Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Finite-Course Drug Regimen Optimization: Model-Based Attribution of Dosing Flexibility and Stopping Opportunities Using LOGOS Theory

Vangelis Karalis

vkaralis@pharm.uoa.gr

National and Kapodistrian University of Athens https://orcid.org/0000-0003-0492-0712

Research Article

Keywords: Finite-course drug regimens, Pharmacokinetic–pharmacodynamic modeling, LOGOS Theory, Dose scheduling, Stopping attribution, Personalized dosing

Posted Date: October 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11484102/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Abstract

Finite-course drug administration can end before exposure or response has ended. Therefore, comparing alternative courses requires an explicit resource rule, an outcome horizon and constraints that are meaningful after the last input. We evaluated the LOGOS Theory B comparison procedure for individual models fitted to the publicly available data on remifentanil and warfarin. The main estimand was the within-profile objective increment between nested administration classes, compared at precisely the same recorded amount. Remifentanil used a concave proxy to represent effect-site coverage. Warfarin used integrated prothrombin complex activity suppression. Neither objective was a validated clinical benefit endpoint.

We used 61 jointly fitted remifentanil profiles and 32 warfarin source-file profiles. Successive remifentanil rate-cap expansions yielded median increments of 0.275900, 0.131507 and 0.131216 min, while the successive warfarin class increments were 0.593123, 0.029600 and 0.007355 h. The gap-adjusted increments were positive in 61/61 and 30/32 profiles, respectively. The median overall improvements were 9.157% and 0.758%. A post hoc comparison removed the hypothetical final quarter continuation floors while retaining all other permissions. Stopping-specific increments were much smaller: 0.001290/0.008721 min for remifentanil and 0.007434/0.053238 h for warfarin at the two floors. A distinct dimensionless injury model demonstrated approximately 20% dose saving at attainable targets and restoration of feasibility above the safe constant-family ceiling.

We have quantified the model-conditional opportunities and show that redistribution, policy-conditional stopping, resource saving and feasibility restoration should be compared separately. These analyses do not establish clinical cessation effects, clinical superiority, or calibrated safety. The distinctive contribution is within-profile attribution of extra temporal permissions and numerical gap checks, and a continuation-policy contrast to isolate stopping-specific value.

Introduction

The finite course has several temporal boundaries. Administration is allowed in a window, benefit is valued over a chosen period and adverse effects may require assessment over a longer time frame. These periods are associated with residual concentration, effect-site equilibration, absorption, and pharmacodynamic turnover. So moving the same amount from one time to another can change a finite window response even if full-tail exposure is kept. The mathematical value varies depending on which response is valued and which choice of administration is allowed.

There are already a few individualized optimal control methods that deal with dose and schedule choice [1]. LOGOS, Longitudinal Optimization of Global Outcomes and Strategies, asks a complementary attribution question: How does each additional decision permission contribute to the same model, resource rule, initial state, endpoint and trajectory constraint? The main LOGOS theory is outlined in the medRxiv preprint [2]; the current application is temporal allocation and stopping. The main difference is between redistribution and stopping. A class with arbitrary nonnegative input already allows terminal

zero input. Calling cessation an additional decision does not increase the class and cannot generate an independent stopping value.

An identifiable stopping comparison requires a smaller class with real continuation obligation. Removing only that obligation while keeping all other permissions defines a policy-conditional increment. The earlier last administration in a flexible schedule is descriptive and does not measure the value of the stopping decision. Similarly, a better objective with fixed amount does not show dose saving. Dose saving requires minimizing amount at a common target, whereas restoration of feasibility requires showing that a target unavailable in the reference class becomes attainable.

Public remifentanil observations provide controlled infusion, concentration, and electroencephalographic responses [3, 4]. Historical warfarin observations provide oral input, concentration, and prothrombin complex activity, or PCA [5]. They support fitted-model comparisons of infusion-rate freedom and dose redistribution. Two pediatric methotrexate datasets provide contextual biomarker evidence of persistence [6, 7], but contain no identified comparison of alternative cessation policies. This paper combines model-based attribution with deliberately separate theoretical and observational evidence. It aims to quantify the opportunities supported by each component without assigning a clinical interpretation that the data cannot establish.

Another contribution is drug-specific application of nested permission attribution, i.e. each additional scheduling freedom is assessed in the same fitted profile with the same realized amount and common exposure guardrails. Numerical upper gap checks distinguish a class increment from an incomplete optimization of its comparator. The continuation floors are removed separately to isolate stopping-specific value and a dimensionless fixed-target analysis distinguishes dose saving from restored feasibility. We can use this model to determine which freedom is important in each reported opportunity and what clinical questions are unanswered.

