Journal of Economic Behavior and Organization 245 (2026) 107529

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

When cures are costly: Budget-constrained epidemic control and strategic drug pricing

Andrea Caravaggio [a], Andrea Di Liddo [b], Angela Martiradonna [b,*]

[a] Department of Economics and Statistics, University of Siena, Piazza San Francesco 7/8, Siena, 53100, Italy
[b] Department of Economics, University of Foggia, via Romolo Caggese 1, Foggia, 71121, Italy

ARTICLE INFO

JEL classification:
C61
C70
I18

Keywords:
Epidemic control
Budget constraints
Optimal treatment strategies
Health policy
Stackelberg game

ABSTRACT

Effective epidemic control policies must balance public health objectives with economic constraints, especially when deploying high-cost curative treatments for infectious diseases. This study develops an optimal control framework for a Health Authority (HA), embedded in a controlled SIR epidemic model, to minimize the societal and economic burden of infection under budget constraints. We analyze two institutional cost specifications: a quadratic cost with an explicit budget constraint and a blow-up cost with endogenous constraint enforcement. Numerical simulations calibrated to the case of hepatitis C in France illustrate how cost structures, budget design, and the health authority's preference over the final state of the epidemic jointly influence treatment intensity, infection dynamics, and long-term health outcomes. As an extension, we incorporate a Stackelberg differential game in which a Pharmaceutical Company (PC) strategically sets the drug price, anticipating the HA's optimal treatment response. We show that the PC's valuation of residual infection decisively shapes price trajectories and public health outcomes. The analysis suggests that effective policies should integrate coherent budget planning, outcome-sensitive pricing schemes, and regulatory incentives to align private behavior with long-term public health objectives.

1. Introduction

Ensuring equitable access to high-cost innovative treatments remains one of the most pressing challenges in contemporary health policy. This is particularly evident in the case of chronic and infectious diseases, where curative pharmaceutical innovations can drastically improve population health but pose significant budgetary and regulatory pressures. At the global level, the WHO's *Consolidated guidelines on HIV, viral hepatitis and STI prevention, diagnosis, treatment and care for key populations* reaffirm the global targets to eliminate viral hepatitis as a public health threat by 2030 (World Health Organization, 2022). The *Pharmaceutical Strategy for Europe* stresses the importance of aligning innovation with affordability, calling for transparent pricing, timely access, and the long-term sustainability of healthcare systems in Europe (European Commission, 2020). Complementary efforts such as the *EU Regulation on Health Technology Assessment* (Parliament, 2021) and the *Council Conclusions on access to medicines* (Council of the European Union, 2021) aim to reinforce collective capacity across Member States to evaluate and negotiate the value of new treatments. These institutional developments reflect a growing recognition that pricing decisions for breakthrough therapies are inherently strategic and must balance private incentives with public health priorities.

* Corresponding author.
E-mail address: angela.martiradonna@unifg.it (A. Martiradonna).

https://doi.org/10.1016/j.jebo.2026.107529
Received 9 July 2025; Received in revised form 2 February 2026; Accepted 16 March 2026
Available online 24 March 2026

In recent decades, the emergence of innovative pharmaceutical treatments for infectious diseases has raised complex economic and regulatory challenges for health authorities. While these innovations often promise curative outcomes, they also involve high upfront costs and uncertain long-term efficacy. Combined with rapidly evolving epidemiological contexts, this creates a tension between immediate budget constraints and long-term health outcomes. This is particularly salient in the case of communicable diseases, where treatment strategies must account not only for individual benefits but also for population-level externalities related to disease transmission and eradication.

The economic modeling of epidemic control has evolved significantly, particularly in understanding how health authorities allocate limited resources to minimize both societal and economic burdens. Classical approaches often formulate this problem as an optimal control framework, where the HA aims to reduce infection prevalence while accounting for treatment costs and budget constraints. For instance, Dubois et al. (2022) develop a dynamic SIR model calibrated to hepatitis C data in France, showing that myopic budget allocation strategies can lead to inferior long-term outcomes, especially when transmission dynamics and diminishing marginal effectiveness are taken into account. Their work highlights the importance of hard budget constraints and credible policy commitments in achieving socially optimal strategies. Moreover, Di Liddo (2016) analyzes the impact of different cost functions (quadratic vs. blow-up) on the feasibility and effectiveness of eradication policies under financial limitations. Other contributions further examine the role of cost function shapes and their influence on treatment timing and intensity, particularly in centralized systems with strict budget enforcement (see Buonomo and Vargas-De-León, 2014; Lashari and Zaman, 2012). Similar control approaches under budget constraints have also been investigated in ecological contexts, including resource management and conservation planning (Marangi et al., 2020, 2023). Finally, some contributions discuss adaptive or endogenous budget adjustments as policy tools to mitigate financial stress during epidemics, although these approaches are often static or limited to short-term horizons (PhRMA, 2017).

Building on this optimal control tradition, a growing literature embeds epidemic dynamics into fully optimizing economic–epidemiological models, typically adopting a centralized planner perspective. Within this framework, Garibaldi et al. (2024) study how forward-looking agents optimally adjust social activity in response to infection risk, and show that decentralized behavior may diverge from socially optimal epidemic paths due to the presence of static and dynamic externalities. Similarly, La Torre et al. (2021) analyze optimal containment policies over a finite planning horizon, emphasizing the intertemporal trade-off between economic costs and long-run health outcomes, and showing that optimal strategies may rationally tolerate endemic infection levels rather than full eradication. Recent extensions further enrich this planner-based approach by allowing for more realistic epidemiological structures and policy instruments. In Hritonenko and Yatsenko (2024), the authors introduce integral epidemic dynamics with realistic infectivity profiles and show that optimal lockdown policies may optimally end with persistent infection due to discounting and intertemporal trade-offs. Focusing on optimal vaccination policies in a SIRS framework with reinfection, Federico et al. (2024) provide a rigorous characterization of optimal feedback controls and show that sustained policy effort may be required to keep infection close to zero in the long run.

While these contributions offer powerful insights into the design of optimal containment, vaccination, and further policies under centralized planning, they typically take treatment prices as exogenous and abstract from the strategic behavior of pharmaceutical firms. As a result, a separate strand of the literature has focused on drug pricing and access, typically relying on static or reduced-form demand specifications and regulatory tools such as cost-effectiveness thresholds, capped pricing, and value-based pricing frameworks, thereby providing only a partial representation of the dynamic strategic interaction between pharmaceutical firms and public payers. Several authors have highlighted the limitations of fixed-price schemes or uncertain listing processes, showing that value-based pricing with risk-sharing components can lead to superior outcomes from both societal and industrial perspectives (Levaggi, 2014; Antonanzas et al., 2011). More recent contributions have incorporated game-theoretic perspectives to explicitly model strategic pricing behavior. For instance, Critchley and Zaric (2019) analyze pharmaceutical marketing and pricing strategies under different access policies, including negotiated pricing and value-based risk-sharing, showing that all policies except the first-best planner solution suffer from inefficiencies in access, pricing, or treatment coverage. Notably, value-based pricing with risk-sharing emerges as the most robust solution, although it yields zero net monetary benefit for the payer. Along the same research lines, several studies adopt static or reduced-form demand specifications, treating demand as exogenous (e.g., Gavious et al., 2014) or as a simple function of price (e.g., Miraldo, 2009).

Only a few studies attempt to integrate dynamic disease progression with pricing decisions. Among these, Di Liddo (2020) develops a stylized SIS model where a pharmaceutical firm and a social planner engage in Stackelberg or Nash bargaining interactions over price and treatment coverage, highlighting how the choice of cost function significantly affects equilibrium outcomes. In addition, Lothan et al. (2022) model the interaction between a pharmaceutical company and a government negotiating HCV treatment prices, with the government setting screening and treatment strategies in response to fixed pricing. Their findings show that under certain conditions, manufacturers may benefit from offering quantity discounts or funding screening programs to induce higher treatment uptake, while risk-sharing agreements (e.g., payment only upon successful treatment) may improve outcomes for the public payer.

Although the literature has made important strides, the strategic interplay between pricing and treatment policies within fully dynamic settings, where both epidemiological dynamics and contractual negotiations co-evolve, remains largely underexplored. In particular, the integration of epidemic models like the SIR framework with dynamic game-theoretic approaches is still limited, despite its importance for evaluating realistic policy interventions. Building on these premises, we propose an optimal control framework to analyze the epidemic management problem faced by a health authority (HA) responsible for delivering high-cost curative treatments for infectious diseases. The epidemic dynamics are described through a controlled SIR model, where the HA determines a time-dependent treatment intensity that balances the societal benefits of reducing infections against the economic burden of pharmaceutical spending, all under explicit or endogenous budget constraints. We examine two institutional cost structures: a quadratic cost with an explicit budget constraint and a blow-up cost that embeds the constraint endogenously. Numerical simulations, calibrated to the

case of hepatitis C in France, illustrate how budget design and cost parameters shape treatment intensity, infection dynamics, and eradication prospects. The analysis is conducted within an open-loop framework, reflecting *ex ante* policy commitment under budget and institutional constraints, while a feedback formulation and its methodological implications are briefly discussed. As an extension, we formulate a Stackelberg differential game in which a pharmaceutical company (PC) strategically sets the drug price, anticipating the HA's optimal treatment response. In this hierarchical framework, the PC acts as the leader and the HA as the follower, enabling us to explore the interplay between strategic pricing and epidemic control outcomes.

The analysis is guided by the following core research questions:

1. How do budget constraints and different cost structures influence the feasibility, regularity, and effectiveness of optimal treatment strategies over time?
2. How does the dynamic interplay between treatment intensity and budget design affect epidemic trajectories and eradication thresholds?
3. How do the pharmaceutical company's strategic pricing incentives reshape public health outcomes when pricing is endogenized through a Stackelberg game?

To address these questions, we first characterize the HA's optimal control problem under both cost specifications, and then extend the analysis to the Stackelberg framework under quadratic costs. Numerical results highlight key policy trade-offs and suggest that coherent budget planning and outcome-sensitive pricing incentives are crucial to align private interests with long-term public health goals.

The paper is structured as follows. Section 2 introduces the epidemiological model and derives analytical results on eradication thresholds. Section 3 formulates the HA's optimal control problem under two alternative cost structures. Section 4 presents numerical simulations illustrating the effects of budget design and cost parameters. Section 5 extends the analysis to a Stackelberg differential game with strategic pricing. Section 6 discusses the policy implications of our findings in light of current European and global frameworks. Section 7 concludes the paper by addressing the research questions and outlining future directions.

2. The epidemic SIR model with treatment

We model the spread of an infectious disease in a population using a classical SIR compartmental framework (Hethcote, 1989). At each time $t \geq 0$, the population is divided into three groups: susceptible, infected, and recovered individuals, with respective densities $S(t)$, $I(t)$, and $R(t)$, all taking values in the interval $[0, 1]$. We assume that the population experiences a constant birth rate $\nu > 0$, with all newborns entering the susceptible class. The death rate is also set equal to ν, ensuring a constant total population density over time. Infection occurs through contact between susceptible and infected individuals, at a transmission rate $\beta > 0$.

At the initial time ($t = 0$), an *innovative drug* is introduced to treat the disease by converting infected individuals into permanently recovered ones. The treatment is administered under the supervision of a *Health Authority* (HA), which sets a time-dependent control variable $\lambda(t) \geq 0$ representing the intensity of treatment applied to the infected population. The removal of treated individuals follows an exponential process at rate $\lambda(t)$.

Under these assumptions, the epidemic SIR model with treatment control $\lambda(t)$ is described by the following system of differential equations:

$$\begin{cases} \dot{S}(t) = \nu\, N(t) - \nu\, S(t) - \beta\, I(t)\, S(t), \\ \dot{I}(t) = \beta\, I(t)\, S(t) - \nu\, I(t) - \lambda(t)\, I(t), \\ \dot{R}(t) = \lambda(t) I(t) - \nu\, R(t), \end{cases} \tag{1}$$

where $N(t) = S(t) + I(t) + R(t)$ denotes the total population density. The term νN accounts for the constant influx of newborns entering the susceptible class, while the terms $-\nu S$, $-\nu I$, and $-\nu R$ represent natural deaths occurring uniformly across the three compartments. The infection term βSI captures the transmission of the disease through contact between susceptible and infected individuals. Finally, $\lambda(t)I$ represents the rate at which infected individuals recover as a consequence of the applied treatment.

The sum of the three equations in (1) yields $\dot{N}(t) = 0$ for all $t \geq 0$, implying that the total population remains constant, i.e., $N(t) = N(0)$ for all $t \geq 0$. We assume the initial conditions:

$$S(0) = S_0 \geq 0, \quad I(0) = I_0 \geq 0, \quad R(0) = R_0 \geq 0,$$

such that $N(0) = S_0 + I_0 + R_0 = 1$. Therefore, $N(t) = 1$ for all t, and we can write $R(t)$ as: $R(t) = 1 - S(t) - I(t)$. Thus, system (1) simplifies to:

$$\begin{cases} \dot{S}(t) = \nu - \nu\, S(t) - \beta\, I(t)\, S(t), \\ \dot{I}(t) = \beta\, I(t)\, S(t) - \nu\, I(t) - \lambda(t)\, I(t), \end{cases} \tag{2}$$

and the resulting equation for recovered density is

$$\dot{R} = (\nu + \lambda(t))\, I + \nu\, (S - 1). \tag{3}$$

As $R(t)$ can be recovered from the knowledge of $S(t)$ and $I(t)$, our analysis can be restricted to system (2) alone.

Observe that the set $\Omega = [0, 1]^2$, representing all biologically meaningful values for the susceptible and infected population shares, is invariant under the dynamics of system (2) for any nonnegative treatment strategy $\lambda(t) \geq 0$. This means that if the initial condition

satisfies $S_0, I_0 \in [0,1]$, then the solutions $S(t)$ and $I(t)$ remain in the interval $[0,1]$ for all $t \geq 0$. Moreover, if $I_0 > 0$, then $I(t) > 0$ for all $t \geq 0$; that is, the infection level remains strictly positive at all times, although it may approach zero asymptotically depending on the treatment policy $\lambda(t)$. In what follows, we assume $I_0 > 0$.

2.1. Equilibria and eradication threshold under constant treatment

To gain insight into the long-term behavior of the system, we focus on the case where the treatment intervention assumes a constant value $\lambda(t) = \lambda \geq 0$. In this setting, we study the existence and stability of the equilibria of system (2), which are governed by the *basic reproduction number*

$$\mathcal{R}_0 = \frac{\beta}{\nu + \lambda}.$$

The basic reproduction number $\mathcal{R}_0$ represents the expected number of secondary infections generated by a single infected individual introduced into a fully susceptible population. It provides a synthetic measure of the epidemic potential of the disease by comparing the rate at which new infections are produced (captured by the transmission parameter β) with the overall rate at which infected individuals are removed from the infectious class. In the present model, removal occurs both through natural turnover at rate ν and through treatment efforts at rate λ (Bailey, 1975; Hethcote, 2000; Diekmann et al., 2010) .

The system always admits the disease-free equilibrium $P_0 = (1, 0)$. To assess its local stability, we compute the Jacobian matrix of system (2):

$$J(S, I) = \begin{pmatrix} -\nu - \beta I & -\beta S \\ \beta I & \beta S - \nu - \lambda \end{pmatrix}.$$

Evaluated at P_0, the Jacobian becomes upper triangular, with eigenvalues $-\nu < 0$ and $(\nu + \lambda)(\mathcal{R}_0 - 1)$. Therefore, the disease-free equilibrium P_0 is locally asymptotically stable if $\mathcal{R}_0 < 1$, and unstable when $\mathcal{R}_0 > 1$.

