



UNIVERSITÀ
DEGLI STUDI
FIRENZE

DISEI

DIPARTIMENTO DI SCIENZE
PER L'ECONOMIA E L'IMPRESA

WORKING PAPERS - ECONOMICS

An Agent-Based Approach to High-cost Drugs for Infectious Diseases

ANDREA CARAVAGGIO, SILVIA LEONI

WORKING PAPER N. 20/2025

*DISEI, Università degli Studi di Firenze
Via delle Pandette 9, 50127 Firenze (Italia) www.disei.unifi.it*

The findings, interpretations, and conclusions expressed in the working paper series are those of the authors alone. They do not represent the view of Dipartimento di Scienze per l'Economia e l'Impresa

An Agent-Based approach to high-cost drugs for infectious diseases

Andrea Caravaggio¹ and Silvia Leoni^{2,*}

¹Department of Economics and Statistics, University of Siena.

²Department of Economics and Management, University of Florence.

*Corresponding author, silvia.leoni@unifi.it.

Abstract

The COVID-19 pandemic has underscored the need for modeling tools that account for the spatial and institutional heterogeneity underlying real-world epidemic dynamics. We develop a spatially structured agent-based model (ABM) in which decentralized health authorities (HAs) allocate costly treatments under local budget constraints to manage the spread of an infectious disease. Individuals are distributed across a grid of locations, with contagion governed by discrete-time Susceptible-Infected-Recovered (SIR) dynamics and spatial spillovers through local interactions. At each time step, HAs choose treatment intensity endogenously based on local infection levels, available resources, and pricing conditions. We analyze how key factors—such as treatment efficacy, pricing schemes, and initial outbreak distribution—shape both local and aggregate outcomes. In addition to a benchmark case with homogeneous pricing, we explore a parsimonious pricing scheme where prices vary across cells. Analytical results identify the threshold conditions for disease eradication, while simulations show how decentralized decisions and spatial feedback can generate persistent inequalities in infection and treatment. Our findings highlight the importance of integrating spatial structure, economic constraints, and pricing design in epidemic policy modeling.

Keywords: Agent-based modeling; health policy; infectious disease control; SIR model; treatment pricing.

JEL Classification: C63 , H51 , I18

1 Introduction

The global experience with COVID-19 and other recent epidemics has underscored the urgent need for robust, flexible frameworks to design and evaluate public health interventions under uncertainty and constraints. In particular, the emergence of innovative but costly pharmaceutical treatments has posed challenging questions for health authorities tasked with containing the spread of infectious diseases while managing scarce budgetary resources.

Traditional epidemiological models, typically grounded in mean-field approaches such as deterministic SIR (Susceptible-Infected-Recovered) systems, have provided useful insights into disease propagation and basic reproduction thresholds. However, such models often fail to capture the spatial, behavioral, and institutional heterogeneity that shapes real-world epidemic dynamics and policy responses. As highlighted by Hunter et al. [1], standard equation-based models underestimate the role of local interactions and decentralized behavioral responses, leading to oversimplified policy prescriptions. In recent years, there has been increasing interest in agent-based modeling (ABM) approaches for studying epidemic control, particularly in contexts characterized by local coordination failures, heterogeneous agents, and spatially distributed interactions [2]. These models allow for a bottom-up representation of complex adaptive systems, enabling researchers to analyze non-linear contagion effects, policy diffusion, and dynamic adaptation to public health interventions. Delli Gatti et al. [3] emphasize the role of behavioral feedback and bounded rationality in epidemic dynamics, proposing ABMs as essential tools for stress-testing policy robustness in uncertain environments. Beyond behavioral dynamics, agent-based models are particularly valuable for their ability to capture spatial interactions, which are essential for understanding disease transmission across heterogeneous geographic areas or interconnected networks [4]. ABMs have been widely used to evaluate the impact of travel restrictions, localized interventions, and individual mobility patterns, highlighting how factors like super-spreading locations and socioeconomic mobility disparities shape epidemic outcomes [4]. Moreover, ABMs can explicitly represent heterogeneous populations, where agents vary in e.g. infectivity, or responsiveness to interventions—features often oversimplified in compartmental models [5]. Modeling individual-level decision-making and interactions also enables the study of emergent collective behaviors, such as compliance with public health guidelines or the spread of misinformation, which can significantly influence epidemic dynamics [6].

With increasing computational power, ABMs are now applied at larger scales, from urban outbreaks to global pandemic simulations, offering insights into the effectiveness of non-pharmaceutical interventions (NPIs) and vaccination strategies [7, 8]. Their flexibility in integrating diverse data sources—including demographic, behavioral, and mobility data—makes ABMs a powerful tool for public health planning and real-time policy evaluation.

Employing a different methodological approach, Dubois and Magnac [9] introduce a dynamic optimization framework for epidemic management, focusing on a centralized planner allocating costly treatments under a budget constraint. Their contribution is notable in highlighting the trade-off between early aggressive intervention and intertemporal budget smoothing. However, their model does not account for spatial

contagion nor decentralized optimization. In contrast, our framework embeds these key elements, showing how localized decision-making and neighborhood spillovers fundamentally reshape the effectiveness and cost-efficiency of epidemic control.

Further research has explored the importance of network structure and contact patterns in shaping epidemic diffusion. For instance, Ajelli and Merler [10] and Ferguson et al. [11] examine how granular contact networks alter epidemic forecasts and the optimal allocation of limited treatment or vaccination resources. More recent agent-based modeling studies have shown that spatial and institutional heterogeneity can produce unintended consequences and tipping points even under well-targeted interventions. For instance, Thomas et al. [12] demonstrate significant local variation in epidemic timing and severity across metropolitan areas, despite consistent average trends. Likewise, Kustudic et al. [13] reveal that contagion dynamics differ markedly depending on the source location and network structure, highlighting how small structural variations can amplify epidemic outcomes beyond average-level predictions.

Our work contributes to this literature by combining a dynamic agent-based SIR model with decentralized treatment decisions, spatial interactions via a Moore neighborhood, and heterogeneous pricing schemes. In doing so, we follow recent calls to move beyond static or representative-agent frameworks and toward spatially explicit models that endogenize both epidemiological dynamics and economic constraints [14]. 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 [15] and Moon et al. [16] have analysed 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.

