under and retired agents are excluded. The layer is active Monday through Saturday, since bus usage in Patras remains substantial for retail and leisure trips on Saturdays. It is deactivated on Sundays, when the church attendance layer becomes active. As a simplifying assumption, the two layers are mutually exclusive on Sundays.

3.3. Church Attendance Layer

Religious attendance is incorporated because of its social importance in the Greek context and its epidemiological relevance during the COVID-19 period. Systematic review evidence confirms that SARS-CoV-2 transmission occurred in religious gathering settings across multiple countries [4], and in Greece, restrictions on Orthodox services encountered organized resistance, creating partial compliance patterns. No published, Patras-specific dataset of Sunday attendance exists. A Pew Research Center survey of religious practice in Western Europe provides broader regional context on attendance patterns, but Greece is not among the countries it covers, and the survey does not report a Greek-specific figure [13]. In the absence of a city-level figure, the authors set the weekly attendance rate for the central district at 5% based on their own estimate, not a value drawn from a cited data source. This figure is explicitly a modeling assumption, not an observed city-level parameter, and its sensitivity is examined directly in Section 5.3.

The central district of Patras contains 20 churches. Agents with ReligionStatus = 1 (comprising 5% of the population) are randomly assigned to one of the 20 churches and assumed to attend the same parish every Sunday. The church layer is active exclusively on Sundays. Attendee assignment in the baseline model is uniform across age groups, not age-weighted. The widely observed tendency for church attendance to skew older is not built into Scenarios 1–3 and 5–8 as a structural feature. Scenario 4 applies this tendency separately, as an explicit, targeted intervention that halves β_CHURCH for agents aged 60 and above. This is not a baseline demographic assumption. The distinction, and the absence of a Patras-specific age-attendance dataset to calibrate it against, is flagged as a limitation in Section 6.

3.4. School Assignment Enhancement

The prior model [2] assigned school identifiers using a uniform random draw across all school types, producing demographically inconsistent configurations. The enhancement constructs a type-stratified school dictionary based on the actual distribution in the central district: 8 lyceums (6 general, 2 vocational), 7 lower secondary schools, 21 primary schools, 18 kindergartens, and 10 nurseries. Identifiers are assigned within age-appropriate institution types so that within-school contacts occur only among agents of comparable ages.

3.5. Transmission Parameter Calibration

Household transmission. The prior model used a daily household transmission probability (β_FAMILY) of 0.400, implying a within-household secondary attack rate (SAR) of approximately 99% over a ten-day infectious period. This substantially exceeds empirically measured COVID-19 household SARs. Madewell et al. [9] meta-analyzed 54 studies and reported a pooled household SAR of 16.6% (95% CI: 13.4–19.6%). Jing et al. [10] reported 17.1% (95% CI: 13.3–21.8%) in a Guangzhou contact-tracing study. We therefore

recalibrate β_FAMILY = 0.025, which yields a within-household SAR of approximately 22% ($1 - (1 - 0.025)^{10}$). The modeled value slightly exceeds the meta-analytic central estimate because infected agents in the model are not assumed to self-isolate within the household—an appropriate first-wave, pre-guidance scenario.

Non-household transmission. The remaining layer-specific beta values are calibrated analytically using the relationship between per-contact transmission probability, mean group size, and effective active days per infectious period. For a layer with mean group size G and T_{eff} active days per 10-day infection, the expected secondary infections per primary case within that layer is $(G - 1) \times \beta \times T_{eff}$. The per-layer β values are selected so that aggregate transmission is consistent with the published SARS-CoV-2 ancestral-strain first-wave basic reproduction number range of $R_0 \in [2.5, 4.0]$. The resulting epidemic growth is then evaluated separately, as described below. The calibrated baseline values are: β_SCHOOL = 0.025 (mean group size 50, active 7.14 days), β_WORK = 0.012 (mean group size 23, active 7.14 days), β_CHURCH = 0.005 (mean group size 49, active 1.43 days), β_PUBTPT = 0.0005 (mean group size 255, active 8.57 days), β_PEERS = 0.001 (mean group size 39, active 10 days), and β_RANDOM = 0.0003 (prevalence-scaled). The $(G - 1) \times \beta \times T_{eff}$ expression gives the secondary infections a primary case would produce if it belonged to that layer. Because each agent belongs to only a subset of layers, every layer term is weighted by the share of the population that participates in it: the household and peer layers include all agents, workplaces about 61%, schools about 17%, public transport about 13%, and the church layer 5%. After weighting, the non-household layers together contribute an order of 3 expected secondary infections per primary case. The unweighted sum of the five-layer terms is about 12, but that figure assumes every agent belongs to every layer. The weighted value is consistent with the Euler-Lotka estimate of $R_0 \approx 3.3$ and with the approximate 0.34 household and 2.96 non-household split used in Sections 5.2 and 5.4.

R_0 estimation. Following calibration, R_0 is estimated from the early-phase exponential growth of I(t) in Scenario 1 using the SEIR Euler-Lotka relationship: $R_0 = (1 + r/\sigma)(1 + r/\gamma)$, where r is the daily exponential growth rate fitted to log(I) over simulation days 5–25. The estimate yields $R_0 \approx 3.3$, consistent with published ancestral-strain estimates for a pre-intervention scenario. This approximation assumes that the early growth phase is dominated by the dominant eigenvalue of the linearized system and that heterogeneity in contact structure does not substantially affect the growth rate during the first 20 days of the epidemic. Formal estimation via next-generation matrix (NGM) methods, which construct a matrix of expected secondary infections across layer- and age-structured subgroups, following Diekmann, Heesterbeek and Roberts [14], is deferred to future work. As a consistency check, we note that for $R_0 = 3.33$, the SIR final-size equation predicts a population attack rate of approximately 96%, whereas the simulated baseline attack rate is 87.1%—a gap of approximately 9 percentage points. This discrepancy is expected in a clustered contact-network model: contact groups (school classes, workplaces, peer groups) reach local saturation before the global epidemic peaks, reducing the effective spreading efficiency relative to the homogeneous-mixing assumption underlying the Euler-Lotka formula. The implied R_0 from the observed final size is approximately 2.3; the Euler-Lotka estimate of 3.3 reflects early growth in high-density clusters and should therefore be interpreted as an upper bound rather than a population-average reproductive number.