Methods

Design and source populations

We conducted secondary analysis of public observations, individual model fitting, and counterfactual simulation. No alternative regimen was administered. The remifentanil source contained 6,046 records from 65 individuals. IDs 9, 24, 33, and 39 lacked EEG observations, leaving 61 joint PK/PD profiles. The warfarin source had 515 records and 32 IDs (including ID 100; ID 10 was absent from the source file). All the source profiles were kept without any inference of merging. The official documentation has conflicting participant descriptions, so 32 is a number of the source-file profiles and not an independently verified count of unique study participants [5].

The methotrexate course has 172 records from 43 children and the rescue dataset has 67 participant records. These data are not linked. Missing markers are counted as unavailable measurements and never zero. Sample sizes were determined by data availability, and all eligible profiles were included.

Successive nested-class increments in the two PK/PD applications were the primary descriptive analysis. Parameter, horizon, initialization, and numerical sensitivities were secondary; continuation-floor and extended robustness analyses were post hoc. We do not claim prospective registration or a prospective power calculation.

The comparison estimand

Let W be the administration window, H the outcome horizon and D the administered amount. For each main comparison we fix D to the amount recorded for that profile: D_i. This matches realized resource within a profile, rather than a common amount across profiles or drugs. Instead, a cap comparison gives us a common ceiling B and may allow different realized amounts. Different rules answer different questions and they are not exchanged during attribution.

For nested feasible families $F_0 \subseteq F_1 \subseteq F_2$ the optimal values are $V_0 \le V_1 \le V_2$ as we have the same goal and common constraints. We calculate each candidate increment $\Delta J = J_{next} - J_{previous}$ within profile before aggregation. The summaries are medians, interquartile ranges where relevant, and counts above 10^{-6} in the application's objective unit. The nonnegative optimum increments are results of nestedness. Positive counts are descriptive numerical results and not tests of clinical treatment effectiveness.

Relative total gain is calculated as $100(J_{flexible}/J_{reference} - 1)$ within each profile. Relative successive increments use the preceding class as denominator. Neither difference between population medians nor ratio of aggregate objectives substitutes for the paired calculation. The two drug applications remain separate because their endpoints, units, histories, and admissible inputs differ.

Remifentanil model and classes

Sequential individual fitting used a positive biexponential plasma kernel, an effect-site equilibration state, and a decreasing Hill EEG response. Plasma concentration is the convolution of input u with $w_1 \exp(-k_1 t) + w_2 \exp(-k_2 t)$. The effect-site state satisfies $C_e' = k_{e0}(C_p - C_e)$. EEG predictions use an individually fitted baseline, maximal fractional decrease, half-effect concentration K, and Hill exponent. These simplified fits are not re-estimates of the original population model [3].

All of the eligible profiles have one positive infusion record at time zero with rate r and amount D_i, so $W = D_i/r$. The file does not record any further interruption or rate change. We now refer to the constant infusion that is represented by that record. Flexible schedules divide W into 20 equal bins, contain no initial loading atom, start from zero plasma and effect-site states and deliver exactly D_i. Rate caps are r, 1.25r, 1.5r and 2r. For each larger class we keep the previous class. At the cap r the exact amount forces all the bins to the reference rate and so we get the algebraic null.

The objective is $J_{Rem} = \int_0^{\square} C_e/(K + C_e)\, dt$, measured in minutes. Its exponent is fixed to one to give a concave coverage proxy under linear state propagation. It will not reproduce all fitted Hill responses and can not measure analgesic benefit or respiratory injury. All classes have ceilings equal to the reference full-path plasma and effect-site maxima and the reference effect-site AUC after W. These are exposure guardrails. The last active bin ends where the final rate above 10^{-4}r ends. The timing depends on binning and threshold.

Warfarin model and classes

For the individual plasma models, we use first-order oral absorption with lag and linear elimination. The common absorption and lag values were empirical medians from 12 early-sampling-eligible profiles in the same source. Such regularization has the effect of dependence of the fitted profiles. For the normalized PCA R, the turnover model is given by $R' = k_{out}[K/(K + C_p) - R]$ and the primary value $R(0) = 1$. The objective $J_{War} = \int_0^{120}(1 - R)dt$ is measured in hours. More integrated suppression is its mathematical direction, not a claim of better anticoagulation.

The source-record reference is the historical amount at time zero. R_{24}, R_{12}, and R_6 distribute exactly that amount across nonnegative continuous dose variables on 24-, 12-, and 6-hour grids over 72 h. Each finer grid includes all the earlier dates and allows zero amounts at added dates. Parameters, initial state, absorption assumptions and reference full-path plasma-peak ceiling remain the same. Tablet increments and clinically qualified dose limits are not there. These are just counterfactual dosing opportunities, not new clinical regimens or advice.