When $\mathcal{R}_0 > 1$, the system also admits an endemic equilibrium $P_e = (S_e, I_e)$, where

$$S_e = \frac{1}{\mathcal{R}_0}, \quad I_e = \frac{\nu}{\beta}(\mathcal{R}_0 - 1).$$

A direct computation of the Jacobian at P_e shows that the equilibrium is locally asymptotically stable, since the eigenvalues are the roots of a second-order polynomial with strictly positive coefficients. We summarize these results in the following proposition:

Proposition 1. *Assume that $\lambda(t) = \lambda \geq 0$ is constant. Then, system (2) admits the disease-free equilibrium $P_0 = (1, 0)$. If $\mathcal{R}_0 < 1$, P_0 is the only equilibrium and it is linearly asymptotically stable. Conversely, when $\mathcal{R}_0 > 1$, P_0 becomes linearly asymptotically unstable, and the asymptotically stable endemic equilibrium $P_e = (S_e, I_e)$ arises.*

Remark 1. In the threshold case $\mathcal{R}_0 = 1$, one eigenvalue vanishes. Then, this critical value corresponds to a transcritical bifurcation at which stability is exchanged between the disease-free equilibrium P_0 and the endemic equilibrium P_e.

According to Proposition 1, disease extinction occurs when the basic reproduction number satisfies $\mathcal{R}_0 < 1$. In the case of constant treatment effort λ, this translates into the condition:

$$\lambda > \lambda_c := \beta - \nu. \tag{4}$$

Here, λ_c defines the critical treatment threshold above which the infection is eradicated. Notice that when $\beta - \nu < 0$, the threshold λ_c is negative, implying that disease extinction is guaranteed for all $\lambda \geq 0$, even in the absence of treatment.

In the more general case where the treatment intensity varies over time, the eradication of the infection requires that the level of treatment eventually becomes (and remains) sufficiently high. In practical terms, this means that the control effort must exceed a certain critical threshold for a long enough period to suppress the infection. If this condition is not met, either because the treatment is too weak or is interrupted for an extended time, the infection may not disappear. Instead, it can persist at low levels or even exhibit prolonged and potentially unpredictable transient dynamics before eventually stabilizing at an endemic equilibrium.

These considerations naturally lead to a broader question: while intensive treatment efforts can rapidly eradicate the disease, they are typically associated with substantial economic and logistical costs. In practice, health authorities must operate under budget constraints and resource limitations, which prevent them from indefinitely sustaining high treatment intensities. As a result, a key challenge arises: how should a HA optimally allocate treatment efforts over time to minimize the impact of the epidemic, while respecting economic constraints? To address this, we extend our analysis by formulating an optimal control problem in which the HA aims to minimize a cost functional that balances the societal cost of infection with the economic cost of treatment. This framework allows us to identify time-dependent treatment strategies that are both effective in suppressing the infection and feasible from a budgetary perspective.

3. The health authority's decision problem: optimizing treatment under economic constraints

In this section, we model the decision-making problem faced by a HA that aims to minimize the total cost associated with the epidemic. The cost includes both the societal burden of infection and the economic cost of implementing treatment interventions. We assume that treatment can be modulated over time via a control function

$$\lambda(t) \geq 0, \tag{5}$$

representing the intensity of effort allocated to recovering infected individuals. We further assume that an innovative drug is available for treatment, which is highly effective in reducing the duration of infection. As discussed in Dubois and Magnac (2024), this is the case for direct-acting antivirals (DAAs) used to treat hepatitis C in France, which achieved high cure rates within short treatment durations but initially involved very high per-patient costs.

Let $p(t) > 0$ denote the unit price at which the Pharmaceutical Company (PC) sells the innovative treatment at time t, corresponding to the cost of providing the therapy to a single infected individual. While the price may vary over time due to negotiations, market dynamics, or policy decisions, we treat $p(t)$ as an exogenous function throughout this analysis, leaving its strategic determination to an extended version of the model presented in Section 5. We assume that the entire treatment is administered at time t, and that the corresponding cost is incurred instantaneously. The corresponding instantaneous treatment expenditure per unit of population is thus defined as:

$$h(t) = p(t)\lambda(t)I(t), \tag{6}$$

which represents the treatment cost incurred by the HA at time t, normalized per unit of total population.

Let $b(t)$ denote the budget available for pharmaceutical spending at time t, expressed per unit of total population. The HA must ensure that the instantaneous treatment expenditure satisfies

$$h(t) \leq b(t), \tag{7}$$

so as not to exceed the available financial resources. Since the unit price $p(t)$ is not directly determined by the HA (its strategic pricing will be addressed in Section 5), for the purposes of the present analysis we treat $p(t)$ as an exogenous input and focus on the optimal allocation of treatment intensity by the HA, subject to the given budget constraints.

3.1. Economic constraints and cost modeling

The HA's objective is to minimize the overall cost associated with the epidemic over a finite planning horizon $[0, T]$. Let $\delta > 0$ denote the discount rate, which reflects the preference for immediate costs and benefits over future ones. Following Di Liddo (2016), the total cost incurred by the HA is represented by the following discounted cost functional:

$$\mathrm{e}^{-\delta T}\theta\, I(T) + \int_0^T \mathrm{e}^{-\delta t}(\eta\, I(t) + \phi(t, h(t)))dt. \tag{8}$$

The first term in (8), $\mathrm{e}^{-\delta T}\theta I(T)$, represents a terminal (or *scrap*) cost incurred at the end of the planning horizon. The parameter $\theta > 0$ measures the weight that the HA assigns to the epidemic outcome at time T: higher values of θ reflect a stronger preference for strategies that achieve low residual infection or eradication by the end of the horizon. Beyond its formal role in a finite-horizon optimal control problem, this terminal term captures the HA's valuation of residual infection as a reduced-form continuation value beyond time T. In practice, nonzero prevalence at the end of a policy cycle entails future healthcare expenditures, political accountability costs, and institutional losses that are not explicitly modeled when the horizon is truncated. The parameter θ therefore represents the shadow cost of transferring an unresolved epidemic to future governments or budget cycles, providing an economic rationale for sustained treatment effort close to the terminal date and mitigating the well-known end-of-horizon under-investment problem in finite-time policy planning.

The second term, $\eta\, I(t)$, represents the *societal cost* associated with the presence of infected individuals in the population. This component captures both direct and indirect consequences of infection, including the burden on healthcare services, productivity losses due to morbidity, social disruption, and potential long-term complications. The parameter $\eta > 0$ summarizes the average cost imposed on society per unit of infected population (Fonkwo, 2008). Finally, the term $\phi(t, h(t))$ captures the economic cost associated with the HA's pharmaceutical spending. It reflects the disutility, or perceived burden, of financing treatment interventions over time. The function $\phi(t, h(t))$ is assumed to be increasing and convex in $h(t)$, indicating that the marginal cost of allocating additional resources to treatment rises with the intensity of expenditure. This formulation allows the model to capture inefficiencies, administrative overhead, and systemic constraints that arise when the health system operates under pressure or near capacity.

Since the HA must comply with the budget constraint (7), which ensures that treatment costs never exceed the available financial resources at any time $t \in [0, T]$, it is crucial to adopt a functional form for $\phi(t, h(t))$ that captures both the increasing disutility of treatment intensity and the economic limitations imposed by the budget. To this end, we consider the following two alternative specifications:

(i) a *quadratic cost function* of the form $\phi_q(t, h(t)) = \dfrac{\alpha_q}{2}(h(t))^2$, which models increasing marginal cost and diminishing returns in pharmaceutical spending (see for example Buonomo and Vargas-De-León, 2014; Lashari and Zaman, 2012). In this formulation, the budget constraint must be explicitly included in the control problem as an inequality constraint.

(ii) a *blow-up cost function* of the form $\phi_b(t, h(t)) = \dfrac{\alpha_b\, b(t)\, h(t)}{b(t) - h(t)}$, with $h(t) < b(t)$, which implicitly incorporates the budget limitation into the cost function itself. In this case, the cost becomes arbitrarily large as $h(t)$ approaches the budget ceiling $b(t)$, effectively discouraging the optimizer from exceeding it. The budget constraint is thus enforced endogenously, and there is no need to impose it as a separate inequality in the control problem. This type of cost function has been previously adopted in Di Liddo (2016) to account for the dramatic increase in cost associated with extremely expensive therapies when the number of patients is high.

In both specifications, we assume $\alpha_q, \alpha_b > 0$ as parameters that scale the total cost up or down (see Buccella et al., 2021). They represent exogenous indexes that measure the cost factor related to the HA's ability to distribute treatment to patients. Specifically, low α_q and α_b can be interpreted as a high ability of the HA to deliver the drug (e.g., through the use of digital platforms on which patients can immediately book to initiate treatment or through direct physician-patient contact) reflected in a reduced cost. Conversely, a high α_q and α_b can be interpreted as less effective delivery of purchased treatments, resulting in higher costs.[1] Considering both of these cost specifications enables us to effectively address differing organizational strategies for pharmaceutical budgeting. The quadratic cost function is particularly suited to contexts where the HA plans and enforces its expenditure limits proactively, which is characteristic of centralized systems with well-defined annual budgets. In these scenarios, increasing marginal costs capture logistical challenges and resource constraints that arise as treatment intensity grows. In contrast, the blow-up cost formulation is more appropriate in settings where budget constraints are not enforced ex ante but emerge as spending rises, a common occurrence in fragmented or decentralized systems. In this case, the sharp increase in costs near the budget threshold effectively acts as a soft constraint, discouraging overspending without requiring strict enforcement mechanisms.

In the following subsections, we analyze in detail the two optimal control problems associated with cost specifications (i) and (ii), highlighting how the formulation of the treatment cost affects the structure and constraints of the optimization problem.

Both problems are formulated and solved within an open-loop framework, in which the HA commits at the initial time to a treatment intensity path over a fixed planning horizon. The main objective is to analyze epidemic control under budget constraints within a strategic and institutional setting. Accordingly, treatment and pricing policies are specified ex ante on the basis of available information, reflecting institutional environments in which policies, budget allocations, and pricing agreements are defined in advance and implemented over predetermined time horizons. The open-loop formulation allows for a focus on strategic commitment, intertemporal trade-offs, and budget-constrained decision making, and provides a transparent characterization of how economic and institutional constraints shape treatment dynamics over time.

3.2. Optimal treatment under quadratic costs and explicit budgets

By substituting the quadratic cost specification

$$\phi_q(t, h(t)) = \frac{\alpha_q}{2}(h(t))^2$$

into the general cost functional (8), and using the definition of $h(t)$ from Eq. (6), the objective functional associated with a given treatment strategy $\lambda(t)$ over the time horizon $[0, T]$ becomes:

$$\mathcal{J}_q^{\mathrm{HA}}(\lambda(\cdot)) = \mathrm{e}^{-\delta T}\theta\, I(T) + \int_0^T \mathrm{e}^{-\delta t}\left[\eta\, I(t) + \frac{\alpha_q}{2}(p(t)\,\lambda(t)\,I(t))^2\right] dt. \tag{9}$$

In this setting, the HA's problem, given the exogenous pricing function $p(t)$, is to determine the optimal treatment strategy that solves the following optimal control problem:

$$\min_{\lambda\in\Lambda_q} \mathcal{J}_q^{\mathrm{HA}}(\lambda(\cdot)) \tag{10}$$

subject to the epidemic dynamics (2) and the budget constraint derived from (7):

$$b(t) - p(t)\,\lambda(t)\,I(t) \geq 0, \quad \forall\, t \in [0, T]. \tag{11}$$

To ensure that the optimal control problem is well posed, in the sense that an optimal treatment strategy exists, we impose the following regularity assumption on the pricing and budget functions:

(A1) The functions $p(t)$ and $b(t)$ are strictly positive, bounded, and piecewise continuous on the interval $[0, T]$.

Given the quadratic cost structure and the epidemic dynamics described by (2), we consider control functions $\lambda(t)$ belonging to the Hilbert space $L^2([0, T])$, and satisfying the nonnegativity and budget constraints almost everywhere. This leads to the following set of admissible controls:

$$\Lambda_q = \left\{ \lambda \in L^2([0,T]) \;\middle|\; \lambda(t) \geq 0,\ p(t)\lambda(t)I(t) \leq b(t),\ \text{a.e. in } [0,T] \right\}.$$

Under assumption (A1), the budget constraint is measurable and bounded, and the cost functional (9) is convex in the control variable. Moreover, the state system (2) is well posed: it admits a unique solution for each admissible control, and the state trajectories depend continuously on $\lambda(t)$. These properties ensure the existence of an optimal solution, as established by standard results in optimal control theory; see, for instance, Cesari (2012) and Seierstad and Sydsaeter (1986).

Turning to the derivation of necessary conditions for optimality in an open-loop setting, we begin by introducing the Hamiltonian function associated with the unconstrained version of the problem, that is, in the absence of the budget limitation:

$$\mathcal{H}_q^{\mathrm{HA}}(S, I, \psi_S, \psi_I, \lambda, t) = \mathrm{e}^{-\delta t}\left(\eta\, I + \frac{\alpha_q}{2}(p(t)\,\lambda\, I)^2\right) + \psi_S^{\mathrm{HA}}(\nu - \nu\, S - \beta\, I\, S) + \psi_I^{\mathrm{HA}}(\beta\, I\, S - \nu\, I - \lambda\, I).$$

[1] For a more in-depth description of the economic and institutional factors that determine α_q and α_b, see Shepard et al. (2013). Furthermore, these parameters ensure consistency in units between the HA's objective functional—expressed in monetary terms over time—and the nonlinear cost terms $\phi_q(t, h(t))$ and $\phi_b(t, h(t))$, respectively.

where $\psi_S^{\mathrm{HA}}(t)$, $\psi_I^{\mathrm{HA}}(t)$ are the adjoint variables associated with the state dynamics of susceptibles and infectives, respectively.

To account for the control constraint (5) and for the budget inequality constraint (11), following the approach in Seierstad and Sydsaeter (1986), we derive the necessary conditions of optimality by introducing two time-dependent Lagrange multipliers $\mu_1(t) \le 0$; $\mu_2(t) \le 0$ associated with these conditions. The corresponding *Lagrangian* (also referred to as the *generalized Hamiltonian*) is defined as:

$$\mathcal{L}^{\mathrm{HA}}(S, I, \psi_S, \psi_I, \lambda, \mu_1, \mu_2, t) = \mathcal{H}_q^{\mathrm{HA}}(S, I, \psi_S, \psi_I, \lambda, t) + \mu_1(t)\,\lambda(t) + \mu_2(t)\,(b(t) - p(t)\,\lambda\,I). \tag{12}$$

The following proposition, derived from Seierstad and Sydsaeter (1986)(Th. 1, Ch. 4), establishes necessary conditions for optimality based on the Pontryagin Maximum Principle.