Mortality calibration. Age-stratified raw hospitalization and fatality probabilities are drawn from Singh et al. [5] (Table 2 of the [5] manuscript). Each agent's mortality fate is determined once at the Exposed → Infectious transition: the agent is flagged as a future fatality if a random draw falls below h_ag × HOSP_FACTOR × d_ag × DEATH_FACTOR, where h_ag and d_ag are the Singh et al. raw probabilities for the agent's age group ag. With the recalibrated β values shifting infection risk toward working-age and younger agents, the

infection-weighted average P(die|infected) fell below the target Greek CFR of approximately 1.20% (EODY, the Greek National Public Health Organization, March 2020–May 2021) [15]. HOSP_FACTOR and DEATH_FACTOR were therefore reset to 0.75 (product 0.5625), producing a simulated Scenario 1 CFR of approximately 1.10% and a range of 1.10–1.18% across scenarios without age-targeted modifications (see Section 5). Throughout this manuscript, simulated CFR is computed as deaths divided by total simulated infections, since the model does not represent case detection or testing. It is therefore not directly comparable to a reported-case CFR such as EODY's [15], which divides deaths by confirmed cases.

3.6. Day-Wise Layer Scheduling

A day-of-week counter in the step() method activates or deactivates each layer according to the following weekly schedule:

- School layer: active Monday through Friday; inactive Saturday and Sunday.
- Work layer: active Monday through Friday; inactive Saturday and Sunday.
- Public transport layer: active Monday through Saturday; inactive Sunday.
- Church layer: active Sunday only; inactive all other days.
- Household, peer, and random-contact layers: active all days.

This mechanism ensures that church and public transport contacts are evaluated in temporally realistic contexts and that no layer is artificially active on days when the corresponding activity does not occur.

4. Simulation Scenarios

Eight scenarios are evaluated, each representing a distinct NPI configuration. Each scenario is run ten times with different random seeds, and results are averaged across runs; standard deviations are computed across the ten replications. Table 2 presents the per-layer beta values.

Table 2. Per-layer beta values across simulation scenarios. FAM = FAMILY, SCH = SCHOOL, CHU = CHURCH, PT = PUBTPT, RND = RANDOM, PEER = PEERS.

Scenario	Description	FAM	WORK	SCH	CHU	PT	RND	PEER
1	Baseline	0.025	0.0120	0.0250	0.00500	0.000500	0.000300	0.001000
2	School closure	0.025	0.0120	0.0000	0.00500	0.000500	0.000300	0.001000
3	Workplace closure	0.025	0.0000	0.0250	0.00500	0.000500	0.000300	0.001000
4	Age-targeted (60+) †	0.025	0.0120	0.0250	0.00250	0.000500	0.000030	0.000100
5	Masks and distancing	0.025	0.0040	0.0083	0.00167	0.000167	0.000100	0.000333
6	Mild lockdown	0.025	0.0030	0.0000	0.00125	0.000125	0.000075	0.000250
7	Moderate lockdown	0.025	0.0020	0.0000	0.00083	0.000083	0.000050	0.000167
8	Strict lockdown	0.025	0.0015	0.0000	0.00000	0.000063	0.000038	0.000125

† In Scenario 4, CHURCH is halved, and RANDOM and PEERS are reduced by a factor of 10 for agents aged 60 and above only. All other agents retain baseline values.

The division factors applied to non-household β in Scenarios 5–8 (3, 4, 6, and 8, respectively) were chosen to construct a monotonically increasing gradient of intervention stringency, from a partial-compliance measure (masks and distancing) to a strict lockdown. This mirrors the scenario-design logic used by Nitzsche and Simm [6] for lockdown-severity comparison. They represent illustrative stringency levels, not values calibrated to empirical compliance or mobility-reduction data for a specific historical Greek intervention (Section 6). School contacts are additionally set to zero in Scenarios 6–8, matching the school-closure policies that accompanied lockdowns in Greece.

In qualitative terms, the Scenario 5 reduction (division by 3) is broadly comparable to a partial, voluntary-compliance regime such as mask use and informal distancing without

mandatory closures. The Scenario 7 and 8 reductions (division by 6 and 8), by contrast, sit in the same qualitative band as the escalating mandatory restrictions of Greece's March–May 2020 lockdown and the moderate-to-strict stringency categories used by trackers such as the Oxford COVID-19 Government Response Tracker [16].

5. Results and Discussion

Results represent means across ten independent runs per scenario, with standard deviations computed across those runs. Table 3 summarizes the main epidemiological outcomes. Scenarios 3 and 5 exhibit bimodal behavior (epidemic establishment versus stochastic extinction); their statistics are reported conditional on epidemic establishment, with the extinction frequency noted.

Table 3. Summary of simulation outcomes (mean ± SD across 10 runs, 365 days, 19,600 agents). Scenarios 3 and 5 show bimodal outcomes; conditional statistics (epidemic-establishing runs only) are reported with extinction frequency in parentheses.

Scenario	R_0 (est.)	Peak I	Peak Day	Cum. Incidence (Mean ± SD)	Final D (Mean ± SD)	CFR (Mean ± SD)
Baseline	3.28 ± 0.85 (median 3.22, IQR 2.86–3.74)	3947	77	87.1% ± 0.2%	189 ± 9	1.10% ± 0.05%
School closure	n/a †	3060	92	77.0% ± 0.5%	178 ± 9	1.18% ± 0.06%
Workplace closure § (3/10 extinct)	n/a †	1273	80	39.9% ± 1.0%	88 ± 10	1.13% ± 0.12%
Age-targeted (60+)	n/a †	3652	75	82.4% ± 0.3%	137 ± 10	0.85% ± 0.06%
Masks and distancing § (1/10 extinct)	n/a †	1110	155	58.0% ± 1.1%	105 ± 11	0.92% ± 0.09%
Mild lockdown ‡	n/a †	21	—	1.5% ± 2.3%	3 ± 3	n/a
Moderate lockdown	n/a †	6	—	0.10% ± 0.07%	0	n/a
Strict lockdown	n/a †	5	—	0.07% ± 0.03%	0	n/a

† R_0 is estimated only for Scenario 1, where the epidemic exhibits early exponential growth. The SD of 0.85 captures the sensitivity of the Euler-Lotka growth-rate estimate to stochastic variation in early-phase I(t) across runs; individual-run estimates ranged from 1.68 to 4.50. In Scenarios 2–8, the trajectory either lacks a clean exponential growth phase or is fully suppressed, making R_0 non-estimable from growth rate alone. § Bimodal distribution detected. Statistics shown for epidemic-establishing runs only. Sc3: 7/10 runs produced a major outbreak; 3/10 achieved stochastic extinction before establishment (seeds 1000, 6000, 9000). Sc5: 9/10 runs produced an epidemic; 1/10 achieved stochastic extinction (seed 9000). These scenarios operate near the epidemic threshold, producing outcomes that are sensitive to early stochastic events. ‡ Near-critical behavior: while mean incidence is 1.5% across all runs, two of ten runs produced notable outbreaks: seed 3000 generated 1409 infections (7.2% attack rate) and seed 4000 generated 701 infections (3.6% attack rate). Five of ten runs exceeded 100 total infections. This heavy-tailed distribution indicates that mild lockdown does not reliably suppress epidemic establishment across stochastic realizations.

