Journal of Evolutionary Economics (2026) 36:57
https://doi.org/10.1007/s00191-026-00974-5

RESEARCH

# Spatial heterogeneity and budget-constrained treatments in epidemic dynamics: An agent-based approach

**Andrea Caravaggio**[1] · **Silvia Leoni**[2]

Received: 29 July 2025 / Accepted: 12 June 2026


**Abstract**
The management of infectious diseases increasingly relies on innovative but costly pharmaceutical treatments, raising complex trade-offs between epidemiological containment, fiscal sustainability, and institutional coordination. We develop a spatially structured agent-based model in which decentralized health authorities allocate treatment under local budget constraints while infection spreads across a two-dimensional lattice through neighborhood spillovers. Within each location, treatment intensity is chosen endogenously, interacting with local GDP dynamics and pricing conditions. Simulation results reveal that purely decentralized optimization mitigates but does not reverse infection growth within policy-relevant horizons, generating persistent spatial heterogeneity in both epidemiological and economic outcomes. We then introduce bounded spatial policy interaction, showing that partial coordination substantially improves containment but may increase the persistence of fiscal engagement. Extending the model to heterogeneous and time-varying pricing, we find that price discrimination amplifies medium-run infection and fiscal pressure under decentralization. However, when surplus revenues finance endogenous R&D, treatment efficacy improves over time, generating a feedback mechanism in which innovation mitigates long-run epidemiological and economic losses. Our findings highlight the critical interplay between spatial structure, decentralized decision-making, pricing design, and innovation incentives in shaping epidemic outcomes. Effective management of high-cost treatments requires not only medical efficacy but also institutional coordination and carefully designed market mechanisms.

**Keywords** Agent-based modeling · Health policy · Infectious disease control · SIR model · Treatment pricing

**JEL Classification** C63 · H51 · I18

✉ Silvia Leoni
silvia.leoni@unifi.it

1 Department of Economics and Statistics, University of Siena, Siena, Italy

2 Department of Economics and Management, University of Florence, Florence, Italy

Published online: 17 July 2026

## 1 Introduction

The global experience with COVID-19 and other recent epidemics has underscored the urgent need for robust and flexible frameworks to design and evaluate public health interventions under significant uncertainty and resource constraints. Central to this challenge is the emergence of innovative but costly pharmaceutical treatments, which pose a fundamental dilemma for health authorities (HAs): how to contain the spread of infectious diseases while managing scarce budgetary resources.

Although HAs manage consistent trade-offs for chronic diseases with regular therapies, novel treatments that can outright cure infectious diseases disrupt this balance. When a novel and effective treatment is introduced at a high price, the HA faces a critical, non-trivial trade-off: the clinical goal of curing the current stock of infected individuals vs. the fiscal necessity of aligning with a strict budget constraint. In the case of communicable diseases, the inability to treat a patient due to budget limitations is not merely a clinical failure but a public health risk, as it leaves active transmission chains in the untreated population.

Recent literature on global health policy has confirmed that inequities in access to treatments and vaccines—including pricing-based disparities—have real consequences for public health outcomes. A comprehensive policy analysis by Gleeson et al. (2023) documents how structural inequities in access to COVID-19 health products undermine disease containment and exacerbate socioeconomic vulnerability. Similarly, computational studies like Gozzi et al. (2023) demonstrate that delayed or uneven distribution of vaccines significantly lowers herd immunity, amplifying both health and economic costs.

Traditional epidemiological models, typically grounded in mean-field approaches such as deterministic Susceptible–Infected–Recovered (SIR) systems, have provided useful insights into disease propagation. However, such models often fail to capture the spatial and institutional heterogeneity that shapes real-world policy responses. Recent advancements in spatial-economic epidemiology, most notably by La Torre et al. (2022, 2024), have successfully integrated geographical heterogeneities and economic externalities into epidemic management frameworks. While these studies provide essential insights into optimal lockdown strategies across space, they operate under the assumption of a centralized social planner. In practice, public health responses—especially regarding pharmaceutical procurement—are often fragmented across decentralized HAs with localized budget constraints.

Our framework takes primary inspiration from the dynamic optimization approach of Dubois and Magnac (2024), who analyze the trade-off between aggressive intervention and budget smoothing for costly treatments. However, while their model offers a critical foundation for understanding drug-related budget constraints, it does not account for spatial contagion or the interactions between decentralized jurisdictions. Our work bridges this gap by embedding these financial trade-offs into a spatially explicit environment where decisions are made not by a single planner, but by localized HAs.

In recent years, there has been increasing interest in agent-based modeling (ABM) approaches for studying epidemic control. To date, this strand of literature has focused predominantly on the individual level, often in the context of HIV among the

blood-borne viruses (Vermeer et al. 2022), or COVID-19 lately (Kerr et al. 2021; Wolfram 2020), to study transmission dynamics among the general population or high-risk subgroups, such as intravenous drug users (see, e.g., Ale et al. 2024). Epidemic ABMs have also been applied to a variety of pathological behaviors that have social origins (Axtell and Farmer 2025), including smoking (Institute of Medicine 2015), drug addiction (Agar and Wilson 2002; Heard et al. 2014), and obesity (Hammond 2009). In these studies, agents represent individuals, and the focus is on behavioral or clinical outcomes.

While traditional models based on ordinary and partial differential equations assume that there are no network effects, agent-based models are particularly valuable for their ability to capture heterogeneous populations and spatial interactions, which are essential for understanding disease transmission across heterogeneous geographic areas or interconnected networks.

Recent ABM studies further suggest that spatial and institutional heterogeneities create “tipping points" and unintended consequences that are invisible in aggregate models. Specifically, localized variations in timing and severity (Thomas et al. 2020) and the sensitivity of contagion to the initial source location (Kustudic et al. 2021) underscore that small structural differences can amplify outcomes far beyond what average-level predictions suggest.

In this paper, we develop an agent-based and spatially structured model of epidemic management in which decentralized health authorities operate under resource constraints to treat the infected population using an innovative but costly drug. The disease propagates through a discrete-time SIR dynamic, where the local transmission rate depends endogenously on the effort spent on treatment, and spatial interactions follow a Moore neighborhood scheme (Shiflet and Shiflet 2014), allowing infections to spread across adjacent cells.[1]

Our work contributes to the literature discussed above in several ways. First, it integrates the traditional SIR compartmental model with the more innovative framework of ABM. Integrating mainstream approaches with ABM is not new (Axtell and Farmer 2025). In this case, the hybrid approach allows us to overcome the inherent limitations of compartmental models in accounting for spatial and institutional heterogeneity. Unlike much of the existing epidemiological literature that adopts ABMs, our model operates at a higher level of institutional abstraction: the “agents” are the HAs themselves, rather than individual people. This shift in perspective enables a specific focus on the economic dimensions of governmental decision-making. By treating HAs as autonomous agents, we can observe how localized coordination failures and spatial spillovers emerge from heterogeneous budget pressures and pricing schemes. To our knowledge, limited research has employed ABM at this institutional level to study either policy diffusion or the explicit impact of drug pricing on the collective capacity to contain communicable diseases.

Second, relative to the literature focusing on the spatial dimension (La Torre et al. 2022, 2024), our approach allows to relax the assumption of a central social planner and models a decentralized optimization process from the institutional side that more

[1] The assumptions regarding the mechanism of illness spread are foundational to the results and dynamics shown in the following sections.

realistically reflects the fragmented nature of public health responses. This framework provides the opportunity to analyze how interactions among decision-makers lead to emergent consequences and systemic effects that would remain unobserved in the absence of such interactions.

In addition, we explore and compare the effects of homogeneous and heterogeneous pricing for the treatment. The pricing dimension of our model relates to a growing strand of literature on differential access to health innovations. In practice, pharmaceutical companies often set prices based on willingness-to-pay, income levels, or effectiveness, a strategy that can introduce inequities when public budgets are constrained. Studies such as Danzon and Towse (2003) and Moon et al. (2011) have analyzed tiered pricing and its implications for access to medicines in low- and middle-income countries. However, few epidemic models incorporate pricing decisions as endogenous factors that interact with local budget constraints and infection dynamics. Furthermore, potential revenues are central determinants of firms' R&D investment decisions (Acemoglu and Linn 2004; Dubois et al. 2015).

Our analysis builds on this literature by endogenizing a mechanism of price discrimination, where the price varies across locations according to a rule based on local GDP, reflecting real-world practice such as external reference pricing (ERP) and value-based pricing (Vogler et al. 2021; Leopold et al. 2012). We include a model specification in which the markup revenue is invested in R&D, which contributes to ameliorating endogenous treatment efficacy. We show how such pricing can distort treatment incentives, leading to increasing spatial inequality in both disease prevalence and economic outcomes, thus undermining both epidemiological effectiveness and equity. This aligns with empirical evidence suggesting that wide price differentials can incentivize parallel trade and the emergence of so-called "gray markets", which can further disrupt supply chains (Kanavos and Costa-Font 2005). Finally, we demonstrate that agent interaction can drastically mitigate the negative effects of price discrimination, effectively internalizing price heterogeneity through coordinated responses.

The remainder of the paper is structured as follows: Section 2 introduces the core structure of the model, presenting the baseline SIR dynamics (without and with spatial contagion) and the formulation of the local optimization problem faced by health authorities. Section 3 details the implementation of the agent-based simulations and the initialization strategy adopted. Section 4 presents the simulation results under uniform pricing, under the specification of null or positive institutional coordination. Section 5 investigates the implications of introducing heterogeneous and time-varying pricing, and compares the resulting dynamics to the baseline scenario. In Section 6 we briefly analyze the effects of coordination and pricing on the spatial distribution of infection. In Section 7, we explore the effect of endogenizing the efficacy of treatment based on an investment in R&D. Finally, Section 8 summarises our findings, and Section 9 concludes.

## 2 The model

We investigate the dynamics and control of a non-lethal infectious disease in a spatially distributed population. The disease can spread locally via agent-to-agent interactions

and can be treated through an innovative but costly therapy. Given the budgetary constraints faced by public health authorities, not all infected individuals can be treated, which gives rise to a control problem: the optimal allocation of limited resources to contain the epidemic over time and space.

To jointly analyze the temporal and spatial dimensions of epidemic control, we implement a two-dimensional toroidal lattice in which each cell represents a local subpopulation evolving over time and following a compartmental discrete-time SIR scheme. The spread of infection occurs both within each cell and across adjacent cells, following a Moore neighborhood structure—that is, each cell interacts with its eight immediate neighbors. Within each cell, a local health authority determines the share of infected individuals to be treated at each time, maximizing a payoff that reflects both the health benefits and economic costs of treatment, infection, and resource use.

To make the entire analysis accessible also to readers who are not specifically familiar with SIR frameworks, we first analyze a simplified, non-spatial version of the model in which treatment intensity is kept constant. This will also allow us to establish some benchmark results concerning the existence and stability of equilibria. We then introduce spatial interactions across cells and examine their implications for local dynamics and epidemic containment. Finally, we incorporate the role of the health authority, formalizing its payoff function and characterizing the optimal treatment policy under a budget constraint.

### 2.1 An epidemic SIR model with treatment

In the absence of spatial interactions, the disease evolves according to the following discrete-time SIR model with treatment. We assume a constant population size $N = S_t + I_t + R_t$ (e.g., via balancing birth and death rates), and let $\lambda_t \in [0, 1]$ denote the fraction of infected individuals treated at time $t$.[2]

$$\begin{cases} S_{t+1} = S_t - \Delta I_t \\ I_{t+1} = I_t + \Delta I_t - \chi \lambda_t I_t \\ R_{t+1} = R_t + \Delta R_t \end{cases} \tag{1}$$

Here, $\Delta I_t = \beta(I_t)\, c\, S_t \frac{I_t}{N}$ denotes new infections at time $t$, where $c > 0$ is the average number of contacts per person per period, while $\Delta R_t = \chi \lambda_t I_t$ represents the individuals who recover due to treatment, with $\chi \in [0, 1]$ measuring treatment efficacy.

The infection rate $\beta(I_t)$ is assumed to decrease in the health authority's treatment effort $h_t = p\lambda_t I_t$, in which $p > 0$ represents the constant unit price of treatment, and it reflects the idea that stronger treatment campaigns reduce effective contagion. Specifically, we consider

$$\beta(I_t) = \frac{\beta_0}{1 + \chi h_t + \overline{s}},$$

[2] Susceptible, infected, and recovered states are assumed to be nonnegative at each time $t$.

where $\beta_0$ is the baseline transmission rate, and $\bar{s} \geq 0$ captures other exogenous mitigation measures (e.g., social distancing). For future reference, it is convenient to express the model in terms of disease prevalence. We define the infection prevalence as $i_t := \frac{I_t}{N}$, which represents the share of the population that is infected at time $t$.

Under the assumption of constant population size, the dynamics of the infected compartment can equivalently be written in terms of $i_t$, which will be particularly useful when comparing regions of different sizes and when introducing spatial interaction.

### 2.1.1 Equilibrium analysis under constant treatment

To gain analytical insights into the long-run behavior of the epidemic, we assume a constant treatment intensity $\lambda_t = \lambda \in [0, 1]$. Under this assumption, the system always admits a disease-free equilibrium (DFE) given by $(S^*, I^*, R^*) = (N, 0, 0)$. Since population size is constant, the system can be reduced to a two-dimensional dynamic system in the $(S_t, I_t)$ plane, with the recovered compartment obtained residually as $R_t = N - S_t - I_t$. To analyze the local stability of the DFE, we linearize the system around the point $(N, 0)$ and derive the corresponding Jacobian matrix:

$$J(S_t, I_t) = \begin{pmatrix} 1 - \beta(h_t)c \cdot \frac{I_t}{N} & -\frac{\beta_0 c(1+\bar{s})}{(1+\chi\, p\lambda\, I_t+\bar{s})^2 N} \\ \beta(h_t)c \cdot \frac{I_t}{N} & \frac{(1-\chi\lambda)(1+\bar{s})^2 N - \beta_0 c(1+\bar{s})N}{(1+\chi\, p\lambda\, I_t+\bar{s})^2 N} \end{pmatrix}. \tag{2}$$

By evaluating $J(S_t, I_t)$ at the disease-free equilibrium $(S^*, I^*) = (N, 0)$, we obtain

$$J_{\text{DFE}} = \begin{pmatrix} 1 & -\dfrac{\beta_0 c}{1+\bar{s}} \\ 0 & 1 + \chi\lambda \left( \dfrac{\beta_0 c}{\chi\lambda(1+\bar{s})} - 1 \right) \end{pmatrix}. \tag{3}$$

Defining the basic reproduction number as

$$\mathcal{R}_0 := \frac{\beta_0 c}{\chi\lambda(1+\bar{s})}, \tag{4}$$

we derive the eigenvalues of $J_{\text{DFE}}$, that is the pair $(\lambda_1, \lambda_2) = (1, 1+\chi\lambda(\mathcal{R}_0 - 1))$. The condition $R_0 < 1$ ensures local decay of the infected compartment around the disease-free state and therefore provides a critical threshold for epidemic control. Indeed, to suppress infection, the treatment rate must satisfy

$$\lambda > \bar{\lambda} = \frac{\beta_0 c}{\chi(1+s)}. \tag{5}$$

Differently, when $\mathcal{R}_0 > 1$, the disease initially spreads and the DFE becomes unstable. However, since there is no demographic renewal and no reinfection, the number of susceptible individuals decreases over time. As a result, the effective reproduction number eventually falls below one, causing the epidemic to peak and then decline. The system ultimately converges to the disease-free equilibrium ($I^* = 0$) not because it is

globally stable, but due to structural depletion of susceptibles. The resulting dynamics describe an epidemic outbreak followed by natural extinction, rather than the emergence of a persistent endemic state.

### 2.2 Local infection dynamics with Moore neighborhood interactions

Consider now an extension of the dynamic model that accounts for local spatial interactions. Each location $v \in \mathcal{V}$ hosts a subpopulation following the same SIR structure, but infection can also spread across neighboring cells. Let $\mathcal{N}_v$ denote the Moore neighborhood of cell $v$, consisting of its eight adjacent locations. The dynamics in each cell $v$ are given by:

$$\begin{cases} S_{t+1}^v = S_t^v - \Delta I_t^v \\ I_{t+1}^v = I_t^v + \Delta I_t^v - \chi \lambda^v I_t^v \\ R_{t+1}^v = R_t^v + \Delta R_t^v \end{cases} \tag{6}$$

where

$$\Delta I_t^v = \beta(I_t^v) \cdot c \cdot S_t^v \cdot \left( i_t^v + \sum_{\ell \in \mathcal{N}_v} i_t^\ell \right), \tag{7}$$

$$\Delta R_t^v = \chi \lambda^v I_t^v. \tag{8}$$

In this spatial setting, the infection rate is still endogenous and defined locally as:

$$\beta(I_t^v) = \frac{\beta_0}{1 + \chi p \lambda^v I_t^v + \bar{s}}. \tag{9}$$

As in the case without local interactions, the system still admits a global DFE defined by $I_t^v = 0$, $R_t^v = 0$, and $S_t^v = N^v$ for all $v$. To analyze its local stability, we linearize the infected compartment around the DFE as

$$I_{t+1}^v \approx I_t^v + \frac{\beta_0 c}{1 + \bar{s}} \left( I_t^v + N^v \sum_{\ell \in \mathcal{N}_v} i_t^\ell \right) - \chi \lambda^v I_t^v,$$

which simplifies to

$$I_{t+1}^v = \left( 1 + \frac{\beta_0 c}{1 + \bar{s}} - \chi \lambda^v \right) I_t^v + \frac{\beta_0 c N^v}{1 + \bar{s}} \sum_{\ell \in \mathcal{N}_v} i_t^\ell.$$

The expression for $I_{t+1}^v$ demonstrates that local infection dynamics are driven by a combination of internal transmission and spillover from the infection prevalence $i_t^\ell$ in neighboring cells $\mathcal{N}_v$. This implies that a cell cannot be viewed in isolation; its stability is contingent on the connectivity and population density of its neighbors. We

define a cell-specific basic reproduction number, $\mathcal{R}_0^v$, which accounts for these spatial interactions:

$$\mathcal{R}_0^v := \frac{\beta_0 c}{\chi \lambda^v (1+\overline{s})} \left(1 + \sum_{\ell \in \mathcal{N}_v} \frac{N^v}{N^\ell}\right). \tag{10}$$

Here, the term $\sum(N^v/N^\ell)$ represents the relative risk posed by neighboring populations. A conservative sufficient condition for local asymptotic stability of the disease-free equilibrium is $\mathcal{R}_0^v < 1$ for all $v$.[3] Consequently, for the entire lattice to remain infection-free, the treatment effort $\lambda^v$ in every cell must exceed a threshold determined by the "weakest link" in the network, i.e., the cell with the highest local transmission potential:

$$\max_{v \in \mathcal{V}} \mathcal{R}_0^v < 1 \quad \Longleftrightarrow \quad \lambda^v > \overline{\lambda}^v := \frac{\beta_0 c}{\chi (1+\overline{s})} \left(1 + \sum_{\ell \in \mathcal{N}_v} \frac{N^v}{N^\ell}\right) \qquad \forall v \in \mathcal{V}. \tag{11}$$

This condition highlights that spatial heterogeneity creates "hotspots". Specifically, cells with high internal density $N^v$ adjacent to low-density neighbors are more prone to instability, as they act as net exporters of the pathogen, requiring a disproportionately higher treatment effort $\overline{\lambda}^v$ to ensure global stability.

