Pharmacophore

ISSN-2229-5402

Journal home page: http://www.pharmacophorejournal.com

THE DOSE TWIN: SIMULATING INDIVIDUALIZED TREATMENT UNDER CHANGING PHYSIOLOGY, DISEASE PROGRESSION, ADHERENCE, AND THERAPEUTIC RESPONSE

Peter Otieno[1]*, Grace Atieno[2], Felix Okoth[3], Caroline Moraa[4], Thomas Müller[5]

1. ***Department of Dose Twin and Individualized Treatment Simulation, Faculty of Pharmacy, Kenya Medical Research Institute, Kisumu, Kenya.***
2. ***Department of Changing Physiology and Disease Progression, Faculty of Pharmacy, University of Nairobi, Nairobi, Kenya.***
3. ***Department of Adherence and Therapeutic Response Modeling, Faculty of Pharmaceutical Sciences, Jomo Kenyatta University of Agriculture and Technology, Nairobi, Kenya.***
4. ***Department of Personalized Dose Simulation, Faculty of Pharmacy, Moi University, Eldoret, Kenya.***
5. ***Department of Treatment Optimization and Digital Twins, Faculty of Veterinary Medicine, University of Zurich, Zurich, Switzerland.***

ARTICLE INFO

Received:
17 January 2025
Received in revised form:
30 March 2025
Accepted:
04 April 2025
Available online:
28 April 2025

Keywords: Model-informed precision dosing, Computational pharmacology, Digital twin, Longitudinal pharmacokinetics, Treatment adherence, Counterfactual simulation

ABSTRACT

Drug dosing is commonly initiated from population evidence and refined through covariates, therapeutic drug monitoring, pharmacokinetic–pharmacodynamic models, or clinician judgment. These approaches remain essential, but they do not necessarily provide a unified representation of an individual whose physiology, disease state, treatment adherence, drug exposure, therapeutic response, and toxicity risk change concurrently. This article proposes the **Dose Twin**, a patient-linked and time-indexed computational pharmacology decision model for reconstructing actual treatment, estimating evolving pharmacological and disease states, and comparing bounded counterfactual dose regimens. The construct integrates population priors with longitudinal patient observations while preserving explicit distinctions among measured variables, latent states, mechanistic assumptions, learned components, uncertainty, and clinical decisions. Its proposed architecture contains a treatment-history reconstruction layer, dynamic physiology and disease representations, a pharmacokinetic–pharmacodynamic or quantitative systems pharmacology core, an adherence model, a sequential updating mechanism, a counterfactual dose evaluator, safety and uncertainty gates, and a clinician authorization layer. The central contribution is not a new optimization metric or autonomous dosing algorithm, but an organizing decision model that treats individualized dosing as repeated inference under changing and incompletely observed conditions. It also prevents apparent treatment failure from being attributed automatically to altered pharmacology when missed doses, delayed administration, disease progression, measurement error, or model misspecification remain plausible explanations. The Dose Twin is presented as a conceptual and methodological contribution requiring context-specific technical verification, pharmacological validation, prospective clinical evaluation, workflow assessment, and governance. It does not establish individual causal effects, clinical benefit, regulatory acceptability, or deployment readiness. Its value lies in defining what must be represented, updated, tested, bounded, and communicated before longitudinal treatment simulation can responsibly influence model-informed precision dosing.

To Cite This Article: Otieno P, Atieno G, Okoth F, Moraa C, Müller T. The Dose Twin: Simulating Individualized Treatment under Changing Physiology, Disease Progression, Adherence, and Therapeutic Response. Pharmacophore. 2025;16(2):53-64. https://doi.org/10.51847/ZY50DJBhRo

Introduction

Drug dosing begins with population knowledge because treatment cannot be designed independently for every patient before therapy starts. Label recommendations, population pharmacokinetic models, exposure–response analyses, organ-function

Corresponding Author: Peter Otieno; Department of Dose Twin and Individualized Treatment Simulation, Faculty of Pharmacy, Kenya Medical Research Institute, Kisumu, Kenya. E-mail: peter.otieno@kemri.go.ke.

adjustments, and therapeutic drug monitoring provide necessary prior information. The unresolved problem is that these resources are usually applied as discrete dosing rules or episodic corrections, whereas the treated patient is a changing biological and behavioral system. Individualized dosing therefore requires more than selecting a subgroup-specific regimen: it depends on fine-grained patient-specific adjustment, qualified pharmacological models, suitable longitudinal data, usable software, clinical workflows, and accountable implementation. The persistence of scientific, infrastructural, evidentiary, and organizational barriers explains why model-informed precision dosing has not yet become routine despite its conceptual maturity [1-3].

The inadequacy of a single performance measure follows directly from this longitudinal problem. Agreement between predicted and observed concentrations may demonstrate a limited aspect of predictive adequacy, while target attainment may show that a modeled exposure objective can be reached. Neither result alone establishes that the model has represented the correct exposure-generating process, distinguished disease change from adherence change, improved a patient-relevant outcome, or remained reliable after physiology and treatment conditions shifted. Improved prediction or simulated target attainment must therefore be separated from prospective clinical usefulness [4]. A dosing model can be numerically accurate over a restricted interval yet become scientifically misleading when it extrapolates beyond its qualified population, assimilates erroneous treatment histories, suppresses structural uncertainty, or treats a changing patient as a fixed parameter vector.

Digital-twin scholarship offers a useful but incomplete conceptual precedent. In medicine, a digital twin is generally understood as a patient-linked computational representation that can be informed by individual observations and used to simulate clinically relevant states or interventions [5]. Direct transplantation of that terminology into dosing would nevertheless be insufficient. A pharmaceutical treatment twin must represent not only the patient but also the administered regimen, the timing and reliability of observations, drug disposition, target engagement, response delays, toxicity, disease evolution, adherence, and the consequences of interventions that themselves alter subsequent data. The object of representation is therefore not a generic virtual patient. It is the evolving relationship among a particular patient, a particular therapy, an intended treatment objective, and a defined clinical context of use.

This article develops the Dose Twin as an original computational pharmacology decision-model construct. The Dose Twin is defined as a patient-linked, time-indexed treatment representation that reconstructs administered therapy, estimates changing latent pharmacological and disease states, simulates bounded counterfactual dose regimens, propagates decision-relevant uncertainty, and presents traceable alternatives for clinician authorization. The contribution is architectural and interpretive rather than empirical: it specifies the components that must be distinguished, the relationships that may be modeled, the evidence required for updating and counterfactual evaluation, and the boundaries that prevent simulated optima from being presented as autonomous prescriptions. The argument proceeds from population dosing rules to individualized treatment simulation, then defines the construct and its data streams before addressing dynamic patient states, sequential updating, uncertainty, safety, oversight, and staged implementation.

From Population Dosing Rules to Individualized Treatment Simulation

Population dosing rules compress evidence into actions that can be applied before extensive individual observations exist. Fixed doses, weight bands, organ-function categories, genotype groups, and covariate equations are therefore not primitive alternatives to individualized dosing; they are prior decision structures. Their limitation is temporal and representational. They usually describe how dose should differ across categories at a particular decision point, not how treatment should be reconsidered when a patient moves between physiological states, receives an interrupted regimen, accumulates toxicity, or exhibits a response inconsistent with the assumed model. Model-informed precision dosing extends population inference by combining prior models with accumulating patient observations, allowing the estimated individual state and the associated dose forecast to change during treatment [6]. The shift is thus from choosing a rule for a patient category to maintaining a conditional simulation of treatment for a changing individual.

