Preprints.org

Article

Not peer-reviewed version

LOGOS Theory: A Novel Framework for Dosage Regimen Design and Evaluation in Pharmacotherapy

Vangelis D. Karalis *

Posted Date: 29 October 2026

Keywords: whole-regimen optimization; pharmacokinetics; pharmacodynamics; optimal control; recovery; toxicity; personalized dosing

Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

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.

Article

LOGOS Theory: A Novel Framework for Dosage Regimen Design and Evaluation in Pharmacotherapy

Vangelis D. Karalis [1,2,*]

[1] School of Health Sciences, National and Kapodistrian University of Athens, 15784 Athens, Greece

[2] Institute of Applied and Computational Mathematics, Foundation for Research and Technology—Hellas (FORTH), 70013 Heraklion, Greece

* Correspondence: vkaralis@pharm.uoa.gr; Tel.: +30-210-727-4267

Abstract

Optimizing pharmacotherapy requires decisions beyond dosage itself but also the timing, duration and discontinuation of the treatment. This study introduces LOGOS (Longitudinal Optimization of Global Outcomes and Strategies) as a new, transparent, and reproducible mathematical framework for assessing and improving drug strategies. LOGOS aims twofold: a) to quantify the additional therapeutic benefit of different dosing strategies in the presence of common limitations in efficacy, resources and safety, and b) to find the best dosing strategies for the given pharmacokinetic/pharmacodynamic setting and its mathematical optimization. LOGOS assesses the safety at the time of treatment and after discontinuation in the absence of long-term drug toxicity. Mathematical analysis shows that the loading-maintenance-drug-free strategy is globally optimal for long treatment horizons whereas the first dose and then drug-free interval is the optimal strategy for shorter treatment horizons. Additional analyses have revealed mechanistic conditions for continuous or intermittent treatment. LOGOS could allow making better decisions in terms of dose adjustments, treatment intervals, planned interruptions, and safe restarting of treatment. The open-source Python implementation independently verified calculations and reproducibility procedures. LOGOS provides a general basis for an individualized and evidence-based treatment plan in the clinic.

Keywords: whole-regimen optimization; pharmacokinetics; pharmacodynamics; optimal control; recovery; toxicity; personalized dosing

1. Introduction

The design of a dosage regimen involves the selection of the amount of drug administered, the timing and structure of administration, and the conditions under which treatment is stopped or restarted. These decisions determine the subsequent course of exposure and response in the biological system. In particular, the effect of a dose depends on the remaining exposure, responsiveness, accumulated injury and disease status, as well as on the time available for benefit or recovery. Therefore, the aim is to select a trajectory that can be achieved through administration and is appropriate for the therapeutic objective.

In the literature, several optimal-control approaches have already been developed to address these decisions. For example, OptiDose formulates individualized pharmacokinetic/pharmacodynamic (PK/PD) dosing as an optimal-control problem [1], while its enhanced formulation includes efficacy and safety targets and state constraints for the design of doses and dosing intervals [2]. In the case of chronic myeloid leukemia, optimization has also been applied to switching among imatinib, dasatinib and nilotinib and to the scheduling of drug holidays under toxicity constraints [3]. Thus, optimization of multiple dosing variables and consideration of toxicity constraints are already established components of regimen design.

In this study, the Longitudinal Optimization of Global Outcomes and Strategies (LOGOS) is proposed as the comparison strategy with respect to PK/PD rather than a new PK/PD model or

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

numerical optimization. It has a specific effect of assessing the incremental value of regimen-design decision-making by means of nested admissible classes under matched goals, models, starting conditions, resource rules, evaluation horizons and safety constraints. This distinguishes the effects of greater treatment choices from changes to the comparison problem in our approach. Furthermore, the evidence supporting each comparison is made explicit with feasible numerical improvements, class improvements in analytically bounded form, the availability of the target to achieve it and the global certification of the optima. Mechanistic bounds, the physical attainability of our treatment and the safety status throughout the treatment period, i.e., after treatment cessation, are also considered in our analysis. Thus, LOGOS provides a systematic way to determine what additional regimen flexibility achieves, under which assumptions, and with what level of mathematical or computational support. Resource saving is considered as a separate optimization question rather than inferring from fixed-amount benefit gains.

The term LOGOS refers to the mathematical comparison framework developed in this study. The term longitudinal describes the evolution of the treatment system over the evaluation and safety horizons. Global outcomes refer to the total therapeutic consequences selected for assessment. It should be mentioned that this terminology does not imply that every numerical solution is globally optimal. Strategies include both predetermined schedules and policies in which treatment decisions are adapted to the available information. Pharmacokinetics (PK) describes drug disposition, whereas pharmacodynamics (PD) describes drug effects. Their coupled description supports these treatment decisions.

Bayesian model-informed precision dosing (MIPD) uses observations to update individualized predictions with respect to specified targets [4–6]. The treatment actions considered in MIPD overlap with those used in dose-schedule optimization, optimal control and observation-based control [1–3,7]. In the Bayesian pharmacokinetics integrated dose-schedule finding (PKIDS) design, adaptive phase I–II dose-schedule selection is performed through joint modeling of PK, toxicity and efficacy [8]. Other tools, such as mapbayr, nlmixr and TUCUXI, provide the infrastructure for estimation, simulation or dosage adaptation [9–11]. These approaches can therefore serve as relevant comparators. However, their documented functions do not establish that each one implements the LOGOS comparison protocol. The use of adaptive abiraterone schedules [12,13], antimicrobial dosing software [14] and Bayesian MIPD in routine practice [15] shows that regimen adaptation already has clinical and computational applications. In addition, analytical studies have identified optimal administration structures for specific objectives. For example, Katriel derived an ideal antimicrobial concentration profile and its approximation through loading and infusion for an eradication/minimum-exposure objective [16]. In this context, LOGOS builds on the existing approaches in order to examine the contribution of an additional regimen choice under matched conditions and the evidence supporting that contribution.

A distinct mathematical contribution of this study is the derivation of a sufficient global optimality certificate for a finite-horizon dosage problem with linear one-compartment pharmacokinetics, direct concave Emax benefit, and fixed total administered amount and zero predose concentration. The certificate is established over the full class of nonnegative administered-dose measures, including both discrete boluses and continuous inputs, rather than only over a selected finite-dimensional schedule family. A supporting marginal-value inequality proves global optimality and identifies a horizon-dependent administration structure: loading, constant-concentration maintenance and an input-free coast for sufficiently long horizons, and an initial bolus followed by coast for shorter horizons. Unlike the minimum-exposure problem subject to an eradication target considered in [16], the present result addresses maximization of integrated Emax benefit at fixed administered amount and treatment horizon. This analytical result provides an exact benchmark for the LOGOS comparisons, allowing feasible numerical candidates to be evaluated against a certified optimum without attributing global optimality to searches that do not establish it. The certificate is specific to the stated model and input class and does not extend automatically to other PK/PD mechanisms or binding constraints. An earlier, shorter presentation of the LOGOS

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

approach appeared as the medRxiv preprint [17]. In the present study, the mathematical formulation is extended and the derivations and model-based comparisons are provided.

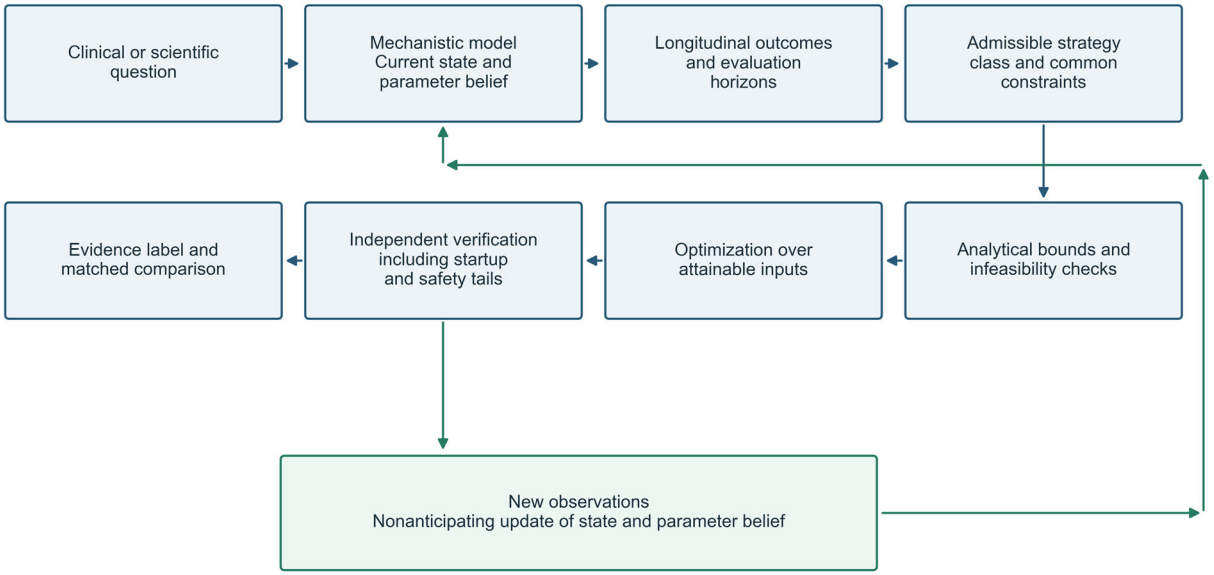


Figure 1. A schematic representation of the Longitudinal Optimization of Global Outcomes and Strategies (LOGOS) procedure. New observations are used to update the state and the belief about model parameters before subsequent optimization. Treatment actions are selected using the information available at that time. The comparison is defined by the outcomes, horizons and admissible actions, while independent verification includes safety at startup and after cessation. If no feasible solution is established, the comparison question is reconsidered. A change in the budget or safety limits defines a new comparison.

The aim of this study is to a) create the LOGOS comparison protocol for systematically evaluating the benefits of additional dosage-regimen choices under common efficacy, resource, and safety conditions and b) develop an optimal dosing strategy for a pharmacokinetic/pharmacodynamic model and show it mathematically. Theoretical work on the same topic focuses on when more flexible treatment design provides additional benefits and how this benefit can be clearly shown. Analytical results and numerical examples are given to illustrate the framework and its potential advantages.

2. Materials and Methods

2.1. States Actions and Observations

Table 1 defines the principal notation. Let x contain the states required by the selected model, potentially including drug amounts, active metabolites, effect-site exposure, sensitivity, injury and disease burden. Let θ contain model parameters. A regimen is denoted s, and x_s is its state trajectory. A schedule combines continuous administration u(t) with discrete doses D_i at times t_i. Between administrations and at jumps, respectively,

$$\dot{x}(t) = f(x(t), u(t), t; \theta). \quad (1)$$
$$x(t_i^+) = \Psi_i(x(t_i^-), D_i). \quad (1b)$$

Continuous administration is expressed as the amount of drug delivered per unit time. In the case of a discrete dose, the administered amount is introduced into the appropriate compartment through the jump map. For example, an intravenous dose may increase the central amount instantaneously, whereas an oral dose usually enters a depot compartment. The subsequent effects of bioavailability, absorption and disposition are described by the model. Thus, each administered amount is counted once and is distinguished from the exposure resulting from the selected route.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

$$D_{tot} = \sum_{i=1}^{N} D_i + \int_0^{T_{adm}} u(t)\, dt. \quad (2)$$

Table 1. LOGOS terminology and notation.