5.1. Baseline and Sector Closures (Scenarios 1–3)

The baseline simulation (Scenario 1) produces a substantial unmitigated epidemic. Cumulative incidence reaches 87.1% (17,078 agents infected on average over 365 days), with an epidemic peak of approximately 3947 concurrently infected agents on day 77 (Figure 2). The estimated $R_0 = 3.28 \pm 0.85$ (mean ± SD across 10 runs; median 3.22, IQR 2.86–3.74, range 1.68–4.50) is broadly consistent with published ancestral-strain SARS-CoV-2 estimates for pre-intervention settings and with Greece's early-wave trajectory. The right-skewed spread

reflects the sensitivity of the Euler-Lotka growth-rate estimator to stochastic variation in early-phase I(t): when the epidemic happens to grow slightly faster in the first 20 days, the fitted growth rate—and hence R_0—is inflated, and vice versa (Section 6 discusses the resulting caution needed when interpreting R_0 as a precise parameter). The recalibrated household $\beta = 0.025$ yields a within-household secondary attack rate of approximately 22%, consistent with the meta-analytic estimates of Madewell et al. [9] and Jing et al. [10], and substantially more realistic than the 99% implied by the prior calibration.

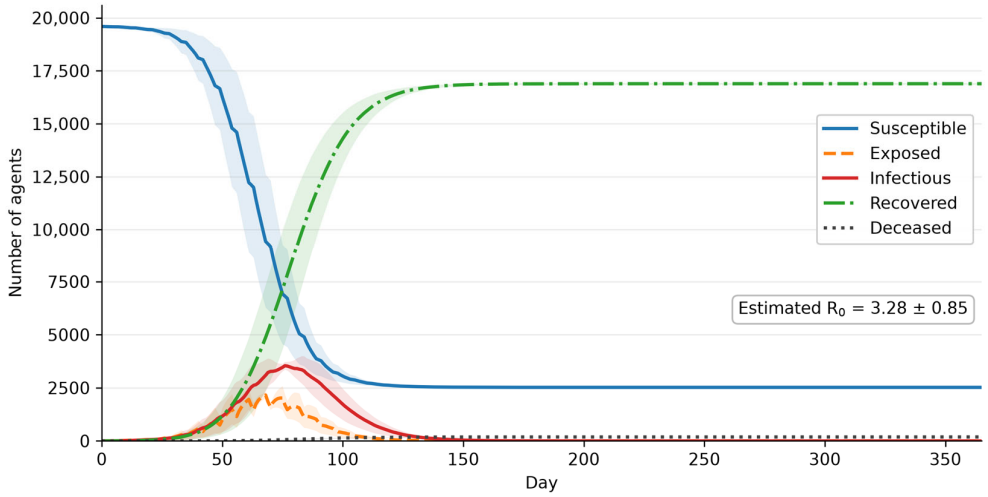

Figure 2. Baseline SEIRD dynamics (Scenario 1). Mean compartment trajectories ± 1 SD band across 10 stochastic replications (365 days, 19,600 agents). Because individual runs peak on different days, the peak of the mean I(t) curve is lower than the mean per-run peak reported in Table 3.

School closure (Scenario 2) reduces cumulative incidence by approximately 10 percentage points (from 87.1% to 77.0%) and delays the epidemic peak from day 77 to day 92 (Figure 3). While this represents a meaningful public health benefit, particularly for healthcare capacity planning, school closure alone is insufficient to contain the epidemic: the effective reproduction number remains above 1, and the epidemic runs its course through workplace, household, peer, and community contacts. The death count in Scenario 2 (178 deaths, CFR 1.18%) is effectively equivalent to the baseline (189 deaths, CFR 1.10%), consistent with the expectation that removing a transmission pathway primarily affecting low-mortality (young) agents does not substantially reduce cumulative mortality.

Workplace closure (Scenario 3) reveals a qualitatively different dynamic from the other partial-measure scenarios: a bimodal outcome distribution. In 7 of 10 replications, the epidemic established and ran its course at a conditional mean attack rate of 39.9% ± 1.0% (peak I ≈ 1273, peak day ≈ 80; 88 ± 10 deaths, CFR 1.13%). In the remaining 3 replications, stochastic extinction occurred before the epidemic established—the five seeded index infections recovered without generating sustained onward transmission, leaving the population essentially intact. This bimodal behavior is diagnostic of near-threshold dynamics: the removal of workplace contacts ($\beta_WORK = 0$) reduces the effective reproduction number to a value near but not reliably below 1. When the epidemic does establish, the reduction in attack rate relative to the baseline (39.9% vs. 87.1%, a 47-percentage-point difference) reflects the contribution of workplace contacts to total transmission; when it fails to establish, complete suppression is achieved by chance. In the two single-sector closures tested here, neither school closure nor workplace closure reliably contained the epidemic across stochastic replications.