This shift cannot be achieved by adding an optimization routine to a population pharmacokinetic model. Individualized treatment simulation requires a computational environment capable of receiving verified treatment and observation data, applying the intended model version, displaying assumptions and uncertainty, and communicating outputs in a form that supports rather than displaces clinical reasoning. Existing precision-dosing software varies in model availability, data handling, transparency, usability, output presentation, and capacity for clinical integration [7]. These differences are scientifically important because the implemented system determines which dose history is reconstructed, which observations are accepted, how missing values are treated, whether model provenance is visible, and whether the clinician can inspect alternative explanations. Software is therefore part of the scientific measurement and decision chain, not a neutral container surrounding the pharmacometric model.

Therapeutic drug monitoring demonstrates the transition from prior dosing knowledge to within-patient updating. A measured drug concentration can be interpreted against a population model, the recorded dose history, sampling time, covariates, and residual-error assumptions to revise the likely individual exposure profile and inform a subsequent regimen [8]. The Dose Twin generalizes this logic without assuming that concentration data alone are sufficient. Concentrations may be unavailable, mistimed, contaminated by uncertain administration records, or weakly informative about delayed pharmacodynamic effects. Conversely, biomarkers, organ-function measurements, adverse events, disease indicators, and documented treatment interruptions may alter the interpretation of an apparently plausible concentration. Individualized simulation must therefore

assimilate heterogeneous observations according to their pharmacological meaning and reliability rather than treating every newly available data point as equivalent evidence.

A further limitation arises when the patient is represented through a single selected population model. Model selection may be fragile when the current individual differs from the development population or when several plausible models encode materially different covariate relationships. Model averaging and selection approaches can reduce dependence on one potentially mismatched model, but they do not eliminate uncertainty about structural suitability [9]. Contemporary model-informed precision dosing consequently depends on the coordinated use of bioanalysis, biomarkers, pharmacometric models, computational tools, and clinical workflows, with applicability remaining conditional on the drug, patient population, decision, and available evidence [10]. The Dose Twin extends this ecosystem by requiring that model plurality, longitudinal state change, adherence, competing explanations, and decision boundaries be represented explicitly rather than collapsed into a single predicted dose. **Table 1** organizes the evidence, constructs, and interpretation boundaries needed to develop from population dosing rules to individualized treatment simulation within the article's central argument.

Table 1. Evidence domains, core questions, scientific requirements, and interpretation boundaries for From population dosing rules to individualized treatment simulation

Construct or Evidence Domain	Core Question	Scientific Basis	Role In The Article	Failure or Overclaiming Risk	Boundary Statement
Population Dosing Prior	What information should guide treatment before informative individual observations are available?	Label dosing, population pharmacokinetics, exposure–response knowledge, covariate effects, and established treatment constraints	Supplies the initial probability model and feasible regimen space	Treating a population mean or subgroup rule as the patient's true state	A population prior initiates individual inference; it does not complete it
Covariate-Based Individualization	Which observed patient characteristics justify modification of the population starting regimen?	Qualified relationships involving organ function, size, maturation, genotype, co-medication, or disease characteristics	Provides interpretable initial stratification and model parameters	Assuming a baseline covariate remains stable or causally determines later change	Covariate adjustment is conditional on measurement quality, model applicability, and time
Therapeutic Drug Monitoring	How should measured exposure alter the estimated patient state and future regimen?	Dose and sampling history, bioanalysis, pharmacokinetic structure, residual error, and Bayesian forecasting	Connects observations obtained during therapy to revised exposure estimates	Interpreting a concentration without verifying dose timing, administration, or assay context	A concentration is evidence within a treatment history, not an isolated dosing instruction
Longitudinal Patient State	Which changes in physiology, disease, adherence, and response must modify the simulation?	Serial clinical observations, treatment events, biomarkers, mechanistic assumptions, and observation models	Converts episodic dose adjustment into time-indexed treatment simulation	Treating dynamic biological and behavioral processes as fixed parameters	Unobserved or weakly observed state changes must remain uncertain
Model Plurality And Applicability	Is one population or mechanistic model adequate for the current patient and decision?	External evaluation, model comparison, model averaging, applicability assessment, and structural sensitivity analysis	Makes model selection and mismatch visible within individual inference	Selecting a numerically favorable model without examining biological or contextual fit	No candidate model is automatically suitable outside its qualified context
Decision Objective	What treatment outcome is being optimized, constrained, or protected?	Defined exposure, response, toxicity, feasibility, and patient-centered treatment objectives	Specifies why counterfactual regimens are compared	Reducing treatment to one surrogate target or concealing trade-offs in an objective function	An optimized surrogate does not alone establish therapeutic benefit
Software and Workflow Layer	Can data, models, outputs, and responsibility be integrated transparently into care?	Data provenance, version control, interface design, interoperability, audit trails, and clinical role definition	Connects computational inference to an accountable decision process	Treating implementation defects as separate from scientific validity	A valid model can still be unsafe or unusable when implemented incorrectly
Individualized Treatment Simulation	Can feasible future regimens be compared under the patient's current estimated state?	Posterior patient-state estimation, mechanistic simulation, uncertainty propagation, and treatment constraints	Forms the decision-model function proposed for the Dose Twin	Presenting the highest-ranked simulation as an autonomous recommendation	Counterfactual comparison informs deliberation; it does not authorize treatment

Defining the Dose Twin and Its Required Data Streams

Systems pharmacology applications provide precedents for constructing individualized virtual counterparts that represent mechanistic heterogeneity and support dose–response or biomarker simulation. Such applications demonstrate that patient-linked parameterization and disease-mechanism modeling can organize virtual treatment trajectories, although their validity remains dependent on the represented population, available observations, model assumptions, and intended use [11, 12]. Building on this precedent, the Dose Twin is proposed as a narrower pharmaceutical construct: a patient-linked, time-indexed computational representation of treatment that joins the actual dose history to evolving physiology, disease, pharmacokinetics,

pharmacodynamics, adherence, response, and toxicity. Its defining feature is not visual resemblance to a patient or continuous data connectivity. It is the capacity to maintain a traceable and revisable treatment-state representation from which bounded future regimens can be simulated.