Baseline PCA was measured in 30 profiles: 21 values were 100% and nine were 82% to 92%. The IDs 6 and 100 did not have a measurement at time zero. In other words, the primary initial fraction of one is an assumption, not a normalization of the baseline for every profile. Post hoc, we fixed measured baseline where available and refitted turnover parameters. Plasma extrema were checked between input and absorption transitions and along the declining tail. No INR target, bleeding outcome, thrombosis outcome, or calibrated clinical injury constraint was analyzed.

Stopping attribution and other extensions

Stopping-specific value was evaluated in the same largest flexible families by adding or removing final-quarter continuation floors. For remifentanil, five final-quarter bins must each have rate at least εr. For warfarin, each of the four opportunities at 54, 60, 66, and 72 h must have dose fraction at least ε/13. The analyzed values were ε = 0.05 and 0.25. Removing these lower bounds preserves dose equality, timing permissions, objectives, initialization, and exposure guardrails.

The floors commit 1.25% and 6.25% of remifentanil amount, versus 1.5385% and 7.6923% of warfarin amount. Identical ε labels therefore do not denote cross-drug-equivalent obligations. The restrictions are

hypothetical analyst-defined policies, not continuation requirements from prescribing records. The contrasts only capture stopping increments in fitted models.

Additional remifentanil analyses reoptimized the fitted-Hill objective from four starts and tested a shared initial plasma and effect-site state equal to 1% of nominal reference peaks. Warfarin extensions tested observed-baseline refitting, four absorption/lag perturbations with individual refits, and leave-one-profile-out estimation of common parameters. Own-profile observations remained in individual fitting. These procedures test explicit assumptions; they do not provide external or alternative-regimen validation.

Uncertainty and numerical verification

The archived bootstrap resampled whole profiles 3,000 times with fixed fits and stored class values. It preserves paired comparisons but excludes parameter fitting and schedule reoptimization. A separate 1,000-replicate warfarin residual-and-profile bootstrap resampled centered within-profile PD residuals, refitted turnover parameters, reoptimized R_6, and then resampled profiles. Its 32,000 fits and optimizations condition on fixed PK, common absorption/lag, initial states, model structure, and residual exchangeability. It propagates only part of scientific uncertainty.

For a concave finite-family objective, the tangent at a candidate provides an upper bound when maximized over the same family or an outer relaxation containing every continuously feasible control. If the smaller-class objective gap is g_{small}, a feasible larger-class candidate yields the increment lower bound $\Delta J - g_{small}$. Independent objective-error allowances are also deducted where specified. The larger-class gap is not needed for that lower bound. Such bounds inherit floating-point, quadrature, and constraint-evaluation assumptions and do not certify an unrestricted continuous-input problem.

Remifentanil verification combined segment quadrature, stationary-point searches, independent DOP853 integration split at input changes, and analytical residual tails. Relative tolerances were 10^{-7} for peak/tail excess and 10^{-8} for amount equality. Warfarin objective agreement was checked at 10^{-6} h and peak ratios at $1 + 10^{-6}$. Fitted-Hill reoptimization can be nonconcave and is reported as multistart candidate evidence but without any global certificate.

Computational environment

The analyses were carried out in Python 3.12 with Jupyter notebooks. The original analysis archive contains CPython 3.12.14, NumPy 2.3.5, SciPy 1.17.1, pandas 2.2.3, and Matplotlib 3.10.9. These versions describe the archived analysis environment.

Code for the present study is available at https://github.com/vkaralis/LOGOS-theory-B-apps and archived on Zenodo at https://doi.org/10.5281/zenodo.23402705 [8]. This DOI identifies the code release. The public examples use synthetic data and do not reproduce the original empirical analyses.

Results

Primary redistribution opportunities

The remifentanil median PK log RMSE is 0.1161 and EEG RMSE is 1.5199. Interleaved omitted-observation diagnostics supported delay of effect site in the 56/61 profiles over direct response. The diagnostic refits the model on the retained observations and evaluates omitted times within the same profile. The diagnostic is used to assess interpolation performance under the observed infusions, but does not support predictions for new regimens. In total, there were seven PD fits that approached their parameter limits, but by excluding them we found that the median primary gain is 8.983% in the 54 profiles we investigated.

Rate-cap factors 1.25, 1.5, and 2 gave median total coverage gains of 4.415%, 6.754%, and 9.157%. The successive paired increments recorded were 0.275900, 0.131507 and 0.131216 min, with positive gap-adjusted lower bounds in all 61 profiles at each transition. The corresponding interquartile ranges were 0.219576–0.318029, 0.111154–0.156611, and 0.119404–0.164307 min. The cap-one implementation exactly reproduced the reference. Earlier last-bin endpoints occurred in 45, 53, and 58 profiles as the cap increased.

Warfarin median PCA RMSE was 2.9728 percentage points. A leave-one-later-observation diagnostic favored turnover over direct response in 30/32 profiles, conditional on full-data PK and common absorption/lag. Common absorption was 2.203766 h^{-1} and lag 0.876524 h. Total median suppression-proxy gains were 0.707%, 0.748%, and 0.758% for R_{24}, R_{12}, and R_{6}. Paired increments were 0.593123, 0.029600, and 0.007355 h, positive after gap adjustment in 30/32 profiles at each transition. The remaining two lower bounds were zero.