Symbol	Meaning	Units or interpretation
x; θ	State vector; model parameters	Component-specific units
u; D_i; t_i	Continuous input; discrete amount; administration time	Amount/time; amount; time
T_{adm}; T_{eff}; T_{safe}	Administration, efficacy and safety horizons	Time; explicitly declared
Finite; Periodic; Relaxed	Finite course; stationary cycle; exposure distribution	Different comparison problems
$\mathcal{A}_i$; $\mathcal{F}_i$	Admissible class; constrained feasible subset	Actions versus trajectory requirements
Π; π	Policy class; a non-anticipating policy	Actions depend only on available information
J; $J_i^\star$	Scalar objective; feasible class supremum	Declared compatible utility units
O; B; Φ	Outcome vector; integrated and terminal benefit	Separate endpoint units and horizons
Z_{max}; Z_{cum}	Peak and cumulative injury	Injury; injury × time
$Y_{disease}$; D_{tot}; N_{doses}	Disease endpoint; total amount; dose count	Endpoint units; amount; count
F_{bio}; CL	Bioavailability; clearance	Dimensionless; volume/time
C; k; V	Concentration; elimination rate; distribution volume in the one-compartment illustrations	Amount/volume; inverse time; volume
a; z	Sensitivity; nonnegative injury state	Dimensionless sensitivity; model-specific injury units
H_i; ΔH_i	Feasible-class indicator; change in this indicator	Dimensionless; zero or one under nesting

The admissible class defines the choices available for regimen design, including dose count, dose amounts, administration times, infusion duration and rate, and cessation and restart rules. Changes in route or formulation can be considered only when the corresponding alternatives and their models are included. In the case of an irregular schedule, frequency follows from the administration times and cannot be specified independently in a contradictory way. Similarly, infusion rate, duration and amount must satisfy their mass relation. Practical restrictions, such as minimum spacing, available infusion equipment and permissible monitoring times, are also included.

For each admissible schedule the dynamics must define a unique physical trajectory. For each measurable admissible input u, let f(x,u(t),t;θ) satisfy the Carathéodory conditions: measurability in time, continuity in the state, and locally integrable bounds on its magnitude and local state-Lipschitz coefficient on compact sets of states. Jump maps are well defined and preserve the physical domain and each finite course has finitely many jumps. An invariant bounded domain or a suitable growth condition can also be used to avoid a finite-time blow up. Continuity of the objective under changes in regimen parameters must be established in the relevant topology, including integral behavior around moving jumps. The measure-input benchmark states its separate linear convolution formulation explicitly.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

Observations provide information about states or parameters rather than access to all model components:

$$y_k = h_{obs}(x(t_k), \theta) + \varepsilon_k. \quad (3)$$

The observation map indicates whether an assay measures free or total drug, plasma or active compartment, response or deviation from baseline and a biomarker or an endpoint in clinical safety. A classical effect compartment is equilibration without adding drug mass to the PK balance [18]. Indirect-response models also represent delayed action [19]. Any deviation in response should be referred to the appropriate untreated baseline.

2.2. Benefit and Safety over the Evaluation Horizons

Benefit is the therapeutically desirable component of a response against a given comparator. Its sign will depend on the goal of the treatment: reducing an unwanted burden may be beneficial. A finite-treatment objective can combine an integrated benefit rate b with a terminal benefit Φ:

$$J(s; \theta) = \Phi(x_s(T_{eff})) + \int_0^{T_{eff}} b(x_s(t), t; \theta)\, dt. \quad (4)$$

The running and terminal terms should have compatible utility units and justified weights. If both describe the same outcome, their combination must avoid unintended double counting. The global outcome vector records the endpoints chosen for the question:

$$O(s) = (B, \Phi, Z_{max}, Z_{cum}, Y_{disease}, D_{tot}, N_{doses}, burden, \dots). \quad (4b)$$

Here B is the integral of the benefit rate over the efficacy horizon, and Φ denotes the terminal benefit evaluated at that horizon. For a single nonnegative injury state z, define peak and cumulative injury by

$$Z_{max} = \sup_{0 \le t \le T_{safe}} z_s(t), Z_{cum} = \int_0^{T_{safe}} z_s(t)\, dt. \quad (4c)$$

The disease endpoint and administration count are distinct from total amount. Each endpoint has its own definition and units. For multiple injury endpoints use corresponding coordinate blocks. Scalar utility requires declared transformations and weights, while constrained optimization retains selected endpoints as hard limits. Pareto efficiency means that no feasible strategy can improve one oriented endpoint without worsening another. These descriptions do not determine acceptable tradeoffs.

We distinguish the administration horizon T_{adm}, the efficacy horizon T_{eff} and the safety horizon T_{safe}. Drug input stops by T_{adm}, but active exposure and injury may continue afterward. Safety constraints apply to the whole relevant path, including both sides of dose jumps and the period after cessation:

$$g_j(x_s(t), t; \theta) \le 0, 0 \le t \le T_{safe}. \quad (5)$$

Feasibility is determined by the hard safety constraints. Applying a finite penalty to the objective may still leave a residual violation and therefore does not provide a safety certificate. If no regimen is feasible at the mandatory administered amount, reducing this amount defines a different problem. Alternatively, a formulation may impose an amount cap, include a dose cost or minimize the administered amount.

2.3. Distinct Finite Periodic and Relaxed Problems

In finite treatment at a fixed total administered amount (Finite) all schedules are compared with the same initial state, total amount and efficacy horizon. Drug remaining at the endpoint needs to be taken into account as well. Keeping the total amount equal allows the contribution of scheduling to be separately examined. However, this is a choice for the design of the comparison and is not a universal clinical requirement.

Periodic steady state at fixed mean administered input, denoted Periodic, evaluates a whole cycle and applies periodic closure to all states that are said to be stationary. For a cycle length τ, the amount administered is mean input multiplied by τ. The objective is mean benefit for that cycle. In periodic numerical comparisons, J denotes mean benefit ⟨b⟩ and mean injury is ⟨z⟩, rather than the finite cumulative injury Z_{cum}. Startup and cessation are separate safety questions. If physiology is

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

externally varying, then a compatible periodic formulation or an explicitly different long-run criterion is needed.

A relaxed exposure distribution problem (Relaxed) maximizes the static mean benefit over probability distributions of concentration with a prescribed mean. It ignores the kinetic transitions required to move between concentrations and generally ignores PD memory. Its solution is an upper bound on an actual schedule only if the schedule meets the same moment constraints and is evaluated with the same static objective.

Under periodic linear PK with constant bioavailability and clearance, mean concentration equals F_{bio} ū/CL, with F_{bio} denoting bioavailability and ū denoting the total administered amount per cycle, including boluses, divided by τ. The problem labels Finite, Periodic and Relaxed are used throughout to avoid ambiguity. This identity does not generally translate to finite treatment, nonlinear elimination, time-varying clearance or target-mediated disposition. We therefore state which quantity is fixed before comparing any exposure strategies.

2.4. Nested Admissible Classes

Let $\mathcal{A}_0$ contain the reference regimen and $\mathcal{A}_1$ allow for optimization of the amounts at fixed administration times. In the next class $\mathcal{A}_2$, selected administration times are also allowed to vary. Class $\mathcal{A}_3$ additionally includes the specified infusion structures and the choice of stopping and restarting. Thus, admissibility determines the actions that are possible and the restrictions on their implementation or use. Feasibility also requires satisfaction of the common trajectory-dependent safety conditions and any remaining resource limitations. Using g(s) to denote the aggregate residuals over all defined horizons,

$$\mathcal{A}_0 \subseteq \mathcal{A}_1 \subseteq \mathcal{A}_2 \subseteq \mathcal{A}_3. \quad (6)$$

$$\mathcal{F}_i = \{s \in \mathcal{A}_i : g(s) \leq 0\}, J_i^\star = \sup_{s \in \mathcal{F}_i} J(s). \quad (6b)$$

The same constraint map is used for all classes in order to keep the feasible sets nested. Each comparison uses the same objective, dose accounting, resource rule, horizons, model and safety requirements. Finite courses are evaluated from the same initial state and stationary comparisons use the same periodic initialization rule. An amount cap matches the available resources. An exact amount also matches the resources actually used. For nonempty sets with finite suprema, nesting means that the optimal values cannot decrease but strict improvement or attainment is not guaranteed. Figure 2 shows the available actions and the common safety requirements with them.

Class indices are present in each comparison and are not fixed biological categories in the LOGOS study. A larger class retains all predecessor actions. If timing, cessation, restart, feedback or other temporal choices are considered, they add value only if the smaller class restricts them. If terminal zero input is already allowed, relabeling cessation as an extra option does not change the class and does not lead to an increment in stopping. Finite, Periodic, and Relaxed label comparison problems. The relaxed input measures in a finite course proof instead describe delivery freedom.

Let $H_i = 1[\mathcal{F}_i \neq \emptyset]$ denote whether the feasible class is nonempty. Define the objective increment and the change in this feasibility indicator by

$$\Delta J_i = J_i^\star - J_{i-1}^\star, \Delta H_i = 1[\mathcal{F}_i \neq \emptyset] - 1[\mathcal{F}_{i-1} \neq \emptyset]. \quad (6c)$$

Objective increments require finite suprema and nonempty feasible sets. Empty sets should not have an infinite therapeutic gain. Under nesting, the feasibility indicator change would be zero or one, where one is the return of feasibility at the common target and constraints. Zero will not exclude a larger safety margin or robustness. A feasible witness proves non-emptiness and search failure does not prove emptiness. Failure of one constant strategy does not exclude the whole constant family.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

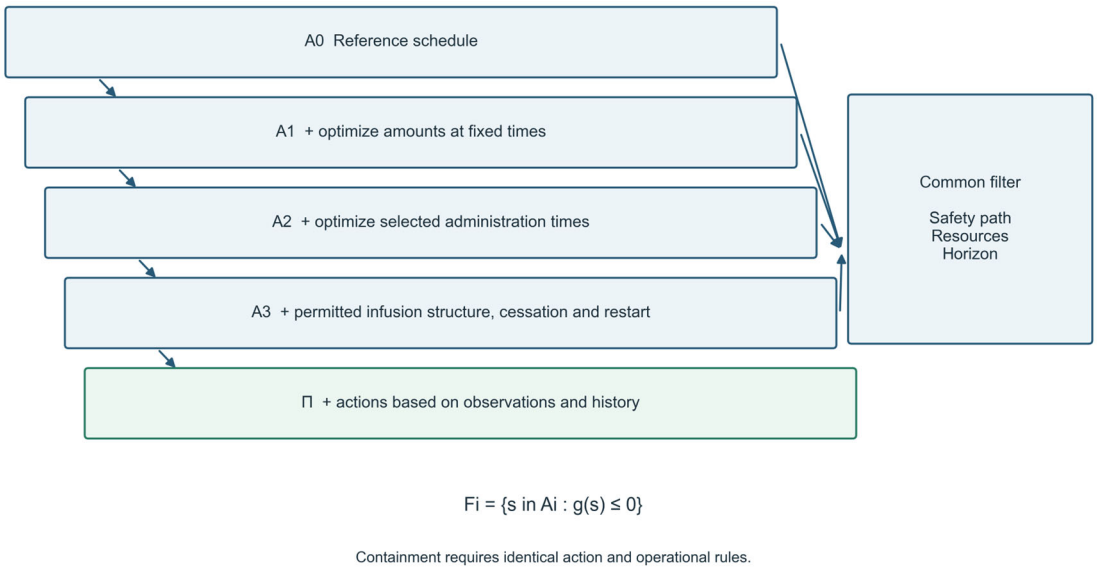


Figure 2. Nested regimen freedom and common feasibility. Every class keeps the actions of its predecessor. Safety and resource constraints filter every class with the same requirements. An open-loop schedule embeds into the policy class if observations are ignored, as long as actions, initial conditions and operations are matched.