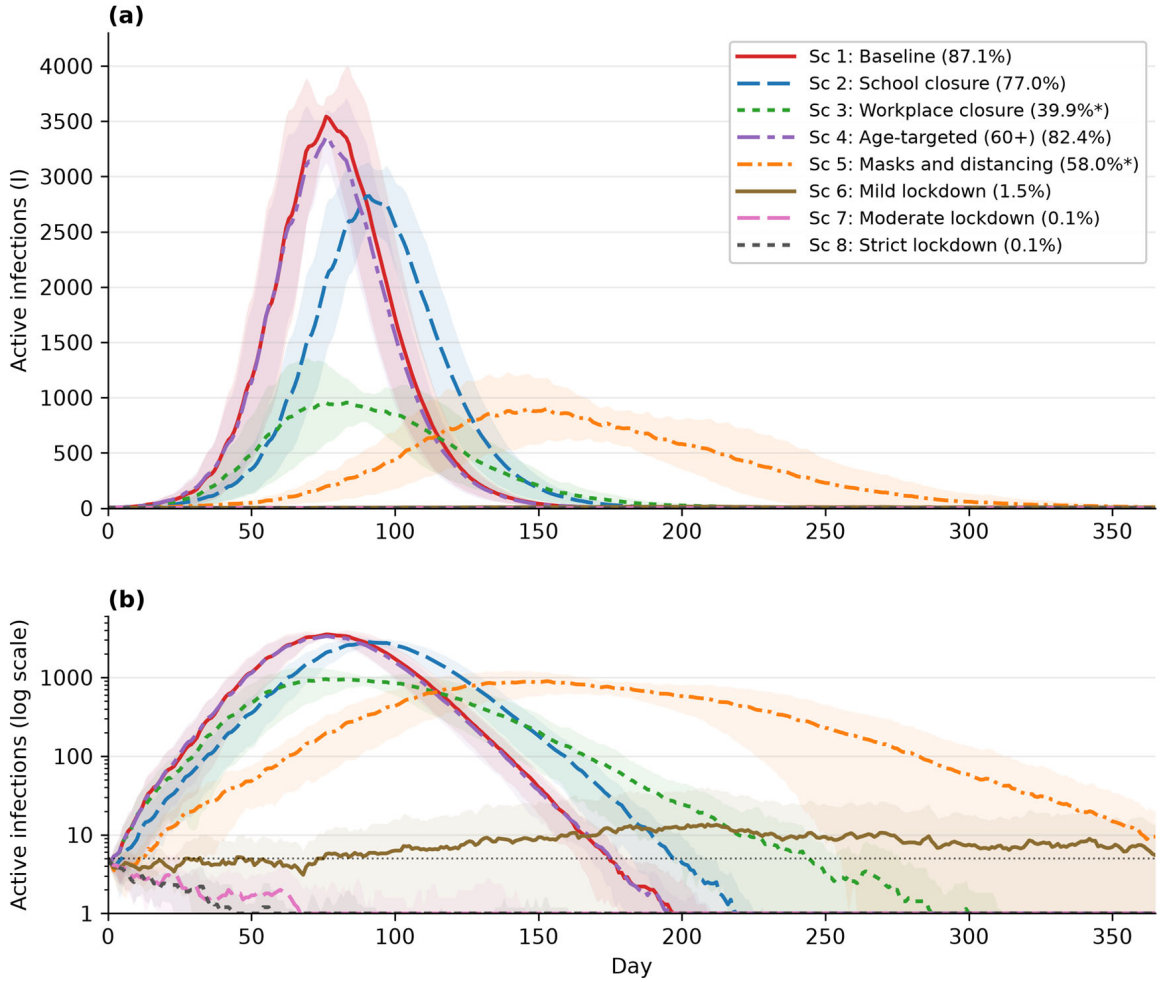

Figure 3. Infectious agent count I(t) across all eight intervention scenarios. (**a**) Linear scale; (**b**) log scale, which makes the lockdown scenarios (6–8) visible. Shaded bands indicate ± 1 SD across 10 replications. Legend values give cumulative incidence; for Scenarios 3 and 5 (marked *), curves and values are conditional on epidemic establishment. The dotted line in the lower panel marks the five initial infections; values below 1 are clipped on the log scale.

5.2. Age-Targeted Restrictions and Masks (Scenarios 4–5)

The age-targeted scenario (Scenario 4) applies reductions only to agents aged 60 and above: church attendance probability halved, and random community and peer contact probabilities reduced tenfold. Overall cumulative incidence decreases only modestly (82.4% vs. 87.1% in the baseline; Figure 4), confirming that the elderly population does not drive transmission at the population level. The protection of the targeted cohort is visible in the mortality statistics: deaths fall from 189 to 137, a reduction of 27.5%, while the overall epidemic size decreases only by 4.7 percentage points. The Scenario 4 CFR of 0.85% falls below the range observed in other epidemic scenarios; this is expected rather than a calibration failure, as the intervention selectively protects the highest-mortality subgroup, reducing the population-weighted average fatality rate below the baseline level. The church layer plausibly contributes to this effect, since β_CHURCH is halved specifically for older agents. This reflects the parameterization, not a measured transmission attribution. Representing the pathway at all is still a distinguishing feature of the present model relative to models without a church layer.

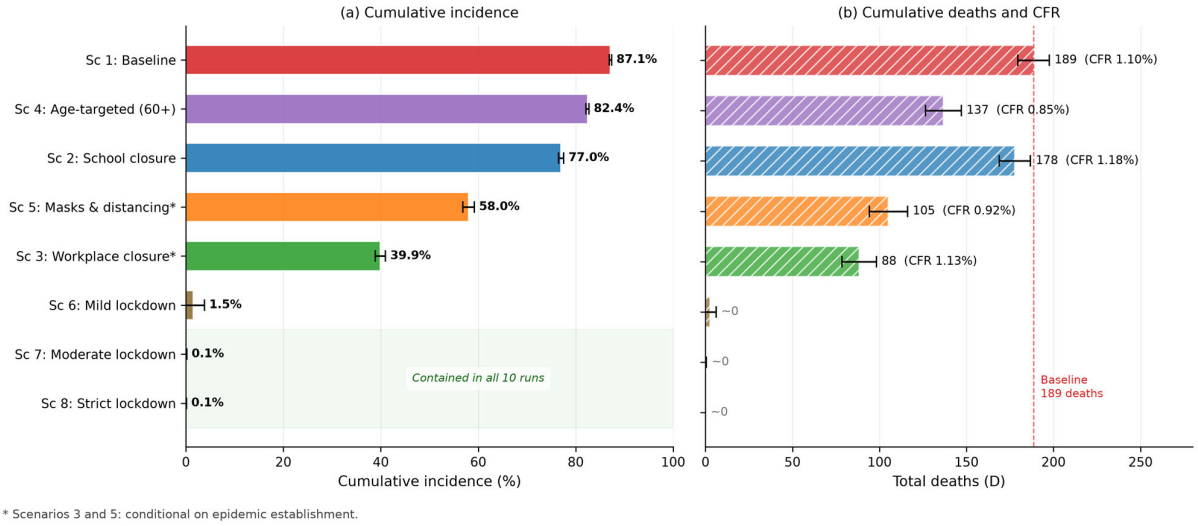

Figure 4. Ranked cumulative incidence and final deaths (mean ± 1 SD) across eight scenarios. Values for Scenarios 3 and 5 are conditional on epidemic establishment.

