



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Precision Pharmacotherapy Requires Dynamic PKPD Updating After Treatment Begins

Lucas Meyer^{1*}, Anna Schmid², Stefan Braun¹, Laura Keller³

¹Department of Precision Pharmacotherapy and Dynamic PKPD, Faculty of Pharmacy, ETH Zurich, Zurich, Switzerland.

²Department of Treatment Update and Drug Response, Faculty of Pharmaceutical Sciences, University of Bern, Bern, Switzerland.

³Department of Clinical Pharmacometrics, Faculty of Pharmacy, EPFL Lausanne, Lausanne, Switzerland.

*Email: lucas.meyer@ethz.ch

ABSTRACT

Precision pharmacotherapy is often framed as the selection of an individualized dose from pretreatment characteristics. That framing is incomplete because treatment itself generates new information. Measured concentrations, early pharmacodynamic responses, toxicity signals, evolving organ function, inflammatory state, adherence, and other time-varying features can make an initially reasonable PKPD prediction progressively less representative of the patient. This Perspective argues that precision dosing should therefore be treated as a sequential inference problem rather than a one-time calculation. Prior population PKPD knowledge remains essential, but it should function as an auditable prior whose transportability and structural assumptions can be challenged by post-dose observations. Early concentration and response data should update different components of uncertainty, while treatment-affected covariates require special handling because they may be predictors, mediators, consequences, or markers rather than independent causal drivers. We propose a dynamic PKPD updating logic that separates prior-model uncertainty, patient-specific parameter uncertainty, observation uncertainty, and evolving clinical-state uncertainty before dose adaptation. The framework does not assume that more frequent updating is always beneficial. Its clinical value depends on observability, response lag, model adequacy, decision consequences, and prospective validation of the complete data-to-dose workflow.

Keywords: Model-informed precision dosing, Bayesian forecasting, Pharmacokinetics, Pharmacodynamics, Therapeutic drug monitoring, Time-varying covariates

Received: 25 June 2026; **Revised:** 21 September 2026; **Accepted:** 25 September 2026

INTRODUCTION

Model-informed precision dosing (MIPD) has moved individualized pharmacotherapy beyond simple covariate-based dose adjustment by combining prior pharmacokinetic or pharmacodynamic knowledge with patient-level data. Contemporary reviews emphasize that the useful unit of prediction is not merely the “average patient with covariates,” but an individual whose state can be revised as treatment proceeds [1]. Yet a persistent translational gap remains between pharmacometric capability and routine clinical use [2]. That gap is not explained by model availability alone. Model qualification, measurement infrastructure, decision-support design, governance, and clinical validation all determine whether a numerically plausible recommendation can become a defensible clinical action [3]. A recent systematic review of implementation further shows that barriers and facilitators are

distributed across technical, organizational, measurement, regulatory, and user-facing domains rather than residing in one component of the dosing algorithm [4].

This Perspective focuses on one inferential problem within that broader implementation challenge: pretreatment prediction is necessarily provisional. Baseline physiology, genotype, body size, comedication, disease characteristics, and prior population models may determine a rational starting dose, but they cannot fully encode future clearance, target exposure, pharmacodynamic sensitivity, adherence, treatment-emergent toxicity, or disease evolution. We therefore propose that precision pharmacotherapy be organized around five analytically distinct operations: prior specification, post-start observation, Bayesian or otherwise model-based updating, sequential dose action, and independent validation. The distinction matters because a better prior is not equivalent to a valid posterior, a good concentration prediction is not equivalent to a validated response prediction, and either can be clinically unhelpful if the treatment state changes before the next decision.

Why pre-treatment prediction is not enough

Pretreatment prediction can be improved substantially without eliminating individual uncertainty. In a very large vancomycin data set, machine-learning-assisted model selection and averaging improved a priori concentration prediction relative to single-model strategies, illustrating that even model choice can be individualized before therapeutic drug monitoring begins [5]. The result is important but bounded: it concerns prediction before patient-specific concentration feedback, not proof that baseline information is sufficient. Earlier external evaluation of vancomycin pharmacometric models similarly showed marked variability in a priori performance and materially improved individual forecasting once measured concentrations were incorporated through Bayesian updating [6]. The practical implication is not that pretreatment models are weak; it is that their uncertainty remains revisable. The same limitation appears at the level of patient representation and population transport. The virtual-twin concept highlights that individualized prediction depends on the availability and relevance of patient-specific information rather than on a nominally individualized model alone [7]. External validation of infliximab population PK models in Korean adults and children with inflammatory bowel disease likewise demonstrated that published models do not transfer uniformly across target populations [8]. A transportability error can reflect differences in population, design, assay, dosing history, covariate distribution, or model structure; it should not be converted automatically into a biological or ethnic mechanism.