Empirical research on the political economy of global health has confirmed that inequities in access to treatments and vaccines—including pricing-based disparities—have real consequences for public health outcomes. A comprehensive study by Gleeson et al. [17] documents how structural inequities in access to COVID-19 health products undermine disease containment and exacerbate socioeconomic vulnerability. Similarly, modeling studies like Gozzi et al. [18] show that delayed or uneven distribution of vaccines significantly lowers herd immunity in lower-income countries, amplifying both health and economic costs.

Our analysis builds on this literature by endogenizing a parsimonious pricing mechanism, where price varies across locations according to the local benefit of recovery and efficacy, but without respect to actual fiscal capacity. We show how such pricing can distort treatment incentives, reducing overall treatment intensity and increasing spatial inequality, thus undermining both epidemiological effectiveness and equity.

In this paper, we develop an agent-based and spatially structured model of epidemic management in which decentralized health authorities (HAs) operate under resource constraints. The population is distributed across a toroidal grid of locations (cells), each with its own HA facing a budget limitation. At every discrete time step, each HA determines the fraction of locally infected individuals to be treated using an

innovative but costly pharmaceutical intervention. Effective treatment with the innovative medicine cures and immunizes patients. 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, allowing infections to spread across adjacent cells. We model a generic infectious disease characterized by relatively slow contagion dynamics, akin to Hepatitis C, rather than fast-spreading pathogens like COVID-19.

Our modeling framework introduces several key innovations with respect to the existing literature. First, it incorporates a decentralized optimization mechanism: treatment intensity is chosen endogenously by each HA based on the current level of infection and subject to a local budget constraint. Second, the model explicitly includes spatial externalities, whereby the infection in one cell may influence epidemiological outcomes in neighboring regions, creating feedback loops that shape aggregate dynamics. Third, we consider heterogeneous pricing strategies for the innovative treatment. In particular, we contrast a benchmark case of uniform pricing with a parsimonious scheme in which the price charged in each cell depends on local recovery benefits and treatment efficacy—without knowledge of the HA’s budget capacity.

Using a combination of analytical and computational methods, we derive several key results. Analytically, we identify the conditions under which the disease-free equilibrium is locally stable in both isolated and interacting settings, characterizing a critical threshold for treatment intensity that guarantees eradication. Computationally, we show that when treatment is effective and prices are fixed, decentralized optimization leads to a rapid reduction in infections across the grid, although spatial spillovers generate heterogeneous and persistent infection patterns. In contrast, parsimonious pricing strategies—intended to reflect local benefit-cost trade-offs—can lead to counterproductive outcomes: higher prices in more vulnerable or treatment-effective regions may suppress treatment effort, thereby amplifying infection rates and increasing inequality in health outcomes across regions.

The simulations also reveal that limited treatment efficacy slows down the system’s convergence to the disease-free state and can generate widespread persistence of the infection, particularly when initial outbreaks occur in highly connected or poorly treated areas. The joint effect of spatial transmission, budget constraints, and pricing heterogeneity produces complex epidemic trajectories, in which early under-treatment in certain cells acts as a self-reinforcing mechanism for disease diffusion. Even when average infection rates fall, variance across locations remains high unless pricing is coordinated and treatment is broadly accessible.

Our findings contribute to current debates on the optimal design of decentralized health responses, the economic trade-offs involved in pharmaceutical pricing, and the broader consequences of spatially heterogeneous epidemic dynamics. In addition, the combination of agent-based modeling with analytical optimization techniques offers a novel methodological contribution to the growing literature relying on this hybrid approach. Such combination allows to enrich behavioral rules or local decision mechanisms. In these hybrid approaches, closed-form solutions or optimal strategies—derived from analytical models—are embedded into agents’ behavior within an ABM framework. This allows for the exploration of dynamic macro outcomes arising from locally

optimal but decentralized decisions. For instance, Gatti et al. [19] integrate optimal firm investment rules into an agent-based macro model to study financial instability, while in [20] individuals populating an overlapping generation economy locally optimize the allocation of their labor income between private goods consumption or the environment. These methods enable researchers to explore how theoretically grounded optimal solutions perform in complex, heterogeneous environments with interactions, feedback loops, and spatial spillovers. Our work contributes to this literature by incorporating a locally optimal treatment rule—derived from a static optimization problem subject to budget constraints—into the decentralized dynamics of a spatial epidemic ABM.

Our modeling framework can be extended in several directions, including behavioral learning by HAs, strategic coordination among jurisdictions, or supplier-side optimization of pricing under incomplete information. In general, our work emphasizes the importance of integrating spatial structure, institutional heterogeneity, and decentralized constraints into epidemic policy modeling.

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 calibration strategy adopted. Section 4 presents the simulation results under uniform pricing, exploring how treatment efficacy and decentralized responses shape epidemic outcomes. Section 5 provides a full sensitivity analysis on the main parameters for all the relevant scenarios proposed. Section 6 investigates the implications of introducing parsimonious, location-specific pricing, and compares the resulting dynamics to the baseline scenario. Finally, Section 7 concludes.

2 The Model

This paper investigates the dynamics and control of an 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 capture these mechanisms, we develop an agent-based model embedded on a two-dimensional toroidal lattice. Each cell (or patch) in the lattice represents a local subpopulation and follows a compartmental SIR scheme in discrete time. 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 .¹

$$\begin{cases} S_{t+1} = S_t - \Delta I_t \\ I_{t+1} = I_t + \Delta I_t - \rho \lambda_t I_t \\ R_{t+1} = R_t + \Delta R_t \end{cases} \quad (1)$$

Here, $\Delta I_t = \beta(I_t) c S_t \frac{I_t}{N}$ denotes the new infections at time t , while $\Delta R_t = \rho \lambda_t I_t$ represents the individuals who recover due to treatment. The infection rate $\beta(I_t)$ is assumed to be decreasing in the health authority's treatment effort $h_t = p \lambda_t I_t$, in which $p > 0$ represents the constant unitary price of treatment, and it reflects the idea that stronger treatment campaigns reduce effective contagion. Specifically, we consider

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

where β_0 is the baseline transmission rate, $\rho > 0$ measures treatment efficacy, and $\bar{s} \geq 0$ captures other exogenous mitigation measures (e.g., social distancing). The parameter $c > 0$ is the average number of contacts per person per period.

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 $(S^*, I^*, R^*) = (N, 0, 0)$. To evaluate its local stability, we linearize the system around this point and derive the corresponding Jacobian matrix:

$$J(S, I) = \begin{pmatrix} 1 - \beta(h_t) c \cdot \frac{I_t}{N} & -\frac{\beta_0 c (1 + \bar{s})}{(1 + \rho p \lambda I_t + \bar{s})^2 N} \\ \beta(h_t) c \cdot \frac{I_t}{N} & \frac{(1 - \rho \lambda)(1 + \bar{s})^2 N - \beta_0 c (1 + \bar{s}) N}{(1 + \rho p \lambda I_t + \bar{s})^2 N} \end{pmatrix} \quad (2)$$

¹Susceptible, infected, and recovered states are assumed to be nonnegative at each time t .