Masks and distancing (Scenario 5), like workplace closure, exhibit near-threshold bimodal behavior. In 9 of 10 replications, the epidemic established with a conditional mean attack rate of 58.0% ± 1.1% (peak I ≈ 1110 on day 155; 105 ± 11 deaths, CFR 0.92%). In the remaining replication (seed 9000), stochastic extinction occurred. The conditional attack rate of 58.0% represents a 29% reduction from the baseline (87.1%), and the peak is delayed by approximately 78 days (day 155 vs. day 77). The 78-day peak delay is particularly relevant for healthcare system planning. The near-threshold behavior—with extinction probability approximately 10%—indicates that masks and distancing bring the effective reproduction number close to but not reliably below 1. The persistence of household transmission at its unconstrained baseline value (β_FAMILY = 0.025 across all scenarios) is the binding constraint. Using the analytical per-layer decomposition introduced in Section 3.5, household contacts alone are estimated to contribute approximately 0.34 expected secondary infections per case. On this approximation, non-household contacts would need to be reduced by approximately 78% (divided by ~4.5) to bring $R_{0,eff}$ to 1, whereas the factor-of-three reduction in Scenario 5 leaves $R_{0,eff} \approx 0.34 + 2.96/3 \approx 1.33$—near but still above the epidemic threshold. This decomposition is analytical rather than measured (Section 6).

5.3. Sensitivity Analysis: Church Attendance Rate

The church attendance rate of 5% is a modeling assumption grounded in the authors' estimate rather than a measured city-level parameter (Section 3.3). To assess the sensitivity of the main results to this assumption, Scenario 1 was repeated with church attendance rates of 2% and 10%. Table 4 reports the outcomes.

Table 4. Sensitivity analysis: church attendance rate (Scenario 1 baseline, mean of 3 runs [seeds 0, 1000, 2000] × 365 days; metrics computed per run then averaged, the same convention used in Table 3).

Church Attendance	Attending Agents	R_0 (est.)	Peak I	Peak Day	Cum. Incidence	Deaths	CFR (%)
2%	392	3.46	3916	78	87.0%	191	1.12
5% (baseline)	980	3.35	3943	78	87.1%	188	1.10
10%	1960	3.53	3999	76	87.9%	219	1.27

Note: this analysis uses 3 replications per attendance level, fewer than the 10 used in Table 3, so its estimates carry more sampling uncertainty than the main scenario comparisons.

The cumulative incidence differs by only 0.9 percentage points across the full 2–10% range. The 5% baseline row uses the same three seeds (0, 1000, 2000) as the first three replications of Table 3's Scenario 1, and reproduces its cumulative incidence and CFR exactly (87.1%, CFR 1.10%), confirming methodological consistency between the two analyses. The church attendance rate is therefore not a primary driver of cumulative incidence; the 5% baseline assumption is unlikely to materially affect the main conclusions regarding the relative ordering of scenario efficacy. Deaths and CFR show a more pronounced sensitivity (188–219 deaths, CFR 1.10–1.27% across the range). Baseline church attendance is assigned uniformly across ages, not age-weighted (Section 3.3), so this sensitivity is not evidence of a church-driven age effect. It is estimated from only three replications per attendance level (Table 4), so it may simply reflect sampling noise at this small replication count, not a systematic relationship between attendance rate and mortality. Mortality outcomes, however, are considerably more sensitive than cumulative incidence to the church attendance rate: over the same 2–10% range, CFR varies by 0.17 percentage points and deaths by 31 (16% of the baseline death count), a proportionally much larger change than the 0.9-point shift in cumulative incidence. Mortality-related conclusions should therefore be read with this uncertainty in mind (Section 6), including the magnitude of Scenario 4's 27.5% death reduction. That reduction reflects three simultaneous age-targeted layer reductions (CHURCH, RANDOM, and PEERS for agents aged 60 and above), not a church-attributable age skew.

5.4. Lockdown Scenarios (Scenarios 6–8)

The lockdown scenarios exhibit a gradient of containment reliability. Moderate and strict lockdown (Scenarios 7–8) reliably suppress the epidemic in all 10 replications. Mean cumulative incidence is 0.10% ± 0.07% and 0.07% ± 0.03%, respectively, with the five seeded index infections recovering or dying without generating sustained onward transmission (mean peak $I \leq 6$ across all runs). These scenarios correspond to effective reproduction numbers well below 1. The mechanism is direct: at baseline $R_0 \approx 3.3$, non-household contacts are estimated to contribute approximately 2.96 expected secondary infections per primary case (the remaining 0.34 from household contacts), using the analytical decomposition of Section 3.5. Reducing non-household transmission by factors of 6 and 8, while eliminating school contacts, brings this approximate effective R_0 to about 0.83 and 0.71, respectively—consistent with the near-certain stochastic extinction observed in the simulations.

Mild lockdown (Scenario 6), which divides non-household β by 4, exhibits near-critical behavior. Applying the same analytical decomposition used above for Scenarios 7–8 (0.34 household + 2.96/4 non-household) gives $R_{0,\text{eff}} \approx 1.08$, marginally above the epidemic threshold, not below it. This rough approximation should not be read as more precise than the "order-of-magnitude" caveat given in Section 6. Mean cumulative incidence across 10 runs is 1.5% ± 2.3%; while five runs produced very small chains (<100 total infections), the remaining five produced outbreaks ranging from 122 to 1409 infections. Figure 5 shows the per-seed outbreak size on a log scale, making the heavy-tail structure directly visible. Outcomes cluster into a low group (5–65 infections, five seeds) and a substantially larger group (122–1409 infections, five seeds spanning more than one order of magnitude), rather than following a narrow unimodal distribution around the mean. The two largest—seed 3000 (1409 infections; 7.2% attack rate) and seed 4000 (701 infections; 3.6% attack rate)—illustrate the heavy-tail behavior characteristic of near-critical branching processes. When $R_{0,\text{eff}}$ is close to 1, most chains quickly die out, but rare early-transmission events can produce outbreaks substantially larger than the mean. From a practical policy perspective, this means mild lockdown achieves reliable containment only

in expectation—in a single implementation with uncertain compliance, the outcomes may be substantially worse.