Treatment should therefore be viewed as an information-generating phase. The initial dose remains anchored in prior evidence, but each post-start observation creates an opportunity to ask whether the patient still resembles the prior state strongly enough for the previous recommendation to remain defensible. Dynamic updating is most useful where the new observation is informative, timely relative to the decision horizon, and interpretable within an adequate model. Where those conditions fail, repeated recalculation can create an illusion of precision rather than improve it.

Prior PKPD information

A dynamic approach begins with a disciplined prior. For tacrolimus, the literature contains many population PK models that differ in structural assumptions, included covariates, transplant populations, and software implementation, demonstrating that “the prior” is often a family of plausible representations rather than a single settled model [9]. Guidance on external evaluation similarly argues that model selection should be tied to intended use, target population, available data, and fit-for-purpose diagnostics [10]. Thus, prior PKPD information should carry provenance: what population generated it, which measurements informed it, what structural assumptions it makes, and which clinical decision it is expected to support.

Model uncertainty can also be represented explicitly. In vancomycin MIPD, model averaging and selection improved predictive performance in the evaluated setting, showing that uncertainty about which population model best represents a patient need not be hidden behind a single chosen prior [11]. Conversely, local evaluation in neonates found that only one of several vancomycin models satisfied the study's acceptability criterion across a priori, a posteriori, and next-course forecasting [12]. A model that is credible in publication can therefore remain unsuitable for a particular clinical population.

For this Perspective, the proposed role of prior information is not to declare a patient's PKPD state, but to define a starting distribution of plausible states and the observations capable of discriminating among them. **Table 1** makes this distinction explicit by separating what can be observed, what dose inference is justified, and what evidence is still needed before a discordant pattern is acted upon.

Table 1. Reconciling prior PKPD information with the first patient-specific evidence

Evidence configuration		Adjudication requirement	
Prior-information case / What is observable after treatment starts	Dose inference that is justified / Inference unjustified	Monitoring signal that resolves the case / Validation criterion	article-level test
Prior-information case: Prior model fits the intended population and early PK agrees What is observable after treatment starts: Concordant concentration trajectory	Dose inference that is justified: Continue the current model-informed dose while uncertainty remains acceptable Inference that remains unjustified: Assume future response will remain concordant	Monitoring signal that resolves the case: Next concentration or response marker at a decision-relevant time Validation criterion: Stable predictive performance under repeated observations	Does new evidence preserve the prior state?
Prior-information case: Prior model is credible but early concentration is discordant What is observable after treatment starts: Concentration deviates beyond the expected individual trajectory	Dose inference that is justified: Re-estimate individual parameters and examine model mismatch Inference that remains unjustified: Attribute the deviation automatically to nonadherence or physiology	Monitoring signal that resolves the case: Repeat or strategically timed sample plus contemporaneous clinical covariates Validation criterion: Posterior prediction improves without systematic residual error	Can parameter updating explain the discordance without changing model structure?
Prior-information case: Several plausible priors remain What is observable after treatment starts: Competing models predict materially different exposure from the same observations	Dose inference that is justified: Retain model uncertainty or use validated model averaging or selection Inference that remains unjustified: Choose a single model because it performs best in another population	Monitoring signal that resolves the case: Observation with high discrimination across candidate priors Validation criterion: External and patient-level predictive adequacy	Is the dosing decision robust to prior-model choice?
Prior-information case: PK is concordant but early response is discordant What is observable after treatment starts: Exposure appears expected while effect or toxicity trajectory differs	Dose inference that is justified: Investigate PD sensitivity, response lag, endpoint validity, and competing causes Inference that remains unjustified: Escalate solely because concentration is “on target”	Monitoring signal that resolves the case: Validated response or toxicity signal at an appropriate lag Validation criterion: Exposure–response interpretation reproduces prospectively	Does the response require a PD-state update rather than a PK dose correction?
Prior-information case: Clinical state changes after the prior was specified What is observable after treatment starts: New organ function, inflammatory state, interacting therapy, or treatment condition	Dose inference that is justified: Reassess whether the original prior remains transportable to the current state Inference that remains unjustified: Treat baseline covariates as permanently valid	Monitoring signal that resolves the case: Time-aligned covariate and concentration or response data Validation criterion: Model remains calibrated across the changed state	Has the prior expired for the current decision horizon?

Note: The cases and article-level tests are a proposed synthesis. They are not validated clinical rules or quantitative thresholds.