The evidence labels are in Table 2. If the class supremum is unknown, the best feasible candidate and the search family and analytical bounds have to be reported separately. Differences between candidates are observed computational gains, not necessarily ΔJ_i. Relative improvements need a positive, meaningful baseline. Equal numbers cannot be used to prove equality of class optima without bounds or certificates.

Table 2. Evidence and claim hierarchy.

Evidence label	Required support	Permitted conclusion
Globally certified optimum	Feasible realization and a sufficient certificate over the declared class	No feasible strategy in that class has a better objective
Analytical bound	Proved inequality under explicit assumptions	Objective cannot exceed the bound under those assumptions
Relaxed upper bound	Valid relaxation and matched moments/objective	Opportunity bound; kinetic attainability remains unresolved
Best feasible candidate	Declared search family and independent feasibility check	Best identified strategy in the reported search
Feasibility result	Feasible witness or a proof of infeasibility	Non-emptiness or emptiness at the stated scope
Implementation check	Independent re-computation or numerical residual check	Computation agrees with the model; no global claim
Empirical hypothesis	Mechanistic rationale and testable endpoint	Requires external or prospective validation

A sufficient test of strict class value is a feasible expanded-class witness whose objective is larger than a valid upper bound for the smaller feasible class. If $J(s_{\mathrm{new}}) > U_{\mathrm{old}}$ and $J^{\star}_{\mathrm{old}} \leq U_{\mathrm{old}}$, then $J^{\star}_{\mathrm{new}} - J^{\star}_{\mathrm{old}} \geq J(s_{\mathrm{new}}) - U_{\mathrm{old}} > 0$. A global optimum for the larger class is not necessary for this existence claim. With numerical upper bounds, the conclusion inherits quadrature, candidate-feasibility and floating-point tolerances. Reporting a practically meaningful prespecified margin also

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

prevents a small increase in numerical value from being misinterpreted as an indication of therapeutic value.

Feasibility restoration requires the same target and path constraints in both classes. A feasible witness proves that the larger set is nonempty. An analytical smaller-class ceiling below the required efficacy target proves that the smaller target-feasible set is empty. The comparison can restore target feasibility even when both safety-only sets contain regimens. Resource saving is a separate question and requires an amount-minimization or resource-efficiency estimand.

2.5. Schedules and Adaptive Policies

Examples of observation-based comparators include warfarin adjustment using the international normalized ratio (INR) in MIPD trial simulations [5] and interruption and restart of abiraterone using prostate-specific antigen (PSA) in adaptive prostate-cancer therapy [12,13]. The monitored states and endpoints differ between these applications. In order to assess feedback in LOGOS, a fixed plan would be compared with an admissible policy using the same information schedule, outcome definitions and resource rules. If additional monitoring or a different drug budget were used, the resulting gain could not be attributed to feedback alone.

A policy makes decisions based on measurements, covariates and treatment history that are already available. Open-loop schedules are embedded only when a policy can ignore observations with the same actions and limitations. Feedback value compares feasible policy and open-loop optima under common expected utility to evaluate the outcomes. Information value compares the observation filtrations with measurement costs taken into account. Nonnegativity requires the richer policy to reproduce the poorer one, including costs and constraints. Neither one is estimated in the self-contained illustrations.

Receding-horizon implementation updates the estimated state and the parameter distribution, optimizes future decisions and performs the next action before repeating the calculation. This is compatible with MIPD and Bayesian forecasting [4]. A controller with only the current point estimate usually approximates the partially observed decision problem. When trying to understand an adaptive policy as optimal it is also important to understand the future value of observations and how actions affect learning.

2.6. Uncertainty Relevant to the Treatment Decision

A nominal optimum is obtained for a given parameter vector. Alternatively, a robust formulation can be used to maximize the worst-case objective over a given uncertainty set, while ensuring safety for each value of the parameters in this set. In an expected utility formulation, the results are averaged over a predictive distribution. A joint chance constraint is used to limit the probability of any violation over the selected safety endpoints and over the entire safety horizon. For ε in [0,1], this requirement is written as

$$P\{\forall j, \forall t \in [0, T_{safe}]: g_j(x_s(t;\theta), t;\theta) \leq 0\} \geq 1 - \varepsilon. \quad (7)$$

The probability covers the declared predictive uncertainty and the whole observation dependent path. Separate endpoint probabilities do not guarantee joint safety at the same level, thus, allocating endpoint error probabilities whose sum is at most ε gives a conservative union bound. Sampled-time checks need control of between-sample maxima. Expected safety, joint chance constraints and hard path guarantees are different requirements. For a feedback policy π, the trajectory in Equation 7 is the random closed-loop trajectory of π and its observation history. The probability must take into account all the modeled sources of uncertainty (including observation noise). For a fixed schedule, the same notation is used for its open-loop trajectory. The chance level is a model-based requirement, and not an empirically known probability of patient safety.

These formulations follow established robust and chance-constrained optimization principles [20,21]. Robust safety is conditional on the declared uncertainty set; chance-constrained safety is conditional on the predictive distribution and event definition. Tractable convex approximations require their own assumptions and are not general certificates for nonlinear PK/PD path constraints.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

$$\max_{s\in\mathcal{A}}\min_{\theta\in\Theta} J(s;\theta)$$
$$\text{subject to } g_j(x_s(t;\theta),t;\theta)\le 0,\forall j,\forall\theta\in\Theta,0\le t\le T_{safe}. \quad (7b)$$

The uncertainty sets and predictive distributions require justification, since the use of average parameters does not generally produce the average trajectory. Structural and practical identifiability determine which parameter combinations can be constrained by the observations [22]. For example, accurate estimation of plasma PK does not necessarily imply accurate estimation of repair or sensitivity states. When a schedule depends on these states, the robustness of the resulting decision must also be assessed.

To design a regimen, model analysis should be made on the predictions that influence the treatment decision. It is not necessary to identify all the microscopic parameters individually if the treatment decision remains the same between any of the admissible alternatives. A narrow concentration prediction alone does not establish certainty about the injury, recovery or disease evolution. Such predictions should be evaluated if they affect the treatment choice.

2.7. Optimization and Independent Feasibility Checking

In a first step, the comparison problem is defined and the model modules are selected. Units and mass balance are then checked, together with any available analytical bounds or infeasibility results. Optimization is performed over a declared finite parameterization or control class. Each feasible reference candidate is retained in the expanded search. In this way, a numerical search failure is not interpreted as a biological disadvantage of the additional design freedom. Following optimization, the selected schedule is evaluated independently with respect to dose jumps, concentration peaks and injury extrema. The same assessment also includes initial transients, cessation tails and restarts. Numerical stability is examined through refinement and the use of other solvers. In the case of nonconvex searches, the starting points and admissible families must be reported. It should be stated that Pontryagin necessary conditions may identify the structure of a solution, but do not by themselves establish global optimality.

For linear PK at fixed administration times, a concave benefit function may retain a convex optimization structure. Suitable convex injury constraints can also be included when repair is linear. However, variable administration times, sigmoidal response, nonlinear disposition and feedback generally remove this structure. Another issue concerns the existence of an optimizer. A continuous objective on a compact parameter class can attain a maximum, whereas an unrestricted class may have only a supremum.

For fixed administration times and a differentiable concave objective on the nonnegative dose simplex with total amount D, a feasible dose vector d gives the upper bound $J^* \le J(d) + D \max_i \partial_i J(d) - \nabla J(d)\cdot d$. This follows from the supporting-hyperplane inequality for a concave function and the difference between this bound and J(d) is an optimality gap bound in exact arithmetic. A numerical evaluation also needs to take into account derivative and integration errors. Solver convergence alone does not give this bound.

2.8. Python Implementation and Numerical Analysis

The numerical analyses were performed in Python 3.12, with Jupyter used as an interactive working environment. The archived reproducibility environment specifies Python 3.12.14, NumPy 2.3.5 for array computations, SciPy 1.17.1 for numerical integration and optimization, and Matplotlib 3.10.9 for plotting. The complete Python implementation and the scripts required to reproduce the numerical results presented in this study are publicly available on GitHub (https://github.com/vkaralis/LOGOS-theory-general) [23]. The repository includes reproducibility instructions, pinned software dependencies, and independent verification procedures, allowing the reported calculations and results to be reproduced and examined.

The illustrative models use normalized concentration, amount and time scales. Their caps are not clinical thresholds. Here EC50 denotes the concentration producing half of the maximum effect,

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

Emax. In our finite calculations, we set V=k=EC50=Emax=D=1, T=3 and start with zero predose concentration. For bolus-only schedules, the discrete intravenous amounts sum to one, with later administration times ordered in [0,3]. When continuous input is included, the combined bolus and infusion amount is one. Concentration and continuous-input caps are one. Injury obeys $\dot{z}=C^2-0.5z$ from zero, with cap 0.5 over the complete future. Bolus objectives use exact exponential-segment integrals. Sequential least squares programming (SLSQP) uses tolerance 10^{-13} and up to 1,000 iterations, while timing uses 24 starts with seed 20261003. The plateau is obtained by a bracketed root solve.

For the periodic calculations, k=V=1 and the mean input is 0.5. Each regular bolus with interval τ has amount 0.5τ and peak concentration $0.5\tau/[1-\exp(-\tau)]$. Sensitivity and injury are initialized on the periodic orbit using affine maps. In the original solve_ivp calculations, the relative and absolute tolerances were 2×10^{-10} and 2×10^{-12}, respectively, and 801 points were used for the peak search followed by local refinement. The interval search starts with 121 points on [0.08,12]. The detected constraint boundaries and objective neighborhoods are then refined and feasible endpoints are kept. This procedure does not exclude every undiscovered component or certify a global periodic optimum. The factorial laws and local parameter grid are shown together with their results.

2.9. Independent Verification and Numerical Stability

The verification process distinguishes exact formulas, runs of the original code, reconstructions and historical comparisons. For finite treatment, trajectories are integrated across each dose jump and infusion transition. Periodic calculations are analyzed by integrating-factor quadrature and more accurate peak searches. Formula checks include the short- and long-horizon Emax branches, coincident injury rates and Hill-envelope caps. Full-precision outputs, controls and provenance records are kept in the repository. For historical campaigns when the original Flagship source was not available, the Recheck scripts reconstruct the printed equations with no claim of the same execution of the missing code. Agreement of the calculations supports the implementation while global and full-tail guarantees follow from the analytical arguments.

The public GitHub repository provides the research implementation and reproduction instructions [23]. It includes independent reconstructions of the declared mathematical models and numerical verification procedures. Public reconstructions and historical executions are distinguished when their verification results are reported.

3. Results

The results are presented in order to examine how the properties of the model affect the choice of regimen. Figure 3 summarizes the corresponding candidate strategies and the assumptions required in each case. When several mechanisms are included in the same model, their combined effects must be assessed.