### 2.3 The health authority's problem in the spatial setting

In the spatial model, we assume that each cell $v \in \mathcal{V}$ hosts a local Health Authority (HA) responsible for managing treatment decisions within its jurisdiction. At each time $t$, the local HA must determine the fraction $\lambda_t^v \in [0, 1]$ of infected individuals in cell $v$ to be treated with the innovative drug.

The treatment is costly, and the HA operates under a health budget $b_t^v$. Since treatment reduces contagion and promotes recovery, but also entails direct and indirect costs, the HA aims to maximize a payoff function that captures both the epidemiological and economic impacts of its decision.

Let $h_t^v = p\lambda_t^v I_t^v$ denote the treatment expenditure in cell $v$ at time $t$, where $p$ is the unit cost of the curative treatment. The HA's optimization problem is given by:

$$\max_{0 \le \lambda_t^v \le 1} \Pi_t^{HA,v} \tag{12}$$

subject to the budget constraint $h_t^v \le b_t^v$, or equivalently $\lambda_t^v \le \frac{b_t^v}{pI_t^v}$ when $I_t^v > 0$.

The payoff function is:

$$\Pi_t^{HA,v} = (\gamma^v \chi - p)\lambda_t^v I_t^v - \alpha_I (I_t^v)^2 - \alpha_T^v \frac{(\lambda_t^v I_t^v)^2}{1 + \delta I_t^v} \tag{13}$$

[3] More generally, local stability of the global DFE in the coupled lattice is characterized by the spectral radius of the linearized infection operator; the condition based on $\max_v \mathcal{R}_0^v$ provides a tractable sufficient bound.

where the positive parameters are described as follows:

- $\gamma^v$ captures the monetary benefit of each recovery due to treatment;
- $\alpha_I$ weighs the (quadratic) cost from the prevalence of infection;
- $\alpha_T^v$ represents the cost of implementing the treatment effort;
- $\delta$ models the decreasing marginal returns of treatment as the number of infections increases.

Although the HA's payoff in Eq. 13 depends only on the local infection stock $I_t^v$, spatial spillovers shape the decision environment indirectly through the dynamics of this state variable. Since $\Delta I_t^v$ depends on both local and neighboring prevalences, infection in adjacent cells affects the future evolution of $I_t^v$ and, therefore, the marginal benefit of treatment. Thus, even if the optimization problem is formulated locally, it is embedded in a spatially interconnected epidemiological system.

Solving the first-order condition for an interior solution yields

$$(\gamma^v \chi - p) I_t^v - 2\alpha_T^v \frac{\lambda_t^v (I_t^v)^2}{1 + \delta I_t^v} = 0, \tag{14}$$

from which we directly obtain

$$\lambda_{int}^{*,v} = \frac{(\gamma^v \chi - p)(1 + \delta I_t^v)}{2\alpha_T^v I_t^v}. \tag{15}$$

Thus, the optimal *internal* policy for each local HA is:

$$\lambda_t^{*,v} = \min \left\{ \max \left\{ 0, \lambda_{int}^{*,v} \right\}, \frac{b_t^v}{p I_t^v}, 1 \right\}. \tag{16}$$

Equation 16 defines the individually optimal treatment decision of the local HA under myopic optimization and budget feasibility. However, we assume that the treatment intensity effectively implemented in each cell may also reflect local policy interaction. In particular, health authorities form *naive expectations* about the treatment policies of neighboring jurisdictions. Let $\mathcal{N}_v$ denote the Moore neighborhood of cell $v$, and define the average effective treatment intensity observed among neighbors at time $t-1$ as:

$$\bar{\lambda}_{t-1}^{\mathcal{N}_v} = \frac{1}{8} \sum_{\ell \in \mathcal{N}_v} \lambda_{t-1}^{\ell}. \tag{17}$$

We assume that HAs expect neighboring policies to persist unchanged, i.e., expectations are purely backward-looking and coincide with the last observed value. The treatment intensity effectively implemented in cell $v$ at time $t$ is then given by a

convex combination of the locally optimal choice and the expected average policy of neighbours:[4]

$$\lambda_t^v = (1-\omega)\lambda_t^{*,v} + \omega\,\overline{\lambda}_{t-1}^{\mathcal{N}_v} \qquad \omega \in [0,1]. \tag{18}$$

Here, $\omega$ measures the degree of local policy interaction. When $\omega = 0$, the model reduces to purely decentralized individual optimization. For $\omega > 0$, treatment policies exhibit bounded spatial interdependence: neighboring policies serve as benchmarks for local decision-making and provide a channel through which policy coordination, imitation, or informational spillovers may affect epidemic dynamics. Since feasibility must be preserved, the adjusted treatment intensity that any HA actually adopts is obtained by projecting $\lambda_t^v$ onto the admissible set defined by the budget constraint and the unit interval:

$$\lambda_t^{adj,v} = \min\left\{\max\{0, \lambda_t^v\}, \frac{b_t^v}{pI_t^v}, 1\right\}. \tag{19}$$

This formulation implies that treatment intensity varies dynamically across space and time as a function of local infection levels and as an adjustment between local decision and neighborhood interactions. Indeed, although each HA solves its optimization problem independently, the effective treatment intensity may exhibit spatial interdependence due to the local interaction mechanism described above. This generates endogenous coordination effects that can influence aggregate epidemic dynamics and the persistence of spatial clusters.

## 3 The agent-based framework

### 3.1 Initialization

The simulation environment employs a toroidal grid topology. Such a topology eliminates edge effects by ensuring all patches have identical neighborhood connectivity. While this abstracts real-world geography, it is computationally preferable for studying global pandemic dynamics, as isolated geographic boundaries would introduce artifacts.

In a toroidal grid, agents crossing one boundary re-enter the opposing boundary. In our case, each cell of the grid represents ideally a world country, with a total of 196 cells. Each cell carries a population drawn from a transformed lognormal distribution that takes values in the interval $[100, 10^6]$. The choice of the distributional model aligns with empirical evidence, which consistently shows that human settlements and population densities are highly skewed (Gabaix 1999; Blank and Solomon 2000; Eeckhout 2004).

At the model setup, the global population is initialized with a fixed aggregate infection share of 1%. This total volume of infected individuals is then distributed

[4] In this ABM framework, spatial policy interaction should not be interpreted as fully strategic behavior in the game-theoretic sense, where agents solve forward-looking best-response problems over payoff functions. Rather, it captures a decentralized behavioral heuristic based on local social learning and policy benchmarking: under uncertainty, HAs may partially adapt their own treatment intensity by observing the policies implemented in geographically adjacent jurisdictions. This interpretation is consistent with ABM approaches, in which agents adjust their behavior using simple rules informed by their peers' actions.

across cells (countries) according to three distinct scenarios, as illustrated in Fig. 1. By holding the total initial infection constant, we are able to isolate the effect of spatial configuration and localized density on the subsequent epidemic trajectory.

In the first scenario, the initial 1% of infected individuals is distributed evenly across all cells in the world. To ensure consistency with local population constraints, if the assigned number of infected individuals for a specific cell exceeds its total population $N^v$, the overflow is redistributed among cells that still have remaining capacity. This scenario represents a post-endemic or undetected widespread transmission context, where the disease is already present globally without a specific point of origin.

The second scenario emulates a classic pandemic emergency with a well-defined epicenter. The 1% infected share is concentrated in a single randomly selected seed country $i_0$. If the population of the seed country is insufficient to contain the total initial infected volume, the remainder is distributed among its immediate neighbors.

The third scenario distributes the 1% infected share across $n = 10$ spatially distinct outbreaks. Each epicenter is assigned a proportionate share of the total infected

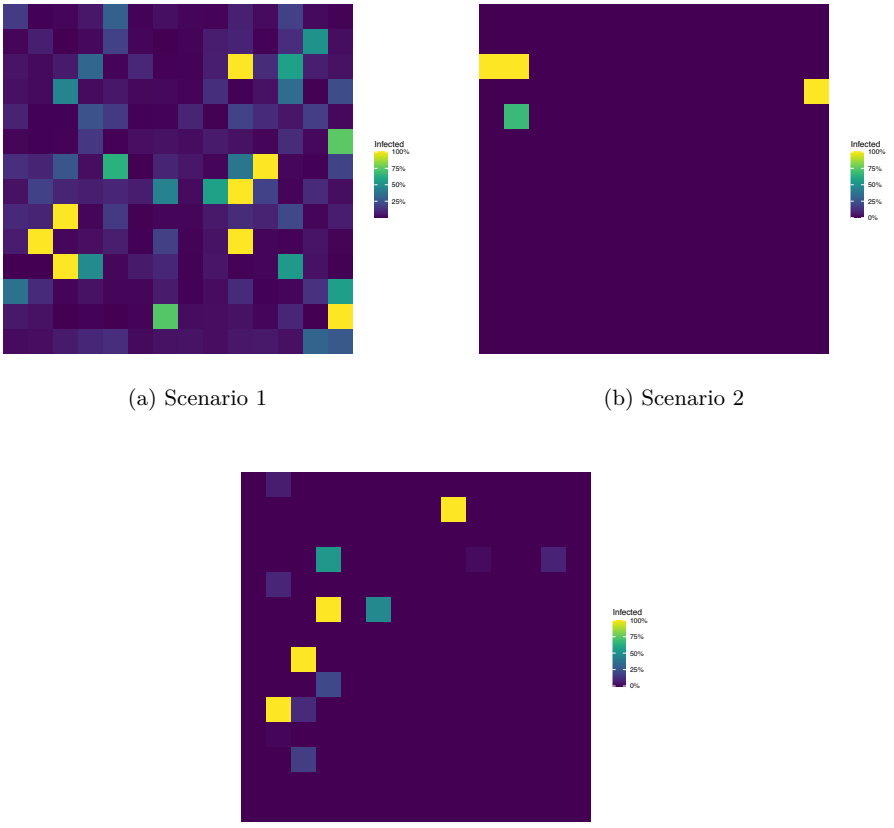


(a) Scenario 1

(b) Scenario 2

(c) Scenario 3

**Fig. 1** The three scenarios at $t = 0$. Seed = 1000

population. Similar to Scenario 2, if an epicenter's population capacity is reached, the infection overflows into adjacent territories.

As mentioned, $\gamma^v$ represents the monetary benefit of recovering from the disease. It is heterogeneous across countries. We imagine it as the average annual salary of those that once recovered and can again productively contribute to their country's economic life.

Therefore, each cell $v$ is initialized with a wage $\gamma^v$ drawn from a double Pareto lognormal (dPLN) distribution. The choice of this specific distributional model is explicitly intended to replicate a well-established stylized fact: while the middle of the income distribution is approximately lognormal, the tails might exhibit power law (Pareto) behavior (Reed and Jorgensen 2004; Clementi and Gallegati 2005). We set the location parameter $\mu = 2.6$, which yields a median annual wage of approximately €13,460 and truncate the distribution at a lower bound of 1.5 (€1,500) and an upper bound of 120 (€120,000). [5]

Since we can interpret the parameter $\gamma^v$ as the average annual wage within a country, we derive a country's GDP, $B^v$, as:

$$B_t^v = \gamma^v (N^v - \sigma I_t^v) \tag{20}$$

where $\sigma \in [0, 1)$ captures the degree to which infected individuals are able to contribute to the country's economic life.[6] By setting $\sigma$ as a measure of diminished productivity, the term $(N^v - \sigma I_t^v)$ effectively represents the productive population equivalent. Consequently, the total output $B_t^v$ fluctuates endogenously with the infection dynamics, creating a direct feedback mechanism between the epidemiological state of the country and the fiscal capacity of its HA.

The constant and homogeneous parameters of the model are given the values reported in Table 1.

### 3.2 Simulating the model

We model a generic infectious disease characterized by relatively slow contagion dynamics, more comparable to chronic infections such as hepatitis C than to acute, fast-spreading viruses like COVID-19 or influenza. Infections of this type typically progress over longer time scales, with slower transmission rates and prolonged infectious periods. This epidemiological profile requires focusing on broader spatial and policy dynamics rather than day-to-day fluctuations.

[5] The double Pareto-lognormal (dPLN) distribution is governed by four parameters: $\mu$ and $\sigma$, representing the location and scale of the underlying lognormal base, and $\alpha$ and $\beta$, which are the shape parameters ruling the upper and lower Pareto tails, respectively. In our implementation, we assume symmetric tails where $\alpha = \beta$. The specific parameters chosen for the wage distribution are listed in Table 1. Also, note that all monetary variables are measured in thousands of euros.

[6] The parameter $\sigma$ is restricted to the interval $[0, 1)$ in the simulations. In the limiting case $\sigma = 1$, infected individuals would contribute zero to productive capacity. If, in addition, the entire population is infected ($I_t^v = N^v$), this would imply $B_t^v = 0$ and therefore a zero health budget $b_t^v$. In such a configuration, treatment effort would become infeasible, and the cost-to-budget ratio would be undefined due to a vanishing denominator. To avoid this degenerate corner case, we impose $\sigma < 1$.

**Table 1** Parameter set employed in the model set-up

| Variable | Values |
|---|---|
| $\chi$ | 0.7 |
| $c$ | 6 |
| $\beta_0$ | 0.1 |
| $\delta$ | 4 |
| $\bar{s}$ | 10 |
| $\sigma$ | 0.3 |
| $N^v$ | $\sim$ Lognormal($\mu = 8.52, \sigma = 0.59$) truncated at [100, $10^6$] |
| $\gamma^v$ | $\sim$ dPLN($\mu = 2.6, \sigma = 1.0, \alpha = 2.5$) subject to $\gamma^v \in [1.5, 120]$ |

As a result, we approximate each discrete time step in the model as representing one year, which provides a meaningful temporal resolution for studying the effects of treatment decisions, budgetary constraints, and spatial interactions over the course of an epidemic. The model is iterated for 60 time steps, therefore each representing a year. While this 60-year span represents a medium-run horizon from an economic policy perspective, it constitutes a long-run period in epidemiological terms, covering the full lifecycle and potential stabilization of the disease. This duration is therefore chosen to provide a reasonable and comprehensive period for observing the dynamics of the type of disease being modeled that would not fully manifest over shorter durations (see, for instance, Dubois and Magnac 2024).

The intervention strategy of each Health Authority ($HA^v$) is governed by a state-dependent activation mechanism. Rather than providing continuous treatment, HAs only enter the pharmaceutical market when local epidemiological conditions warrant intervention.

Let $a_t^v \in \{0, 1\}$ be a binary indicator representing the active treatment status of country $v$ at time $t$. The policy follows a hysteresis-based trigger: the $HA^v$ activates the health policy once prevalence surpasses a critical 5% threshold. To reflect institutional persistence and ensure the outbreak is sufficiently contained before withdrawing resources, the policy is only deactivated when prevalence falls below a "safe" low-threshold of 2%. Within the intermediate range, the $HA^v$ maintains its current status. This logic is formalized as follows:

$$a_t^v = \begin{cases} 1 & \text{if } i_t^v > 0.05 \\ 0 & \text{if } i_t^v < 0.02 \\ a_{t-1}^v & \text{if } 0.02 \leq i_t^v \leq 0.05. \end{cases} \tag{21}$$

When the status is inactive ($a_t^v = 0$), no treatment is provided ($\lambda_t^{adj,v} = 0$). Conversely, when the policy is active ($a_t^v = 1$), the $HA^v$ determines the share of the infected population to treat, $\lambda_t^{adj,v}$, on the basis of a linear combination of its optimal

strategy and its expectation about the average policy adopted by neighbors (as shown in Eq. 18):

$$\lambda_t^{adj,v} = \begin{cases} 0 & \text{if } a_t^v = 0 \\ \min\left\{\max\{0, \lambda_t^v\}, \frac{b^v}{pI_t^v}, 1\right\} & \text{if } a_t^v = 1. \end{cases} \tag{22}$$

The health budget in each cell $v$ at time $t$ is represented by $b_t^v$. This budget is determined by a state-dependent share $r_t^v$ of the local GDP ($B_t^v$):

$$b_t^v = r_t^v \cdot B_t^v \tag{23}$$

The expenditure share $r_t^v$ is endogenously adjusted according to the local disease prevalence $i_t^v$. We assume a baseline allocation $\bar{r} = 0.05$, which remains constant until a critical prevalence threshold $\widetilde{i} = 0.05$ is reached. The fiscal rule is defined as follows:

$$r_t^v = \begin{cases} \bar{r} & \text{if } i_t^v < \widetilde{i} \\ \min\left(\bar{r} + \ln(1 + i_t^v), 1\right) & \text{if } i_t^v \geq \widetilde{i}. \end{cases} \tag{24}$$

As seen in Eq. 13, the HA's payoff takes into account a monetary cost, $\alpha_T^v$, indirectly associated with the treatment, representing, for instance, distribution and personnel costs. We assume this cost to be highly dependent on exogenous macroeconomic shocks and inflation dynamics. Therefore, $\alpha_T^v$ is updated at each time step according to the following rule:

$$\alpha_T^v \sim \mathcal{U}(50, 100). \tag{25}$$

Once cells set their $\alpha_T^v$, the infection rate $\beta^v(I_t^v)$ evolves according to Eq. (9), and the number of new infections, following Eq. (7). Finally, cells update their population compartments accordingly. The algorithm followed at each time step is displayed in Fig. 2. We run 2000 Monte Carlo simulations for each setup scenario and model specification adopted.

