Available online www.ijpras.com

International Journal of Pharmaceutical Research & Allied Sciences, 2025, 14(4):102-112
https://doi.org/10.51847/tWB6hEb1rZ

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Population Pharmacokinetics and Pharmacodynamics under Treatment Evolution, Time-Varying Covariates, Informative Sampling, and Changing Disease State during Clinical Care

Sara Ben Youssef[1]*, Amal Trabelsi[1], Karim Boudiaf[2], Nabil Jebali[3]

[1] Department of Population PK/PD and Treatment Evolution, Faculty of Pharmacy, University of Tunis, Tunis, Tunisia.
[2] Department of Time-Varying Covariates and Sampling, Faculty of Pharmacy, University of Sousse, Sousse, Tunisia.
[3] Department of Changing Disease State Modeling, Faculty of Pharmacy, University of Sfax, Sfax, Tunisia.

***Email:** sara.benyoussef@inat.ucar.tn*

ABSTRACT

Population pharmacokinetic and pharmacodynamic models commonly support inference by relating dose, exposure, response, and patient characteristics, yet their application during clinical care can become problematic when treatment, physiology, disease severity, observation intensity, and clinical decisions change together. This article develops an original longitudinal modelling framework that treats clinical PK/PD as a coupled system of biological states, treatment events, observation processes, and model-mediated decisions. The framework distinguishes latent pharmacokinetic and pharmacodynamic states from measured covariates, recorded doses, laboratory observations, and clinician actions. Treatment is represented as a timestamped history rather than a fixed regimen, while physiological variables and disease markers are treated according to their temporal availability and their possible roles as predictors, mediators, treatment consequences, or imperfect measurements of latent disease. Missing data are positioned within the observation process because absence may arise from visit timing, test-ordering decisions, documentation failure, assay limitations, or selective clinical attention. Dynamic prediction is defined as sequential updating from accumulating patient information, whereas population-model relearning is treated as a separate, governed activity. Evaluation therefore requires more than conventional goodness-of-fit assessment: calibration, uncertainty coverage, temporal validation, external transportability, observation-process sensitivity, update stability, and prospective workflow evaluation must be considered. The original contribution is an integrated state–event–observation architecture connecting treatment evolution, changing disease, time-varying covariates, informative sampling, and model updating without assuming that any component is sufficient alone. The framework remains conceptual and requires drug-specific, population-specific, computational, prospective, and implementation validation before decision use.

Key words: *Clinical pharmacokinetics, Pharmacodynamics, PK/PD, Dose–exposure–response, Population modelling, Precision dosing*

Received: 11September 2025; **Revised:** 05 December 2025; **Accepted:** 08 December 2025

INTRODUCTION

Population pharmacokinetics characterizes variability in drug disposition across and within patients, while pharmacodynamics relates exposure to biomarkers, physiological effects, therapeutic response, or toxicity. Their integration provides a structured account of the dose–exposure–response chain, but clinical use requires more than estimating a population parameter distribution. Model-informed precision dosing must connect an

appropriate model to reliable patient data, a defined therapeutic objective, an interpretable prediction, and a feasible decision process. Persistent implementation barriers show that scientific model development, individual forecasting, workflow integration, and stakeholder confidence are interdependent rather than separable stages [1].
Therapeutic drug monitoring and model-informed precision dosing are related but not interchangeable. Conventional therapeutic drug monitoring often interprets measured concentrations against predefined ranges, whereas model-informed precision dosing combines prior population information with patient-specific dose, concentration, covariate, and response data to forecast exposure or support regimen evaluation. This distinction matters because a concentration result is not an isolated biological fact: its interpretation depends on sampling time, prior doses, adherence, assay performance, the exposure target, and the clinical context in which the test was ordered. Contemporary antibiotic applications illustrate how sampling strategy, pharmacokinetic forecasting, pharmacodynamic targets, and clinical interpretation must be coordinated [2].
Digital health infrastructure creates opportunities to assemble longitudinal dose records, laboratory measurements, clinical assessments, and model outputs, but data availability does not itself establish temporal validity or decision readiness. Electronic records may contain delayed documentation, conflicting timestamps, omitted administrations, incomplete adherence information, and measurements collected preferentially during clinical deterioration. A digitally integrated model must therefore distinguish the time at which a biological event occurred, the time at which it was observed, the time at which it became available to the model, and the time at which a decision was made. Recommendations for model-informed dosing in the digital-health era similarly emphasize the need to integrate scientific, computational, clinical, and implementation requirements [3].
This Original Longitudinal Modelling Framework Article proposes that population PK/PD during clinical care should be represented as a coupled state–event–observation system. The framework integrates latent PK and PD states, evolving disease, time-varying physiology, treatment and dose history, visit and sampling processes, missingness, sequential updating, and decision feedback. It distinguishes prediction from causal inference, patient-specific updating from population-model relearning, and analytical output from clinical action. The proposed relationships organize testable hypotheses and validation requirements; they do not establish that dynamic modelling is universally superior, that observed associations are causal, or that the architecture is clinically validated, regulator-endorsed, or deployment-ready.