3.1. Analytical Bounds and Mechanistic Insights

Let us consider periodic linear PK at fixed mean input, with directly acting concave benefit b(C) and convex toxic drive v(C). Let us suppose that the same constant concentration can be achieved in the input class and there is no additional memory, adaptation or time-dependent weighting. With angle brackets indicating cycle averages, the clearance balance fixes the mean concentration and Jensen's inequality gives

$$\langle b(C)\rangle \le b(\langle C\rangle), \langle v(C)\rangle \ge v(\langle C\rangle). \quad (8)$$

Constant exposure attains both bounds. With linear injury repair, $\dot{z} = v(C) - \varrho z$ and $\varrho > 0$, periodic balance also gives mean injury $\langle z\rangle = \langle v(C)\rangle/\varrho$. So smoothing maximizes the mean benefit and minimizes mean drive and mean injury in this system. Strict concavity gives a strictly lower mean benefit for truly variable exposure. In this same periodic linear repair model, the peak injury is at

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

least mean injury, so the constant exposure also minimizes periodic peak injury: $\sup z \geq \langle z \rangle \geq v(\langle C \rangle)/\varrho$, with equality for constant exposure.

This negative control applies to periodic linear PK with direct concave benefit and attainable constant input. This does not account for finite startup, terminal outcome, injury peaks under other repair laws, nonlinear clearance or exposure-dependent repair. Thus, loading and final cessation might still be important even when steady constant exposure does limit the benefit.

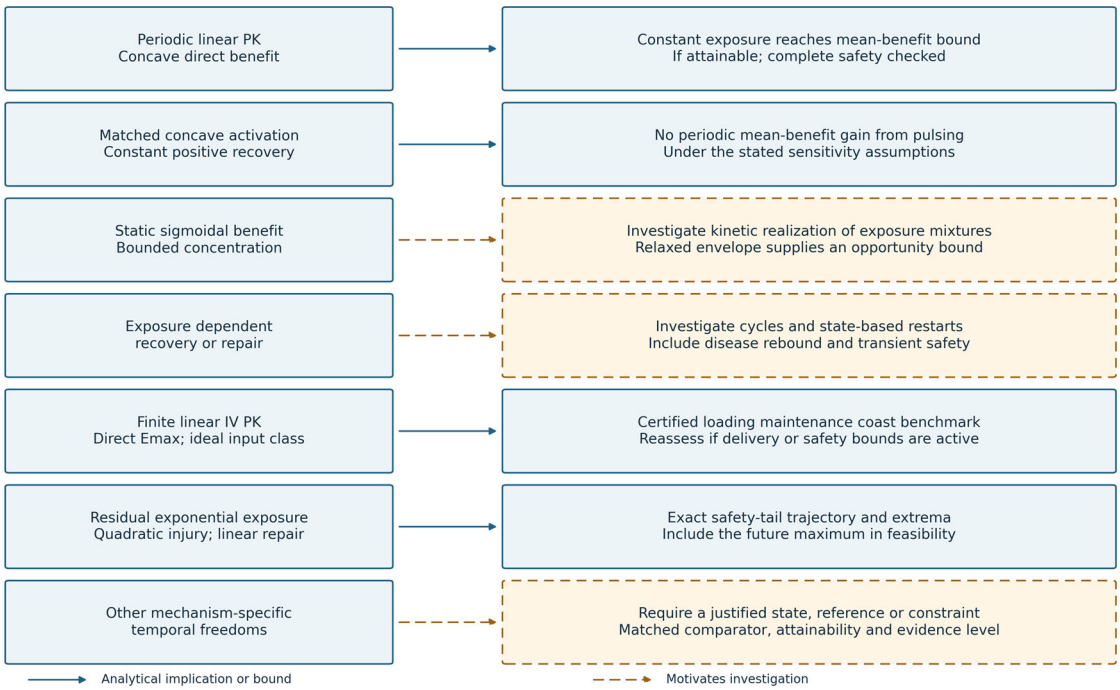


Figure 3. Mechanism-to-strategy map. Solid blue arrows and blue result boxes indicate analytical implications or bounds under the stated assumptions. Dashed amber arrows and amber boxes indicate directions for exploration and do not establish superiority of an achievable schedule. There are more temporally structured freedoms that can be explored if a mechanistically meaningful state, reference or constraint is given. These extensions require their own matched comparator, attainability analysis, and evidence level. For the sigmoidal static benefit, the concave envelope provides a relaxed bound, and its kinetic realization needs to be optimized and verified. PK, pharmacokinetics; IV, intravenous; Emax, maximum effect.

The interval [0,1] is the invariant physical domain for sensitivity a. Suppose recovery is independent of exposure, benefit and desensitization share a bounded nonnegative concave activation q(C) and exposure is periodic. Assume $k_{rec} > 0, k_{des} > 0$ and, for the periodic solution, $a(t) > 0$ for every t in the cycle. This strict positivity permits division by a and integration of its logarithmic derivative:

$$\dot{a} = k_{rec}(1 - a) - k_{des}q(C)a, b = aq(C). \quad (9)$$

Periodic sensitivity balance and logarithmic balance give an upper bound on mean benefit. Define $G(y) = \frac{k_{rec}y}{k_{rec}+k_{des}y}$ for $y \geq 0$. Then

$$\langle b \rangle \leq G(\langle q(C) \rangle) \leq G(q(\langle C \rangle)). \quad (10)$$

At the same mean concentration, attainable constant exposure with stationary sensitivity attains the final bound (Appendix A). Under these assumptions, desensitization supplies no mean-benefit justification for pulsing. The result depends on the activation and recovery laws: different benefit and desensitization activations, exposure-dependent recovery, nonconcave activation or finite transients can change it. Independent injury dynamics can still require a pause for safety.

For a static Hill benefit, assume $h > 1$, $EC50 > 0$, $Emax > 0$ and a concentration interval $[0, C_{max}]$ with $C_{max} > 0$. The response and its tangent concentration are

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

$$b(C) = \frac{E_{max}C^h}{EC_{50}^h+C^h}, C^\dagger = EC_{50}(h-1)^{1/h}. \quad (11)$$

The Relaxed problem identifies a favorable concentration by maximizing the benefit per unit exposure. If the interval contains C†, the concave envelope joins zero to C† and travels in the concave part of the Hill curve. For means below C†, a mixture of zero and C† reaches the envelope; for larger means, constant exposure reaches it. If the imposed concentration cap C_{max} is below C†, the envelope on $[0,C_{max}]$ is instead the straight chord joining zero to $b(C_{max})$ and the relaxed optimum is a zero/C_{max} mixture. This is because b(C)/C is increasing up to C†. The result holds for mean concentrations within the stated interval. A mean above C_{max} is infeasible.

For h = 3, EC50 = Emax = 1 and mean concentration 0.3, the relaxed envelope is 0.158740 vs. constant benefit 0.026290. This difference bounds an opportunity, and it is not an attainable bolus gain. Elimination prevents immediate removal of drugs and transitions consume exposure. Memory, injury, and administration limits may also affect the result. The static bound only transfers with the same mean exposure and objective. Kinetic realization is another optimization problem.

Recovery can be suppressed in active exposure while low exposure can restore responsiveness or tissue function. For a prescribed exposure trajectory, a useful model has nonnegative rates r, d, v and ϱ:

$$\dot{a} = r(C)(1-a) - d(C)a, \dot{z} = v(C) - \rho(C)z, b = af_B(C). \quad (12)$$

Here $f_B(C) \geq 0$ denotes the benefit activation function. For locally integrable rates on a prescribed periodic exposure of period τ, sensitivity remains in [0,1] and injury remains nonnegative. Their cycle maps are $a(\tau) = A_a\ a(0) + B_a$ and $z(\tau) = A_z\ z(0) + B_z$, with

$$A_a = \exp[-\int_0^\tau (r(C(t)) + d(C(t)))dt], A_z = \exp[-\int_0^\tau \rho(C(t))dt]. \quad (12b)$$

The nonnegative forcing terms are

$$B_a = \int_0^\tau \exp[-\int_\xi^\tau (r(C(t)) + d(C(t)))dt] r(C(\xi))d\xi. \quad (12c)$$

$$B_z = \int_0^\tau \exp[-\int_\xi^\tau \rho(C(t))dt] v(C(\xi))d\xi. \quad (12d)$$

A unique attracting periodic state is $a(0) = \frac{B_a}{1-A_a}$ and $z(0) = \frac{B_z}{1-A_z}$ if each corresponding multiplier is less than one (that is, its nonnegative rate integral is strictly positive). For sensitivity, $A_a = 1$ implies r = d = 0 almost everywhere, hence $B_a = 0$ and every initial sensitivity in [0,1] is periodic. For injury, $A_z = 1$ implies ϱ = 0 almost everywhere. If $B_z > 0$, the injury increases by B_z per cycle and there is no finite periodic state. If $B_z = 0$, then every nonnegative initial injury is periodic. These cycle reductions do not establish a globally optimal cycle.

It should be stated that stopping administration does not necessarily mean low active exposure or biological recovery. The changes may be delayed by depot absorption, metabolites, equilibration or persistent binding. Therefore, a restart threshold is a candidate time rather than an optimum. The safety of the next dose and the possibility of disease rebound need to also be considered. In addition, state-dependent rates may invalidate the affine reduction. Even when periodic closure is established, startup, each restart and final withdrawal require separate safety checks.

3.2. Globally Optimal Finite-Treatment Strategy

The ideal model is used as an analytical benchmark in order to study the optimization process overall. Let V, k, EC50, Emax, D and T be positive. Volume V is constant, k = CL/V, predose concentration is zero ($C(0^-) = 0$) and benefit is described by Emax C/(EC50 + C). A nonnegative administered-dose measure M of total D is assumed, including an initial intravenous bolus. No adaptation or terminal payoff is taken into account and there is no active rate, concentration or injury constraint. Over the efficacy horizon T, concentration is described by dC = −kC dt + dM/V.

If T is more than k^{-1} log(2 + EC50 V/D), the optimal concentration is: the initial load to c, the maintenance at c for time L and the input-free coast for time δ. The parameters satisfy

$$\delta = \frac{1}{k}\log(2 + \frac{EC_{50}}{c}), L = T - \delta, D = Vc + CLcL. \quad (13)$$

Loading amount is Vc and maintenance rate is CLc. For a shorter time horizon, the optimum is a single initial bolus followed by coast at the threshold, the maintenance period has zero duration.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

This is a global optimum over all nonnegative administered-dose measures of the stated total amount, including competing bolus and infusion schedules. The supporting marginal-value inequality in Appendix B is sufficient to establish this result. Thus, the conclusion follows from the analytical certificate and does not depend on convergence of a numerical optimizer. Strict concavity also establishes uniqueness of the concentration trajectory up to sets of zero-time measure. The coast exploits residual drug for benefit, with no recovery mechanism in this model.

Coast is a benefit-bearing time interval without new input. Residual benefit alone is not sufficient to establish a favorable net benefit-harm outcome, but therapeutic interpretation requires an explicitly defined criterion for net value. The present global certificate applies to the integrated Emax-benefit objective, fixed administered amount and model assumptions stated above.

A course starting drug free has boundary conditions absent from the periodic constant-exposure problem. In this case, the finite solution provides an algorithmic benchmark in the mentioned measure input class. Active constraints, oral absorption, adaptation, terminal rewards and uncertain/nonlinear PK can alter the optimal structure and need to be considered separately.

Although both analyses yield loading/infusion structures, Katriel minimizes antimicrobial exposure subject to an eradication target [16], whereas this benchmark maximizes integrated Emax response at fixed administered amount and horizon. The certificates are specific to their goals.

