entropy

MDPI

Article

Optimal Chemotherapy Scheduling for Chronic Lymphocytic Leukemia Under Immune and Allergy Constraints

Rawan Abdullah [1,2,*], Andrei Halanay [2] and Lara Abou Orm [1]

[1] Department of Arts and Sciences, Rafik Hariri University, P.O. Box 10, Damour-Chouf 2010, Lebanon; abouormla@rhu.edu.lb
[2] Department of Mathematics and Informatics, Faculty of Applied Sciences, National University of Science and Technology POLITEHNICA Bucharest, 060042 Bucharest, Romania; andrei.halanay@upb.ro
* Correspondence: abdallahrb@rhu.edu.lb

Abstract

We study an optimal control framework for chemotherapy administration in patients with chronic lymphocytic leukemia (CLL) while accounting for immune regulation and treatment-induced allergic reactions. The analysis is based on a previously developed nonlinear delay differential equation model describing the interactions between leukemic cells, immune populations, antigen-presenting cells, and cytokine dynamics, with three distinct biological delays. The chemotherapy infusion rate is introduced as a time-dependent control variable and optimized to reduce leukemic burden, shift the helper T-cell balance toward a Th1-dominant configuration associated with lower hypersensitivity risk, and preserve immune competence. Existence of an optimal control is established for arbitrary delays and horizon, without the commensurability hypothesis required by reductions in delay systems to higher-dimensional delay-free ones; the argument uses only that the control enters the dynamics affinely and the running cost concavely. Necessary optimality conditions are derived via Pontryagin's Maximum Principle for systems with delays, and the resulting eleven-dimensional adjoint system, which carries advanced arguments generated by the three delays, is written out explicitly. A contraction estimate for the associated sweep operator yields both uniqueness of the optimal control on a short horizon and geometric convergence of the numerical scheme. The optimality system is solved by a forward–backward sweep adapted to the delayed setting, with documented convergence and grid independence and sensitivity analysis over kinetic parameters, delays, initial conditions and objective weights. The optimized schedule is compared not only with the untreated case and a low constant dose, but also with a constant infusion delivering the same cumulative exposure, so that the reported benefit is attributable to the temporal distribution of the dose rather than to its total amount. At equal exposure, the optimal schedule reaches each therapeutic milestone earlier—Th1 dominance 0.9 days sooner and a 90% leukemic reduction 1.6 days sooner—and attains a terminal leukemic burden lower by a factor of 2.25; a constant infusion of the same total dose reaches a comparable configuration later. The benefit of adaptive scheduling in this model is therefore principally one of rate of response at fixed drug exposure. We emphasize that the absolute Th2 population is not reduced by treatment; the reduction in hypersensitivity risk arises from the resulting Th1-dominant *relative* balance rather than from direct Th2 suppression. To characterize the therapeutic outcome in information-theoretic terms, we describe the two competing goals as distributional balances: an *allergy axis*, given by the Th1/Th2/Treg distribution, and a *leukemia axis*, given by the immune/leukemic distribution, each measured by its Shannon entropy and its Kullback–Leibler divergence to a healthy reference profile. These quantities are used in two roles. As diagnostics, they are evaluated along the computed

Check for updates

Academic Editor: Yoh Iwasa

Received: 25 June 2026
Revised: 7 August 2026
Accepted: 13 August 2026
Published: 17 August 2026

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

trajectories, and the ordering of dosing strategies is shown to be robust across twenty alternative reference profiles. As an objective, the combined divergence is then taken as the running cost of a second optimal control problem; because it depends on the leukemic population only through a normalized fraction, it prescribes a markedly gentler schedule that administers 37% of the drug and still achieves a 93% leukemic reduction, against 98% for the population-based formulation. These results suggest that adaptive, immune-aware chemotherapy scheduling may accelerate disease control at fixed drug exposure, and that information-theoretic objectives offer a scale-free alternative formulation of the therapeutic goal.

Keywords: optimal control; delay differential equations; chronic lymphocytic leukemia; chemotherapy scheduling; Pontryagin's Maximum Principle; forward–backward sweep method; Kullback–Leibler divergence; Shannon entropy; Th1/Th2 balance; sensitivity analysis

1. Introduction

Cancer has long been a major focus of scientific research, and mathematical modeling continues to provide valuable tools for understanding disease progression and improving treatment strategies. In particular, dynamical models can assist in evaluating therapeutic protocols, optimizing drug administration, and balancing treatment efficacy against adverse side effects.

Chronic lymphocytic leukemia (CLL) is a hematological malignancy characterized by the abnormal accumulation of dysfunctional mature B lymphocytes. Although several therapies are available, chemotherapy remains an important treatment option in many clinical settings. However, treatment outcomes may be complicated by toxicity, immune suppression, tumor lysis syndrome (TLS), and drug-induced hypersensitivity reactions.

The authors of [1] investigated the cytotoxic effects of several chemotherapy agents, including chlorambucil, melphalan, and cytarabine, against leukemic cells. Their findings motivated mathematical formulations in which tumor growth and death rates depend explicitly on drug concentration.

A serious complication during CLL treatment is tumor lysis syndrome, which may arise when a large number of leukemic cells are destroyed rapidly, leading to dangerous metabolic imbalances. Clinical studies such as [2,3] highlight the need for carefully designed treatment schedules that control leukemic burden while avoiding excessive toxicity. That the schedule itself, and not only the agent or the total dose, influences outcome is supported by clinical experience: a retrospective comparison of alternative chlorambucil administration schedules in previously untreated CLL reported differences in response and tolerability between regimens [4]. The present study addresses the same question in a modeling framework, by treating the infusion profile as a decision variable and comparing schedules at fixed cumulative exposure.

In addition, hypersensitivity reactions to chlorambucil have been reported in [5,6], and cytokine release has been implicated in the infusion-related reactions observed when chlorambucil is combined with monoclonal antibody therapy [7]. From an immunological viewpoint, allergic responses are strongly linked to the balance between Th1 and Th2 helper T-cell populations, together with the regulatory role of Treg cells and cytokine signaling pathways; interleukin-6 in particular has been shown to play non-redundant roles in allergic inflammation [8].

Mathematical models in [9,10] demonstrated that a shift toward Th2 dominance may promote allergic inflammation. Motivated by these biological mechanisms, our previous

studies developed delay differential equation models describing CLL treatment under desensitization protocols and analyzed several qualitative properties of the resulting systems; see [11–14].

In particular, the model presented in [13] established positivity, boundedness, global existence, and stability-related properties for a delayed CLL–immune interaction system under prescribed chemotherapy administration. These results provide the mathematical basis for the present work.

Optimal control has become a standard tool for the design of anticancer treatment schedules, and it is useful to situate the present contribution within that literature. Classical formulations optimize cytotoxic dosing against tumor burden and toxicity in ordinary differential equation models, and a recurrent qualitative finding is that schedules of bang–bang or bang-then-taper type arise whenever the control enters the dynamics linearly and the penalty on the control is quadratic. Extensions to delayed dynamics are comparatively recent: Gollmann and Maurer [15] developed the theory and numerical treatment of optimal control problems with multiple constant time delays in both state and control, and Rodrigues, Silva and Torres [16] applied delayed optimal control to an HIV model, obtaining necessary conditions of the same structure as those used here. On the numerical side, McAsey, Mou and Han [17] established convergence of the forward–backward sweep method under a Lipschitz-and-horizon-length condition, and Saleem et al. [18] surveyed its practical implementation. Relative to this body of work, the present paper differs in three respects. First, the controlled dynamics form an eleven-dimensional delay system with three distinct delays in which the control acts on a pharmacokinetic compartment rather than directly on the tumor equation, so that the drug influences the leukemic and immune populations only indirectly. Second, the objective couples disease control to an immunological side-effect constraint—the Th1/Th2/Treg balance governing hypersensitivity—rather than to a generic toxicity term. Third, the therapeutic goal is formulated in two ways, once in terms of cell populations and once in information-theoretic terms, and the two resulting optimal schedules are compared.

The aim of this paper is different from the previous desensitization-based approach. Here, we seek an *optimal time-dependent chemotherapy strategy* by introducing the drug infusion rate as a control variable. The objective is to determine dosing schedules that simultaneously:

- Reduce leukemic burden.
- Shift the helper T-cell balance toward a Th1-dominant configuration, which is associated with reduced allergy-related hypersensitivity risk.
- Preserve essential immune cell populations.
- Avoid excessive treatment exposure.

To achieve this goal, we formulate an optimal control problem governed by a nonlinear delay differential equation system. The existence of an optimal control is established for arbitrary delays and horizons; the argument requires only that the control enter the dynamics affinely and the running cost concavely, and therefore dispenses with the commensurability hypothesis on which reductions in delay systems to higher-dimensional delay-free systems rely. Necessary optimality conditions are derived using Pontryagin's Maximum Principle for systems with delays, and the resulting eleven-dimensional adjoint system, which carries advanced arguments generated by the three delays, is written out explicitly. A contraction estimate for the associated sweep operator yields both uniqueness of the optimal control on a short horizon and geometric convergence of the numerical scheme.

The optimality system is solved numerically using a forward–backward sweep procedure adapted to the delayed setting. A methodological point governs how the results are reported. Since the computed optimal schedule delivers substantially more drug than

a low constant infusion, comparing the two conflates the effect of dose *timing* with that of dose *amount*. We therefore include a constant infusion calibrated to deliver exactly the cumulative exposure of the optimal schedule, and it is against this exposure-matched protocol that the contribution of adaptive scheduling is assessed. The simulations show that at equal cumulative exposure the optimal schedule reaches every therapeutic milestone earlier—Th1 dominance, a fifty and a ninety percent leukemic reduction, and the minimum of the combined therapeutic divergence—while attaining a lower terminal leukemic burden. The benefit of adaptive dosing in this model is therefore principally one of rate of response at fixed drug exposure rather than of eventual position. The computed schedule does, however, administer the maximal dose over an initial phase of approximately three days, which stands in tension with the concern about rapid leukemic destruction noted above; this is discussed among the limitations.

Beyond population-level trajectories, the therapeutic goal of the treatment is itself naturally described in information-theoretic terms. The treatment must balance two competing objectives, each of which can be expressed as a distributional balance between cell populations. The first is an *allergy axis*: the relative proportions of Th1, Th2, and regulatory T cells determine the hypersensitivity risk, since a Th1-dominant configuration is associated with a low allergic response, whereas a shift toward Th2 dominance promotes allergic inflammation. The second is a *leukemia axis*: the relative proportions of immune effector and leukemic cells determine disease control, an immune-dominant configuration corresponding to effective tumor suppression. Each balance forms a probability distribution whose Shannon entropy measures its diversity and regulatory order, while the Kullback–Leibler divergence to a healthy reference profile (Th1-favorable on the allergy axis, immune-dominant on the leukemia axis) quantifies how far the current state lies from the desired one.

These information-theoretic quantities are used in two distinct roles, and it is worth separating them at the outset. In the first they are *diagnostics*: the optimal control is determined by the objective functional of Section 4, which is expressed directly in terms of cell populations, and the entropy and relative entropy are then evaluated *along* the resulting trajectories, providing a dimensionless summary comparable across strategies and across the two axes. In the second, developed in Section 6.7, the combined divergence itself becomes the running cost of a distinct optimal control problem, so that the information-theoretic quantity is the object being optimized rather than a measure of the outcome. Comparing the two resulting schedules answers a control-theoretic question rather than a descriptive one: whether the population-level and information-theoretic formulations of the therapeutic goal induce the same optimal policy. They do not. Because the divergence depends on the leukemic population only through a normalized fraction, it becomes insensitive to further absolute reduction once that fraction is small, and the divergence-driven schedule consequently administers substantially less drug while forgoing only a small part of the leukemic reduction achieved by the population-based schedule.

The main contributions of this paper are as follows:

1. Formulation of a delay optimal control framework for chemotherapy dosing in CLL with immune and allergy constraints, with an explicitly weighted objective functional.
2. Proof of existence of an optimal treatment strategy for arbitrary delays and horizons, without the commensurability hypothesis used in earlier reductions in delayed control systems, together with uniqueness, Lipschitz regularity and absence of singular arcs for the optimal schedule.

3. Derivation of necessary optimality conditions via Pontryagin's Maximum Principle, including the explicit eleven-dimensional adjoint system with its advanced arguments, and a contraction estimate that simultaneously yields uniqueness of the optimizer and convergence of the forward–backward sweep.
4. Numerical solution by a forward–backward sweep adapted to the delayed setting, with documented convergence, grid independence, and sensitivity analysis over kinetic parameters, delays, initial conditions and objective weights.
5. Comparison of the optimized schedule against untreated, low-dose and *exposure-matched* constant protocols, isolating the effect of dose timing from that of cumulative dose.
6. Information-theoretic characterization of both the allergy axis (Th1/Th2/Treg) and the leukemia axis (immune/leukemic distribution) via the Shannon entropy and Kullback–Leibler divergence, in two roles: as diagnostics whose ordering of strategies is shown to be robust to the choice of reference profile, and as the running cost of a second optimal control problem whose solution is compared with the first.

2. Controlled Delay Model for CLL Dynamics

We consider the delayed mathematical model for chronic lymphocytic leukemia (CLL) treatment introduced in [13], where positivity, boundedness, global existence, and qualitative dynamical properties were previously investigated. In the present work, this framework is extended to an optimal control setting by introducing a time-dependent chemotherapy administration rate $u(t)$, representing the normalized drug infusion intensity. The state equations are otherwise unchanged; the sole modification is that the infusion term in the pharmacokinetic equation, previously a prescribed function of time, is now a decision variable.

The model describes the interactions among leukemic cells, immune populations, antigen-presenting cells, cytokines, and the chemotherapeutic drug. In particular, it incorporates helper T-cell subpopulations (Th1, Th2), regulatory T cells, naive immune cells, activated antigen-presenting cells, immune effector cells, and living/dead leukemic cells.

Three biologically relevant delays are included:

- τ_1: propagation delay associated with transport from the central to the peripheral compartment;
- τ_2: delay corresponding to cytokine production and signaling;
- τ_3: leukemic cell-cycle duration.

The control variable satisfies

$$u(t) \in [0,1], \qquad t \in [0, t_f],$$

where $u(t) = 0$ corresponds to no treatment and $u(t) = 1$ represents maximal normalized drug administration. Note that u enters the system only through the additive infusion term of the pharmacokinetic equation, so that the right-hand side is affine in the control; this structural fact is used in the existence proof of Section 4.

Let the state variables be defined as follows:

x_1 : naive T-helper cells,
x_2 : Th1 cells,
x_3 : Th2 cells,
x_4 : regulatory T cells,
x_5 : naive APCs,
x_6 : mature APCs,
x_7 : drug concentration,
x_8 : cytokine concentration,
x_9 : living leukemic cells,
x_{10} : dead leukemic cells,
x_{11} : immune effector cells.

All concentrations are expressed per microlitre, as set out in Section 8.

The controlled delay differential system is therefore written as

$$\begin{cases} \dot{x}_1 = \alpha - \beta_1 x_1 - x_1 x_6(t-\tau_1)\dfrac{x_2}{1+\mu_2 x_3} - \phi x_1 x_6(t-\tau_1)x_3 - \kappa x_1 x_6(t-\tau_1)x_4 \\ \dot{x}_2 = -\beta_2 x_2 + v\dfrac{x_1 x_6(t-\tau_1)x_2}{(1+\mu_r x_4)(1+\mu_2 x_3)} \\ \dot{x}_3 = -\beta_3 x_3 + \phi v\dfrac{x_1 x_6(t-\tau_1)x_3(1+\mu_2 x_3)}{(1+\mu_r x_4)(1+\mu_1 x_2+\mu_2 x_3)} \\ \dot{x}_4 = -\beta_4 x_4 + \kappa v x_1 x_6(t-\tau_1)x_4 - \eta_r\dfrac{x_8 x_4}{1+x_8} \\ \dot{x}_5 = a_{APC} - \beta_0 x_7 x_5 - \gamma_{11} x_5 \\ \dot{x}_6 = \beta_0 x_7 x_5 - \gamma_{12} x_6 - \mu_0 x_6 x_4 \\ \dot{x}_7 = u(t) - \gamma_D x_7 - \dfrac{\mu_D x_9 x_7}{a+x_7} \\ \dot{x}_8 = -\gamma_2 x_8 + k_1[x_1(t-\tau_2) + x_2(t-\tau_2) + x_3(t-\tau_2) + x_4(t-\tau_2) + x_6(t-\tau_2)] \\ \dot{x}_9 = -\gamma_L x_9 - \Psi(x_9) + 2e^{-\gamma_L \tau_3}\Psi\big(x_9(t-\tau_3)\big) - c_1 x_9 x_{11} - \dfrac{\mu_L x_9 x_7}{a+x_7} \\ \dot{x}_{10} = \gamma_L x_9 - d x_{10} + \dfrac{\mu_L x_9 x_7}{a+x_7} \\ \dot{x}_{11} = s - m x_{11} + \dfrac{\rho x_9 x_{11}}{\gamma + x_9} - \dfrac{\delta x_{11} x_7}{c+x_7} - c_2 x_9 x_{11} \end{cases} \tag{1}$$

The leukemic self-renewal flux is taken as

$$\Psi(x) = \beta_L \frac{\theta^2 x^2}{\theta^2 + x^2}, \tag{2}$$

a Hill-type function of the leukemic load with half-saturation θ. The factor $2e^{-\gamma_L \tau_3}$ accounts for the doubling of cells completing the cycle, discounted by mortality over its duration τ_3. The corresponding per-cell renewal rate is $\Psi(x)/x = \beta_L \theta^2 x/(\theta^2 + x^2)$, which increases with leukemic density for $x < \theta$ and decreases beyond it, attaining its maximum at $x = \theta$. This density-dependent form differs from the monotonically decreasing per-cell rate of [19]; the leukemic burdens attained in the present simulations remain well below θ, so the model operates throughout in the regime where the flux is approximately $\beta_L x^2$, as discussed in Section 6.2 and in the limitations.