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

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

Defining the basic reproduction number as

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

we derive the eigenvalues of J_{DFE} , that is the pair $(\lambda_1, \lambda_2) = (1, 1 + \rho\lambda(\mathcal{R}_0 - 1))$. Therefore, the DFE is locally stable if and only if $\mathcal{R}_0 < 1$ and this condition provides a critical threshold for disease eradication. Indeed, to suppress infection, the treatment rate must satisfy

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

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 which 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:

$$\begin{cases} S_{t+1}^v = S_t^v - \Delta I_t^v \\ I_{t+1}^v = I_t^v + \Delta I_t^v - \rho\lambda I_t^v \\ R_{t+1}^v = R_t^v + \Delta R_t^v \end{cases} \quad (6)$$

where

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

$$\Delta R_t^v = \rho\lambda I_t^v. \quad (8)$$

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

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

As in the case without local interactions, the system still admits a global Disease-Free Equilibrium (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

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

which simplifies to

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

The above expression reveals that the local infection dynamics are influenced not only by the infected individuals in v , but also by the infection prevalence in neighboring cells, scaled by population ratios. By defining a cell-specific basic reproduction number as

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

The DFE is locally asymptotically stable in cell v if and only if $\mathcal{R}_0^v < 1$. Therefore, the global DFE is stable across the lattice if

$$\max_{v \in \mathcal{V}} \mathcal{R}_0^v < 1 \quad \Rightarrow \quad \lambda^v > \bar{\lambda}^v \frac{\beta_0 c}{\rho (1 + \bar{s})} \cdot \max_{v \in \mathcal{V}} \left(1 + \sum_{\ell \in \mathcal{N}_v} \frac{N^v}{N^\ell} \right). \quad (11)$$

This condition generalizes the non-spatial threshold by incorporating the spatial heterogeneity in population distribution. Cells surrounded by lower-population neighbors are more prone to instability, requiring stronger treatment effort to prevent epidemic persistence.

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 fixed per-period budget $b_t^v = b$. Since treatment reduces contagion and promotes recovery, but also entails direct and indirect costs, the HA chooses λ_t^v to maximize a payoff function that captures both the epidemiological and economic impacts of its decision. Importantly, HAs are myopic

and do not form expectations about the state or actions of neighboring cells. 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:

$$\max_{0 \leq \lambda_t^v \leq 1} \Pi_t^{HA,v} \quad (12)$$

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

The payoff function is:

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

where the positive parameters are described as follows:

- γ^v captures the monetary benefit of each recovery due to treatment;
- α_I weighs the (quadratic) cost from the prevalence of infection;
- α_T^v represents the cost of implementing the treatment effort;
- δ models the decreasing marginal returns of treatment as the number of infections increases.

Solving the first-order condition for an interior solution yields

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

from which we directly obtain

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

Thus, the complete policy rule for each local HA is:

$$\lambda_t^{*,v} = \min \left\{ \max \{0, \lambda_{int}^{*,v}\}, \frac{b^v}{pI_t^v}, 1 \right\} \quad (16)$$

This rule ensures that the treatment decision remains feasible (bounded between 0 and 1) and respects the budget constraint. It also implies that treatment intensity varies dynamically across space and time as a function of local infection levels. Since each HA operates independently and without coordination across the lattice, treatment policies may vary substantially across locations, depending on local epidemic conditions and resource constraints. This spatial heterogeneity can affect the aggregate disease dynamics and the likelihood of cross-cell contagion.

3 The Agent-Based Framework

3.1 Initialization

The simulation world employs a toroidal grid topology. Such topology eliminates edge effects by ensuring all patches have identical neighborhood connectivity. While this abstracts real-world geography, this is computationally preferable for studying global pandemic dynamics where 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 a world country, with a total of 196 cells. Each cell carries a population drawn from a transformed lognormal distribution that takes values $[100, 10^6]$. At the setup, a cell (country) population can either be entirely susceptible, or it can have a share of infected, according to three different scenarios, displayed in Figure 1. The selected scenarios are designed to represent a strategically relevant initial spread configuration and to highlight specific aspects of the disease’s evolution.

In the first scenario, each cell is initialized with an infected share i^v drawn independently from a continuous uniform distribution over the interval $[0, 1)$:²

$$i^v \sim \mathcal{U}(0, 1) \quad (17)$$

where $i^v = \frac{I^v}{N^v}$ represents the proportion of infected individuals.

This approach simulates a situation where the disease is already widespread or endemic at various levels globally, without a single isolated point of origin. The objective is to study propagation in a context of diffuse pre-existence, which is useful for understanding diseases with long incubation periods or asymptomatic transmissions leading to undetected spread.

In the second scenario, initial conditions simulate a single localized world outbreak where the seed country represents a high-infection core ($> 50\%$ infected) and its adjacent territories are partially infected ($\leq 50\%$ infected). In detail:

- Seed country (randomly selected patch i_0):

$$i_0^k \sim \mathcal{U}(0.5, 0.8)$$

- Adjacent neighbors of cell k (Moore neighborhood):

$$i_0^\ell \sim \mathcal{U}(0, 0.5) \quad \forall \ell \in \mathcal{N}_k$$

- All other countries:

$$i_0^v = 0 \quad \forall v \notin \mathcal{N}_k$$

This emulates a classic pandemic emergency with a well-defined epicenter, allowing for the analysis of spatial spread dynamics and the impact of geographical connectivity in the early phase of an epidemic.

²The choice of this range reflects the behavior of the command used in NetLogo (the programming environment adopted for simulations), which generates a pseudo-random floating-point number greater than or equal to 0 and strictly less than the specified upper bound.




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

The third scenario extends Scenario 2 by establishing $n = 10$ spatially distinct outbreaks. Each epicenter and its adjacent territories are initialized identically to the single-outbreak case. This configuration simulates the near-simultaneous emergence of the disease in multiple regions, providing insight into how the interaction and coalescence of multiple foci influence the overall global pandemic trajectory.

Together, these three scenarios offer a diverse range of starting conditions, from widespread pre-endemic to isolated and multifocal emergence, enabling us to thoroughly explore the impact of spatial heterogeneity and the initial distribution of infected individuals on the dynamics of a large-scale epidemic.