The first required data stream is treatment identity and history. Drug, formulation, route, prescribed regimen, administered doses, infusion changes, delays, interruptions, omissions, and relevant co-medications determine the exposure-generating input. These data must be time-stamped and accompanied by provenance because an uncertain dose history can be confounded with altered absorption, clearance, distribution, or adherence. The second stream represents dynamic physiology, including organ function, body size and composition where relevant, maturation or pregnancy state, inflammatory status, extracorporeal or other supportive interventions, and biochemical determinants of drug disposition. Physiology-based pharmacokinetic modeling illustrates why explicit information concerning organs, enzymes, transporters, drug properties, covariates, and qualified assumptions is necessary when physiological structure is used to explain or predict exposure [13]. The Dose Twin does not require maximal physiological detail; it requires a level of mechanistic resolution proportionate to the dosing question. A third stream represents disease and therapeutic response. Disease burden, target status, progression or remission indicators, pharmacodynamic biomarkers, clinical response, adverse effects, and competing clinical events must be separated conceptually even when they are statistically associated. A fourth stream contains observation metadata: assay method, sampling time, measurement error, missingness, data latency, and evidence provenance. A fifth stream records adherence or administration reliability, ideally through time-stamped direct administration records where available and otherwise through explicitly uncertain indirect evidence. Mechanistic and machine-learning components may contribute complementary functions—for example, a pharmacokinetic model may represent disposition while a learned component extracts features, estimates missing inputs, detects patterns, selects among models, or corrects a bounded residual process [14]. Their interfaces must remain visible, however. A learned estimate must not silently become a biological state, and improved prediction must not be interpreted automatically as mechanistic confirmation.

Operationally, the Dose Twin contains nine linked elements. First, patient identity and context define the indication, treatment objective, care setting, and qualified population. Second, treatment-history reconstruction supplies the exposure-driving input. Third and fourth, dynamic physiology and disease-state representations describe changing patient conditions. Fifth, a pharmacokinetic–pharmacodynamic, physiologically based, or quantitative systems pharmacology core maps treatment and patient states to predicted exposure, response, and toxicity at an appropriate level of resolution. Sixth, bounded learned components may support specific tasks without overriding mechanistic and uncertainty constraints. Seventh, a sequential updater assimilates new observations while preserving model and data provenance. Eighth, a counterfactual evaluator compares only clinically feasible candidate regimens under the current estimated state. Ninth, safety and uncertainty gates determine when the system must abstain, restrict its influence, or escalate the case for clinician review. The clinician authorization layer remains external to computational optimization and retains accountable ownership of the treatment decision. Together, these elements distinguish the proposed Dose Twin from a static dose calculator, an unqualified patient replica, an electronic-health-record dashboard, a single population model, or an autonomous prescribing system.

Dynamic Physiology, Disease Evolution, Adherence, and Response

A Dose Twin must represent physiology as a time-varying state rather than assuming that baseline covariates remain valid throughout treatment. Renal and hepatic function, inflammatory status, body composition, fluid balance, protein binding, enzyme activity, concomitant therapy, and supportive interventions may change at different rates and may be measured with unequal frequency. Time-varying clearance models demonstrate that drug disposition can evolve during therapy and cannot always be represented adequately by one fixed individual parameter [15]. Within the proposed construct, physiological observations therefore update a latent disposition state through an explicit observation model. The model must preserve the difference between a measured variable, such as a serum biomarker, and the pharmacological state inferred from that variable. A transient laboratory change should not automatically trigger a new dose unless its timing, measurement reliability, expected pharmacological relevance, and consistency with other observations support such an interpretation.

Disease evolution constitutes a separate but interacting process. Changes in disease burden may alter drug disposition, target availability, biomarker production, treatment tolerance, or the therapeutic objective itself. Conversely, treatment response may alter the disease processes that initially influenced pharmacokinetics. Analyses of time-varying monoclonal-antibody clearance illustrate how disease state and pharmacokinetic behavior can change together, making observed exposure–response associations difficult to interpret as simple one-directional effects [16]. The Dose Twin should therefore contain distinct disease and pharmacological state variables connected only through relationships justified for the drug and context of use. A declining drug concentration could reflect increased clearance, altered administration, changing distribution, assay error, or an intervention-induced change in disease burden. The architecture must retain these competing explanations until the available evidence can discriminate among them.

Adherence must be represented as an exposure-generating process, not merely as a baseline behavioral covariate. The prescribed regimen, the administered regimen, and the regimen actually received may diverge through missed doses, timing deviations, incomplete infusions, dose reductions, treatment interruptions, or unrecorded self-administration. Model-based analyses show that non-adherence can alter both apparent effectiveness and safety by changing cumulative exposure and the temporal pattern of treatment [17]. A Dose Twin must therefore reconstruct treatment history before attributing unexpected concentrations or responses to altered pharmacokinetics or pharmacodynamics. Direct administration records, pharmacy data,

electronic monitoring, clinician documentation, and patient reports may contribute different forms of adherence evidence, but none should be treated as infallible. When treatment history is uncertain, counterfactual simulations should include plausible administration scenarios rather than hiding uncertainty inside an estimated clearance or sensitivity parameter.

Therapeutic response and toxicity complete the longitudinal treatment state. Response can be immediate or delayed, reversible or cumulative, directly measurable or represented through imperfect biomarkers. Toxicity may modify future tolerability, alter organ function, necessitate interruption, or change the acceptable balance between efficacy and exposure. Therapeutic proteins further illustrate how disease burden, target-mediated disposition, immunogenicity, treatment history, and response may interact in patient-specific ways [18]. The proposed Dose Twin consequently treats exposure, response, disease, and toxicity as related but non-interchangeable processes. It does not assume that higher exposure is always beneficial, that biomarker improvement proves causal treatment effect, or that observed failure reflects pharmacological resistance. **Table 2** organizes the evidence, constructs, and interpretation boundaries needed to develop dynamic physiology, disease evolution, adherence, and response within the article's central argument.

Table 2. Components, relationships, uncertainties, and validation needs within Dynamic physiology, disease evolution, adherence, and response

Component or Layer	Inputs	Core Function	Expected Output	Uncertainty or Limitation	Validation Requirement
Dynamic Physiology State	Serial organ-function measures, body size and composition, inflammatory markers, co-medications, supportive interventions, and relevant life-stage variables	Represents time-varying determinants of absorption, distribution, metabolism, and elimination	A time-indexed distribution of disposition-relevant physiological states	Sparse measurements, delayed recording, biomarker surrogacy, and confounding among physiological processes	External evaluation of time-varying relationships, temporal calibration, and sensitivity to missing or delayed observations
Disease-State Trajectory	Disease burden, severity indicators, target expression, progression or remission measures, and clinical events	Represents changes in disease that may affect pharmacology, treatment objectives, or outcome interpretation	Estimated current disease state and plausible future trajectories	Disease measures may be indirect, treatment-responsive, and non-identifiable from pharmacological change	Longitudinal qualification against clinically meaningful disease observations and competing trajectory models
Adherence and Administration History	Prescriptions, medication-administration records, infusion logs, pharmacy data, electronic monitoring, clinician notes, and patient reports	Reconstructs the dose actually received and its timing	Probabilistic administered-dose history rather than a presumed prescribed regimen	Discordant sources, undocumented omissions, inaccurate timestamps, and no universal adherence gold standard	Prospective triangulation of adherence sources and stress testing under alternative administration scenarios
Pharmacokinetic State	Treatment history, physiological state, concentrations, drug properties, and relevant covariates	Maps administered treatment and patient state to exposure	Exposure distributions and time-varying individual pharmacokinetic parameters	Structural mismatch, assay error, model non-identifiability, and extrapolation	Prospective predictive checks, calibration, plausibility review, and qualification in the intended population
Pharmacodynamic and Therapeutic-Response State	Exposure history, biomarkers, clinical outcomes, disease state, and response delays	Represents treatment effects and the temporal relationship between exposure and response	Predicted response trajectories with explicit delay and uncertainty	Biomarkers may not be validated surrogates; association may be mistaken for mechanism	Evaluation against independent response data and comparison with simpler exposure–response representations
Toxicity and Tolerability State	Adverse events, laboratory abnormalities, patient-reported symptoms, organ injury markers, and treatment modifications	Represents current and accumulating treatment burden	Predicted toxicity or tolerability distributions and safety constraints	Under-reporting, competing causes, delayed toxicity, and subjective grading	Prospective safety calibration, adjudication procedures, and clinically defined escalation thresholds
Interaction and Confounding Layer	Joint histories of physiology, disease, adherence, exposure, response, and toxicity	Preserves multiple explanations for observed treatment change	Ranked or weighted competing state explanations	Correlated trajectories may be statistically indistinguishable and causality may remain unresolved	Identifiability analysis, negative-control scenarios, perturbation studies, and explicit reporting of unresolved alternatives