3.3. Safety After Treatment Cessation

Assume C_0, z_0, $\alpha \geq 0$, $k > 0$ and $\varrho \geq 0$; constant parameters. One-compartment residual exposure $C(t) = C_0\exp(-kt)$; injury $\dot{z} = \alpha C^2 - \varrho z$ and no further administered or absorbed input. Time zero is cessation, with injury z_0. For $\varrho \neq 2k$, the unique injury trajectory is

$$z(t) = (z_0 + A)e^{-\rho t} - Ae^{-2kt}, A = \frac{\alpha C_0^2}{2k-\rho}, \rho \neq 2k. \quad (14)$$

The future maximum must include the current injury, every positive stationary time, and the limiting endpoint. Appendix C is devoted to the coincident-rate case and a conservative bound. Even a short observation window may miss the maximum even if no drug is being administered. This safety tail is part of the feasibility analysis of the whole regimen, including loading, cessation and each proposed restart. With $\varrho = 0$ injury is nondecreasing towards $z_0 + \alpha C_0^2/(2k)$. For $\alpha C_0^2 > 0$, the supremum is approached only asymptotically, otherwise the injury is constant. If $\varrho > 0$, the limiting injury is zero. Failure scenarios are continued depot input, active metabolites, delayed effect-site equilibration or nonlinear repair. The advice for our implementation of LOGOS is to verify the relevant full tail based on the correct model instead of extrapolating this formula.

3.4. Incremental Value of Regimen Optimization

We set $V = k = EC50 = Emax = D = 1$ and $T = 3$ with zero predose concentration. All schedules share the objective of cumulative direct Emax benefit on [0,3]. Discrete doses are intravenous, total amount is one, concentration must not exceed one, and injury $\dot{z} = C^2 - 0.5z$ starts at zero and must remain below or equal to 0.5 over the complete future tail. The continuous administration rate is bounded by one. These caps are nonbinding in this illustration and isolate design freedom rather than a safety-driven advantage.

For this benchmark, $\mathcal{A}_0$ is the equal three-bolus schedule at times 0, 1 and 2; $\mathcal{A}_1$ allows nonnegative amounts at those times, including zero amounts. $\mathcal{A}_2$ retains an initial dose at zero and allows the two later times to satisfy $0 \leq t_2 \leq t_3 \leq 3$. Coincident doses are combined. Class $\mathcal{A}_3$ contains these schedules and additionally permits nonnegative continuous input up to rate one. All classes retain exact total amount one. The unrestricted measure-input optimum is feasible in $\mathcal{A}_3$, so its certificate also applies to this constrained class. The result of Section 3.5 therefore supplies a certified reference for this comparison. The loading–maintenance–coast schedule attains the maximum integrated benefit over the declared class, while the bolus schedules in Table 3 show the values obtained with more restricted administration choices. The analytical certificate establishes the final

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

optimum. The amount-and-time result remains a feasible numerical candidate. The full-tail safety bound in Appendix D verifies injury feasibility independently of the optimization grid.

Table 3. Incremental value in the finite-treatment comparison.

Class and freedom added	Benefit J	Observed gain	Relative gain	Evidence	Feasibility	Resource rule	Realized D_{tot}
$\mathcal{A}_0$: equal doses at 0, 1, 2	0.629606	—	—	Reference	Certified	exact-dose	1
$\mathcal{A}_1$: optimize amounts	0.672945	0.043339	6.88%	finite-dimensional numerical bound	Certified	exact-dose	1
$\mathcal{A}_2$: add two variable times	0.677043	0.004098	0.61%	feasible numerical witness	Certified	exact-dose	1
$\mathcal{A}_3$: add continuous input	0.679553	0.002509	0.37%	analytical global certificate	Certified	exact-dose	1

The adjacent observed gains are shown in Table 3. Each candidate is compared to the previous candidate. If the class optimum is unknown, these differences cannot be seen as certified class increments. The common mass-kernel bound certifies feasibility and the feasible-class indicator is still one in every row. The expanded benchmark gains about 0.98% over optimized fixed-time amounts and 0.37% over the timing candidate. These are modest, model-specific effects, calculated from unrounded objectives. For the fixed-time candidate, the supporting-hyperplane bound in Section 2.7 agrees with J = 0.6729453961 to the displayed precision. The numerical marginal values are approximately 0.49879045, 0.49879045 and 0.47836668. These floating-point checks do not provide an interval-arithmetic certificate.

With administration times fixed, the optimizer allocates approximately 0.628402 and 0.371598 at times zero and one, while the third amount is reduced to zero. When timing is also varied, the search gives amounts of approximately 0.529424, 0.235288 and 0.235288 at times 0, 0.587748 and 1.175496. This result is the best candidate from 24 successful multistart runs and has no global timing certificate. For the expanded regimen, the globally certified solution within the stated admissible class consists of a load of 0.400534, maintenance at that concentration for 1.496664, and coast for 1.503336. Its value is established by the analytical result. Figure 4 shows the corresponding concentration and cumulative-benefit trajectories. All these schedules satisfy the injury cap analytically. For any nonnegative dose measure of total mass one with k = V = 1, concentration is at most one and the full-tail integral of C^2 is at most one half. Injury cannot exceed this integral because repair removes a nonnegative quantity. Appendix D gives the mass-kernel bound. So feasibility is not only deduced from optimizer time points.

3.5. Periodic Regimen Comparisons and Numerical Verification

In the matched-activation periodic model, q = C/(1+C), k_{rec} = 0.7 and k_{des} = 2. Constant exposure gives mean benefit 0.170732, and every examined regular-bolus interval lies below the analytical bound. In a different recovery/repair model, $q = C^3/(1+C^3)$, $r = 0.7(1-q)^4$, d = 2q and b = aq. The drive for injury is $q_T = C^8/(0.35^8+C^8)$, and $\dot{z} = q_T-(1-q_T)z$. The common caps are concentration four and injury 3.5. The bolus interval chosen 2.870281 gives mean benefit 0.091278, peak concentration 1.521377 and peak injury 1.738049. Constant exposure gives mean benefit 0.073656 but stationary injury 17.346653 breaches the cap. The chosen cycle is periodically feasible but the specified constant comparator is infeasible. This comparison does not show that each constant input is

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

unsuccessful or that the cycle is optimal globally. We then change the injury law with $\dot{z} = C^2 - 0.5z$ and the injury comparison is reversed. The same bolus leads to a mean injury of 0.803807 instead of 0.5 for the constant exposure. In Figure 5 we show how our matched activation with constant recovery is compared to the exposure-dependent recovery/repair model. Startup and cessation require separate safety assessments.

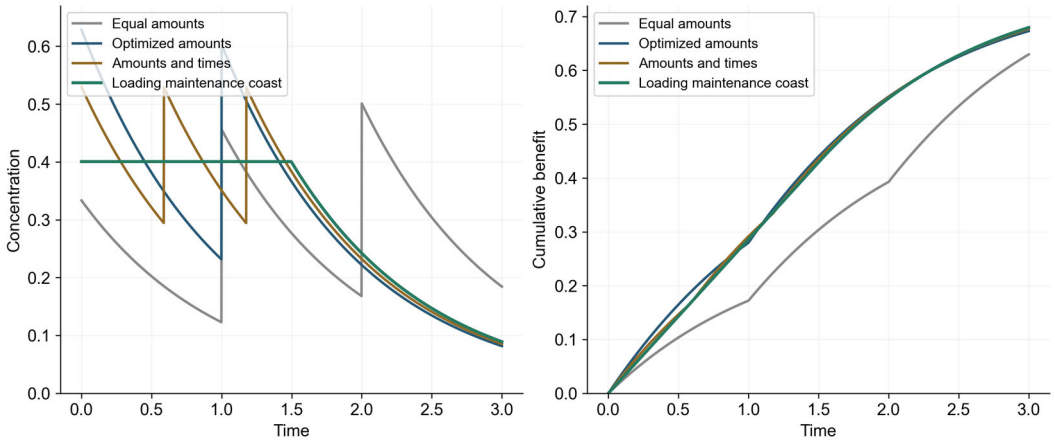


Figure 4. Finite-treatment comparison at total amount one and efficacy horizon three. Left, concentration trajectories for the reference and expanded design classes. Right, cumulative maximum-effect (Emax) benefit. Safety and continuous-rate limits are the same, the continuous-rate bound is for infusion rather than ideal dose atoms. The $\mathcal{A}_3$ trajectory has a global analytical certificate in the class mentioned. The timing result is a multi-start candidate.

In this study, it is performed a 2 × 2 numerical comparison with the same one-compartment PK, mean input 0.5, concentration cap four, injury cap 3.5 and regular-bolus intervals [0.08,12]. The factors are concave versus sigmoidal activation and exposure-independent versus exposure-dependent recovery/repair. Each cycle has dose 0.5τ and is evaluated at its periodic state. The activation laws and the common injury-generation law are

$$q_{conc}(C) = \frac{C}{1+C};\ q_{sig}(C) = \frac{C^3}{1+C^3};\ q_T(C) = \frac{C^8}{0.35^8 + C^8}$$

All cases use sensitivity loss $d = 2q$ and benefit $b = aq$. In Eq. (12), $f_B = q$ and $v = q_T$. The exposure-independent factor has $r = r_0, \rho = \rho_0$ the exposure-dependent factor has

$$r(C) = r_0(1 - q(C))^4, \rho(C) = \rho_0(1 - q_T(C)), r_0 = 0.7, \rho_0 = 1$$

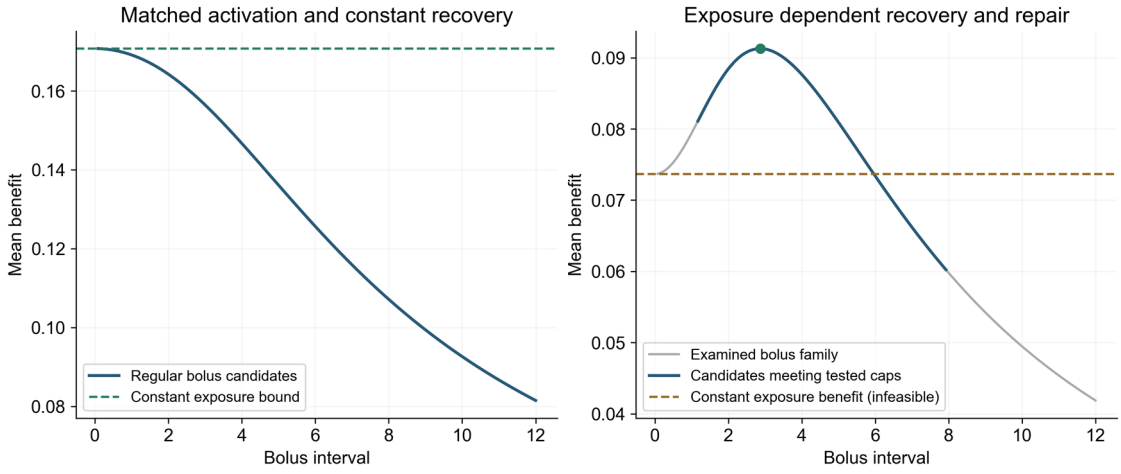


Figure 5. Regular-bolus scans at the same mean input. Left, we matched activation with constant recovery, where constant exposure achieves the analytical mean-benefit bound. Right, a separate model with sigmoidal activation

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

and exposure-dependent recovery and repair. Note that the darker part satisfies the tested concentration and periodic injury caps. Constant exposure in the right panel is infeasible under the injury cap. The point marks the locally refined candidate in the examined family.

Table 4 reports both strategies in each case, and Figure 6 shows the searched benefit curves. In case A, feasible constant exposure gives mean benefit 0.170732, slightly above the best tested bolus value 0.170721 at the lower search limit $\tau = 0.08$. This order is consistent with the matched activation analytical bound. Case B also has feasible constant exposure, but the bolus we have chosen increases the average benefit from 0.084337 to 0.111507 (32.22%) at $\tau = 3.159490$. An intermittent benefit advantage therefore occurs here without exposure-dependent recovery or repair.