Initial data are prescribed through history functions on the interval

$$[-\tau_{\max}, 0], \qquad \tau_{\max} = \max\{\tau_1, \tau_2, \tau_3\},$$

namely,

$$x_i(t) = \varphi_i(t), \qquad t \in [-\tau_{\max}, 0], \quad i = 1, \ldots, 11.$$

The qualitative properties of the uncontrolled dynamics and the biological relevance of the model parameters were established in [13]. In the present study, the main novelty lies in treating $u(t)$ as a decision variable and determining an optimal chemotherapy schedule that balances treatment efficacy, immune preservation, and reduction in allergy-related adverse responses. For the general theory of delay differential systems underlying the analysis that follows—existence, continuation and stability of solutions of functional differential equations—we refer to [20,21].

3. Well-Posedness and Preliminary Properties

The optimal control analysis developed below relies on three structural properties of delayed System (1): existence and uniqueness of solutions for every admissible control, invariance of the nonnegative orthant, and boundedness of the state on the finite horizon. These were established in [13] for the uncontrolled dynamics with a prescribed infusion profile; we restate them in the form required here, since the control now ranges over an entire admissible class rather than being a fixed function of time.

Throughout, admissible controls are the measurable functions

$$\mathcal{U} = \{u : [0, t_f] \to [0, 1] \ : \ u \text{ measurable}\},$$

and we write $\tau_{\max} = \max\{\tau_1, \tau_2, \tau_3\}$ and

$$\mathcal{C}^+ = \left\{\varphi \in C\big([-\tau_{\max}, 0]; \mathbb{R}^{11}\big) \ : \ \varphi_i(s) \geq 0 \ \ \forall s, \ i = 1, \ldots, 11\right\}$$

for the cone of nonnegative continuous initial histories.

Proposition 1 (Well-posedness of the controlled system)**.** *Let $\varphi \in \mathcal{C}^+$ and let $u \in \mathcal{U}$. Then, the following apply:*

(i) (Local existence and uniqueness.) *The right-hand side F of (1) is continuous in t and locally Lipschitz in (x, x_τ) on $\mathbb{R}^{11}_+ \times \mathbb{R}^{11\times 3}_+$, uniformly for $u \in [0, 1]$. Consequently, Equation (1) admits a unique solution on a maximal interval $[0, t_{\max})$, obtained by the method of steps on successive intervals of length $\min\{\tau_1, \tau_2, \tau_3\}$.*

(ii) (Positivity.) *$x_i(t) \geq 0$ for all $t \in [0, t_{\max})$ and all i. Each equation of Equation (1) has the form $\dot{x}_i = P_i - x_i Q_i$ with $P_i \geq 0$ and Q_i bounded on bounded sets whenever the state is nonnegative,*
so that

$$x_i(t) = \left[x_i(0) + \int_0^t P_i(r) e^{\int_0^r Q_i \, dr}\right] e^{-\int_0^t Q_i} \ \geq \ 0.$$

(iii) (Boundedness and global existence.) *On every finite horizon $[0, t_f]$ there is a constant $K = K(\varphi, t_f)$, independent of $u \in \mathcal{U}$, with $\|x(t)\|_\infty \leq K$ for all $t \in [0, t_f]$. In particular $t_{\max} > t_f$.*

Sketch. Parts (i) and (ii) are as in ([13], Section 3); the control does not affect them, entering only additively and boundedly in the x_7 equation. For (iii) we chain a priori estimates, each uniform in u. From $\dot{x}_7 \leq 1 - \gamma_D x_7$ we get $x_7(t) \leq \max\{x_7(0), 1/\gamma_D\}$, uniform because $u \leq 1$. Then $\dot{x}_5 \leq a_{APC} - \gamma_{11} x_5$ and $\dot{x}_6 \leq \beta_0 x_7 x_5 - \gamma_{12} x_6$ bound the antigen-presenting compartments. Self-renewal flux Equation (2) is bounded above by $\beta_L \theta^2$, since $\Upsilon(x) \leq \beta_L \theta^2$ for every $x \geq 0$; the x_9 equation therefore gives $\dot{x}_9 \leq 2e^{-\gamma_L \tau_3} \beta_L \theta^2$, so that x_9 grows at most linearly on $[0, t_f]$, and the same bounds x_{10} and, through the linear supply term, x_{11}. Finally,

$x_1, \ldots, x_4$ are bounded because their gain terms are dominated by $v\, x_1 x_6(t-\tau_1)$ with x_6 already bounded, and x_8 by the linear cytokine equation driven by bounded arguments. □

Remark 1. *The uniformity in u in part (iii) is what the existence argument of Theorem 1 consumes: it makes the family of state trajectories generated by $\mathcal{U}$ uniformly bounded and equicontinuous on $[0, t_f]$.*

4. Optimal Control Formulation

We formulate the chemotherapy scheduling problem for chronic lymphocytic leukemia (CLL) as a finite-horizon optimal control problem, where the chlorambucil infusion rate is treated as a time-dependent control variable. The objective is to determine a dosing strategy that balances therapeutic efficacy, immune preservation, and reduction in treatment-related adverse effects.

The controlled dynamics are governed by the delay differential system (Equation (1)), subject to the initial history conditions

$$x_i(t) = \varphi_i(t), \qquad t \in [-\tau_{\max}, 0], \quad i = 1, \ldots, 11.$$

Recall from Section 3 the admissible control set

$$\mathcal{U} = \left\{ u : [0, t_f] \to [0,1] \;\middle|\; u \text{ is measurable} \right\},$$

where $u(t) = 0$ corresponds to no drug administration and $u(t) = 1$ denotes the maximal normalized dose.

The therapeutic goals are as follows:

- Enhance anti-tumor immune activity through Th1 cells $x_2(t)$ and effector cells $x_{11}(t)$.
- Shift the helper T-cell balance toward Th1 dominance by penalizing the Th2 population $x_3(t)$.
- Reduce the leukemic cell population $x_9(t)$.
- Avoid excessive chemotherapy exposure.

Accordingly, we consider the weighted objective functional

$$J(u) = \int_0^{t_f} \left[A_1 x_{11}(t) + A_2 x_2(t) - A_3 x_3(t) - A_4 x_9(t) - A_5 u^2(t) \right] dt, \qquad A_1, \ldots, A_5 > 0. \tag{3}$$

The positive terms reward immune reinforcement, whereas the negative terms penalize leukemic burden, allergy-related responses, and excessive drug administration. The coefficients $A_1, \ldots, A_4$ are trade-off weights between the four therapeutic goals and A_5 sets the price of drug exposure relative to them. Throughout the numerical experiments the baseline choice is

$$A_1 = A_2 = A_3 = A_4 = A_5 = 1,$$

weighting the four population objectives equally; the dependence of the computed schedule on this choice is examined in Section 6.2.

Remark 2 (Dimensional consistency)**.** *All state variables in Equation (1) are expressed in normalized concentration units, and u is a dimensionless normalized infusion intensity in $[0, 1]$. The weights $A_1, \ldots, A_4$ therefore carry the reciprocal of the corresponding concentration unit and A_5 is dimensionless, so every term of the integrand of Equation (3) has the same physical dimension. The values quoted above are those of the weights in this normalized system of units. This is why the weights must be displayed rather than absorbed into the integrand: the value of J is meaningful*

only relative to a stated normalization, and comparisons of J across dosing strategies are made at fixed weights.

The optimal control problem is to determine $u^* \in \mathcal{U}$ such that

$$J(u^*) = \max_{u \in \mathcal{U}} J(u).$$

To derive necessary optimality conditions, we apply Pontryagin's Maximum Principle for systems with discrete delays [22,23]. This yields a coupled state–adjoint system together with a characterization of the optimal control.

The resulting optimality system is solved numerically by a forward–backward sweep iterative procedure adapted to incorporate delay terms in both the state and adjoint equations. This framework provides a practical methodology for designing adaptive chemotherapy schedules that improve treatment effectiveness while limiting adverse immune responses.

Theorem 1 (Existence of an optimal control)**.** *Let $\varphi \in \mathcal{C}^+$ and let $t_f > 0$ be arbitrary. Consider delayed System (1) with the objective functional Equation (3) over the admissible set $\mathcal{U}$. Then there exists $u^* \in \mathcal{U}$ with*

$$J(u^*) = \max_{u \in \mathcal{U}} J(u),$$

and the associated trajectory x^* is the unique solution of Equation (1) corresponding to u^*. No arithmetic relation among the delays τ_1, τ_2, τ_3 and the horizon t_f is required.

Proof. The admissible set is nonempty and J is bounded above on $\mathcal{U}$, since by Proposition 1(iii) all admissible trajectories remain in a compact set $\mathcal{K} \subset \mathbb{R}^{11}_+$ depending only on φ and t_f; write $J^* := \sup_{u \in \mathcal{U}} J(u) < \infty$. Let $(u^n) \subset \mathcal{U}$ be a maximizing sequence with associated states x^n.

Step 1 (compactness of controls). $\mathcal{U}$ is a bounded, convex, weak-$*$ closed subset of $L^\infty(0, t_f)$, hence weak-$*$ sequentially compact by the Banach–Alaoglu theorem together with the separability of $L^1(0, t_f)$. Passing to a subsequence, $u^n \rightharpoonup^* u^*$ with $u^* \in \mathcal{U}$, the pointwise bounds being preserved because $\mathcal{U}$ is convex and weak-$*$ closed.

Step 2 (convergence of the states). Since $x^n(t) \in \mathcal{K}$ and $|u^n| \le 1$, the right-hand side F is uniformly bounded along the sequence, so $(\dot{x}^n)$ is uniformly bounded and (x^n) is uniformly bounded and equicontinuous on $[0, t_f]$. By the Arzelà–Ascoli theorem a further subsequence converges uniformly, $x^n \to x$ in $C([0, t_f]; \mathbb{R}^{11})$; the same convergence holds for the shifted arguments $x^n(\cdot - \tau_j)$, since the initial histories are fixed and uniform convergence is preserved under translation.

Step 3 (admissibility of the limit). Write Equation (1) in integral form. The control appears only in the seventh component and only additively:

$$x_7^n(t) = \varphi_7(0) + \int_0^t u^n(r)\, dr - \int_0^t \left[\gamma_D x_7^n + \frac{\mu_D x_9^n x_7^n}{a + x_7^n}\right] dr.$$

The first integral converges to $\int_0^t u^*$ for each t by weak-$*$ convergence, since $\chi_{[0,t]} \in L^1$. The remaining integrals, and all components $i \neq 7$, involve u nowhere, and their integrands converge uniformly by Step 2 together with the local Lipschitz continuity of F on $\mathcal{K}$. Hence x satisfies Equation (1) with control u^* and history φ, and by Proposition 1(i) it is the unique such solution: $x = x^*$.

Step 4 (upper semicontinuity of the objective). Decompose

$$J(u^n) = \underbrace{\int_0^{t_f} \left[A_1 x_{11}^n + A_2 x_2^n - A_3 x_3^n - A_4 x_9^n\right] dt}_{=:J_1(x^n)} - A_5 \int_0^{t_f} \left(u^n\right)^2 dt.$$

The term J_1 is linear in the state and converges to $J_1(x^*)$ by uniform convergence. The map $u \mapsto \int_0^{t_f} u^2$ is convex and strongly continuous on $L^2(0, t_f)$, hence weakly lower semicontinuous, so

$$\int_0^{t_f} \left(u^*\right)^2 dt \;\le\; \liminf_{n\to\infty} \int_0^{t_f} \left(u^n\right)^2 dt.$$

Combining, $J(u^*) \ge \limsup_{n\to\infty} J(u^n) = J^*$, and since $u^* \in \mathcal{U}$ also $J(u^*) \le J^*$. Therefore $J(u^*) = J^*$. □

Remark 3 (On the role of the structural hypotheses)**.** *The proof uses two structural features of the problem and nothing else. First, F is* affine *in u: the control enters only as the additive infusion term of the pharmacokinetic equation. Second, the integrand of Equation (3) is* concave *in u, the control appearing only through* $-A_5 u^2$*. Together these are precisely the convexity condition on the extended velocity set*

$$N(t, x, x_\tau) = \left\{ \left(F(t, x, x_\tau, u),\, z\right) :\; z \le A_1 x_{11} + A_2 x_2 - A_3 x_3 - A_4 x_9 - A_5 u^2,\; u \in [0,1] \right\}$$

required by the Filippov–Cesari theory [24]: for $u_a, u_b \in [0,1]$ *and* $\mu \in [0,1]$*, affineness of F gives* $\mu F(u_a) + (1-\mu) F(u_b) = F(\mu u_a + (1-\mu) u_b)$*, while concavity of the integrand gives the corresponding inequality in the last coordinate, so N is convex. If the control acted nonlinearly on the dynamics, or if the penalty on u were replaced by an* L^1 *term, this convexity could fail, and a relaxed formulation would be required.*

Earlier treatments of delayed optimal control problems of this type obtain existence by exploiting rational dependence of the delays to transform the system into an equivalent delay-free system of higher dimension [25,26], which requires $\tau_j = k_j \tau$ *and* $t_f = L\tau$ *with integer* k_j, L*. The argument above shows that for the present problem that hypothesis is dispensable: the delays enter only through arguments handled by the Arzelà–Ascoli step, and the compactness required is supplied by the uniform bounds of Proposition 1(iii) rather than by a finite-dimensional reduction. The delays of Table 8 are consequently used exactly as tabulated, with no rounding to a commensurate grid.*

5. Adjoint System and Optimality Conditions

To derive the necessary conditions for optimality, we apply Pontryagin's Maximum Principle for systems with discrete delays [22,23]. Let

$$x(t) = \left(x_1(t), \ldots, x_{11}(t)\right)^\top,$$

and denote the delayed state variables by

$$x_\tau(t) = \left(x(t-\tau_1), x(t-\tau_2), x(t-\tau_3)\right).$$

Controlled System (1) can be written in compact form as

$$\dot{x}(t) = F\left(t, x(t), x_\tau(t), u(t)\right).$$

We introduce the adjoint vector

$$\lambda(t) = \left(\lambda_1(t), \ldots, \lambda_{11}(t)\right)^\top,$$

and define the Hamiltonian by

$$H(t, x, x_\tau, u, \lambda) = A_1 x_{11} + A_2 x_2 - A_3 x_3 - A_4 x_9 - A_5 u^2 + \lambda^\top F(t, x, x_\tau, u). \tag{4}$$

For an optimal control $u^*(t)$ with associated trajectory $x^*(t)$, there exists an adjoint function $\lambda(t)$ satisfying

$$\dot{\lambda}_i(t) = -\left(\frac{\partial H}{\partial x_i}(t) + \sum_{j=1}^{3} \chi_{[0, t_f - \tau_j]}(t) \frac{\partial H}{\partial y_{i\tau_j}}(t + \tau_j) \right), \qquad i = 1, \ldots, 11, \tag{5}$$

where $y_{i\tau_j}$ denotes the argument of the Hamiltonian corresponding to the delayed state $x_i(t - \tau_j)$, with terminal conditions

$$\lambda_i(t_f) = 0, \qquad i = 1, \ldots, 11.$$

We draw attention to the structure of Equation (5): the delays of the state system appear in the adjoint system as *advanced* arguments evaluated at $t + \tau_j$, which is what dictates both the explicit form given in Section 5.1 and the numerical scheme of Section 5.4.