Updating, Recalibration, and Counterfactual Dose Evaluation

Updating within a Dose Twin occurs at two different levels that must not be conflated. Within-patient updating revises the estimated state of one treated individual when new evidence becomes available. Across-patient recalibration modifies the population prior, model parameters, structural assumptions, or software behavior using information accumulated from multiple patients. Continued-learning methods provide a basis for controlled revision of population pharmacometric models when clinical populations diverge from the cohorts used for original development [19]. Such recalibration may improve future inference, but it changes the model on which all subsequent patient-level decisions depend. It must therefore occur through a version-controlled and separately governed process rather than through silent adaptation during routine treatment.

Clinical data streams may also support iterative refinement of local models, provided that each update is monitored for data quality, temporal drift, parameter instability, subgroup effects, and degradation outside the setting in which learning occurred

[20]. The Dose Twin should maintain a complete distinction among the immutable record of past observations, the current patient-state estimate, the version of the population or mechanistic model used for that estimate, and any proposed recalibrated model. Reanalysis under a new model should create a traceable new state estimate rather than overwrite the previous decision record. This separation is essential for auditability: a later model revision must not make it appear that an earlier clinician had access to information or model behavior that did not exist when the treatment decision was made.

Sequential inference provides the mathematical bridge between observations and counterfactual treatment evaluation. Bayesian data assimilation can revise distributions over latent physiological, pharmacokinetic, disease, adherence, and response states as new measurements arrive, while policy-learning approaches can explore adaptive treatment choices in simulated environments [21]. The proposed Dose Twin does not require reinforcement learning, nor does it permit an unrestricted algorithm to discover dosing behavior directly in clinical care. Policy-learning methods are included as one possible computational approach for comparing sequential actions under uncertainty. Their use is conditional on a qualified simulator, constrained action space, explicit objective function, adequate representation of delayed outcomes, and protection against extrapolative or unsafe actions.

Counterfactual dose evaluation asks what could plausibly occur under a limited set of alternative future regimens, given the estimated current state and its uncertainty. Candidate regimens should be clinically feasible and should preserve constraints involving formulation, route, dose increments, treatment schedules, maximum exposure, toxicity, monitoring capacity, and clinician-defined objectives. Adaptive-dosing research shows that candidate policies can be compared prospectively in silico, but simulated advantage does not establish clinical utility [22]. A counterfactual output should therefore report distributions of exposure, response, and toxicity; the assumptions supporting each scenario; and sensitivity to alternative patient-state explanations. It should not return one decontextualized "optimal dose." **Table 3** organizes the evidence, constructs, and interpretation boundaries needed to develop updating, recalibration, and counterfactual dose evaluation within the article's central argument.

Table 3. Components, relationships, uncertainties, and validation needs within Updating, recalibration, and counterfactual dose evaluation

Evaluation Dimension	What Must be Tested	Suitable Evidence or Assessment	Failure Signal	Interpretive Limitation	Decision Relevance
Within-Patient State Updating	Whether new observations revise the correct latent state without destabilizing unrelated components	State-recovery simulations, posterior predictive checks, temporal holdout evaluation, and sensitivity to observation timing	Implausible parameter jumps, overreaction to one measurement, or failure to respond to informative evidence	A narrow posterior may reflect restrictive assumptions rather than informative data	Determines whether current simulations represent the patient's evolving treatment state
Population-Model Recalibration	Whether accumulated clinical data justify changes to population priors, covariate effects, or structural assumptions	Versioned recalibration studies, external validation, drift analysis, subgroup assessment, and model-comparison procedures	Improved local fit accompanied by poorer transportability or subgroup calibration	Updating a model does not establish that the revised structure is biologically correct	Determines whether later patients should be analyzed with a modified prior model
Treatment-History Reconstruction	Whether prescribed, administered, interrupted, and missed doses are distinguished correctly	Reconciliation across medication records, infusion systems, pharmacy data, electronic monitoring, and sensitivity scenarios	Unexplained concentrations are absorbed into clearance or response parameters	Complete reconstruction may remain impossible outside directly observed administration	Determines whether exposure simulations begin from a credible input history
Sequential Data Assimilation	Whether evidence is incorporated in the correct temporal order with appropriate error models	Simulation-based calibration, delayed-data tests, missingness experiments, and comparison with batch estimation	Future information leaks into past state estimates or stale data dominate current inference	Accurate state estimation remains conditional on the observation model	Determines when the twin should be updated and which observations should influence the decision
Counterfactual Regimen Generation	Whether candidate regimens are clinically feasible, pharmacologically interpretable, and within qualified support	Constraint checking, expert review, overlap diagnostics, physiological plausibility testing, and scenario libraries	The model proposes unsupported doses, schedules, or treatment combinations	Feasibility does not imply therapeutic desirability	Defines which alternatives may be presented for clinician comparison
Counterfactual Outcome Simulation	Whether alternative regimens generate credible exposure, response, and toxicity distributions	Posterior predictive simulation, structural sensitivity analysis, external challenge cases, and comparison with simpler models	Rankings reverse under minor assumptions or depend on poorly identified states	Simulated outcomes are not observed individual causal effects	Supports bounded comparison without claiming the future outcome is known

Objective-Function Adequacy	Whether the modeled objective reflects the clinical treatment problem and relevant trade-offs	Multistakeholder specification, sensitivity to alternative utility functions, and explicit constraint review	A surrogate target dominates safety, feasibility, or patient priorities	No mathematical objective fully captures clinical judgment	Determines what "better" means within the simulated comparison
Robustness and Abstention	Whether the system detects extrapolation, model disagreement, poor data quality, and decision-relevant uncertainty	Out-of-distribution tests, model-ensemble disagreement, uncertainty calibration, and prespecified abstention scenarios	High-confidence output persists despite missing, contradictory, or unsupported evidence	Abstention thresholds are context dependent and require prospective assessment	Determines when no computationally favored regimen should be presented

Safety Constraints, Uncertainty, and Clinician Oversight

The evidentiary burden for a Dose Twin should increase with the influence that its outputs are permitted to exert on treatment. Risk-informed credibility frameworks connect context of use, model consequence, verification, validation, applicability, and uncertainty rather than treating model performance as an isolated property [23-25]. This principle applies to mechanistic, empirical, and hybrid components alike. A model used for exploratory retrospective simulation may require a different level of evidence from one that orders candidate regimens for review during an active treatment episode. Credibility is therefore not a universal certificate attached permanently to a model. It is a context-dependent argument showing that the implemented model, data, assumptions, and uncertainty are adequate for a specified question and a bounded degree of decision influence.