## 4 Analysis under homogeneous and constant pricing

In this section, we investigate the epidemiological and economic implications of the model under a benchmark pricing regime in which the unit price of treatment is homogeneous across countries and constant over time. The three epidemiological initialization scenarios introduced earlier are analyzed comparatively, and their dynamic trajectories are evaluated in terms of infection prevalence, treatment intensity, and spatial and distributional outcomes.

Within this pricing framework, we distinguish between two institutional configurations governing the behavior of local HAs. In the first configuration, there is no spatial policy interaction: each HA implements its individually optimal treatment decision independently of neighboring jurisdictions, so that treatment effort is purely decentralized and determined exclusively by local conditions (that is, $\lambda_t^{adj,v} = \lambda_t^{*,v}$ for every $v$).

In the second configuration, we allow for local spatial interdependence in policy implementation: the effective treatment intensity in each cell is determined as a convex

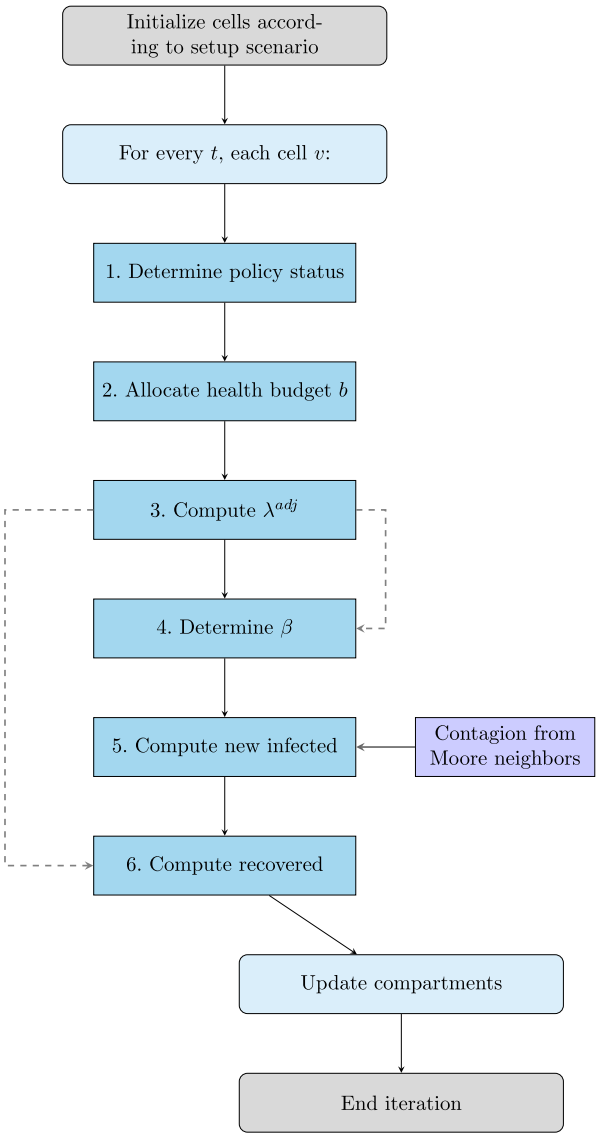


**Fig. 2** Flow of actions within a single cell $v$ in a single time step $t$

combination of the locally optimal decision and the average treatment effort observed in neighboring cells. The parameter $\omega$ governs the strength of this local interaction mechanism, capturing the extent to which treatment policies display spatial imitation or coordination.

Throughout Section 4, we assume a homogeneous and constant unit price of treatment, denoted by $p$. We set $p = 5$, which corresponds to a cost of €5000 per treatment unit (as monetary variables are measured in thousands of euros). This order of

magnitude is consistent with reported prices of antiviral regimens for chronic infections such as hepatitis B, hepatitis C, and HIV in high-income countries, where treatment costs per patient can amount to several thousand euros even after price negotiations (see Dubois and Magnac 2024; Yao et al. 2025). However, the purpose of this calibration is not to replicate a specific drug price, but to adopt a realistic benchmark for high-cost innovative treatments under budget constraints.

### 4.1 Purely decentralized treatment policies ($\varpi = 0$)

We begin by considering the benchmark case in which treatment policies are purely decentralized and no spatial policy interaction occurs ($\omega = 0$). In this configuration, each Health Authority implements its individually optimal treatment intensity, so that $\lambda_t^{adj,v} = \lambda_t^{*,v}$ for all $v$ and $t$.

This case provides a clean institutional baseline: spatial spillovers arise exclusively from epidemiological contagion, while policy decisions remain entirely local. Any observed spatial patterns or aggregate regularities therefore emerge solely from the interaction between local optimization, budget constraints, and infection dynamics.

Panels (a), (b), (c) in Fig. 3 show the final state ($t = 60$) for the three scenarios considered for one simulation run. The comparison with the initial state $t = 0$ displayed in Fig. 1 immediately shows that initial conditions are crucial for the disease spread and the final outcome over the time span considered.[7] Substantial spatial heterogeneity persists. Clusters of infection remain visible in Scenarios 1 and 3, whereas Scenario 2, which is characterized by a single localized outbreak, displays rapid containment and near-eradication.

Figure 4 reports the cross-sectional density of compartment sizes across cells at $t = 60$. A first robust pattern across scenarios is the strong right-skewness of the infected distribution. In all three cases, the majority of cells exhibit very low levels of active infection, while a relatively small subset of cells retains non-negligible infection stocks. This heavy-tailed structure indicates that decentralized containment generates substantial heterogeneity across regions even when aggregate infection is declining. In Scenario 1, the recovered compartment displays a pronounced bimodal structure. A mass of cells is clustered near zero, indicating regions that experienced limited contagion, while a second peak appears at higher values, corresponding to regions that underwent significant infection waves before converging. This polarization reflects path dependence induced by initial conditions: small stochastic differences in early infection trajectories translate into persistent cross-sectional divergence. Scenario 2 (single epicenter) exhibits a markedly different distribution. Both infected and recovered densities concentrate sharply near zero, indicating rapid containment and limited spatial diffusion. The absence of a pronounced secondary mode suggests that decentralized optimization is sufficient to prevent widespread propagation when the outbreak is geographically concentrated at the outset. Scenario 3 (multiple localized outbreaks) occupies an intermediate position. While most cells converge toward low infection levels, a visible secondary mass emerges in the recovered distribution. This

[7] Comparison is ensured by setting the same seed, i.e., the numerical value used to initialize a pseudo-random number generator (PRNG).

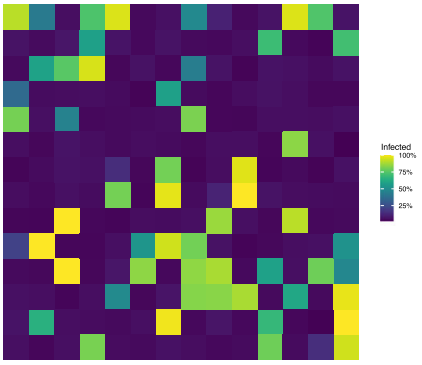


(a) Scenario 1, $\omega = 0$

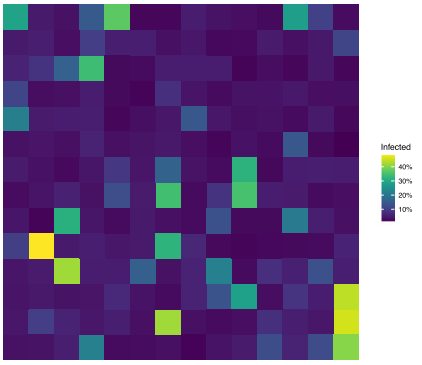


(d) Scenario 1, $\omega = 0.3$

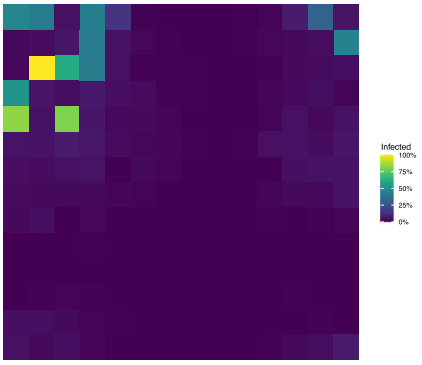


(b) Scenario 2, $\omega = 0$

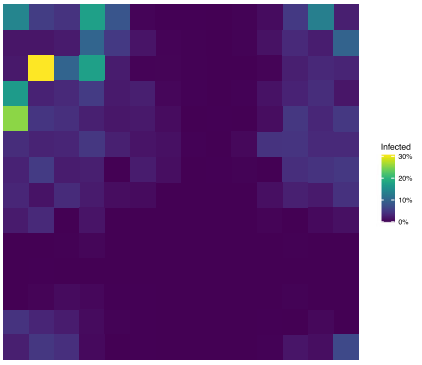


(e) Scenario 2, $\omega = 0.3$

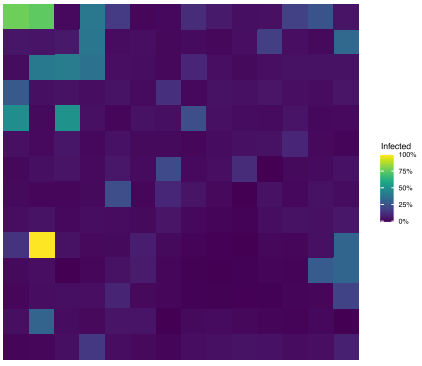


(c) Scenario 3, $\omega = 0$

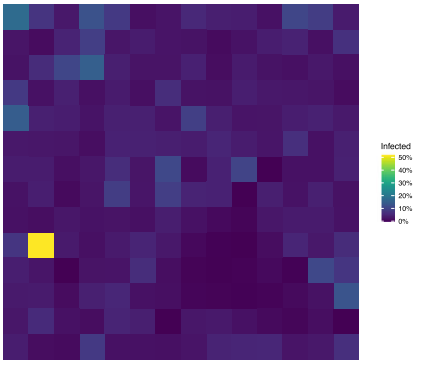


(f) Scenario 3, $\omega = 0.3$

**Fig. 3** Final spatial distribution of infection across cells at $t = 60$ (Seed = 1000)

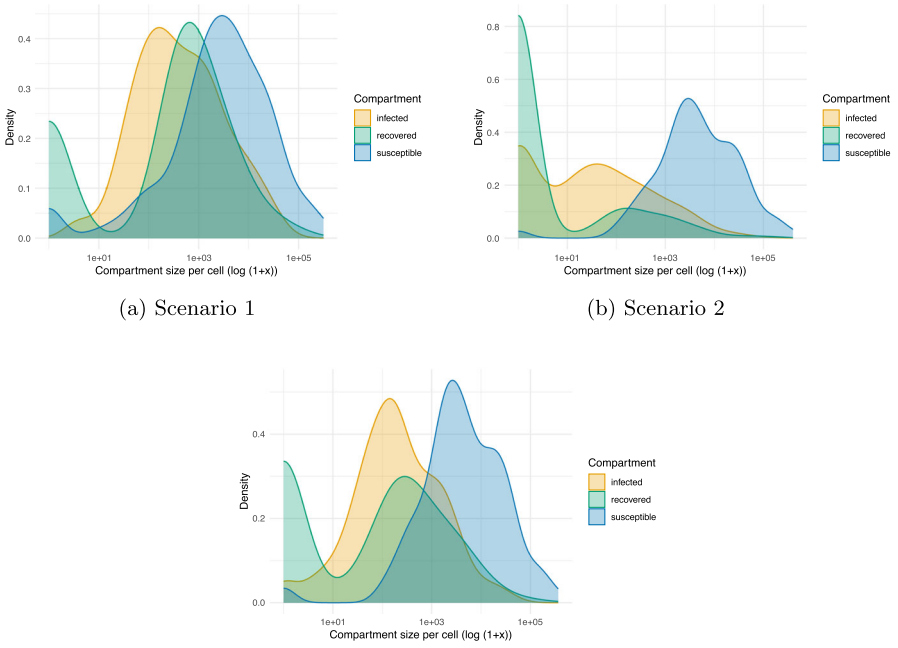


(a) Scenario 1

(b) Scenario 2

(c) Scenario 3

**Fig. 4** Distribution of population compartments across cells at ($t = 60$). Seed = 1000.
Note to the figure: The data are presented on a $\log(1 + x)$ scale. This transformation serves two key purposes: firstly, the logarithmic scale significantly enhances visualization by compressing wide ranges of values, making trends, and patterns more discernible; secondly, the addition of '1' to the argument $(1 + x)$ prevents the loss of observations that would otherwise be zero, which would be undefined under a simple logarithmic transformation

indicates partial but uneven diffusion across space, consistent with the coexistence of locally contained regions and clusters that experienced temporary surges.

Reading these results from an epidemiological perspective, we observe that decentralized policies are capable of reducing aggregate prevalence but do not eliminate spatial dispersion. Infection dynamics remain locally persistent in a subset of cells due to spatial spillovers and heterogeneous budget constraints. From an economic standpoint, the cross-sectional dispersion in recovered and infected stocks implies heterogeneous GDP trajectories across jurisdictions. Since treatment expenditure and output losses are endogenously linked to infection prevalence, the emergent bimodality translates into asymmetric economic impacts despite uniform pricing and identical institutional rules. This macroeconomic dispersion is the direct consequence of heterogeneous epidemiological trajectories interacting with budget-constrained local decisions. Since output evolves according to Eq. 20, even moderate differences in infection prevalence translate into persistent differences in productive capacity. Regions that experience higher intermediate infection peaks face both larger temporary output losses and tighter budget constraints, which in turn limit their treatment intensity and slow convergence toward full containment. The spatial patterns reported in Panels (a), (b), (c) of Fig. 5 confirm this mechanism. In Scenario 1, GDP contractions

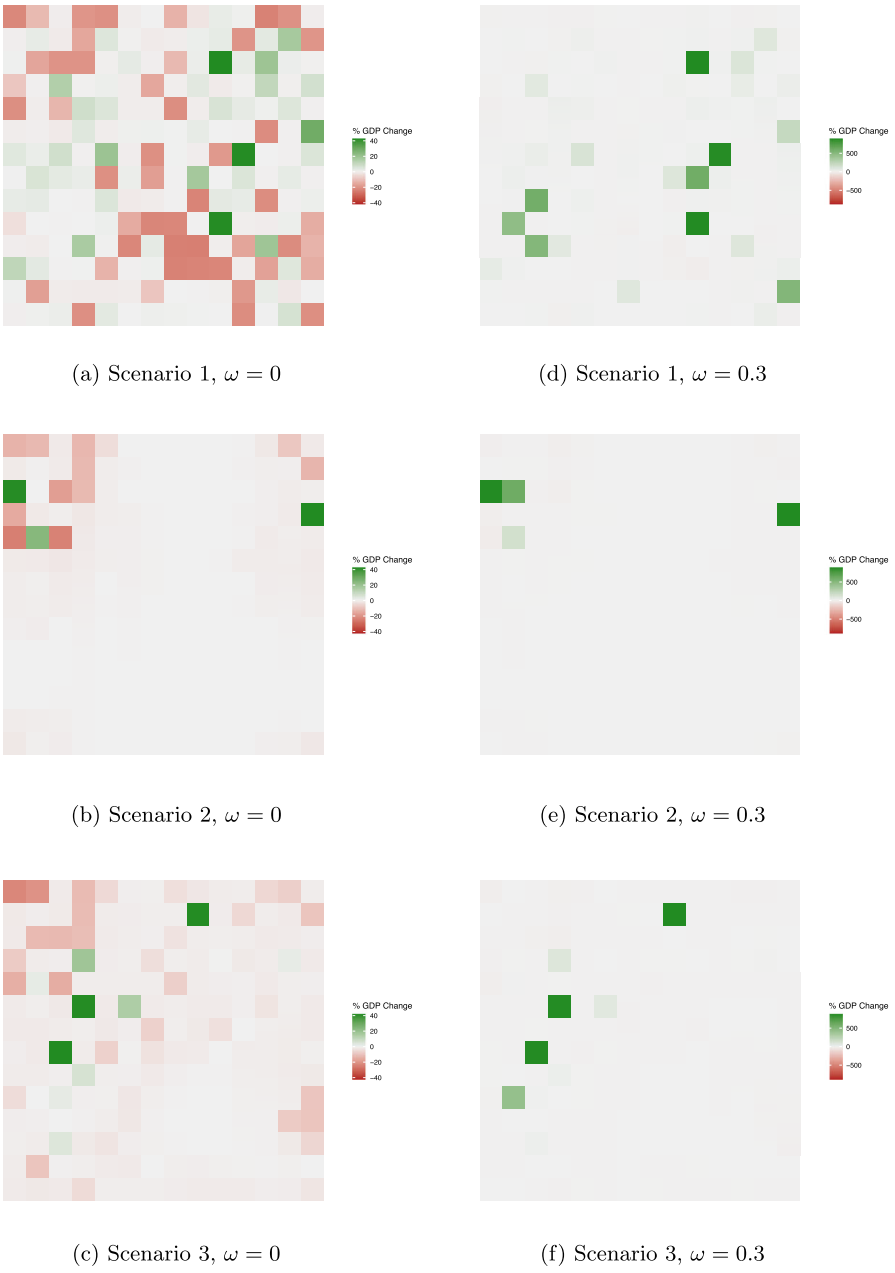


(a) Scenario 1, $\omega = 0$ (d) Scenario 1, $\omega = 0.3$

(b) Scenario 2, $\omega = 0$ (e) Scenario 2, $\omega = 0.3$

(c) Scenario 3, $\omega = 0$ (f) Scenario 3, $\omega = 0.3$

**Fig. 5** Percentage change in GDP at $t = 60$ with respect to $t = 0$. Seed = 1000

are widespread but uneven: some countries display sizeable negative deviations, while others remain close to their initial output level. A subset of cells (the green ones) shows relative output gains. These regions likely started with a high initial share of infected