The optimal control is characterized by the maximization condition

$$H\big(t, x^*(t), x^*_\tau(t), u^*(t), \lambda(t)\big) = \max_{0 \le u \le 1} H\big(t, x^*(t), x^*_\tau(t), u, \lambda(t)\big),$$

which is made explicit in the characterization given in Section 5.3 below.

5.1. Explicit Form of the Adjoint System

Because the model is eleven-dimensional with three distinct delays, we display the adjoint system in full; this is also what is implemented in the backward sweep of Section 5.4. Abbreviate

$$y_1(t) := x_6(t - \tau_1), \qquad W := 1 + \mu_2 x_3, \qquad R := 1 + \mu_r x_4, \qquad M := 1 + \mu_1 x_2 + \mu_2 x_3,$$

$$G(z) := \Psi'(z) = \frac{2\beta_L \theta^4 z}{\left(\theta^2 + z^2\right)^2}, \qquad \chi_j(t) := \chi_{[0, t_f - \tau_j]}(t), \quad j = 1, 2, 3,$$

and write

$$\Phi := \lambda_1 \Big(\frac{x_1 x_2}{W} + \phi\, x_1 x_3 + \kappa\, x_1 x_4 \Big) - \frac{\lambda_2\, v\, x_1 x_2}{WR} - \frac{\lambda_3\, \phi\, v\, x_1 x_3\, W}{RM} - \kappa\, \lambda_4\, v\, x_1 x_4, \tag{6}$$

all quantities in Equation (6) being evaluated at the same instant. With Hamiltonian Equation (4), adjoint Equation (5) reads

$$\dot{\lambda}_1 = \lambda_1 \Big(\beta_1 + \phi x_3 y_1 + \kappa x_4 y_1 + \frac{x_2 y_1}{W} \Big) - \frac{\lambda_2 v x_2 y_1}{WR} - \frac{\lambda_3 \phi v x_3 y_1 W}{RM} - \kappa \lambda_4 v x_4 y_1 \; - \; \chi_2\, k_1 \lambda_8 (t + \tau_2), \tag{7}$$

$$\dot{\lambda}_2 = \; - A_2 + \frac{\lambda_1 x_1 y_1}{W} + \lambda_2 \Big(\beta_2 - \frac{v x_1 y_1}{WR} \Big) + \frac{\lambda_3 \mu_1 \phi v x_1 x_3 y_1 W}{RM^2} \; - \; \chi_2\, k_1 \lambda_8 (t + \tau_2), \tag{8}$$

$$\begin{aligned} \dot{\lambda}_3 = A_3 &- \lambda_1 \Big(\frac{\mu_2 x_1 x_2 y_1}{W^2} - \phi x_1 y_1 \Big) + \frac{\lambda_2 \mu_2 v x_1 x_2 y_1}{W^2 R} \\ &+ \lambda_3 \Big(\beta_3 + \frac{\mu_2 \phi v x_1 x_3 y_1 W}{RM^2} - \frac{\mu_2 \phi v x_1 x_3 y_1}{RM} - \frac{\phi v x_1 y_1 W}{RM} \Big) \; - \; \chi_2\, k_1 \lambda_8 (t + \tau_2), \end{aligned} \tag{9}$$

$$\dot{\lambda}_4 = \kappa\lambda_1 x_1 y_1 + \frac{\lambda_2 \mu_r v x_1 x_2 y_1}{WR^2} + \frac{\lambda_3 \mu_r \phi v x_1 x_3 y_1 W}{R^2 M} + \lambda_4\left(\beta_4 + \frac{\eta_r x_8}{1+x_8} - \kappa v x_1 y_1\right) + \lambda_6 \mu_0 x_6 - \chi_2\, k_1 \lambda_8(t+\tau_2), \tag{10}$$

$$\dot{\lambda}_5 = \lambda_5(\beta_0 x_7 + \gamma_{11}) - \lambda_6 \beta_0 x_7, \tag{11}$$

$$\dot{\lambda}_6 = \lambda_6(\gamma_{12} + \mu_0 x_4) + \chi_1\, \Phi(t+\tau_1) - \chi_2\, k_1 \lambda_8(t+\tau_2), \tag{12}$$

$$\dot{\lambda}_7 = \lambda_7\left(\gamma_D + \frac{a\,\mu_D x_9}{(a+x_7)^2}\right) + \beta_0 x_5(\lambda_5 - \lambda_6) + (\lambda_9 - \lambda_{10})\frac{a\,\mu_L x_9}{(a+x_7)^2} + \frac{c\,\delta\,\lambda_{11} x_{11}}{(c+x_7)^2}, \tag{13}$$

$$\dot{\lambda}_8 = \gamma_2 \lambda_8 + \frac{\eta_r \lambda_4 x_4}{(1+x_8)^2}, \tag{14}$$

$$\dot{\lambda}_9 = A_4 + \frac{\lambda_7 \mu_D x_7}{a+x_7} + \lambda_9\left(\gamma_L + G(x_9) + c_1 x_{11} + \frac{\mu_L x_7}{a+x_7}\right) - \lambda_{10}\left(\gamma_L + \frac{\mu_L x_7}{a+x_7}\right) + \lambda_{11}\left(c_2 x_{11} - \frac{\rho\,\gamma\, x_{11}}{(\gamma + x_9)^2}\right) - \chi_3\, 2e^{-\gamma_L \tau_3}\, \lambda_9(t+\tau_3)\, G\big(x_9(t)\big), \tag{15}$$

$$\dot{\lambda}_{10} = d\,\lambda_{10}, \tag{16}$$

$$\dot{\lambda}_{11} = -A_1 + c_1 \lambda_9 x_9 + \lambda_{11}\left(m + c_2 x_9 + \frac{\delta x_7}{c+x_7} - \frac{\rho x_9}{\gamma + x_9}\right). \tag{17}$$

Unless indicated otherwise all state and adjoint variables are evaluated at t.

Three structural features dictate the numerical treatment. First, the advanced arguments $\lambda_8(t+\tau_2)$, $\lambda_9(t+\tau_3)$ and $\Phi(t+\tau_1)$ are already available when required, since the adjoint system is integrated backward from t_f; the indicators χ_j switch them off on the terminal window $(t_f - \tau_j,\, t_f]$. Second, only λ_7 enters the control characterization, but λ_7 is coupled to $\lambda_5, \lambda_6, \lambda_9, \lambda_{10}, \lambda_{11}$, so the full system must be integrated. Third, $\dot{\lambda}_{10} = d\lambda_{10}$ with $\lambda_{10}(t_f) = 0$ gives $\lambda_{10} \equiv 0$: the dead-cell compartment neither appears in the objective nor feeds back into the other equations, so it is decoupled from the optimization. We retain it in the implementation as a consistency check on the backward integrator.

5.2. Transversality Conditions

Because the objective functional is purely integral, with no terminal cost term $h(x(t_f), t_f)$, and the final state $x(t_f)$ is free, the transversality conditions of Pontryagin's Maximum Principle reduce to the homogeneous form

$$\lambda_i(t_f) = 0, \qquad i = 1, \ldots, 11.$$

In the absence of a terminal penalty, variations in the final state do not affect the value of the objective functional, so the costates vanish at the terminal time. This reasoning follows the standard treatment in ([27], [Chapter 5, p. 192]). We note for later reference that this is the mechanism by which the optimized infusion is withdrawn as $t \to t_f$, a feature discussed in Section 6.

5.3. Characterization of the Optimal Control

According to Pontryagin's Maximum Principle, the optimal control maximizes the Hamiltonian almost everywhere over the admissible set $0 \le u(t) \le 1$. Since

$$\frac{\partial H}{\partial u} = \lambda_7(t) - 2A_5\, u(t), \qquad \frac{\partial^2 H}{\partial u^2} = -2A_5 < 0,$$

the Hamiltonian is strictly concave in u, the pointwise maximizer is unique, and the stationary condition yields the unconstrained candidate $u = \lambda_7/(2A_5)$. Taking the pointwise bounds into account, the optimal chemotherapy administration rate is obtained by projection onto the admissible interval:

$$u^*(t) = \min\left\{1, \max\left\{0, \frac{\lambda_7(t)}{2A_5}\right\}\right\}. \tag{18}$$

This expression guarantees feasibility while satisfying the first-order optimality condition. In particular u^* is continuous whenever λ_7 is, so the optimal schedule contains no jump discontinuities: the saturated phase $u^* \equiv 1$ observed in Section 6 is the interval on which $\lambda_7(t) \geq 2A_5$, and the subsequent taper follows the decay of λ_7 toward its zero terminal value.

Two further properties follow, which are not automatic for delayed optimal control problems and which determine the behavior of the numerical scheme.

Proposition 2 (Regularity and absence of singular arcs)**.** *Let u^* be optimal with adjoint λ. Then, the following apply:*

(i) *The maximizer of $u \mapsto H(t, x^*, x^*_\tau, u, \lambda)$ over $[0,1]$ is unique for every t, and is given by Equation (18).*
(ii) *u^* is Lipschitz continuous on $[0, t_f]$, with $\mathrm{Lip}(u^*) \leq \|\dot{\lambda}_7\|_\infty/(2A_5)$.*
(iii) *The problem admits no singular arcs.*
(iv) *$\lambda_{10} \equiv 0$, so the dead-cell compartment is decoupled from the optimization.*

Proof. Since $\partial^2 H/\partial u^2 = -2A_5 < 0$, the Hamiltonian is strictly concave in u and its maximizer over the compact interval $[0,1]$ is unique, which gives (i) and rules out (iii): a singular arc would require $\partial H/\partial u \equiv 0$ on an interval with $\partial^2 H/\partial u^2 = 0$, which cannot occur for $A_5 > 0$. Assertion (ii) follows because $z \mapsto \min\{1, \max\{0, z\}\}$ is 1-Lipschitz and λ_7 is C^1 by Equations (7)–(17). Assertion (iv) is immediate from $\dot{\lambda}_{10} = d\lambda_{10}$ with $\lambda_{10}(t_f) = 0$. □

Proposition 3 (Uniqueness for a short horizon and convergence of the sweep)**.** *Let $\mathcal{K}$ be the compact set of Proposition 1(iii) and let L_F, L_Λ denote Lipschitz constants for the right-hand sides of Equation (1) and of Equations (7)–(17) on $\mathcal{K}$. Define the sweep operator $\mathcal{T} : \mathcal{U} \to \mathcal{U}$ by*

$$\mathcal{T}(u) = \min\left\{1, \max\left\{0, \tfrac{\lambda_7[u]}{2A_5}\right\}\right\},$$

where $\lambda[u]$ is the adjoint associated with the state trajectory generated by u. Then there exists $C = C(\mathcal{K}, L_F, L_\Lambda) > 0$ such that

$$\|\mathcal{T}(u) - \mathcal{T}(v)\|_\infty \leq \frac{C}{2A_5}\, t_f\, e^{Ct_f}\, \|u - v\|_\infty, \qquad u, v \in \mathcal{U}.$$

Consequently, if t_f is small enough that $\kappa := Ct_f e^{Ct_f}/(2A_5) < 1$, then $\mathcal{T}$ is a contraction, the optimal control is unique, and the forward–backward sweep of Section 5.4 converges geometrically to it at rate κ from any admissible initialization. Under the damped update (Equation (19)), the corresponding rate is $|1 - \omega + \omega\kappa|$, which is less than one for every $\kappa < 1$ and every $\omega \in (0,1]$, and also for $\kappa \geq 1$ provided $\omega < 2/(1+\kappa)$.

Sketch. Standard Gronwall estimates on the forward system give $\|x[u] - x[v]\|_\infty \leq c_1 t_f e^{L_F t_f}\|u - v\|_\infty$, the delayed arguments contributing through the method of steps without altering the form of the bound because the histories coincide. A backward Gronwall

estimate on Equations (7)–(17), in which the advanced arguments are evaluated on already-determined portions of the solution, gives $\|\lambda[u] - \lambda[v]\|_\infty \le c_2 t_f e^{L_\Lambda t_f} \|x[u] - x[v]\|_\infty$. Composing, and using the fact that the projection in Equation (18) is 1-Lipschitz, yields the stated estimate with C depending on c_1, c_2, L_F, L_Λ. The contraction mapping theorem then gives uniqueness of the fixed point, which by Proposition 2(i) coincides with the optimal control; the rate for the damped iteration follows from $\mathcal{T}_\omega = (1-\omega)\mathrm{Id} + \omega\mathcal{T}$. □

Remark 4. *The smallness condition of Proposition 3 is sufficient but far from necessary, and we do not claim it is satisfied for the horizon $t_f = 10$ used in Section 6; this is the same limitation identified for the undelayed case in [17]. What the proposition provides is a single estimate governing both the well-posedness of the optimization and the behavior of the numerical scheme, together with the explicit role of the exposure weight A_5: increasing A_5 contracts the sweep. The convergence observed at the parameter values used is reported in Section 6.3 and is consistent with $\kappa < 1$.*

5.4. Numerical Computation of the Optimal Control

Because the control depends on both the state and adjoint variables, the optimality system is solved numerically by a forward–backward sweep iterative procedure. The delayed setting requires two modifications of the classical scheme: the forward solve must carry a history buffer, and the backward solve must access adjoint values at the advanced arguments $t + \tau_j$. The latter is the reason a fixed-step scheme on a stored grid is used rather than an adaptive solver, which cannot supply values at times it has not yet reached.

5.4.1. Discretization

The horizon $[0, t_f]$ is discretized with a uniform step h; delayed and advanced values are obtained from the stored trajectory by interpolation, the history being constant and equal to φ for $t \le 0$. Both sweeps use a classical fourth-order Runge–Kutta scheme. The baseline step is $h = 10^{-3}$ days, giving $N = 10^4$ steps; grid independence is documented in Section 6.

5.4.2. Algorithm

1. **Initialization.** Set $k = 0$ and choose $u^{(0)} \equiv 0.1$ on $[0, t_f]$:
2. **Forward sweep.** Using $u^{(k)}$, integrate state System (1) forward from $t = 0$ to $t = t_f$ with the prescribed history, storing $x^{(k)}$ at every node.
3. **Backward sweep.** Using $x^{(k)}$, integrate adjoint Systems (7)–(17) backward from $t = t_f$ subject to
$$\lambda_i(t_f) = 0, \qquad i = 1, \ldots, 11.$$
On the terminal window $(t_f - \tau_j, t_f]$ the corresponding advanced term is set to zero through χ_j; elsewhere $\lambda_8(t + \tau_2)$, $\lambda_9(t + \tau_3)$ and $\Phi(t + \tau_1)$ are read from nodes already computed in the current backward pass.
4. **Control update with damping.** Form the pointwise maximizer
$$\tilde{u}^{(k)}(t) = \min\left\{1, \max\left\{0, \frac{\lambda_7^{(k)}(t)}{2A_5}\right\}\right\}$$
and update by convex combination
$$u^{(k+1)} = \omega\, \tilde{u}^{(k)} + (1-\omega)\, u^{(k)}, \qquad \omega \in (0, 1]. \tag{19}$$
The undamped scheme corresponds to $\omega = 1$; we use $\omega = 0.5$ throughout.

5. **Stopping criterion.** Terminate when the relative sup-norm increment satisfies

$$\frac{\|u^{(k+1)} - u^{(k)}\|_\infty}{\max\{\|u^{(k+1)}\|_\infty, \epsilon_0\}} < \varepsilon, \tag{20}$$

with tolerance $\varepsilon = 10^{-6}$ and $\epsilon_0 = 10^{-12}$ guarding the degenerate case $u \equiv 0$.

The final iterate is taken as an approximation of the optimal control and is used in the simulations presented in Section 6.

5.4.3. Convergence

Convergence of the forward–backward sweep is not automatic: McAsey, Mou and Han [17] showed that the iteration is a contraction provided the right-hand sides and the control characterization are Lipschitz and the horizon is sufficiently short relative to the Lipschitz constants, and that divergence can occur otherwise. Proposition 3 gives the corresponding estimate for the present delayed problem, and Section 6 reports the observed convergence history together with a grid-refinement study.

5.4.4. A Numerical Detail

The immune death term $\delta x_{11} x_7/(c + x_7)$ has half-saturation $c = 7.14 \times 10^{-7}$, so its derivative with respect to x_7 is of order $\delta/c \approx 2 \times 10^3$ at $x_7 = 0$ and negligible thereafter. Since $x_7(0) = 0$ exactly, the adjoint right-hand side carries a removable spike at the initial node, which a fixed-step scheme cannot resolve; we therefore set $u^{(k+1)}(0) = u^{(k+1)}(h)$ at each iteration. The effect is confined to a single node and does not affect the reported results.