Limitations of static population PK/PD assumptions

Static population PK/PD models often represent patient characteristics through baseline covariates and assume stable parameter–covariate relationships. This may be adequate over short, controlled periods, but renal and hepatic function, body composition, inflammation, protein concentrations, disease burden, concomitant therapy, and treatment response can change during prolonged care. Time-varying covariates may also precede treatment, mediate treatment effects, reflect disease progression, or result from prior treatment. Indiscriminate adjustment may therefore remove part of the causal pathway or introduce overadjustment bias [4].
Static assumptions also apply to model structure. Several plausible structural or covariate models may fit the same development data, yet predictions are commonly based on a single selected model. Vancomycin research showed that accounting for uncertainty across candidate models can improve individualized prediction [5]. Model averaging is not always necessary, but structural uncertainty should remain visible when predictions inform treatment.
Within-patient pharmacokinetic change is another concern. Clearance may vary with organ function, disease status, catabolism, immunogenicity, or treatment response. Time-varying nivolumab clearance illustrates that temporal change differs from interindividual variability [6]. Parameter stability should therefore be tested rather than assumed.
Static models may fail by fixing changing parameters, treating treatment-affected covariates as independent, replacing treatment history with current dose, ignoring informative sampling, or concealing structural uncertainty. Apparent change may also reflect assay error, misspecification, dropout, undocumented dosing, adherence, or selective sampling. The solution is not universal complexity, but a longitudinal architecture that makes assumptions explicit and compares dynamic models with simpler alternatives.

Proposed longitudinal modelling framework

The framework separates states, events, observations, and decisions. Latent states include clearance, distribution, exposure, pharmacodynamic sensitivity, toxicity risk, and disease severity. Events include dosing, treatment interruption, infusion changes, concomitant therapy, and supportive care. Observations include concentrations,

laboratory values, biomarkers, clinical assessments, and outcomes, while decisions represent clinical interpretation and treatment adaptation. Precision dosing requires these elements to remain connected while preserving diagnostics, uncertainty, and scientific oversight [7].

The architecture follows clinical time. Treatment alters exposure; exposure affects pharmacodynamic and disease trajectories; changing physiology and disease may alter later pharmacokinetics; and latent states generate imperfect observations through the sampling process. EHR integration therefore requires more than data transfer: provenance, timestamps, model versioning, uncertainty display, and workflow placement must also be managed [8]. An explicit observation layer separates biological change from the process determining when and how it is measured.

Figure 1 presents the original conceptual synthesis for longitudinal population pk/pd system under evolving treatment and disease, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