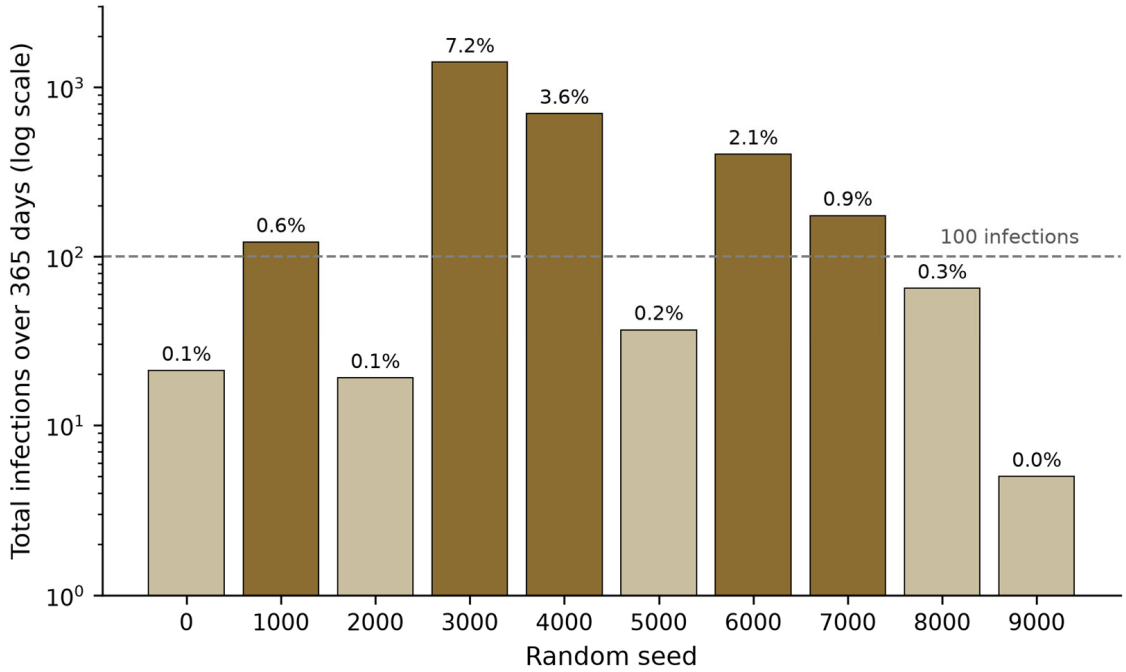

Figure 5. Per-run outbreak size for Scenario 6 (mild lockdown) across the 10 random seeds used in Table 3, log-scaled. Dark bars mark runs exceeding 100 total infections; the dashed reference line marks this threshold. Percentage labels give final cumulative incidence for each run.

Within the model, and conditional on compliance achieving the modeled reductions, lockdown measures at these intensities suppress the simulated epidemic and do not merely flatten its trajectory. This is a finding about the simulated dynamics under the model's assumptions, not an independent empirical claim about real-world lockdown efficacy, and not a formal comparison to the observed Greek time series (Section 6). The functional differentiation between mild, moderate, and strict lockdown in terms of containment reliability is a central finding of the 10-replication analysis: while all three scenarios produce near-zero mean incidence, only moderate and strict lockdown provide robust suppression across all stochastic realizations.

Across epidemic Scenarios 1–5 (conditional on epidemic establishment for Scenarios 3 and 5), the simulated CFR ranges from 0.85% to 1.18%, broadly consistent with the EODY-derived calibration target of approximately 1.20% [15]. Scenario 4 is the lowest for the structural reason described in Section 5.2.

5.5. Comparison with Related Agent-Based Models

Table 5 positions the present model against its direct baseline [2] and the two other agent-based COVID-19 studies most central to its design: Singh et al. [5], the source of the age-stratified mortality structure used in Section 3.5, and Nitzsche and Simm [6], whose lockdown-severity scenario design motivated the Scenario 5–8 stringency gradient in Section 4. The comparison is structural and methodological, not a quantitative outcome benchmark. The four studies differ in population, scale, and reporting conventions, so their published epidemiological outcomes are not directly comparable on a single numeric scale.

Table 5. Comparative positioning of the present model relative to its baseline and to related agent-based COVID-19 models.

Study	Setting/Scale	Contact Layers	Compartments	Mortality Calibration	Intervention Design
Thomopoulos & Tsichlas [2]	Patras central district, 19,600 agents	5: household, workplace, school, peer, random	SEIR	Not modeled	Not evaluated in [2]
Singh et al. [5]	Madrid metropolitan area	Not detailed in this comparison	SEIR-type, with hospitalization/fatality tracking	Age-stratified hospitalization and fatality probability tables (source of Section 3.5's mortality structure)	Not compared on this axis
Nitzsche and Simm [6]	Abstract city ABM	Abstract city model, layers not specified	Not detailed in this comparison	Not compared on this axis	Lockdown scenarios of varying severity, with a monotonic restriction-intensity/suppression relationship (basis for Section 4's stringency gradient)
This study	Patras central district, 19,600 agents (extends [2])	7: adds public transport and church attendance to [2]'s five layers	SEIRD (adds Deceased, using [5]'s mortality structure)	Age-stratified probabilities from [5], rescaled via HOSP_FACTOR/DEATH_FACTOR to a Greek CFR target (Section 3.5)	Eight-scenario stringency gradient following the qualitative logic of [6] (Section 4)

The comparison highlights the specific contribution of this study relative to its three points of reference: it extends [2]'s synthetic-population architecture with two empirically motivated interaction layers absent from that baseline. It adopts [5]'s mortality structure instead of re-deriving one. And it adapts [6]'s scenario-design logic to a Greek intervention context instead of reproducing [6]'s abstract-city parameterization directly.

6. Limitations

The single largest limitation of this study is the absence of external validation. The model is calibrated against a single scalar target—the population-level case-fatality rate—together with a comparison of the estimated R_0 to published literature values. It is not validated against empirical Greek epidemic timing, peak timing, cumulative incidence, hospitalization trajectories, mortality trajectories, age-stratified case distributions, intervention timing, or mobility data, and no simulated trajectory is compared against an observed Patras or Greek COVID-19 time series. The model should therefore be interpreted as an exploratory scenario generator that supports internally consistent comparisons between intervention configurations, not as a predictive model of the actual Patras outbreak. All quantitative findings in Section 5 should be read with this qualification in mind.