6. Numerical Results

6.1. Numerical Implementation

The simulations were performed in MATLAB R2016a using the fixed-step scheme described in Section 5.4. The state and adjoint systems were integrated with a classical fourth-order Runge–Kutta method on a uniform grid of step $h = 10^{-3}$ days, giving $N = 10^4$ steps on the horizon $[0, t_f]$ with $t_f = 10$ days; delayed and advanced arguments were obtained from the stored trajectory, the initial history being taken constant on $[-\tau_{\max}, 0]$. The forward–backward sweep used the damped update (Equation (19)) with $\omega = 0.5$ and terminated at the relative tolerance $\varepsilon = 10^{-6}$ after 23 iterations. Model parameters are those of Table 8; the initial history was

$$\varphi \equiv (1.0,\ 0.1,\ 0.5,\ 0.1,\ 0.3,\ 0.2,\ 0,\ 0.1,\ 1.5,\ 0,\ 0.5),$$

representing a biologically plausible immune and leukemic state in the absence of treatment.

Comparison Scenarios

A comparison against an untreated baseline and a single low constant dose does not by itself establish that the *shape* of the optimal schedule is responsible for any observed improvement, since the optimal schedule delivers a larger cumulative dose than a low constant infusion. Writing

$$E(u) := \int_0^{t_f} u(t)\, dt$$

for the cumulative normalized drug exposure, the computed optimal control satisfies $E(u^*) = 4.4596$, so that a constant infusion delivering the same total drug is

$$\bar{u} = \frac{E(u^*)}{t_f} = 0.4460.$$

This is more than four times the constant dose $u = 0.1$ considered previously. We therefore compare four strategies:

- **No treatment:** $u(t) = 0$, $E = 0$;
- **Low constant dose:** $u(t) = 0.1$, $E = 1.0$;
- **Exposure-matched constant dose:** $u(t) = \bar{u} = 0.4460$, $E = 4.4596$;
- **Optimal control:** $u^*(t)$, $E = 4.4596$.

The third scenario isolates the contribution of dose *timing*: any difference between it and the optimal schedule cannot be attributed to a difference in total drug administered.

Figure 1 shows the four dosing profiles. The optimal control applies the maximal admissible dose $u^* = 1$ during the initial phase, up to approximately day 3, and then decreases smoothly toward zero, becoming negligible after about day 8. This bang-then-taper structure reflects the optimality condition (Equation (18)): the control saturates on the interval where $\lambda_7(t) \geq 2A_5$, and the subsequent taper follows the decay of λ_7 toward its zero terminal value. The exposure-matched protocol distributes the same total drug uniformly at $\bar{u} = 0.4460$.

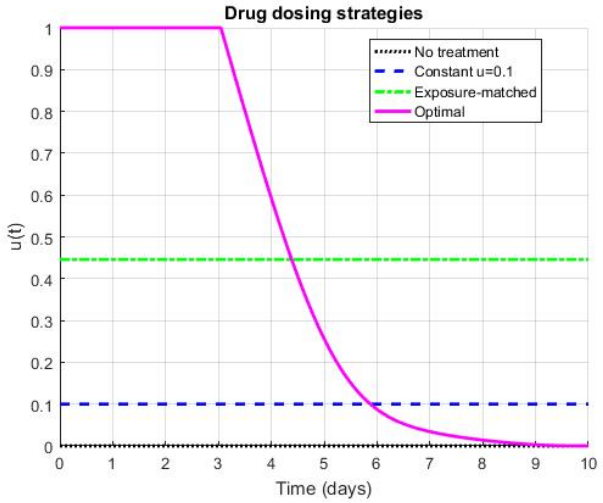


Figure 1. Comparison of drug dosing strategies. The optimal control and the exposure-matched constant dose deliver identical cumulative exposure $E = 4.4596$.

Figure 2 shows the Th1 effector population. Under optimal control the Th1 cells remain low during the first days and then expand rapidly, reaching 7.011 at the end of the horizon; the exposure-matched constant dose produces a similar terminal level of 6.803 but with the expansion displaced later in time. The low constant dose produces only a modest late increase to 2.073, and the untreated case remains essentially flat at 0.286. The distinguishing feature of the optimal schedule is therefore not the terminal Th1 level, which the exposure-matched protocol nearly attains, but the timing of the expansion.

Figure 3 shows the Th2 population. In absolute terms the Th2 level rises modestly under all four scenarios, from 0.5 to between 0.527 and 0.599, and is in fact slightly highest under optimal control. The therapeutic benefit therefore lies not in suppressing Th2 in absolute terms but in the *relative* balance between Th1 and Th2. The terminal ratio x_2/x_3 is 0.54 without treatment, 3.60 under the low constant dose, 11.37 under the exposure-matched dose and 11.71 under optimal control. More informative than the terminal ratio is the time at which Th1 overtakes Th2: this occurs at day 3.74 under optimal control against

day 4.61 under the exposure-matched protocol and day 6.96 under the low constant dose, and never within the horizon in the untreated case.

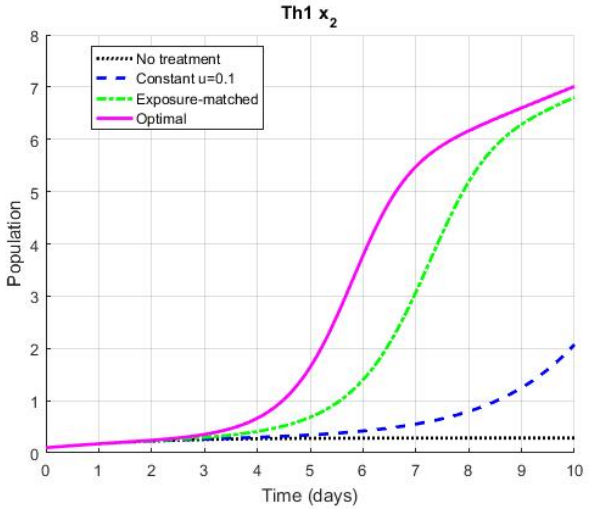


Figure 2. Effector T cells $x_2(t)$ (Th1) under the four dosing strategies.

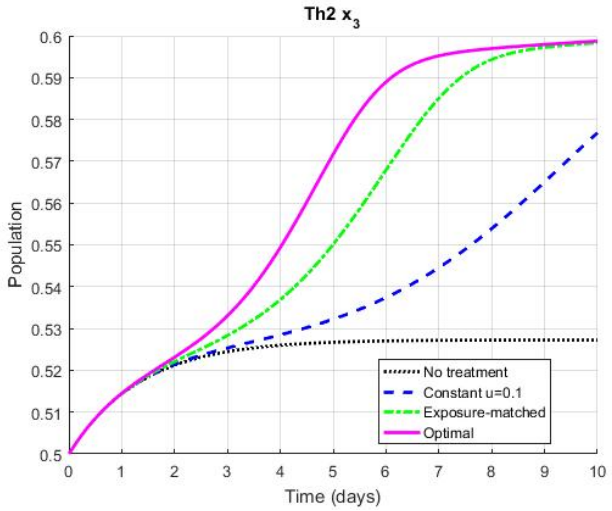


Figure 3. Suppressor T cells $x_3(t)$ (Th2) under the four dosing strategies.

Figure 4 shows the leukemic burden. Under optimal control the population decreases from its initial value of 1.5 to 0.0267, a reduction of 98.3% relative to the untreated terminal value of 1.56; the exposure-matched constant dose reaches 0.0601, a reduction of 96.1%, while the low constant dose reaches only 0.4474 and the untreated case shows slight continued progression. Thus, at equal cumulative exposure, the optimal schedule attains a terminal leukemic burden lower by a factor of 2.25, but the larger part of the improvement over the low constant dose is attributable to total exposure rather than to scheduling.

The separation between the two exposure-matched strategies is more pronounced in the rate of response than in the endpoint. Writing t_{50} and t_{90} for the times at which the leukemic burden falls below 0.75 and 0.15, corresponding to reductions of 50% and 90% from its initial value, the optimal schedule attains $t_{50} = 2.74$ days and $t_{90} = 6.43$ days, against $t_{50} = 3.78$ and $t_{90} = 8.05$ days for the exposure-matched constant dose—reductions of 28% and 20%, respectively. The low constant dose reaches $t_{50} = 7.36$ days and does not achieve a 90% reduction within the horizon.

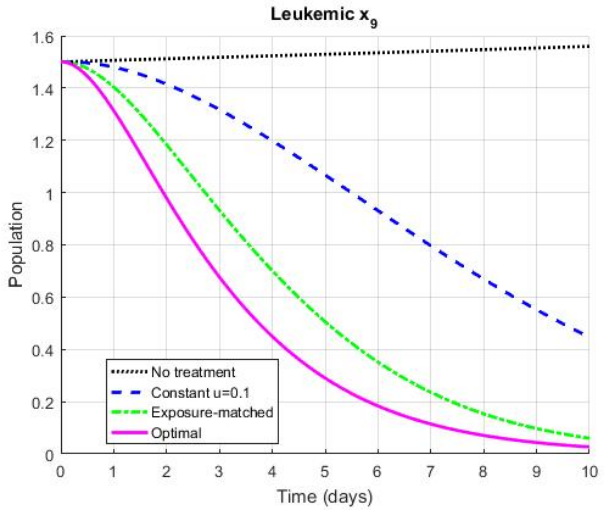


Figure 4. Leukemic burden $x_9(t)$ under the four dosing strategies.

Figure 5 shows the immune effector population. The four scenarios produce nearly indistinguishable trajectories, the population growing steadily from 0.5 to approximately 1.48 in all cases, with the untreated case marginally highest. As discussed in the limitations, this reflects the fact that in the present parameter regime the interaction coefficients c_1, c_2 and ρ are several orders of magnitude smaller than the supply and decay rates s and m, so that the effector equation is dominated by its intrinsic supply–decay balance. The practical implication is favorable but must be stated precisely: the optimal strategy achieves substantial leukemic reduction without measurable loss of immune effector competence, and the reduction is attributable to the direct cytotoxic action of the drug rather than to enhanced immune-mediated clearance.

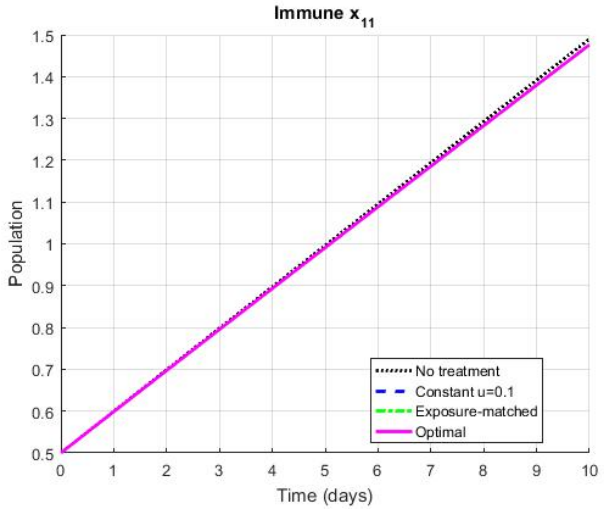


Figure 5. Immune effector population $x_{11}(t)$ under the four dosing strategies. The four trajectories are nearly coincident and therefore overlap; this near-identity across dosing strategies is the point of the figure and is discussed in the text.

Table 1 summarizes the comparison, and read together with the exposure-matched arm it supports a more precise statement of the contribution than a comparison against an untreated baseline alone would allow.

The objective value provides the primary ordering: J increases from -8.03 under no treatment to -0.65 under the low constant dose, 17.18 under the exposure-matched constant dose, and 24.15 under optimal control. The last two are the informative pair, since they are attained at identical cumulative exposure: the optimal schedule improves the objective by

40.5% over a constant infusion delivering the same total drug. This ordering is structural rather than empirical—u^* maximizes Equation (3) by construction, so no admissible control can exceed it—and it therefore serves as a consistency check on the computed solution as much as a comparison between protocols. What the exposure-matched arm establishes is the *magnitude* of the gap: redistributing a fixed quantity of drug over time is worth some forty percent in the quantity the problem was posed to maximize.

Table 1. Quantitative comparison of the four dosing strategies. $E(u)$ is the cumulative normalized drug exposure; t_{cross} is the first time at which x_2 exceeds x_3; t_{50} and t_{90} are the times at which the leukemic burden falls below 0.75 and 0.15, corresponding to reductions of 50% and 90% from the initial value; $D_{\text{tot}}^{\min}$ is the minimum of the combined therapeutic divergence and $t_{\min}$ the time at which it is attained; J is the value of the objective functional (Equation (3)) at $A_1 = \cdots = A_5 = 1$.

Strategy	$E(u)$	$x_9(t_f)$	$x_2(t_f)$	$x_3(t_f)$	x_2/x_3	t_{cross}	t_{50}	t_{90}	$D_{\text{tot}}(t_f)$	$D_{\text{tot}}^{\min}$ $(t_{\min})$	J
No treatment	0	1.560	0.286	0.527	0.54	—	—	—	1.292	1.292 (10.00)	−8.0289
Constant $u = 0.1$	1.000	0.4474	2.073	0.577	3.60	6.96	7.36	—	0.2661	0.2661 (10.00)	−0.6452
Exposure-matched $\bar{u}$	4.4596	0.0601	6.803	0.598	11.37	4.61	3.78	8.05	0.2924	0.2032 (7.19)	17.1836
Optimal u^*	4.4596	0.0267	7.011	0.599	11.71	3.74	2.74	6.43	0.3327	**0.1942** (5.82)	**24.1465**

Bold marks the best value attained across the four strategies.

The population-level measures show where that gap originates. At equal cumulative exposure the optimal schedule attains a lower terminal leukemic burden (0.0267 against 0.0601, a factor of 2.25), a marginally higher terminal Th1/Th2 ratio, and a lower minimum combined divergence; but the differences in terminal quantities are modest, and the exposure-matched protocol reaches a broadly comparable final configuration. The consistent separation is in timing. Th1 dominance is established 0.87 days sooner, a 50% leukemic reduction 1.04 days sooner, a 90% reduction 1.62 days sooner, and the minimum of the combined divergence 1.37 days sooner. The benefit of adaptive scheduling in this model is therefore principally one of *rate of response at fixed drug exposure* rather than of eventual position.

One quantity runs the other way and should be recorded explicitly. The terminal combined divergence is higher under optimal control (0.3327) than under the exposure-matched constant dose (0.2924). This is a consequence of the free terminal state: as λ_7 approaches its zero terminal value, the optimized infusion is withdrawn, and the leukemic population begins to recover after approximately day 6, whereas the constant infusion continues to act until t_f. The minimum of the divergence, attained at day 5.82 under optimal control against day 7.19 under the exposure-matched dose, is both lower and earlier, and Section 6.5 shows that both orderings—optimal lower at the minimum, higher at t_f—persist across twenty alternative reference profiles. The effect is thus a property of the formulation rather than of the diagnostic: a terminal cost on the leukemic population would remove it, and we regard this as a natural extension rather than a defect of the present formulation.

6.2. Sensitivity Analysis

Because the model parameters are drawn from heterogeneous sources and several were calibrated rather than measured (Section 8), the conclusions must be shown to be robust to plausible variation. We performed a one-at-a-time study in four groups: kinetic parameters, delays, initial leukemic load, and objective weights. For each perturbed configuration the full optimality system was re-solved to the tolerance of Section 5.4, and five outcomes were recorded: the terminal leukemic burden $x_9(t_f)$, the cumulative exposure $E(u^*)$, the minimum combined divergence $D_{\text{tot}}^{\min}$, and the times t_{50}, t_{90} at which

the leukemic burden falls below 0.75 and 0.15, respectively. For the first three, we report the normalized sensitivity index

$$S_{\pi}^{Q} \approx \frac{Q(\pi(1+\epsilon)) - Q(\pi(1-\epsilon))}{2\epsilon\, Q(\pi)}, \qquad \epsilon = 0.2,$$

so that $|S| > 1$ indicates amplification of the relative perturbation; for t_{50} and t_{90} we report the absolute shift in days, which is more directly interpretable.

6.2.1. Kinetic Parameters

Table 2 reports the sixteen kinetic parameters examined. The results are strongly concentrated: a single parameter dominates every outcome.

Table 2. Normalized sensitivity indices under ±20% perturbation of the kinetic parameters, with absolute shifts in t_{50} and t_{90} in days. Parameters for which all indices vanish to four decimal places are grouped in the final row.