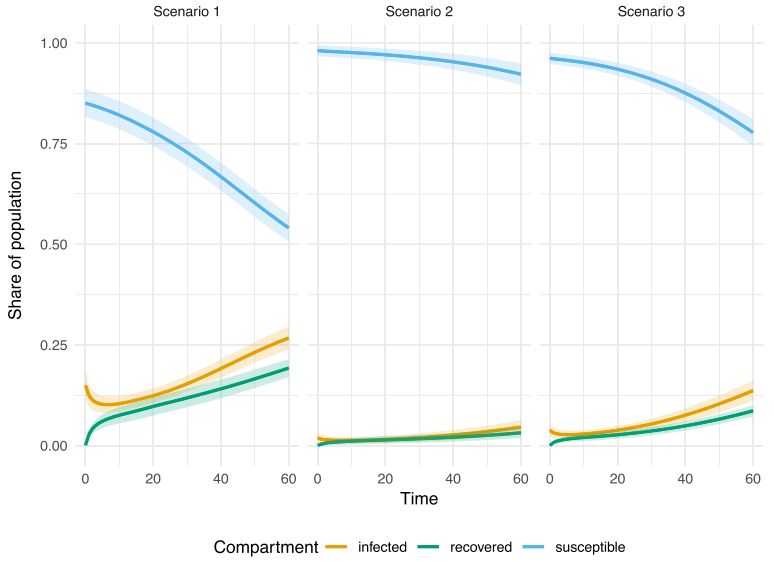


(a) $\omega = 0$

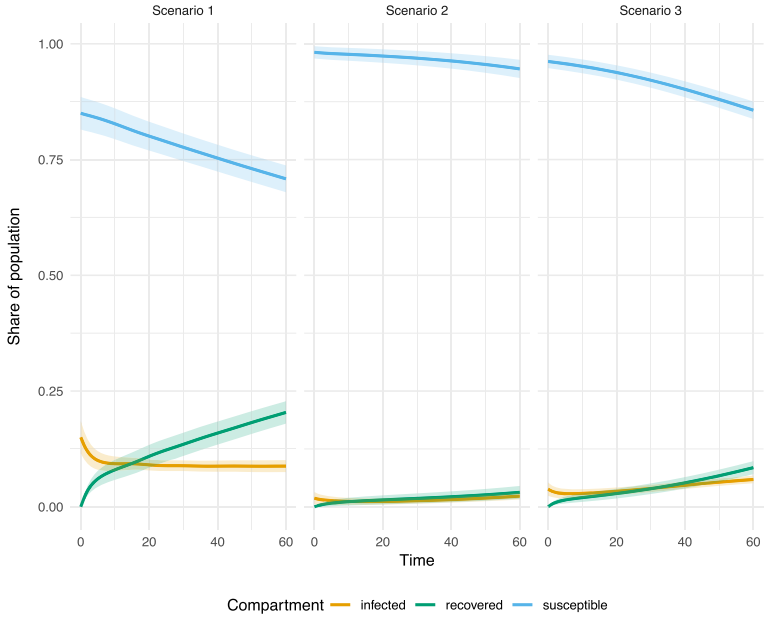


(b) $\omega = 0.3$

**Fig. 6** SIR dynamics. Averages and standard deviations of population compartments across 2000 simulations

individuals, which depressed their baseline GDP. As decentralized treatment successfully reduces local prevalence, these cells experience a "rebound" effect, converging toward their full productive capacity. Consequently, the observed green cells do not necessarily represent superior economic performance, but rather the normalization of output in previously high-burden areas once the infection is contained.

Scenario 2 exhibits a much more concentrated macroeconomic impact, confined to cells in proximity to the original epicenter, with most jurisdictions showing negligible GDP variation. Scenario 3 again produces an intermediate configuration, characterized by a patchwork of moderate contractions corresponding to localized outbreaks. Importantly, since $\omega = 0$, these economic disparities are not driven by policy coordination or imitation effects. They emerge endogenously from the interaction between local infection dynamics and budget-constrained optimization. In this decentralized benchmark, uniform pricing and identical institutional rules do not imply uniform economic outcomes. Instead, macroeconomic heterogeneity arises as an emergent property of spatial contagion combined with locally rational but uncoordinated decision-making.

Panel (a) in Fig. 6 summarizes the aggregate Monte Carlo dynamics under purely decentralized policies ($\omega = 0$), reporting averages across 2000 simulations over a horizon of $t = 60$ ticks. A notable result is that, in all three scenarios, the share of infected individuals increases over time (see Fig. 7a). Although treatment mitigates the speed of contagion, decentralized optimal decisions are not sufficient to reverse the upward trajectory of infection within the considered time window. This pattern indicates that, in the absence of spatial coordination, locally optimal treatment choices are only partially effective in counteracting inter-cell contagion. Each jurisdiction responds rationally to its own prevalence and budget constraint, but because infection spreads across neighboring cells, decentralized optimization generates a structural gap between local containment and global control. The disease therefore continues to diffuse across space, albeit at a moderated pace.[8]

At the same time, treatment activity produces a non-negligible accumulation of recovered individuals (see Fig. 7c). In all scenarios, the recovered share increases steadily over time. This reflects both the direct curing effect of treatment and the endogenous incentive for health authorities to purchase treatment in order to restore productive capacity. Given endogenous output evolution (Eq. 20), curing infected individuals yields an immediate macroeconomic benefit. The treatment decision therefore embodies a dual motive: epidemiological containment and GDP stabilization. The dynamics of the adjusted treatment rate $\lambda_t^{adj,v}$ (see Panel (a) in Fig. 8) further clarify this mechanism. Treatment intensity rises during phases of increasing infection and remains persistently positive throughout the horizon. However, the magnitude of $\lambda_t^{adj,v}$ is insufficient to offset the cumulative effect of spatial spillovers. In other words, treatment slows the epidemic but does not eliminate its growth under decentralized decision-making. Fiscal pressure evolves consistently with this interpretation.

[8] It should be noted that complete disease eradication is a rare historical outcome; to date, smallpox remains the only human infectious disease officially declared eradicated worldwide by the World Health Organization (WHO) (World Health Organization 1980). Consequently, the model's focus remains on containment and prevalence reduction rather than total elimination, which represents an exceptionally high bar for global health policy.

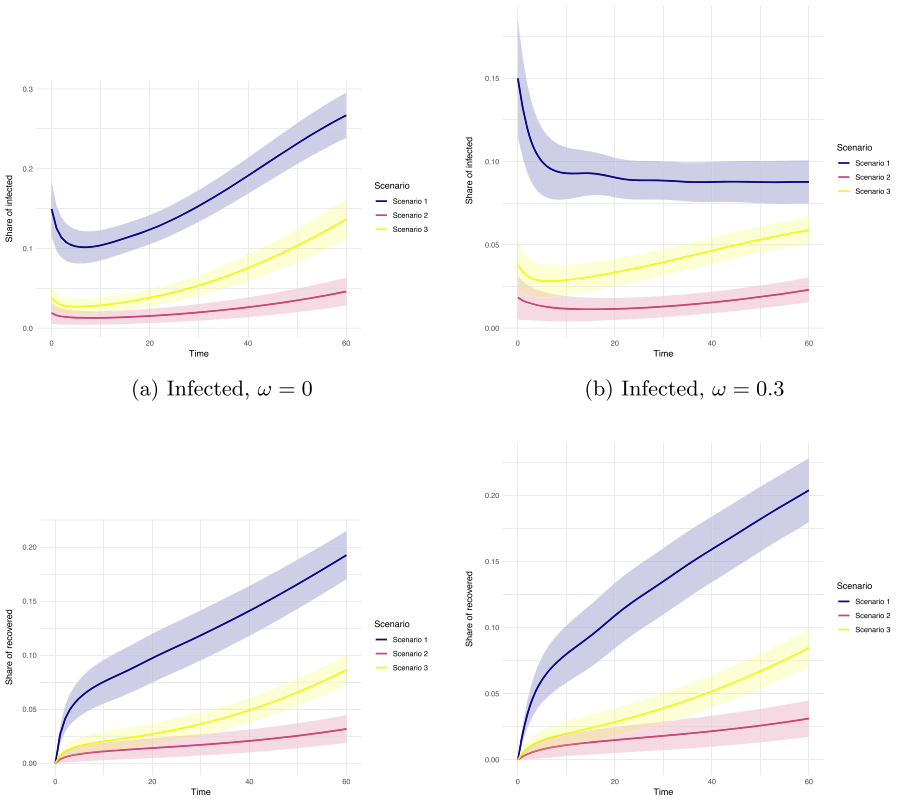


(a) Infected, $\omega = 0$ (b) Infected, $\omega = 0.3$

(c) Recovered, $\omega = 0$ (d) Recovered, $\omega = 0.3$

**Fig. 7** Average share of infected and recovered across cells for 2000 Monte Carlo simulations. *Shaded areas* represent $\pm$ one standard deviation

The cost-to-budget ratio,

$$\frac{p\lambda_t^{v,*} I_t^v}{b_t^v} = \frac{h_t^v}{b_t^v},$$

displays sustained positive values (Panel (c) in Fig. 8), reflecting ongoing expenditure rather than a temporary emergency response. Because infection does not collapse, treatment remains necessary, and fiscal resources remain engaged, although the ratio's level remains quite small as a consequence of the optimization solution. The number of policy-active cells (Panel (e) in Fig. 8) confirms that treatment is not episodic but structurally embedded in the medium-run dynamics. Activation remains widespread across regions, particularly in Scenarios 1 and 3, indicating that decentralized authorities continuously react to rising local prevalence. Taken together, the decentralized benchmark highlights a structural limitation of purely local optimization. Health authorities respond rationally to their own infection rates and budget constraints, and treatment yields a substantial number of recoveries. However, the aggregate trajectory of infection remains upward over the considered horizon. The system thus operates in a regime

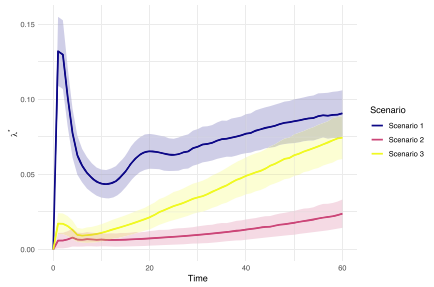


(a) $\lambda^{adj}, \omega = 0$

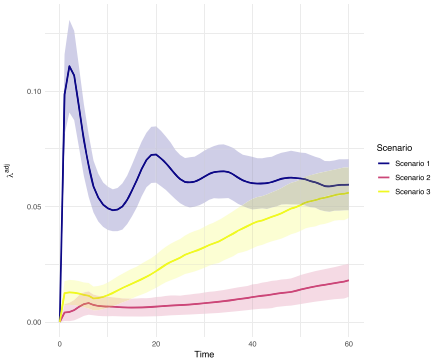


(b) $\lambda^{adj}, \omega = 0.3$

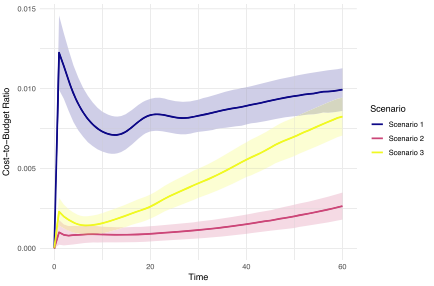


(c) Cost-to-budget ratio, $\omega = 0$

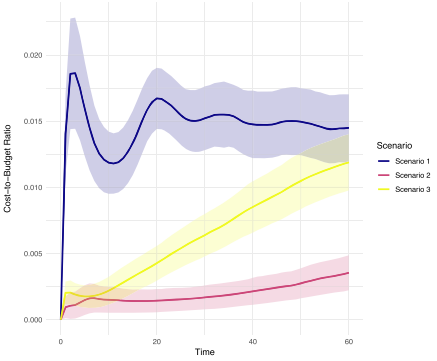


(d) Cost-to-budget ratio, $\omega = 0.3$

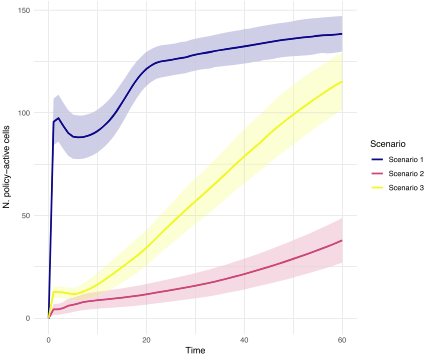


(e) N. policy-active cells, $\omega = 0$

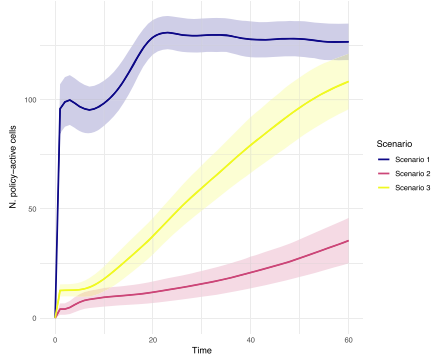


(f) N. policy-active cells, $\omega = 0.3$

**Fig. 8** Average share of $\lambda^{adj}$ across cells, cost-to-budget ratio across cells, and number of policy-active cells for 2000 Monte Carlo simulations. *Shaded areas* represent ± one standard deviation

of endogenous mitigation rather than containment: infection growth is slowed, but not halted, and cross-sectional disparities in both epidemiological burden and GDP persist.

This raises a natural question. If the limitation of the benchmark lies in the absence of coordination across neighboring jurisdictions, can even a bounded form of spatial policy interaction internalize part of the contagion externality? We address this question next by introducing local policy interdependence ($\omega > 0$).

### 4.2 Bounded spatial interaction ($\omega > 0$)

We now run the model under the specification of bounded spatial interdependence in treatment decisions, i.e., $\omega > 0$. Specifically, we set $\omega = 0.3$, so that each health authority implements a convex combination of its locally optimal treatment and the expected average treatment intensity observed among neighboring jurisdictions.[9] The introduction of spatial policy interaction produces a marked improvement in epidemiological outcomes in Scenarios 1 and 3. A comparison of the heat maps in Panels (d), (e), (f) of Fig. 3 with those obtained under $\omega = 0$ shows a notable reduction in the volume and spatial concentration of infected individuals. In both scenarios, infection levels are visibly lower and spatial clusters appear attenuated. We interpret the mechanism's effectiveness as stemming from two channels. First, by partially internalizing neighboring behavior, treatment effort becomes more synchronized across space, reducing the inter-jurisdictional gaps that previously allowed contagion to propagate. Second, it introduces a temporal anticipatory effect driven by a positive information externality. By incorporating the parameter $\omega > 0$, the intensity of neighboring interventions serves as a leading indicator—or a proxy signal—of the impending epidemiological risk. This allows an HA to preemptively scale its own response, effectively internalizing the information generated by its neighbors' actions to mitigate an incoming wave before it fully manifests in domestic data. Consequently, the coordination parameter serves as a channel for policy learning, in which the "shadow" of future infections is translated into immediate, proactive containment.

The effect is more nuanced in Scenario 2. Here, the initial outbreak is geographically concentrated, and many health authorities remain inactive because local prevalence stays below the activation threshold. When policy decisions incorporate expectations about neighboring treatment intensity, jurisdictions located near inactive neighbors may under-react relative to the purely decentralized case. As a consequence, infection can spill into additional districts that would otherwise have been more aggressively treated under $\omega = 0$. While the increase in infected cells is not dramatic, it highlights that spatial interaction can transmit not only containment effort but also complacency, depending on the initial conditions.

The benefits of policy interaction are even more evident on the economic side (Fig. 5, Panels d–f). In the reported run, no country exhibits a deterioration of GDP under $\omega > 0$ across the three scenarios. This reflects the stronger accumulation of recovered individuals, who return to productive activity and expand the effective labor

[9] We also conducted this analysis assuming HAs' expectation over the global average as the benchmark for policy interdependence. Since the results do not deviate significantly from the local neighborhood specification, they are omitted here for brevity but are available in the Appendix.

base. By synchronizing treatment intensity across neighboring jurisdictions, policy interaction accelerates recovery flows and dampens persistent productivity losses.

Monte Carlo simulations confirm and generalize these findings. In Scenario 1, as highlighted in Fig. 6, the average infected share no longer exhibits a monotonically increasing trajectory as under $\omega = 0$. Instead, it trends downward during the first years and subsequently levels off at a significantly lower baseline. The epidemic clearly enters a regime of regression rather than expansion. This inversion of the infection curve is accompanied by a pronounced acceleration in the recovered compartment. In Scenarios 2 and 3, although the infected curve remains upward sloping over part of the horizon, both the level and growth rate of infection are significantly reduced relative to the decentralized benchmark. At the same time, the recovered share rises more rapidly and reaches higher levels. Thus, even when infection does not fully reverse, bounded coordination generates a meaningful compression of epidemiological dispersion.