Proposition 2. *Let $\lambda^*(t) \ge 0$ be an optimal treatment strategy for the HA problem, and let $S^*(t)$ and $I^*(t)$ denote the corresponding state variables that solve system (2) under $\lambda^*(t)$. Then, there exist piecewise differentiable adjoint functions $\psi_S^{HA}(t)$, $\psi_I^{HA}(t)$, and piecewise continuous multipliers $\mu_1(t) \le 0$, $\mu_2(t) \le 0$, such that the following conditions hold for all $t \in [0, T]$:*

(a) Hamiltonian minimization under budget feasibility:

$$\mathcal{H}_q^{HA}(S^*(t), I^*(t), \psi_S^{HA}(t), \psi_I^{HA}(t), \lambda^*(t), t) \le \mathcal{H}_q^{HA}(S^*(t), I^*(t), \psi_S^{HA}(t), \psi_I^{HA}(t), \lambda(t), t)$$

for all $\lambda(t) \ge 0$ such that $b(t) - p(t)\,\lambda(t)\,I^*(t) > 0$;

(b) First-order optimality condition:

$$\frac{\partial \mathcal{L}^{HA}}{\partial \lambda}(S^*(t), I^*(t), \psi_S^{HA}(t), \psi_I^{HA}(t), \lambda^*(t), \mu_1(t), \mu_2(t), t) = 0$$

(c) Complementary slackness condition:

$$\begin{cases} \mu_1(t) \le 0 \ \text{ and } \ \mu_1(t) = 0 \ \text{ if } \ \lambda^*(t) > 0, \\ \mu_2(t) \le 0; \ \text{ and } \ \mu_2(t) = 0 \ \text{ if } \ b(t) - p(t)\lambda^*(t)I(t) > 0. \end{cases}$$

(d) Adjoint equations:

$$\begin{cases} \dot{\psi}_S^{HA} = -\dfrac{\partial \mathcal{L}^{HA}}{\partial S}(S^*(t), I^*(t), \psi_S^{HA}(t), \psi_I^{HA}(t), \lambda^*(t), \mu_1(t), \mu_2(t), t), \\ \dot{\psi}_I^{HA} = -\dfrac{\partial \mathcal{L}^{HA}}{\partial I}(S^*(t), I^*(t), \psi_S^{HA}(t), \psi_I^{HA}(t), \lambda^*(t), \mu_1(t), \mu_2(t), t); \end{cases}$$

(e) Transversality conditions: $\psi_S^{HA}(T) = 0, \psi_I^{HA}(T) = \mathrm{e}^{-\delta T}\theta$.

Based on the necessary conditions derived above, we now summarize the full optimality system satisfied by the state and adjoint variables, the optimal control, and the Lagrange multiplier. This system characterizes any solution to the HA's problem under quadratic treatment costs and a budget constraint. Specifically, since $I^*(t)$ remains strictly positive due to the Assumption $I_0 > 0$, and $p(t) > 0$ by assumption (A1), the first-order optimality condition (b) admits a closed-form expression. This leads to a pointwise characterization of the optimal control. Accordingly, the HA's problem reduces to solving the following forward–backward system:

$$\begin{cases} \dot{S}^* = \nu - \nu S^* - \beta\, S^*\, I^*, & S^*(0) = S_0, \\ \dot{I}^* = \beta\, S^*\, I^* - \nu\, I^* - \lambda^*(t)\, I^*, & I^*(0) = I_0, \\ \dot{\psi}_S^{\mathrm{HA}} = (\nu + \beta\, I^*)\,\psi_S^{\mathrm{HA}} - \beta\, I^*\,\psi_I^{\mathrm{HA}}, & \psi_S^{\mathrm{HA}}(T) = 0, \\ \dot{\psi}_I^{\mathrm{HA}} = -\mathrm{e}^{-\delta t}\left(\eta + \alpha_q\,(p(t)\,\lambda^*(t))^2\, I^*\right) + \beta\, S^*\,\psi_S^{\mathrm{HA}} & \\ \qquad\qquad -(\beta\, S^* - \nu - \lambda^*(t))\,\psi_I^{\mathrm{HA}} + \mu_2(t)\,p(t)\,\lambda^*(t), & \psi_I^{\mathrm{HA}}(T) = \mathrm{e}^{-\delta T}\theta, \\ \lambda^*(t) = \dfrac{\mathrm{e}^{\delta t}\left(\psi_I^{\mathrm{HA}}(t) + \mu_2(t)\,p(t)\right)}{\alpha_q\, p(t)^2\, I^*(t)} - \mu_1 \dfrac{\mathrm{e}^{\delta t}}{\alpha_q\, p(t)^2\,(I^*(t))^2}, & \\ \mu_1(t) \le 0, \quad \mu_1(t) = 0 \quad \text{if } \lambda^*(t) > 0. & \\ \mu_2(t) \le 0, \quad \mu_2(t) = 0 \quad \text{if } p(t)\,\lambda^*(t)\,I^*(t) < b(t). & \end{cases} \tag{13}$$

Note that the expression for $\lambda^*(t)$ can also be equivalently written as:

$$\lambda^*(t) = \max\left\{0,\ \frac{\mathrm{e}^{\delta t}\left(\psi_I^{\mathrm{HA}}(t) + \mu_2(t)\,p(t)\right)}{\alpha_q\, p(t)^2\, I^*(t)}\right\}.$$

The expression for the optimal control $\lambda^*(t)$ in the quadratic-cost case reflects the balance between the marginal benefit of reducing infections and the marginal economic burden of treatment. The structure of the solution can be interpreted through two distinct regimes, governed by the activation of the budget constraint:

- *Unconstrained interior solution*: When the optimal level of treatment satisfies the budget constraint $p(t)\lambda^*(t)I^*(t) < b(t)$, the Lagrange multiplier vanishes ($\mu_2(t) = 0$), and the control is given by the interior formula

$$\lambda^*_{int}(t) = \frac{\mathrm{e}^{\delta t}\psi_I^{\mathrm{HA}}(t)}{\alpha_q\, p(t)^2\, I^*(t)}. \tag{14}$$

  In this case, the optimal strategy increases with the marginal value of reducing infections, $\psi_I^{\mathrm{HA}}(t)$, and decreases with treatment cost and disease prevalence.

- *Constrained solution*: When the budget constraint becomes binding, the multiplier $\mu(t) < 0$ adjusts the optimal control downward to ensure that $p(t)\lambda^*(t)I^*(t) = b(t)$. The control still follows the same formula but now incorporates the additional marginal cost captured by $\mu_2(t)$. In this regime, the available budget limits the intensity of treatment, regardless of its epidemiological desirability.

On sufficiency and local uniqueness. Due to the bilinear infection term SI and the bilinear coupling λI, the Hamiltonian fails to be concave in the joint state–control variables. As a consequence, classical sufficiency results based on global concavity, such as those in Seierstad and Sydsaeter (1986), do not apply. In this setting, sufficiency is interpreted in accordance with existence and characterization results for optimal control problems based on the Pontryagin Maximum Principle in Bressan and Piccoli (2007). More precisely, whenever the solution of the optimality system (13) is unique on a given time horizon, the corresponding extremal is the only admissible trajectory satisfying the necessary conditions and therefore coincides with the optimal control.

Uniqueness of the optimality system (13) can be established on sufficiently short time horizons T by means of a direct energy argument. The proof relies on uniform a priori bounds and on a contraction property for the coupled state and adjoint variables, following approaches similar to those used in the analysis of state control optimality systems in Baker et al. (2018), Marangi et al. (2020). For sufficiently short horizons, the epidemiological dynamics ensure that the infected fraction remains uniformly bounded away from zero. In particular, under the assumption that the unit price satisfies $p(t) \geq p_1 > 0$, the budget constraint implies a uniform bound on treatment expenditure, namely

$$0 \leq \lambda(t)I(t) \leq \frac{b(t)}{p(t)} \leq \frac{1}{p_1}\|b\|_{L^\infty(0,T)}.$$

This bound allows one to derive a uniform lower bound on the infected population over sufficiently short time horizons, ensuring that $I(t)$ does not approach zero. This property makes it possible to compare candidate solutions of the optimality system through suitable energy estimates. The resulting inequalities imply a contraction property for the coupled state and adjoint system on short horizons, which in turn ensures uniqueness of the solution. As a consequence, in this regime the Pontryagin necessary conditions also provide a sufficient characterization of the optimal control. The proof is omitted for brevity.

3.3. Implicit budget enforcement via blow-up costs

We now turn to the case in which the treatment cost is modeled using a *blow-up* cost function, given by

$$\phi_b(t, h(t)) = \frac{\alpha_b\, b(t)\, h(t)}{b(t) - h(t)},$$

where $h(t) = p(t)\,\lambda(t)\,I(t)$ denotes the instantaneous treatment expenditure per unit of population, and $b(t)$ is the corresponding pharmaceutical budget. This formulation captures the steep rise in marginal cost as spending nears the available resources, effectively discouraging the system from reaching that threshold. Substituting this specification into the general cost functional (8) yields:

$$\mathcal{J}_b^{\mathrm{HA}}(\lambda(\cdot)) = \mathrm{e}^{-\delta T}\theta\, I(T) + \int_0^T \mathrm{e}^{-\delta t}\left[\eta\, I(t) + \frac{\alpha_b\, b(t)\, p(t)\,\lambda(t)\, I(t)}{b(t) - p(t)\,\lambda(t)\, I(t)}\right] dt. \tag{15}$$

Since the cost becomes unbounded as spending approaches $b(t)$, this structure enforces the resource constraint implicitly. The resulting formulation excludes the need for an explicit budget inequality, as any attempt to exceed the limit leads to an infinite penalty.

To ensure feasibility of the cost functional (15), we define the admissible control set as:

$$\Lambda_b = \left\{ \lambda \in L^\infty([0,T]) \;\middle|\; \lambda(t) \geq 0,\ p(t)\lambda(t)I(t) < b(t) \text{ a.e. in } [0,T] \right\}.$$

The HA's problem then reduces to determining the optimal treatment policy $\lambda(t) \in \Lambda_b$ that minimizes the cost functional (15), subject to the epidemic dynamics described by system (2). Formally, the control problem reads:

$$\min_{\lambda \in \Lambda_b} \mathcal{J}_b^{\mathrm{HA}}(\lambda(\cdot)) \tag{16}$$

subject to the state dynamics (2).

Although the cost functional contains a singularity as $h(t) \to b(t)^-$, the admissible control set Λ_b is explicitly constructed to avoid this blow-up. In this setting, the integrand remains finite along all admissible trajectories, and the cost functional is convex with respect to the control variable. In addition, the state equations Eq. (2) are regular and depend continuously on the control, ensuring that the state-control mapping is well defined and stable. Together, these properties allow us to apply standard existence results for optimal control problems with singular integrands (see, e.g., Seierstad and Sydsaeter, 1986), thus guaranteeing the existence of an optimal solution.

To derive the necessary conditions for optimality, we define the Hamiltonian associated with the blow-up cost functional (15). It takes the form:

$$\mathcal{H}_b^{\mathrm{HA}}(S, I, \psi_S, \psi_I, \lambda, t) = \mathrm{e}^{-\delta t}\left(\eta\, I + \frac{\alpha_b\, b(t)\, p(t)\,\lambda\, I}{b(t) - p(t)\,\lambda\, I}\right) + \psi_S^{\mathrm{HA}}(\nu - \nu S - \beta SI) + \psi_I^{\mathrm{HA}}(\beta SI - \nu I - \lambda I),$$

where $\psi_S^{\mathrm{HA}}(t)$ and $\psi_I^{\mathrm{HA}}(t)$ are the adjoint variables associated with the dynamics of susceptibles and infectives, respectively.

Based on the necessary conditions and the characterization of the blow-up cost structures, the HA's optimal control problem leads to the following forward–backward system:

$$\begin{cases} \dot{S}^* = \nu - \nu\, S^* - \beta\, S^*\, I^*, & S^*(0) = S_0, \\ \dot{I}^* = \beta\, S^*\, I^* - \nu\, I^* - \lambda^*\, I^*, & I^*(0) = I_0, \\ \dot{\psi}_S^{\mathrm{HA}} = (\nu + \beta\, I^*)\, \psi_S^{\mathrm{HA}} - \beta\, I^*\, \psi_I^{\mathrm{HA}}, & \psi_S^{\mathrm{HA}}(T) = 0, \\ \dot{\psi}_I^{\mathrm{HA}} = -\mathrm{e}^{-\delta t}\left(\eta + \dfrac{\alpha_b\, b(t)^2\, p(t)\, \lambda^*}{(b(t) - p(t)\lambda^*\, I^*)^2}\right) + \beta\, S^*\, \psi_S^{\mathrm{HA}} & \\ \qquad\qquad -(\beta\, S^* - \nu - \lambda^*)\, \psi_I^{\mathrm{HA}}, & \psi_I^{\mathrm{HA}}(T) = \mathrm{e}^{-\delta T}\theta, \\ \lambda^*(t) = \max\left\{0,\ \dfrac{b(t)}{p(t) I^*(t)}\left(1 - \sqrt{\dfrac{\alpha_b\, \mathrm{e}^{-\delta t}\, p(t)}{\psi_I^{\mathrm{HA}}(t)}}\right)\right\}, & \end{cases} \tag{17}$$

provided that $\psi_I^{\mathrm{HA}}(t) > 0$ for all $t \in [0, T]$. This condition ensures that the square root in the expression for $\lambda^*(t)$ is well defined, and that the optimal control remains finite and real throughout the time horizon. Moreover, the term $1 - \sqrt{\cdots}$ is always strictly less than 1, which guarantees that the resulting treatment expenditure $h^*(t) = p(t)\lambda^*(t)I^*(t)$ remains strictly below the available budget $b(t)$. The constraint is thus satisfied by construction. In the case where $\lambda^*(t) = 0$, the constraint is trivially satisfied.

The expression for the optimal control $\lambda^*(t)$ under blow-up costs reflects the trade-off between the marginal benefit of treatment and the associated blow-up cost. The structure of the optimal solution is divided into two regimes:

- *No treatment ($\lambda^*(t) = 0$)*: This occurs when $\mathrm{e}^{-\delta t}\, \alpha_b\, p(t) \geq \psi_I^{\mathrm{HA}}(t)$, that is, when the marginal benefit of reducing the infected population, represented by $\psi_I^{\mathrm{HA}}(t)$, is less than or equal to the discounted marginal cost of treatment, $\mathrm{e}^{-\delta t}\alpha_b\, p(t)$. In this case, treatment is not cost-effective, and the optimal policy is to abstain from intervention.
- *Interior solution ($\lambda^*(t) > 0$)*: This arises when the marginal benefit of treatment exceeds its discounted marginal cost, i.e., $\psi_I^{\mathrm{HA}}(t) > \mathrm{e}^{-\delta t}\alpha_b\, p(t)$. In this case, the optimal control takes a strictly positive value that balances the gain from reducing infections with the rising cost due to the proximity to the budget ceiling. The expression for $\lambda^*(t)$ ensures that the instantaneous expenditure remains strictly below $b(t)$, thereby avoiding the blow-up in the cost functional.

The optimal treatment strategy adapts dynamically to both economic conditions (via $p(t)$, $b(t)$) and the epidemiological state (via $I(t)$ and the costate $\psi_I^{\mathrm{HA}}(t)$), in order to balance public health objectives with financial and institutional constraints. These dynamics are clearly illustrated in the numerical simulations presented in Section 4, where the outcomes under quadratic and blow-up cost structures are compared in terms of treatment intensity, infection trajectories, and total cost.

On sufficiency and local uniqueness. The blow-up cost specification does not alter the conceptual interpretation of sufficiency discussed for the quadratic case. The main difference is the enforcement of the budget constraint. Although the mixed constraint is not imposed explicitly, the blow-up structure of the running cost prevents optimal trajectories from approaching the budget threshold. As a consequence, the same a priori bounds used in the quadratic case can be recovered. In particular, on sufficiently short time horizons the infected fraction remains uniformly bounded away from zero, ensuring that the pointwise representation of the optimal control is well defined and uniformly bounded. Uniqueness of the associated optimality system then follows by arguments analogous to those employed in the quadratic case, based on weighted energy estimates and a contraction property for the coupled state–adjoint variables.

3.4. Economic rationale for alternative cost structures

The adoption of quadratic versus blow-up cost structure reflects different institutional views on how health authorities manage pharmaceutical expenditures under budget constraints.

