Preprints.org

Article

Not peer-reviewed version

Finite Course Dosage Regimen Design in Therapeutics with LOGOS Theory B

Vangelis D. Karalis *

Posted Date: 26 October 2026

Keywords: finite-course optimization; pharmacokinetics; temporal allocation; cessation; dosage regimen; personalized medicine

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Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1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

Finite Course Dosage Regimen Design in Therapeutics with LOGOS Theory B

Vangelis D. Karalis [1,2,*]

[1] Department of Pharmacy, 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

Temporal allocation and cessation are distinct design freedoms in finite treatment courses. We study Theory B, the finite-course special case of the LOGOS framework, under matched resource and safety assumptions. We establish an endpoint-specific timing null, a global solution for a concave linear-elimination benchmark, and a policy-conditional stopping contrast. The analytical solution has no-treatment, bolus–coast, and loading–maintenance–coast regimes. At the canonical matched dose, flexible allocation improves integrated benefit over an optimal no-loading constant reference by 13.218%; removing the declared continuation floor contributes a separate increment of 0.139 benefit units. A delayed-injury counterexample shows why safety assessment may need to extend beyond the efficacy horizon. An analytical constant-family ceiling and a conservative flexible witness establish restoration of target feasibility for $11.271 < J_{req} \leq 13.689$. Fixed-target calculations distinguish about 20% dose savings at targets 10 and 11 from restored feasibility at targets 12 and 13. Analytical certificates, feasible witnesses and floating-point bounds within a finite-dimensional family have different scopes. Together, these results provide an attribution framework that separates temporal allocation, policy-conditional cessation, resource savings and target feasibility. They support conditional model-based results and do not demonstrate clinical efficacy or safety gains.

Keywords: finite-course optimization; pharmacokinetics; temporal allocation; cessation; dosage regimen; personalized medicine

Introduction

A planned treatment course may not be the same as the period of administration. Exposure and downstream effects can persist after the last dose and injury can go beyond the efficacy endpoint. Finite course design therefore has three boundaries (last administration time, benefit evaluation horizon and safety evaluation horizon). The separation of these boundaries can give temporal placement and cessation independent value.

The mathematics of these dosing structures predates LOGOS. Murray suggested chemotherapy control with toxicity limits and bolus followed by infusion [1] and Ledzewicz and Schättler investigated how PK and PD would change the optimal control structures [2]. Effect-site models can account for delayed action [3] and transit models can mimic delayed myelosuppression [4]. OptiDose is a tool for personalized discrete-dose optimization [5]; robust repeated-administration control with parameter uncertainty has been studied [6]. Fernández and Pola provide model-specific classification in terms of PK/PD variants and integral constraints [7]. These precedents define the scope of our work and we do not claim to have a new optimizer nor the first manifestation of loading and maintenance.

The foundational LOGOS (Longitudinal Optimization of Global Outcomes and Strategies) framework was introduced in Beyond the Dose: Introducing LOGOS for Dosage Regimen Design in Therapeutics, now available as a medRxiv preprint [8]. The present paper develops its finite-course

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

special case Theory B that considers additional regimen design freedoms under the same model, resource cap, objective and safety requirements. Its main results include timing null conditions, global classification in a solvable benchmark, an endogenous cessation boundary, stopping value identifiability, and the difference between objective improvement and restored feasibility. The benchmark itself is not a reconstruction of the initial preprint calculation for finite courses. In that calculation, we use a direct Emax response and an exact total amount, while in the present benchmark we use a quadratic response, dose cap and budget shadow value. Their threshold formulas, coast lengths and reported benefits therefore apply to different optimization problems.

The general LOGOS framework supplies principles for comparison and verification [8]. Theory A describes recovery during withdrawal and restarting conditions; Theory B describes finite-course allocation and policy-conditional cessation. The companion Theory B application manuscript is devoted to nested time window opportunities in fitted remifentanil and warfarin profiles and the Theory A application manuscript is devoted to observable withdrawal and restart components. The current benchmark is based on persistence of the state, finite evaluation time horizons and downstream consequences without requiring responsiveness recovery. A benefit-bearing terminal coast and recovery interval are therefore different mechanisms in LOGOS.

We build these comparisons using the endpoint-specific timing null, the globally certified measure input benchmark, the continuation policy attribution contrast, and a feasible safety witness against an analytical reference ceiling. Strict structural value is achieved only when a feasible and richer strategy is found which goes beyond the reference ceiling or a certified smaller-class value. Early termination is used as a descriptor of the schedule, except when the two classes being compared differ only in their obligation to continue. Each pharmacological application must assess its endpoint and constraints; the benchmark optimum cannot be transferred directly.

The contribution is the attribution and verification framework that connects the analytical benchmark with resource and feasibility comparisons. Loading, maintenance and the underlying optimal control methods are well-known concepts; there are explicit comparison rules in place to assess the value of these.

Methods

Course Horizons and Admissible Controls

For a system $\dot{x} = f(x,u,t;\theta)$, $x(0)=x_0$, T_eff and T_safe must be specified before optimization. Administration is allowed up to T_eff in the analyses below but other cases can have a different maximum window. Let D be an administered resource and B its cap. Benefit and path constraints are defined on their horizons.

$$J(u) = \int_0^{T_{eff}} \ell(x(t),u(t),t)dt + \Phi(x(T_{eff})), \quad D(u) \le B \qquad (1)$$

$$h_j(x(t),u(t)) \le 0, \quad 0 \le t \le T_{safe} \qquad (2)$$

For nonzero input measure μ, we define the cessation τ to be the supremum of the support. For zero input treatment, we assume τ = 0 because its support is empty. So too an initial atom alone gives τ=0, but its positive dose makes it different from no treatment. The last active infusion-bin time is a threshold-dependent numerical descriptor, not an exact free boundary.