The Monte Carlo evidence on treatment intensity, cost-to-budget ratio, and policy activation further clarifies the role of bounded spatial interaction. A comparison with the decentralized benchmark ($\omega = 0$) reveals two systematic differences. First, the adjusted treatment rate $\lambda_t^{adj,v}$ becomes noticeably more oscillating under $\omega > 0$ (see Panel (b) in Fig. 8). In Scenario 1 in particular, treatment no longer follows the smooth rise-and-stabilize pattern observed under purely decentralized decisions. Instead, it exhibits visible oscillations over time. These fluctuations are mirrored in the cost-to-budget ratio (Panel (d) in Fig. 8), which now displays recurrent increases and declines rather than monotonic convergence toward a steady level. Second, the number of policy-active cells displayed in Panel (f) of Fig. 8 evolves in waves. The effects of bounded coordination differ markedly across scenarios. In Scenarios 2 and 3, where infection originates from one or a few localized outbreaks, the introduction of spatial policy interaction does not significantly alter either the level or the slope of the curve describing the number of policy-active cells. The trajectory remains increasing over time. Despite the higher average treatment intensity induced by coordination, the progressive diffusion of contagion from the epicenter(s) continues to activate new cells. In these configurations, spatial interaction amplifies intervention intensity but does not prevent the geographical expansion of the disease within the considered horizon. The contagion front advances outward, and an increasing number of countries cross the activation threshold defined in Eq. 21. The mechanism is different in Scenario 1. Because infection is initially widespread, coordination produces a stronger collective intervention response. The synchronized increase in treatment intensity reduces prevalence more effectively and earlier than under $\omega = 0$. As infection regresses, fewer cells exceed the activation threshold, and the number of policy-active cells begins to decline. In this case, bounded coordination not only intensifies treatment but also compresses the spatial domain of the epidemic, generating a weak reversal in the activation curve. Thus, while coordination always increases intervention effort, its effect on the extensive margin of policy activation depends critically on the spatial structure of the outbreak. When contagion is spatially concentrated, coordination strengthens local containment but does not halt outward diffusion. When contagion is initially pervasive, coordination induces a sufficiently strong synchronized response to push the system into regression, thereby reducing both infection levels and the number of active jurisdictions.

The oscillatory dynamics exhibited by treatment effort and cost-to-budget ratio can be understood directly from the structure of the decision rule introduced in Eq. 18. The presence of spatial expectations introduces both a temporal lag and a cross-regional feedback mechanism. Treatment decisions are no longer purely reactive to local prevalence but are now interdependent across cells and over time. Such lagged spatial feedback naturally generates cyclical adjustment. In Scenario 1, where infection is initially widespread and spatially homogeneous, this mechanism is particularly strong. When prevalence rises across a cluster of jurisdictions, coordinated treatment intensifies simultaneously due to shared expectations. The resulting synchronized intervention sharply reduces local infection. As prevalence falls, however, the convex combination rule induces a collective relaxation of effort. Because contagion continues to diffuse through neighboring cells, infection may re-emerge, triggering another coordinated response. This alternating amplification and relaxation produces wave-like oscillations in both $\lambda_t^{adj,v}$ and the cost-to-budget ratio (see Fig. 8, panels (b) and (d)). Importantly, these oscillations do not signal instability, but rather endogenous policy cycling.

In epidemiological terms, bounded coordination transforms monotonic diffusion into a regime of intermittent suppression, where the disease transitions from expansion to regression significantly earlier. The economic implications are equally important. Because treatment restores productive capacity through the recovery channel, synchronized policy responses generate temporary surges in expenditure that are offset by subsequent GDP gains. The fluctuations observed in the cost-to-budget ratio therefore represent an adaptive economic engagement process. In Scenarios 2 and 3, where infection remains more geographically localized, the oscillatory mechanism is weaker but still visible. Even when the infected share continues to grow over part of the time interval, its level and growth rate are substantially lower than under $\omega = 0$, while the recovered share increases more rapidly. Thus, bounded spatial interaction compresses epidemiological dispersion and accelerates macroeconomic recovery across cells. These results indicate that even modest spatial interdependence among decentralized health authorities can internalize part of the contagion externality, without requiring centralization, leading to stronger containment and improved macroeconomic outcomes.

### 4.3 Sensitivity analysis with respect to spatial interaction

To confirm the robustness of these results, we conduct a sensitivity analysis on the coordination parameter $\omega$. For increasing values of $\omega \in [0, 1]$, we perform 20 simulation runs for each scenario and compute the average values at $t = 60$ of three key variables: the adjusted treatment intensity $\lambda^{adj,v}$, the share of infected individuals, and the cost-to-budget ratio. Across all scenarios, moving from $\omega = 0$ to strictly positive values of $\omega$ leads to a reduction in the average final treatment intensity (see Fig. 9, Panel (a)). This confirms the mechanism discussed above: spatial coordination enhances the effectiveness of intervention, so that a lower level of individual treatment effort is sufficient to achieve improved epidemiological outcomes. At intermediate values of $\omega$, the model displays a substantial degree of robustness. The final

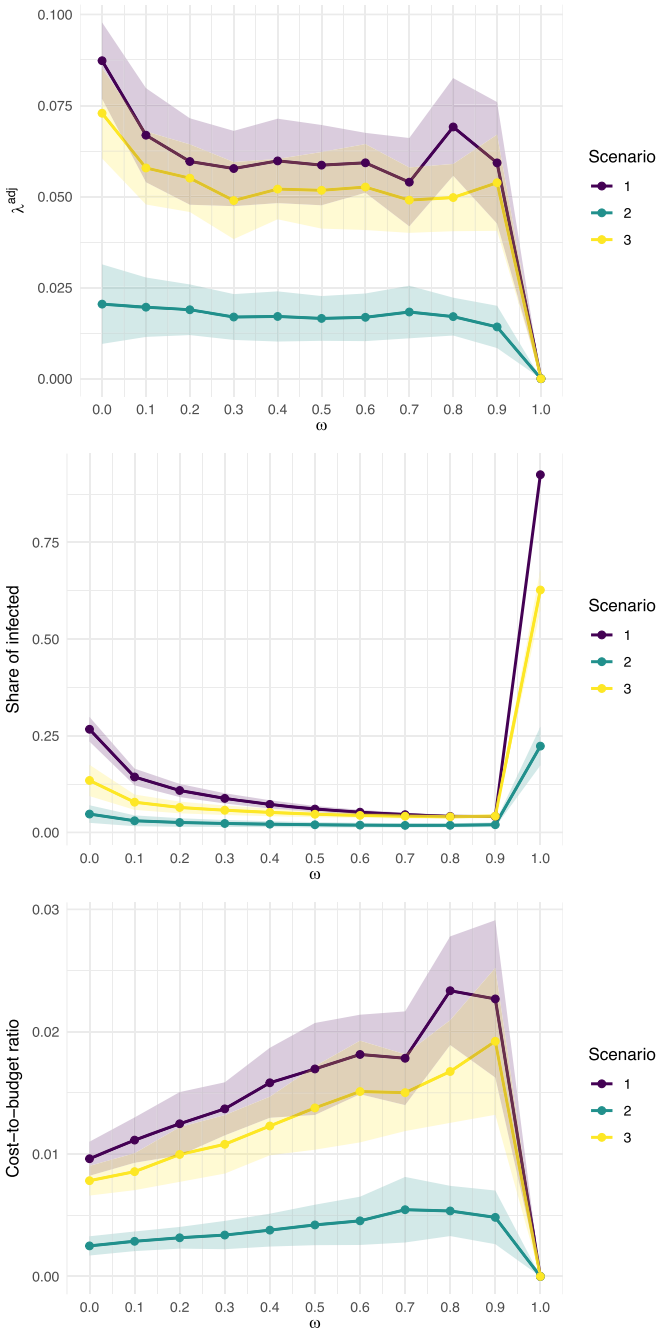


**Fig. 9** Sensitivity analysis for the parameter $\omega$. From *top to bottom*: final value at $t = 60$ of $\lambda^{adj}$, share of infected, and cost-to-budget ratio. Results are averaged across 20 runs with fixed seed with $\pm$ one standard deviation represented by the *shaded areas*

treatment effort, infection levels, and economic pressure vary only moderately, suggesting that modest degrees of coordination are sufficient to generate most of the epidemiological gains observed in the previous subsection. However, when $\omega$ becomes excessively large, meaning that treatment decisions rely predominantly on expectations about neighbors rather than on local optimization, new distortions emerge. In this regime, health authorities tend to underestimate local infection dynamics when neighboring jurisdictions appear relatively inactive. This generates coordination-induced complacency, which ultimately requires higher treatment effort at the final date among active cells. From a purely epidemiological standpoint, the sensitivity analysis reinforces the beneficial role of coordination. In all three scenarios, the average final share of infected individuals declines monotonically as $\omega$ increases (see Panel (b) in Fig. 9). Even when coordination induces temporary misalignments in local decisions, the overall diffusion of treatment effort across space improves containment and reduces medium-run prevalence. The economic implications, however, are more nuanced. Indeed, Panel (c) of Fig. 9 suggests that increasing $\omega$ generates a progressively larger share of jurisdictions that remain policy-active for longer periods. While coordination reduces infection, it also prolongs collective engagement in treatment, particularly at higher levels of $\omega$. As a consequence, the average cost-to-budget ratio at $t = 60$ increases with stronger coordination. This reflects a trade-off between epidemiological gains and fiscal exposure: better containment is achieved at the cost of more sustained budgetary commitment. Overall, the sensitivity analysis highlights a non-linear relationship between coordination and efficiency. Moderate levels of spatial interaction deliver significant epidemiological improvements with limited fiscal distortion. Excessive reliance on expectation-based decisions, however, may generate economic inefficiencies by weakening the informational role of local prevalence and inducing over-persistent activation. These findings underscore that while bounded coordination can substantially improve decentralized epidemic management, its intensity must remain balanced to avoid expectation-driven misallocation of resources.

An extreme case is $\omega = 1$, where local optimization is entirely replaced by neighborhood imitation. In this limit, since jurisdictions only implement a treatment intensity $\lambda^{adj,v}$ based on the expected behavior of their neighbors, the system becomes susceptible to an inertia trap. Since the initial treatment intensity at $t = 0$ is zero, the absence of an internal optimization signal means that no cell takes the first step to scale up intervention. Consequently, the treatment rate $\lambda^{adj,v}$ remains at zero, the cost-to-budget ratio vanishes, and the infection is allowed to propagate unchecked until prevalence reaches its maximum. This limit case illustrates a fundamental risk of extreme coordination: when policy becomes purely mimetic and decoupled from local incentives, the "information externality" that proved beneficial at lower values of $\omega$ collapses into a collective inaction trap, where jurisdictions essentially coordinate on a zero-effort equilibrium.

## 5 Spatially heterogeneous and time-varying pricing

In Section 4, the unit price of treatment was assumed to be homogeneous across countries and constant over time. While useful as a benchmark, this abstraction

neglects institutional features of pharmaceutical markets. In practice, procurement and reimbursement agreements for innovative treatments are typically structured over multi-year horizons rather than renegotiated annually. Multi-year supply contracts are widely used in global health markets to reduce uncertainty and administrative costs, and strategic price negotiations have played a central role in expanding access to antiretroviral therapy (ARV) and direct-acting antiviral (DAA) therapies (Assefa et al. 2018). Similarly, subscription-based and managed entry agreements often involve multi-year commitments and negotiated pricing arrangements (Auty et al. 2022). Motivated by these institutional features, we allow treatment prices to vary across countries and to be periodically updated. Specifically, prices are revised every five periods (interpreted as years), representing a stylized renegotiation cycle. Between revision dates, prices remain fixed, reflecting contractual rigidity. In addition, prices vary across countries according to their relative economic capacity. Differential pricing based on income levels is a well-documented practice in pharmaceutical markets and is often rationalized through Ramsey-type arguments balancing access and innovation incentives (Danzon and Towse 2003). Countries with below-average GDP face a baseline price, while richer jurisdictions pay an income-dependent surcharge. Because GDP evolves endogenously with infection dynamics, pricing becomes indirectly linked to epidemiological conditions, introducing an additional feedback channel between health and economic outcomes.

Formally, let $t_k = 5 \cdot k$, $k = 1, 2, \ldots$, denote renegotiation dates. At each $t_k$, the manufacturer sets country-specific prices based on relative income levels. Let the cross-country average GDP at time $t$ be

$$\overline{B}_t = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} B_t^v.$$

At each renegotiation date $t_k$, the country-specific treatment price is defined by the following piecewise rule:

$$p_{t_k}^v = \begin{cases} \bar{p} & \text{if } B_{t_k}^v \leq \overline{B}_{t_k} \\ \bar{p} + \eta \left( \dfrac{B_{t_k}^v}{\overline{B}_{t_k}} \right) & \text{if } B_{t_k}^v > \overline{B}_{t_k} \end{cases} \tag{26}$$

where $\bar{p} > 0$ is the baseline price and $\eta > 0$ measures the intensity of income-based price discrimination. Between renegotiation dates, prices are constant:

$$p_{t+1}^v = p_t^v \qquad \text{for } t \notin \{t_k\}.$$

We set $\bar{p} = 5$ and $\eta = 2$. A first robust result, visible when comparing the dynamics in Figs. 10, 11, 12 with those in the homogeneous-price benchmark described by Figs. 6, 7, 8, is that price discrimination tends to worsen (at least in levels, across all scenarios) the joint epidemiological and fiscal performance of decentralized containment. In the purely decentralized regime ($\omega = 0$), Section 4 already shows that

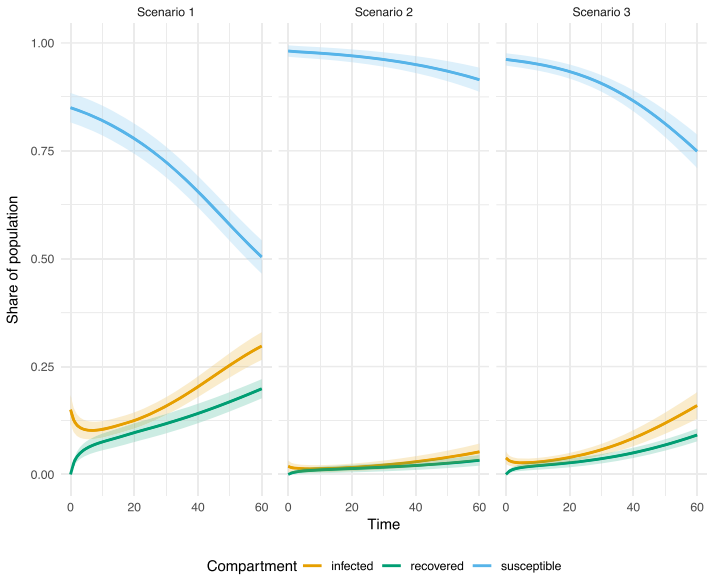


(a) $\omega = 0$

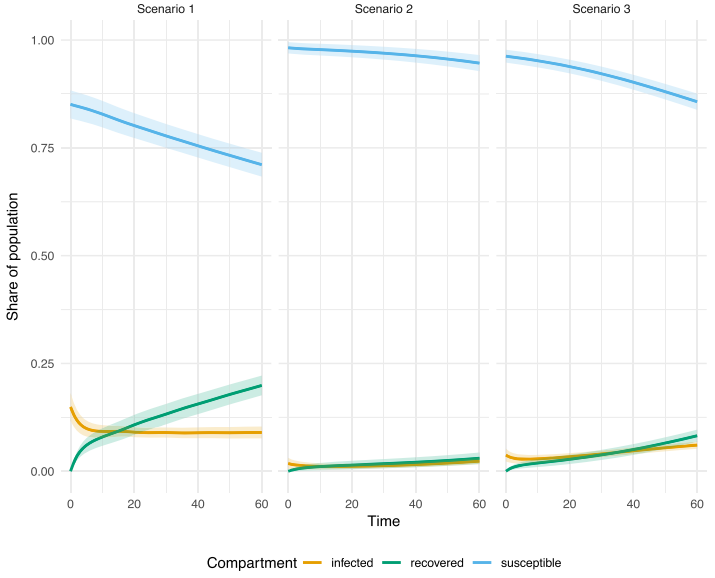


(b) $\omega = 0.3$

**Fig. 10** SIR dynamics under price discrimination. Averages and standard deviations across 2000 simulations

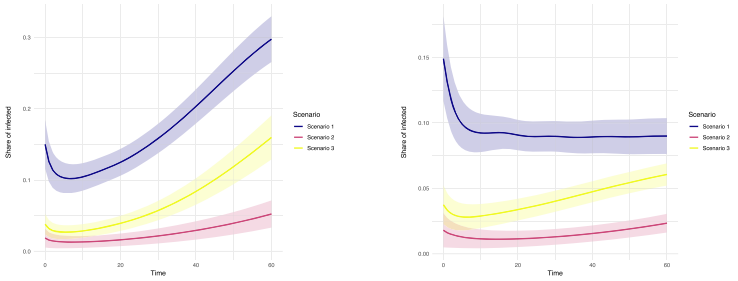


(a) Infected, $\omega = 0$




(b) Infected, $\omega = 0.3$

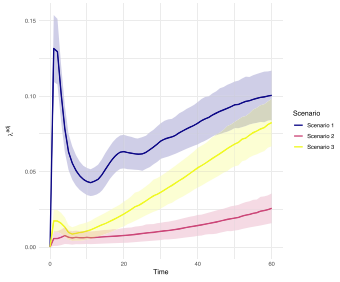


(c) $\lambda^{adj}, \omega = 0$

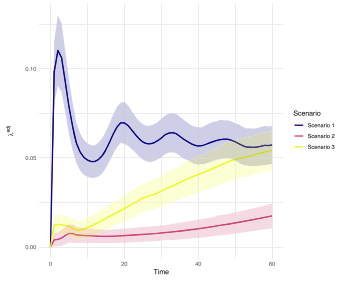


(d) $\lambda^{adj}, \omega = 0.3$

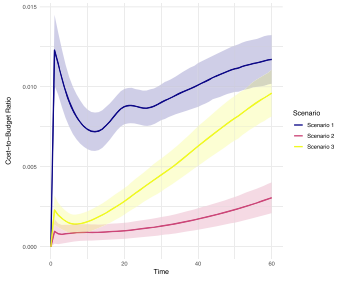


(e) Cost-to-budget ratio, $\omega = 0$

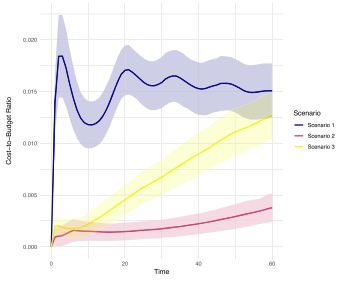


(f) Cost-to-budget ratio, $\omega = 0.3$

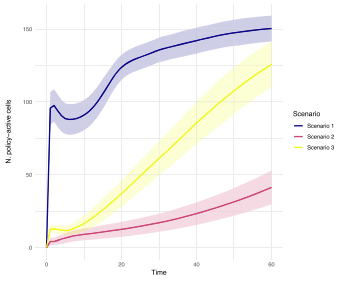


(g) N. policy-active cells, $\omega = 0$

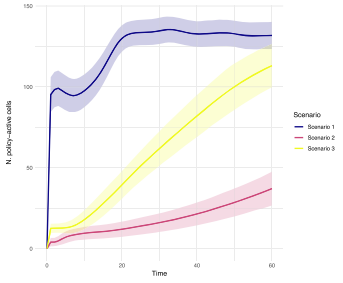


(h) N. policy-active cells, $\omega = 0.3$

**Fig. 11** Average share of infected over population, $\lambda^{adj}$, cost-to-budget ratio across cells, and number of policy-active cells for 2000 Monte Carlo simulations, under heterogeneous pricing. *Shaded areas* represent $\pm$ one standard deviation