Technical safety begins before any clinical output is generated. The selected pharmacometric or mechanistic model must match the drug, formulation, patient population, sampling context, and intended decision; the software must implement that model correctly; units, covariates, dose timing, and parameter transformations must be verified; and model versions must be traceable. Precision-dosing tutorials emphasize that an apparently functional dosing engine can remain unreliable when an unsuitable model is selected or a correct model is implemented incorrectly [26]. The Dose Twin should therefore include automated input validation, reproducible test cases, solver diagnostics, model-signature checks, and warnings for unsupported population or regimen combinations. These controls verify implementation integrity but cannot establish biological validity or clinical benefit.

Decision-relevant uncertainty must be decomposed rather than compressed into a single confidence score. Relevant sources include measurement error, missing treatment history, adherence ambiguity, parameter uncertainty, between-model disagreement, uncertain disease state, structural misspecification, future physiological change, and unsupported counterfactual extrapolation. Safety constraints should combine hard exclusions with softer escalation rules. Hard constraints may prohibit physiologically impossible states, incompatible formulations, or regimens outside an authorized action space. Soft constraints may trigger additional sampling, specialist consultation, conservative alternatives, or abstention when uncertainty is too large relative to the therapeutic margin. Uncertainty reporting is useful only when it changes the permissible interpretation or influence of the output; a visually narrow interval should never substitute for evidence that all material uncertainty sources have been represented.

Clinician oversight is not a decorative final step appended to an otherwise autonomous system. The clinician must be able to inspect the treatment history used, the current estimated state, the model and software versions, alternative explanations, excluded scenarios, assumptions, uncertainty, and the rationale for comparing particular regimens. The interface should distinguish observed facts from model-derived states and distinguish safety rules from probabilistic preferences. Authorization should remain attributable to a named clinical role, with escalation pathways for disagreement, missing information, unusual patients, or outputs that conflict with bedside evidence. Human oversight does not validate a weak model, however, and poorly designed explanations can produce automation bias rather than informed scrutiny. The proposed governance objective is accountable collaboration: computational simulation organizes evidence and alternatives, while clinicians retain responsibility for treatment interpretation and action.

Implementation Pathways in Model-Informed Precision Dosing

Implementation begins with reliable integration of clinical data, model execution, decision support, provenance, and workflow ownership. Electronic-health-record-integrated precision-dosing tools show how patient characteristics, treatment histories, laboratory observations, and pharmacometric models may be connected within clinical decision support [27]. A Dose Twin would require a broader longitudinal interface because it must preserve actual administration, adherence evidence, physiology, disease, response, toxicity, and model-state history. Data ingestion should therefore be selective rather than indiscriminate. Each input needs a defined semantic meaning, acceptable latency, source hierarchy, unit convention, missingness policy, and validation rule. The system must also write back a transparent decision record without converting uncertain model states into unqualified clinical facts.

Interoperability, automation, monitoring, and multidisciplinary governance are central requirements for digital-health-enabled precision dosing [28]. Implementation responsibility would extend across clinical pharmacology, pharmacy, treating specialties, nursing, laboratory medicine, bioinformatics, information technology, software engineering, data governance, quality assurance, and institutional oversight. Automation may reduce transcription and calculation burdens, but it also increases the speed and scale at which erroneous data or model behavior can propagate. Interfaces should therefore expose

model applicability, data freshness, unresolved discrepancies, and version changes. Institutions also need procedures for downtime, manual override, model retirement, post-update testing, and review of unexpected recommendations or systematic disagreement between simulated and observed trajectories.

Bedside adoption is constrained not only by technical performance but also by evidence, resources, usability, trust, workflow compatibility, role clarity, and access to trained personnel [29]. The proposed implementation pathway should therefore progress through bounded stages. An initial research environment can evaluate data availability and retrospective state reconstruction without influencing care. Silent prospective evaluation can then test whether the Dose Twin operates reliably on real-time data while its outputs remain hidden from decision-makers. A consultative pilot may permit specialists to review outputs under predefined cases and stop conditions. Comparative prospective evaluation would be required before claiming added clinical value. Governed routine use, where justified, would remain restricted to the qualified drug, population, data environment, and decision type rather than being treated as a universal platform capability.

Routine model-informed precision dosing requires alignment among model development, clinical expertise, software, institutions, and prospective evaluation [30]. For the Dose Twin, progression between implementation stages should depend on prespecified evidence thresholds rather than technological availability or apparent simulation sophistication. Required evidence includes reliable treatment-history reconstruction, state-estimation calibration, model and software verification, acceptable workflow burden, subgroup assessment, uncertainty-triggered abstention, and demonstration that outputs improve decision quality or patient-relevant outcomes for the intended context. **Table 4** organizes the evidence, constructs, and interpretation boundaries needed to develop implementation pathways in model-informed precision dosing within the article's central argument. **Figure 1** presents the individualized dose twin observatory, showing how the article's key components and boundaries are connected within implementation pathways in model-informed precision dosing.

Table 4. Evaluation, implementation, and research priorities arising from The Dose Twin Simulating Individualized Treatment under Changing Physiology, Disease Progression, Adherence

Research or implementation priority	Unresolved gap	Required methodological work	Evidence needed	Responsible actors	Expected contribution
Common Dose Twin ontology	Patient state, treatment history, adherence, model state, and decision outputs may be represented inconsistently	Define interoperable concepts, temporal relationships, provenance fields, and boundaries between observed and inferred information	Cross-disciplinary consensus, use-case testing, and semantic interoperability studies	Clinical pharmacologists, pharmacometricians, clinicians, pharmacists, informaticians, and standards organizations	Enables comparable implementations and prevents concept drift
Treatment-history reconstruction	Prescribed treatment is often substituted for treatment actually received	Develop source-reconciliation methods and probabilistic administration histories	Prospective comparison of administration records, devices, pharmacy data, and patient reports	Health systems, pharmacists, nurses, patients, data scientists, and vendors	Reduces exposure-estimation error and false attribution to altered pharmacology
Dynamic-state identifiability	Physiology, disease, adherence, response, and toxicity may generate similar observable patterns	Develop joint state-space models, identifiability diagnostics, informative sampling strategies, and competing-explanation analysis	Simulation studies, longitudinal cohorts, perturbation data, and external challenge cases	Pharmacometricians, systems pharmacologists, statisticians, laboratory scientists, and clinical teams	Clarifies which latent changes can be inferred and when the system should abstain
Counterfactual evaluation standards	Simulated regimen rankings may depend on unsupported assumptions or objectives	Specify intervention support, feasible action spaces, objective functions, causal assumptions, and robustness tests	Prospective simulation protocols, sensitivity analyses, simpler-model comparisons, and adaptive trial designs	Methodologists, clinicians, trialists, ethicists, and regulators	Establishes when simulated alternatives are interpretable without equating them with benefit
Credibility and software assurance	Model validity and implementation integrity are often assessed separately	Integrate model qualification, code verification, input validation, version control, and regression testing	Independent audits, manufactured cases, external validation, and lifecycle monitoring	Model developers, software engineers, quality units, institutions, and oversight bodies	Produces traceable context-of-use credibility arguments
Uncertainty and abstention design	Confidence outputs may omit material data, model, and scenario uncertainty	Develop decomposed uncertainty reporting and context-specific escalation or abstention thresholds	Calibration studies, out-of-distribution tests, clinician interpretation studies, and safety stress tests	Statisticians, pharmacometricians, clinicians, human-factors specialists, and governance committees	Limits decision influence when evidence is insufficient
Workflow and human-factors evaluation	Technically valid outputs may be unusable, misunderstood, or overtrusted	Conduct task analysis, interface testing, alert-burden evaluation, explanation studies, and role-definition work	Silent trials, usability studies, simulation exercises, and qualitative implementation research	Clinicians, pharmacists, nurses, human-factors researchers, informaticians, and patients	Supports accountable use without creating automation bias