Most warfarin gain arose from releasing the single-dose restriction. Later grid refinements added smaller increments. This pattern describes the selected model and endpoint, rather than a general rule about dosing frequency. Across drugs, the numerical values cannot be ranked because a minute of coverage and an hour of PCA suppression are different quantities. The conditional 95% profile-bootstrap interval for the largest-class median gain was 8.273–9.736% for remifentanil and 0.515–1.097% for warfarin. Table 1 and Fig. 1 summarize the paired distributions. The large warfarin range suggests that there is a large heterogeneity between profiles.

Table 1
Distributions of paired increments between successive nested classes

Application and transition	Median ΔJ	IQR	Range	n+	Median relative increment %
Remifentanil 1 to 1.25	0.275900 min	0.219576–0.318029	0.127118–0.529572	61/61	4.4149
Remifentanil 1.25 to 1.5	0.131507 min	0.111154–0.156611	0.078483–0.212551	61/61	2.2138
Remifentanil 1.5 to 2	0.131216 min	0.119404–0.164307	0.064908–0.219337	61/61	2.2659
Warfarin R0 to R24	0.593123 h	0.355160–1.059016	0.000000–7.391979	30/32	0.7073
Warfarin R24 to R12	0.029600 h	0.020903–0.047924	0.000000–0.308671	30/32	0.0364
Warfarin R12 to R6	0.007355 h	0.004725–0.010451	0.000000–0.067845	30/32	0.0085

IQR, interquartile range. IQR and range use minutes for remifentanil and hours for warfarin. n+ counts candidate increments above 10^{-6} in the corresponding unit; relative increments are medians of $100\Delta J/J_{previous}$. Gap-adjusted counts are reported in the text.

Policy conditional stopping

Removing ε = 0.05 and 0.25 floors, median remifentanil increments were 0.001290 and 0.008721 min, both of which had positive lower bounds in the 61/61 profiles. Warfarin increments were 0.007434 and 0.053238 h with positive lower bounds in 28/32 profiles for each of the floors. Four profiles did not resolve a strictly positive warfarin increment at the specified tolerance. Small negative numerical lower bounds were left as unresolved bounds not interpreted as negative optimal stopping value. Table 2 reports these increments alongside the minimum amount committed by each continuation floor.

Table 2
Additional value from removing final-quarter continuation floors

Drug	Floor ε	Profiles	Median stopping ΔJ	Positive lower bounds	Minimum committed amount %
Remifentanil	0.05	61	0.001290 min	61/61	1.2500%
Remifentanil	0.25	61	0.008721 min	61/61	6.2500%
Warfarin	0.05	32	0.007434 h	28/32	1.5385%
Warfarin	0.25	32	0.053238 h	28/32	7.6923%

Positive lower bounds exceed 10^{-6} in the objective's units after subtracting the continuation-family supporting gap and both independent ODE objective errors. Unresolved bounds are retained. Identical ε values impose different committed fractions across drugs.

These increments are smaller than the wider redistribution gains. They quantify the removal of the stated obligation while keeping the other decisions common. The earlier-input counts in the primary analysis and the new stopping increments answer different questions and cannot be interchanged. Higher floor strength changes the policy being eliminated; it does not establish a measured clinical dose-response relationship for cessation.

Sensitivity and post input contribution

Fitted-Hill remifentanil reoptimization achieved a median gain of 8.838%, with positive gains in all profiles. It improved on evaluating the nominal schedules under that objective in 13 profiles, with a largest gain change of 0.558 percentage points. The same 1% initial state sensitivity kept the gains positive, with a median of 8.802%. Refinement to 40 and 80 bins produced modest changes in six selected profiles; this does not establish grid independence across the cohort.

At remifentanil outcome horizon 1.5W, median gain drops to 3.566% and 31 profiles had an earlier last bin. Gain grew to 13.131% and no profile had an earlier last bin at 0.75W. Administration was still allowed through W so exact dose could result in input after evaluation. The results from this short-horizon analysis do not support the view of regimen stopping. If we only checked our exposure in the administration window W, that could allow for a subsequent effect-site peak ratio of up to 1.007683. Note that stable objective summaries do not guarantee that all path constraints are met.

In our study we found that warfarin refitting based on the observed baseline results in a median R_6 gain of 0.757%, and restricting the analysis to the 30 profiles with measured baseline yielded a gain of 0.671%. The four pooled-parameter perturbations yielded median R_6 gains of 0.672–0.761%. We found that the leave-one-profile-out common parameter estimation with the individual refitting provided 0.788%. The conditional 95% residual-and-profile bootstrap interval was 0.53–1.13%; without profile resampling, it was 0.67–0.82%. All 32,000 fits and optimizations reported success. These are conditional intervals for model proxy gains and not treatment effects.