No new administration means μ=0; it does not imply C=E=Z=0 when residual states are present at initialization. The no-treatment classification below assumes C(0−)=0 and the stated linear elimination model, so zero input yields zero exposure and zero benefit.

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

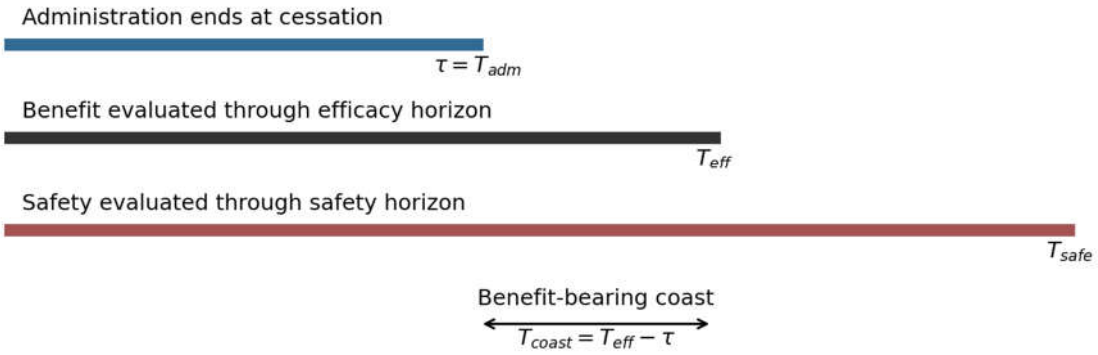


Figure 1. Administration, efficacy, and safety horizons. The coast interval has zero input after cessation and can retain positive benefit. Bar lengths are conceptual; safety can extend beyond efficacy.

A matched Budget Benchmark and Its Dual Parameter

All the canonical quantities are dimensionless mechanism demonstration values. Exposure follows linear elimination with a finite nonnegative input measure for relaxed loading. The main problem is to maximize the concave integrated benefit under a common dose cap. We write the input measure and total resource as concentration equivalents. If the mass input is intravenous with a fixed physical distribution volume V_d, we write dμ_mass/V_d as input to the state and the full-tail concentration exposure as D_mass/(kV_d).

$$dC = -kCdt + d\mu, \quad C(0^-) = 0, \quad g(C) = aC - \frac{b}{2}C^2 \qquad (3)$$

$$V(B,T) = \sup_{\mu \geq 0, D(\mu) \leq B} \int_0^T g(C(t))dt, \quad D(\mu) = \mu([0,T]) \qquad (4)$$

Here a,b,k>0, T=T_eff>0 and B≥0. The instantaneous loading atom is a theoretical upper-bound freedom, not necessarily a physically achievable action. Bounded-rate implementations are part of this relaxed class. The quadratic g is only a representation, and not a calibrated clinical efficacy response; the positive-shadow-value optima stay below a/b.

Introduce q≥0 as the budget dual variable or shadow value. Maximizing the penalized functional traces budget-optimal courses; q is not a toxicity parameter. The Lagrangian differs from the following functional only by the constant qB.

$$F_q(\mu) = \int_0^T g(C(t))dt - qD(\mu) \qquad (5)$$

The analytical theorem uses q>0 to obtain the interior active-budget family. A companion injury state $\dot{Z}=\eta C-\varrho Z$ is used only in the explicitly declared safety campaigns. It is not part of the unconstrained benefit objective. The effect-site extension uses $\dot{E}=k_e(C-E)$ and applies g to E, and the analytical benchmark has E=C.

Regimen Classes and Policy Conditional Attribution

For nested nonempty classes $R_0 \subseteq R_1 \subseteq R_2$ we may define V_i as the finite supremum of the same objective under the same resource cap and safety conditions. The classes will limit the additional design flexibility. If R_1 already has terminal zero input, then adding a cessation label makes $R_2=R_1$ and the additional stopping value is zero.

Class indices and symbols are local to this paper: B is the administered-resource cap, D is the administered resource, and q is its shadow value. In the flagship, B is the integrated benefit; in Theory A, D is disease burden, and q is the normalized mean input. These endpoints and parameter values are not interchangeable across articles. All finite comparisons have the same initial state, the evaluation horizons and resource rule. A dose-cap comparison matches the available resource ceiling and need not match realized dose; exact-dose comparisons are identified separately.

$$\Delta V_{shape} = V_1 - V_0, \quad \Delta V_{stop|policy} = V_2 - V_1 \qquad (6)$$

$$\Delta V_{structure} = V_2 - V_0 = \Delta V_{shape} + \Delta V_{stop|policy} \qquad (7)$$

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

For the attribution construction we use ε=0.5, T=4, B=3[5−arcosh(4/3)]≈12.614 and the same dimensionless exposure model and benefit objective. There is no safety cap and no finite input-rate upper bound in this construction. R_0 consists of constant rates r on [0,T], no loading atom, and ε≤r≤B/T. R_1 permits μ=$d_0\delta_0$+u(t)dt with d_0≥0, u≥ε almost everywhere on [0,T], and d_0+∫u≤B. Only an initial atom is allowed. R_2 allows for the same initial atom and rate restrictions on a prefix [0,τ], where τ∈[0,T], followed by exactly zero input; d_0+$\int_0$^τu≤B. At τ=0, the prefix condition is vacuous. We get $R_0 \subseteq R_1 \subseteq R_2$.