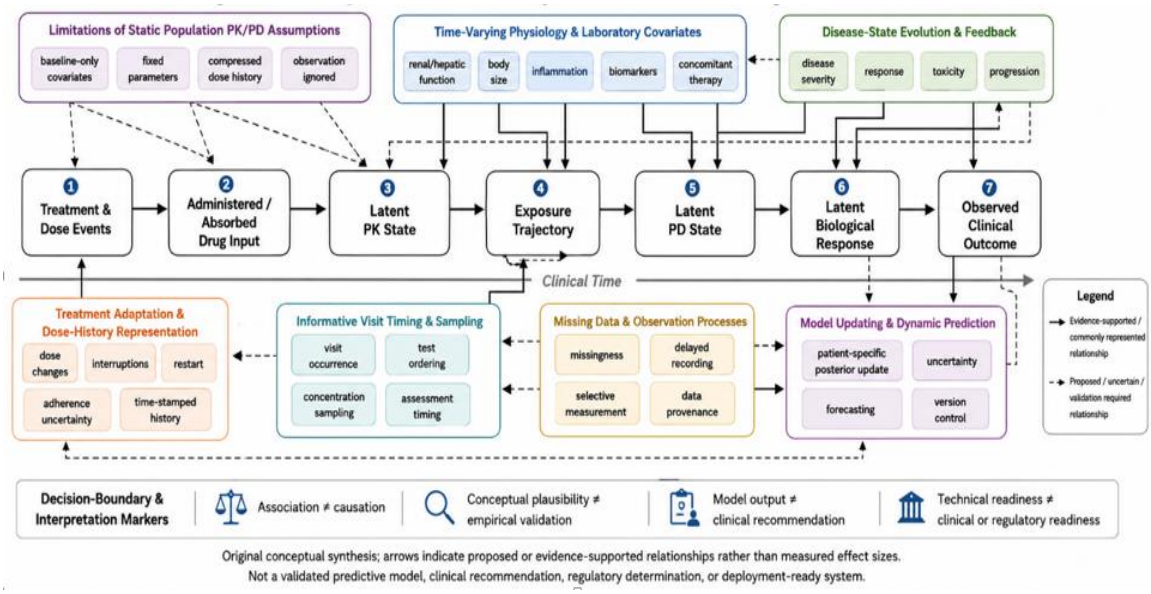

Figure 1. Longitudinal Population PK/PD System under Evolving Treatment and Disease.

The figure is an original conceptual synthesis developed for "Population Pharmacokinetics and Pharmacodynamics under Treatment Evolution, Time-Varying Covariates, Informative Sampling, and Changing Disease State during Clinical Care." Arrows and grouping indicate proposed or evidence-supported relationships rather than measured effect sizes. The figure does not represent a validated predictive model, clinical recommendation, regulatory determination, or deployment-ready system.

The patient-specific update component receives the information that was available at a defined prediction time and estimates a posterior distribution over current PK, PD, and disease states. This distribution should preserve uncertainty arising from sparse sampling, residual variability, measurement error, uncertain dosing, and model structure. It may then support prediction of future exposure or response under specified treatment scenarios. Continuous-learning work in paediatric vancomycin dosing shows how accumulating clinical data can be used to evaluate and update a population model over time [9]. The framework nevertheless separates individual Bayesian updating from population-level relearning because they create different validation and governance obligations.

An individual update changes inference for one patient under a specified population model. Population relearning changes the model that will later be applied to multiple patients. The latter can alter parameter distributions, covariate relationships, uncertainty estimates, and predictions for groups not represented in the newest data. It therefore requires prespecified eligibility rules, data-quality review, version control, temporal validation, monitoring for drift, and a rollback mechanism. A model should not modify its population structure merely because new data are available, particularly when the new data were generated under previous model-informed decisions.

The framework also distinguishes model output from pharmaceutical consequence. A predicted exposure, probability distribution, or simulated regimen is an analytical product; a dose change is a clinical action conditioned on therapeutic objectives, available alternatives, toxicity, urgency, feasibility, and professional

judgement. The architecture may organize evidence for supervised decision support, but it does not prescribe treatment. Its value must be tested against simpler models and within explicit contexts of use. Where temporal change, informative observation, or treatment adaptation is negligible, a static model may remain preferable because it is more identifiable, interpretable, and operationally stable.

Time-varying physiology and laboratory covariates

A time-varying covariate is a measured or derived patient characteristic whose value and predictive relevance may change during the modelling interval. Examples include body weight, renal and hepatic function, serum proteins, inflammatory markers, disease burden, immunogenicity, and measures of treatment response. Longitudinal analysis of certolizumab pegol showed that disease-related factors and antidrug antibodies can contribute to pharmacokinetic variability in Crohn's disease [10]. Within the proposed architecture, such variables are not treated as interchangeable predictors. Their interpretation depends on whether they represent a determinant of disposition, a measurement of latent disease, a treatment consequence, or a combination of these roles.

Time-varying clearance in avelumab provides a related example in which pharmacokinetic change was associated with evolving disease state [11]. The conceptual difficulty is that a biomarker may improve prediction while remaining unsuitable for a causal interpretation. A contemporaneous association between disease severity and clearance does not by itself show whether disease changed clearance, exposure changed disease, treatment altered both, or selective observation produced the apparent relationship. Lagged specifications, joint latent-state models, mediation analysis, and sensitivity analyses may clarify alternatives, but causal claims require assumptions about temporal ordering, confounding, measurement, and intervention mechanisms.