The quadratic formulation is suitable for centralized systems that operate through structured, anticipatory planning. In these settings, treatment policies are typically designed within fixed annual budgets, and increasing marginal costs capture logistical limitations, administrative burdens, or resource saturation. The explicit incorporation of a budget constraint into the control problem reflects a transparent and proactive approach to financial governance, consistent with systems where overspending is institutionally or politically infeasible.

By contrast, the blow-up specification models situations where financial limits are not imposed ex ante but emerge progressively as spending intensifies. This setting is representative of fragmented or reactive health systems, in which expenditures can initially grow unconstrained but become sharply penalized as they approach the available budget. The endogenous steepening of marginal costs captures institutional frictions, emergency reallocations, or procurement bottlenecks that arise under financial stress. In this context, the constraint is enforced indirectly by the cost function itself, which discourages excessive spending through rapidly escalating penalties.

These two approaches, though distinct in formulation and interpretation, serve complementary purposes. The quadratic model facilitates analysis under stable financial planning, while the blow-up model offers a more realistic representation of systems operating near or beyond capacity. Their joint inclusion within the decision framework allows for a flexible exploration of how treatment intensity should adapt under different cost architectures, ranging from predictable budget regimes to crisis-prone environments with endogenous constraints.

Table 1
Model parameters and units used in the numerical simulations.

Parameter	Description	Value	Unit
S_0	Initial susceptible population density	0.95	–
I_0	Initial infected population density	0.05	–
R_0	Initial recovered population density	0	–
ν	Birth/death rate	0.0125	1/year
β	Transmission rate	0.0015	1/year
δ	Discount rate	0.03	1/year
α_q	Scaling parameter for quadratic cost	0.0454	year/€
α_b	Scaling parameter for blow-up cost	1.135	year/€
η	Societal cost per infected individual	40,000	€/year

4. Simulated policy dynamics under budget constraints

In this section, we present a set of numerical simulations calibrated to the case of hepatitis C in France, using selected parameter values from Dubois and Magnac (2024) together with additional realistic and empirically grounded calibrations. The baseline parametrization is reported in Table 1 and serves as the reference scenario throughout the analysis.[2] The initial conditions assume that 5% of the population is infected at time zero, while $R_0 = 0$, reflecting the fact that no individuals have recovered prior to the introduction of the innovative treatment. The remaining 95% of the population is considered susceptible. This setup is consistent with an early-intervention scenario, in which curative therapies have not yet been widely administered and natural recovery is assumed to be negligible.

Using the baseline parameter values reported in Table 1, we can compute the critical treatment threshold defined in Eq. (4) as $\lambda_c = 0.025$. This value represents the minimum constant treatment intensity required to reduce the basic reproduction number $\mathcal{R}_0$ below 1 and ensure disease extinction in the long run. In the following simulations, this threshold will serve as a benchmark to assess whether the optimal treatment intensity $\lambda^*(t)$ determined by the HA under different cost specifications and budget constraints is sufficient to eradicate the infection or only to stabilize it at endemic levels. This means that, in order to eradicate the infection, the HA must maintain a treatment intensity of at least $\lambda^*(t) > 0.025$ for a sufficiently long period. If treatment levels fall below this threshold or are applied inconsistently over time, eradication is no longer guaranteed, and the infection may persist at endemic levels. This benchmark will guide the interpretation of the optimal treatment policies computed in the following.

In the following, numerical resolution of the optimal control problems is performed using a forward–backward algorithm where state and costate equations are integrated iteratively until convergence. The implementation relies on symplectic exponential Lawson-type integrators (Diele et al., 2011), which preserve structural properties of the Hamiltonian system and improve numerical stability. The overall procedure builds on the algorithmic framework developed in Martiradonna et al. (2020).

4.1. Comparative dynamics of treatment policies under budget constraints

In this subsection, we present numerical experiments illustrating how the optimal treatment strategy adapts to different budget levels. The aim is to show how cost structure, resource availability, and drug prices affect treatment intensity and epidemic dynamics. We focus on the two alternative formulations discussed in Section 3 for the cost of treatment: a quadratic cost structure with an explicit budget constraint, and a blow-up cost formulation that embeds the budget limitation directly into the cost function. In all simulations, the drug price is assumed to be constant and equal to $p = €18{,}000$. The epidemic and economic parameters are taken from the baseline case study and are summarized in Table 1. Moreover, we set the terminal penalty parameter to $\theta = 0$, so that the HA does not explicitly target the final infection level.

Fig. 1 illustrates the time profiles of the optimal treatment intensity $\lambda(t)$, the epidemiological outcomes $I(t)$ and $R(t)$, and the instantaneous expenditure $h(t)$, under both quadratic and blow-up cost structures and for different levels of the budget constraint b. As the budget increases from $b = €100$ (top row) to $b = €200$ (middle row) and $b = €300$ (bottom row), treatment becomes more intensive. For lower budgets, both cost specifications induce moderate treatment levels, and the budget constraint is not fully saturated. When the budget becomes more generous, the quadratic cost formulation promotes a more aggressive and front-loaded treatment strategy, with sharper initial increases in $\lambda(t)$, faster reductions in $I(t)$, and rapid saturation of the budget constraint. In contrast, the blow-up specification leads to a smoother and more gradual ramp-up, discouraging abrupt spending peaks through its endogenous penalization mechanism. As a result, expenditures under the blow-up cost remain below the ceiling throughout the horizon, avoiding early saturation and enforcing a softer allocation profile. Across all scenarios, treatment intensities remain well above the critical threshold λ_c for most of the horizon, leading to a progressive reduction in infection prevalence. The differences lie mainly in the intensity and duration of the treatment effort: higher budgets allow treatment levels to remain elevated for longer periods, accelerating the decline of infections and moving the system closer to eradication. Comparing the effectiveness of the two cost structures,

[2] In all the numerical exercises, the time horizon is set to 5 years. This reflects two motivations: analytically, we focus on short- to medium-term effects of HA decisions on epidemic dynamics; institutionally, T represents the contract duration between the HA and the pharmaceutical company, with 10 years being a realistic upper bound.

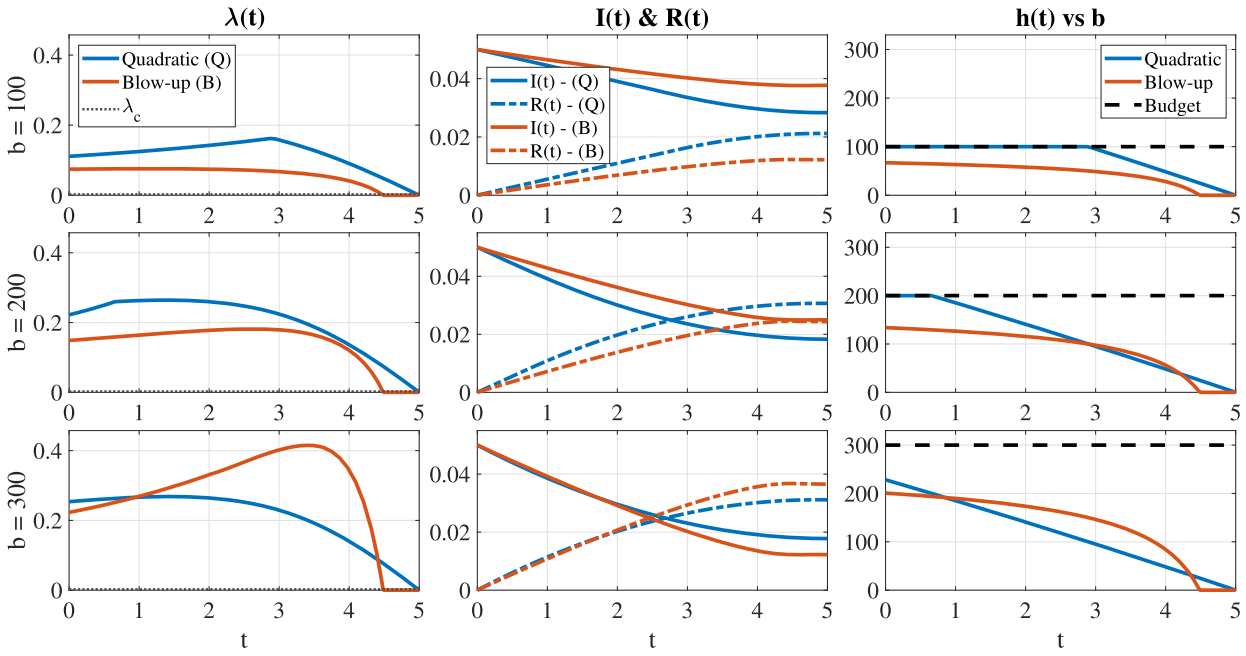

Fig. 1. Dynamics of the optimal treatment policy and epidemic evolution under quadratic (Q) and blow-up (B) cost formulations, for different budget levels b. Left column: optimal treatment intensity $\lambda(t)$ over time, with the eradication threshold λ_c indicated by a dotted line. Center column: trajectories of infected $I(t)$ (solid lines) and recovered $R(t)$ (dashed lines). Right column: instantaneous treatment expenditure $h(t)$ compared to the budget ceiling b (dashed black line). Rows correspond to increasing budget levels: b = €100 (top), b = €200 (middle), and b = €300 (bottom).

Table 2
Comparison of key epidemiological and economic outcomes under quadratic and blow-up cost specifications for different budget levels over a 5-year horizon.

Per capita budget	Index	Quadratic	Blow-up
€100	Final infected share	2.83%	3.77%
	Aggregate recovered share	43.97%	25.30%
	Aggregate per capita expenditure	7720 €	6450 €
€200	Final infected share	1.83%	2.50%
	Aggregate recovered share	63.85%	50.61%
	Aggregate per capita expenditure	7360 €	3570 €
€300	Final infected share	1.78%	1.22%
	Aggregate recovered share	64.89%	75.91%
	Aggregate per capita expenditure	7360 €	695 €

Fig. 1 and the outcomes provided in Table 2 allow us to observe how the two specifications lead to different trade-offs between cost efficiency and epidemic control. For low budget levels (e.g., €100), the quadratic formulation proves more effective in curbing the epidemic, achieving a significantly higher recovery rate and a lower final prevalence of infection, albeit at the expense of greater per capita spending. In contrast, the blow-up model, which strongly penalizes treatment intensities approaching the budget ceiling, results in more conservative treatment and worse health outcomes. As the available budget increases, both cost structures allow the HA to intensify treatment, improving outcomes. However, an important divergence emerges. While the quadratic model maintains high expenditure even at higher budget levels—due to the smooth nature of the cost function—the blow-up model shows a marked reduction in spending once the budget is sufficiently large. This is because, in the blow-up formulation, treatment becomes increasingly cost-efficient as long as it remains far from triggering the singularity at $h(t) = b$. Strikingly, at €300 per capita, the blow-up model achieves better epidemiological outcomes than the quadratic one (e.g., lower final infection share and higher cumulative recovery), while incurring dramatically lower costs. These findings indicate that, although the quadratic cost function supports stronger early intervention when resources are limited, the blow-up specification becomes more advantageous when budgetary conditions allow for sustained and moderate treatment without incurring excessive marginal costs.

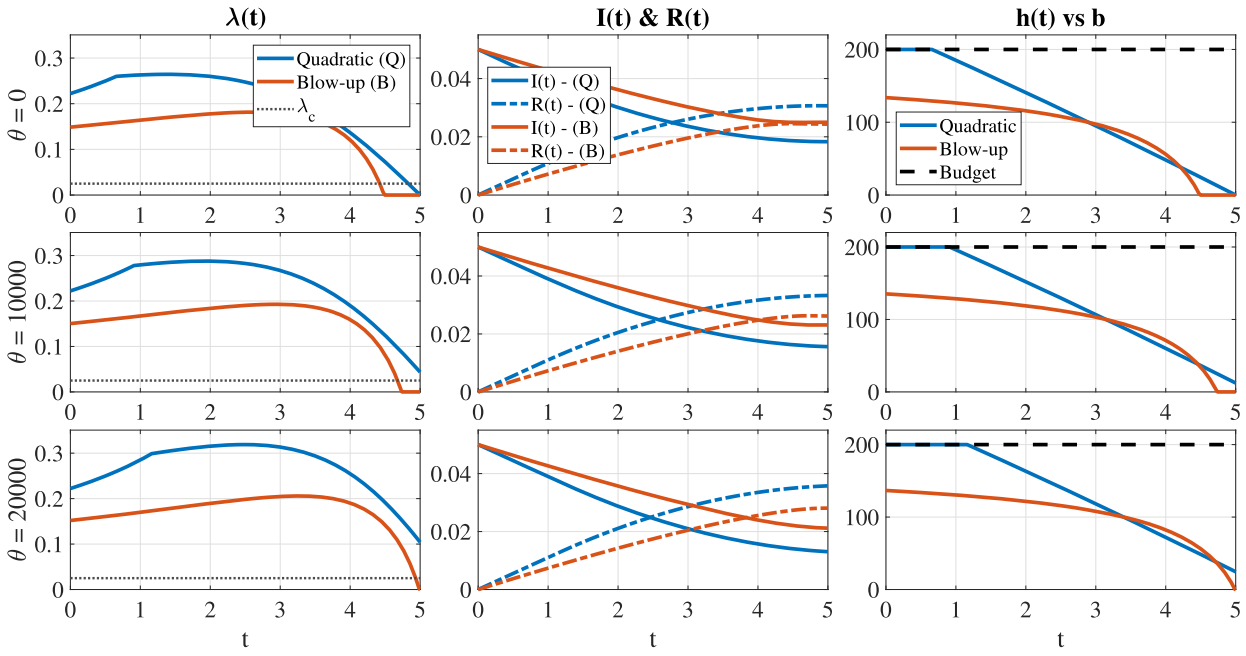

Fig. 2. Dynamics of the optimal treatment policy and epidemic evolution under quadratic (Q) and blow-up (B) cost formulations, for different terminal penalty levels θ, with fixed budget $b = €200$. Rows correspond to increasing θ: $\theta = 0$ (top), $\theta = 10{,}000$ (middle), and $\theta = 20{,}000$ (bottom). Left column: treatment intensity $\lambda(t)$; center column: infection $I(t)$ and recovered $R(t)$; right column: per capita expenditure $h(t)$ versus budget ceiling.

4.2. Policy preferences and economic trade-offs driven by θ

In this subsection, we investigate how the terminal penalty parameter θ affects the optimal treatment dynamics and epidemic outcomes. As shown in Section 3, θ weighs the residual number of infected individuals at the end of the planning horizon, effectively encouraging the HA to prioritize long-term eradication when set to higher values. To isolate the effects of θ, we fix the per capita budget at $b = €200$ and maintain all other parameters at their baseline values (see Table 1). We explore three representative scenarios: $\theta = 0$ (no terminal penalty), $\theta =$€ 10,000, and $\theta =$ € 20,000.

Fig. 2 displays the dynamics of the optimal treatment intensity $\lambda(t)$, the infection share $I(t)$, the recovered share $R(t)$, and the instantaneous per capita expenditure $h(t)$, under different values of θ over a 5-year horizon. As θ increases from 0 to 20,000 (moving from top to bottom rows), we observe that the HA becomes progressively more aggressive in its treatment strategy, especially toward the end of the time horizon. In the scenario with $\theta = 0$, the focus is entirely on minimizing cumulative infections and costs during the planning period, resulting in a moderate decline of $I(t)$ and a more gradual treatment profile. In this case, $\lambda(t)$ drops to zero toward the final period, indicating no concern for the future state of the epidemic beyond the time horizon. When θ is increased to 10,000 and 20,000, the authority explicitly prioritizes reducing the final infected share $I(T)$. This leads to sustained or even intensified treatment levels in the later phases, as visible in the higher and more persistent $\lambda(t)$. Crucially, for large θ, $\lambda(t)$ remains consistently above the critical threshold λ_c until the end of the horizon. This sustained effort supports the convergence toward a disease-free state and prevents resurgence after the planning period. In terms of expenditures $h(t)$, we see that higher θ values lead to a slower decline in spending, reflecting the policy's effort to achieve near-eradication by the end of the horizon. Despite the fixed budget constraint, the quadratic formulation tends to saturate the budget earlier and more sharply than the blow-up cost, which enforces a smoother expenditure path.