Early concentration data

A concentration becomes useful only when its timing, assay, model, and intended PK quantity make it informative. Bayesian estimation of vancomycin exposure in adults with hematologic malignancy shows that sampling strategy changes the accuracy with which AUC can be reconstructed [13]. The implication is broader than vancomycin: “an early level” is not a generic unit of evidence. A sample obtained at a poorly informative time may narrow the posterior less than expected or may preferentially identify one parameter while leaving another weakly determined.

Endpoint-specific observability is equally important. In bone and joint infection, a levofloxacin MIPD application identified sparse sampling capable of supporting AUC estimation, while the evaluated strategies did not satisfactorily estimate C_{max} [14]. The same observations can therefore be adequate for one dosing objective and inadequate for another. For 5-fluorouracil, evaluation across infusion regimens further showed that limited-sampling strategies and model performance depend on the administration design; a schedule optimized for one infusion structure should not be assumed transportable to another [15].

Even after a concentration is measured, uncertainty about the prior model can remain. In patients receiving sunitinib or pazopanib, model averaging was more robust across tested prediction scenarios than reliance on a single model, whereas flattened priors rarely improved prediction [16]. Early concentration data should therefore update a structured uncertainty state rather than merely replace a population estimate with an individual number.

Early response data

Concentration is only one post-treatment signal. In pediatric Crohn's disease, real-world infliximab PK analyses linked patient and disease-related information to clearance and supported an electronic dashboard for precision dosing [17]. Such information can indicate that a patient is moving into a pharmacokinetic state not anticipated from baseline alone. Response-side information can be even more distinct. A carboplatin PKPD study used longitudinal neutrophil and platelet trajectories from electronic health records to inform individualized prediction of myelosuppression, demonstrating how repeated toxicity measurements can contribute to patient-specific PD characterization [18].

This distinction matters because PK concordance does not guarantee PD concordance. A patient can achieve expected exposure while showing delayed, attenuated, or exaggerated response. Conversely, an unexpected response may reflect disease evolution, concomitant treatment, endpoint noise, or response lag rather than altered drug exposure. Prospective eculizumab PK work in paroxysmal nocturnal hemoglobinuria also illustrates clinically relevant heterogeneity under fixed dosing and the potential value of individualized schedules [19], but it does not justify a general rule that every biologic should be adjusted from the same response signal.

We therefore propose parallel PK and PD update lanes. Concentration data primarily challenge assumptions about exposure and individual PK parameters; response data challenge assumptions about sensitivity, delay, disease state, and clinical consequence. Discordance between the two is not a nuisance to be averaged away. It is diagnostic information that determines which layer of the model should be questioned next.

Bayesian updating

Bayesian forecasting provides the natural inferential engine for combining prior PK information with patient observations, but the posterior is conditional on the adequacy of both. Bedside MIPD discussions have long emphasized that patient-specific prediction requires appropriate models, relevant data, and scientifically defensible implementation rather than Bayesian computation in isolation [20]. The first boundary is therefore model adequacy: if the prior model is structurally inappropriate for the patient or current treatment state, a numerically precise posterior can still be misleading.

Updating also occurs at more than one scale. Continuous-learning work in pediatric vancomycin has shown how accumulating clinical data can be used to detect and reduce mismatch between an existing population model and a new population [21]. Methodological work on continued learning formalizes this broader process of model updating in clinical practice [22]. These operations should not be conflated. Patient-level Bayesian updating estimates the current individual's parameters conditional on a population prior; population-level learning modifies the prior itself as evidence accumulates across patients. The latter needs governance because indiscriminate online updating can propagate measurement bias, treatment-selection bias, or local practice artifacts into the model.

Model uncertainty can persist at both scales. External evaluation of cyclosporine population PK models and multi-model strategies in adult renal transplantation demonstrates that model choice remains a live problem even when rich therapeutic drug monitoring information is available [23]. We therefore propose decomposing uncertainty during Bayesian updating into four analytically separate components: patient-parameter uncertainty, structural or prior-model uncertainty, observation uncertainty, and evolving clinical-state uncertainty. The decomposition is not itself a validated estimator; it is a decision architecture intended to prevent one narrow posterior interval from being interpreted as certainty about the whole treatment state.

Table 2 places the boundaries of Bayesian interpretation next to the evidence needed to cross them, rather than relegating limitations to a final generic column.