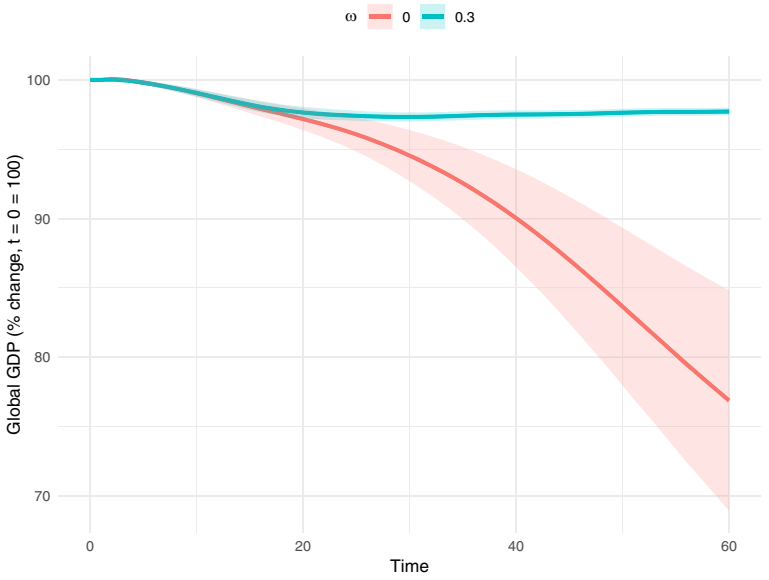


**Fig. 12** Percentage global GDP variation with respect to value at $t = 0$, average across 2000 runs. *Shaded areas* represent $\pm$ one standard deviation

infection continues to increase over the $t = 60$ horizon and that locally optimal treatment choices are only partially effective against inter-cell contagion (Figs. 6 and 7). Under heterogeneous pricing, the same qualitative failure of decentralization is amplified: compared to Figs. 7 and 8, Fig. 11 shows higher average infected shares, higher adjusted treatment effort, and a higher cost-to-budget ratio. This deterioration is not explosive in the plotted ranges, and this is consistent with the fact that heterogeneity in unit prices induces a larger mass of policy-active health authorities (see panel (g)), which partially offsets the contractionary impact of higher prices on treatment intensity and prevents the system from drifting into an even worse epidemiological state. The deterioration can also be observed by analyzing a single simulation run in detail. A comparison between the left-hand panels of Fig. 3 and those of Fig. 13 reveals a clear shift in the latter, which are characterized by a higher frequency of yellow-green cells, indicating a more widespread penetration of the disease.

Turning to $\omega > 0$ ($\omega = 0.3$), the comparison between the right panels of Fig. 11 and the corresponding coordinated benchmark in Figs. 7 and 8 confirms again that bounded spatial interaction remains beneficial even in the presence of price discrimination: in Scenario 1 the infected trajectory returns to a declining pattern within the 60-year horizon, while in Scenarios 2 and 3 infection becomes clearly "less increasing" than under $\omega = 0$, and the economic side improves accordingly through a lower fiscal burden and a more efficient evolution of the extensive margin of intervention, i.e. the number of cells actively engaged in the health policy (see panel (h) of Fig. 11).

The key finding is that price heterogeneity and temporal volatility undermine the benefits of treatment. By introducing persistent mismatches between local infection levels, available budget, and unit costs, these price dynamics erode the gains seen under

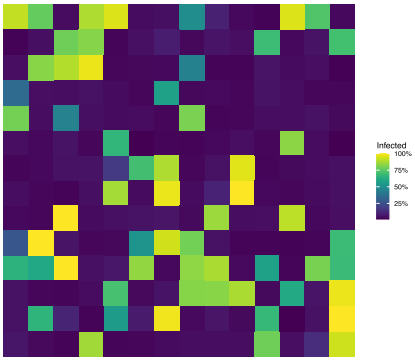


(a) Scenario 1, $\omega = 0$

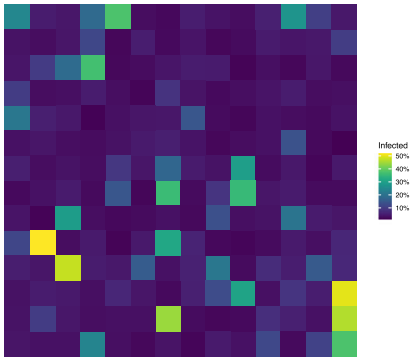


(d) Scenario 1, $\omega = 0.3$

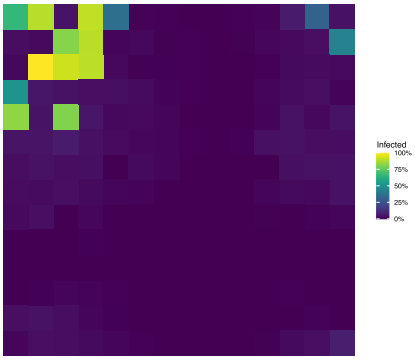


(b) Scenario 2, $\omega = 0$

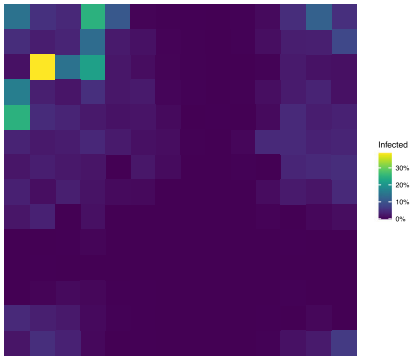


(e) Scenario 2, $\omega = 0.3$

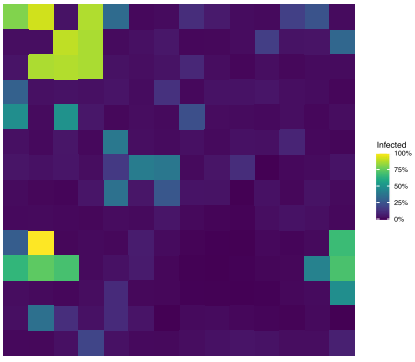


(c) Scenario 3, $\omega = 0$

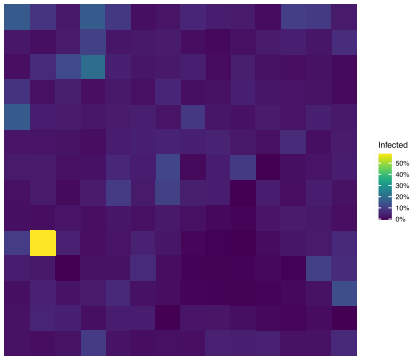


(f) Scenario 3, $\omega = 0.3$

**Fig. 13** Final spatial distribution of infection across cells under heterogeneous pricing at $t = 60$ (Seed = 1000)

uniform pricing. Although from an epidemiological point of view the deterioration appears to be minor, this comes with an important economic cost and, in fact, it is particularly evident in the aggregate economic dynamics: The figure shows that global GDP declines in both regimes, but with $\omega > 0$ it stabilizes at levels not too distant from the baseline, whereas under $\omega = 0$ the contraction is monotone throughout the horizon. Therefore, we observe that price discrimination combined with multi-year contractual rigidity acts as a vehicle of amplification for the inability of pure decentralization to curb diffusion: it moderately raises the medium-run infection burden and greatly worsens the country's economic outlook, making the mitigating role of even limited coordination (via neighborhood expectations) far more salient in sustaining recovery and preventing persistent GDP losses.

## 6 A focus on spatial concentration

While the focus so far has been on the temporal evolution of the disease, understanding the geographical footprint of the infection is crucial to evaluating the effectiveness of containment policies. In this section, we shift our focus to the spatial dimension, analyzing how the infected population is distributed across the grid.

To quantify the degree of spatial concentration, we employ the Herfindahl–Hirschman Index (HHI), typically used to measure market concentration. In our spatial framework, we define the index as

$$HHI = \sum_{v=1}^{\mathcal{V}} \left( \frac{I^v}{\sum_{v=1}^{\mathcal{V}} I^v} \right)^2 . \tag{27}$$

The index ranges from $1/\mathcal{V}$ (perfectly uniform distribution) to 1 (complete concentration in a single patch). In our results, an HHI close to 1 at $t = 0$ signifies that the disease starts from a highly localized cluster, while a decline toward values below 0.2 at $t = 60$ indicates a significant spatial dispersion of the pathogen. Figure 14 illustrates the distribution of this index across the three different scenarios and interaction levels ($\omega$).

At the beginning of the simulation ($t = 0$), the HHI shows a high variance in Scenario 2 and a much higher median than the other scenarios, while in Scenarios 1 and 3, the HHI is significantly lower, indicating a more dispersed initial distribution (see Fig. 14a). This reflects the known differences in the initial conditions.

By $t = 60$, the spatial concentration reaches a more stable state across all scenarios, though the direction of the trend depends on the initial setup (Fig. 14b and c). In Scenario 1 we observe an upper shift and a stretch of the box displaying the distribution of the HHI. A stretch is also visible for Scenario 3. This suggests that as the disease progresses, the infected population tends to polarize into specific cell types. Therefore, even when treatment and spatial interactions among institutions ($\omega > 0$) are effective in controlling the spread of the disease along the temporal dimension as seen in the previous Sections, they do not necessarily prevent the emergence of local health disparities. In these scenarios, the infection persists in a state of moderate

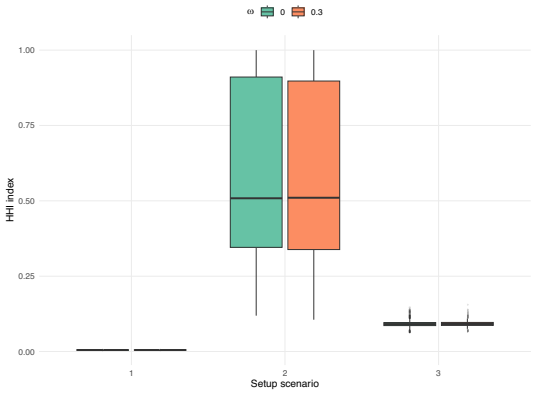


(a) $t = 0$

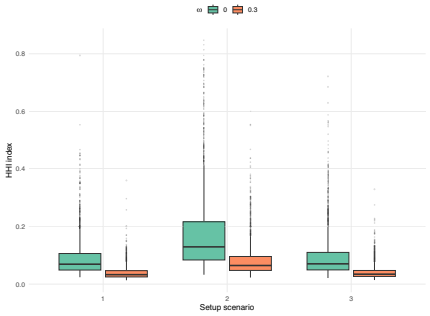


(b) Homogeneous price, $t = 60$

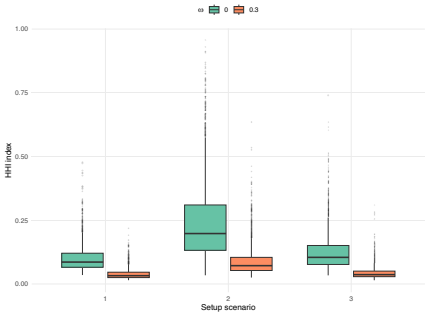


(c) Heterogeneous price, $t = 60$

**Fig. 14** Distribution of the Herfindahl–Hirschman Index (HHI) across 2000 independent simulations. The plot illustrates the degree of spatial concentration of the infected population

spatial heterogeneity, with a small subset of patches hosting a disproportionate share of cases.

Conversely, in Scenario 2, the high initial concentration (median HHI $\approx$ 0.50, indicating intense concentration) is significantly mitigated by $t = 60$. Here, effective treatment and local coordination dilute the original hotspot, bringing the HHI down to a moderate level below 0.20, with a stronger effect of $\omega$ in the case of heterogeneous and time-varying price (Fig. 14c). This is in line with the findings discussed in Section 4.2: when contagion is spatially concentrated at setup, coordination does not halt the disease transmission. As it spreads, the initial concentration weakens.

Across all scenarios, in fact, coordination among HAs ($\omega = 0.3$, orange) consistently maintains the HHI at lower levels than the case of $\omega = 0$ (green). Internalizing a certain degree of homophily within the decision over $\lambda^{adj}$ helps homogenize epidemic outcomes across space and thus prevents disparities across countries.

## 7 Endogenous treatment efficacy and R&D investment

In the previous sections, treatment efficacy was treated as an exogenous parameter. We now relax this assumption and introduce an endogenous innovation channel whereby the pharmaceutical manufacturer invests in research and development (R&D) to improve treatment effectiveness over time. The inclusion of R&D is grounded in a well-established economic literature linking market incentives to pharmaceutical innovation. Both theoretical and empirical contributions show that expected market size and potential revenues are central determinants of firms' R&D investment decisions (see Acemoglu and Linn 2004; Dubois et al. 2015). In the context of infectious diseases—where social returns are often large, and demand may fluctuate over time—additional mechanisms such as advance market commitments and subscription-based procurement schemes have been proposed to strengthen innovation incentives (Kremer et al. 2014; Sciarretta et al. 2016).

To capture this innovation channel parsimoniously, we assume that the manufacturer uses the extra revenues generated by income-based price discrimination (relative to a baseline price) to finance R&D. Higher R&D spending increases treatment efficacy, which we model as an endogenous, time-varying parameter. This extension introduces a feedback mechanism linking market outcomes and innovation to epidemiological dynamics: pricing affects manufacturer revenues, revenues fund R&D, R&D improves efficacy, and higher efficacy affects recovery and subsequent economic performance.

Let $\bar{p}$ denote the baseline price introduced in Section 5 and let $p_t^v$ be the effective price paid by country $v$ at time $t$. Define the extra-price component as

$$\psi_t^v := p_t^v - \bar{p}. \tag{28}$$

By construction, $\psi_t^v \geq 0$ and $\psi_t^v > 0$ only for above-average GDP countries.

At time $t$, country $v$ purchases $\lambda_t^{adj,v} I_t^v$ treatment units. Therefore, the manufacturer's extra revenue generated by price discrimination in country $v$ is $\psi_t^v \lambda_t^{adj,v} I_t^v$. Aggregating across all countries yields total extra revenues:

$$R_t := \sum_{v \in \mathcal{V}} \psi_t^v \, \lambda_t^{adj,v} \, I_t^v. \tag{29}$$

We assume that the manufacturer reinvests a fraction $\theta \in [0, 1]$ of extra revenues into R&D. Hence, the R&D budget is

$$F_t := \theta \, R_t. \tag{30}$$

Treatment efficacy is now time-varying and evolves endogenously according to

$$\chi_{t+1} = 1 - (1 - \chi_t) \left(1 + \frac{F_t}{\bar{F}}\right)^{-k}, \tag{31}$$

where $k > 0$ controls the responsiveness of innovation to R&D spending and the parameter $\bar{F}$ represents the technological scale of pharmaceutical innovation. It cap-

tures the magnitude of annual R&D resources required to generate a meaningful improvement in treatment efficacy. When annual premium revenues $F_t$ are small relative to $\bar{F}$, improvements in efficacy are negligible. Only sustained and sufficiently large investment flows produce gradual convergence toward the technological frontier.

The introduction of endogenous R&D fundamentally alters the interpretation of the heterogeneous pricing regime analyzed in Section 5. While price discrimination was previously shown to exacerbate epidemiological and fiscal imbalances, it now generates an additional dynamic channel through which higher revenues finance innovation and progressively enhance treatment efficacy.