Temporal covariate analysis is further complicated when measurements are unavailable at clinically important times. Population PK modelling of paracetamol in neonates demonstrated that missing time-dependent body-weight information can affect covariate handling and model performance [12]. Simple last-observation-carried-forward procedures may create artificial plateaus, while linear interpolation can imply unobserved trajectories and future-informed interpolation can leak information into earlier predictions. Model-based imputation introduces additional assumptions about the covariate process. The choice should therefore reflect biological plausibility, observation timing, intended prediction time, uncertainty propagation, and sensitivity to alternative missing-data mechanisms.

Across the selected PK examples, longitudinal disease markers, immunogenicity, body size, and time-varying clearance can function as time-indexed predictors, but their interpretation depends on measurement timing and causal position [10–12]. The proposed framework accordingly assigns each covariate a temporal-role declaration specifying when it becomes available, whether it is measured or latent, whether prior treatment may affect it, which state it is intended to predict, and whether its use is predictive or causally interpretive. This declaration does not solve confounding or identifiability. It makes the assumptions inspectable and prevents a changing laboratory value from entering the model as though it were an error-free, externally determined baseline characteristic.

Treatment adaptation and dose-history representation

Treatment during clinical care is an evolving sequence of actions rather than a fixed baseline category. A longitudinal dose history should distinguish prescribed, dispensed, administered, and plausibly received doses; route; amount; infusion duration; administration time; interruptions; restarts; and concomitant therapy. Dose–exposure–response evaluation in paediatric rivaroxaban development illustrates the importance of explicitly connecting administered regimens, resulting exposure, and clinical response across time [13]. The proposed framework therefore treats treatment as a timestamped event process that drives, but is analytically distinct from, the latent PK and PD states.

Accurate representation also requires temporal provenance. A dose may be administered at one time, documented later, corrected retrospectively, or omitted from the electronic record. Similarly, a concentration may be collected before its recorded sampling time or become available only after a clinical decision. An electronic health record-embedded infliximab dashboard demonstrated how longitudinal dose and concentration information can support individualized pharmacokinetic assessment in children with Crohn's disease [14]. The framework extends this principle by recording event time, documentation time, model-availability time, and decision time separately whenever the data permit.

Treatment adaptation introduces feedback because the current regimen is partly determined by previous observations and clinical judgements. Apparent associations between dose and subsequent response may consequently reflect disease severity, prior toxicity, clinician concern, or model-informed intervention rather than

an unconfounded treatment effect. Prospective validation of a vancomycin precision-dosing tool in intensive care patients shows why evaluation must extend beyond retrospective model fit to the operation of the complete dosing workflow [15]. Prospective feasibility, however, does not by itself establish improved patient outcomes or universal applicability.

The proposed dose-history representation therefore preserves the sequence of treatment opportunities and constraints rather than compressing them into a single current-dose variable. It can include uncertain adherence, withheld doses, delayed administrations, rescue therapy, and treatment changes made for reasons not captured by the model. This structure supports prediction and sensitivity analysis, but it does not determine an appropriate regimen. Drug-specific therapeutic objectives, clinical contraindications, available alternatives, and professional judgement remain outside the authority of the conceptual architecture.

Disease-state evolution and feedback

Disease state is represented as a latent or partially observed process that changes during treatment and generates multiple imperfect measurements. Depending on the condition, these measurements may include symptoms, viral load, tumour burden, inflammatory markers, organ-function indices, functional assessments, or composite scores. Joint population PK/PD analysis of bamlanivimab and etesevimab linked antibody pharmacokinetics with longitudinal viral-load response, illustrating how exposure and disease-related trajectories may be represented within one model [16]. No single biomarker should automatically be equated with the full disease state.

Disease evolution may also modify pharmacokinetics. Population analysis of durvalumab identified associations between longitudinal biomarkers of disease status and changing clearance [17]. Such findings are compatible with several mechanisms, including altered protein turnover, inflammatory activity, disease burden, nutritional status, or treatment response. They do not independently establish the direction of causation. The proposed framework consequently separates the observed biomarker trajectory from the latent disease process and allows alternative directional structures to be compared.

A disease-progression component should be defined by its context of use. It may describe natural history, predict future status, quantify treatment-associated change, support trial simulation, or provide an input to clinical decision support. These uses require different assumptions and validation evidence. Disease-progression modelling has been proposed as a means of strengthening clinical development and decision making, but its value remains conditional on biological plausibility, data quality, model qualification, and the question being addressed [18].