The first enlargement releases temporal redistribution including relaxed initial loading. A separate loading contribution is not isolated. The second enlargement removes the need to continue input at rate ε. The minimum-rate construction is an attribution device and not a clinical maintenance requirement or a causal estimate from patients. Its stopping value depends on the declared policy.

Motifs and Nondimensional Groups

Loading brings exposure to some operating level, either through a relaxed atom or an early bounded-rate segment. Maintenance preserves a constant exposure using u=kC. A benefit-bearing coast is a positive-length zero-input interval after cessation with g(E)>0. We reserve therapeutic coast for net positive value when benefit and harm are in the same criterion (e.g. g(E)−h(Z)>0). Positive concentration or a positive benefit proxy does not give net therapeutic value.

$$H = kT, \quad \gamma = \frac{qk}{a}, \quad \beta = \frac{bB}{a}, \quad y = \frac{bC}{a} \qquad (8)$$

Further groups include ϱ/k, k_e/k, k_a/k, bu_max/(ak), and T_safe/T. They govern extensions but are not axes of the proven unconstrained analytical classification.

Table 1. Notation and definitions.

Symbol	Definition
T_adm = τ	Last input time; zero input uses τ=0 convention
T_eff = T and T_safe	Benefit horizon and full-path safety horizon
C, E and Z	Exposure, effect-site exposure, and injury
μ, u, D and B	Input measure, rate, administered resource, and cap
a, b, k and q	Benefit parameters, elimination rate, and budget shadow value
H, γ and β	Horizon, dual-value, and budget dimensionless groups
H_0, H_c and L	Treatment-onset boundary, maintenance boundary, and coast length
ε and ΔV_stop|policy	Persistence obligation and its conditional stopping value

Numerical Extension

The fixed-target extension minimizes administered amount at J_req=10, 11, 12, or 13 under the common injury cap 2.5 and the canonical amount cap B=3[5−arcosh(4/3)]≈12.614. It uses zero initial states and benefit horizon 4, injury checks through time 20 with a decreasing analytical tail, an initial atom, 80 equal infusion bins and rate bound 20. The initial atom is also bounded above by 20 in our version, which is a redundant restriction as the common resource cap is smaller. Concavity gives supporting-LP dose lower bounds; candidates satisfying the constraints to the declared numerical tolerance give numerical upper bounds. Independent DOP853 and Radau integration, quadrature checks and stationary-point localization check the stored controls. This is a post hoc extension planned before its numerical runs, not prospective registration.

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

Computational Environment

The numerical calculations were performed in Python 3.12 using NumPy 2.3.5 for numerical array operations, SciPy 1.17.1 for optimization and numerical integration, and Matplotlib 3.10.9 for visualization. These package versions are pinned in the archived computational release [12]. The fixed-target results reported here correspond to the archived theoretical dose-frontier outputs. Analytical proofs do not depend on the numerical optimizer or on the separate synthetic pharmacological examples in the archive.

The computational code is available at https://github.com/vkaralis/LOGOS-theory-B and archived on Zenodo as release v1.0.1, DOI 10.5281/zenodo.23402692 [12].

Use of Artificial Intelligence

Grammarly was used for English language editing. The author is responsible for the manuscript's scientific content and conclusions.

Results

Timing Null Conditions

Theorem 1. Cumulative-only timing null. If J(u)=Ψ(D(u)) and every regimen has D=B, all feasible regimens have objective Ψ(B). Temporal redistribution has no objective value between nonempty comparison classes. Safety may change class feasibility without changing the value of the feasible members.

Proof. Substituting the common dose into Ψ gives the same value for every member. This statement concerns a cumulative-only endpoint and does not apply to an arbitrary memoryless nonlinear response.

Corollary 1. Full-tail linear-exposure equivalence. With zero initial exposure, finite total input, k>0, and vanishing terminal exposure, unweighted exposure integrated over the entire tail is timing-invariant. On a finite window the corresponding kernel depends on input time.

$$\int_0^\infty C(t)dt = \frac{D}{k} \quad (9)$$

$$\int_0^T C(t)dt = \int_{[0,T]} \frac{1-e^{-k(T-s)}}{k} d\mu(s) \quad (10)$$

Proof. Integration of the linear state equation over the full tail gives kʃC=D; convolution followed by interchanging integrals gives the finite-window identity. Thus equal administered dose implies equal full-tail linear exposure, but not equal finite-window exposure.

Observation 1. If benefit is g(u), g(0)=0, and there is no persistent benefit-bearing state or terminal benefit, zero input after cessation produces zero residual benefit. Earlier stopping may still follow from cost or safety; it is not a benefit-bearing coast.

Global Three Regime Classification

Theorem 2. Global classification theorem for the canonical benchmark. For finite nonnegative input measures, a,b,k,q>0, and the penalized benchmark, define the following boundaries when 0<γ<1.

$$c = \frac{a-qk}{b}, \quad H_0 = -\log(1-\gamma), \quad H_c = \operatorname{arcosh}\left(\frac{1}{1-\gamma}\right) \quad (11)$$

If γ≥1 the optimum is no treatment for every finite H. If 0<γ<1 then there is no treatment when H≤H_0. If H_0<H≤H_c then the optimum is one initial relaxed loading atom of size d followed by benefit-bearing coast. If H>H_c it is initial loading to c, maintenance at kc until τ and coast thereafter.

$$d = \frac{2[a(1-e^{-H})-qk]}{b(1-e^{-2H})}, \quad C(t) = de^{-kt} \quad (12)$$

$$d\mu^* = c\delta_0 + kc1_{(0,\tau)}dt, \quad \tau = T-L, \quad L = \frac{H_c}{k} \quad (13)$$

$$C^*(t) = c \ \text{ for } 0 \le t \le \tau; \quad C^*(t) = ce^{-k(t-\tau)} \ \text{ for } \tau < t \le T \quad (14)$$

At H=H_c, maintenance has zero duration and d=c. For q>0 the exposure trajectory is unique and the final input atom is inefficient. Appendix A verifies a global sufficient condition over all finite

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

nonnegative input measures, as opposed to only in a loading–maintenance family. These three regimes describe this benchmark rather than every pharmacological model.