The stationary constant injury in cases C and D is 17.346653 and violates the cap, whereas the selected bolus candidates are feasible. In case C we use concave activation and the mean benefit is 0.085510 at $\tau = 4.442336$. In case D we repeat the above example with the benefit being 0.091278 at $\tau = 2.870281$. So we do not need the sigmoidal activation to account for the recovery/repair-driven feasibility difference in this model. In case B we compare the benefit between the feasible strategies and in cases C and D we recover the periodic feasibility of a bolus against a fixed mean constant comparator. None of these searches gives a global periodic optimum or a statistical interaction. Response levels and rates are different for each case as indicated.

Table 4. Periodic factorial comparisons at injury cap 3.5 and mean input 0.5. A–D are concave/independent, sigmoidal/independent, concave/dependent and sigmoidal/dependent. Listed intervals are selected feasible witnesses in the searched regular-bolus family, not global cycle optima.

Case and strategy	Mean benefit	Peak injury	Mean injury	Peak C	Interval τ	Feasible in cycle	Resource rule	Amount per cycle
A constant	0.170732	0.945494	0.945494	0.500000	—	Yes	fixed mean input	0.5τ
A bolus	0.170721	0.944709	0.944598	0.520267	0.080000	Yes	fixed mean input	0.5τ
B constant	0.084337	0.945494	0.945494	0.500000	—	Yes	fixed mean input	0.5τ
B bolus	0.111507	0.766421	0.490731	1.649774	3.159490	Yes	fixed mean input	0.5τ
C constant	0.057260	17.346653	17.346653	0.500000	—	No	fixed mean input	0.5τ
C bolus	0.085510	1.789967	0.862802	2.247619	4.442336	Yes	fixed mean input	0.5τ
D constant	0.073656	17.346653	17.346653	0.500000	—	No	fixed mean input	0.5τ
D bolus	0.091278	1.738049	1.124215	1.521377	2.870281	Yes	fixed mean input	0.5τ

A local grid varies baseline recovery over {0.56,0.70,0.84}, repair over {0.8,1.0,1.2} and injury cap over {2.8,3.5,4.2} while keeping activation, PK, mean input and interval family fixed. In all 27 combinations, constant exposure is infeasible and a selected bolus is feasible. Constant injury ranges from 14.455544 to 21.683316. Selected bolus peak injury ranges from 1.583994 to 1.998673, interval from 2.764844 to 3.012607 and mean benefit from 0.083470 to 0.097223. Selected candidates lie below the strictest local cap.

The 18 broader cap stress tests show where these classifications change. At cap one, no feasible bolus is detected in nine recovery/repair pairs. At cap twenty, both strategies are feasible in six pairs, while three still admit only a detected cycle. Search failure cannot establish universal infeasibility, and the deterministic ±20% grid is not an estimated uncertainty distribution. Recovery has no direct effect on injury in this model. Disease rebound or another injury law could change the classifications (Figure 7).

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

For the cessation-tail model, set $C_0 = \alpha = k = 1$, $\varrho = 0.2$ and $z_0 = 0$. Injury at time 0.05 is 0.047340, below the illustrative cap of 0.1. The exact future maximum reaches 0.387132 at time 1.279214, above the cap after the early assessment window. In this example, we check the analytical safety-tail calculation without numerical regimen optimization. Figure 8 illustrates this delayed injury peak.

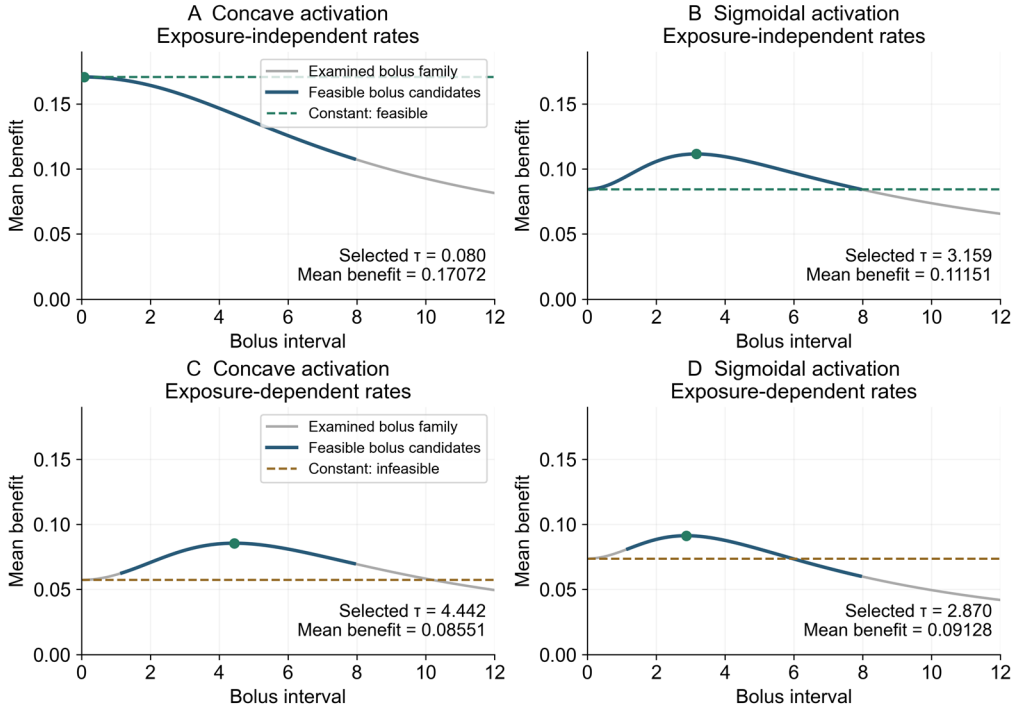
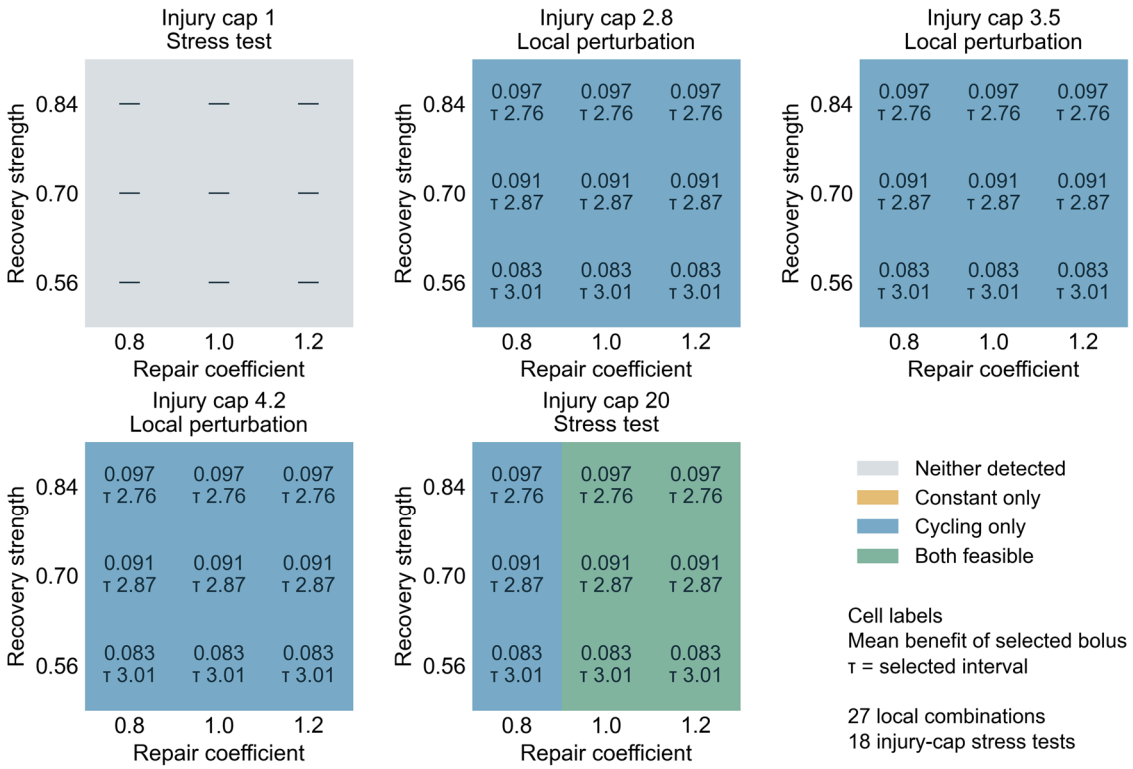


Figure 6. Factorial separation of activation shape and recovery/repair. Gray curves show the family of regular-bolus we considered. Blue segments satisfy the common concentration and injury caps. Dashed lines indicate constant-exposure benefit, with feasibility stated in the legend. The selected feasible candidates are indicated by the points. Case A selects the lower interval bound. We are not making global cycle certificates but conditional periodic comparisons.



© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

Figure 7. Feasibility map for constant exposure and the bolus family analyzed. Caps 2.8, 3.5 and 4.2 define the 27-point local grid. Caps one and twenty are stress tests. Labels indicate selected bolus mean benefit and interval τ and a dash indicates that there is no feasible candidate detected. Constant-only cells do not occur. Sampled boundaries are not analytical bifurcation curves or global infeasibility certificates.

In the original verification record, agreement was obtained within 5.02×10^{-13} for the finite objectives and within 2.89×10^{-9} for the periodic candidate metrics. Grid refinement did not change the feasibility classifications. These statistics refer to the historical campaigns. A separate archived reconstruction from 11 October 2026 passed 37 periodic checks with factorial scans and 27 local/18 stress test classifications. The 13-check benchmark suite provides independent reconstructions of the finite objectives, supporting conditions and injury tails. These reconstructions support the printed values, but do not establish identity with the unavailable historical source, arrays or residual statistics.

4. Discussion

4.1. Methodological Contribution of LOGOS

The contribution of LOGOS is to make the value of an additional regimen decision interpretable under a common comparison design. Established optimal-control and PK/PD optimization methods can be used to solve the individual problems. LOGOS then links their results through nested admissible classes, matched resource and safety requirements, mechanistic explanation and an explicit distinction between candidate improvement and certified optimality. For example, optimizing dose amounts at fixed times and subsequently allowing administration times to vary addresses two different decision classes. Keeping the model, objective, resources and horizons unchanged allows the resulting difference to be attributed to the additional choices within that model. Independent checks determine whether the proposed input is attainable and whether safety holds throughout the relevant trajectory.