Within the proposed architecture, disease affects future PK, PD sensitivity, treatment decisions, and observation intensity, while exposure and treatment may alter future disease. These feedback pathways are hypotheses to be tested rather than presumed causal mechanisms. Identifiability may be limited when measurements are sparse, treatment changes rapidly, or the same biomarker influences both care decisions and model estimation. **Table 1** organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in population pharmacokinetics and pharmacodynamics under treatment evolution.

Table 1. Definitions, components, relationships, and intended uses in Population Pharmacokinetics and Pharmacodynamics under Treatment Evolution

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Latent PK state	Fixed parameters may obscure within-person change	Dose history, concentrations, physiology, disease status	**Proposed synthesis:** individual disposition evolves and generates exposure; time-varying clearance has been observed in treatment settings [6]	More explicit representation of temporal PK heterogeneity	Longitudinal concentrations, reliable dose timing, structural-model comparison	Non-identifiability, assay error, sparse sampling, misspecified temporal trajectory	Dynamic PK structure is justified only when supported relative to a simpler model

Latent PD state	Observed response may incompletely represent drug effect	Exposure, biomarkers, clinical outcomes, effect delays	**Proposed synthesis:** exposure drives an unobserved or partially observed response state; joint PK/PD trajectories are feasible [16]	Separation of exposure, effect, and measured outcome	Repeated response measures, temporal alignment, mechanistic plausibility	Confounding by disease evolution or concomitant treatment	A fitted exposure–response association does not establish treatment causality
Disease-state process	Disease changes during care and may affect both PK and PD	Disease severity, biomarkers, symptoms, organ function	**Proposed synthesis:** disease may affect disposition and be altered by treatment; longitudinal biomarkers have been associated with clearance [17]	Joint representation of disease, response, and disposition	Disease-specific longitudinal measurements and competing model structures	Biomarker surrogacy, reverse causation, latent-state misspecification	The disease state must not be reduced automatically to one observed marker
Treatment-event history	Current dose cannot capture previous administrations and adaptations	Prescribed and administered doses, interruptions, route, adherence	**Proposed synthesis:** timestamped treatment events generate exposure and condition future decisions; dose–exposure–response evaluation supports temporal linkage [13]	Reconstructable treatment context for forecasting	Verified administration records and adherence sensitivity analyses	Missing administrations, delayed documentation, uncertain adherence	The event history supports analysis but does not prescribe a dose
Time-varying covariates	Baseline values may become obsolete	Physiology, laboratory values, disease markers, concomitant therapies	**Proposed synthesis:** covariate effects depend on timing and causal role; treatment-affected covariates can create overadjustment [4]	More defensible temporal covariate use	Timestamped measurements, role declaration, lag and sensitivity analyses	Reverse causation, measurement error, future-information leakage	Predictive usefulness must not be confused with causal interpretation
Observation process	Collected data may depend on patient state	Visit timing, test ordering, sampling, documentation	**Proposed synthesis:** latent state influences whether and when observations occur	Recognition of informative sampling and selective visibility	Visit and order timestamps, reasons for testing, observation-process models	Unmeasured reasons for observation and selective follow-up	Irregular sampling is not necessarily informative, but informativeness must be assessed

Patient-specific update	A population model alone may not represent current individual state	Prior model, dose history, observations, covariates	**Proposed synthesis:** new data update posterior patient states; controlled continuous learning has been demonstrated [9]	Sequential individualized prediction with uncertainty	Temporally ordered prospective data and update diagnostics	Overconfident posteriors, model drift, data-quality failures	Individual Bayesian updating is distinct from population-model modification
Uncertainty and decision boundary	Point predictions can obscure decision risk	Parameter, structural, residual, measurement, and transport uncertainty	**Proposed synthesis:** uncertainty accompanies every state estimate and remains separate from clinical action	More transparent interpretation of forecasts	Calibration, uncertainty coverage, sensitivity and external validation	False precision, omitted structural uncertainty, context shift	Model output is analytical evidence, not an autonomous pharmaceutical decision

Informative visit timing and sampling

Clinical measurements are not collected on a neutral schedule. Patients may be assessed more often when symptoms worsen, toxicity is suspected, organ function changes, adherence is uncertain, or treatment modification is considered, whereas stable patients may generate fewer observations. Routine-data research distinguishes informative presence from informative observation: both inclusion in the dataset and measurement of a variable may depend on the underlying clinical state [19].