These findings highlight the importance of the terminal penalty parameter θ in shaping not only immediate epidemic control but also long-term eradication goals. By assigning sufficient weight to $I(T)$, policymakers can enforce a persistent treatment effort that ensures the system remains above λ_c and supports disease elimination, rather than merely stabilizing the infection at low endemic levels.

5. An extension to the strategic pricing of innovative treatment: a Stackelberg differential game approach

Building upon the baseline model developed in the previous sections, where the price $p(t)$ of the innovative treatment was assumed to be exogenously given, we now extend the framework to a more realistic and strategic setting. In this extended formulation, the price is endogenously determined by the Pharmaceutical Company (PC), which optimizes its pricing strategy to maximize profits while anticipating the optimal treatment response of the Health Authority (HA). This interaction is modeled as a Stackelberg differential game over the finite horizon $[0, T]$, where the PC acts as the leader and the HA as the follower.

The PC chooses the price trajectory $p(t) > 0$ over the time horizon $[0, T]$ to maximize the total discounted profit:

$$\mathcal{J}^{\mathrm{PC}}(p(\cdot)) = \mathrm{e}^{-\delta T}\,\kappa\, I(T) + \int_0^T \mathrm{e}^{-\delta t}(p(t) - c)\,\lambda(t; p(t))\, I(t)\, dt, \tag{18}$$

where $c > 0$ denotes the constant unit cost of drug production, and $\lambda(t; p(t))$ is the treatment intensity optimally chosen by the HA in response to the price path $p(t)$. That is, $\lambda(t; p) = \lambda^*(t)$, as characterized in system (13) for the quadratic cost case, or in system (17) for the blow-up case.

The first term in $\mathcal{J}^{\mathrm{PC}}$, $\mathrm{e}^{-\delta T}\kappa I(T)$, plays the role of a terminal (or *scrap*) value and captures the strategic relevance, for the PC, of the epidemic status at the end of the planning horizon. The economic meaning of this term depends on the sign of the parameter $\kappa \in \mathbb{R}$. When $\kappa > 0$, the firm benefits from persistent disease prevalence, for instance because it sustains future demand, delays competitive entry, or enhances bargaining power in subsequent negotiations. Conversely, when $\kappa < 0$, a high terminal infection level represents a reputational or strategic cost, reflecting regulatory pressure, public scrutiny, or the risk of exclusion from future procurement and reimbursement processes. The case $\kappa = 0$ corresponds to a firm that is indifferent to the terminal epidemic outcome and focuses solely on contemporaneous profit flows. Beyond its formal role in the objective function, the terminal term $\mathrm{e}^{-\delta T}\kappa I(T)$ captures the intertemporal link between current pricing decisions and future strategic payoffs. In this sense, κ provides a reduced-form representation of the long-run policy and regulatory environment faced by the firm, summarizing how forward-looking incentives shape dynamic pricing behavior beyond static profit maximization.

The control variable of the PC is the price trajectory $p(t)$, which is constrained from below by a minimum price equal to the production cost c. Accordingly, the admissible control set is defined as:

$$\mathcal{P} = \left\{ p \in L^{\infty}([0, T]) \;\middle|\; p(t) \geq c > 0 \quad \text{a.e. in } [0, T] \right\}.$$

The PC aims to maximize its total discounted profit by selecting a time-dependent pricing strategy $p(t)$, while anticipating the treatment response $\lambda(t; p(t)) = \lambda^*(t)$ of the HA. This leads to a hierarchical (Stackelberg) control problem in which the leader's decision influences the optimal behavior of the follower. The resulting optimization is subject to the epidemiological dynamics of the system and to the HA's optimality conditions, which depend on the assumed cost structure. Specifically, the HA's best-response $\lambda^*(t)$ is implicitly defined by the solution to either the quadratic-cost system (13) or the blow-up-cost system (17). In both cases, the treatment intensity is a nontrivial function of the price trajectory and state variables, reflecting the trade-off between epidemiological control and economic feasibility.

If $\lambda^*(t) = 0$, the PC earns no profit at time t, regardless of the price level as $(p(t) - c)\,\lambda^*(t)\, I(t) = 0$. From the PC's perspective, such intervals are strategically undesirable, as they suppress demand entirely and nullify revenue generation. For this reason, we assume throughout the analysis that the HA's optimal policy satisfies the interiority condition:

$$\lambda^*(t) > 0 \quad \text{for all } t \in [0, T],$$

which ensures that the price always induces a strictly positive treatment response and that the profit functional remains active over the entire horizon.

5.1. Pricing behavior under HA quadratic response

We now turn to the analysis of the PC's optimization problem, assuming that the HA operates under a quadratic cost specification. To ensure smooth dependence of the treatment intensity $\lambda^*(t)$ on the price trajectory $p(t)$, we focus on the *interior regime*, where the HA's budget constraint remains inactive throughout the planning horizon. Specifically, we assume:

$$\lambda^*(t) > 0 \quad \text{and} \quad p(t)\lambda^*(t)I(t) < b(t) \quad \text{for all } t \in [0, T],$$

with the associated Lagrange multipliers $\mu_1(t)$ and $\mu_2(t)$ equal to zero, ensuring that the HA's budget constraint remains non-binding throughout the planning horizon. These assumptions eliminate potential nonsmooth behavior (kink points) in the treatment response and guarantee that the HA's optimal control remains continuously differentiable in both time and price. Based on the first-order condition of the HA's problem, the optimal treatment intensity is given by:

$$\lambda(t; p(t)) = \frac{\mathrm{e}^{\delta t}}{\alpha_q\, p(t)^2\, I(t)}\, \psi_I^{\mathrm{HA}}(t) > 0 \quad \text{for all } t \in [0, T],$$

which guarantees that the interior regime is active at all times. Moreover, this expression defines a continuous and monotonic relationship between the price and the treatment intensity: higher prices reduce the affordability of treatment and thereby depress the HA's optimal response.

The PC, acting as Stackelberg leader, seeks to maximize its profit functional:

$$\mathcal{J}_q^{\mathrm{PC}}(p(\cdot)) = \mathrm{e}^{-\delta T}\,\kappa\, I(T) + \int_0^T \frac{p(t) - c}{\alpha_q\,(p(t))^2}\, \psi_I^{\mathrm{HA}}(t)\, dt,$$

subject to the HA's forward–backward system:

$$\begin{cases} \dot{S} = \nu - \nu S - \beta S I, & S(0) = S_0 \\ \dot{I}(t) = \beta S I - \nu I - \dfrac{\mathrm{e}^{\delta t}}{\alpha_q\, p^2}\, \psi_I^{\mathrm{HA}}, & I(0) = I_0 \\ \dot{\psi}_S^{\mathrm{HA}} = (\nu + \beta\, I)\, \psi_S^{\mathrm{HA}} - \beta\, I\, \psi_I^{\mathrm{HA}}, & \psi_S^{\mathrm{HA}}(T) = 0 \\ \dot{\psi}_I^{\mathrm{HA}} = -\mathrm{e}^{-\delta t}\eta + \beta\, S\, \psi_S^{\mathrm{HA}} - (\beta\, S - \nu)\, \psi_I^{\mathrm{HA}}, & \psi_I^{\mathrm{HA}}(T) = \mathrm{e}^{-\delta T}\theta \end{cases} \tag{19}$$

Under the given assumptions, the admissible price set $\mathcal{P}$ is closed and bounded from below, and all variables remain finite over the planning horizon. Although the absence of an upper bound means that the set is not compact, the structure of the problem and the fact that the profit term becomes negligible as the price grows sufficiently large help ensure that the leader's optimization problem is well-defined. These features support the existence of an optimal pricing strategy according to standard principles in dynamic optimization.

We adopt an open-loop Stackelberg formulation of the leader's problem, assuming full precommitment to the entire price trajectory over the planning horizon. In this setting, the PC commits at the initial time $t = 0$ to a complete price trajectory $p(t)$ over $[0, T]$, taking into account the dynamic response of the HA.

To derive the first-order necessary conditions for optimality, we introduce the Hamiltonian associated with the PC's problem. The state variables are $S(t), I(t), \psi_S^{\mathrm{HA}}(t), \psi_I^{\mathrm{HA}}(t)$, while the corresponding costate variables introduced by the PC are $\psi_S^{\mathrm{PC}}(t), \psi_I^{\mathrm{PC}}(t), \chi_S^{\mathrm{PC}}(t), \chi_I^{\mathrm{PC}}(t)$. The Hamiltonian function $\mathcal{H}^{\mathrm{PC}}$ captures both the instantaneous profit of the PC and the indirect effects of pricing on the state and costate variables through the HA's dynamics:

$$\begin{aligned} \mathcal{H}^{\mathrm{PC}} = & \frac{p - c}{\alpha_q\, p^2}\, \psi_I^{\mathrm{HA}} + (\nu - \nu S - \beta S I)\, \psi_S^{\mathrm{PC}} + \left(\beta S I - \nu I - \frac{\mathrm{e}^{\delta t}}{\alpha_q\, p^2}\, \psi_I^{\mathrm{HA}} \right) \psi_I^{\mathrm{PC}} \\ & + \left((\nu + \beta I)\, \psi_S^{\mathrm{HA}} - \beta I\, \psi_I^{\mathrm{HA}} \right) \chi_S^{\mathrm{PC}} + \left(-\mathrm{e}^{-\delta t}\eta + \beta\, S\, \psi_S^{\mathrm{HA}} - (\beta\, S - \nu)\, \psi_I^{\mathrm{HA}} \right) \chi_I^{\mathrm{PC}}, \end{aligned} \tag{20}$$

The derivatives of the Hamiltonian with respect to p are given by:

$$\frac{\partial \mathcal{H}^{\mathrm{PC}}}{\partial p} = \frac{1}{\alpha_q} \psi_I^{\mathrm{HA}} \frac{1}{p^3} \left(2c + 2\mathrm{e}^{\delta t} \psi_I^{\mathrm{PC}} - p \right), \quad \frac{\partial^2 \mathcal{H}^{\mathrm{PC}}}{\partial p^2} = \frac{2}{\alpha_q} \psi_I^{\mathrm{HA}} \frac{1}{p^4} \left(p - 3c - 3\mathrm{e}^{\delta t} \psi_I^{\mathrm{PC}} \right).$$

These expressions indicate that the Hamiltonian is increasing in p when $p \le 2(c + \mathrm{e}^{\delta t}\psi_I^{\mathrm{PC}})$ and decreasing otherwise. Moreover, it is concave for $p \le 3(c + \mathrm{e}^{\delta t}\psi_I^{\mathrm{PC}})$. Consequently, the candidate optimum price is the maximum between c and $2(c + \mathrm{e}^{\delta t}\psi_I^{\mathrm{PC}})$.

The full set of necessary conditions for the PC's optimal pricing problem consists of the following forward–backward system:

$$\begin{cases} \dot{S}(t) = \nu - \nu S(t) - \beta S(t) I(t), & S(0) = S_0, \\ \dot{I}(t) = \beta S(t) I(t) - \nu I(t) - \dfrac{\mathrm{e}^{\delta t}}{\alpha_q\, (p^S(t))^2}\, \psi_I^{\mathrm{HA}}(t), & I(0) = I_0, \\ \dot{\psi}_S^{\mathrm{HA}}(t) = (\nu + \beta I(t))\, \psi_S^{\mathrm{HA}}(t) - \beta I(t)\, \psi_I^{\mathrm{HA}}(t), & \psi_S^{\mathrm{HA}}(T) = 0, \\ \dot{\psi}_I^{\mathrm{HA}}(t) = -\mathrm{e}^{-\delta t}\eta + \beta S(t)\, \psi_S^{\mathrm{HA}}(t) - (\beta S(t) - \nu)\, \psi_I^{\mathrm{HA}}(t), & \psi_I^{\mathrm{HA}}(T) = \mathrm{e}^{-\delta T}\theta, \\ \dot{\psi}_S^{\mathrm{PC}}(t) = \nu\, \psi_S^{\mathrm{PC}}(t) + \beta I(t) \left(\psi_S^{\mathrm{PC}}(t) - \psi_I^{\mathrm{PC}}(t) \right) - \beta I(t) \chi_S^{\mathrm{PC}}(t) + \beta I(t) \chi_I^{\mathrm{PC}}(t), & \psi_S^{\mathrm{PC}}(T) = 0, \\ \dot{\psi}_I^{\mathrm{PC}}(t) = \nu\, \psi_I^{\mathrm{PC}}(t) - \beta S(t) \left(\psi_S^{\mathrm{PC}}(t) - \psi_I^{\mathrm{PC}}(t) \right) - \beta S(t) \chi_S^{\mathrm{PC}}(t) + (\beta S(t) - \nu)\, \chi_I^{\mathrm{PC}}(t), & \psi_I^{\mathrm{PC}}(T) = \mathrm{e}^{-\delta T}\kappa, \\ \dot{\chi}_S^{\mathrm{PC}}(t) = -\beta\, \psi_I^{\mathrm{HA}}(t)\, \psi_I^{\mathrm{PC}}(t) - \beta I(t)\, \chi_I^{\mathrm{PC}}(t), & \chi_S^{\mathrm{PC}}(0) = 0, \\ \dot{\chi}_I^{\mathrm{PC}}(t) = \dfrac{p^S(t) - c}{\alpha_q\, (p^S(t))^2} - \dfrac{\mathrm{e}^{\delta t}}{\alpha_q\, (p^S(t))^2}\, \psi_I^{\mathrm{PC}}(t) + \beta S(t) \chi_I^{\mathrm{PC}}(t) - (\nu + \beta S(t))\, \chi_S^{\mathrm{PC}}(t), & \chi_I^{\mathrm{PC}}(0) = 0, \\ p^S(t) = \max \left\{ c,\ 2\left(c + \mathrm{e}^{\delta t} \psi_I^{\mathrm{PC}}(t) \right) \right\}. \end{cases} \tag{21}$$

The optimal pricing strategy derived in this Stackelberg framework reflects a balance between production costs and the marginal value of treatment, as captured by the adjoint variable ψ_I^{PC}. The optimal price $p^S(t)$ increases with the perceived economic benefit of reducing infections and the discounting of future profits.

Correspondingly, the follower's (HA) optimal treatment response under the quadratic cost structure is given by:

$$\lambda^S(t) = \frac{\mathrm{e}^{\delta t}}{\alpha_q \left(p^S(t) \right)^2 I(t)}\, \psi_I^{\mathrm{HA}}(t).$$

An important analytical insight emerges when we evaluate the optimal pricing strategy at the terminal time T. The optimal price function $p^S(t)$ satisfies:

$$p^S(T) = \max \{ c,\ 2(c + \kappa) \}.$$

This formula highlights the direct dependence of the terminal price on the strategic weight κ assigned by the PC to the final infection level $I(T)$. The interpretation of this result is as follows:

- If $\kappa = 0$, the PC is neutral with respect to the terminal infection level, and the optimal terminal price simplifies to $p^S(T) = 2c$. This provides a benchmark reference level, driven solely by production cost considerations.