Figure 15 illustrates this mechanism. As R&D investment accumulates over time, the efficacy parameter $\chi_t$ increases steadily across scenarios (see Fig. 15f), approaching higher effectiveness levels within the 60-period horizon. This endogenous improvement translates into visible epidemiological gains. Compared to the heterogeneous-price case without R&D (Section 5), infection dynamics become more favorable: in Scenario 1, the infected share declines more rapidly and stabilizes at a lower level (Fig. 15a), while in Scenarios 2 and 3 the growth of infection is markedly dampened. The recovered share correspondingly increases (Fig. 15b), reflecting both direct curing and improved efficacy.

The interaction between innovation and decentralized decision-making is particularly evident in treatment intensity and fiscal variables. Figures 15c and d show that although $\lambda_{adj}$ remains positive throughout the horizon and may initially be volatile (particularly in Scenario 1), the cost-to-budget ratio progressively stabilizes as treatment efficacy improves. This effect arises from two reinforcing mechanisms. A higher $\chi_t$ increases the direct recovery rate per treated individual and simultaneously reduces transmission through the infection channel. As a result, the infection stock $I_t^v$ declines more rapidly than under exogenous efficacy. Since treatment expenditure is proportional to $p\,\lambda_{adj,t}^v I_t^v$, the shrinking infection base offsets the need for continuously rising treatment intensity. At the same time, improved containment supports GDP stabilization, preventing further erosion of fiscal capacity. Hence, endogenous gains in efficacy raise the marginal productivity of health expenditure and dampen the upward fiscal pressure observed in the heterogeneous-pricing regime without R&D. The effect of endogenous innovation is also evident in the dynamics of policy-active cells (Fig. 15e). Compared to the heterogeneous-pricing regime without R&D, the number of active jurisdictions tends to decline over time as treatment efficacy improves. Higher $\chi_t$ accelerates the reduction in local prevalence, so that fewer cells exceed the activation threshold defined in Eq. 21. As a result, innovation affects not only the intensity of intervention but also its spatial extension, progressively shrinking the geographical footprint of active treatment. This contrasts with the no-innovation case of Section 5, where price heterogeneity sustained widespread and persistent policy activation.

Monte Carlo simulations of SIR dynamics displayed in Fig. 16 confirm that this innovation channel systematically compresses infection trajectories across all scenarios. Relative to the heterogeneous pricing regime of Section 5, the system now exhibits a clear medium-run improvement driven by the co-evolution of revenues and technological progress. Therefore, this mechanism mitigates the previously identified amplification effect of price discrimination. The reinvestment of surplus revenues into R&D generates a delayed but persistent corrective force.

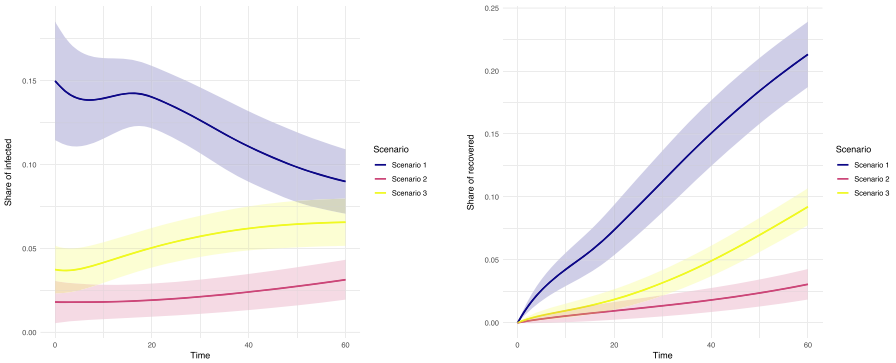


(a) Infected




(b) Recovered

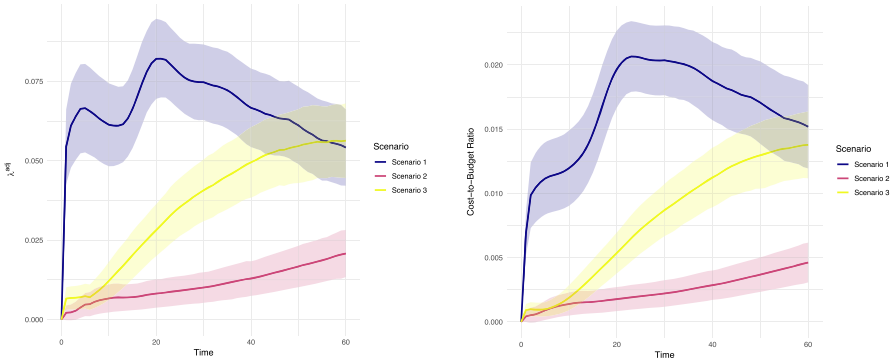


(c) $\lambda^{adj}$




(d) Cost-to-Budget ratio

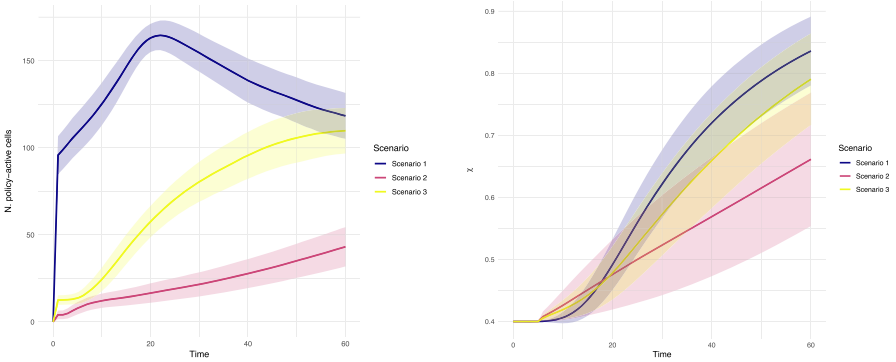


(e) N. policy-active cells




(f) Efficacy of treatment

**Fig. 15** Average share of infected over population, $\lambda^{adj}$, cost-to-budget ratio across cells, and number of policy-active cells for 2000 Monte Carlo simulations, under heterogeneous pricing and endogenous $\chi$. *Shaded areas* represent $\pm$ one standard deviation. $\chi_{t=0} = 0.4$, $\omega = 0.3$, $\theta = 0.5$, $\bar{F} = 1000$, $k = 0.01$

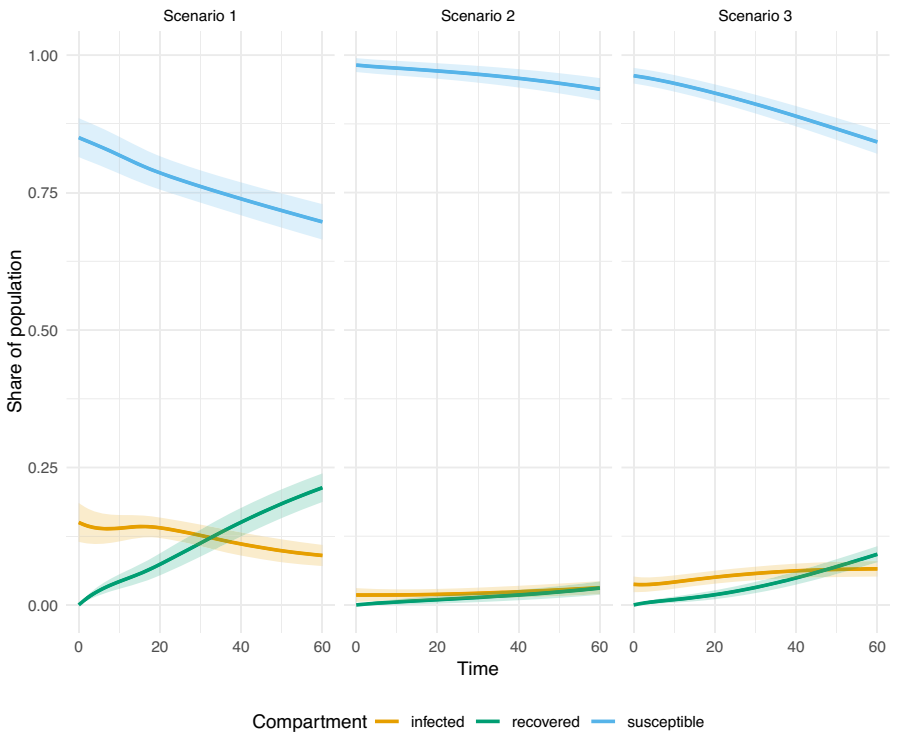


**Fig. 16** SIR dynamics under heterogeneous pricing and endogenous treatment efficacy. Averages and standard deviations of population compartments across 2000 simulations. $\chi_{t=0} = 0.4$, $\omega = 0.3$, $\theta = 0.5$, $\bar{F} = 1000$, $k = 0.01$

In economic terms, the model embeds a Schumpeterian feedback mechanism in which market outcomes drive technological progress. Spatially differentiated prices generate heterogeneous revenues; these revenues finance R&D; innovation enhances treatment efficacy; and improved efficacy reduces infection and restores productive capacity. This logic at least partially parallels Schumpeterian growth models, where expected profits provide incentives for innovation (see Aghion and Howitt 1992). In our framework, the effective "market size" is endogenously determined by epidemiological and pricing dynamics, so that higher revenues expand R&D investment and increase the productivity of health expenditure through rising $\chi_t$. This mechanism partially reconciles the tension identified in Section 5 between equity and efficiency. While price discrimination may worsen medium-run containment, sustained reinvestment of surplus revenues can generate long-run epidemiological gains via endogenous technological progress.

## 8 Summary of results

Given the number of simulation exercises considered in the paper, Table 2 provides a compact synthesis of the main mechanisms and outcomes emerging from

**Table 2** Summary of the main simulation results

| Regime | Institutional configuration | Epidemiological outcomes | Macroeconomic and fiscal outcomes | Main mechanism |
|---|---|---|---|---|
| Uniform and constant pricing (Section 4) | Pure decentralization ($\omega = 0$) | Infection increases over time in all scenarios, although treatment mitigates its growth. Spatial clusters and right-skewed infection distributions persist, especially in Scenarios 1 and 3. | GDP effects are heterogeneous across jurisdictions. Treatment expenditure remains positive but locally constrained, generating persistent fiscal engagement. | Local optimization slows contagion but does not internalize spatial spillovers. |
| Uniform and constant pricing (Section 4) | Bounded spatial interaction ($\omega = 0.3$) | Scenario 1 shifts toward a declining infection path, while Scenarios 2 and 3 display lower infection levels and compressed growth rates. | Cross-sectional GDP dispersion across jurisdictions is reduced, but treatment and cost-to-budget ratios may display oscillatory policy cycles due to lagged spatial expectations. | Local policy interaction acts as a coordination device and partially internalizes neighboring epidemiological risks. |
| Spatially heterogeneous pricing (Section 5) | Pure decentralization ($\omega = 0$) | Price heterogeneity amplifies infection dynamics relative to uniform pricing, increasing medium-run prevalence and spatial disparities. | Higher treatment prices tighten local budget constraints, reduce feasible treatment intensity, and contribute to more persistent GDP losses. | Pricing rules interact with fiscal constraints and distort local containment incentives. |
| Spatially heterogeneous pricing (Section 5) | Bounded spatial interaction ($\omega = 0.3$) | Coordination substantially dampens the negative epidemiological effects of price heterogeneity. Scenario 1 returns to a declining path, while Scenarios 2 and 3 experience slower diffusion. | GDP remains closer to initial levels, and fiscal pressure is reduced along the extensive margin of policy activation. | Spatial interaction mitigates both contagion externalities and price-induced treatment inequalities. |

**Table 2** continued

| Regime | Institutional configuration | Epidemiological outcomes | Macroeconomic and fiscal outcomes | Main mechanism |
|---|---|---|---|---|
| Spatial concentration of infection (Section 6) | Across pricing and coordination regimes | Outcomes depend strongly on the initial spatial distribution of infections. Single-epicenter outbreaks are more easily contained than widespread or multiple-outbreak configurations. | Economic losses are uneven and reflect the geography of contagion, recovery, and local budget capacity. | Initial spatial conditions generate path dependence and persistent heterogeneity in epidemiological and economic outcomes. |
| Endogenous efficacy and R&D (Section 7) | Bounded spatial interaction ($\omega = 0.3, \theta = 0.5$) | Rising treatment efficacy progressively compresses infection trajectories across all scenarios and reduces the persistence of active outbreaks. | Price premia create short-run fiscal pressure but finance innovation that improves long-run containment. | Surplus revenues invested in R&D operate as a Schumpeterian corrective mechanism linking market design to endogenous treatment effectiveness. |

the different model specifications. The table is not intended to replace the detailed analysis presented above, but rather to support the comparison of the epidemiological and economic implications of spatial interaction, pricing rules, and endogenous innovation across scenarios.

## 9 Conclusions

This paper develops a spatially structured agent-based model in which decentralized health authorities (HAs) allocate a costly treatment under local budget constraints to manage an infectious disease evolving on a two-dimensional lattice with Moore-neighborhood spillovers. Within each cell, disease dynamics follow a discrete-time SIR structure, while cross-cell interactions make local prevalence depend on neighboring infection levels; at each time step, HAs endogenously choose treatment intensity as a function of local epidemiological conditions, fiscal capacity, and the prevailing price of treatment. The framework is designed to jointly capture spatial contagion, institutional fragmentation, and the economic constraints that shape real-world epidemic responses.

The simulation analysis yields three main results. First, under homogeneous and time-invariant pricing, decentralized decision-making is generally insufficient to eliminate infection within the finite and policy-relevant horizon considered: spatial spillovers sustain transmission across cells and generate persistent cross-location heterogeneity in both infection and treatment. This combination of institutional fragmentation and spatial contagion is precisely what makes the agent-based framework informative in this context. In a standard mean-field SIR model, where the population is assumed to be homogeneous and well-mixed, only aggregate infection dynamics would be observed. Such an approach would conceal the local feedback loops generated by heterogeneous budgets, asynchronous treatment decisions, and neighborhood spillovers. In contrast, the ABM reveals how localized fiscal constraints and spatial transmission interact to produce persistent regional disparities, clusters of infection, and uneven containment paths that would be invisible in an aggregate representation.

Second, introducing local policy interaction through neighborhood-based adjustment of decisions (institutional coordination) improves outcomes, but with non-trivial trade-offs. When HAs partially incorporate neighborhood behavior into implemented treatment, the epidemiological burden can be reduced, and the spatial diffusion of infection attenuated, especially in scenarios where outbreaks are initially widespread. At the same time, stronger interaction may increase the persistence of collective intervention and shift fiscal pressure over time, underscoring that coordination affects both the intensity and the spatial extension of policy responses.

Third, moving from uniform to spatially heterogeneous and time-varying prices (interpreted as multi-period procurement renegotiations) predictably worsens epidemiological and economic performance in the medium run. This result aligns with standard theory: differential pricing interacts with decentralized budget constraints by distorting treatment incentives and amplifying spatial inequality in both disease prevalence and economic outcomes, thereby reinforcing the baseline inefficiency of uncoordinated

control. Crucially, however, the model demonstrates that institutional interaction can act as a stabilizing force. By partially internalizing these heterogeneities through coordinated responses, policy interaction helps mitigate the amplification of the outbreak. In practical terms, such coordination may also reduce the incentives for parallel trade or informal cross-border patient flows—phenomena that typically arise when wide price disparities make formal local procurement unsustainable.

Finally, the last extension endogenizes treatment efficacy by introducing an innovation channel: extra revenues generated by GDP-based price discrimination finance R&D, which increases efficacy over time. This mechanism alters the interpretation of price discrimination. While heterogeneous pricing may initially raise infection and fiscal stress, reinvestment can generate a delayed corrective force: as $\chi_t$ improves, infection paths compress across scenarios, and the geographic footprint of active intervention shrinks, because fewer regions remain above the activation threshold.

From a broader perspective, the main message of the work is that epidemic policy design cannot be evaluated separately from institutional structure and market conditions. Spatial contagion and decentralization generate persistent inequalities and make eradication harder within realistic horizons; coordination can substantially improve containment but may increase the persistence of intervention; and pricing schemes shape both access and fiscal pressure, with the possibility that innovation-financed efficacy improvements partially reconcile short-run distortions with longer-run gains. These results point to a policy implication consistent with the model's structure: effective management of high-cost drugs for infectious diseases requires simultaneously addressing (i) spatial spillovers, (ii) fragmented decision-making, and (iii) the incentives created by pricing and procurement institutions.

Future work could extend this framework to account for more realistic mobility patterns, information frictions across local authorities, strategic pricing behavior by pharmaceutical firms, and learning or adaptive expectations among agents. The model may also serve as a computational framework to explore the performance of decentralized versus coordinated epidemic policies under alternative institutional and fiscal architectures.

## Appendix

## Appendix A: SIR dynamics

Figure 17 shows the SIR dynamics for a single run under the baseline specification of homogeneous price (left panels) in comparison with the model in the absence of treatment, i.e., $\lambda_t^v = 0$. Figure 18 displays the evolution of the global share of infected (panel a) and the average share of infected across cells for 2000 runs, both in the absence of treatment.

These figures highlight the validity of the treatment assumption and show the predictable SIR dynamics in absence of cure.