In addition, the warfarin objective was extended to 240 and 336 h, and the median gains were 10.494% and 17.487%. All profiles were then extrapolated beyond their last PCA measurement by 96–192 and 192–288 h, respectively. Even the main 120-h horizon is greater than the last PCA measurement of one profile by 72 h. The long-horizon gains indicate objective and turnover memory sensitivity but no better efficacy.

The candidate objective fraction after final active input had median 12.34% for remifentanil, and 65.86% for warfarin. Warfarin's fraction includes ongoing absorption and previously established suppression. These fractions indicate persistence within a candidate; they do not compare that candidate with an otherwise matched continuation policy.

Theoretical resource and safety results

A separate dimensionless benchmark uses $C' = -C + u$, $g(C) = 4C - C^2/2$, benefit horizon 4, and available cap $B = 12.6139036163$. The relaxed input class permits an initial atom. Without the injury constraint introduced below, its analytical optimum loads to concentration 3, maintains, and ceases at 3.204635. The benefit is 28.826262 vs 25.460925 for the best constant no-loading reference at that amount. Note that the continuation-constrained class has the same initial loading permission and the same total amount but we need $u(t) \geq 0.5$ over (0,4). Its optimum $J = 28.686784$. Removing only the floor adds 0.139478 and the continuation-constrained gain of 3.225858 over the constant reference reflects redistribution and loading.

The quadratic benchmark is self-contained. For plateau c and the remaining horizon δ, the marginal dose value on plateau is $4 - c$. On the zero input coast it is $4[1 - \exp(-\delta)] - (c/2)[1 - \exp(-2\delta)]$. For equality $c = 4/\cosh(\delta)$. For $c = 3$ this gives $\delta = \operatorname{arcosh}(4/3)$, the cessation time is $4 - \delta$ and $B = 3[1 + 4 - \delta]$. Since g is concave and PK is linear, the supporting marginal inequality is sufficient to get the global optimum in the ideal nonnegative measure-input class. The plateau is $\varepsilon + (4 - \varepsilon)/\cosh(\delta)$ under the continuation floor ε. The exact amount determines δ. These are analytical model comparisons, separate from fitted profile analyses. Figure 2A shows the matched-amount concentration profiles; Fig. 2B illustrates why delayed injury requires a separate assessment horizon.

A distinct injury model adds $Z' = C - 0.2Z$ and cap $Z \leq 2.5$ over the complete future path. Its safe constant-family ceiling is 11.270930290. The flexible candidate at target 13 described below is feasible with a positive numerical injury margin. Thus targets above the constant ceiling and up to 13 are attainable by that witness. This is an existence result against an analytical reference ceiling, not a certified unrestricted flexible optimum.

For the constant family, u is constant on [0,4], zero thereafter, and no loading atom is permitted. Its objective is $4uI_1 - u^2I_2/2$, where $I_1 = 4 - 1 + \exp(-4)$ and $I_2 = 4 - 2[1 - \exp(-4)] + [1 - \exp(-8)]/2$. Linearity gives max $Z = uM$, where M is the full-tail peak for unit input. That peak occurs at $4 + \log\{[1 - \exp(-4)]/[1 - \exp(-0.8)]\}/0.8$. Maximizing the quadratic over $0 \leq u \leq \min(B/4, 2.5/M)$ gives the stated ceiling. Flexible dose candidates use an initial nonnegative atom and 80 equal infusion bins on [0,4], each with rate at most 20, under the same resource cap B and complete-path injury cap.

At fixed targets 10 and 11, flexible dose candidates were 2.928587 and 3.266454, versus constant-family minima 3.666055 and 4.081999. Savings were 20.1161% and 19.9791%. At targets 12 and 13, flexible doses were 3.614249 and 3.973756 while the constant family was infeasible. Percentage saving is undefined in those cases. Dose bounds concern the stated 80-bin-plus-atom family. These dimensionless results are not drug-specific dose reductions or clinically calibrated safety findings.

Biomarker context and computational agreement

Among 35 children with paired first- and fourth-course measurements, erythrocyte MTX-PG2–5 increased by a median of 71.50 nmol/L with a 95% pair-bootstrap interval of 59.10–95.10. A 12–16-day sampling restriction retained 26 pairs, with a median increase of 67.55 nmol/L. Eight children did not have complete endpoint pairs, and no complete pair had recorded transfusion code zero at both measurements. Course order does not indicate a continuous effect trajectory. Biochemical changes in the rescue cohort were paired but neither data set identifies the value of stopping.