- If $\kappa > 0$, the PC has an incentive to maintain some residual infection in the population (e.g., to sustain future demand or delay market competition). As a result, $p^S(T)$ increases linearly with κ, making treatment more expensive and reducing the HA's ability to eradicate the disease.
- If $\kappa < 0$, the firm incurs a reputational or strategic penalty from persistent infection, such as reduced approval in future negotiations or pressure from public institutions. In this case, the optimal price may drop significantly, and if κ is sufficiently negative to satisfy $2(c + \kappa) \leq c$, then the lower bound becomes active and the price is set at the minimum feasible level $p^S(T) = c$.

This structure shows that the terminal pricing decision incorporates not only economic fundamentals (such as production cost), but also strategic positioning over long-term health outcomes. It also provides a practical rationale for price drops observed in real-world pharmaceutical negotiations, especially in the context of large-scale eradication campaigns or public-private partnerships. The trajectory of the pricing function $p^S(t)$ over time reflects a complex interplay between infection dynamics, the adjoint variables of both players, and intertemporal trade-offs shaped by discounting. Prices may decrease as the epidemic subsides, remain high when eradication incentives are weak, or even rise near the horizon when persistent infection retains strategic value. In some cases, revenues may be front-loaded to exploit early demand.

On existence, sufficiency and uniqueness in the Stackelberg setting. Under the interiority assumptions and on sufficiently short time horizons, the follower's optimal response is well defined and depends continuously on the leader's pricing strategy. This regularity implies local well-posedness of the coupled forward–backward system (21) and guarantees the existence of at least one solution.

In the same regime, uniqueness of the solution of the optimality system follows by arguments analogous to those employed for the HA problem, based on weighted energy estimates and a contraction property for the coupled state–adjoint variables. Accordingly, the first-order optimality conditions provide a locally sufficient characterization of the Stackelberg equilibrium. In this local regime, multiplicity of equilibria and unstable behavior are ruled out. Possible global bifurcations, loss of stability, or coexistence of multiple equilibria may arise outside this regime, depending on parameter values and on the length of the planning horizon, but their systematic analysis is beyond the scope of the present paper.

Remark on precommitment, institutional pricing agreements, and time consistency. The Stackelberg pricing problem analyzed in this section is formulated under an open-loop precommitment assumption. In this setting, the HA and the PC commit at the initial time $t = 0$ to their respective strategies over the planning horizon. Specifically, the PC commits to a complete price trajectory, while the HA commits to the optimal treatment policy induced by its control problem in response to the announced price. The open-loop formulation adopted here is consistent with the institutional design of Managed Entry Agreements (MEAs), which are contractual arrangements between manufacturers and payers or Health Authorities that define pricing and reimbursement conditions in advance in order to manage clinical and financial uncertainty associated with innovative therapies (Wenzl and Chapman, 2019; Ciulla et al., 2023). Under these agreements, price-related conditions such as price–volume rules, budget caps, or outcome-based clauses are fixed at market entry and remain binding over the duration of the agreement.

In the absence of a precommitment assumption, the leader's problem may suffer from time inconsistency, a well known issue in open-loop Stackelberg differential games (Dockner, 2000). Intuitively, if the PC were allowed to revise its pricing strategy at an intermediate date, the continuation of the initially announced price path would generally no longer be optimal. As a result, the pricing policy obtained from a re optimization would differ from the original plan, undermining the credibility of ex ante commitment.

In this paper, we focus on the benchmark case of full precommitment, which allows us to clearly characterize the strategic interaction between the HA and the PC and to highlight the role of institutional pricing agreements. Relaxing the precommitment assumption is expected to affect the credibility and the intertemporal structure of the announced price trajectory, with potential implications for treatment incentives and epidemic dynamics. Investigating these effects within a feedback or renegotiation framework is left for future research.

5.2. Simulation-based insights on costs and incentives

We present numerical simulations illustrating the dynamic interplay between price, treatment intensity, and epidemic outcomes, adopting the baseline epidemiological and economic parameters introduced in Section 4, and varying key variables such as production cost c, budget levels, and the terminal infection valuation parameter κ.

5.2.1. Cost-driven pricing and treatment dynamics

In this analysis, we fix both terminal weights to zero, i.e., $\kappa = 0$ and $\theta = 0$, so that neither player assigns value to the final state of the epidemic at time $T = 5$ years. The strategic interaction is therefore governed solely by the stream of intermediate payoffs and costs. Fig. 3 illustrates the dynamic response of the system for three representative values of the unit production cost: €8,000, €12,000, and €16,000. The figure reports the time paths of key outcome variables, including the optimal price $p(t)$ set by the PC, the treatment intensity $\lambda(t)$ chosen by the HA, and the resulting epidemiological dynamics.

The first panel of Fig. 3 shows that the drug price $p(t)$ remains approximately constant over time for each scenario and grows linearly with the production cost c. This behavior reflects the Stackelberg pricing rule, under which the PC sets the price as a markup over marginal cost, stabilizing around $p(t) \approx 2c$. In the absence of price regulation or terminal penalties ($\kappa = 0$), this direct proportionality holds robustly, confirming that higher production costs translate immediately into higher prices. As c increases, the higher price induces a pronounced reduction in the treatment intensity $\lambda(t)$ chosen by the HA, as shown in the second panel. When

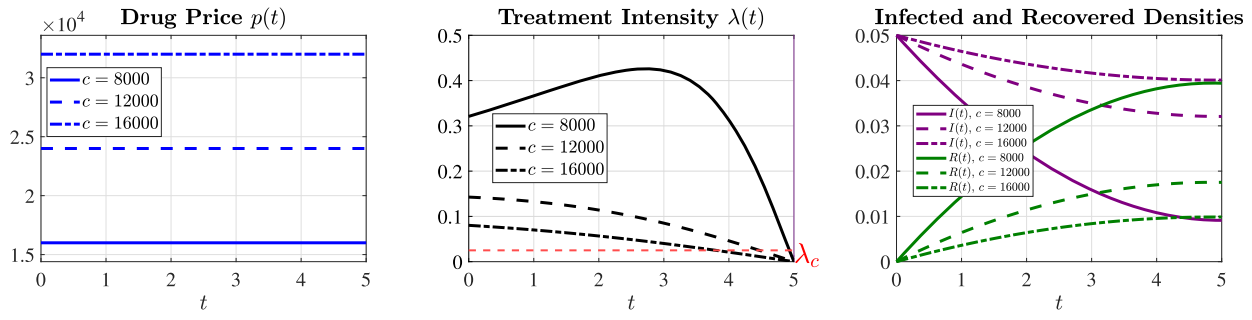

Fig. 3. Dynamic behavior of the Stackelberg system under three values of the unit production cost, $c = 8{,}000$, $12{,}000$, and $16{,}000$, with $\kappa = \theta = 0$, over a time horizon of $T = 5$ years. The panels show: (left) drug price $p(t)$, (center) treatment intensity $\lambda(t)$, and (right) infected and recovered densities $I(t)$ and $R(t)$.

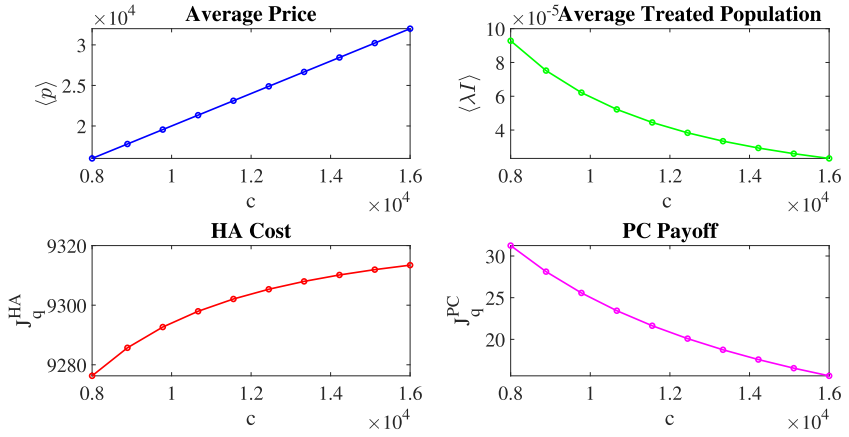

Fig. 4. Sensitivity analysis of Stackelberg outcomes with respect to the unit production cost $c \in [8000, 16000]$, with $\kappa = \theta = 0$. The four panels display: (top left) average drug price $\langle p \rangle$, (top right) average treated population density $\langle \lambda I \rangle$, (bottom left) HA's per capita total cost $\mathcal{J}_q^{\mathrm{HA}}$, and (bottom right) PC's per capita payoff $\mathcal{J}_q^{\mathrm{PC}}$, computed over a simulation horizon of $T = 5$.

$c = €8{,}000$, treatment begins at a relatively high level (around 0.3), peaks above 0.4, and then gradually declines over time. For $c = €12{,}000$, the initial intensity is substantially lower, and the peak is significantly dampened. At $c = €16{,}000$, treatment remains markedly suppressed, staying below 0.15 for most of the horizon. In all scenarios, $\lambda(t)$ initially exceeds the eradication threshold λ_c, enabling early strong suppression of infections. However, over time, $\lambda(t)$ approaches or even drops below λ_c, reflecting a loss of momentum in sustaining aggressive intervention without explicit incentives for eradication at the final horizon. These dynamics in price and treatment intensity jointly shape the epidemiological outcomes, depicted in the third panel. When treatment is intense (low c), infections decline rapidly, and a large share of the population transitions to recovery. As c rises, the diminished treatment effort results in slower infection decline and lower cumulative recovery. In the high-cost scenario ($c = €16{,}000$), infections persist at higher levels, and the recovered fraction remains limited, illustrating the public health risks associated with limited affordability and low coverage.

A transversal reading across the three panels reveals a cascading chain of effects driven by the production cost c: higher c leads to higher $p(t)$, which discourages treatment ($\lambda(t)$), ultimately resulting in more persistent infections and lower recovery. This sequence underscores how economic fundamentals directly shape epidemiological trajectories in a strategic environment. The analysis highlights the importance of integrating cost containment, pricing policy, and incentive design to maintain treatment coverage and prevent unfavorable epidemic scenarios.

To synthesize the dynamic results and facilitate comparison across a broader range of scenarios, Fig. 4 presents an aggregate sensitivity analysis. Each panel reports a key performance indicator, either a cumulative payoff or a time-averaged quantity, computed over the full simulation horizon $T = 5$ years, as a function of the unit production cost $c \in [8{,}000, 16{,}000]$. Time-averaged quantities are denoted using angle brackets $\langle \cdot \rangle$, and are defined as

$$\langle f \rangle = \frac{1}{T} \int_0^T f(t)\, dt. \tag{22}$$

In the top-left panel, we observe that the average drug price $\langle p \rangle$ increases linearly with c, reflecting the optimal markup strategy adopted by the PC. This confirms that the PC fully internalizes production cost in its pricing decision, leading to a proportional price adjustment even without external constraints. The top-right panel shows that the average treated population $\langle \lambda I \rangle$ decreases steeply

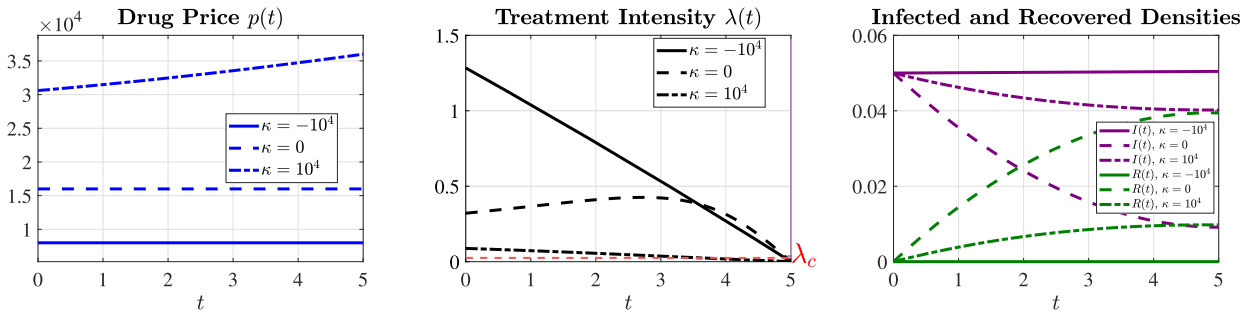

Fig. 5. Dynamic behavior of the system under three values of the terminal weight parameter: €–10,000, €0, and €10,000, with c =€8,000, and $\theta = 0$, over a time horizon of $T = 5$ years. The panels show: (left) drug price $p(t)$, (center) treatment intensity $\lambda(t)$, and (right) infected and recovered population shares $I(t)$ and $R(t)$.

as c increases. This outcome arises from the HA's optimal response: higher prices imply greater economic burden, thus reducing treatment intensity. The nonlinearity in this relationship is shaped by the quadratic cost structure, which heavily penalizes intense interventions when prices are high. The bottom-left panel illustrates the per capita total cost incurred by the HA, $\mathcal{J}_q^{\mathrm{HA}}$. Interestingly, this cost increases with c, but at a decreasing rate. Although higher c leads to higher prices, it also induces lower treatment intensity, partially offsetting the total expenditure. However, when treatment is too limited, persistent infections generate additional societal costs, maintaining upward pressure on $\mathcal{J}_q^{\mathrm{HA}}$. Finally, the bottom-right panel reports the PC's per capita payoff $\mathcal{J}_q^{\mathrm{PC}}$, which declines sharply with c. While higher prices could theoretically increase revenues, they simultaneously reduce demand through the HA's endogenous response. Consequently, the volume of treated individuals shrinks faster than the price rises, leading to overall lower profits. This highlights a form of endogenous demand erosion: excessive pricing undermines intervention coverage and ultimately harms the PC's own payoff.

The figure clearly illustrates the delicate balance between pricing, affordability, and epidemic control in a strategic context. The unit production cost emerges as a key factor that shapes not only market dynamics but also the ability of public health systems to respond effectively. These findings emphasize the need for integrated policy strategies that combine cost containment, fair pricing negotiations, and incentive alignment to support large-scale treatment programs and achieve lasting public health goals.

5.2.2. *The role of long-term incentives*

We now investigate the role of the terminal value parameter κ, which captures the strategic importance assigned by the PC to the final infection level $I(T)$. Positive values of κ model scenarios in which the firm benefits from residual disease prevalence at the end of the horizon, for instance, due to sustained demand, delayed competition, or increased leverage. Conversely, negative values of κ represent reputational or regulatory penalties associated with high long-term infection levels. In this analysis, we fix the unit production cost at c =€8,000 and set the terminal weight of the HA to zero, i.e., $\theta = 0$, so that only the PC is forward-looking.

Fig. 5 illustrates how the dynamic Stackelberg equilibrium is affected by variations in κ. The left panel shows that the drug price $p(t)$ is strongly influenced by the PC's valuation of the epidemic end state. When $\kappa > 0$, the firm strategically prefers maintaining a higher infection level and thus increases the price over time: we observe a clear upward-sloping trajectory. In contrast, when $\kappa < 0$, the firm is penalized for high residual infection and reduces the price dynamically to stimulate treatment uptake, resulting in a decreasing price path. For $\kappa = 0$, the price remains approximately constant around $2c$, reflecting a purely static profit motive without terminal concerns. The center panel shows the corresponding treatment intensity $\lambda(t)$ chosen by the HA. In the $\kappa < 0$ scenario, lower prices incentivize more aggressive treatment, with $\lambda(t)$ starting above 1 and gradually decreasing over time. When $\kappa > 0$, the higher price discourages intervention, resulting in lower initial and overall treatment levels. In all scenarios, the treatment intensity initially exceeds the eradication threshold λ_c, but eventually declines and may fall below it depending on the price trajectory and the PC's strategic incentives. The right panel displays the evolution of the infected and recovered population shares. Under $\kappa < 0$, aggressive treatment leads to a rapid decline in infections and a higher final share of recovered individuals. Conversely, when $\kappa > 0$, infections persist at higher levels, and the recovered fraction remains lower, highlighting the impact of pricing strategy on long-term epidemic control. This analysis confirms that the PC's valuation of the terminal epidemic state, encoded in κ, critically shapes pricing behavior, which in turn alters the HA's response and the resulting epidemiological outcomes. Even in the absence of terminal incentives for the HA ($\theta = 0$), forward-looking strategies by the PC are sufficient to significantly shift the equilibrium dynamics. This underscores the importance of incorporating long-term incentives and potential regulatory penalties into policy design to better align private pricing decisions with public health objectives.