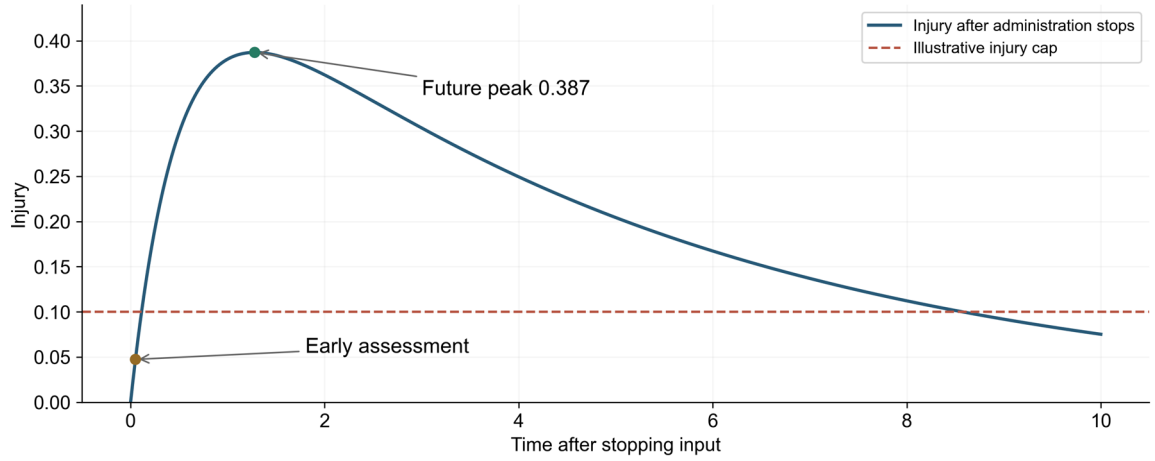


Figure 8. Injury after administration has stopped. The residual exponential exposure still generates injury that peaks after initial assessment. The plotted safety horizon is $T_{safe} = 10$. The exact formula also shows that injury decreases after its peak and tends to zero outside this horizon. The cap is illustrative and dimensionless. Safety of the whole regimen must include the relevant future trajectory or a valid tail certificate.

4.2. Applications in Pharmacotherapy

For example, in a possible docetaxel application, the amounts could first be optimized at fixed administration dates. In a subsequent step, the dates and observation-based treatment holds could also be allowed to vary. The tumor objective, resource rule, neutrophil constraint, monitoring and safety horizon would remain the same throughout the comparison. Existing myelosuppression

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

models provide relevant states for delayed injury [24]. This example describes a proposed application of the procedure and is not a clinical LOGOS result.

4.3. Mechanistic Interpretation of Regimen Choices

The analytical examples show that it is important to specify the mechanism and comparison conditions. Under the periodic assumptions studied here, constant exposure is sufficient for concave benefit and this result is preserved in the matched desensitization model. In the case of sigmoidal curvature, a relaxed opportunity may exist even without recovery. For a finite course, loading followed by coast can be preferred even when concave response is present. The injury mechanism also plays a role in interpretation. For example, exposure-dependent repair may favor cycling, whereas another injury endpoint can reverse the comparison. In the factorial experiment, recovery and repair were changed jointly, so that they cannot be separately accounted for from this comparison.

4.4. Mathematical Contribution and Global Optimality

The finite-treatment result in Sections 3.5 and 3.7 is the strongest mathematical result in this paper. The loading–maintenance–coast regimen is the global optimum in the declared class of measure inputs, and the same is true for the constrained numerical example, because that regimen is within the restrictions. The value of 0.679553 is therefore a certified reference for the more restricted schedules. The last adjacent numerical gain is only 0.37%, while the importance of the result lies in the determination of the optimum and the administration structure under explicit assumptions. The significance of such a gain depends on the treatment burden and uncertainty, neither of which is estimated in this dimensionless model. The ideal dose atoms and abstract rate bounds also have to be evaluated as to whether they can be implemented in practice.

4.5. Limitations and Clinical Translation

The conclusions of each comparison depend on the model, objective, resources, horizons and type of evidence. Increasing the available regimen choices may leave a simple optimum unchanged. Similarly, a relaxed mixture may not be physically attainable, and periodic feasibility alone does not guarantee safety at startup or withdrawal. It should therefore be stated that normalized response gains and model-based safety states do not establish patient benefit or clinical safety.

For clinical applications, the comparator should be consistent with current practice, and the observations should constrain the PK, benefit, injury and recovery predictions relevant to the decision. Agreement with fitted treatment histories alone is not sufficient to qualify predictions for other regimens that were not observed. The need to distinguish clinical objectives from model assumptions is also illustrated by the literature. In the Beta-Lactam Infusion Group (BLING) III trial, continuous β-lactam infusion in sepsis did not significantly reduce 90-day mortality in the primary unadjusted analysis, although clinical cure was higher than with intermittent infusion [25]. These results refer to specific endpoints and do not validate the direct-concavity assumptions used in the present study. Similarly, salbutamol tolerance supports consideration of a sensitivity state but does not validate the matched-activation law [26]. Concentration–kill data provide a reason to examine response curvature without establishing an attainable exposure mixture [27]. In the case of vancomycin, loading and maintenance also motivate comparisons of delivery strategies using qualified exposure and renal-safety criteria [28].

4.6. Novelty, Advantages, and Potential Applications of LOGOS

The main novelty of LOGOS is the unification of a mathematical framework that goes beyond just dosage optimization. The key original results and their clinical implications are: I) LOGOS assesses whether therapeutic outcomes can be enhanced by adjusting drug dose, administration time, infusions, or treatment interruptions. In practice, this could help determine if adjusting dose alone is

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

enough or if timing, duration, or pattern of treatment can help improve patient safety without increasing drug use or exposure to the drug. II) Making better use of each administered dose namely, LOGOS mathematically provides an optimal treatment plan consisting of an initial loading dose, maintenance period, and a last drug-free period under certain PK/PD assumptions. III) Protecting patients from toxicity that develops after treatment stops since LOGOS extends safety assessment beyond the active treatment period by accounting for residual drug exposure and delayed adverse effects. IV) LOGOS considers how drug response and tolerance, recovery of drug sensitivity, and tissue repair are modulated by the drug-response characteristics and the differential benefits of continuous and intermittent use. V) Supporting safer and more individualized treatment decisions as it provides a systematic way to compare alternative treatment strategies, determine whether they meet the safety criteria, and determine whether their expected benefits are supported by mathematical and computational evidence.

In a typical finite-treatment case, LOGOS identifies a globally optimal treatment combination that not only increases the integrated therapeutic benefit but also keeps the same amount and the same safety limits. While this is a model-based change and does not necessarily lead to clinical benefit, it is an example of how optimizing the treatment schedule can improve the therapeutic response without increasing drug consumption. The wider clinical role of LOGOS is to shift treatment planning from simply finding the right dose to finding the best combination of dose and timing to treatment duration, interruptions, and stopping. Possible applications include personalized pharmacotherapy, anticancer treatment scheduling, antimicrobial therapy, therapeutic drug monitoring, and adaptive treatment strategies. In conjunction with PK/PD models and precision dosing, LOGOS opens up a new tool to test for treatment alternatives based on their expected therapeutic benefit, safety, and practical utility. Drug-specific modeling and clinical validation are needed to confirm the theoretical results before these theoretical findings can be applied to treatment recommendations.

5. Conclusions

This study introduced LOGOS as a novel, open and reproducible approach to optimizing pharmacotherapy beyond dose selection, that takes into account timing, treatment duration and treatment interruption or discontinuation of treatment. Its main contribution is to explore the ways that the flexibility on treatment decision-making can enhance treatment benefit and also share common resource and safety constraints. Whereas most approaches focus on finding optimal schedules, LOGOS measures the value of different treatment options and distinguishes between numerical improvement from mathematically certified optimality. An important mathematical result is to derive a globally optimal loading-maintenance-coast strategy for a finite-horizon PK/PD model with an initial dose followed by a drug-free period for shorter treatment times. It is shown that there are benefits to taking daily or periodic (i.e., continuous) drugs, and the recovery mechanism for the drug and the toxicity limit are also important variables to be considered. Most importantly, LOGOS extends the safety evaluation beyond treatment cessation and recognizes that drug exposure may still cause side effects. As a clinical approach, LOGOS could potentially help to make decisions on dose adjustments, administration times, treatment holidays, discontinuation and restarting safely. Its wider impact is to find out not only how treatment can be improved, but also which therapeutic choices provide a meaningful benefit in what circumstances and with what level of evidence. Open-source Python implementation, independent numerical verification and reproducibility procedures are available for transparent assessment and replication of the findings.

Author Contributions: Conceptualization, V.D.K.; methodology, V.D.K.; software, V.D.K.; validation, V.D.K.; formal analysis, V.D.K.; investigation, V.D.K.; resources, V.D.K.; data curation, V.D.K.; writing—original draft preparation, V.D.K.; writing—review and editing, V.D.K.; visualization, V.D.K.; supervision, V.D.K.; project administration, V.D.K. The author has read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

Institutional Review Board Statement: Not applicable. This study presents mathematical models and numerical simulations. No participants were recruited, no intervention was administered and no new patient observations were collected.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new patient observations or experimental datasets were collected for this study. The literature sources informing the models and their context are cited in the article. The analytical results and numerical illustrations are presented in the manuscript, including its appendices. The research code and reproduction instructions are publicly available at https://github.com/vkaralis/LOGOS-theory-general [23]. No separate dataset accompanies this submission.

Acknowledgments: The author would like to express his gratitude to Professor Georgia Karali, Professor of Mathematics at the Mathematics Department of the National and Kapodistrian University of Athens, for her valuable scientific discussions, insightful suggestions, and constructive guidance regarding the mathematical formulation, analytical development, and solution of the mathematical models underlying the LOGOS theory. During the preparation of this manuscript, the author used Grammarly (accessed in September 2026) for language editing and clarity improvements. The author has reviewed and edited the content of this publication.

Conflicts of Interest: The author declares no conflicts of interest.

Appendix A. Matched Activation Bound

For a positive periodic sensitivity satisfying Equation 9, integration over one cycle gives $\langle b\rangle = (k_{rec}/k_{des})(1-\langle a\rangle)$. Dividing the sensitivity equation by a and integrating the logarithmic derivative gives $0 = k_{rec}\langle 1/a\rangle - k_{rec} - k_{des}\langle q\rangle$. Convexity of the reciprocal implies $\langle 1/a\rangle \geq 1/\langle a\rangle$. Hence $\langle a\rangle \geq k_{rec}/(k_{rec}+k_{des}\langle q\rangle)$, and the benefit identity yields $\langle b\rangle \leq G(\langle q\rangle)$.

The function G is increasing for nonnegative scalar argument y. Concavity of q then gives $G(\langle q(C)\rangle) \leq G(q(\langle C\rangle))$. Constant concentration with stationary sensitivity attains this bound. The proof concerns the specified periodic benefit model. It does not compare an independent injury process or finite-treatment transients.

Appendix B. Global Optimality Certificate for the Finite-Treatment Benchmark

Write E = EC50 and b(C) = Emax C/(E+C). For the candidate concentration C and a dose measure M, define the marginal value of administered amount at time s by

$$p(s) = \int_s^T b'(C(t))e^{-k(t-s)}\,dt, m(s) = \frac{p(s)}{V}. \quad (B1)$$

The adjoint satisfies $\dot{p} = kp - b'(C)$, with $p(T) = 0$. On a maintained plateau, p is constant and equals $b'(c)/k$. During coast, direct integration gives

$$p(t) = \frac{E_{max}E}{kC(t)}\left[\frac{1}{E+C(t)e^{-k(T-t)}} - \frac{1}{E+C(t)}\right]. \quad (B2)$$

Matching plateau and coast at $t = L$ yields $\exp(-k\delta) = c/(E+2c)$, hence Equation 13. During coast, the remaining time $T-t$ decreases while $k^{-1}\log(2+E/C(t))$ increases. Substitution in Equation B2 gives $kp(t) < b'(C(t))$ for $t > L$ and hence p decreases. Therefore m(t) is strictly below its plateau value $\mu = b'(c)/(kV)$. On the loading atom and maintenance support, m equals μ. If $T \leq k^{-1}\log(2+EV/D)$, the all-coast solution $C(t) = (D/V)\exp(-kt)$ instead has its maximum marginal value at time zero and supports an initial atom.