The main audits were 244 remifentanil schedules and 96 flexible warfarin schedules. We found that the residuals of the constraints are within the specified tolerances. The largest supporting gaps were 1.464×10^{-6} min and 6.411×10^{-6} h. Our initial largest class reconstruction agreed with archived objectives within 2.44×10^{-8} min and 6.13×10^{-7} h. These calculations provide some insight into the numerical consistency of our models but do not validate the counterfactual clinical response.

Discussion

The main contribution is a matched attribution procedure combining within profile nested scheduling comparisons, gap-adjusted increments and continuation-floor removal in two PK/PD systems. The main comparisons are to measure the time freedom at the same realized amount and common exposure guardrails. The continuation-floor comparisons are to isolate a smaller stopping opportunity. The theoretical question is to separate resource saving from feasibility restoration. Combining these quantities into a single claim of advantageous cessation would discard the conditions that make each comparison interpretable.

Persistence can explain why schedule timing matters, but does not prove stopping is beneficial. Full-tail linear exposure can be set by amount and clearance of the material while a finite-window nonlinear objective changes. Also, the effect site delay, oral absorption, and turnover play a role in time-weighting. The outcome horizon is then part of the estimand and not a neutral numerical setting. The large extended-warfarin gains illustrate how a change in the evaluation window changes the question while increasing extrapolation.

Exposure guardrails also have limited interpretation. Matching reference peaks and tail AUC limits selected model quantities but does not validate respiratory harm, bleeding, thrombosis or other clinical outcomes. Complete path verification may prevent missed model violations after administration but clinical safety also requires an adequate injury model and observed follow-up. For example, after a unit loading atom with $C(t) = \exp(-t)$ and $Z' = C - 0.2Z$, $Z(0) = 0$, the injury is $Z(t) = [\exp(-0.2t) - \exp(-t)]/0.8$ and shows a peak at $\log(5)/0.8$ even if there is no subsequent treatment. This shows the delayed injury but does not provide a clinical calibration for either application. In Fig. 2B, the injury peak occurs after the short outcome window, illustrating the risk of truncating constraint assessment.

There are also uncertainties in fitting under one historical input. Such fitting does not qualify doubled infusion rates or divided oral doses. The main bootstrap has the parameters fixed, while the extended bootstrap propagates only PD residual and profile variation. Common-parameter dependence, PK

uncertainty, model choice, residual assumptions, missing baseline states and transport to other regimens are still in the early stages of their assessment. There is no global nonconcave certificate for multistart Hill candidates. We were not able to estimate delivery constraints and clinically meaningful stopping requirements.

A definitive application would prespecify a clinical endpoint, nested attainable classes, an exact-amount or cap resource rule, and separate benefit and safety horizons. It would validate alternative-regimen predictions with independent patients or genuinely separated temporal data, propagate relevant fitting uncertainty through optimization and test delayed outcomes beyond final treatment. An observational cessation analysis would also need to consider indication, informative stopping, and time-varying confounding. Our data support the fitted-model attribution part of that program.

Why finite window timing differs from total exposure

For a zero-state linear elimination model with a finite nonnegative input, we can integrate the state equation across the whole decreasing tail and get $k\int_0^\infty C(t)dt = D/V$ where D is the intravenous amount and V is a constant volume. For normalized concentration-equivalent units with V = 1 this is simply D. An equal amount then determines the unweighted linear exposure in the full tail if k is common but this does not fix exposure in a fixed window of time since the input at time s is weighted $[1 - \exp(-k(H - s))]/(kV)$ for $s \le H$ and zero for $s > H$. The later inputs have less time to contribute before H; if we assume nonlinear response to concentration then the response also depends on how the exposure is distributed. Note that the timing null is valid for cumulative-only endpoints but not for every instantaneous response model.

Warfarin provides an explicit delayed-response version of this distinction. With initial suppression zero, integrating the turnover solution yields $J(H) = \int_0^\square [C_p(s)/(K + C_p(s))][1 - \exp(-k_{out}(H - s))]ds$. If H is changed, the same plasma history is under different weights. Increasing H is a change in endpoint specification, even if the dosing dates and amount are stable. It may increase the opportunity for the model but also force it to predict outside observation support. The horizon stress tests reveal that dependence rather than resolving the appropriate clinical horizon.

Concavity is useful for numerical attribution but it is not sufficient for clinical adequacy. For exponent one coverage, the marginal gain decreases with exposure and therefore we can give an upper supporting tangent. For warfarin suppression with common initialization, the objective is a positive temporal weighting of a concave concentration response. Linear PK maps are used to transfer concavity to the dose. A fitted-Hill objective does not need to have this property. So mathematical tractability is endpoint-specific and a certificate obtained for one proxy cannot be assigned to another readout simply because the two use the same fitted concentrations.