Cessation Boundary and Horizon Extension

The marginal benefit of input at s is λ(s). Global optimality is achieved if λ≤q throughout, with equality on the input support. On maintenance, λ=q and λ̇=kλ−(a−bC) require C=c. At cessation, the marginal future benefit of input equals the budget shadow value; it subsequently falls below that value.

$$\lambda(s) = \int_s^T e^{-k(t-s)}[a - bC^*(t)]dt \quad (15)$$

$$\lambda(\tau) = \frac{a}{k}(1-r) - \frac{bc}{2k}(1-r^2) = q, \quad r = e^{-kL} \quad (16)$$

$$r = \frac{1-\sqrt{1-(1-\gamma)^2}}{1-\gamma} = e^{-H_c} \quad (17)$$

The residual exposure remains strictly between zero and a/b, so g(C)>0 on the entire final coast. This proves residual benefit, not a harm-adjusted therapeutic outcome.

Corollary 2. When broadening the horizon while maintaining a constant shadow value, we fix a, b, k and q and increase T without changing the maintenance regime. The coast length is L=H_c/k and the plateau c is constant. The τ variable increases one-for-one with T. To obtain the matched-budget optimum, we need additional resources of kcΔT.

Proof. H_c and c depend only on the fixed model and dual parameter. Substituting τ=T−H_c/k and B=c(1+kτ) we arrive at the conclusion. The shadow value remains constant and the budget grows.

At fixed total budget the coast generally changes. In the active maintenance branch, β=(1+H−$\mathcal{L}$)/cosh($\mathcal{L}$) where $\mathcal{L}$=kL. Implicit differentiation gives the following positive derivative at fixed β. A longer horizon may thus lead to a longer coast if the budget is fixed.

$$\frac{d\mathcal{L}}{dH}\Big|_\beta = \frac{1}{1+(1+H-\mathcal{L})\tanh(\mathcal{L})} > 0 \quad (18)$$

Budget Optimal Courses and the Structure Map

Corollary 3. Budget correspondence. Let μ_q maximize F_q with dose B(q). For each input of dose at most B(q), its benefit cannot exceed that of μ_q. Thus μ_q is globally optimal for the matched-cap problem at B(q).

$$J(\mu) \le F_q(\mu_q) + qD(\mu) \le F_q(\mu_q) + qB(q) = J(\mu_q) \quad (19)$$

At fixed H, the dual-to-budget correspondence is strictly decreasing on the positive-dose, positive-value branch. In dimensionless form it is piecewise explicit.

$$\beta(\gamma, H) = \frac{2[(1-e^{-H})-\gamma]}{1-e^{-2H}} \quad \text{on the bolus branch} \quad (20)$$

$$\beta(\gamma, H) = (1-\gamma)[1 + H - H_c(\gamma)] \quad \text{on the maintenance branch} \quad (21)$$

The zero-dose region begins at γ≥1−e^(−H). The active branch covers 0<β<1+H. The bolus–maintenance boundary is β_c=sech(H). For 0<β≤sech(H), the budget optimum is an initial bolus of B. For sech(H)<β<1+H we solve β=(1+H−$\mathcal{L}$)/cosh($\mathcal{L}$) for 0<$\mathcal{L}$<H and use c=(a/b)sech($\mathcal{L}$). For β=1+H, the smallest dose optimum exhausts the budget and for β>1+H, the same optimum leaves the excess resource unutilized. This representative loads to a/b and maintains that exposure through T with zero final coast. A terminal atom cannot improve integrated benefit.

The budget panel therefore has no positive-budget no-treatment region for this benefit-only objective. Its zero-budget axis is no input. Unlike the dual panel, it also displays the saturation branch. The boundary statements follow by inserting H_c=H into the dual map and taking q ↓ 0. Appendix B gives monotonicity and the corresponding qualification of the horizon result.

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

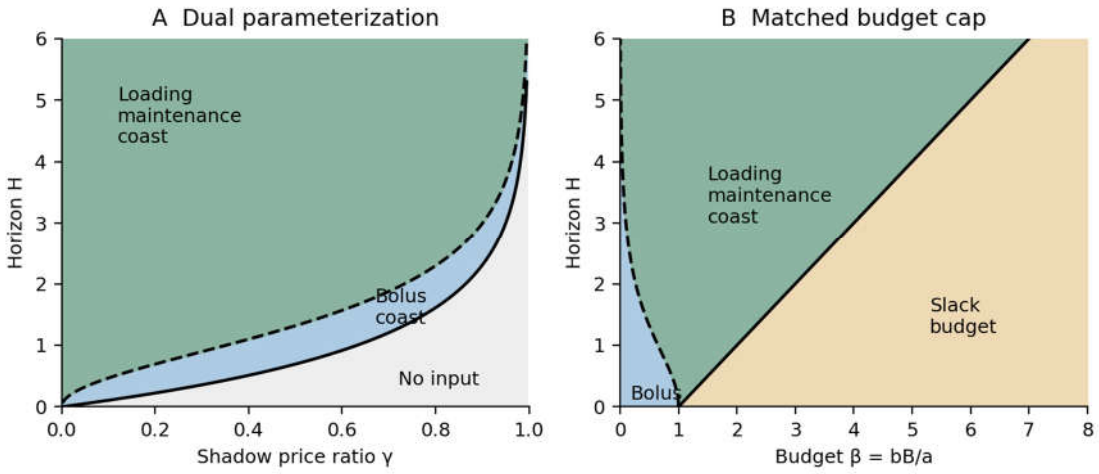


Figure 2. Analytical Structure Map. A uses horizon H and budget shadow-value ratio γ. B uses H and resource cap β. Blue denotes bolus–coast, green loading–maintenance–coast, and tan a slack or saturated budget with a minimal-dose full-window maintenance representative. No input in B is confined to β=0. Active safety constraints are outside this proven map.