Research or implementation priority	Unresolved gap	Required methodological work	Evidence needed	Responsible actors	Expected contribution
Prospective clinical utility	Predictive adequacy does not establish improved treatment decisions or outcomes	Design context-specific comparative and pragmatic prospective evaluations	Decision-quality measures, patient-relevant outcomes, safety events, workflow burden, and resource use	Clinical investigators, health systems, patients, funders, and oversight bodies	Determines whether the Dose Twin adds value beyond simpler alternatives
Equity and transportability	Data quality, model calibration, and workflow access may differ across sites and populations	Evaluate subgroup performance, missingness, access barriers, and cross-site transfer	Representative multisite data, subgroup calibration, fairness analyses, and implementation comparisons	Health systems, community partners, patients, methodologists, and policymakers	Defines where the construct can be used fairly and where local redevelopment is required
Staged governance and lifecycle management	Model updates, software changes, and expanding use may outpace validation	Establish stage gates, stop conditions, update review, incident reporting, and retirement procedures	Prospective monitoring, audit logs, change-impact assessments, and post-implementation review	Institutions, clinical governance groups, quality teams, developers, and regulators	Prevents unreviewed expansion from conceptual or pilot use to operational authority

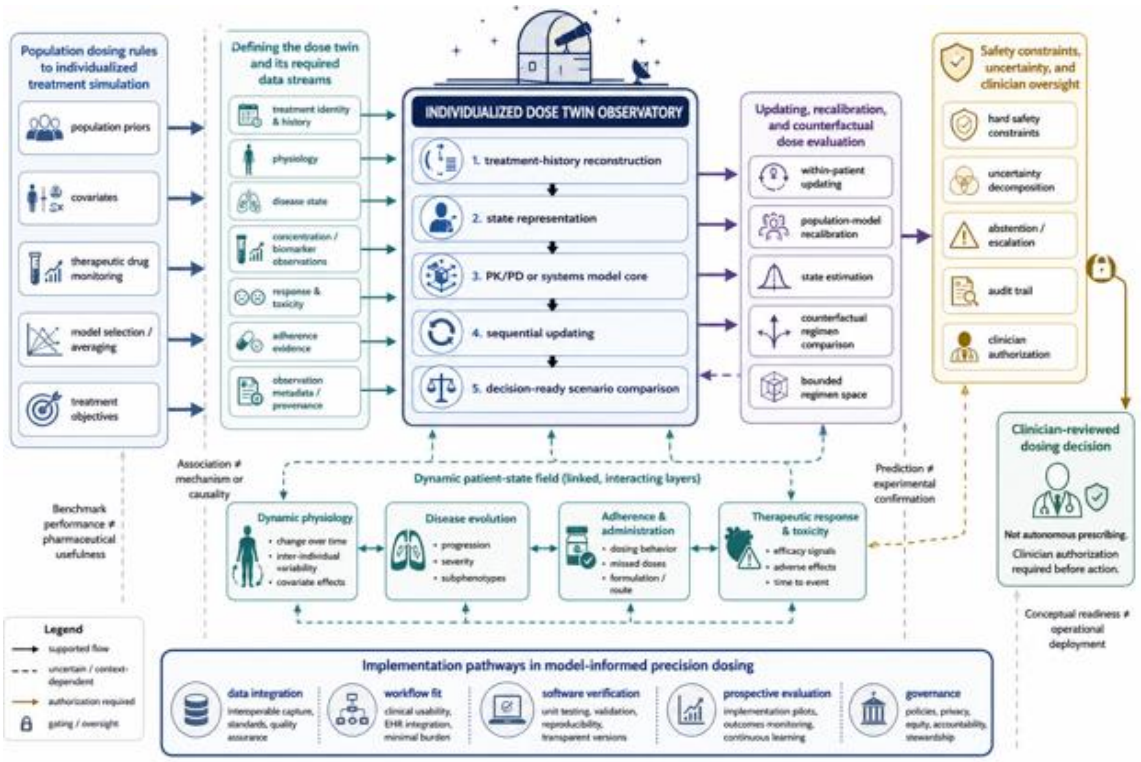

Figure 1. Individualized Dose Twin Observatory

The figure is an original conceptual synthesis that organizes the article's central contribution across population dosing rules to individualized treatment simulation, defining the dose twin and its required data streams, dynamic physiology, disease evolution, adherence, and response, dynamic physiology, disease evolution, updating, recalibration, and counterfactual dose evaluation. Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, regulatory determination, clinical recommendation, or deployment-ready system.

Alt-Text Description

A conceptual scientific diagram of the individualized dose twin observatory showing population dosing rules to individualized treatment simulation, defining the dose twin and its required data streams, dynamic physiology, disease evolution, adherence, and response, dynamic physiology, disease evolution, updating, recalibration, and counterfactual dose evaluation, counterfactual dose evaluation, safety constraints, uncertainty, and clinician oversight, their relationships, uncertainties, and decision boundaries.

Q1-Journal-Publication-Ready Figure Prompt

Create a publication-grade conceptual figure titled “Individualized Dose Twin Observatory.” Build the composition around one unmistakable central visual anchor: The Individualized Dose Twin Observatory. Use a landscape, flat, vector-first, scientifically restrained design suitable for a Q1 journal. The figure must express the logic of Implementation pathways in model-informed precision dosing and include clearly differentiated elements for population dosing rules to individualized treatment simulation, defining the dose twin and its required data streams, dynamic physiology, disease evolution, adherence, and response, dynamic physiology, disease evolution, updating, recalibration, and counterfactual dose evaluation, counterfactual dose evaluation, safety constraints, uncertainty, and clinician oversight, safety constraints, clinician oversight, implementation pathways in model-informed precision dosing. Show directional relations only where the manuscript supports them; use dashed connectors for uncertain, context-dependent, or proposed relations. Add visible boundary markers separating benchmark performance from pharmaceutical usefulness, association from mechanism or causality, prediction from experimental confirmation, and conceptual readiness from operational deployment where relevant. Use a clean white background, precise alignment, professional sans-serif typography, thin uniform connectors, generous white space, and a restrained color-blind-safe palette. Do not make the figure resemble a generic flowchart, dashboard, pipeline of rectangular boxes, circular wheel, marketing infographic, or decorative molecular collage. Do not use journal, publisher, corporate, software, database, regulatory, or product logos; copied scientific figures; interface screenshots; stock laboratory photography; proprietary molecular renderings; or invented numerical outputs. Ensure all labels remain legible at journal-column width and the design is suitable for high-resolution export.