Conditions for a fair comparison

Matching total amount is necessary for the main question, but that is not enough. The comparison must also have the same initial state, parameter estimates, concentration scale, outcome direction, outcome horizon and exposure constraints. Otherwise, an apparent increase in class might be due to a different starting state or a relaxed guardrail. If the sensitivity changes its initial state or parameters, its reference and flexible classes must both use the changed assumptions and the corresponding recalculated reference ceilings. Its within-scenario contrast remains interpretable, while direct comparison of unmatched objective levels does not.

Attainability is another boundary. A continuous nonnegative dose variable is a modeling permission, while tablet strength, pump rate and minimum practical intervals impose implementation limitations. These limitations can lower the feasible class or make impossible the exact historical amount matching. The problems that arise from this analysis should be presented as a new attainable comparison with a well-defined resource rule. This new comparison should not inherit the results of the more lenient family. We need to remember that the historical reference for warfarin is a context-sensitive reference since it is the origin of the experiment and is not a current treatment protocol.

The numerical uncertainty question is also smaller than the pharmacological uncertainty question. A positive candidate difference can arise if the smaller class was not fully optimized. Subtracting the smaller class's upper objective gap addresses that concern. Independent integration checks propagation and quadrature agreement, while stationary-point checks address maxima between grid nodes. These procedures do not establish prediction accuracy for an unobserved regimen. Tight numerical agreement can coexist with incomplete counterfactual qualification.

Finally, stopping attribution is conditional on the obligation being removed. If a free optimizer already satisfies a very small continuation floor, removing it can be of zero value. A stronger floor can prevent any good redistribution and lead to a larger increment. With our same model, we are seeing an order of feasible policies rather than a biological value of cessation. A clinical policy contrast would require an observed or justified continuation rule and an endpoint to assess the effect of releasing that rule. The hypothetical floors define a model attribution experiment rather than a clinically validated continuation policy.

Conclusion

We provide a within-profile approach to distinguish the value of additional scheduling permissions from the narrower value of releasing a continuation obligation. LOGOS Theory B provides one conditional framework for attributing finite-course opportunities. Remifentanil and warfarin models show positive redistribution increments for declared resources and guardrails and smaller stopping increments for hypothetical continuation floors. Another model for injury distinguishes dose saving from restoring target feasibility. These results establish model-specific numerical opportunities while leaving clinical benefit, clinical stopping effects, operational attainability, and calibrated delayed safety unvalidated.

Declarations

Data and code availability

Original data sources are identified in [4–7]. Code for the present study is available at https://github.com/vkaralis/LOGOS-theory-B-apps and archived as version 1.0.1 at https://doi.org/10.5281/zenodo.23402705 [8]. The Zenodo DOI points to the code release, but not to the main theory preprint or clinical data. The public release contains synthetic examples and theoretical benchmarks, but no original clinical data and patient-fitted parameters, nor does the public release reproduce the empirical gains and sensitivity analysis, or full bootstrap campaigns. True verification of synthetic workflows is not the same as empirical reproduction.

Patient and public involvement

Not applicable.

Ethics and consent

This study analyzed publicly released, previously collected data and mathematical models. No new participants were recruited, no intervention was administered and no new patient contact or consent collection took place. The original studies and data sources are cited in the Methods and References.

Funding

This research received no funding.

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;189:46–65. https://doi.org/10.1007/s10957-021-01819-w
2. 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
3. Minto CF, Schnider TW, Egan TD, et al. Influence of age and gender on the pharmacokinetics and pharmacodynamics of remifentanil. I. Model development. Anesthesiology. 1997;86(1):10–23. https://doi.org/10.1097/00000542-199701000-00004
4. Simulations Plus. Remifentanil data set. MonolixSuite Documentation, 2024R1. https://monolixsuite.slp-software.com/monolix/2024R1/remifentanil-data-set. Data retrieval recorded 8 October 2026.
5. Simulations Plus. Warfarin joint data set. MonolixSuite Documentation, 2024R1. https://monolixsuite.slp-software.com/monolix/2024R1/warfarin-joint-data-set. Data retrieval recorded 8 October 2026.

6. Oosterom N, Fiocco M, Kloos RQH, et al. The evaluation of red blood cell folate and methotrexate levels during protocol M in childhood acute lymphoblastic leukemia. BMC Cancer. 2020;20:940. https://doi.org/10.1186/s12885-020-07422-y. Dataset: https://doi.org/10.6084/m9.figshare.12909395.
7. Oosterom N, de Jonge R, Smith DEC, et al. Changes in intracellular folate metabolism during high-dose methotrexate and Leucovorin rescue therapy in children with acute lymphoblastic leukemia. PLOS ONE. 2019;14(9):e0221591. https://doi.org/10.1371/journal.pone.0221591
8. Karalis VD. LOGOS-theory-B-apps. Version 1.0.1 [Computer software]. Zenodo; 2026. https://doi.org/10.5281/zenodo.23402705. Source code: https://github.com/vkaralis/LOGOS-theory-B-apps