Matched dose Objective Gain and Policy Identifiable Stopping Value

With a=4, b=k=q=1 and T=4, the plateau is 3, cessation is approximately 3.205, coast is 0.795, and the associated dose is 12.614. Integrated benefit is approximately 28.826. The best no-loading constant reference under the same cap uses the full amount and achieves 25.461: a 13.218% gain at exactly matched realized dose. The absolute gain is 3.365. The benefit during the final coast is 4.792. Resource cost cancels in this exact-dose comparison.

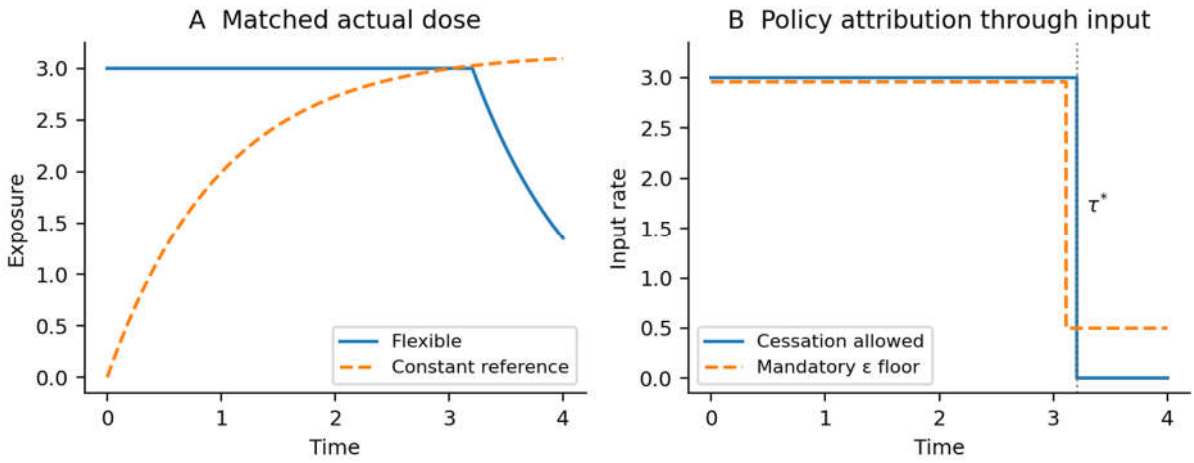


Figure 3. Canonical course and policy attribution. A compares the analytical optimum with the optimized no-loading constant reference at identical realized dose. B compares the input rates of the two flexible classes with identical initial-loading permission: zero input after cessation versus mandatory rate ε. Both classes also permit the initial atom, which is not represented as a finite input rate. The vertical line marks the canonical cessation τ*. The minimum-rate policy is an attribution device.

For the explicit ε=0.5 attribution construction, the globally optimal R_1 course has plateau 2.961 and ends input above the mandatory floor at 3.11, but continues mandatory rate ε until T. Its benefit is 28.687. The unrestricted canonical course is a member of R_2 and is therefore also globally optimal in that class. Appendix C proves the floor-constrained result.

$$\Delta V_{shape} \approx 3.226, \quad \Delta V_{stop|policy} \approx 0.139, \quad \Delta V_{structure} \approx 3.365 \qquad (22)$$

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

These numbers attribute value to the declared restrictions. They are not an intrinsic stopping gain of linear PK and do not contradict the zero gain obtained when both classes already permit terminal zero input.

Safety Horizons and a Feasibility Restoration Interval

Proposition 1. Analytical delayed-injury counterexample. After a unit bolus, let C=exp(−kt), Z′=C−ϱZ, and Z(0)=0, with k>ϱ>0. The injury trajectory and its unique finite maximum are explicit.

$$Z(t) = \frac{e^{-\rho t} - e^{-kt}}{k-\rho}, \quad t_{peak} = \frac{\log(k/\rho)}{k-\rho} \qquad (23)$$

Proof. The integrating-factor solution gives Z; setting Ż=0 gives the unique maximum. Any cap strictly between Z(T_eff) and Z(t_peak) is passed on [0,T_eff] but violated subsequently if T_eff<t_peak≤T_safe.

For k=1, ϱ=0.2, T_eff=0.2 and T_safe=8 the injury is 0.178 at the efficacy endpoint and peaks at 0.669 at 2.012. A cap of 0.3 is violated after the endpoint. This is an analytical counterexample to checking safety only when evaluating the efficacy.

A separate constrained campaign uses the same canonical cap B≈12.614, T_eff=4, T_safe=20, Ż=C−0.2Z and Z≤2.5. We can define the reference family as a constant input rate without any initial loading. The flexible class allows for an initial loading atom and nonnegative time-varying rates. The model, benefit, dose cap, initialization state and safety requirement are shared; realized doses need not be equal. For a constant rate u on [0, 4] and zero thereafter, we have J(u) = $4uI_1 - u^2I_2/2$ with $I_1 = 4 - 1 + \exp(-4)$ and $I_2 = 4 - 2[1 - \exp(-4)] + [1 - \exp(-8)]/2$. Let M be the complete-tail peak injury created by unit constant input. The linearity gives us max Z=uM so the admissible interval is 0≤u≤min(B/4,2.5/M). Maximizing the stated concave quadratic over this interval gives the reference ceiling.