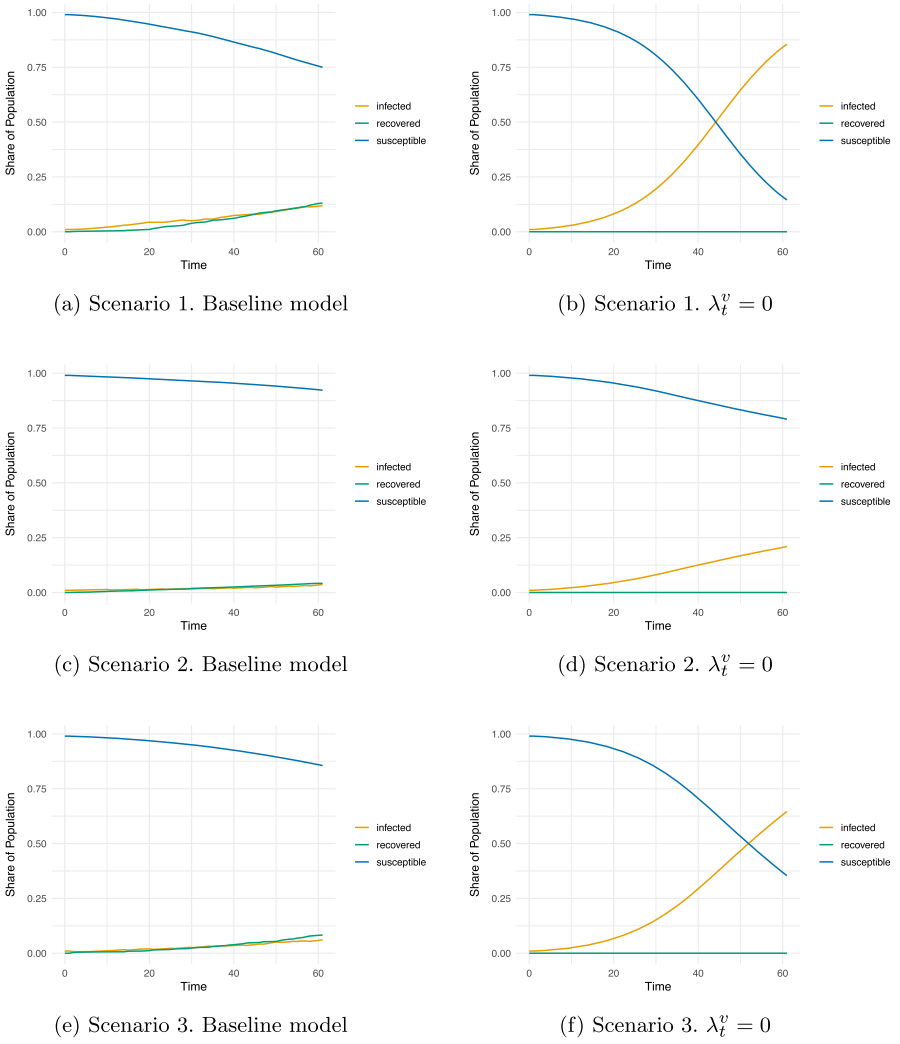


(a) Scenario 1. Baseline model

(b) Scenario 1. $\lambda_t^v = 0$

(c) Scenario 2. Baseline model

(d) Scenario 2. $\lambda_t^v = 0$

(e) Scenario 3. Baseline model

(f) Scenario 3. $\lambda_t^v = 0$

**Fig. 17** Comparison of infection dynamics across the three different scenarios. Seed = 1000

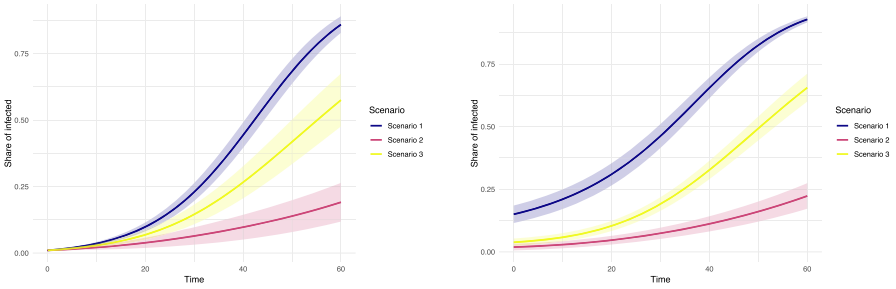


(a) Global share of infected. $\lambda^* = 0$ (b) Average share of infected. $\lambda^* = 0$

**Fig. 18** Average share of infected across 2000 Monte Carlo simulations, $\lambda^* = 0$

## Appendix B: Sensitivity analysis

We conducted a sensitivity analysis on a set of model parameters, displayed in Fig. 19. While the model behavior shown in the figure is expected in response to variation

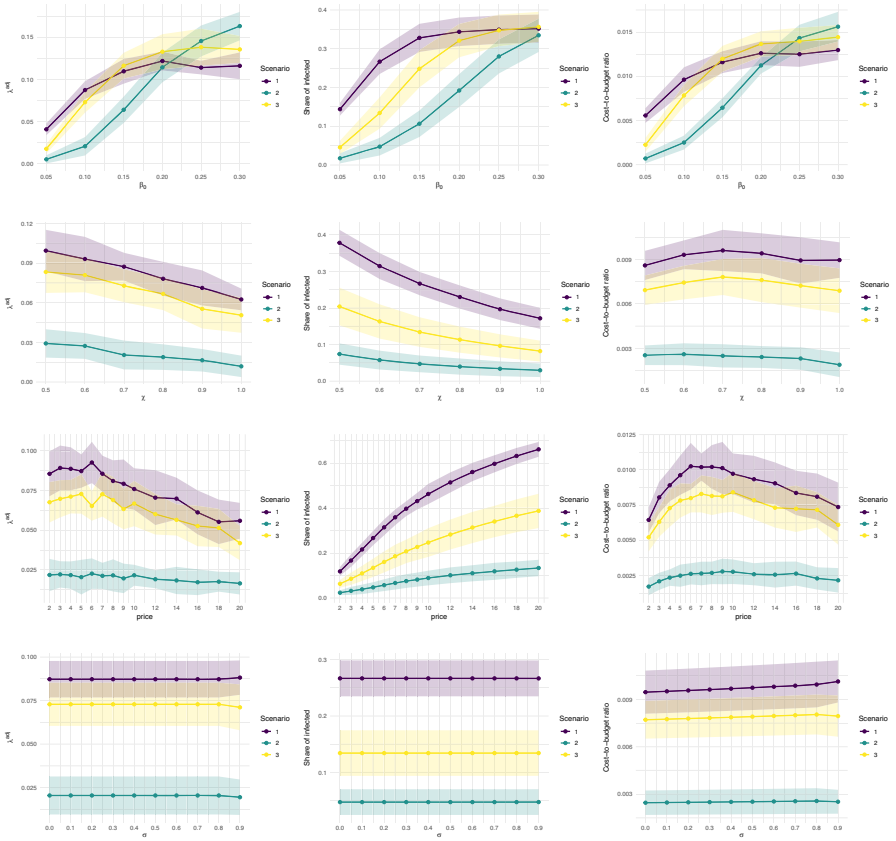


**Fig. 19** Sensitivity analysis. *Rows* correspond to parameters tested; columns report (*left to right*) λ, share of infected, and cost-to-budget ratio at the final step (mean across 20 seeds with ± SD)

in $\beta_0$, $\chi$, and price, the analysis of $\sigma$, the parameter representing the idiosyncratic productivity loss of infected individuals, is somewhat less trivial.

As illustrated in the bottom row of Fig. 19, the model exhibits a notable lack of sensitivity to this parameter; the trajectories for the adjusted treatment rate $\lambda^{adj}$ follow a stable path across the different values tested, with obvious gaps across scenarios. This suggests that the aggregate macroeconomic impact—and the subsequent policy response—is primarily driven by the volume of infected individuals rather than the specific magnitude of their individual productivity loss. Furthermore, this lack of sensitivity is reinforced by the way the health budget is built (recall Eqs. 23 and 24): because the fiscal allocation for treatment is structurally decoupled from short-run economic fluctuations, the health budget remains relatively stable compared to the more volatile variations in GDP. Consequently, the epidemiological state of the system acts as the dominant signal for policy calibration, as the budget's relative inertia prevents minor changes in productivity loss from significantly altering the capacity for intervention. This is then reflected in the likewise stable level of the share of infected and the cost-to-budget ratio. The only visible change occurs as $\sigma$ approaches 1, a region that we do not explore further to avoid computational issues in the cost-to-budget ratio.

## Appendix C: Global spatial interaction

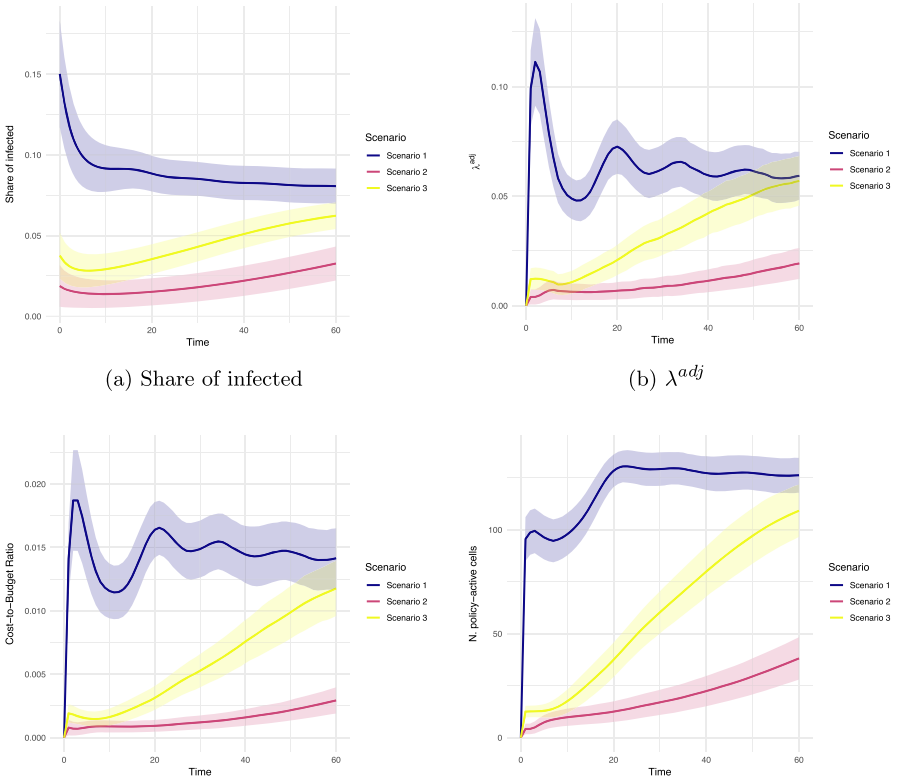


(a) Share of infected (b) $\lambda^{adj}$

(c) Cost-to-budget ratio (d) N. of policy-active cells

**Fig. 20** Average share of infected over population, $\lambda^{adj}$, cost-to-budget ratio across cells, and number of policy-active cells for 2000 Monte Carlo simulations, under homogeneous pricing and $\omega = 0.3$ with expectations over the average global $\lambda^{adj}$. *Shaded areas* represent $\pm$ one standard deviation

**Acknowledgements** We acknowledge the two anonymous reviewers and the Leading Guest Editors for their valuable comments on earlier drafts, allowing a substantial improvement in the quality of the article. We thank the organizers and participants of the 3rd International Workshop on Heterogeneity, Evolution, and Networks in Economics held at the University of Florence in September 2025 and of the 2nd Workshop in Health, Environment, Development and Growth Economics held at the University of Pisa in February 2026 for their insightful comments. We are also grateful to the participants of the DISEI department seminars in Florence.

**Author Contributions** All authors contributed to the study conception and model design. A.C. conducted the analytical study, while S.L. performed the computational analysis, including simulations and data processing. Both authors contributed to writing the manuscript. All authors read and approved the final version.

**Funding** Open access funding provided by Università degli Studi di Firenze within the CRUI-CARE Agreement. No funding was received for conducting this study.

**Data Availability** The authors will make the data and code used in this study available upon request or via a personal repository.

## Declarations

**Competing interests** The authors declare no competing interests.



## References

Acemoglu D, Linn J (2004) Market size in innovation: Theory and evidence from the pharmaceutical industry. Quart J Econ 119(3):1049–1090

Agar MH, Wilson D (2002) Drugmart: Heroin epidemics as complex adaptive systems. Complexity 7(5):44–52

Aghion P, Howitt P (1992) A model of growth through creative destruction. Econometrica 60:323–351

Ale S, Hunter E, Kelleher JD (2024) Agent based modelling of blood borne viruses: a scoping review. BMC Infect Dis 24(1):1411

Assefa Y, Hill PS, Williams OD (2018) Access to hepatitis c virus treatment: lessons from implementation of strategies for increasing access to antiretroviral treatment. Int J Infect Dis 70:65–68

Auty SG, Griffith KN, Shafer PR, Gee RE, Conti RM (2022) Improving access to high-value, high-cost medicines: the use of subscription models to treat hepatitis c using direct-acting antivirals in the united states. J Health Polit Policy Law 47(6):691–708

Axtell RL, Farmer JD (2025) Agent-based modeling in economics and finance: Past, present, and future. J Econ Literat 63(1):197–287

Blank A, Solomon S (2000) Power laws in cities population, financial markets and internet sites (scaling in systems with a variable number of components). Physica A 287(1–2):279–288

Clementi F, Gallegati M (2005) Power law tails in the italian personal income distribution. Physica A 350(2–4):427–438

Danzon PM, Towse A (2003) Differential pricing for pharmaceuticals: Reconciling access, r&d, and patents. Int J Health Care Finance Econ 3(3):183–206

Dubois P, Magnac T (2024) Optimal intertemporal curative drug expenses: The case of hepatitis c in france. J Health Econ 94:102861

Dubois P, De Mouzon O, Scott-Morton F, Seabright P (2015) Market size and pharmaceutical innovation. Rand J Econ 46(4):844–871

Eeckhout J (2004) Gibrat's law for (all) cities. American Economic Review 94(5):1429–1451

Gabaix X (1999) Zipf's law and the growth of cities. American Economic Review 89(2):129–132

Gleeson D, Townsend B, Tenni BF, Phillips T (2023) Global inequities in access to covid-19 health products and technologies: A political economy analysis. Health Place 83:103051

Gozzi N, Chinazzi M, Dean NE, Longini IM, Halloran ME, Perra N, Vespignani A (2023) Estimating the impact of covid-19 vaccine inequities: A modeling study in lower-income countries. Nat Commun 14:225

Hammond RA (2009) Complex systems modeling for obesity research. Prev Chronic Dis 6(3):97

Heard D, Bobashev GV, Morris RJ (2014) Reducing the complexity of an agent-based local heroin market model. PLoS ONE 9(7):102263

Institute of Medicine (2015). Assessing the use of agent-based models for tobacco regulation. The National Academies Press. https://doi.org/10.17226/19018
Kanavos P, Costa-Font J (2005) Pharmaceutical parallel trade in europe: stakeholder and competition effects. Economic policy 20(44):758–798
Kerr CC, Stuart RM, Mistry D, Abeysuriya RG, Rosenfeld K, Hart GR, Núñez RC, Cohen JA, Selvaraj P, Hagedorn B et al (2021) Covasim: an agent-based model of covid-19 dynamics and interventions. PLoS Comput Biol 17(7):1009149
Kremer M, Levin J, Snyder CM (2020) Advance market commitments: insights from theory and experience. In: AEA Papers and Proceedings, vol 110, pp 269–273. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203
Kustudic M, Niu B, Liu Q (2021) Agent-based analysis of contagion events according to sourcing locations. Sci Rep 11(1):16032
La Torre D, Liuzzi D, Marsiglio S (2022) Geographical heterogeneities and externalities in an epidemiological-macroeconomic framework. Journal of Public Economic Theory 24(5):1013–1037. https://doi.org/10.1111/jpet.12561
La Torre D, Liuzzi D, Marsiglio S (2024) Epidemic outbreaks and the optimal lockdown area: a spatial normative approach. Econ Theor 77(2):475–503. https://doi.org/10.1007/s00199-023-01501-4
Leopold C, Vogler S, Mantel-Teeuwisse AK, Joncheere K, Leufkens HG, Laing R (2012) Differences in external price referencing in europe–a descriptive overview. Health Policy 104(1):50–60
Moon S, Jambert E, Childs M, Von Schoen-Angerer T (2011) A win-win solution?: A critical analysis of tiered pricing to improve access to medicines in developing countries. Glob Health 7:39
Reed WJ, Jorgensen M (2004) The double pareto-lognormal distribution–a new parametric model for size distributions. Communications in Statistics-Theory and Methods 33(8):1733–1753
Sciarretta K, Røttingen J-A, Opalska A, Van Hengel AJ, Larsen J (2016) Economic incentives for antibacterial drug development: literature review and considerations from the transatlantic task force on antimicrobial resistance. Clin Infect Dis 63(11):1470–1474
Shiflet, A, Shiflet, G (2014). Introduction to computational science: modeling and simulation for the sciences, 2nd edn. Princeton University Press
Thomas LJ, Huang P, Yin F, Luo XI, Almquist ZW, Hipp JR, Butts CT (2020) Spatial heterogeneity can lead to substantial local variations in covid-19 timing and severity. Proc Natl Acad Sci 117(39):24180–24187
Vermeer W, Gurkan C, Hjorth A, Benbow N, Mustanski BM, Kern D, Brown CH, Wilensky U (2022) Agent-based model projections for reducing hiv infection among msm: Prevention and care pathways to end the hiv epidemic in chicago, illinois. PLoS ONE 17(10):0274288
Vogler S, Schneider P, Zimmermann N, Zuba M, Habimana K, Boehm K, Litzberg A (2021) Study on enhanced cross-country coordination in the area of pharmaceutical product pricing. Technical report, European Commission, CHAFEA, Luxembourg. Final Report. https://jasmin.goeg.at/id/eprint/218/
Wolfram C (2020) An agent-based model of covid-19. Complex Systems 29(1)
World Health Organization (1980) Declaration of global eradication of smallpox= proclamation de l'éradication de la variole dans le monde entier. Weekly Epidemiological Record= Relevé épidémiologique hebdomadaire 55(20):148–148. https://iris.who.int/handle/10665/223060
Yao L, Ying X, Malik G, Tsai C, Jesudian AB, Brown RS Jr, Congly SE (2025) A global comparison of hepatitis b & c drug pricing. Ann Hepatol 30(2):101928