Each scenario is averaged across ten stochastic runs. This is sufficient to detect bimodal outcome distributions (Scenarios 3 and 5) and to estimate standard deviations, but it falls short of the 30+ replications conventionally required for formal confidence intervals or significance testing. The baseline cumulative incidence of 87.1% is substantially above any measured single-wave COVID-19 incidence in Greece. This follows from a fully susceptible population with no behavioral adaptation, vaccination, or prior immunity over 365 days; the framing suits a hypothetical unmitigated worst case but would overstate transmission burden under partial immunity or behavioral change. The church attendance rate of 5% is likewise the authors' own estimate, not a measured figure. As Section 5.3 shows, varying it from 2% to 10% changes cumulative incidence by under 1 percentage point, but shifts deaths by up to 31 and CFR by up to 0.17 points. This is a sensitivity that should be weighed

when interpreting mortality-specific findings such as Scenario 4's mortality benefit, and one estimated from only three replications per attendance level rather than the ten used elsewhere in this paper. The model does not represent vaccination, reinfection, variant evolution, seasonal forcing, or geospatial movement between neighborhoods.

R_0 estimation is a second significant source of uncertainty. The reported value ($R_0 = 3.28 \pm 0.85$; median 3.22, IQR 2.86–3.74, across ten replications of Scenario 1) comes from a log-linear fit to the early exponential growth of I(t) (Euler-Lotka), an estimator developed for homogeneously mixing compartmental models. The present model is stochastic, network-structured, multilayer, and demographically heterogeneous, none of which the Euler-Lotka estimator assumes; classical exponential-growth estimators are therefore not a natural fit for this model class. The right-skewed run-to-run distribution (median below mean; range 1.68–4.50) reflects sensitivity to stochastic early-growth variation rather than a stable underlying constant, and the gap between this estimate and the R_0 implied by the observed final epidemic size (approximately 2.3, Section 3.5) reinforces the point. Reported R_0 values, including the "effective R_0" figures for the lockdown scenarios in Section 5.4 ($\approx$0.83, $\approx$0.71), should accordingly be read as order-of-magnitude indicators rather than precisely estimated parameters. Formal R_0 estimation via next-generation matrix aggregation across layers, better suited to this model's structure, is deferred to future work.

Section 5 offers two mechanistic explanations: the decomposition of expected secondary infections into household and non-household components, and the attribution of part of Scenario 4's mortality reduction to the church layer. Both are analytical approximations derived from the calibrated β values and mean contact-group sizes, not measurements from a formal transmission-attribution analysis. They indicate the likely order of magnitude and direction of each mechanism, not a precise, independently verified decomposition. A layer-resolved transmission attribution analysis is identified as useful future work.

The intervention stringency levels used in Scenarios 5–8 (dividing non-household β by 3, 4, 6, or 8) were chosen, as described in Section 4, to construct a monotonically increasing gradient of NPI intensity. They are not calibrated to measured compliance or mobility-reduction data for any specific Greek policy episode. The finding that all three lockdown scenarios contain the epidemic is consequently inherent to this parameterization: any configuration that brings the effective reproduction number below 1 will produce this outcome. Mapping empirical compliance data onto these stringency levels is left to future work.

Finally, conclusions describing interventions as "suppressing" or "containing" the epidemic (Sections 5.4 and 7) refer to suppression within this simulated model, under its specific parameterization and assumptions. They are not claims about the real-world effectiveness of the corresponding public health measures, which would require the external validation discussed above.

7. Conclusions

This study presents an extended multilayer ABM for COVID-19 transmission in the central district of Patras, Greece, built on the 19,600-agent synthetic population and five-layer contact architecture of [2]. The principal contributions are: (i) a public transport layer grounded in local ridership data; (ii) a church attendance layer motivated by Greek social context and supported by systematic review evidence [4]; and (iii) a recalibration of household and non-household transmission parameters to be consistent with empirical household secondary attack rate estimates [9,10] and to yield a biologically plausible $R_0 \approx 3.3$.

The simulation results yield four principal epidemiological findings. First, sector closures vary substantially in efficacy depending on the population share they target:

school closure reduces cumulative incidence by approximately 10 percentage points, while workplace closure, when the epidemic establishes, achieves a 47-percentage-point reduction. Second, workplace closure and masks-and-distancing operate near the epidemic threshold, producing bimodal outcome distributions: stochastic extinction was observed in 30% and 10% of replications, respectively, indicating that these interventions reduce the effective reproduction number to near but not reliably below 1. Third, age-targeted isolation of individuals aged 60 and above provides a 27.5% reduction in mortality with minimal impact on overall transmission. This intervention lowers contact rates in the church, random-community, and peer layers at the same time for the targeted cohort, so the effect cannot be attributed to any single layer. Fourth, lockdown scenarios exhibit a gradient of suppression reliability: moderate and strict lockdown reliably contain the epidemic across all replications, while mild lockdown exhibits near-critical behavior with two of ten runs producing attack rates of 7.2% and 3.6%. Near the epidemic threshold, the average outcome across runs can differ substantially from what a single run produces. A deterministic compartmental model does not capture this variability. Sensitivity analysis confirms that the church attendance rate parameter (Section 5.3, Table 4) is not a primary driver of these findings: varying attendance from 2% to 10% changes cumulative incidence by less than 1 percentage point.

The most practically relevant finding is the distinction between interventions that reduce R_0 substantially but not below 1 (school closure, workplace closure, masks) and those that bring R_0 reliably below 1 (moderate and strict lockdown). Near-threshold scenarios produce bimodal outcomes that a small number of replications may mischaracterize as either reliably epidemic or reliably containing. The 10-replication design was sufficient to reveal this pattern in the present experiment, though a larger ensemble would be needed to characterize it with more precision. As discussed in Section 6, the model has not been validated against empirical Greek epidemic time series, and its intervention stringency levels and R_0 estimate are analytical approximations rather than externally calibrated measurements. The findings above should therefore be read as internally consistent comparisons within an exploratory scenario generator, not as validated predictions of the real Patras epidemic. Future work will prioritize external validation against observed Greek case, hospitalization, and mortality time series, most plausibly by comparing simulated day-by-day case and death trajectories against EODY's weekly surveillance reports [15], already used here for CFR calibration. It will also pursue formal R_0 estimation via next-generation matrix methods, an expanded geographic scope, and a larger replication count to support formal uncertainty quantification.