Limitations and Boundary Conditions

The Dose Twin is a proposed conceptual and methodological construct, not a validated model class. Its components can be instantiated through population pharmacokinetics, pharmacokinetic–pharmacodynamic models, physiologically based pharmacokinetics, quantitative systems pharmacology, statistical state-space models, or bounded hybrid approaches, but the selection of a sophisticated architecture does not establish that its internal states correspond to true biological mechanisms. Mechanistic and learned components may improve representation or prediction while remaining non-identifiable, context dependent, or sensitive to unobserved assumptions [14]. The construct cannot guarantee that sufficient data exist to estimate every proposed state, and it does not require maximal model complexity. In many settings, a simpler covariate rule, therapeutic drug-monitoring model, or clinician-led protocol may be more transparent, more robust, and more appropriate.

The available evidence is also uneven across drugs, populations, disease settings, treatment durations, and observation patterns. Dynamic clearance, disease progression, adherence, response, and toxicity can be correlated without being distinguishable from the observations available during routine care. Disease-linked pharmacokinetic change, for example, may represent multiple biological and statistical relationships rather than one proven causal pathway [16]. Adherence evidence may be sparse, delayed, or contradictory, and longitudinal measurements may be collected because the patient is deteriorating, thereby creating informative observation and missingness processes. The Dose Twin cannot recover an accurate individualized trajectory from unsuitable data merely by applying Bayesian updating, machine learning, or mechanistic simulation.

The construct must remain bounded by its qualified context of use. It cannot establish individual causal effects from observational updating, predict unobserved treatment responses with certainty, justify regimens outside the supported intervention space, or replace experimental and prospective clinical evaluation. It also cannot establish regulatory acceptability, institutional readiness, cost-effectiveness, equitable access, or improved outcomes without corresponding evidence. Credibility requirements should scale with the consequence of model influence, and use should be restricted or suspended when data quality, applicability, uncertainty, or implementation integrity is inadequate [23]. The most serious misuse would be to present the Dose Twin as an autonomous prescriber or to interpret a simulated optimum as a treatment recommendation independent of clinician judgment and patient context.

Research Agenda and Implementation Priorities

The first research priority is to determine whether the proposed state architecture is identifiable and useful for specific drug–disease–decision contexts. Studies should compare alternative representations of physiology, disease, adherence, exposure, response, and toxicity; determine which observations are sufficiently informative; and identify when competing explanations cannot be resolved. Evaluation should include code verification, model qualification, state-recovery simulation, temporal calibration, missing-data stress tests, model disagreement, and comparison with simpler alternatives. Counterfactual validity requires special attention because plausible simulated trajectories do not prove individual causal effects. Risk-informed credibility principles can organize these studies by connecting each technical assessment to the intended decision and consequence of model failure [24].

The second priority is prospective evaluation in workflows that preserve staged decision influence. Silent prospective studies should first test data latency, treatment-history reconstruction, update stability, abstention behavior, and discrepancies between modeled and clinical interpretations. Controlled consultative studies can then assess usability, clinician understanding, decision quality, workload, and automation bias. Comparative studies should examine whether the Dose Twin adds value over established dosing rules, standard therapeutic drug monitoring, specialist pharmacometric consultation, or simpler Bayesian forecasting. Current implementation evidence indicates that resources, trust, usability, workflow integration, and stakeholder roles are likely to influence adoption as strongly as algorithmic performance [29]. Patient-relevant benefit, safety, equity, and operational burden must therefore remain central evaluation outcomes.

The third priority is creation of durable infrastructure and reporting standards. Required elements include interoperable temporal data models, provenance-aware treatment histories, validated model repositories, version-controlled software, transparent uncertainty representations, audit logs, incident review, and governance for recalibration or retirement. Research reports should distinguish within-patient updating from population-model learning, observed variables from inferred states, prediction from mechanism, and simulated policy value from clinical utility. Multidisciplinary implementation frameworks should define responsibility for model selection, software verification, clinical authorization, monitoring, and response to failure [28]. Progress should be staged: conceptual coherence should precede technical verification, technical verification should precede prospective decision support, and routine influence should remain conditional on context-specific evidence rather than assumed from success in another drug, institution, or population.

Conclusion

The Dose Twin is proposed as a patient-linked, time-indexed computational pharmacology decision model for reconstructing treatment actually received, estimating changing physiological, disease, pharmacological, adherence, response, and toxicity states, and comparing bounded counterfactual dose regimens under explicit uncertainty and safety constraints. Its original contribution is to organize individualized dosing as repeated inference about an evolving treatment system rather than optimization of a single predicted exposure or performance measure. The construct separates within-patient updating from governed population-model recalibration, distinguishes prescribed treatment from administered treatment, and places clinician authorization outside computational optimization. It also defines boundaries that are essential to responsible interpretation: association is not mechanism, prediction is not experimental confirmation, calibrated uncertainty is not certainty, simulated policy advantage is not clinical utility, and conceptual coherence is not deployment readiness. The Dose Twin may provide a useful framework for designing and evaluating future model-informed precision-dosing systems, but its value remains conditional on drug-specific model qualification, suitable longitudinal data, technical verification, prospective clinical assessment, usable workflows, equitable implementation, and accountable governance.

Acknowledgments: None

Conflict of interest: None

Financial support: None

Ethics statement: None

References

1. Darwich AS, Ogungbenro K, Vinks AA, Powell JR, Reny JL, Marsousi N, et al. Why has model-informed precision dosing not yet become common clinical reality? Lessons from the past and a roadmap for the future. Clin Pharmacol Ther. 2017;101(5):646-56. doi:10.1002/cpt.659
2. Neely M. Scalpels not hammers: The way forward for precision drug prescription. Clin Pharmacol Ther. 2017;101(3):368-72. doi:10.1002/cpt.593
3. Darwich AS, Polasek TM, Aronson JK, Ogungbenro K, Wright DFB, Achour B, et al. Model-informed precision dosing: Background, requirements, validation, implementation, and forward trajectory of individualizing drug therapy. Annu Rev Pharmacol Toxicol. 2021;61:225-45. doi:10.1146/annurev-pharmtox-033020-113257
4. Wright DFB, Martin JH, Cremers S. Spotlight commentary: Model-informed precision dosing must demonstrate improved patient outcomes. Br J Clin Pharmacol. 2019;85(10):2238-40. doi:10.1111/bcp.14050
5. Corral-Acero J, Margara F, Marciniak M, Rodero C, Loncaric F, Feng Y, et al. The 'Digital Twin' to enable the vision of precision cardiology. Eur Heart J. 2020;41(48):4556-64. doi:10.1093/eurheartj/ehaa159
6. Keizer RJ, ter Heine R, Frymoyer A, Lesko LJ, Mangat R, Goswami S. Model-informed precision dosing at the bedside: Scientific challenges and opportunities. CPT Pharmacometrics Syst Pharmacol. 2018;7(12):785-7. doi:10.1002/psp4.12353
7. 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? Front Pharmacol. 2020;11:620. doi:10.3389/fphar.2020.00620
8. Wicha SG, Märtson AG, Nielsen EI, Koch BCP, Friberg LE, Alffenaar JWC, et al. From therapeutic drug monitoring to model-informed precision dosing for antibiotics. Clin Pharmacol Ther. 2021;109(4):928-41. doi:10.1002/cpt.2202
9. Uster DW, Stocker SL, Carland JE, Brett J, Marriott DJE, Day RO, et al. A model averaging/selection approach improves the predictive performance of model-informed precision dosing: Vancomycin as a case study. Clin Pharmacol Ther. 2021;109(1):175-83. doi:10.1002/cpt.2065
10. 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. Adv Drug Deliv Rev. 2024;215:115421. doi:10.1016/j.addr.2024.115421