The largest safe constant rate is approximately 1.049, with a post-administration injury peak at time 4.723. Its globally maximal benefit is approximately 11.271. The nominal 80-bin flexible witness has benefit approximately 13.691 and reaches the injury cap to within numerical tolerance. For the feasibility claim, we multiply its initial atom and every input rate by $1-10^{-4}$. This conservative witness has benefit at least 13.689, dose about 4.563 and peak injury below 2.5 with a strict positive cap margin of about 2.5×10^{-4}. Endpoints and stationary points are checked through time 20 and the subsequent input-free tail is decreasing. Independent ODE integration checks this witness. The near-cap stationary-time differences between implementations do not change the target interval.

Corollary 4. Restoration of target feasibility. Let J_ref* be the analytical safe constant-family ceiling and J_witness the full-path-checked conservative flexible value. Any target above J_ref* and no larger than J_witness is infeasible in the reference family and can be reached by this witness. Feasibility also includes the common efficacy requirement and injury cap even if both safety-only classes have regimens. Using outward-safe rounded endpoints, the verified interval is 11.271 < J_req ≤ 13.689. The matched resource is the available cap; realized doses need not be equal.

$$J^*_{ref} < J_{req} \le J_{witness} \quad \Leftarrow \quad 11.271 < J_{req} \le 13.689 \qquad (24)$$

Proof. The analytical reference maximum is the upper limit for all members of the safe constant family. This conservative witness satisfies both resource and injury constraints and achieves an objective value of at least 13.689. So every target in the rounded interval is achievable using this witness and is above the constant family ceiling. This result shows restoration of target feasibility, although it does not establish global optimality of the safety-constrained flexible class or a stopping-only contribution.

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

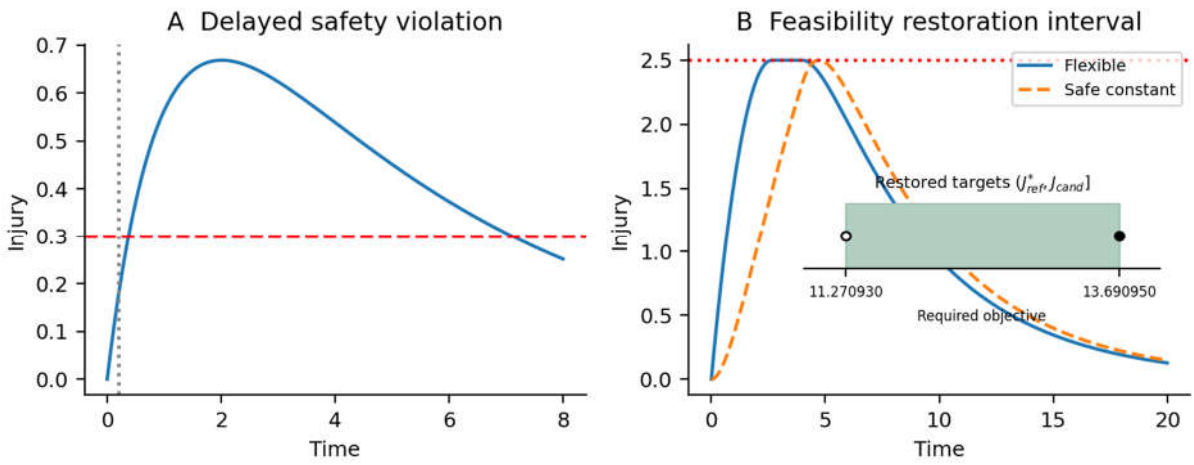


Figure 4. Safety-horizon mismatch and restoration of target feasibility. A shows injury after the efficacy horizon. B shows the nominal safe paths under the common cap. The existence claim uses the slightly scaled conservative flexible witness with margin and target interval in the text, but the nominal near-cap plotting does not provide this margin. Display sampling is separated from the stationary-point and analytical-tail verification.

Fixed Target Dose Saving and Restored Feasibility

At targets 10 and 11, flexible minimum-dose candidates used 2.929 and 3.266, versus analytic constant-family minimum amounts 3.666 and 4.082, corresponding to 20.116% and 19.979% dose savings. The flexible doses at targets 12 and 13 were 3.614 and 3.974. No constant input can attain either target under the same injury cap because its analytic ceiling is approximately 11.271. Those two cases demonstrate target-feasibility restoration, rather than a finite percentage saving against an unavailable constant solution. Dose lower bounds differ from candidate amounts by less than 3×10^{-5} within the 80-bin-plus-atom family. This does not establish a minimum over unrestricted continuous inputs or a drug-specific dose reduction.

Table 2. Minimum dose candidates under injury cap 2.5.

Target J	Realized D flexible	Realized D constant	Dose saving	Interpretation	Comparison basis
10	2.929	3.666	20.116%	Dose saving	Fixed target J; common safety cap and available upper cap B
11	3.266	4.082	19.979%	Dose saving	Fixed target J; common safety cap and available upper cap B
12	3.614	Infeasible	Not defined	Restored feasibility	Fixed target J; common safety cap and available upper cap B
13	3.974	Infeasible	Not defined	Restored feasibility	Fixed target J; common safety cap and available upper cap B

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

Doses are in the dimensionless model's resource units. Displayed results are rounded to at most three decimal places; calculations and constraint checks use the unrounded controls. Conservative rounding is used for the target-feasibility interval. In the archived fixed-target calculations, target residuals have absolute magnitude below 2×10^{-12} and the safety margins are positive. The smallest safety margin, at target 13, is approximately 10^{-6}. These are floating-point checks, not exact-arithmetic certificates. Both the available resource cap and realized use are specified. Ideal loading atoms and the stated rate bounds are part of this numerical family.