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

each cell v is initialized with γ^v drawn from a transformed pareto-lognormal distribution that takes values $[10, 500]$ so to represent a realistic distribution of wages.³ Since we can interpret the parameter γ^v as the average annual wage within a country, we derive a country's GDP, B^v , as:

$$B^v = \gamma^v N^v \quad (18)$$

We exploit the GDP to obtain the national health budget in the following way:

$$b^v = r N^v \quad (19)$$

where $r \sim \mathcal{U}(0, 1)$, such that cells are endowed with a heterogeneous random share of their GDP as budget devoted to health expenses, in other words, the HA budget. Besides these heterogeneous constant parameters, the model includes constant parameters which are homogeneous across cells, ρ , c , β_0 , δ , and are given the values reported in Table 1.

Price, p , is assumed homogeneous and time-invariant following the idea that a monopolistic drug manufacturer sets a fixed global price.

Variable	Values
ρ	0.7
c	6
β_0	0.1
δ	4
$\bar{s}$	10
N^v	$\sim \text{Lognormal}(\ln 5000, (\ln 1.8)^2)$ subject to $N^v \in [10, 10^6]$
γ^v	$\sim \text{Pareto-Lognormal}(100, 1.2, 1.8)$ subject to $\gamma^v \in [10, 500]$
p	18

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

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 timescales, with slower transmission rates and prolonged durations of infectiousness. This epidemiological profile allows us to abstract away from rapid day-to-day fluctuations and instead focus on longer-term spatial and policy dynamics. 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, each representing a year. This time span is chosen to provide a reasonable and comprehensive period for observing

³Note that a Pareto-Lognormal distribution is governed by a set of three parameters, μ , the location parameter of the lognormal distribution, σ , the scale parameter of the lognormal distribution, and α , the shape parameter ruling the tails. The parameters chosen for the wage distribution are listed in Table 1.

the long-term dynamics of the type of disease being modeled.

Every year, each HA^v must decide over the optimal share of the infected population to treat with the innovative drug, $\lambda^{*,v}$, on the basis of its budget. As seen in 12, the optimization takes into account a monetary cost, α_T^v , indirectly associated with the treatment, representing for instance distribution and personnel cost. We assume this cost to be highly dependent on exogenous macroeconomic shocks and inflation dynamics, therefore α_T^v is updated at each time step according to the following rule:

$$\alpha_T^v \sim \mathcal{U}(50, 100) \quad (20)$$

Once cells set their α_T^v , the infection rate $\beta^v(I_t^v)$ evolves according to equation (9), and the number of new infections, following equation (7).

Finally, cells update their population compartments accordingly. The algorithm followed at each time step is displayed in Figure 2.

We run 2000 Montecarlo simulation for each setup scenario discussed.




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

4 Results

Figure 3 shows the final state ($t = 60$) for the three scenarios considered for one simulation run. The comparison with the initial state $t = 0$ displayed in Figure 1 shows that initial conditions are crucial on the disease spread and the final outcome over the time span considered.⁴

Figure 4 displays the distribution of population compartments (infected, recovered, and susceptible) across cells at $t = 60$ time steps for the three different initial scenarios.

Scenarios 1 and 3 clearly show a bimodal distribution of recovered which follows the double bump observed for the infected compartment.

In the first scenario, the disease has spread and a substantial number of individuals have recovered, but a large portion of the population remains susceptible. The graph also shows that most cells either have a very low number of currently infected individuals or are entirely free of active infections. The broader tail extending towards higher $\log(1+x)$ values indicates that a small proportion of cells might still harbor a larger number of infected individuals, but at much lower densities. This pattern is characteristic of a disease that has spread widely and then largely receded, with only residual or very low-level ongoing transmission in many areas. The presence of a significant recovered peak confirms that the disease has progressed through a considerable portion of the global population.

Similarly to the first scenario, the third one shows a very high density of cells with a small number of infected or recovered individuals (approaching zero on the linear scale), although the second peak of the recovered distribution is far less prominent indicating that the recovered populations, while present, are less concentrated at the very highest compartment sizes across the grid compared to the first scenario where the disease diffused more widely from an initial pervasive state.

The bimodal pattern is less visible in the second scenario where early treatment adoption allows to significantly contain the infection process generated by a single outbreak. The infected curve displays an extremely sharp and high peak at $\log(1+x) \approx 0$. This indicates that the vast majority of cells have effectively zero or extremely few currently infected individuals after 60 years. This strong concentration at zero is indicative of a successful eradication or a highly localized initial spread that has since faded out in most of the grid.

Similarly, the recovered curve also shows a very high peak at $\log(1+x) \approx 0$, reinforcing the idea that most cells either had no significant infection to begin with, or the infected individuals recovered and their numbers subsequently fell to very low levels while a significant portion of cells remained largely untouched by the infection.

Figure 5 shows the dynamics of the global population compartments for a single simulation for the baseline model (left panels) in comparison with the model simulated with $\lambda_t^v = 0$ (right panels), i.e. assuming that HAs do not treat any share of their infected population. The comparison shows the success of the drug adoption in slowing down the infection process for all three scenarios. Note that although the share of infected population in the third scenario of the baseline model appears increasing for

⁴Comparison is ensured by setting the same seed, i.e. the numerical value used to initialize a pseudo-random number generator (PRNG).




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

the 60 years analyzed, in the very long run it stabilizes concurrently to an increase of the recovered population (see Figure 20 in the Appendix).

We can draw similar interpretations for Figure 6 which shows the average share of infected across cells for the three scenarios. Instant treatment adoption results in a quick and remarkable drop in the number of infected in the first scenario. The third scenario follows a similar trend while the second one reflects the instant infection containment with average share of infected across cells approaching zero (Figure 6a). The third scenario can be interpreted as a middle ground between the more extreme first and second scenarios. Notably, the results from the first and third scenarios show greater alignment, whereas the second scenario diverges more substantially. This is particularly evident in the graph illustrating disease spread in the absence of a cure




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.

(i.e. $\lambda_t^v = 0$ for all v and t), where convergence to full infection occurs more slowly in the second scenario (Figure 6b).

In Figure 7 we report the average λ^* across cells and average cost-budget ratio across cells, defined for every cell v at each time step t as

$$\frac{p\lambda_t^{v,*}I_t^v}{b^v} = \frac{h_t^v}{b^v}. \quad (21)$$

By analyzing Figures 6 and 7 jointly, we can draw three key conclusions regarding the economic-epidemiological dynamics arising from decentralized treatment decisions under budget constraints. These insights results from the interaction between the initial spatial distribution of infections and induced treatment response, with important long-term fiscal implications.