Figures

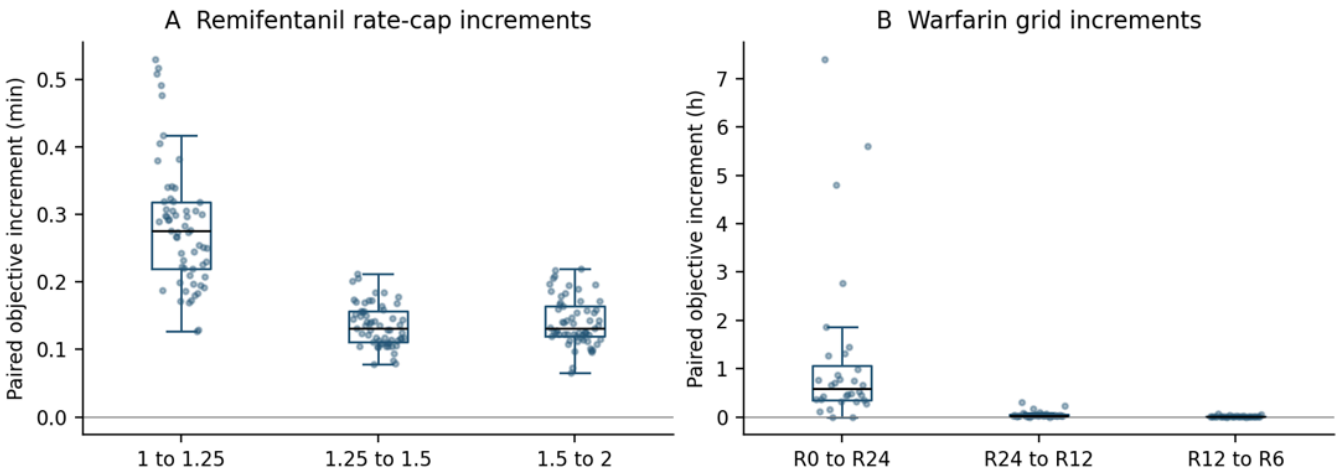


Figure 1

Individual paired increments between nested classes. Boxes show quartiles and medians; whiskers follow the 1.5-IQR rule, and points show all profiles. Panel units differ. Changes concern model proxies, not clinical benefit.

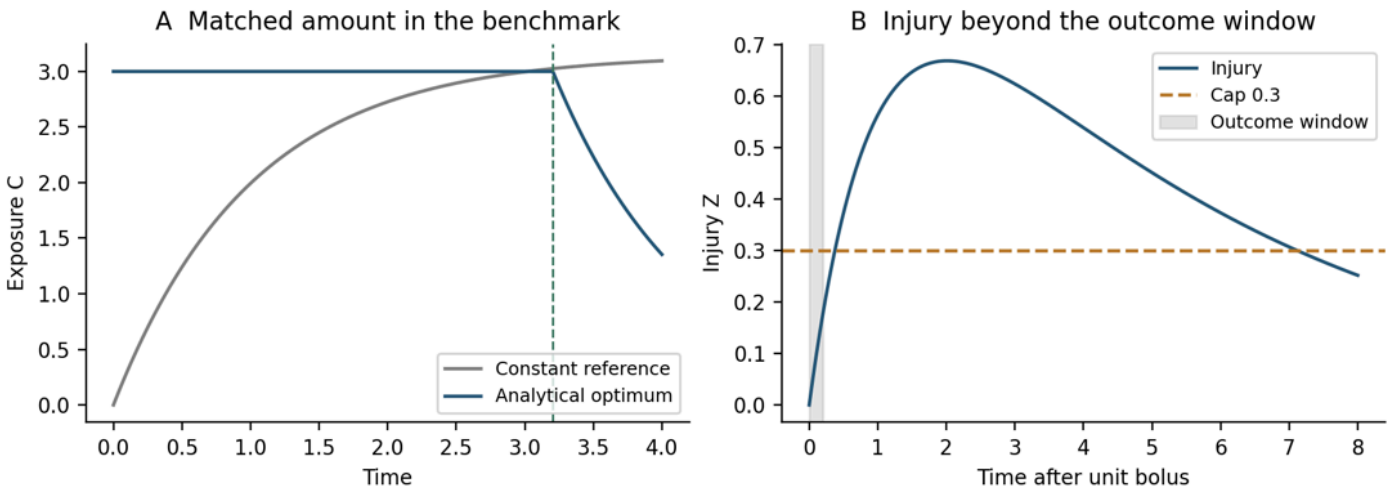


Figure 2

Dimensionless theoretical examples. (A) Analytical quadratic optimum and optimized no-loading constant reference at matched amount; the dashed line marks cessation. (B) Delayed injury after a separate unit bolus exceeds the illustrative cap of 0.3 beyond the shaded outcome window. The panels represent separate comparisons and do not establish clinical safety.