This dependence can bias population PK/PD inference if measurement timing is treated as independent of the process being estimated. A concentration obtained because of suspected toxicity is not equivalent to one collected at a protocol-defined time. EHR studies show that evolving patient characteristics and outcomes can influence visit timing and distort inference, although the direction and magnitude of bias are context-specific [20].

Possible responses include joint modelling of visits and outcomes, inverse-intensity weighting, shared random effects, sensitivity analyses, and explicit models of sample-ordering decisions [21]. None eliminates bias when important determinants of sampling remain unmeasured. The framework therefore distinguishes visit occurrence, test ordering, sample collection, assay completion, and result recording. Protocol, routine, symptom-triggered, and opportunistic samples should not be pooled without qualification, and sensitivity analyses should address the observation process itself [19–21].

Missing data and observation processes

Missingness in longitudinal PK/PD data has multiple forms. Doses may be undocumented, sample times unreliable, assay results below quantification, laboratory tests unrequested, or patients lost to follow-up. These situations arise through different mechanisms and should not be managed by uniform deletion or imputation [22]. EHR missingness is especially complex because clinical care determines what is measured. A value may be absent because testing was considered unnecessary, the patient was stable, the result existed elsewhere, or the care process failed. Missingness should therefore be understood as a product of clinical, organizational, and technical processes, and missing-at-random assumptions must be justified [23].

Observed datasets may also be selectively generated when inclusion depends on disease severity, health-care use, or outcome-related factors, potentially altering estimated associations [24]. Sensitivity analyses should consider both missing values and missing opportunities for observation. The framework therefore records ordering, collection, processing, and result times, together with the reason for testing where available. Missingness indicators may inform the observation model but should not be treated as biological variables without justification.

Model updating and dynamic prediction

Dynamic prediction estimates future exposure, response, toxicity, or disease state using information available at a specified time and may be revised as new doses, concentrations, laboratory results, or outcomes become

available. Sequential prediction methods illustrate how forecasts can use an expanding longitudinal history rather than a fixed baseline, although evidence outside pharmacometrics does not validate their use in PK/PD [25].

A digital twin is more demanding than an updated forecast. It requires a patient-specific computational representation supported by repeated data assimilation and mechanistic or data-driven modelling [26]. A PK/PD digital twin would therefore need validated patient-state, treatment-effect, observation, uncertainty, and updating components. A Bayesian forecast should not be labelled a digital twin simply because it changes after a new measurement.

Individual updating must also be distinguished from population relearning. Individual updating estimates a patient-specific state using a fixed or version-controlled model, whereas population relearning changes the underlying model. Continuous-learning studies in paediatric vancomycin dosing show that accumulating clinical data can support model refinement [9], but such changes require governance because earlier models may have influenced treatment and sampling.

The proposed architecture separates a locked analytical model, a patient-specific posterior state, an evaluation environment for candidate revisions, and a governance gate for model replacement. Updates should be tested on later data, assessed across subgroups, and compared with the current model before release. Updating frequency should reflect clinically meaningful change rather than data availability, because rapid updates may amplify transient errors or local practice patterns.

Every prediction must also specify its prediction time and available information. Data unavailable at that moment must not enter the forecast. Prospective replay and silent evaluation can detect temporal leakage. Updating is not assumed to improve performance; a locked model may be preferable when data are sparse, drift is uncertain, or governance is inadequate.

Evaluation, calibration, and external transportability

Evaluation should begin with the intended use of the model. Parameter recovery, goodness-of-fit plots, residual diagnostics, and internal predictive checks are necessary but insufficient when the output will support future prediction or treatment evaluation. Calibration concerns agreement between predicted and observed quantities or risks and is distinct from discrimination or apparent fit. Predictive-methodology research identifies poor calibration as a central weakness when models are translated into decision contexts [27].

External transportability asks whether the model remains adequate in populations, sites, disease stages, dosing practices, assays, and observation regimes that differ from development conditions. External evaluation of published imatinib population PK models demonstrated that predictive performance can vary when models are applied to an independent clinical population [28]. Recalibration, covariate revision, structural modification, or rejection may therefore be appropriate responses; external evaluation is not merely a procedural confirmation exercise.

Fit-for-purpose evaluation requires both calibration assessment and testing in independent populations or settings [27, 28]. For the proposed framework, this should include population-level calibration, individual predictive error, uncertainty coverage, temporal stability, sensitivity to missingness and informative sampling, and performance under treatment adaptation. Evaluation should compare the longitudinal architecture with simpler alternatives so that additional complexity is retained only when it improves relevant prediction or inference without unacceptable instability.