Discussion

The central contribution is an attribution framework for finite-course freedom under explicit comparison rules. Its distinctive feature is the explicit connection of a global measure-input solution to a stopping-identifiability criterion and separate fixed-dose, fixed-target and target-feasibility comparisons. Established PK/PD optimal-control work already describes structured dosing [1,2,7]. The present results connect a timing null, a globally solved benchmark, and policy-specific cessation to separate resource and feasibility questions. The canonical loading permission is relaxed, the response is illustrative, and the injury state is dimensionless. These assumptions make the comparison transparent but do not establish a clinical dosing policy.

A benefit-bearing coast describes residual contribution to the selected objective. This does not mean a net therapeutic benefit, nor does the integrated benefit mean the stopping increment. Identifying the latter requires two classes that are different only in their continuation commitments. A common cap does not necessarily ensure the same dose realized: the safety witness and the fixed target frontier address different questions from the canonical exact-dose comparison.

The global result is for the given concave measure-input benchmark. The safety witness proves existence beyond this reference limit and the fixed-target supporting bounds apply for an 80-bin plus atom family. Neither of these results offers an unrestricted safety-constrained optimum. Pharmacological applications are made in the companion application study and their endpoints require separate interpretation. The clinical translation requires calibrated benefit and harm models and evidence for the associated policy comparison.

Conclusion

The contribution is a unified and verifiable attribution of finite-course design freedoms and not a new dosing motif. Theory B distinguishes between temporal redistribution, policy-specific stopping, dose saving and restoring target feasibility. The timing null and global benchmark identify the consequences of our model structure. A conservative safety witness and a fixed-target frontier define feasibility and resource results. These conclusions are based on the comparator, endpoint, permissions and evidence level and do not imply clinical benefit or unrestricted safety-constrained numerical optimality.

Appendix A Global Sufficient Condition Proof

Let C* be a proposed exposure path and $\lambda(s)=\int_s^T e^{-k(t-s)}[a-bC^*(t)]dt$. Concavity of g gives a global upper supporting inequality for any alternative input measure μ.

$$F_q(\mu) - F_q(\mu^*) \leq \int_{[0,T]} [\lambda(s) - q] d(\mu - \mu^*)(s) \qquad (25)$$

If $\lambda \leq q$ and equality holds on the support of μ*, the right-hand side is nonpositive. This is enough to ensure global optimality because the nonnegative inputs form a convex cone and the exposure is linearly related to the input. The strict concavity of the function ensures that there is a unique exposure trajectory almost everywhere. Its input is given by the state equation. If $q > 0$, the terminal atom strictly decreases F_q and is excluded.

For zero input, $\lambda(s)=a[1-e^{-k(T-s)}]/k$. Its maximum is at s=0, and it is at most q precisely when $\gamma \geq 1-e^{-H}$, including $\gamma \geq 1$. This proves the no-treatment branch and its boundary.

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

For a pure initial atom of amount d, the equality $\lambda(0)=q$ gives us the bolus formula which is positive exactly for $H>H_0$. To check that $\lambda(s)\leq\lambda(0)$ on the other branch we have r_T=e^(−H) and z=e^(−ks) with r_T≤z≤1.

$$k\lambda(s) = a(1 - r_T/z) - \frac{bd}{2}(z - r_T^2/z) \quad (26)$$

$$\frac{d(k\lambda)}{dz} = \frac{ar_T - bd\, r_T^2/2}{z^2} - \frac{bd}{2} \quad (27)$$

The derivative is minimized at z=1 and is nonnegative if $bd\leq 2ar_T/(1+r_T^2)$. Substituting d reduces the condition to $(1-\gamma)(1+r_T^2)\leq 2r_T$, i.e. $H\leq H_c$. So $\lambda(s)\leq q$ in the bolus–coast case. At the boundary d=c.

For maintenance, $\lambda=q$ and $\dot{\lambda}=k\lambda-(a-bC)$ force C=c. On a coast of length L, $\lambda(\tau)=q$ gives the quadratic $(1-\gamma)r^2-2r+(1-\gamma)=0$, with r=e^(−kL). Its admissible root in (0,1) is e^(−H_c). For elapsed coast time v, let z=e^(−kv). The derivative of $k\lambda$ with respect to z is $bc(1/z^2-1)/2\geq 0$ considering the cessation boundary relation. As z decreases with v, λ decreases from q to zero on the coast. Before the cessation it is exactly q. This proves the global inequality for the loading-maintenance course.

The optimal exposure is positive and below a/b whenever input is nonzero and q>0. Hence the entire post-cessation segment before T has g(C)>0. The bolus regime has cessation at zero with the same positive-benefit property. This is the formal benefit-bearing coast statement.

Appendix B Dual Monotonicity and Resource Qualifications

For fixed H, the bolus-branch budget β is affine and strictly decreasing in γ. On the maintenance branch differentiate $\beta=(1-\gamma)(1+H-H_c)$ and use $dH_c/d\gamma=1/[(1-\gamma)\sqrt{(1-(1-\gamma)^2)}]$.

$$\frac{\partial\beta}{\partial\gamma} = -[1 + H - H_c] - \frac{1}{\sqrt{1-(1-\gamma)^2}} < 0 \quad (28)$$

Thus every active cap $0<\beta<1+H$ has one associated positive value. The treatment onset has β = 0, and q=0 gives β=1+H for the minimal dose saturated exposure. At $H=H_c$, $\beta=1-\gamma=\operatorname{sech}(H)$, proving the budget-map boundary. Above saturation the optimal exposure a/b maximizes g pointwise. The relaxed initial atom and rate ka/b reach it. A larger cap does not increase the maximum.