Parameter	$S^{x_9(t_f)}$	$S^{E(u^*)}$	$S^{D_{\text{tot}}^{\min}}$	Δt_{50} (d)	Δt_{90} (d)
μ_L	**−4.5263**	−0.0239	−0.0839	−0.380	−1.016
β_L	0.3164	0.0030	0.0081	0.036	0.105
γ_L	−0.1059	−0.0010	−0.0022	−0.015	−0.028
γ_D	0.0775	0.0231	0.0019	0.003	0.012
v	0.0654	−0.0867	−0.0980	0.000	0.006
η_r	−0.0009	0.0006	0.2122	0.000	0.000
κ	0.0008	−0.0006	−0.0470	0.000	0.000
ϕ	0.0015	−0.0008	0.0256	0.000	0.000
s	≈0	≈0	−0.0210	0.000	0.000
μ_D	0.0033	0.0015	0.0002	0.000	0.000
θ	0.0001	0.0000	0.0000	0.000	0.000
μ_r	−0.0009	0.0012	0.0002	0.000	0.000
c_1, c_2, ρ, m, δ	≈0	≈0	≤0.0002	0.000	0.000

Bold marks the only index whose magnitude exceeds unity, indicating amplification of the relative perturbation.

Three features of Table 2 deserve comment.

First, the drug-induced leukemic death rate μ_L is the only parameter whose perturbation is amplified: $S_{\mu_L}^{x_9(t_f)} = -4.53$, so a 20% uncertainty in μ_L produces roughly a 90% change in terminal leukemic burden, and shifts t_{90} by about a day. This is expected, since μ_L governs the direct cytotoxic term through which the control acts, and it identifies μ_L as the parameter most in need of experimental determination. Notably, the effect on $E(u^*)$ is two orders of magnitude smaller (-0.024): the *schedule* is largely insensitive to μ_L even though the *outcome* is not.

Second, all remaining parameters have $|S| < 0.32$, and most are below 0.01. In particular the parameters governing the immune–leukemic interaction, c_1, c_2, ρ, m and δ, have sensitivity indices indistinguishable from zero across all five outcomes. This provides quantitative confirmation of the structural observation recorded in the limitations: in the present parameter regime the immune effector compartment is dynamically decoupled from the leukemic dynamics, and the reported therapeutic effect is attributable entirely to the direct cytotoxic action of the drug.

Third, θ has $S \approx 10^{-4}$ across all outcomes. Since $\theta = 100$ while the leukemic burden never exceeds 1.6, the Hill factor $\left(1 + (x_9/\theta)^2\right)^{-1}$ appearing in Equation (2) remains within 3×10^{-4} of unity throughout the simulation, so the saturating nonlinearity in the self-renewal flux is inactive in this regime and $\Psi(x_9)$ reduces to $\beta_L x_9^2$ to four significant figures. The leukemic equation therefore operates in its effectively quadratic range, and the results

reported here do not probe the density-dependence that Ψ was introduced to represent; this should be borne in mind when interpreting the model.

A separate observation concerns the allergy axis. The Treg inhibition rate η_r has a negligible effect on the leukemic outcomes but the largest effect of any kinetic parameter on the minimum combined divergence, $S_{\eta_r}^{D^{\min}} = 0.212$, with κ and ϕ next. These are precisely the parameters governing the Th1/Th2/Treg balance, so the two therapeutic axes are sensitive to disjoint sets of parameters—a structural separation that supports treating them as distinct objectives.

6.2.2. Delays

Table 3 reports the sensitivity of the five outcomes to the three delays. The leukemic cell-cycle delay τ_3 is the only influential delay, with $S_{\tau_3}^{x_9(t_f)} = 0.45$; this is unsurprising, since τ_3 enters the leukemic equation both through the shifted argument and through the amplification factor $2e^{-\gamma_L \tau_3}$. The transport and cytokine delays τ_1 and τ_2 have indices below 0.01 on every outcome.

Table 3. Sensitivity to the three delays under $\pm 20\%$ perturbation.

Delay	$S^{x_9(t_f)}$	$S^{E(u^*)}$	$S^{D^{\min}_{\text{tot}}}$	Δt_{50} (d)	Δt_{90} (d)
τ_1	0.0056	-0.0061	0.0010	0.000	0.000
τ_2	0.0000	-0.0000	-0.0008	0.000	0.000
τ_3	0.4500	0.0030	0.0086	-0.001	0.121

6.2.3. Initial Leukemic Load

Table 4 reports the dependence of the computed schedule on the initial leukemic burden. Across a sixfold variation in initial disease burden, the bang-then-taper structure of the optimal control is preserved, and the switching time varies only between 2.97 and 3.18 days—a range of 7% for a 500% change in initial condition. The cumulative exposure likewise varies by less than 3%. The optimal *schedule* is therefore remarkably insensitive to the initial leukemic load, while the terminal burden scales approximately linearly with it. The minimum combined divergence is non-monotonic, attaining its lowest value at the baseline load $x_9(0) = 1.5$; this reflects the two axes trading off, since a smaller initial burden leaves the allergy axis relatively further from its target.

Table 4. Dependence on the initial leukemic burden $x_9(0)$. The switching time is the instant at which the optimal control leaves saturation. Note that t_{50} and t_{90} are times to reach the *absolute* levels 0.75 and 0.15, not relative reductions, so the first row reflects an initial state already below the first threshold.

$x_9(0)$	$x_9(t_f)$	$E(u^*)$	$D^{\min}_{\text{tot}}$	t_{50} (d)	Switching Time (d)
0.50	0.00719	4.4119	0.2366	0	2.97
1.00	0.01603	4.4345	0.2001	1.538	3.01
1.50	0.02670	4.4596	0.1942	2.734	3.05
2.00	0.03936	4.4872	0.2014	3.554	3.09
3.00	0.07148	4.5492	0.2149	4.798	3.18

6.2.4. Objective Weights

Table 5 reports the effect of the four objective weights. The weights separate sharply into two groups. The exposure weight A_5 and the Th1 reward A_2 substantially reshape the schedule: increasing A_5 from 0.5 to 5 reduces cumulative exposure from 5.05 to 2.79 and shortens the saturated phase from four days to nothing, the control ceasing to saturate at

all at $A_5 = 5$; increasing A_2 over the same range has the opposite effect, raising exposure from 3.84 to 5.68. The behavior of A_5 is the one predicted directly by optimality Condition (18), in which A_5 rescales the threshold $\lambda_7 \geq 2A_5$ at which the control saturates, and its appearance in the numerics is a useful consistency check on the optimality system.

Table 5. Effect of the objective weights. All weights other than the one listed are held at unity. A switching time of zero indicates that the control never saturates.

Weight	Value	$x_9(t_f)$	$E(u^*)$	$D_{\text{tot}}^{\min}$	t_{50} (d)	Switching Time (d)
A_5	0.5	0.02385	5.0461	0.19367	2.734	4.00
	1	0.02670	4.4596	0.19422	2.734	3.05
	2	0.03243	3.7704	0.19580	2.748	1.50
	5	0.05465	2.7865	0.20079	3.130	0
A_2	0.5	0.03170	3.8350	0.19558	2.742	1.76
	1	0.02670	4.4596	0.19422	2.734	3.05
	2	0.02395	5.0184	0.19368	2.734	3.98
	5	0.02183	5.6759	0.19346	2.734	4.82
A_4	0.5	0.02687	4.4297	0.19426	2.734	3.00
	1	0.02670	4.4596	0.19422	2.734	3.05
	2	0.02638	4.5164	0.19416	2.734	3.14
	5	0.02560	4.6675	0.19401	2.734	3.36
A_3	0.5	0.02668	4.4616	0.19422	2.734	3.05
	1	0.02670	4.4596	0.19422	2.734	3.05
	2	0.02672	4.4556	0.19423	2.734	3.04
	5	0.02680	4.4433	0.19425	2.734	3.00

By contrast, the leukemic penalty A_4 and the Th2 penalty A_3 have little effect. A tenfold increase in A_4 changes the terminal leukemic burden by only 4%, and a tenfold increase in A_3 changes it in the fourth decimal place; A_3 is essentially inert because, as Figure 3 shows, the Th2 population varies by less than 15% across all strategies, so the term $-A_3 x_3$ contributes almost no gradient to the optimization. This is worth stating plainly: the Th2 penalty, though motivated by the allergy objective that partly defines the problem, does not materially influence the computed schedule in this parameter regime.

Finally, t_{50} is invariant at 2.734 days across every weight configuration except $A_5 = 5$. The rate of early response is therefore a property of the dynamics and the saturation constraint rather than of the particular trade-off encoded in the weights, which reinforces the interpretation of the timing advantage reported in Table 1 as a robust feature of the optimal schedule.

6.2.5. Summary

Across all perturbations considered, the qualitative structure of the optimal schedule—an initial saturated phase of approximately three days followed by a smooth taper—and the ordering of the four dosing strategies with respect to every outcome measure were preserved, the single exception being $A_5 = 5$, at which the exposure penalty is high enough that saturation never occurs. Quantitatively, only μ_L amplifies its relative perturbation, and only μ_L, β_L, γ_L and τ_3 have indices above 0.1 on any outcome. The calibrated parameters β_L, γ_L and θ discussed in Section 8 have indices of 0.32, -0.11 and 10^{-4}, respectively, so the conclusions are not sensitive to the calibration choices at the level of a $\pm 20\%$ uncertainty.

6.3. Convergence of the Numerical Scheme

Figure 6 reports the convergence history. The control increment decreases from 9.0×10^{-1} at the first iteration to 5.6×10^{-7} at the twenty-third, at which point the relative tolerance $\varepsilon = 10^{-6}$ was met. The decay is linear on the logarithmic axis over six orders of magnitude, with an observed contraction factor of approximately 0.55 per iteration. This is consistent with the estimate of Proposition 3: for the damped iteration the predicted rate is $|1 - \omega + \omega\kappa|$, so that the observed value corresponds to $\kappa \approx 0.1$ at the parameter values used, comfortably inside the contraction regime. We emphasize that this is empirical verification at one parameter set and not a proof that the sufficient condition of Proposition 3 holds at $t_f = 10$.

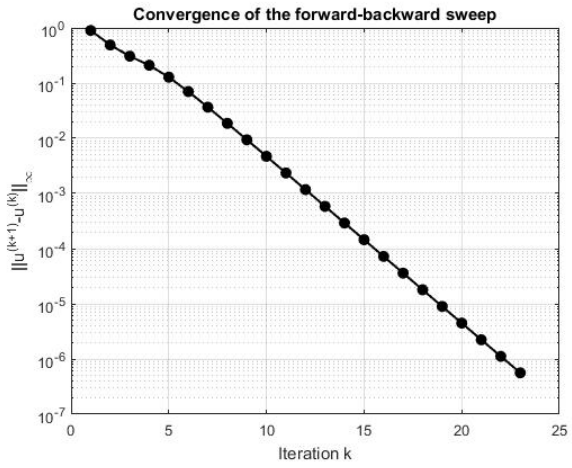


Figure 6. Convergence of the forward–backward sweep: sup-norm of the control increment $\|u^{(k+1)} - u^{(k)}\|_\infty$ against iteration index k, logarithmic vertical axis, at the baseline parameter set with damping $\omega = 0.5$.

Table 6 reports the effect of successive grid refinement. The terminal leukemic burden and the cumulative exposure agree to six significant figures across a fourfold change in step size, and the sup-norm difference between the controls computed on successive grids falls from 1.43×10^{-4} to 2.71×10^{-7}; the corresponding L^2 differences are 9.66×10^{-6} and 2.92×10^{-7}, a reduction consistent with at least fourth-order behavior. The larger difference at the coarsest refinement is localized at the instant the control leaves saturation, where u^* has a corner whose resolution improves rapidly with the mesh; away from it the agreement is uniform. All quantities reported in this paper are therefore insensitive to the discretization at the working step $h = 10^{-3}$.

Table 6. Grid independence of the computed optimal control. Differences are measured against the next coarser grid on the interior interval $[0.5, 9.5]$, after linear interpolation onto a common mesh.

N	h	Iterations	$x_9(t_f)$	$E(u^*)$	$\|u_h - u_{h/2}\|_\infty$
2500	4×10^{-3}	23	0.0266961	4.45963	—
5000	2×10^{-3}	23	0.0266962	4.45962	1.43×10^{-4}
10,000	1×10^{-3}	23	0.0266962	4.45962	2.71×10^{-7}

A further observation concerns the iteration count, which is 23 on all three grids. The number of sweeps required to meet Tolerance (20) is thus independent of the discretization, as the estimate of Proposition 3 predicts: the contraction factor is a property of the continuous optimality system and of the damping parameter, not of the mesh on which it is approximated.

6.4. Information-Theoretic Characterization of the Therapeutic Goal

The therapeutic goal admits a natural information-theoretic description, which we develop in this section. Before introducing the relevant quantities, we make their status precise since it bears directly on what the control problem is taken to optimize. Objective Functional (3) aggregates four populations measured in different units, and its numerical value is therefore interpretable only relative to a stated normalization (Remark 2). A divergence between normalized distributions, by contrast, is dimensionless, invariant under rescaling of the underlying populations, and directly comparable across strategies and across the two therapeutic axes. In the present section and in Sections 6.5 and 6.6 these quantities are evaluated *along* the trajectories generated by each dosing strategy and serve as diagnostics; the control being diagnosed is the one determined by Equation (3), in which they play no part. Section 6.7 then treats the complementary problem in which the combined divergence is itself the running cost, and compares the two resulting schedules.

The treatment studied here pursues two competing objectives that can each be expressed as a balance between cell populations. The first is the control of treatment-induced allergic reactions, governed by the adaptive T-helper compartment: a Th1-dominant configuration is associated with low hypersensitivity risk, whereas a shift toward Th2 dominance promotes allergic inflammation, with regulatory T cells (Treg) modulating this balance. The second is the control of the disease itself, governed by the relative abundance of immune effector cells and leukemic cells: an immune-dominant configuration corresponds to effective tumor suppression. Both balances admit a natural information-theoretic formulation, which we develop on two distinct axes.

6.4.1. Allergy Axis

Restricting attention to the adaptive T-helper compartment $\{x_2, x_3, x_4\}$ (Th1, Th2, Treg), we define the normalized subpopulation distribution

$$p_i^A(t) = \frac{x_i(t)}{\sum_{j\in\mathcal{J}_A} x_j(t)}, \qquad \mathcal{J}_A = \{2,3,4\}, \quad i \in \mathcal{J}_A,$$

with Shannon entropy

$$S_A(t) = -\sum_{i\in\mathcal{J}_A} p_i^A(t)\,\ln p_i^A(t), \qquad 0 \le S_A(t) \le \ln 3,$$

which quantifies the diversity and regulatory order of the helper T-cell response. To measure proximity to an allergy-protective profile, we introduce a reference distribution $q_A = (0.70,\ 0.10,\ 0.20)$ and evaluate the relative entropy

$$D_A(t) = D_{\mathrm{KL}}\Big(p^A(t)\,\|\,q_A\Big) = \sum_{i\in\mathcal{J}_A} p_i^A(t)\,\ln\frac{p_i^A(t)}{q_{A,i}}.$$

The three components of q_A encode Th1 predominance, since a Th1-skewed helper compartment is associated with low hypersensitivity risk [9,10]; a small but nonzero Th2 fraction, since complete Th2 depletion is neither physiological nor a therapeutic goal; and an intermediate regulatory fraction, since Treg cells must remain sufficient to restrain the response without suppressing anti-tumor immunity. These are qualitative immunological targets rather than calibrated values, and Section 6.5 verifies that the conclusions do not depend sensitively on them.

6.4.2. Leukemia Axis

Analogously, restricting attention to the disease-control pair $\{x_{11}, x_9\}$ (immune effector and leukemic cells), we define

$$p_i^L(t) = \frac{x_i(t)}{x_{11}(t) + x_9(t)}, \qquad i \in \mathcal{J}_L = \{11, 9\},$$

and evaluate the relative entropy to an immune-dominant reference $q_L = (0.85, 0.15)$, in the order (immune, leukemic),

$$D_L(t) = D_{\mathrm{KL}}\Big(p^L(t) \,\|\, q_L\Big) = \sum_{i \in \mathcal{J}_L} p_i^L(t) \, \ln \frac{p_i^L(t)}{q_{L,i}}.$$

This reference encodes immune predominance over a residual leukemic fraction rather than complete eradication, which is not attainable within the horizon considered.

6.4.3. Combined Therapeutic Divergence

Each divergence is nonnegative and vanishes only when the corresponding balance coincides with its healthy target. To summarize both goals in a single quantity we define the combined therapeutic divergence

$$D_{\mathrm{tot}}(t) = D_A(t) + D_L(t),$$

which measures the total information-theoretic distance of the immune–leukemic state from a simultaneously allergy-protective and disease-controlling configuration.

The optimal control $u^*(t)$ derived in Section 5.4 does not act on D_{tot} directly, since the divergences do not appear in Equation (3); what follows is therefore an assessment of how far the schedule obtained from the population-based objective happens to advance the information-theoretic one. The entropy $S_A(t)$ is examined separately in Section 6.6, where it is shown to carry information that the divergence alone does not.