Fig. 5: Comparison of infection dynamics across the three different scenarios. Seed = 1000.

First, Figure 6a highlights the significant containment effect of endogenous treatment policies. When HAs adaptively determine the share of infected individuals to treat—subject to local budget constraints—the epidemic curve is significantly flat-tened compared to the case without treatment (Figure 6b). In the no-treatment simulations, the average infection levels remain markedly higher and show persistent




Fig. 6: Average share of infected over population across all countries, across 2000 Montecarlo simulations




Fig. 7: Average λ^* (panel a) and Cost-to-Budget ratio (panel b) over population across all countries, across 2000 Montecarlo simulations

growth over time, indicating the inability of natural recovery alone to suppress the epidemic. This contrast clearly illustrates the effectiveness of decentralized but responsive intervention in reducing infection prevalence and transmission speed.

Second, Figure 7b reveals that in all simulated scenarios, a strong and immediate fiscal reaction occurs: the cost-to-budget ratio peaks early on, reflecting the urgency with which HAs respond to initial outbreaks. This front-loaded treatment effort is instrumental in curbing the epidemic's spread and is especially pronounced in more severe or widespread initial configurations. Importantly, the panel also shows that despite the heterogeneity of initial conditions, and resulting differences of government decision over the share of infected to treat (i.e. λ^* , see Figure 7a) across the three scenarios, the cost-to-budget ratio converges to roughly the same medium- to long-run level across all scenarios. This outcome demonstrates that once the disease is effectively contained, the required treatment effort stabilizes at a low and comparable level—underscoring the self-correcting nature of optimal local interventions.

Third, the comparison across scenarios reinforces the nonlinearity of the epidemic-management process. When the outbreak is spatially concentrated (as in Scenario 2), early aggressive intervention proves particularly efficient: it quickly eliminates contagion in most cells and drives both infection and treatment rates close to zero. In contrast, scenarios with widespread or multifocal outbreaks (Scenarios 1 and 3) exhibit more persistent infection dynamics and higher ongoing treatment levels. Nonetheless, the convergence of fiscal burden over time (Figure 7b) reveals that effective local containment ultimately harmonizes the economic cost structure, even across initially disparate contexts.

These results emphasize the strategic importance of early, targeted, and adaptive responses. While local HAs act independently and without global coordination, their endogenous treatment policies—driven by local infection levels and bounded by fiscal constraints—generate emergent dynamics that can suppress the epidemic effectively. In particular, the uniformity of long-run economic outcomes across diverse scenarios points to the robustness of decentralized responses, provided that initial outbreaks are met with sufficient intensity. From a policy perspective, this suggests that early allocation of treatment capacity and decentralized autonomy can jointly ensure both epidemiological and fiscal sustainability.

5 Sensitivity Analysis

To assess the robustness of our findings, we conducted sensitivity analyses on key model parameters across all three scenarios. Parameters were selected based on their relevance to the dynamics of disease transmission and intervention effectiveness. Specifically, we test the model sensitivity to variation of the marginal return of treatment, δ ; the starting infection rate, β_0 ; the efficacy of treatment, ρ ; and price p . For each scenario, we systematically varied these parameters within plausible ranges to examine the stability of the observed patterns for three main variables: the average λ^* , the average share of infected, and the average cost-to-budget ratio across cells. To ensure comparability, each simulation was run using the same random seed, so that all other conditions remained constant across parameter changes.⁵

5.1 Scenario 1

As known from equation (13), the marginal return of treatment is negatively related to the number of infections. We observe this negative relation in Figure 9a. It follows that as the marginal return to treatment is decreasing in the level of infected population, an additional treatment unit becomes cheaper and therefore we observe an increasing trend of the share of treated population, λ^* (Figure 8a), from which it follows an increasing cost-to-budget ratio (Figure 10a). However, besides the jumps (or drop in the case of infected) observed with respect to $\delta = 0$, changes observed in the selected variables remain extremely small.

Variations in the parameter β_0 produce unsurprising changes in the share of infected which increases in the level of the starting infection rate (Figure 9b), from which

⁵The seed is set to 1000 for all simulation runs in the sensitivity analysis. This ensures compatibility with other Figures in the paper.

follows an increase in the share of population to treat (Figure 8b), and therefore a growing cost-to-budget ratio (Figure 10b).

In Figures 8, 9, and 10 we test values of ρ between $[0.5, 1]$. Increasing the efficacy of treatment leads to smaller shares of infected population (Figure 9c), and therefore smaller λ^* (Figure 8c) and cost over national budget (Figure 10c), with these relations showing a non-linear pattern.

The chosen range of values tested for ρ comes from the consideration that for values $\rho \rightarrow 0$ the treatment is close to complete inefficacy and thus it would not make a reasonable investment for HAs. We highlight this aspect in Figure 11 for Scenario 1 (Scenarios 2 and 3 are shown in Figures 21 and 22 in the Appendix).

By conducting a comparative statics exercise on the relationship between the optimal treatment share λ_v^* and the level of treatment efficacy ρ , one observes that higher values of ρ lead to a greater fraction of infected individuals being treated. However, the figures go beyond this by showing the spatial and epidemiological consequences of low, intermediate, or high levels of efficacy, as reflected in the system's final state at time $t = 60$.

In particular, when comparing these results to the initial distribution of infection in Scenario 1 (as shown in Figure 1), it is evident, especially for $\rho = 0.3$ (panel (a)) that even though the disease is eradicated in cells where initial infection levels were extremely low, the combined effect of disease spread and insufficient treatment efficacy results in a final state where infection remains widely distributed—visibly indicated by a significant number of yellow-marked cells, with respect to the green cells observable in Figure 1.

This undesirable outcome diminishes as ρ increases: higher treatment efficacy corresponds to more effective containment and faster suppression of the disease. In the extreme case where $\rho = 1$, representing a fully effective treatment, the epidemic is essentially eradicated across the entire spatial domain within just 60 time steps.