The fixed-value horizon statement is given by $B=c(1+kT-H_c)$. For a constant active β we differentiate $\beta=(1+H-\mathscr{D})\operatorname{sech}(\mathscr{D})$ to get $d\mathscr{D}/dH=1/[1+(1+H-\mathscr{D})\tanh(\mathscr{D})]$. Note that these derivatives are defined only in the maintenance branch and should not be extrapolated beyond the onset or saturation boundaries.

Appendix C The Continuation Policy Attribution Device

For R_1 with mandatory rate ε through T, let $A_\varepsilon=a-b\varepsilon/k$ and seek $0<L_\varepsilon<T$ satisfying the following equations. They describe a maintenance branch rather than all possible ε,B,T combinations.

$$c_\varepsilon = \frac{\varepsilon}{k} + \frac{A_\varepsilon}{b\cosh(kL_\varepsilon)} \quad (29)$$

$$B = c_\varepsilon[1 + k(T - L_\varepsilon)] + \varepsilon L_\varepsilon \quad (30)$$

The input is an initial atom c_ε, rate kc_ε until $T-L_\varepsilon$, and rate ε thereafter. Its tail exposure is $\varepsilon/k+(c_\varepsilon-\varepsilon/k)e^{-k(t-T+L_\varepsilon)}$. Assume $A_\varepsilon>0$ and $\varepsilon/k<c_\varepsilon<a/b$. Let $q_\varepsilon=(a-bc_\varepsilon)/k>0$ and write $\mu=\varepsilon dt+\nu$, where ν is the additional nonnegative input.

On the terminal mandatory-rate segment, we can perform the marginal-input calculation with a replaced by A_ε and exposure replaced by $C-\varepsilon/k$. It follows from the cosh relation that $\lambda=q_\varepsilon$ at the end of maintenance and $\lambda\leq q_\varepsilon$ afterwards. Before the boundary, $\lambda=q_\varepsilon$. The equality is satisfied on the support of ν*. The common mandatory component cancels out in the concavity inequality and there is a global optimum of benefit minus $q_\varepsilon D$. The candidate exhausts B and then the matched cap benefit optimum is given by Corollary 3.

For the declared ε=0.5 case, $L_\varepsilon\approx 0.890$, $c_\varepsilon\approx 2.961$ and $q_\varepsilon\approx 1.039$. The unrounded root satisfies both the budget relationship and the marginal input inequality. A separate 80-bin constrained solution yields a benefit approximately 6×10^{-6} below the implicit analytical solution, both rounding to 28.687. R_2 contains the unrestricted optimum, with a maintenance rate of 3 exceeding ε. In the

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

initial expansion, initial loading and redistribution are permitted; the subsequent expansion eliminates mandatory continuation.

Code and Data Availability

Code for the computational illustrations is available at https://github.com/vkaralis/LOGOS-theory-B and archived as release v1.0.1 at https://doi.org/10.5281/zenodo.23402692 [12]. The fixed-target controls, dose bounds and numerical residuals are provided in reference_results/theoretical_controls.json and reference_results/theoretical_dose_frontier.csv. The release also contains separate synthetic remifentanil-like and warfarin-like examples, which are not clinical evidence. The analytical results are stated in the text and proved in Appendices A–C. This study uses dimensionless mathematical models and no patient observations. The archived fixed-target numerical files retain full precision for reproducibility; the manuscript displays at most three decimal places.

Patient and Public Involvement: Not applicable.

Ethics and Consent: Not applicable. This mathematical and computational study did not involve human participants, patient records or animal experiments.

Funding: This research received no funding.

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

References

1. Murray JM. Some optimal control problems in cancer chemotherapy with a toxicity limit. Mathematical Biosciences. 1990;100(1):49–67. https://doi.org/10.1016/0025-5564(90)90047-3
2. Ledzewicz U, Schättler H. The influence of PK/PD on the structure of optimal controls in cancer chemotherapy models. Mathematical Biosciences and Engineering. 2005;2(3):561–578. https://doi.org/10.3934/mbe.2005.2.561
3. Fuseau E, Sheiner LB. Simultaneous modeling of pharmacokinetics and pharmacodynamics with a nonparametric pharmacodynamic model. Clinical Pharmacology and Therapeutics. 1984;35(6):733–741. https://doi.org/10.1038/clpt.1984.104
4. 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
5. 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
6. Vitková Z, Dodek M, Miklovičová E, Pavlovičová J, Babinec A, Vitko A. Robust control of repeated drug administration with variable doses based on uncertain mathematical model. Bioengineering. 2023;10(8):921. https://doi.org/10.3390/bioengineering10080921
7. Fernández LA, Pola C. Catalog of the optimal controls in cancer chemotherapy for the Gompertz model depending on PK/PD and the integral constraint. Discrete and Continuous Dynamical Systems B. 2014;19(6):1563–1588. https://doi.org/10.3934/dcdsb.2014.19.1563
8. Karalis VD. Beyond the Dose: Introducing LOGOS for Dosage Regimen Design in Therapeutics. medRxiv [Preprint]. 2026 Oct 19. https://doi.org/10.64898/2026.10.17.26365797
9. Karalis VD. LOGOS-theory-B. Version 1.0.1. Computer software. Zenodo; 2026. https://doi.org/10.5281/zenodo.23402692

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)

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

Preprints.org | NOT PEER-REVIEWED | Posted: 26 October 2026 https://doi.org/10.20944/preprints202610.2391.v1

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.