For any competing nonnegative dose measure M' with equal total mass, let C' denote its concentration trajectory. The prime on C denotes this competing trajectory, whereas b' is the derivative of the benefit function. Concavity and the linear concentration convolution give

$$J(M') - J(M) \leq \int_0^T b'(C(t))[C'(t) - C(t)]\,dt = \int_{[0,T]} m(s)\,d(M' - M)(s) \leq 0. \quad (B3)$$

The last inequality follows from $m \leq \mu$ and equality wherever M places amount. It proves global optimality rather than relying only on a necessary switching condition. Strict concavity yields

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

uniqueness of the concentration trajectory up to sets of zero time measure. The equation D = Vc [1+kT−log(2+E/c)] has a unique plateau root on L ≥ 0 because its derivative is V [1+kL+E/(2c+E)] > 0.

Appendix C. Exact and Conservative Cessation Checks

Equation 14 is obtained by integrating the linear injury equation with its exponential repair kernel. When $\alpha C_0^2 = 0$, injury is $z(t) = z_0\exp(-\varrho t)$ and therefore its supremum is z_0 and no stationary time division is needed. For $\alpha C_0^2 > 0$, $\varrho > 0$ and $\varrho \neq 2k$, a positive stationary time satisfies $\exp((2k-\varrho)t) = 2kA/[\varrho(z_0+A)]$. Use this relation only when the denominator is nonzero, the ratio is positive and the resulting time is positive. Otherwise, compare the present injury and limiting value directly. For $\varrho = 0$, injury increases toward $z_0 + \alpha C_0^2/(2k)$, as stated in Section 3.6.

For $\varrho = 2k$, the solution is $z(t) = (z_0+\alpha C_0^2 t)\exp(-2kt)$. If $\alpha C_0^2 > 0$, a positive interior stationary time exists exactly when $\alpha C_0^2 > 2kz_0$, and is $1/(2k) - z_0/(\alpha C_0^2)$. If $\alpha C_0^2 \leq 2kz_0$, injury is not increasing and its maximum value is z_0. For $k > 0$ and $\varrho \geq 0$, dropping repair gives the conservative full-tail bound $z(t) \leq z_0 + \alpha C_0^2/(2k)$. The continued absorption, active metabolites or a different repair law requires a different tail calculation.

Appendix D. Complete Tail Bound for the Finite Illustration

Extend the finite-treatment concentration into its input-free future. With k = V = 1 and nonnegative administered measure M of mass one, its convolution is $C(t) = \int[0,t] \exp(-(t-s))dM(s)$. Tonelli's theorem gives

$$\int_0^\infty C(t)^2\,dt \;=\; \frac{1}{2}\int_{[0,T]}\int_{[0,T]} e^{-|s-r|}\,dM(s)\,dM(r) \;\leq\; \frac{1}{2}. \quad (D1)$$

Concentration is at most one because the exponential kernel is at most one and the available total mass is one. Injury with nonnegative repair is bounded by the integral of its generating drive, so $z(t) \leq 0.5$ at every future time. This bound applies to all compared nonnegative schedules and avoids reliance on a truncated safety grid. It is specific to the normalized linear PK and convex injury model used in Section 3.7.

References

1. Bachmann F, Koch G, Pfister M, Szinnai G, Schropp J. OptiDose: Computing the Individualized Optimal Drug Dosing Regimen Using Optimal Control. Journal of Optimization Theory and Applications. 2021;189(1):46–65. https://doi.org/10.1007/s10957-021-01819-w
2. Bachmann F, Koch G, Bauer RJ, Steffens B, Szinnai G, Pfister M, et al. Computing optimal drug dosing regarding efficacy and safety: the enhanced OptiDose method in NONMEM. Journal of Pharmacokinetics and Pharmacodynamics. 2024;51(6):919–934. https://doi.org/10.1007/s10928-024-09940-9
3. He Q, Zhu J, Dingli D, Foo J, Leder KZ. Optimized Treatment Schedules for Chronic Myeloid Leukemia. PLOS Computational Biology. 2016;12(10):e1005129. https://doi.org/10.1371/journal.pcbi.1005129
4. Kantasiripitak W, Van Daele R, Gijsen M, Ferrante M, Spriet I, Dreesen E. Software Tools for Model-Informed Precision Dosing: How Well Do They Satisfy the Needs? Frontiers in Pharmacology. 2020;11:620. https://doi.org/10.3389/fphar.2020.00620
5. Augustin D, Lambert B, Robinson M, Wang K, Gavaghan D. Simulating clinical trials for model-informed precision dosing: using warfarin treatment as a use case. Frontiers in Pharmacology. 2023;14:1270443. https://doi.org/10.3389/fphar.2023.1270443
6. Minichmayr IK, Dreesen E, Centanni M, Wang Z, Hoffert Y, Friberg LE, et al. Model-informed precision dosing: State of the art and future perspectives. Advanced Drug Delivery Reviews. 2024;215:115421. https://doi.org/10.1016/j.addr.2024.115421
7. Chen T, Kirkby NF, Jena R. Optimal dosing of cancer chemotherapy using model predictive control and moving horizon state/parameter estimation. Computer Methods and Programs in Biomedicine. 2012;108(3):973–983. https://doi.org/10.1016/j.cmpb.2012.05.011

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

8. Lu M, Yuan Y, Liu S. A Bayesian pharmacokinetics integrated phase I–II design to optimize dose-schedule regimes. Biostatistics. 2025;26(1):kxae034. Published online 13 September 2024. https://doi.org/10.1093/biostatistics/kxae034
9. Le Louedec F, Puisset F, Thomas F, Chatelut É, White-Koning M. Easy and reliable maximum a posteriori Bayesian estimation of pharmacokinetic parameters with the open-source R package mapbayr. CPT: Pharmacometrics & Systems Pharmacology. 2021;10(10):1208–1220. https://doi.org/10.1002/psp4.12689
10. Fidler M, Wilkins JJ, Hooijmaijers R, Post TM, Schoemaker R, Trame MN, et al. Nonlinear Mixed-Effects Model Development and Simulation Using nlmixr and Related R Open-Source Packages. CPT: Pharmacometrics & Systems Pharmacology. 2019;8(9):621–633. https://doi.org/10.1002/psp4.12445
11. Thoma Y, Buclin T. TUCUXI-Core: An Open Source generic and efficient computing engine for Model-Informed Precision Dosing. Smart Health. 2025;38:100600. https://doi.org/10.1016/j.smhl.2025.100600
12. Zhang J, Cunningham JJ, Brown JS, Gatenby RA. Integrating evolutionary dynamics into treatment of metastatic castrate-resistant prostate cancer. Nature Communications. 2017;8(1):1816. https://doi.org/10.1038/s41467-017-01968-5
13. Gallagher K, Strobl MA, Gatenby RA, Zhang J, Maini PK, Anderson AR. Mathematical Biomarkers of Adaptive Therapy Outcomes in Prostate Cancer. JAMA Oncology. 2026;12(10):1079–1087. https://doi.org/10.1001/jamaoncol.2026.2781
14. Jager NGL, Chai MG, van Hest RM, Lipman J, Roberts JA, Cotta MO. Precision dosing software to optimize antimicrobial dosing: a systematic search and follow-up survey of available programs. Clinical Microbiology and Infection. 2022;28(9):1211–1224. https://doi.org/10.1016/j.cmi.2022.03.041
15. Alghamdi WA. Clinical Use of Bayesian Model-Informed Precision Dosing in Routine Practice: A Focused Systematic Review. Journal of Clinical Medicine. 2026;15(10):3838. https://doi.org/10.3390/jcm15103838
16. Katriel G. Optimizing Antimicrobial Treatment Schedules: Some Fundamental Analytical Results. Bulletin of Mathematical Biology. 2024;86(1):1. Published online 23 November 2023. https://doi.org/10.1007/s11538-023-01230-8
17. Karalis V. Beyond the Dose: Introducing LOGOS for Dosage Regimen Design in Therapeutics. medRxiv [Preprint]. 2026. https://doi.org/10.64898/2026.10.17.26365797.
18. Sheiner LB, Stanski DR, Vozeh S, Miller RD, Ham J. Simultaneous modeling of pharmacokinetics and pharmacodynamics: Application to d-tubocurarine. Clinical Pharmacology & Therapeutics. 1979;25(3):358–371. https://doi.org/10.1002/cpt1979253358
19. Dayneka NL, Garg V, Jusko WJ. Comparison of four basic models of indirect pharmacodynamic responses. Journal of Pharmacokinetics and Biopharmaceutics. 1993;21(4):457–478. https://doi.org/10.1007/bf01061691
20. Ben-Tal A, Nemirovski A. Robust Convex Optimization. Mathematics of Operations Research. 1998;23(4):769–805. https://doi.org/10.1287/moor.23.4.769
21. Nemirovski A, Shapiro A. Convex Approximations of Chance Constrained Programs. SIAM Journal on Optimization. 2007;17(4):969–996. Published online 22 November 2006. https://doi.org/10.1137/050622328
22. Raue A, Kreutz C, Maiwald T, Bachmann J, Schilling M, Klingmüller U, et al. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. Bioinformatics. 2009;25(15):1923–1929. https://doi.org/10.1093/bioinformatics/btp358
23. Karalis VD. LOGOS theory general: reproducible dimensionless PK/PD examples [Research software]. Version 0.5.0. https://github.com/vkaralis/LOGOS-theory-general.
24. Friberg LE, Henningsson A, Maas H, Nguyen L, Karlsson MO. Model of Chemotherapy-Induced Myelosuppression With Parameter Consistency Across Drugs. Journal of Clinical Oncology. 2002;20(24):4713–4721. https://doi.org/10.1200/jco.2002.02.140
25. Dulhunty JM, Brett SJ, De Waele JJ, Rajbhandari D, Billot L, Cotta MO, et al. Continuous vs Intermittent β-Lactam Antibiotic Infusions in Critically Ill Patients With Sepsis: The BLING III Randomized Clinical Trial. JAMA. 2024;332(8):629–637. https://doi.org/10.1001/jama.2024.9779
26. Wraight JM, Hancox RJ, Herbison GP, Cowan JO, Flannery EM, Taylor DR. Bronchodilator tolerance: the impact of increasing bronchoconstriction. European Respiratory Journal. 2003;21(5):810–815. https://doi.org/10.1183/09031936.03.00067503

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.

Preprints.org | NOT PEER-REVIEWED | Posted: 29 October 2026 https://doi.org/10.20944/preprints202610.2916.v1

27. Regoes RR, Wiuff C, Zappala RM, Garner KN, Baquero F, Levin BR. Pharmacodynamic Functions: a Multiparameter Approach to the Design of Antibiotic Treatment Regimens. Antimicrobial Agents and Chemotherapy. 2004;48(10):3670–3676. https://doi.org/10.1128/aac.48.10.3670-3676.2004
28. Rybak MJ, Le J, Lodise TP, Levine DP, Bradley JS, Liu C, et al. Therapeutic monitoring of vancomycin for serious methicillin-resistant Staphylococcus aureus infections: A revised consensus guideline and review by the American Society of Health-System Pharmacists, the Infectious Diseases Society of America, the Pediatric Infectious Diseases Society, and the Society of Infectious Diseases Pharmacists. American Journal of Health-System Pharmacy. 2020;77(11):835–864. https://doi.org/10.1093/ajhp/zxaa036

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.

© 2026 by the author(s). Distributed under a Creative Commons CC BY license.