Finally, we test the model sensitivity to variations of treatment price, while keeping it homogeneous across cells. The relation to all three variables analyzed appears to be non-linear, with a decreasing overall trend of the average λ^* across cells for increasing values of price: as treatment is assumed more expensive, the HA treats a smaller share of population (Figure 8d). As mentioned, the relationship is non-linear with a clear concave shape for values of price up to about 30. While this might appear counterintuitive, we remind that in addition to contributing to determining the optimal number of infected to treat, λ^* , at local/national level, price enters equation 9 thus contributing to the infection spread through interaction among cells. Overall, the share of infected in the model increases with increasing price as we expect smaller λ^* (Figure 9d). The cost-to-budget ratio clearly shows concavities reflecting the ambiguous effect of price as it participates in a double endogenous process while reasonably increasing overall as the treatment is assumed more expensive (Figure 10d).







Fig. 8: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average λ^* across simulations.




(a) δ



(b) β_0




(c) ρ




(d) price


Fig. 9: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average share of infected across simulations.




(a) δ



(b) β_0




(c) ρ




(d) price


Fig. 10: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average cost over budget across simulations.




(a) $\rho = 0.3$



(b) $\rho = 0.5$



(c) $\rho = 0.7$



(d) $\rho = 1$

Fig. 11: Final spatial distribution of infection at $t = 60$ in Scenario 1, under different treatment efficacies (Seed 1000).


5.2 Scenario 2

The second scenario is somewhat a more extreme case with respect to the first one and this results in less linear reactions of the model to changes in its parameters, particularly λ^* and price, for which clear and multiple concavities are observed in Figures 12, 13, and 14.


For values of δ in the range $[3, 5]$, the decreasing trend of λ^* might in fact appear non-trivial. Our interpretation is based on two main points. First, one should remember that mean values of λ^* across cells include a great majority of cases in which $\lambda^* = 0$, limiting the observed pattern of Figure 12a to a very localized decision-making process. Second, as shown in (13), δ interacts with the cost α_T^v which changes according to exogenous macroeconomic shocks (see Equation 20). Therefore, an increase in δ , which can be interpreted as a stronger convexity in the treatment cost function, reduces the marginal gain from curing additional infected individuals, especially in areas where the infection has already spread. From an economic perspective, this mimics real-world situations where rising marginal treatment costs, due to congestion or reduced treatment efficiency under pressure, discourage full utilization of available resources, particularly when budget constraints are tight.

Reactions of the model to different values of β_0 and ρ are as expected and similar to what observed for Scenario 1. Despite the limited infection spread given from the single-outbreak setting of Scenario 2, Figure 15 shows that different infection rates at the start result in drastically different final states of the disease spread. This reinforces the idea that even isolated initial shocks can lead to large-scale epidemics if local interactions are not effectively controlled. In economic terms, this underscores the importance of early containment strategies and how differences in intrinsic transmissibility (β_0) can enlarge the aggregate burden of disease even under seemingly favorable starting conditions.


A major difference with respect to Scenario 1 is the overall increasing trend of the share of treated individuals λ^* with regards to increasing levels of price. We interpret the non-linear pattern as the joint effect of the contribution of price to both λ^* and the infectivity rate β through λ^* , as shown in (9). More specifically, while higher prices naturally constrain the HA's budget and may reduce the feasible intensity of treatment, they also indirectly increase β via a decrease in treatment coverage, reinforcing infection spread. The feedback mechanism thus generates a non-monotonic relation: in some intermediate price ranges, rising prices induce HAs to react by allocating more aggressively within constrained environments. This mechanism illustrates how pricing policies, even when exogenous, can interact with decentralized optimization in non-obvious ways, amplifying spatial inequalities and distorting treatment strategies.




(a) δ



(b) β_0




(c) ρ




(d) price


Fig. 12: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average λ^* across simulations.




(a) δ



(b) β_0




(c) ρ




(d) price


Fig. 13: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average share of infected across simulations.




(a) δ



(b) β_0




(c) ρ




(d) price


Fig. 14: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average cost over budget across simulations.




(a) $\beta_0 = 0.1$



(b) $\beta_0 = 0.3$



(c) $\beta_0 = 0.5$



(d) $\beta_0 = 0.7$

Fig. 15: Final spatial distribution of infection at $t = 60$ in Scenario 2, under different baseline transmission rates (Seed 1000).


5.3 Scenario 3

As mentioned in previous sections, the third scenario represents a sort of middle way between Scenarios 1 and 2. Figures 16, 17, 18 show similar results as compared to Scenario 1, however levels of infected are smaller at all values tested for the four parameters and consequently lead to smaller levels of the share of treated population and corresponding budget devoted to treat the disease.


From an economic point of view, this outcome can be understood by considering the infection configuration of Scenario 3: the disease is initially seeded in a moderate number of cells, allowing for some spatial diffusion, but not as explosively as in Scenario 1. This intermediate setting enables health authorities to react in a more targeted and cost-efficient manner. With fewer cases per cell and weak spatial externalities, the perceived marginal benefit of treating infected individuals declines, particularly when compared to scenarios where contagion risks are more extreme and require urgent intervention to prevent cascades.

Moreover, in a context where the infection is neither minimal nor huge, local health authorities may face a relatively balanced optimization problem: the infection is significant enough to justify treatment, but not so severe as to require intensive or widespread effort. This translates into more cautious use of the available budget, reflected in lower λ^* and treatment costs. This behavior aligns with rational resource allocation under uncertainty and diminishing returns: as the infection remains under partial control, authorities avoid over-investment and instead preserve resources for future spikes or shifts in epidemiological conditions.


In addition, the smoother sensitivity profiles observed in Scenario 3, compared to the more nonlinear and discontinuous responses in Scenario 2, suggest that intermediate disease prevalence reduces the occurrence of threshold dynamics in policies. When infection levels are extreme, as in Scenario 2, small changes in parameters like price or treatment efficacy can produce large behavioral discontinuities. In contrast, the moderate setting of Scenario 3 produces more stable marginal incentives, allowing the system to respond gradually to parametric variation.




(a) δ



(b) β_0




(c) ρ




(d) price


Fig. 16: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average λ^* across simulations.




(a) δ



(b) β_0



(c) ρ



(d) price

Fig. 17: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average share of infected across simulations.




Fig. 18: Sensitivity analysis on δ , β_0 , ρ , and p . Mean and sd of average cost over budget across simulations.

6 The case of heterogeneous parsimonious prices

In contrast to the preceding sections, where the unit price of the innovative treatment was assumed to be fixed and exogenous across all countries, we now explore a scenario in which the treatment provider (the manufacturer) adopts a form of *parsimonious pricing*, i.e., a strategy of price differentiation across regions. In this framework, the supplier sets country-specific prices for the treatment based on observable economic and medical characteristics of each region, but without direct knowledge of the local health authority’s budget b^v .