Figure 7 reports the two axis divergences and their sum for the four dosing strategies. On the allergy axis (left), all cases begin from a large divergence, $D_A \approx 1.13$, reflecting an initially mixed helper T-cell configuration with elevated allergic potential. The optimal control produces the fastest decrease, falling to approximately 0.23 by day 5 and remaining close to that level thereafter; the exposure-matched constant dose follows the same descent displaced by roughly a day, the low constant dose declines gradually and reaches a comparable value only at the end of the horizon, and the untreated trajectory initially decreases but subsequently drifts *away* from the target, rising to 0.93 as the disease persists. It is worth noting that D_A does not approach zero under any strategy: as discussed in Section 6.6, the treated trajectories overshoot the Th1 fraction prescribed by q_A rather than converging to it.

On the leukemia axis (center), all cases start at $D_L \approx 0.91$. The optimal control drives the immune/leukemic balance to the target fastest, with D_L reaching a minimum near zero at approximately day 6.3; thereafter, the divergence increases again, to about 0.10 at t_f, as the tapering infusion permits partial leukemic recovery. The exposure-matched constant dose reaches its minimum about a day and a half later but, because it continues to deliver drug until t_f, remains close to the target at the terminal time. The low constant dose declines steadily throughout, and the untreated case remains far from the target, consistent with continued tumor progression.

The combined divergence (right) inherits both features, and the comparison between the two exposure-matched strategies depends on which instant is examined. The optimal schedule attains the lower *minimum*, $D_{\mathrm{tot}}^{\min} = 0.194$ at day 5.82, against 0.203 at day 7.19 for

the exposure-matched constant dose: lower, and reached 1.37 days earlier. At the *terminal* time, the ordering reverses, $D_{\text{tot}}(t_f) = 0.333$ under optimal control against 0.292 under the exposure-matched dose, for the reason given above. Both orderings are stable under variation of the reference profiles, as documented in Section 6.5. The untreated case remains highest throughout by a wide margin, ending at 1.292.

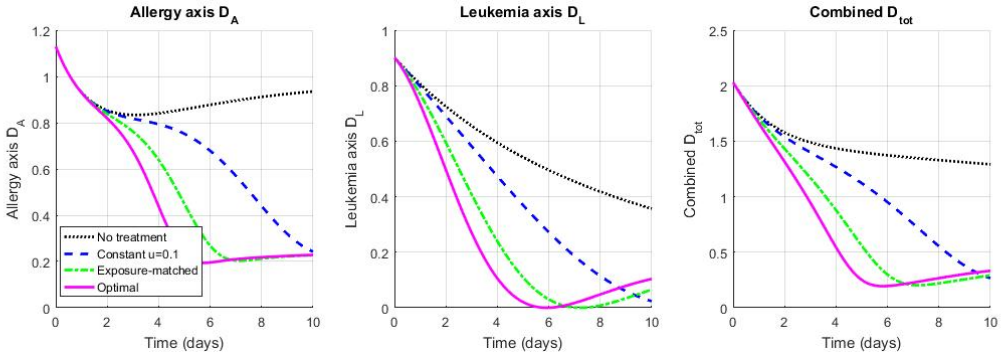


Figure 7. Information-theoretic diagnostics of the therapeutic goal for the four dosing strategies. **Left**: relative entropy $D_A(t)$ on the allergy axis (Th1/Th2/Treg distribution to a Th1-favorable reference q_A). **Center**: relative entropy $D_L(t)$ on the leukemia axis (immune/leukemic distribution to an immune-dominant reference q_L). **Right**: combined therapeutic divergence $D_{\text{tot}}(t) = D_A(t) + D_L(t)$.

These trajectories admit a direct biological interpretation, and one qualification. The optimal control, by adjusting the infusion rate in response to the evolving state, steers the helper T-cell distribution toward a Th1-favorable configuration and the immune/leukemic balance toward immune dominance more rapidly than any of the alternatives, and it does so at cumulative drug exposure identical to that of the exposure-matched protocol. The qualification is that this advantage is one of rate rather than of eventual position: given the same total drug, a constant infusion reaches a comparable configuration roughly a day and a half later, and—because it is still acting at t_f—holds it slightly better at the end of a fixed horizon. The diagnostics therefore support the conclusion that adaptive dosing advances both therapeutic objectives faster, but not the stronger conclusion that it advances them further within this model and horizon.

6.5. Robustness with Respect to the Reference Profiles

Since q_A and q_L are prescribed rather than estimated, the comparison between dosing strategies must be shown not to depend on the particular choice. We recomputed D_A, D_L and D_{tot} along all four trajectories for every combination of

$$q_A \in \{(0.70, 0.10, 0.20),\ (0.60, 0.20, 0.20),\ (0.80, 0.10, 0.10),\ (0.50, 0.20, 0.30),\ (0.85, 0.05, 0.10)\}$$

$$q_L \in \{(0.85, 0.15),\ (0.75, 0.25),\ (0.90, 0.10),\ (0.95, 0.05)\},$$

a total of twenty reference pairs spanning Th1 fractions from 0.50 to 0.85 and immune fractions from 0.75 to 0.95.

One invariance can be established analytically. On the leukemia axis the divergence

$$D_L(p^L \,\|\, q_L(\vartheta)) = p^L_{11} \ln \frac{p^L_{11}}{\vartheta} + p^L_{9} \ln \frac{p^L_{9}}{1-\vartheta}$$

is, for each fixed t, a strictly decreasing function of the immune fraction p^L_{11} on the range $p^L_{11} < \vartheta$. Consequently any reference whose immune fraction exceeds those attained along the simulated trajectories ranks the strategies in the same order, and only a reference representing an *unhealthy* target could reverse it. The numerical sweep confirms that this

stability extends to the combined divergence. Table 7 reports the minimum combined therapeutic divergence attained by each dosing strategy under all twenty reference pairs.

Table 7. Minimum combined therapeutic divergence $D_{\text{tot}}^{\min}$ over $[0, t_f]$ for each dosing strategy under twenty reference pairs; the time at which the minimum is attained is given in parentheses for the two exposure-matched strategies. The optimal schedule attains the lowest minimum in all twenty cases. Bold marks the lowest minimum divergence in each row, attained by the optimal schedule in all twenty cases.

q_A	q_L	No Treatment	Constant $u = 0.1$	Exposure-Matched	Optimal	Ordering
(0.70, 0.10, 0.20)	(0.85, 0.15)	1.292	0.2661	0.2032 (7.2)	**0.1942** (5.8)	preserved
(0.70, 0.10, 0.20)	(0.75, 0.25)	1.092	0.2434	0.2213 (6.8)	**0.2100** (5.5)	preserved
(0.70, 0.10, 0.20)	(0.90, 0.10)	1.471	0.3166	0.2107 (7.7)	**0.2031** (6.3)	preserved
(0.70, 0.10, 0.20)	(0.95, 0.05)	1.799	0.4363	0.2240 (9.2)	**0.2175** (7.7)	preserved
(0.60, 0.20, 0.20)	(0.85, 0.15)	0.8989	0.2360	0.2000 (6.3)	**0.1863** (5.1)	preserved
(0.60, 0.20, 0.20)	(0.75, 0.25)	0.6989	0.2057	0.1790 (6.0)	**0.1635** (4.8)	preserved
(0.60, 0.20, 0.20)	(0.90, 0.10)	1.078	0.2864	0.2371 (6.6)	**0.2253** (5.3)	preserved
(0.60, 0.20, 0.20)	(0.95, 0.05)	1.406	0.4061	0.3040 (8.6)	**0.2932** (7.2)	preserved
(0.80, 0.10, 0.10)	(0.85, 0.15)	1.250	0.1636	0.0914 (7.3)	**0.0849** (6.0)	preserved
(0.80, 0.10, 0.10)	(0.75, 0.25)	1.049	0.1408	0.1158 (6.9)	**0.1078** (5.6)	preserved
(0.80, 0.10, 0.10)	(0.90, 0.10)	1.429	0.2140	0.0942 (7.9)	**0.0886** (6.5)	preserved
(0.80, 0.10, 0.10)	(0.95, 0.05)	1.757	0.3338	0.1028 (9.3)	**0.0979** (7.8)	preserved
(0.50, 0.20, 0.30)	(0.85, 0.15)	0.9600	0.3757	0.3310 (6.1)	**0.3129** (4.9)	preserved
(0.50, 0.20, 0.30)	(0.75, 0.25)	0.7599	0.3349	0.2982 (5.8)	**0.2780** (4.6)	preserved
(0.50, 0.20, 0.30)	(0.90, 0.10)	1.139	0.4277	0.3763 (6.4)	**0.3605** (5.1)	preserved
(0.50, 0.20, 0.30)	(0.95, 0.05)	1.467	0.5474	0.4630 (7.1)	**0.4517** (5.7)	preserved
(0.85, 0.05, 0.10)	(0.85, 0.15)	1.675	0.2668	0.1188 (7.9)	**0.1155** (6.5)	preserved
(0.85, 0.05, 0.10)	(0.75, 0.25)	1.475	0.2440	0.1642 (7.4)	**0.1591** (6.0)	preserved
(0.85, 0.05, 0.10)	(0.90, 0.10)	1.854	0.3172	0.1083 (8.4)	**0.1060** (6.9)	preserved
(0.85, 0.05, 0.10)	(0.95, 0.05)	2.182	0.4370	0.1055 (9.6)	**0.1042** (8.1)	preserved

Three conclusions follow, and we state all three, including the one that does not favor the optimized schedule.

First, the *ordering* of the strategies with respect to the minimum combined divergence is preserved in all twenty reference pairs: the optimal control attains the lowest minimum, followed by the exposure-matched constant dose, the low constant dose, and finally the untreated case. The margin between the two exposure-matched strategies is modest, ranging from 1.3% to 7.4%, but its sign never changes.

Second, the *timing* advantage is both larger and equally stable. The optimal schedule reaches its minimum divergence between 1.2 and 1.5 days earlier than the exposure-matched constant dose in every reference pair, while the untreated and low-dose trajectories are in most cases still decreasing at t_f and therefore attain no interior minimum at all. This is consistent with the population-level milestones reported in Table 1 and reinforces the interpretation that the benefit of adaptive scheduling in this model is principally one of rate of response at fixed cumulative exposure.

Third, the same sweep evaluated at the *terminal* time reverses the comparison between the two exposure-matched arms: $D_{\text{tot}}(t_f)$ is higher under optimal control than under the exposure-matched constant dose in all twenty reference pairs. As discussed above, this reflects the withdrawal of the optimized infusion as λ_7 approaches its zero terminal value, after which the leukemic population begins to recover, whereas the constant infusion continues to act until t_f. The effect is therefore a property of the free terminal state in Formulation (3) rather than of the reference profiles, and it is as robust as the minimum-divergence ordering.

Finally, the absolute magnitudes are strongly reference-dependent: the minimum divergence under optimal control ranges from 0.085 to 0.452 across the twenty pairs, a variation of more than a factor of five. The entropy and divergence values reported in this paper should accordingly be read comparatively, as a means of ranking strategies on a common scale, and not as absolute measures of immunological or disease states.

6.6. Entropy of the Helper T-Cell Repertoire

The relative entropy D_A measures distance from a prescribed target, whereas the Shannon entropy S_A measures how evenly the helper compartment is distributed, independently of any target. The two are complementary, and reading them together resolves a feature of Figure 7 that D_A alone leaves ambiguous.

Figure 8 shows $S_A(t)$. All four strategies begin near 0.80 and rise slightly to approximately 0.86 by day 1.3, as the initially Th2-leaning compartment becomes more evenly mixed. From day 4 the optimal control drives a steep decrease, reaching approximately 0.29 at t_f; the exposure-matched dose follows the same descent about a day later and reaches a comparable terminal value, while the low constant dose declines only to 0.54 and the untreated case to 0.68.

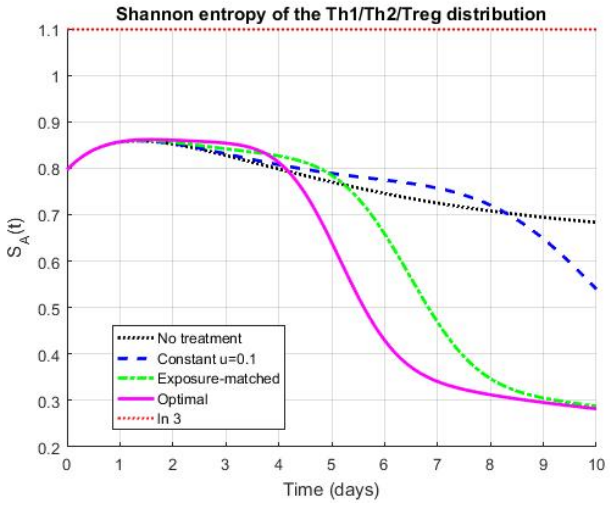


Figure 8. Shannon entropy $S_A(t)$ of the normalized Th1/Th2/Treg distribution under the four dosing strategies. The upper dotted line marks the maximum $\ln 3 \approx 1.099$ attained by the uniform distribution.

The decrease is a loss of diversity, and in this setting it is therapeutically desirable: it reflects a directed concentration of the helper compartment on the protective Th1 subset rather than an unstructured collapse. This is precisely why S_A is not by itself an adequate diagnostic and must be read alongside D_A; entropy registers the concentration but not its direction.

The comparison with the reference profile is also informative. The entropy of $q_A = (0.70, 0.10, 0.20)$ is

$$S(q_A) = -\sum_i q_{A,i} \ln q_{A,i} \approx 0.802,$$

so the treated trajectories terminate at an entropy well *below* that of the target profile. The helper distribution under effective treatment is therefore more strongly Th1-concentrated than q_A prescribes, which explains why D_A plateaus near 0.23 rather than approaching zero in Figure 7: the trajectories do not stop at the reference but pass beyond it. Whether such an overshoot is desirable is a question the present model cannot settle, since q_A encodes a

qualitative immunological target rather than a calibrated one; we record it as an observation that a patient-calibrated reference profile would need to address.

6.7. Divergence-Driven Optimal Control

The diagnostics of Section 6.4 summarize the therapeutic goal in a single dimensionless quantity, which raises the question of whether that quantity could serve as the objective itself, and whether doing so would produce a materially different dosing schedule. We therefore solved a second optimal control problem in which the combined divergence replaces the population-based integrand.

6.7.1. Formulation

Keeping the maximization convention of Section 4, we consider

$$J_{\mathrm{IT}}(u) = -\int_0^{t_f} \Big[D_A(t) + D_L(t) + A_5\, u^2(t) \Big]\, dt, \qquad J_{\mathrm{IT}}(u^*_{\mathrm{IT}}) = \max_{u \in \mathcal{U}} J_{\mathrm{IT}}(u), \tag{21}$$

subject to the same delayed dynamics (Equation (1)) and the same admissible set. Existence follows from Theorem 1 without modification, since the proof uses only that F is affine in u and that the integrand is concave in u and continuous in the state.

6.7.2. Adjoint System

The optimality system changes only in its state-dependent forcing terms. The required gradients are available in closed form.

Lemma 1 (Gradients of the axis divergences)**.** *Let* $\Sigma_A := x_2 + x_3 + x_4$ *and* $\Sigma_L := x_9 + x_{11}$. *Then for* $k \in \mathcal{J}_A$ *and* $\ell \in \mathcal{J}_L$,

$$\frac{\partial D_A}{\partial x_k} = \frac{1}{\Sigma_A}\left[\ln \frac{p^A_k}{q_{A,k}} - D_A\right], \qquad \frac{\partial D_L}{\partial x_\ell} = \frac{1}{\Sigma_L}\left[\ln \frac{p^L_\ell}{q_{L,\ell}} - D_L\right],$$

and the derivatives with respect to the remaining state variables vanish.

Proof. Since $p^A_i = x_i/\Sigma_A$ we have $\partial p^A_i/\partial x_k = (\delta_{ik} - p^A_i)/\Sigma_A$, whence

$$\frac{\partial D_A}{\partial x_k} = \sum_{i \in \mathcal{J}_A} \frac{\partial p^A_i}{\partial x_k}\left[\ln \frac{p^A_i}{q_{A,i}} + 1\right] = \frac{1}{\Sigma_A}\left[\ln \frac{p^A_k}{q_{A,k}} + 1 - \big(D_A + 1\big)\right],$$

the bracketed sum being $D_A + 1$ because $\sum_i p^A_i = 1$. The statement for D_L is identical. □

Accordingly, the adjoint equations of Section 5.1 apply verbatim with the constant cost terms replaced as follows:

$$\text{in } \dot\lambda_2 : \ -A_2 \ \mapsto \ \frac{\partial D_A}{\partial x_2}, \quad \text{in } \dot\lambda_3 : \ +A_3 \ \mapsto \ \frac{\partial D_A}{\partial x_3}, \quad \text{in } \dot\lambda_4 : \ \ 0 \ \mapsto \ \frac{\partial D_A}{\partial x_4},$$
$$\text{in } \dot\lambda_9 : \ +A_4 \ \mapsto \ \frac{\partial D_L}{\partial x_9}, \quad \text{in } \dot\lambda_{11} : \ -A_1 \ \mapsto \ \frac{\partial D_L}{\partial x_{11}}.$$

Two features deserve note. First, the regulatory compartment x_4 now enters the optimization: it appears in the allergy-axis distribution but not in Equation (3), so the divergence-based formulation couples Treg dynamics into the control problem where the population-based one did not. Second, since $\partial H_{\mathrm{IT}}/\partial u = \lambda_7 - 2A_5 u$ exactly as before, the control characterization (Equation (18)) and the numerical scheme of Section 5.4 carry over unchanged.