Author Contributions: Conceptualization, V.T.; methodology, V.T. and P.T.; software, P.T.; validation, V.T.; formal analysis, P.T. and V.T.; investigation, P.T.; data curation, P.T.; writing—original draft preparation, P.T. and V.T.; writing—review and editing, V.T.; visualization, P.T.; supervision, V.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. This study is based entirely on computational modeling and simulation and does not involve human subjects or animal experimentation.

Informed Consent Statement: Not applicable.

Data Availability Statement: Code, scenario configurations, and simulation random seeds are publicly available at https://github.com/vorlon83/Extending-MultilayerABM-COVID-19-Patras (accessed on 9 November 2026).

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ABM	Agent-Based Model
CFR	Case-Fatality Rate
ELSTAT	Hellenic Statistical Authority (Greek: ΕΛΣΤΑΤ)
EODY	Greek National Public Health Organization (Greek: ΕΟΔΥ)
IQR	Interquartile Range
NGM	Next-Generation Matrix
NPI	Non-Pharmaceutical Intervention
ODE	Ordinary Differential Equation
SAR	Secondary Attack Rate
SD	Standard Deviation
SEIRD	Susceptible-Exposed-Infected-Recovered-Deceased

References

1. Bonabeau, E. Agent-Based Modeling: Methods and Techniques for Simulating Human Systems. *Proc. Natl. Acad. Sci. USA* **2002**, *99*, 7280–7287. [CrossRef]
2. Thomopoulos, V.; Tsichlas, K. An Agent-Based Model for Disease Epidemics in Greece. *Information* **2024**, *15*, 150. [CrossRef]
3. Luo, Q.; Liu, W.; Liao, J.; Gu, Z.; Fan, X.; Luo, Z.; Zhang, X.; Hang, J.; Ou, C. COVID-19 Transmission and Control in Land Public Transport: A Literature Review. *Fundam. Res.* **2024**, *4*, 417–429. [CrossRef]
4. Goumballa, N.; Hoang, V.T.; Al-Tawfiq, J.A.; Sokhna, C.; Gautret, P. Evidence for Transmission of SARS-CoV-2 at Religious Mass Gatherings: A Systematic Review. *New Microbes New Infect.* **2024**, *60–61*, 101442. [CrossRef]
5. Singh, D.E.; Marinescu, M.-C.; Guzmán-Merino, M.; Durán, C.; Delgado-Sanz, C.; Gomez-Barroso, D.; Carretero, J. Simulation of COVID-19 Propagation Scenarios in the Madrid Metropolitan Area. *Front. Public Health* **2021**, *9*, 636023. [CrossRef]
6. Nitzsche, C.; Simm, S. Agent-Based Modeling to Estimate the Impact of Lockdown Scenarios and Events on a Pandemic Exemplified on SARS-CoV-2. *Sci. Rep.* **2024**, *14*, 13391. [CrossRef]
7. Wan, J.; Ichinose, G.; Small, M.; Sayama, H.; Moreno, Y.; Cheng, C. Multilayer Networks with Higher-Order Interaction Reveal the Impact of Collective Behavior on Epidemic Dynamics. *Chaos Solitons Fractals* **2022**, *164*, 112735. [CrossRef]
8. Von Hoene, E.; Roess, A.; Kavak, H.; Anderson, T. Synthetic Population Generation with Public Health Characteristics for Spatial Agent-Based Models. *PLoS Comput. Biol.* **2025**, *21*, e1012439. [CrossRef]
9. Madewell, Z.J.; Yang, Y.; Longini, I.M.; Halloran, M.E.; Dean, N.E. Household Transmission of SARS-CoV-2: A Systematic Review and Meta-Analysis. *JAMA Netw. Open* **2020**, *3*, e2031756. [CrossRef]
10. Jing, Q.-L.; Liu, M.-J.; Zhang, Z.-B.; Fang, L.-Q.; Yuan, J.; Zhang, A.-R.; Dean, N.E.; Luo, L.; Ma, M.-M.; Longini, I.; et al. Household Secondary Attack Rate of COVID-19 and Associated Determinants in Guangzhou, China: A Retrospective Cohort Study. *Lancet Infect. Dis.* **2020**, *20*, 1141–1150. [CrossRef]
11. Kazil, J.; Masad, D.; Crooks, A. Utilizing Python for Agent-Based Modeling: The Mesa Framework. In *Social, Cultural, and Behavioral Modeling*; Thomson, R., Bisgin, H., Dancy, C., Hyder, A., Hussain, M., Eds.; Lecture Notes in Computer Science; Springer International Publishing: Cham, Switzerland, 2020; Volume 12268, pp. 308–317. [CrossRef]
12. He, X.; Lau, E.H.Y.; Wu, P.; Deng, X.; Wang, J.; Hao, X.; Lau, Y.C.; Wong, J.Y.; Guan, Y.; Tan, X.; et al. Temporal Dynamics in Viral Shedding and Transmissibility of COVID-19. *Nat. Med.* **2020**, *26*, 672–675. [CrossRef]
13. Pew Research Center. *Being Christian in Western Europe*; Pew Research Center: Washington, DC, USA, 2018. Available online: https://www.pewresearch.org/religion/2018/05/29/being-christian-in-western-europe/ (accessed on 9 November 2026).
14. Diekmann, O.; Heesterbeek, J.A.P.; Roberts, M.G. The Construction of Next-Generation Matrices for Compartmental Epidemic Models. *J. R. Soc. Interface.* **2010**, *7*, 873–885. [CrossRef]
15. National Public Health Organization. *Epidemiological Surveillance of COVID-19*; EODY: Athens, Greece, 2026. Available online: https://eody.gov.gr/el/ (accessed on 9 November 2026).
16. Hale, T.; Angrist, N.; Goldszmidt, R.; Kira, B.; Petherick, A.; Phillips, T.; Webster, S.; Cameron-Blake, E.; Hallas, L.; Majumdar, S.; et al. A Global Panel Database of Pandemic Policies (Oxford COVID-19 Government Response Tracker). *Nat. Hum. Behav.* **2021**, *5*, 529–538. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.