More precisely, we assume that the treatment provider has access to information about (i) the *monetary benefit* of recovery γ^v , and (ii) the *effectiveness* of the treatment, measured by the local value of the parameter ρ . These two variables are used to determine the perceived value of the treatment in each country, and prices are set accordingly. This approach reflects a form of profit-oriented price discrimination, whereby the supplier seeks to extract more surplus in regions where the value of treatment is higher—either because the recovery yields greater economic returns or because the treatment is more effective. Specifically, we assume that for every cell v , the price p^v is drawn from a uniform distribution with minimum value 10 and maximum value $\gamma^v \rho$:

$$p^v \sim \mathcal{U}(10, \gamma^v \rho). \quad (22)$$

Importantly, because the supplier is assumed to lack precise knowledge of local health budgets, pricing is not tailored to maximize accessibility or epidemiological impact, but rather to align with perceived willingness-to-pay inferred from local benefit and efficacy levels. This distinction introduces a novel layer of complexity: local health authorities must react to prices that are endogenously determined by factors they (at least partially) do not control and that may not reflect their fiscal capacity.


In terms of results, Figure 19 shows that the adoption of heterogeneous pricing across countries—unrelated to the actual budget available to local Health Authorities—leads to a clear deterioration in epidemiological outcomes. In all simulated scenarios, the introduction of price dispersion results in higher infection levels (panel (a)) and a reduced capacity to eradicate the disease (panel (b)) compared to the benchmark case with uniform pricing (see Figures 6a and 7).

Moreover, Figure 19a reveals a striking increase in the spatial variability of infection prevalence, which not only persists but amplifies over time. This heterogeneity in local epidemic trajectories contributes to the broader diffusion of the disease, creating a self-reinforcing mechanism that sustains elevated infection rates across many cells until the final simulation period at $t = 60$.


This additional layer of heterogeneity introduced by the pricing scheme also affects other key variables. As shown in Panels (b) and (c), the same volatility is observed in the trajectories of treatment intensity λ and the cost-to-budget ratio. While the latter remains broadly comparable to the uniform price case in aggregate terms, the treatment share λ is consistently lower across most cells.

This outcome is largely driven by the fact that, although the assigned price may reflect the local value of recovery and efficacy, it does not necessarily ensure affordability. In many cases, the assigned price is such that only very low or even null values of λ satisfy


the local budget constraint. As a result, the inability to respond adequately in high-need areas further amplifies both spatial inequality and long-run infection persistence. As observed for the baseline case, comparison among scenarios reveals that Scenarios 1 and 3 tend to align in the HAs' decision outcome λ^* , while different initial conditions of Scenario 2 do not allow the average λ^* across cells to converge with other scenarios. However, considering the average cost-budget ratio across cells, the three scenarios all align and show similar levels compared to the baseline model (Figure 7). Despite price heterogeneity, the parsimonious condition allows all countries to deal with and contain the infection spread at the start of the simulation time. The ratio then quickly drops to values close to zero. Note that since this is given by $\frac{p^v \lambda^{*,v} I^v}{b^v}$ (see eq. 21), price directly and indirectly enters at the numerators given its contribution to the determination of $\lambda^{*,v}$ resulting in a way in a coordinating effect between price and $\lambda^{*,v}$ through price: higher levels of price correspond to smaller $\lambda^{*,v}$ and viceversa (see also Figure 8d).



(a) Average share of infected



(b) Average λ^*



(c) Average cost-to-budget ratio

Fig. 19: Average values across all countries for 2000 Montecarlo simulations

7 Conclusions

This paper develops and analyzes a spatially distributed agent-based SIR model with local optimization and budget-constrained treatment policies. Our framework captures key elements of epidemic management in real-world settings: heterogeneous local health authorities, costly but effective pharmaceutical treatments, and spatial spillovers that propagate infection across neighboring regions. The model departs from standard mean-field approaches by introducing decentralized decision-making, endogenous treatment intensity, and dynamically evolving infection risk at the cell level.

From a theoretical point of view, we characterize the stability properties of the disease-free equilibrium under both constant and optimal treatment strategies, highlighting the critical role played by treatment efficacy (ρ), the basic reproduction number ($\mathcal{R}_0$), and budget availability. The existence of a threshold treatment share $\bar{\lambda}$, above which the epidemic is eradicated, provides a clear policy target. However, once spatial interactions are introduced through a Moore neighborhood structure, the stability of the disease-free equilibrium becomes conditional on the entire spatial configuration, not just on local parameters. Infection spillovers can destabilize cells that would otherwise suppress the disease, illustrating the limits of purely local control.

Our numerical simulations confirm and extend these insights. We show that decentralized treatment decisions—when constrained by budget and guided by local infection levels—can lead to effective epidemic control in many scenarios. Nevertheless, spatial heterogeneity in infection, treatment intensity, and cost dynamics naturally emerges and persists, particularly in early periods. The local reproduction number $\mathcal{R}_0$ and the corresponding treatment behavior evolve in a highly asymmetric fashion across the grid, generating clusters of risk that delay eradication.

We also examine two important dimensions of heterogeneity: the efficacy of the treatment (ρ) and the pricing strategy adopted by the supplier of the innovative drug. A static comparative analysis of different values of ρ reveals that low efficacy can severely hinder the system’s ability to reach a low-infection equilibrium, even when optimal treatment behavior is applied. Conversely, higher efficacy allows for convergence to near-eradication within relatively few periods. When the price is fixed across cells, treatment intensity adapts to the infection level and budget constraint in a relatively smooth and effective manner. However, when the supplier adopts a parsimonious pricing strategy—setting the price based on local recovery benefits and treatment efficacy—this generates wide dispersion in treatment effort, ultimately exacerbating infection levels and inequality across regions.

Taken together, these results provide several policy-relevant insights. First, decentralized control mechanisms can be effective, but only if spatial spillovers are contained and treatment efficacy remains sufficiently high. Second, uniform pricing of life-saving treatments may in fact be preferable to recovery-based price discrimination, especially when health authorities operate under strict budget constraints. Finally, efforts to coordinate health responses spatially—particularly during early stages of contagion—may be crucial to avoiding persistent regional disparities and to accelerating convergence toward a disease-free state.

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 testbed to explore the performance of decentralized versus coordinated epidemic policies under alternative institutional and fiscal architectures.

Acknowledgements

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 for their insightful comments.