Remark 5 (Numerical safeguard). *The gradients of Lemma 1 are singular where a component of p^A or p^L vanishes. The helper fractions remain bounded away from zero in all simulations, but the leukemic fraction p_9^L approaches zero under effective treatment. The logarithms were therefore evaluated using fractions floored at $\epsilon_p = 10^{-10}$ and renormalized; the computed schedule was verified to be insensitive to this choice.*

6.7.3. Comparison of the Two Schedules

Figure 9 compares the two schedules, and the difference is substantial. The divergence-driven control u^*_{IT} never saturates: it begins at approximately 0.87 and decreases smoothly to zero by day 5.5, whereas u^* holds the maximal dose for three days before tapering. The cumulative exposures differ accordingly,

$$E(u^*_{\mathrm{IT}}) = 1.636 \qquad \text{against} \qquad E(u^*) = 4.460,$$

so the divergence-driven schedule administers 37% of the drug. The resulting terminal leukemic burdens are $x_9(t_f) = 0.106$ and 0.027 respectively—reductions of 93.2% and 98.3% relative to the untreated terminal value of 1.560.

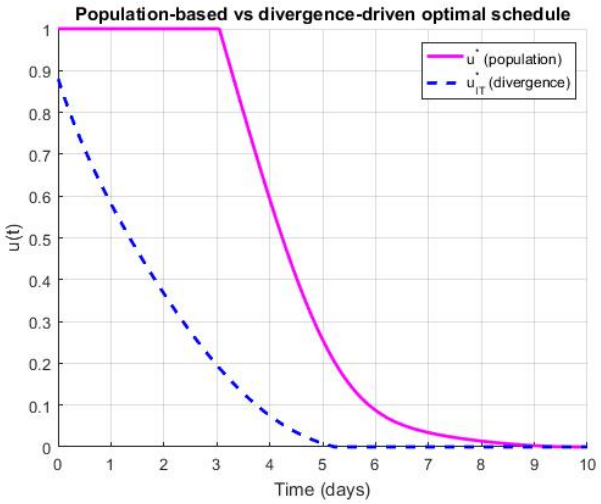


Figure 9. Optimal schedules obtained from the population-based objective (Equation (3)) and from the divergence-based objective (Equation (21)).

This is a genuine therapeutic trade-off rather than a numerical artefact, and Proposition 4 explains it. The divergence D_L depends on the leukemic population only through the fraction $p_9^L = x_9/(x_9 + x_{11})$, so once that fraction is small the objective is nearly insensitive to further absolute reduction, and the quadratic exposure penalty dominates. The population-based objective, by contrast, contains x_9 linearly and continues to reward absolute reduction until the terminal time. The information-theoretic formulation therefore prescribes a markedly gentler schedule that forgoes the last five percentage points of leukemic reduction in exchange for a 63% reduction in cumulative drug exposure.

6.7.4. Relation Between the Two Formulations

Proposition 4 (Scale invariance of the divergence-based objective). *Let $\sigma_A, \sigma_L > 0$ and consider the rescaling $x_i \mapsto \sigma_A x_i$ for $i \in \mathcal{J}_A$ and $x_i \mapsto \sigma_L x_i$ for $i \in \mathcal{J}_L$. Then S_A, D_A, D_L and D_{tot} are unchanged, and hence so is J_{IT}. By contrast J is not invariant, transforming as*

$$J \longmapsto \int_0^{t_f} \left[\sigma_A\left(A_2 x_2 - A_3 x_3\right) + \sigma_L\left(A_1 x_{11} - A_4 x_9\right) - A_5 u^2\right] dt,$$

so that a change of measurement units acts on J exactly as a reweighting of $A_1, \ldots, A_4$.

Proof. p^A and p^L are normalized fractions, hence invariant under a common positive scaling of the populations entering each axis, and S_A, D_A, D_L depend on the state only through them. The transformation of J follows from the linearity of its state-dependent part. □

Remark 6. *Proposition 4 identifies the practical difference between the two formulations. In Equation (3), the choice of weights cannot be separated from the choice of units in which the compartments are measured, so a weighting that is reasonable under one normalization need not be under another; this is one reason the sensitivity study of Section 6.2 is required. Formulation (21) removes that ambiguity at the cost of discarding information about absolute magnitude: it is by construction blind to a proportional reduction in an entire axis. Neither formulation dominates the other. Which is preferable depends on whether the clinical objective is understood as minimizing absolute residual disease or as restoring a favorable immune–leukemic balance at the least cumulative exposure, and the two schedules computed here quantify what that choice costs in each currency.*

6.8. Model Assumptions and Limitations

The proposed model is based on several simplifying assumptions that should be considered when interpreting the results. Nonlinear interactions among Th1, Th2, and regulatory T cells are represented through Hill-type suppression functions [9], which capture the principal mechanisms of immune cross-regulation but do not account for stochastic fluctuations that may arise in real immune responses. Three fixed delays, τ_1, τ_2 and τ_3, describe antigen presentation, cytokine production, and leukemic cell-cycle progression; these correspond to average physiological timescales and do not reflect patient-specific variability. Cytokine activity is aggregated into a single compartment, and the immune effector population combines cytotoxic T lymphocytes and natural killer cells into one variable, which may mask distinct functional roles [28,29]. The chemotherapeutic effect follows the classical log-kill hypothesis [1]; drug resistance, pharmacogenomic variability, and individualized pharmacokinetics are not incorporated. Accordingly, the numerical simulations should be interpreted as exploratory and hypothesis-generating rather than as direct clinical predictions.

A related observation follows from the sensitivity study. The half-saturation constant θ of the leukemic self-renewal flux (Equation (2)) has a normalized sensitivity index of order 10^{-4} on every outcome, because θ exceeds the leukemic burden attained in the simulations by nearly two orders of magnitude. The saturating nonlinearity is therefore inactive in the present regime, and the leukemic equation operates in its effectively quadratic range, with $\Psi(x_9) \approx \beta_L x_9^2$ throughout. Two consequences should be recorded. The results reported here do not probe the saturation of leukemic self-renewal at high density, so a parameter regime in which that mechanism is active would be required to test it; and, over the range of burdens simulated, the per-cell renewal rate $\Psi(x_9)/x_9$ increases with leukemic density rather than decreasing, which departs from the density-limiting form of [19] and should be borne in mind when comparing the present model with that literature.

A further limitation concerns the immune–leukemic coupling. In the parameter regime of Table 8 the interaction coefficients c_1, c_2 and ρ are several orders of magnitude smaller than the supply and decay rates s and m, with the consequence that the effector equation is dominated by its intrinsic supply–decay balance and the effector trajectory is nearly identical across all dosing strategies (Figure 5); the sensitivity indices of these coefficients are indistinguishable from zero on every outcome measure. Two implications follow and should be stated plainly. The reduction in leukemic burden reported here is attributable to the direct cytotoxic term rather than to enhanced immune-mediated clearance; and the finding that immune competence is preserved under optimal dosing is, in this regime, a structural feature of the parameterization rather than an independent prediction of the

model. A parameter set in which the immune–tumor interaction is dynamically active would be required to test whether adaptive dosing additionally recruits immune-mediated clearance, and we regard this as the most informative direction for calibration against clinical data.

Table 8. Model parameters used in the simulations. Values taken directly from the cited source are given as such; values marked calibrated were adjusted so that the untreated system reproduces the slow progression characteristic of chronic lymphocytic leukemia, with the source that informed the choice given for comparison.

The production rate of naive CD4+ cells. [30]	α	0.1
The strength of suppression rate of Th1 by Th2 [9]	μ_2	0.1
The strength of suppression of Th2 by Th1 [9]	μ_1	0.2
The strength of suppression rate by Treg [9]	μ_r	0.25
The differences in the autocrine action of the three subsets [9]	ϕ	0.05
The differences in the autocrine action of the three subsets [9]	κ	0.1
The death rate of immature APCs [30]	γ_{11}	$0.08\ \text{day}^{-1}$
The death rate of mature APCs [30]	γ_{12}	$0.8\ \text{day}^{-1}$
The death rate of Naive T cells [30]	β_1	$0.03\ \text{day}^{-1}$
The death rate of T_1 cells, calibrated; cf. [31]	β_2	$4.16 * 10^{-5}\ \text{day}^{-1}$
The death rate of T_2 cells, calibrated; cf. [31]	β_3	$4.16 * 10^{-5}\ \text{day}^{-1}$
The death rate of T_{reg} cells, calibrated; cf. [31]	β_4	$4.16 * 10^{-5}\ \text{day}^{-1}$
The proliferation rate of stimulated T-cells [9]	v	4
Natural decay of induced cytokine during chemotherapy [29]	γ_2	$0.4152\ \text{day}^{-1}$
Inhibition rate of Treg cells by the induced cytokines [32]	η_r	$0.4\ \text{day}^{-1}$
First time delay [33]	τ_1	0.0794 day
Second time delay [29]	τ_2	0.25 day
Third time delay [34]	τ_3	2.8 day
The production of induced cytokines by other cells [35]	k_1	1
The supply rate of naive APCs, calibrated; cf. [30]	a_{APC}	2
Rate of APC activation by the antigen [36]	β_0	0.01
Rate of APC inhibition by regulatory T cells [36]	μ_0	10^{-2}
Chemical deactivation rate of drug, calibrated; cf. [1]	γ_D	$0.0192\ \text{day}^{-1}$
Deactivation rate of drug due to killing of tumor cells, calibrated; cf. [1]	μ_D	$7.48 * 10^{-3}\ \text{day}^{-1}$
Drug dose that produces 50% maximum effect [1]	a	$2 * 10^3$ mL = 2 in microliter
Natural death rate of leukemic cells, calibrated; cf. [19]	γ_L	$0.01\ \text{day}^{-1}$
Death rate resulting from the action of drug on cancer cells estimated from [1]	μ_L	$0.74\ \text{day}^{-1}$
Rate of dissolution of dead tumor cells, calibrated; cf. [1]	d	$7.07 * 10^{-4}\ \text{day}^{-1}$
The supply rate of immune cells [37]	s	$7 * 10^5\ \text{cells day}^{-1} = 0.1\ \text{cells}\ \mu\text{L}^{-1}\ \text{day}^{-1}$
Natural death rate of immune cells [37]	m	$10^{-3}\ \text{day}^{-1}$
Rate of immune cell population activated by the tumor [37]	ρ	$10^{-12}\ \text{day}^{-1}$
Half-saturation in immune activation [37]	γ	$10^2\ \text{cells} = 1.42 * 10^{-5}\ \text{cells}\ \mu\text{L}^{-1}$
The mortality rate of immune cells due to the chemotherapeutic drug [37]	δ	$1.42 * 10^{-3}\ \text{day}^{-1}$
Half-saturation parameter [37]	c	$7.14 * 10^{-7}$ in microliter
Rate of elimination of tumor cells by immune cells [37]	c_1	$5 * 10^{-11}\ \mu\text{L cells}^{-1}\ \text{day}^{-1}$
Rate of immune cells which directly eliminate tumor cells [37]	c_2	$10^{-13}\ \mu\text{L cells}^{-1}\ \text{day}^{-1}$
Component of the rate of self-renewal, calibrated; cf. [19]	β_L	$0.01\ \text{day}^{-1}$
Component of the rate of self-renewal, calibrated; cf. [19]	θ	$100\ \text{cells}\ \mu\text{L}^{-1}$

Note on calibration. All concentrations are expressed per microlitre; conversions from absolute cell counts use a total blood volume of 7×10^6 µL.

A tension between the optimized schedule and a clinical consideration raised in the introduction should also be acknowledged. The computed control administers the maximal admissible dose for approximately the first three days, and the leukemic population falls by half within 2.7 days. Rapid destruction of a large leukemic burden is precisely the circumstance associated with tumor lysis syndrome, a complication the present objective functional does not represent. A formulation intended for clinical interpretation would require either a state constraint bounding the rate of leukemic decline or an explicit term penalizing it; the schedules computed here should therefore be read as characterizing the efficacy–exposure trade-off in isolation, and not as directly implementable protocols.

The information-theoretic quantities introduced in Section 6.4 depend on prescribed reference distributions q_A and q_L, which reflect qualitative immunological targets rather than patient-calibrated profiles. This affects both roles in which they are used. As diagnostics, the ordering of dosing strategies is preserved across all twenty alternative reference pairs tested in Section 6.5, but the absolute divergence values vary by more than a factor of five and must be read comparatively. As an objective, in Section 6.7, the reference profiles

become design parameters of the control problem, so that the schedule u^*_{IT} inherits whatever clinical judgement is embedded in them; calibrating these profiles against patient data is consequently a prerequisite for any translational use of the divergence-driven formulation, and a more demanding requirement than for the diagnostic use.

From a practical perspective, implementation of adaptive chemotherapy protocols guided by immune dynamics would require repeated monitoring of relevant biomarkers, including cytokine assays, flow cytometry measurements, and immune profiling. Such requirements may increase logistical and financial demands, particularly in resource-limited settings, though continued advances in diagnostic technologies may improve their feasibility.

Finally, while the present study compares the optimized schedule against untreated, low-dose and exposure-matched constant protocols, and against the divergence-driven schedule of Section 6.7, systematic comparison against alternative optimal control formulations in oncology—objective functionals with terminal cost, L^1 penalties on the control inducing genuinely bang–bang schedules, or explicit state constraints representing toxicity thresholds—was not undertaken and remains an important topic for future investigation. The first of these is of particular interest here, since a terminal cost on the leukemic population would remove the late rebound documented in Section 6.5. Future extensions may also include patient-specific parameter estimation, multiple cytokine subclasses, stochastic immune fluctuations, drug-resistance mechanisms, and calibration using clinical datasets.

7. Conclusions and Recommendations

The numerical simulations indicate that the proposed optimal control strategy improves therapeutic outcomes within the model when compared with the absence of treatment, with a low constant dose, and—most informatively—with a constant infusion delivering the same cumulative drug exposure. Because the last of these holds the total drug administered fixed at $E = 4.4596$, the differences between it and the optimized schedule are attributable to the temporal distribution of the dose alone. On that comparison the optimal schedule improves the objective functional by 40.5% ($J = 24.15$ against 17.18) and reaches every therapeutic milestone between 0.9 and 1.6 days earlier, while attaining a terminal leukemic burden lower by a factor of 2.25. The benefit of adaptive scheduling in this model is therefore principally one of *rate of response at fixed drug exposure* rather than of eventual position: given the same total drug, a constant infusion reaches a broadly comparable configuration, only later.

The principal observations may be summarized as follows:

- **Effector Th1 cells (x_2).** Under the optimized strategy the Th1 population expands rapidly from approximately day 4, overtaking Th2 at day 3.74 and reaching 7.011 by the end of the horizon. The exposure-matched constant dose attains a comparable terminal level of 6.803 but crosses over at day 4.61, while the low constant dose produces only a limited late increase to 2.073 and the untreated case remains at 0.286. The distinguishing feature is thus the timing of the expansion rather than its eventual magnitude.
- **Suppressor Th2 cells (x_3).** In absolute terms the Th2 population rises under all scenarios and is in fact slightly highest under optimal control. The therapeutic benefit therefore lies in the Th1/Th2 *ratio* rather than the absolute Th2 level. Within the Th1/Th2 framework adopted here, this shift in proportion is the mechanism by which the model represents a reduced risk of allergic hypersensitivity; the model does not represent hypersensitivity events directly, and the association should be read as a modeling assumption inherited from [9,10] rather than a validated clinical endpoint.

The sensitivity study further shows that the Th2 penalty A_3 has almost no influence on the computed schedule, since Th2 varies by less than 15% across all strategies.