11. Susilo ME, Li CC, Gadkar K, Hernandez G, Huw LY, Jin JY, et al. Systems-based digital twins to help characterize clinical dose-response and propose predictive biomarkers in a Phase I study of bispecific antibody, mosunetuzumab, in NHL. Clin Transl Sci. 2023;16(7):1134-48. doi:10.1111/cts.13501
12. Kaddi C, Tao M, Bergeler S, George K, Geerts H, van der Graaf PH, et al. Quantitative systems pharmacology-based digital twins approach supplements clinical trial data for enzyme replacement therapies in Pompe disease. Clin Pharmacol Ther. 2025;117(2):579-88. doi:10.1002/cpt.3498
13. Lin W, Chen Y, Unadkat JD, Zhang X, Wu D, Heimbach T. Applications, challenges, and outlook for physiologically based pharmacokinetic modeling and simulation: A regulatory, industrial and academic perspective. Pharm Res. 2022;39(8):1701-31. doi:10.1007/s11095-022-03274-2
14. Aghamiri SS, Amin R, Helikar T. Recent applications of quantitative systems pharmacology and machine learning models across diseases. J Pharmacokinet Pharmacodyn. 2022;49(1):19-37. doi:10.1007/s10928-021-09790-9
15. Chen J, Houk B, Pithavala YK, Ruiz-Garcia A. Population pharmacokinetic model with time-varying clearance for lorlatinib using pooled data from patients with non-small cell lung cancer and healthy participants. CPT Pharmacometrics Syst Pharmacol. 2021;10(2):148-60. doi:10.1002/psp4.12585
16. Wilkins JJ, Brockhaus B, Dai H, Vugmeyster Y, White JT, Brar S, et al. Time-varying clearance and impact of disease state on the pharmacokinetics of avelumab in Merkel cell carcinoma and urothelial carcinoma. CPT Pharmacometrics Syst Pharmacol. 2019;8(6):415-27. doi:10.1002/psp4.12406
17. Hill-McManus D, Soto E, Marshall S, Lane S, Hughes D. Impact of non-adherence on the safety and efficacy of uric acid-lowering therapies in the treatment of gout. Br J Clin Pharmacol. 2018;84(1):142-52. doi:10.1111/bcp.13427
18. Strik AS, Wang YMC, Ruff LE, Yashar W, Messmer BT, Mould DR. Individualized dosing of therapeutic monoclonal antibodies—a changing treatment paradigm? AAPS J. 2018;20(6):99. doi:10.1208/s12248-018-0257-y
19. Maier C, de Wiljes J, Hartung N, Kloft C, Huisinga W. A continued learning approach for model-informed precision dosing: Updating models in clinical practice. CPT Pharmacometrics Syst Pharmacol. 2022;11(2):185-98. doi:10.1002/psp4.12745
20. Hughes JH, Tong DMH, Lucas SS, Faldasz JD, Goswami S, Keizer RJ. Continuous learning in model-informed precision dosing: A case study in pediatric dosing of vancomycin. Clin Pharmacol Ther. 2021;109(1):233-42. doi:10.1002/cpt.2088
21. Maier C, Hartung N, Kloft C, Huisinga W, de Wiljes J. Reinforcement learning and Bayesian data assimilation for model-informed precision dosing in oncology. CPT Pharmacometrics Syst Pharmacol. 2021;10(3):241-54. doi:10.1002/psp4.12588
22. Tosca EM, De Carlo A, Ronchi D, Magni P. Model-informed reinforcement learning for enabling precision dosing via adaptive dosing. Clin Pharmacol Ther. 2024;116(3):619-36. doi:10.1002/cpt.3356
23. Kuemmel C, Yang Y, Zhang X, Florian J, Zhu H, Tegenge M, et al. Consideration of a credibility assessment framework in model-informed drug development: Potential application to physiologically based pharmacokinetic modeling and simulation. CPT Pharmacometrics Syst Pharmacol. 2020;9(1):21-8. doi:10.1002/psp4.12479
24. Musuamba FT, Skottheim Rusten I, Lesage R, Russo G, Bursi R, Emili L, et al. Scientific and regulatory evaluation of mechanistic in silico drug and disease models in drug development: Building model credibility. CPT Pharmacometrics Syst Pharmacol. 2021;10(8):804-25. doi:10.1002/psp4.12669
25. Skottheim Rusten I, Musuamba FT. Scientific and regulatory evaluation of empirical pharmacometric models: An application of the risk-informed credibility assessment framework. CPT Pharmacometrics Syst Pharmacol. 2021;10(11):1281-96. doi:10.1002/psp4.12708
26. Taylor ZL, Poweleit EA, Paice K, Somers KM, Pavia K, Vinks AA, et al. Tutorial on model selection and validation of model input into precision dosing software for model-informed precision dosing. CPT Pharmacometrics Syst Pharmacol. 2023;12(12):1827-45. doi:10.1002/psp4.13056
27. Vinks AA, Peck RW, Neely M, Mould DR. Development and implementation of electronic health record-integrated model-informed clinical decision support tools for the precision dosing of drugs. Clin Pharmacol Ther. 2020;107(1):129-35. doi:10.1002/cpt.1679
28. Kluwe F, Michelet R, Mueller-Schoell A, Maier C, Klopp-Schulze L, van Dyk M, et al. Perspectives on model-informed precision dosing in the digital health era: Challenges, opportunities, and recommendations. Clin Pharmacol Ther. 2021;109(1):29-36. doi:10.1002/cpt.2049
29. Dibbets AC, Koldeweij C, Osinga EP, Scheepers HCJ, de Wildt SN. Barriers and facilitators for bringing model-informed precision dosing to the patient's bedside: A systematic review. Clin Pharmacol Ther. 2025;117(3):633-45. doi:10.1002/cpt.3510
30. Polasek TM, Rostami-Hodjegan A, Yim DS, Jamei M, Lee H, Kimko H, et al. What does it take to make model-informed precision dosing common practice? Report from the 1st Asian Symposium on Precision Dosing. AAPS J. 2019;21(2):17. doi:10.1208/s12248-018-0286-6