Transportability must also cover the observation and decision processes. A model developed in protocolized sampling may not transfer to symptom-triggered clinical monitoring, even if biological relationships are similar. Changes in test-ordering, assay platforms, clinical thresholds, or model-supported treatment can alter the data distribution. Sensitivity analysis for informative visits and presence is therefore part of transportability assessment rather than an optional secondary analysis [20].

Limitations and implementation priorities

The framework is limited by the information available in clinical care. Adherence may be unknown, dose timestamps may be inaccurate, important physiological states may be unmeasured, and treatment changes may depend on clinician reasoning that is absent from the dataset. Implementation evidence also indicates that model-informed precision dosing faces scientific, technical, organizational, workflow, usability, and professional barriers [29]. A coherent architecture cannot eliminate these constraints; it can only expose where assumptions and responsibilities reside.

Advanced machine-learning components may support representation learning, data harmonization, missing-data modelling, or extraction from heterogeneous records. Foundation models have been proposed as a basis for generalist medical artificial intelligence, but their scale and flexibility introduce concerns about calibration, provenance, bias, interpretability, and external validity [30]. Within longitudinal PK/PD, such models should remain auxiliary unless their inputs, outputs, uncertainty, failure modes, and contribution beyond validated pharmacometric approaches are independently established.

Implementation should proceed through bounded stages. Computational testing should assess identifiability, temporal leakage, structural uncertainty, and failure under simulated misspecification. Retrospective temporal evaluation should precede prospective silent testing, independent external evaluation, supervised workflow use, and formal decision-impact assessment. Each model version should have a declared context of use, accountable owners, documented data dependencies, audit trails, and withdrawal criteria. These priorities are proposed governance requirements, not regulator-approved readiness levels, and they do not imply that every application should progress to clinical deployment.

CONCLUSION

This article proposes a longitudinal population PK/PD framework in which treatment events, latent pharmacokinetic and pharmacodynamic states, evolving disease, time-varying covariates, informative observation, missingness, patient-specific updating, and clinician-mediated feedback are represented as connected but analytically distinct processes. Its principal contribution is to move beyond the assumption that baseline covariates, fixed parameters, complete dose records, neutral sampling, or one-time model fitting are sufficient during evolving clinical care. The architecture generates testable propositions concerning temporal covariates, dose history, disease–PK/PD coupling, observation-process modelling, dynamic prediction, calibration, and transportability. It remains a conceptual synthesis. Drug-specific evidence, identifiable models, reliable provenance, prospective validation, external evaluation, uncertainty assessment, clinical oversight, and controlled governance are required before any decision use. The framework therefore provides a structured research agenda and a vocabulary for separating biological plausibility, predictive performance, causal interpretation, implementation feasibility, and clinical readiness without treating any one of them as proof of the others.

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. 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
3. 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
4. Goulooze SC, Snelder N. Time-varying covariates, overadjustment bias and mediation in pharmacokinetic/pharmacodynamic modeling. CPT Pharmacometrics Syst Pharmacol. 2024;13(8):1285-8. doi:10.1002/psp4.13200
5. 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