References

- [1] Hunter, E., Mac Namee, B., Kelleher, J.D.: A comparison of agent-based models and equation based models for infectious disease epidemiology. In: Proceedings of the 26th AIAI Irish Conference on Artificial Intelligence and Cognitive Science (2018)
- [2] Axtell, R.L., Farmer, J.D.: Agent-based modeling in economics and finance: Past, present, and future. *Journal of Economic Literature* **63**(1), 197–287 (2025)
- [3] Delli Gatti, D., Reissl, S., Russo, A., Gallegati, M.: Adaptive expectations, social learning and epidemic dynamics. *Physica A: Statistical Mechanics and its Applications* **537**, 122701 (2020)
- [4] Chang, S., Pierson, E., Koh, P.W., Gerardin, J., Redbird, B., Grusky, D., Leskovec, J.: Mobility network models of covid-19 explain inequities and inform reopening. *Nature* **589**(7840), 82–87 (2021)
- [5] Hunter, E., Mac Namee, B., Kelleher, J.D.: A taxonomy for agent-based models in human infectious disease epidemiology. *Journal of Artificial Societies and Social Simulation* **20**(3) (2017)
- [6] Squazzoni, F., Polhill, J.G., Edmonds, B., Ahrweiler, P., Antosz, P., Scholz, G., Chappin, É., Borit, M., Verhagen, H., Giardini, F., *et al.*: Computational models that matter during a global pandemic outbreak: A call to action. *JASSS-The Journal of Artificial Societies and Social Simulation* **23**(2), 10 (2020)
- [7] Kerr, C.C., Stuart, R.M., Mistry, D., Abeysuriya, R.G., Rosenfeld, K., Hart, G.R., Núñez, R.C., Cohen, J.A., Selvaraj, P., Hagedorn, B., *et al.*: Covasim: an agent-based model of covid-19 dynamics and interventions. *PLOS Computational Biology* **17**(7), 1009149 (2021)
- [8] Hinch, R., Probert, W.J., Nurtay, A., Kendall, M., Wymant, C., Hall, M., Lythgoe, K., Bulas Cruz, A., Zhao, L., Stewart, A., *et al.*: Openabm-covid19—an agent-based model for non-pharmaceutical interventions against covid-19 including contact tracing. *PLoS computational biology* **17**(7), 1009146 (2021)
- [9] Dubois, P., Magnac, T.: Managing an epidemic with budget-constrained policies. mimeo (2024). Working paper
- [10] Ajelli, M., Merler, S.: Comparing large-scale computational approaches to epidemic modeling: agent-based versus structured metapopulation models. *BMC Infectious Diseases* **10**(1), 190 (2010)
- [11] Ferguson, N.M., Cummings, D.A., Fraser, C., Cajka, J.C., Cooley, P.C., Burke, D.S.: Strategies for mitigating an influenza pandemic. *Nature* **442**(7101), 448–452 (2006)


- [12] Thomas, L.J., Huang, P., Yin, F., Luo, X.I., Almquist, Z.W., Hipp, J.R., Butts, C.T.: Spatial heterogeneity can lead to substantial local variations in covid-19 timing and severity. *Proceedings of the National Academy of Sciences* **117**(39), 24180–24187 (2020)
- [13] Kustudic, M., Niu, B., Liu, Q.: Agent-based analysis of contagion events according to sourcing locations. *Scientific Reports* **11**(1), 16032 (2021)
- [14] Epstein, J.M.: Modelling to contain pandemics. *Nature* **460**(7256), 687 (2009)
- [15] Danzon, P.M., Towse, A.: Differential pricing for pharmaceuticals: Reconciling access, r&d, and patents. *International Journal of Health Care Finance and Economics* **3**(3), 183–206 (2003)
- [16] Moon, S., Jambert, E., Childs, M., Von Schoen-Angerer, T.: A win-win solution?: A critical analysis of tiered pricing to improve access to medicines in developing countries. *Globalization and Health* **7**, 39 (2011)
- [17] Gleeson, D., Townsend, B., Tenni, B.F., Phillips, T.: Global inequities in access to covid-19 health products and technologies: A political economy analysis. *Health and Place* **83**, 103051 (2023)
- [18] Gozzi, N., Chinazzi, M., Dean, N.E., Longini, I.M., Halloran, M.E., Perra, N., Vespignani, A.: Estimating the impact of covid-19 vaccine inequities: A modeling study in lower-income countries. *Nature Communications* **14**, 225 (2023)
- [19] Gatti, D.D., Di Guilmi, C., Gallegati, M., Giulioni, G.: Financial fragility, industrial dynamics, and business fluctuations in an agent-based model. *Macroeconomic Dynamics* **11**(S1), 62–79 (2007)
- [20] Caravaggio, A., Sodini, M.: Local environmental quality and heterogeneity in an olg agent-based model with spatial externalities. *Journal of Economic Interaction and Coordination* **17**(1), 287–317 (2022)

A Appendix: Long run dynamics


Figure 20 shows the dynamic of the global share of infected population in the very long run for the baseline model. In the first and third scenario (see panel (a) and (c)) we observe a downward sloping curve of susceptible population while the share of recovered grows as the infection remains controlled, respectively, in hundreds and thousands of years. However, this is not observed for the second scenario, where the limited infection diffusion at setup prevents disease spread even in the very long run (5000 years). Note that the three graphs display a transient phase leading to a disease-free equilibrium, which is achieved due to the system's stability condition—not because of saturation. As discussed in Section 2, this condition implies sustained containment of the disease over time. Importantly, the disease-free state does not correspond to the entire population being susceptible (i.e., $S = N$), but rather to $S = N - R$, where a share of the population has been infected and then recovered.



(a) Scenario 1



(b) Scenario 2



(c) Scenario 3

Fig. 20: Long-run comparison of global infection dynamics across the three initialization scenarios (seed = 1000).

B Appendix: Efficacy of Treatment

Figures 21 and 22 show the final spatial distribution of infection ($t = 60$) respectively in the second and third scenarios for different values of the efficacy of treatment (ρ). As ρ increases, the containment of the disease is more effective, particularly for Scenario 2 where the epidemic is essentially eradicated across the entire spatial domain within the time span considered.







Fig. 21: Final spatial distribution of infection at $t = 60$ in Scenario 2, under different treatment efficacies (Seed 1000).




(a) $\rho = 0.3$



(b) $\rho = 0.5$



(c) $\rho = 0.7$



(d) $\rho = 1$

Fig. 22: Final spatial distribution of infection at $t = 60$ in Scenario 3, under different treatment efficacies (Seed 1000).