Fig. 6 shows how the average price $\langle p \rangle$, the average treated population $\langle \lambda I \rangle$, the HA's per capita total cost $\mathcal{J}_q^{\mathrm{HA}}$, and the PC's per capita payoff $\mathcal{J}_q^{\mathrm{PC}}$ vary with the terminal weight parameter κ, ranging from –10,000 to 10,000 euros. The top-left panel shows that the average drug price increases monotonically with κ, confirming that the PC internalizes the terminal valuation of residual infection: higher values of κ incentivize the firm to raise prices, anticipating long-term benefits from sustained prevalence. As prices rise, the HA reduces treatment intensity in response to affordability constraints, leading to a marked decline in $\langle \lambda I \rangle$, as shown in the top-right panel. The bottom-left panel reports the HA's total cost, which increases steadily with κ. This outcome reflects a dual burden: higher prices directly increase pharmaceutical expenditures, while reduced treatment intensity slows infection containment,

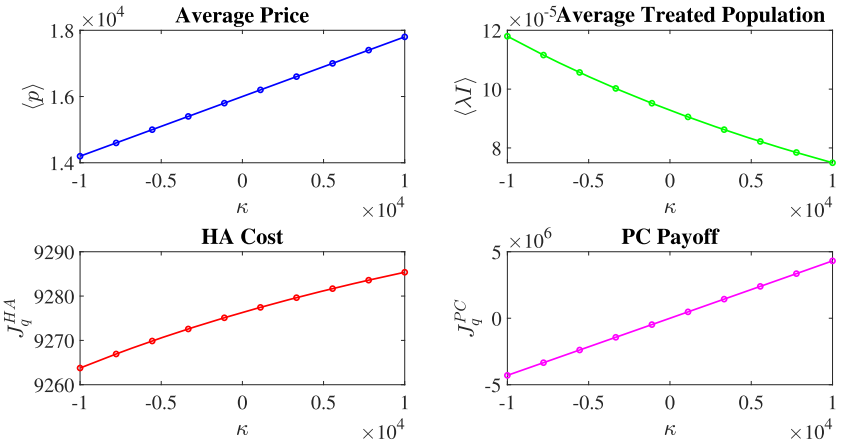

Fig. 6. Sensitivity analysis of Stackelberg outcomes with respect to the terminal weight κ in the PC's objective, with $c = €8{,}000$ and $\theta = 0$. The four panels display: (top left) average drug price $\langle p \rangle$, (top right) average treated population density $\langle \lambda I \rangle$, (bottom left) HA's per capita total cost $\mathcal{J}_q^{\mathrm{HA}}$, and (bottom right) PC's per capita payoff $\mathcal{J}_q^{\mathrm{PC}}$, computed over a simulation horizon of $T = 5$ years.

Table 3
Discounted average annual pharmaceutical expenditure B_T (in € per capita per year) sustained by the HA over a 5-year horizon, under the unconstrained optimal policy. Rows vary by the unit production cost c, columns by the PC's terminal weight κ, with $\theta = 0$.

	$\kappa = -10000$	$\kappa = -5000$	$\kappa = 0$	$\kappa = 5000$	$\kappa = 10,000$
$c = 8000$	249.99	155.88	125.00	83.08	62.25
$c = 10,000$	200.00	112.28	100.00	71.24	55.35
$c = 12,000$	166.66	125.87	83.33	62.35	49.83
$c = 14,000$	142.85	100.53	71.43	55.44	45.31
$c = 16,000$	122.02	83.69	62.50	49.90	41.54

amplifying social costs. In contrast, the PC's payoff $\mathcal{J}_q^{\mathrm{PC}}$ grows strongly with κ, as shown in the bottom-right panel. For sufficiently negative values of κ, the firm's payoff becomes negative—indicating that reputational penalties or regulatory pressures can offset the short-term economic gains from high pricing.

The analysis highlights that assigning a high strategic value to residual infection ($\kappa > 0$) incentivizes the PC to set higher prices and reduce effective coverage, ultimately worsening epidemic outcomes. Conversely, negative values of κ act as implicit regulatory mechanisms, forcing the firm to lower prices and support broader access. From a policy perspective, when the PC assigns positive value to residual infection, it tends to maintain high prices and discourage treatment, undermining eradication efforts. In the absence of explicit regulation or long-term alignment of objectives, such strategic pricing may significantly hinder public health goals. By contrast, negative values of κ, reflecting reputational or regulatory disincentives for high infection levels at the end of the horizon, encourage price reductions and broader treatment coverage. This underscores the potential effectiveness of policy tools such as outcome-based agreements, public–private partnerships, or reputational scoring systems that explicitly link pricing behavior to long-term epidemiological outcomes.

In all scenarios analyzed above, the treatment intensity $\lambda(t)$ starts above the eradication threshold $\lambda_c = 0.025$, as shown in the central panels of Figs. 3 and 5. However, it eventually falls below λ_c before time T. This behavior reflects the fact that the HA assigns no value to the final infection level ($\theta = 0$) in the current simulations. In the absence of terminal incentives, the HA optimally reduces treatment intensity over time, even when early intervention is strong.

This outcome is consistent with the structure of the HA's optimal control problem, where the terminal cost term $\theta I(T)$ directly incentivizes sustained treatment when $\theta > 0$. Theoretically, larger values of θ shift the trade-off toward higher treatment intensity near the end of the horizon, increasing the likelihood of eradication.

5.2.3. Budget benchmarks for sustaining optimal interventions

In the theoretical analysis, the PC's optimization problem was solved under the assumption that the HA's budget constraint remains inactive throughout the planning horizon, i.e., $\mu(t) = 0$ and $p(t)\lambda(t)I(t) < b(t)$ for all $t \in [0, T]$. This interiority condition ensures that the HA's treatment intensity responds smoothly to price variations, avoiding discontinuities in the optimal control. To assess the plausibility of this assumption, we compute the ex post per capita pharmaceutical expenditure sustained by the HA under the

Stackelberg equilibrium. This quantity is defined as:

$$B_T := \frac{1}{T}\int_0^T \mathrm{e}^{-\delta t} b(t)\,dt,$$

where $b(t) = p(t)\lambda(t)I(t)$ is the instantaneous pharmaceutical spending rate. The resulting value B_T corresponds to the average annual expenditure, in present value terms, required to implement the unconstrained optimal policy over the horizon $[0, T]$. This benchmark helps determine whether the assumption of an inactive budget constraint is plausible in practical settings, given the financial resources realistically available to the HA.

Table 3 reports the minimum per capita budgets B_T required to sustain an optimal treatment strategy in the absence of binding constraints. If the actual pharmaceutical resources allocated to the HA fall below these thresholds, the optimal treatment intensity $\lambda^*(t)$ derived from the unconstrained model becomes unfeasible. In such cases, the HA must revise its policy to remain within its financial means, potentially resulting in truncated, delayed, or discontinuous treatment profiles. As shown in the table, the required pharmaceutical expenditure declines both with increasing production costs c and with higher terminal valuation of residual infection by the PC (i.e., increasing κ). In both cases, higher drug prices discourage intervention, prompting the HA to scale back treatment intensity and reduce cumulative spending. Conversely, lower values of c and κ support more aggressive intervention, which entails substantially higher budgetary needs.

This pattern suggests that cost-containment policies alone may be insufficient: to ensure effective epidemic control, they must be accompanied by institutional mechanisms that preserve adequate incentives for broad treatment coverage. From a public health perspective, the expenditure levels reported in Table 3 serve as critical thresholds for assessing the financial feasibility of optimal treatment strategies. Falling below these levels may trigger deviations from the interior regime, where the budget constraint becomes binding. Such departures can have significant epidemiological and economic consequences: reduced treatment intensity may slow infection decline, raise long-term societal costs, and jeopardize eradication goals. In addition, budget-induced distortions may weaken the HA's responsiveness to pricing strategies, reduce the effectiveness of Stackelberg dynamics, and shift bargaining power toward the PC.

To maintain consistency between strategic objectives and financial capacity, it is therefore important to align public budgets with the spending levels implied by the unconstrained optimum. If tighter budgets are unavoidable, the theoretical framework should be extended to explicitly account for active budget constraints and their influence on both players' decisions. This highlights the broader need to embed financial feasibility into the design of pricing and procurement policies.

6. Discussion and policy implications

The dynamic modeling framework presented in this study provides new insights into the complex interaction between pharmaceutical pricing strategies, treatment intensity, and budget constraints in epidemic control. In particular, our analysis shows how the HA's optimal policy dynamically adapts to different economic and institutional parameters — such as budget availability, price levels, and long-term eradication incentives — and highlights the potential tensions between immediate budgetary priorities and long-term epidemiological goals.

The sensitivity analyses, conducted under a setting with exogenous and constant treatment price, reveal that the HA adopts different strategies depending on the available budget. In particular, we observe that under tight budget constraints, the HA is forced to operate well below its full treatment capacity, limiting its ability to effectively suppress the epidemic. Conversely, when the per capita budget is sufficiently high, the HA is able to respond more aggressively, significantly reducing the share of infected individuals at the terminal time. A cross-comparison of the two cost structures further refines these insights. Under severe budget constraints, the quadratic specification proves more effective in promoting early and sustained treatment interventions, resulting in better containment of the disease. However, for a sufficiently high budget, the blow-up cost structure becomes increasingly advantageous. This shift occurs because, when sufficient budget space is available, the HA can deploy moderate treatment intensities without approaching the cost singularity near the budget ceiling. In such settings, the blow-up formulation achieves superior epidemiological outcomes, such as lower terminal infection rates and higher cumulative recovery, while incurring substantially lower expenditures.

An important lesson from the simulations concerns the critical threshold λ_c: while both specifications often maintain treatment intensities above this threshold during most of the horizon, the trajectory of $\lambda(t)$ in the final phase is decisive. When the HA does not explicitly value the final state (e.g., $\theta = 0$), treatment intensity tends to decline toward zero at the end, ignoring future risk and potentially allowing infections to stabilize at low endemic levels. Conversely, increasing θ induces a more persistent effort, keeping $\lambda(t)$ above λ_c up to the terminal time. This strategic maintenance supports convergence toward a disease-free state, emphasizing the role of θ as a policy lever for long-term eradication.

These simulation-based findings resonate with earlier studies emphasizing the importance of sustained treatment efforts and dynamic budget planning. For example, it has been shown that myopic or overly conservative budget allocations can lead to suboptimal long-term outcomes, even when early intervention appears promising (Dubois and Magnac, 2024). Other works highlight that the shape of the treatment cost function critically affects the feasibility of eradication strategies, especially under resource constraints (Di Liddo, 2016). Our results extend these insights by explicitly quantifying how different cost structures (quadratic versus blow-up), terminal valuation parameters θ, and budget levels jointly influence treatment trajectories and epidemic outcomes. Furthermore, our analysis echoes recent policy discussions on performance-based agreements and flexible budget mechanisms to support access to high-cost therapies (PhRMA, 2017).

In studying the extended version of the model, where the HA and the PC play a Stackelberg game, a central issue emerging from the analysis is the alignment between private incentives and public health objectives. When pharmaceutical companies assign positive value to residual infections (high κ), they are incentivized to maintain high prices, limiting treatment coverage and undermining eradication efforts. Conversely, negative values of κ, representing reputational or regulatory penalties, encourage price reductions and broader access. This underscores the potential effectiveness of policy tools such as outcome-based agreements, public–private partnerships, or reputational scoring systems that explicitly link pricing behavior to long-term epidemiological outcomes (European Commision, 2018; Forum, 2020). Another critical dimension concerns price transparency and affordability. Several studies have highlighted that confidential discounts and managed entry agreements (MEAs) are increasingly used to manage costs, yet often fail to ensure equitable access (Vogler, 2022; Thanimalai et al., 2021). Our results indicate that relying solely on indirect price controls or negotiated caps, without explicitly accounting for public health externalities (e.g., through penalties for residual infections), may be insufficient to achieve optimal coverage and eradication. A further problem is the fragmentation of pricing and reimbursement frameworks across Europe. Differences in value assessment methods and evidence standards create opportunities for strategic pricing behaviors that erode treatment coverage (Kaló et al., 2024). Harmonization of assessment frameworks and the adoption of shared risk-sharing agreements could help address these discrepancies and promote sustainable access.

These policy considerations naturally raise the question of whether the mechanisms highlighted by the model are consistent with observed treatment dynamics in real-world settings. Although our simulations are not intended as a formal empirical validation exercise, several qualitative patterns generated by the model are consistent with empirical evidence on the rollout of direct-acting antivirals (DAAs), i.e. highly effective interferon-free therapies for hepatitis C, in France. The focus on the French case is natural in our setting, as the numerical analysis builds on the parameterization proposed by Dubois and Magnac (2024), which is explicitly calibrated to the French institutional and epidemiological context. First, our simulations highlight that treatment expansion is gradual and constrained by budgetary limits, even in the presence of highly effective therapies. This pattern is qualitatively consistent with the evolution of the hepatitis C care cascade documented by Brouard et al. (2020), who show that treatment uptake increased progressively following the introduction of DAAs. In our model, treatment intensity rises gradually and depends critically on available resources, as illustrated by the treatment paths in left panels of Fig. 1. Second, the persistence of substantial healthcare expenditures despite large-scale treatment adoption, which emerges endogenously in our framework when pharmaceutical spending is front-loaded, is consistent with real-world evidence from nationwide administrative data. Using French claims data, Lam et al. (2024) document that managing chronic hepatitis C generated a large cumulative economic burden and that healthcare costs remained elevated even several years after DAA treatment. In our simulations, infection prevalence and pharmaceutical expenditures remain positive over the planning horizon (see right panels of Fig. 1), and final outcomes reported in Table 2 show that high treatment intensity does not immediately eliminate economic costs. Third, our results emphasize the role of explicit policy objectives in sustaining treatment efforts over time. Fig. 2 shows that when the Health Authority assigns positive weight to the terminal infection level, treatment intensity remains above the eradication threshold until the end of the horizon, supporting convergence toward disease elimination. This mechanism is consistent with the French policy shift toward explicit elimination targets and universal access to DAAs. Finally, our numerical analysis predicts non-monotonic treatment trajectories in response to institutional changes and strategic incentives. This feature is strongly supported by the evidence in Brouard et al. (2024), who identify distinct phases in DAA initiation in France and sharp temporary contractions associated with exogenous shocks such as the COVID-19 pandemic. Similar non-monotonic patterns arise in our simulations under strategic pricing and changing incentives, as illustrated by the dynamic treatment paths in Fig. 3. All these insights, combining the simulation-based analysis with empirical evidence on treatment rollout and expenditure dynamics, reinforce the need for a coherent policy framework that integrates transparent pricing, dynamic budget allocation, and explicit alignment of private and public incentives. Beyond simple cost containment, future strategies should focus on long-term sustainability, equitable access, and the attainment of strategic eradication goals, while remaining robust to institutional constraints and exogenous shocks. The quantitative results provided here, interpreted in light of observed real-world patterns, may support health authorities and policymakers in negotiating pricing agreements, designing risk-sharing schemes, and defining realistic budget priorities to achieve comprehensive epidemic control.