- **Leukemic cells (x_9).** The leukemic burden falls to 0.0267 under the optimal schedule and 0.0601 under the exposure-matched constant dose, against 0.4474 for the low constant dose and continued slight progression when untreated. In the parameter regime considered this reduction is driven by the direct cytotoxic term: the interaction coefficients c_1, c_2 and ρ have sensitivity indices indistinguishable from zero, so the model does not support an additional claim of enhanced immune-mediated clearance.
- **Immune effector cells (x_{11}).** The effector population grows almost identically under all four scenarios, indicating that within the present model it is governed by its intrinsic supply–decay dynamics and is not measurably suppressed by chemotherapy over this horizon. The optimal strategy therefore achieves substantial leukemic reduction without compromising immune competence, though this should be understood as a structural feature of the parameterization rather than an independent prediction.
- **Information-theoretic diagnostics.** The relative-entropy diagnostics on the allergy and leukemia axes, together with their combined divergence, characterize both treatment goals in a single scale-free framework. Under optimal control both axes approach their healthy targets fastest, and the minimum combined divergence is both lowest (0.194 against 0.203) and earliest (day 5.82 against day 7.19); this ordering is preserved across all twenty alternative reference profiles tested. At the terminal time, however, the comparison reverses—$D_{\mathrm{tot}}(t_f) = 0.333$ under optimal control against 0.292 under the exposure-matched dose—because the optimized infusion is withdrawn as the costate approaches its zero terminal value while the constant infusion continues to act. This reversal is equally robust across reference profiles and is a property of the free terminal state rather than of the diagnostic.
- **Divergence-driven control.** When the combined divergence is taken as the running cost of a second optimal control problem, the resulting schedule differs materially: it never saturates, administers only 37% of the drug ($E = 1.636$ against 4.460), and still achieves a 93.2% leukemic reduction against 98.3% for the population-based formulation. Because the divergence depends on the leukemic population only through a normalized fraction, it becomes insensitive to further absolute reduction once that fraction is small. The two formulations of the therapeutic goal therefore do not induce the same optimal policy, and the comparison quantifies what the choice between them costs in each currency.

Overall, these findings suggest that adaptive, immune-aware chemotherapy scheduling may accelerate disease control and the establishment of a Th1-favorable helper balance at fixed cumulative drug exposure, and that information-theoretic objectives offer a scale-free alternative formulation of the therapeutic goal that prescribes substantially gentler dosing. The numerical results are supported by documented convergence of the optimality system—geometric at a rate of approximately 0.55 per iteration, with the iteration count independent of the discretization—and by grid independence to six significant figures. A one-at-a-time sensitivity study over kinetic parameters, delays, initial conditions and objective weights leaves the qualitative structure of the optimal schedule and the ordering of the four strategies unchanged; the drug-induced leukemic death rate μ_L is the only parameter whose relative perturbation is amplified, with a normalized sensitivity index of -4.53 on the terminal leukemic burden.

The limitations set out in Section 6.8 should nevertheless be borne in mind, in particular that the immune–leukemic coupling is dynamically inactive in the present parameter regime, that the schedules computed here administer maximal dose over an initial phase

without representing tumor lysis risk, and that the conclusions rest on simulation without direct clinical validation.

Recommendations. Future therapeutic protocols for CLL may benefit from incorporating immune monitoring to guide individualized chemotherapy dosing. In particular, the Th1/Th2 ratio—and, more generally, information-theoretic measures such as the entropy and relative entropy of the helper T-cell and immune/leukemic distributions—could serve as useful biomarkers for balancing treatment efficacy and safety. Three directions follow most directly from the present results. First, experimental determination of μ_L would be the single most valuable refinement, since it is the only parameter whose uncertainty is amplified in the outcome. Second, a formulation incorporating either a terminal cost on the leukemic population or a state constraint bounding its rate of decline would address, respectively, the late rebound documented here and the tumor lysis risk associated with an initial maximal-dose phase. Third, calibration of the reference profiles q_A and q_L against patient data is a prerequisite for any translational use of the divergence-driven formulation, since in that formulation they cease to be diagnostic conventions and become design parameters of the control problem. Further work should also include patient-specific parameter estimation, comparison with alternative optimal control strategies, and extension to combination therapies or multi-drug protocols.

8. Parameters

Parameter values used in the simulations are listed below. Values taken directly from the cited source are given as such; values that were calibrated so that the untreated system reproduces the slow progression characteristic of chronic lymphocytic leukemia are marked *calibrated*, with the source that informed the choice given for comparison.

The parameters governing leukemic self-renewal and clearance (β_L, θ, γ_L) were calibrated rather than adopted directly from [19], which reports values for chronic myelogenous leukemia. Adopted directly, those values cause the untreated system to clear the leukemic population within approximately two days, which is inconsistent with the indolent progression characteristic of chronic lymphocytic leukemia. The calibrated values instead reproduce a near-stationary untreated trajectory: over the ten-day horizon the untreated leukemic burden rises from 1.500 to 1.560 (Figure 4), corresponding to an exponential growth rate of 3.9×10^{-3} day^{-1} and a lymphocyte doubling time of 177 days, or approximately 5.8 months. Lymphocyte doubling time has been established as a prognostic marker in chronic lymphocytic leukemia since [38], and a doubling time below six months is one of the criteria for active disease indicating that treatment is warranted [39]. The calibration therefore places the untreated model in the progressive-disease regime for which the treatment schedules studied here would in practice be considered, rather than being chosen to produce a particular treatment outcome.

The sensitivity study of Section 6.2 indicates that this calibration does not drive the conclusions: the normalized sensitivity indices of β_L, γ_L and θ on the terminal leukemic burden are 0.32, -0.11 and 10^{-4}, respectively, so a $\pm 20\%$ uncertainty in any of them leaves the reported comparisons unchanged.

The drug pharmacokinetic parameters (γ_D, μ_D, d), the helper-cell death rates ($\beta_2, \beta_3, \beta_4$) and the naive APC supply rate (a_{APC}) are likewise calibrated values informed by the cited sources rather than direct adoption of the reported figures. Two half-saturation constants appear in the model with very different magnitudes: $a = 2$ in the drug-elimination and cytotoxic terms, and $c = 7.14 \times 10^{-7}$ in the immune-mortality term, although both refer to the same drug concentration x_7. The latter follows the normalization of [37], in which the corresponding parameter is reported on a scale that differs from the one used here; with $c \ll x_7$ over almost the whole horizon, the factor $x_7/(c + x_7)$ is effectively unity once

treatment begins, so the immune-mortality term reduces to δx_{11}. Since $\delta = 1.42 \times 10^{-3}$ is two orders of magnitude below the effector supply rate, this term is in any case negligible, as the sensitivity indices of both δ and c confirm.

Author Contributions: Conceptualization, R.A., A.H. and L.A.O.; methodology, R.A. and A.H.; formal analysis, L.A.O. and R.A.; software, R.A.; validation, R.A. and L.A.O.; investigation, R.A.; writing—original draft preparation, R.A.; writing—review and editing, L.A.O., A.H. and R.A.; visualization, R.A.; supervision, A.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The MATLAB code implementing the forward–backward sweep and generating all figures and tables reported in this article is available from the corresponding author upon reasonable request. No new data were created.

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

References

1. Guzev, E.; Luboshits, G.; Mendrazitsky, S.B.; Firer, M.A. Experimental Validation of a Mathematical Model to Describe the Drug Cytotoxicity of Leukemic Cells. *Symmetry* **2021**, *13*, 1760. [CrossRef]
2. Tambaro, F.P.; Wierda, W.G. Tumour lysis syndrome in patients with chronic lymphocytic leukaemia treated with BCL-2 inhibitors: risk factors, prophylaxis, and treatment recommendations. *Lancet Haematol.* **2020**, *7*, e168–e176. [CrossRef] [PubMed]
3. Fischer, K.; Al-Sawaf, O.; Hallek, M. Preventing and monitoring for tumor lysis syndrome and other toxicities of venetoclax during treatment of chronic lymphocytic leukemia. *Hematol. Am. Soc. Hematol. Educ. Program.* **2020**, *2020*, 357–362. [CrossRef] [PubMed]
4. Bašić-Kinda, S.; Hude, I.; Ranković, E.; Ostojic-Kolonic, S. Comparison of different schedules of rituximab and chlorambucil in previously untreated chronic lymphocytic leukemia: A retrospective study of KROHEM. *J. Cancer Sci. Clin. Ther.* **2021**, *5*, 426–433. [CrossRef]
5. Torricelli, R.; Kurer, S.B.; Kroner, T.; Wuthrich, B. Delayed allergic reaction to Chlorambucil (Leukeran). Case report and literature review. *Schweiz. Med. Wochenschr.* **1995**, *125*, 1870–1873. [PubMed]
6. Levin, M.; Libster, D. Allergic reaction to chlorambucil in chronic lymphocytic leukemia presenting with fever and lymphadenopathy. *Leuk Lymphoma.* **2005**, *46*, 1195–1197. [CrossRef] [PubMed]
7. Freeman, C.L.; Gribben, R.; Eichhorst, P.; Hallek, M. Cytokine release in chronic lymphocytic leukemia patients treated with obinutuzumab and chlorambucil and possible relationship with infusion-related reactions. *Blood* **2015**, *126*, 2646–2649. [CrossRef]
8. Gubernatorova, E.; Gorshkova, E.; Namakanova, O.; Zvartsev, R.; Hidalgo, J.; Drutskaya, M.; Tumanov, A.; Nedospasov, S. Non-redundant functions of IL-6 produced by macrophages and dendritic cells in allergic airway inflammation. *Front. Immunol.* **2018**, *9*, 2718. [CrossRef] [PubMed]
9. Gross, F.; Behn, U. Mathematical modeling of allergy and specific immunotherapy: Th1, Th2 and Treg interactions. *J. Theor. Biol.* **2011**, *269*, 70–78. [CrossRef] [PubMed]
10. Segel, L.A.; Fishman, M.A. Modeling immunotherapy for allergy. *Bull. Math. Biol.* **1996**, *58*, 1099–1121. [CrossRef] [PubMed]
11. Abdullah, R.; Amin, K.; Halanay, A.; Mghames, R. Partial Stability in a Model for Allergic Reactions Induced by Chemotherapy of Acute Lymphoblastic Leukemia. *Ann. Acad. Rom. Sci. Ser. Math. Appl.* **2023**, *15*, 1–2. [CrossRef]
12. Abdullah, R.; Badralexi, I.; Halanay, A. Stability Analysis in a Mathematical Model for Allergic Reactions. *Axioms* **2024**, *13*, 102. [CrossRef]
13. Abdullah, R.; Badralexi, I.; Halanay, A. Stability Analysis in a New Model for Allergic Reactions Induced by Chemotherapy of Chronic Lymphocytic Leukemia. *Mathematics* **2023**, *11*, 3225. [CrossRef]
14. Abdullah, R.; Badralexi, I.; Fakih, L.; Halanay, A. Mathematical Modeling of Immune Dynamics in Chronic Myeloid Leukemia Therapy: Unraveling Allergic Reactions and T Cell Subset Modulation by Imatinib. *Axioms* **2024**, *13*, 464. [CrossRef]
15. Gollmann, L.; Maurer, H. Theory and applications of optimal control problems with multiple time-delays. *J. Ind. Manag. Optim.* **2014**, *10*, 413–441. [CrossRef]
16. Rodrigues, F.; Silva, C.J.; Torres, D.F.M. Optimal control of a delayed HIV model. *Discret. Contin. Dyn. Syst. Ser. B* **2018**, *23*, 443–458. [CrossRef]
17. McAsey, M.; Mou, L.; Han, D. Convergence of the forward–backward sweep method in optimal control. *Comput. Optim. Appl.* **2012**, *53*, 207–226. [CrossRef]

18. Saleem, M.; Hameed, N.; Kara, A.H.; Younis, M.; Furati, K.M. A review of forward–backward sweep optimality technique in optimal control problems. *Numer. Algebr. Control. Optim.* **2014 4**, 387–402.
19. Colijn, C.; Mackey, M.C. A mathematical model of hematopoiesis I-Periodic chronic myelogenous leukemia. *J. Theor. Biol.* **2005**, *237*, 117–132. [CrossRef] [PubMed]
20. Kharitonov, V.L. *Time-Delay Systems, Lyapunov Functionals and Matrices*; Birkhäuser: Basel, Switzerland, 2013.
21. Kolmanovskii, V.; Myshkis, A. *Applied Theory of Functional Differential Equations*; Kluwer Academic Publishers: Dordrecht, The Netherlands, 1992.
22. Pontryagin, L.S.; Boltyanskii, V.G.; Gamkrelidze, R.V.; Mishchenko, E.F. *The Mathematical Theory of Optimal Processes*; Interscience Publishers, New York, NY, USA, 1962.
23. Kolmanovskii, V.; Myshkis, A. *Introduction to the Theory and Applications of Functional Differential Equations*; Kluwer Academic Publishers: Dordrecht, The Netherlands, 1999.
24. Fleming, W.H.; Rishel, R.W. *Deterministic and Stochastic Optimal Control*; Springer: Berlin/Heidelberg, Germany, 1975.
25. Zubov, V. *Oscillations in Non-Linear and Controllable Systems*; Gosudarstvennoe Soyuznoe Izdatelstvo Sudostroitelnoi Promyshlennosti: Leningrad, Russia, 1962. (In Russian)
26. Brandt-Pollmann, U.; Winkler, R.; Sager, S.; Moslener, U.; Schlöder, J. Numerical solution of optimal control problems with constant control delays. *Comput. Econ.* **2008**, *31*, 181–206. [CrossRef]
27. Kirk, D.E. *Optimal Control Theory: An Introduction*; Prentice-Hall: Hoboken, NJ, USA, 1970.
28. Greil, R.; Tedeschi, A.; Moreno, C.; Anz, B.; Larratt, L.; Simkovic, M.; Gill, D.; Gribben, J.G.; Flinn, I.W.; Wang, Z.; et al., Pretreatment with ibrutinib reduces cytokine secretion and limits the risk of obinutuzumab-induced infusion-related reactions in patients with CLL: Analysis from the iLLUMINATE study. *Ann. Hematol.* **2021**, *100*, 1733–1742. [CrossRef] [PubMed]
29. Nazari, F.; Pearson, A.T.; Noë, J.E.; Jackson, T.L. A mathematical model for IL-6-mediated, stem cell driven, tumor growth and targeted treatment. *PloS. Comput. Biol.* **2018**, *14*, e1005920. [CrossRef] [PubMed]
30. Kim, P.; Lee, P.; Levy, D. A theory of immunodominance and adaptive regulation. *Bull. Math. Biol.* **2011**, *73*, 1645–1665. [CrossRef] [PubMed]
31. Kogan, Y.; Agur, Z.; Elishmereni, M. A mathematical model for the immunotherapeutic control of the TH1/TH2 imbalance in melanoma. *Discret. Contin. Dyn. Syst. Ser. B* **2013**, *18*, 1017–1030. [CrossRef]
32. Hong, T.; Xing, J.; Li, L.; Tyson, J. A Mathematical Model for the Reciprocal Differentiation of T Helper 17 Cells and Induced Regulatory T Cells. *PLoS Comput. Biol.* **2011**, *7*, e1002122. [CrossRef] [PubMed]
33. Wu, G. Calculation of steady-state distribution delay between central and peripheral compartments in two-compartment models with infusion regimen. *Eur. J. Drug Metab. Pharmacokinet.* **2002**, *27*, 259–264. [CrossRef] [PubMed]
34. Radulescu, I.R.; Candea, D.; Halanay, A. A study on stability and medical implications for a complex delay model for CML with cell competition and treatment. *J. Theor. Biol.* **2014**, *363*, 30–40. [CrossRef] [PubMed]
35. Kareva, I.; Berezovskaya, F.; Karev, G.; Mathematical model of a cytokine storm. *bioRxiv* **2022**. [CrossRef] [PubMed]
36. Fouchet, D.; Regoes, R. A Population Dynamics Analysis of the Interaction between Adaptive Regulatory T Cells and Antigen Presenting Cells. *PLoS ONE* **2008** *3*, e2306. [CrossRef] [PubMed]
37. Gil, W.; Carvalho, T.; Mancera, P.; Rodrigues, D.S. A Mathematical Model on the Immune System Role in Achieving Better Outcomes of Cancer Chemotherapy. *Trends Comput. Appl. Math.* **2019**, *20*, 343–357. [CrossRef]
38. Montserrat, E.; Sánchez-Bisono, J.; Nolas, N.V.; Rozman, C. Lymphocyte doubling time in chronic lymphocytic leukaemia: analysis of its prognostic significance. *Br. J. Haematol.* **1986**, *62*, 567–575. [CrossRef] [PubMed]
39. Hallek, M.; Cheson, B.D.; Catovsky, D.; Caligaris-Cappio, F.; Dighiero, G.; Döhner, H.; Hillmen, P.; Keating, M.; Montserrat, E.; Chiorazzi, N.; et al. iwCLL guidelines for diagnosis, indications for treatment, response assessment, and supportive management of CLL. *Blood* **2018**, *131*, 2745–2760. [CrossRef] [PubMed]

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