6. Bajaj G, Wang X, Agrawal S, Gupta M, Roy A, Feng Y. Model-based population pharmacokinetic analysis of nivolumab in patients with solid tumors. CPT Pharmacometrics Syst Pharmacol. 2017;6(1):58-66. doi:10.1002/psp4.12143
7. 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
8. 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
9. 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
10. Vande Casteele N, Mould DR, Coarse J, Hasan I, Gils A, Feagan BG, et al. Accounting for pharmacokinetic variability of certolizumab pegol in patients with Crohn's disease. Clin Pharmacokinet. 2017;56(12):1513-23. doi:10.1007/s40262-017-0535-3
11. 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
12. Krzyzanski W, Cook SF, Wilbaux M, Sherwin CMT, Allegaert K, Vermeulen A, et al. Population pharmacokinetic modeling in the presence of missing time-dependent covariates: Impact of body weight on pharmacokinetics of paracetamol in neonates. AAPS J. 2019;21(4):68. doi:10.1208/s12248-019-0331-0
13. Young G, Lensing AWA, Monagle P, Male C, Thelen K, Willmann S, et al. Rivaroxaban for treatment of pediatric venous thromboembolism: An EINSTEIN-Jr phase 3 dose-exposure-response evaluation. J Thromb Haemost. 2020;18(7):1672-85. doi:10.1111/jth.14813
14. Xiong Y, Mizuno T, Colman R, Hyams J, Noe JD, Boyle B, et al. Real-world infliximab pharmacokinetic study informs an electronic health record-embedded dashboard to guide precision dosing in children with Crohn's disease. Clin Pharmacol Ther. 2021;109(6):1639-47. doi:10.1002/cpt.2148
15. ter Heine R, Keizer RJ, van Steeg K, Smolders EJ, van Luin M, Derijks HJ, et al. Prospective validation of a model-informed precision dosing tool for vancomycin in intensive care patients. Br J Clin Pharmacol. 2020;86(12):2497-506. doi:10.1111/bcp.14360
16. Chigutsa E, O'Brien L, Ferguson-Sells L, Long A, Chien J. Population pharmacokinetics and pharmacodynamics of the neutralizing antibodies bamlanivimab and etesevimab in patients with mild to moderate COVID-19 infection. Clin Pharmacol Ther. 2021;110(5):1302-10. doi:10.1002/cpt.2420
17. Baverel PG, Dubois VFS, Jin CY, Zheng Y, Song X, Jin X, et al. Population pharmacokinetics of durvalumab in cancer patients and association with longitudinal biomarkers of disease status. Clin Pharmacol Ther. 2018;103(4):631-42. doi:10.1002/cpt.982
18. Starling MS, Kehoe L, Burnett BK, Green P, Venkatakrishnan K, Madabushi R. The potential of disease progression modeling to advance clinical development and decision making. Clin Pharmacol Ther. 2025;117(2):343-52. doi:10.1002/cpt.3467
19. Sisk R, Lin L, Sperrin M, Barrett JK, Tom BDM, Diaz-Ordaz K, et al. Informative presence and observation in routine health data: A review of methodology for clinical risk prediction. J Am Med Inform Assoc. 2021;28(1):155-66. doi:10.1093/jamia/ocaa242
20. Goldstein BA, Phelan M, Pagidipati NJ, Peskoe SB. How and when informative visit processes can bias inference when using electronic health records data for clinical research. J Am Med Inform Assoc. 2019;26(12):1609-17. doi:10.1093/jamia/ocz148
21. Sun Y, McCulloch CE, Marr KA, Huang CY. Recurrent events analysis with data collected at informative clinical visits in electronic health records. J Am Stat Assoc. 2021;116(534):594-604. doi:10.1080/01621459.2020.1801447
22. Irby DJ, Ibrahim ME, Dauki AM, Badawi MA, Illamola SM, Chen M, et al. Approaches to handling missing or problematic pharmacology data: Pharmacokinetics. CPT Pharmacometrics Syst Pharmacol. 2021;10(4):291-308. doi:10.1002/psp4.12611
23. Tsiampalis T, Panagiotakos DB. Methodological issues of the electronic health records' use in the context of epidemiological investigations, in light of missing data: A review of the recent literature. BMC Med Res Methodol. 2023;23(1):180. doi:10.1186/s12874-023-02004-5

24. Harton J, Mitra N, Hubbard RA. Informative presence bias in analyses of electronic health records-derived data: A cautionary note. J Am Med Inform Assoc. 2022;29(7):1191-9. doi:10.1093/jamia/ocac050
25. Jarrett D, Yoon J, van der Schaar M. Dynamic prediction in clinical survival analysis using temporal convolutional networks. IEEE J Biomed Health Inform. 2020;24(2):424-36. doi:10.1109/JBHI.2019.2929264
26. 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
27. Van Calster B, McLernon DJ, van Smeden M, Wynants L, Steyerberg EW, Bossuyt PMM, et al. Calibration: The Achilles heel of predictive analytics. BMC Med. 2019;17(1):230. doi:10.1186/s12916-019-1466-7
28. Corral Alaejos Á, Zarzuelo Castañeda A, Jiménez Cabrera S, Sánchez-Guijo F, Otero MJ, Pérez-Blanco JS. External evaluation of population pharmacokinetic models of imatinib in adults diagnosed with chronic myeloid leukaemia. Br J Clin Pharmacol. 2022;88(4):1913-24. doi:10.1111/bcp.15122
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. Moor M, Banerjee O, Abad ZSH, Krumholz HM, Leskovec J, Topol EJ, et al. Foundation models for generalist medical artificial intelligence. Nature. 2023;616(7956):259-65. doi:10.1038/s41586-023-05881-4