7. Conclusions and future directions

This paper developed a dynamic optimization framework to study how a centralized Health Authority (HA) manages an innovative, curative treatment in the face of an evolving infectious disease, whose evolution is based on a continuous-time SIR model. In addition, we extended the model to a strategic setting where a forward-looking Pharmaceutical Company (PC) sets the treatment price anticipating the HA's optimal behavior. Drawing on the results obtained from both versions of the model, we summarize our main findings as follows:

1. Higher drug prices, resulting either from production costs or strategic behavior, reduce treatment intensity and slow epidemic control. Financial constraints further amplify this effect: when public budgets are insufficient, the HA is forced to limit intervention, leading to persistent infection levels. Conversely, adequate funding allows for sustained treatment and potential eradication.
2. The structure of healthcare costs plays a critical role in shaping the behavior of the optimal policy. Quadratic costs permit smooth treatment paths but require explicit enforcement of budget constraints. Blow-up costs embed these limits endogenously, ensuring compliance but generating steeper and less predictable responses. The choice of cost structure has direct implications for policy feasibility and implementation.

3. The PC's long-term strategic incentives, captured via a terminal valuation of residual infection, significantly affect pricing. When residual infection is valuable to the firm ($\kappa > 0$), it raises prices, reducing access and weakening eradication prospects. When persistent infection is penalized ($\kappa < 0$), prices fall, and broader coverage becomes feasible. In both cases, optimal pricing reflects a balance between short-term profits and long-term positioning.

These insights highlight the need for coherence between budget design, pricing rules, and health objectives. Cost-containment must be complemented by institutional mechanisms that sustain coverage incentives, whether through long-term planning, regulatory oversight, or negotiated pricing frameworks.

The Stackelberg analysis in this paper focused on quadratic treatment costs, which, despite their simplicity, provided clear insights into the interplay between pricing, intervention, and institutional constraints. Extending the framework to blow-up cost structures, where financial limits arise endogenously, remains a promising avenue. Preliminary results indicate that this leads to more complex equilibrium conditions, including cubic pricing equations, raising new questions about existence, uniqueness, and strategic behavior under financial stress.

In the present framework, both treatment and pricing decisions are determined ex ante and remain fixed over the planning horizon. As a consequence, neither the HA nor the PC adjusts their strategies in response to evolving epidemiological conditions, unexpected shocks, or parameter uncertainty. This structure is consistent with institutional settings characterized by strategic commitment, predefined budget horizons, and contractual pricing arrangements, but abstracts from fully adaptive policy responses. A natural extension of the framework would be to allow for state-dependent or feedback-based treatment and pricing policies. This would make it possible to explicitly account for real-time information, adaptive decision making, and time inconsistency in dynamic policy environments, and is left for future research.

Future work could also incorporate uncertainty (Carratelli et al., 2024), behavioral or epidemiological heterogeneity (Fenichel et al., 2011), and explicit market regulation. These extensions would help the model better reflect the complexity of real-world health systems, including stochastic dynamics, multi-firm competition, and institutional procurement constraints. Another direction involves modeling price-setting as a bargaining process between the HA and the PC, as in Dubois et al. (2022), where negotiated outcomes replace unilateral pricing. In particular, uncertainty could be incorporated through stochastic perturbations acting directly on key epidemiological or economic parameters, while preserving their boundedness. In realistic settings, quantities such as transmission intensity, treatment effectiveness, or behavioral response fluctuate within constrained ranges due to biological, institutional, or regulatory limits. Modeling such effects via bounded stochastic processes, as in d'Onofrio et al. (2018), Lacitignola and Martiradonna (2025), is particularly appealing from a methodological perspective. For instance, bounded perturbations based on sine–Wiener processes allow key parameters to fluctuate within prescribed ranges, capturing realistic uncertainty and preserving feasibility, while still accounting for regime-dependent sensitivity and robustness properties of nonlinear dynamics.

The terminal value parameters κ and θ have been introduced to capture long-term strategic considerations of both PC and HA within a finite-horizon setting. While this formulation enables the incorporation of forward-looking incentives and residual infection valuation into the analysis, future work could extend the framework to an explicit infinite-horizon formulation. Such an approach would allow for a more rigorous investigation of sustained control strategies, stationary equilibria, and the long-term interactions between pricing decisions and public health objectives.

Another promising research direction concerns price discrimination under budget asymmetry and partial market connectivity. One could consider a setting with two health authorities (e.g., from a high-income and a low-income country) facing different budget constraints, where the PC offers differentiated prices. If the two buyers are partially connected—meaning one can (legally or informally) resell treatments to the other—then arbitrage opportunities emerge, potentially undermining price discrimination strategies (Danzon, 2018). This setting introduces a new layer of complexity when combined with cost structures, as marginal costs become endogenous and sensitive to inter-agent dynamics. Formalizing this problem would provide useful insights for global procurement policy, equitable access strategies, and market segmentation under financial asymmetry. Similar concerns arise in other domains such as energy markets, environmental permits, or vaccine distribution (Miklós-Thal and Shaffer, 2021).

Funding

This publication was supported by funding from "PRA - Progetti di Ricerca di Ateneo" of the University of Foggia, under the project "Strategic Interactions in Economics, Ecology and Health". Open access funding provided by University of Foggia within the CRUI-CARE Agreement

Data availability

Data will be made available on request.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to thank the two anonymous Reviewers for their careful reading of the manuscript and for their constructive comments and suggestions, which significantly improved the paper. A.C. is a member of the Gruppo Nazionale per l'Analisi Matematica, la Probabilità e le loro Applicazioni (GNAMPA–INdAM). A.M. is a member of the Gruppo Nazionale per il Calcolo Scientifico (GNCS–INdAM).

References

Antonanzas, F., Juarez-Castello, C., Rodriguez-Ibeas, R., 2011. Should health authorities offer risk-sharing contracts to pharmaceutical firms? A theoretical approach. Health Econ. Policy Law 6 (3), 391–403.
Bailey, N. T.J., 1975. The Mathematical Theory of Infectious Diseases and its Applications.
Baker, C.M., Diele, F., Marangi, C., Martiradonna, A., Ragni, S., 2018. Optimal spatiotemporal effort allocation for invasive species removal incorporating a removal handling time and budget. Nat. Resour. Model. 31 (4), e12190.
Bressan, A., Piccoli, B., 2007. Introduction to the Mathematical Theory of Control. Vol. 1. American institute of mathematical sciences Springfield.
Brouard, C., Pillonel, J., Boussac, M., de Lédinghen, V., Rachas, A., Silvain, C., Lydié, N., Chevaliez, S., Pioche, C., Durand, J., et al., 2020. French hepatitis C care cascade: substantial impact of direct-acting antivirals, but the road to elimination is still long. BMC Infect. Dis. 20 (1), 759.
Brouard, C., Schwager, M., Expert, A., Drewniak, N., Laporal, S., de Lagasnerie, G., Lot, F., 2024. Impact of public policy and COVID-19 pandemic on hepatitis C testing and treatment in France, 2014–2021. Viruses 16 (5), 792.
Buccella, D., Fanti, L., Gori, L., 2021. To abate, or not to abate? A strategic approach on green production in Cournot and Bertrand duopolies. Energy Econ. 96, 105164.
Buonomo, B., Vargas-De-León, C., 2014. Effects of mosquitoes host choice on optimal intervention strategies for malaria control. Acta Appl. Math. 132, 127–138.
Carratelli, G.P., Cheng, X., Parag, K.V., Lestas, I., 2024. Optimal feedback policies in stochastic epidemic models. In: 2024 IEEE 63rd Conference on Decision and Control (CDC). IEEE, pp. 7062–7066.
Cesari, L., 2012. Optimization–Theory and Applications: Problems with Ordinary Differential Equations. Vol. 17. Springer Science & Business Media.
Ciulla, M., Marinelli, L., Di Biase, G., Cacciatore, I., Santoleri, F., Costantini, A., Dimmito, M.P., Di Stefano, A., 2023. Healthcare systems across europe and the US: the managed entry agreements experience, Healthcare 11, 447.
Council of the European Union, 2021. Council conclusions on access to medicines and medical devices – strengthening the balance of the pharmaceutical system in the EU and its member states. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021XG0707(01). OJ C 269, 7.7.2021, p. 1–6.
Critchley, G.J., Zaric, G.S., 2019. The impact of pharmaceutical marketing on market access, treatment coverage, pricing, and social welfare. Health Econ. 28 (8), 1035–1051.
Danzon, P.M., 2018. Differential pricing of pharmaceuticals: theory, evidence and emerging issues. Pharmacoeconomics 36 (12), 1395–1405.
Di Liddo, A., 2016. Optimal control and treatment of infectious diseases. the case of huge treatment costs. Mathematics 4 (2), 21.
Di Liddo, A., 2020. Price and treatment decisions in epidemics–Stackelberg and bargaining games: the role of cost functions. Int. Game Theory Rev. 22 (01), 1950007.
Diekmann, O., Heesterbeek, J. A.P., Roberts, M.G., 2010. The construction of next-generation matrices for compartmental epidemic models. J. R. Soc. Interface 7 (47), 873–885.
Diele, F., Marangi, C., Ragni, S., 2011. Exponential Lawson integration for nearly Hamiltonian systems arising in optimal control. Math. Comput. Simul. 81 (5), 1057–1067.
Dockner, E., 2000. Differential Games in Economics and Management Science. Cambridge University Press.
Dubois, P., Gandhi, A., Vasserman, S., 2022. Bargaining and International Reference Pricing in the Pharmaceutical Industry. Technical Report. National Bureau of Economic Research.
Dubois, P., Magnac, T., 2024. Optimal intertemporal curative drug expenses: the case of hepatitis C in France. J. Health Econ. 94, 102861.
d'Onofrio, A., Caravagna, G., de Franciscis, S., 2018. Bounded noise induced first-order phase transitions in a baseline non-spatial model of gene transcription. Physica A Stat. Mech. Appl. 492, 2056–2068.
European Commision, 2018. Innovative payment models for high-cost innovative medicines–a report by the expert panel on effective ways of investing in health Available at: https://health.ec.europa.eu/system/files/2019-11/opinion_innovative_medicines_en_0.pdf.
European Commission, 2020. Pharmaceutical strategy for Europe. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52020DC0761. COM(2020) 761 final.
Federico, S., Ferrari, G., Torrente, M.-L., 2024. Optimal vaccination in a SIRS epidemic model. Econ. Theory 77 (1), 49–74.
Fenichel, E.P., Castillo-Chavez, C., Ceddia, M.G., Chowell, G., Parra, P. A.G., Hickling, G.J., Holloway, G., Horan, R., Morin, B., Perrings, C., et al., 2011. Adaptive human behavior in epidemiological models. Proc. Natl. Acad. Sci. 108 (15), 6306–6311.
Fonkwo, P.N., 2008. Pricing infectious disease: the economic and health implications of infectious diseases. EMBO Rep. 9 (S1), S13–S17.
Forum, E.P., 2020. The value and pricing of innovative medicines: position paper. Available online: https://www.eu-patient.eu/globalassets/library/position-papers--briefings/executive-summary---pricing---finalversion.pdf.
Garibaldi, P., Moen, E.R., Pissarides, C.A., 2024. Static and dynamic inefficiencies in an optimizing model of epidemics. Econ. Theory 77 (1), 9–48.
Gavious, A., Greenberg, D., Hammerman, A., Segev, E., 2014. Impact of a financial risk-sharing scheme on budget-impact estimations: a game-theoretic approach. Eur. J. Health Econ. 15, 553–561.
Hethcote, H.W., 1989. Three basic epidemiological models. In: Applied Mathematical Ecology. Springer, pp. 119–144.
Hethcote, H.W., 2000. The mathematics of infectious diseases. SIAM Rev. 42 (4), 599–653.
Hritonenko, N., Yatsenko, Y., 2024. Analysis of optimal lockdown in integral economic–epidemic model. Econ. Theory 77 (1), 235–259.
Kaló, Z., et al., 2024. Importance of aligning the implementation of new payment models for innovative pharmaceuticals in European countries. Expert Rev. Pharmacoeconomics Outcomes Res. 24 (2), 181–187.
La Torre, D., Liuzzi, D., Marsiglio, S., 2021. Epidemics and macroeconomic outcomes: social distancing intensity and duration. J. Math. Econ. 93, 102473.
Lacitignola, D., Martiradonna, A., 2025. Can we enhance trust in the circular economy through referral marketing control? PLoS ONE 20 (9), e0332348.
Lam, L., Carrieri, P., Hejblum, G., Bellet, J., Bourlière, M., Carrat, F., 2024. Real-world economic burden of hepatitis C and impact of direct-acting antivirals in France: a nationwide claims data analysis. Liver Int. 44 (5), 1233–1242.
Lashari, A.A., Zaman, G., 2012. Optimal control of a vector borne disease with horizontal transmission. Nonlinear Anal. Real World Appl. 13 (1), 203–212.
Levaggi, R., 2014. Pricing schemes for new drugs: a welfare analysis. Social Sci. Med. 102, 69–73.
Lothan, R., Gutman, N., Yamin, D., 2022. Country versus pharmaceutical company interests for hepatitis C treatment. Health Care Manag. Sci. 25, 725–749.
Marangi, C., Casella, F., Diele, F., Lacitignola, D., Martiradonna, A., Provenzale, A., Ragni, S., 2020. Mathematical Tools for Controlling Invasive Species in Protected Areas. Mathematical Approach to Climate Change and its Impacts: MAC2I , 211–237.
Marangi, C., Martiradonna, A., Ragni, S., 2023. Optimal resource allocation for spatiotemporal control of invasive species. Appl. Math. Comput. 439, 127614.
Martiradonna, A., Diele, F., Marangi, C., 2020. Optimal control of invasive species with budget constraint: qualitative analysis and numerical approximation. Curr. Trends Dyn. Syst. Biol. Nat. Sci., 147–164.
Miklós-Thal, J., Shaffer, G., 2021. Input price discrimination by resale market. Rand J. Econ. 52 (4), 727–757.
Miraldo, M., 2009. Reference pricing and firms' pricing strategies. J. Health Econ. 28 (1), 176–197.
Parliament, E., 2021. Regulation (EU) 2021/2282 on health technology assessment and amending directive 2011/24/eu. https://eur-lex.europa.eu/eli/reg/2021/2282/oj. OJ L 458, 22.12.2021, p. 1–32.

Pharmaceutical Research and Manufacturers of America (PhRMA), 2017. Barriers to value-based contracts for innovative medicines: PhRMA member survey results. Available at: https://www.statnews.com/wp-content/uploads/2017/03/PhRMA_ValueBased_MemberService_R23.pdf.
Seierstad, A., Sydsaeter, K., 1986. Optimal Control Theory with Economic Applications. Elsevier North-Holland, Inc.
Shepard, M., Baicker, K., Skinner, J., 2013. Optimal healthcare spending with redistributive financing. http://www.dartmouthatlas.org/downloads/papers/ShepardBaickerSkinner_OptimalHealthSpending.pdf
Thanimalai, S., et al., 2021. Stakeholder views of managed entry agreements: A literature review of national studies, Health Policy Open 2, 100032.
Vogler, S., 2022. Prices of new medicines: international analysis and policy options. ZEFQ 175, 96–102.
Wenzl, M., Chapman, S., 2019. Performance-based managed entry agreements for new medicines in OECD countries and EU member states: how they work and possible improvements going forward, OECD health working papers. https://doi.org/10.1787/6e5e4c0f-en
World Health Organization, 2022. Consolidated guidelines on HIV, viral hepatitis and STI prevention, diagnosis, treatment and care for key populations. World Health Organization. Geneva. Available at: https://www.who.int/publications/i/item/9789240052390.