Table 2. Boundary conditions that determine what a Bayesian PKPD update can legitimately change

Source state and scale		Transfer conditions	Qualification endpoint
Patient-variability scenario / Posterior interpretation supportable	PD or target-response status / Model-assumption requirement	Boundary that weakens or reverses that interpretation / Time lag or feedback concern / Covariate trajectory	synthesis consequence
Patient-variability scenario: Informative concentration, stable physiology, fit-for-use prior Posterior interpretation that remains supportable: Update individual PK parameters and exposure prediction	PD or target-response status: Response may still be unknown Model-assumption requirement: Prior remains locally predictive	Boundary that weakens or reverses that interpretation: Assay error or unrecognized structural misspecification Time lag or feedback concern: PK information precedes PD Covariate trajectory: Approximately stable	Permit PK adaptation but preserve PD uncertainty
Patient-variability scenario: Concentration discordant across candidate models Posterior interpretation that remains supportable: Treat prior-model identity as part of uncertainty	PD or target-response status: Response interpretation depends on exposure model Model-assumption requirement: Candidate priors externally plausible	Boundary that weakens or reverses that interpretation: Model selection driven by the same noisy observation Time lag or feedback concern: Repeated sampling may discriminate models Covariate trajectory: Stable or known	Avoid false precision from one selected posterior
Patient-variability scenario: Exposure concordant, response discordant Posterior interpretation that remains supportable: Preserve PK posterior while reopening PD state	PD or target-response status: Unexpected effect or toxicity Model-assumption requirement: Exposure–response model must be applicable	Boundary that weakens or reverses that interpretation: Invalid endpoint, delayed response, disease progression, or competing therapy Time lag or feedback concern: Explicit response lag Covariate trajectory: May be treatment affected	Shift the next update toward PD rather than PK
Patient-variability scenario: Physiology changes after the last update Posterior interpretation that remains supportable: Reassess both covariate effect and prior transportability	PD or target-response status: May change independently or in parallel Model-assumption requirement: Time-varying relationship is specified appropriately	Boundary that weakens or reverses that interpretation: Covariate is itself altered by treatment or disease Time lag or feedback concern: State transition may precede measured response Covariate trajectory: Nonstationary	Treat the previous posterior as time-limited
Patient-variability scenario: Repeated local data show systematic model mismatch Posterior interpretation that remains supportable: Maintain individual safety monitoring while evaluating population recalibration	PD or target-response status: Drug-specific Model-assumption requirement: Model-updating governance is required	Boundary that weakens or reverses that interpretation: Local practice or selection process contaminates learning data Time lag or feedback concern: Cross-patient learning is slower than individual adaptation Covariate trajectory: Population distribution may drift	Separate patient rescue from population-model revision

Note: The scenarios and synthesis consequences are proposed analytical categories, not validated dosing instructions.

The uncertainty structure is summarized visually in **Figure 1** because Bayesian updating should not be interpreted as a single arrow from measurement to a new dose.




Figure 1. Where certainty can and cannot increase during Bayesian PKPD updating

Time-varying covariates

Dynamic updating becomes most difficult when the covariate itself changes because of treatment, disease evolution, or both. Simulation work on time-varying covariates in PKPD modeling demonstrates that adjustment for a treatment-affected covariate can create overadjustment bias when the variable lies on the causal pathway between treatment and outcome [24]. A postbaseline laboratory value is therefore not automatically an independent explanatory factor. Depending on context, it may function as a predictor of clearance, a mediator of treatment effect, a marker of latent disease state, a consequence of exposure, or measurement noise. We propose preserving these roles explicitly rather than forcing every changing variable into the same covariate function. The clinical importance of this distinction is clearest when physiology changes rapidly. Antibiotic precision-dosing literature emphasizes that critical illness can produce marked and evolving PK variability, making an earlier parameter estimate progressively less representative as organ function, fluid balance, extracorporeal support, inflammation, or other determinants change [25]. A recent vancomycin model combining population PK with machine-learning inputs further illustrates the predictive value that current concentration, laboratory, and dosing information can provide in a specific data set [26]. These findings support dynamic reassessment, but they do not establish a universal monitoring frequency or prove that every dynamic laboratory variable should enter a causal PKPD model.

For the proposed framework, each posterior state therefore has a decision horizon. Its relevance expires when enough state-defining information changes that the assumptions supporting the previous update are no longer credible. **Figure 2** visualizes this as an exposure–response feedback field rather than a linear chain: dose changes exposure, exposure and disease state shape response, treatment may alter covariates, and those covariates can in turn change subsequent exposure or the interpretation of response. The central safeguard is that feedback does not equal causality.




Figure 2. Treatment-induced state change can invalidate an otherwise coherent exposure–response update

Sequential dose adaptation

Once treatment has generated credible patient-specific information, precision dosing becomes a sequential decision problem. The relevant question is no longer simply which dose best fits the baseline model, but which action is justified by the current state, what uncertainty remains, and what observation should follow. Model-informed reinforcement-learning work illustrates how dose selection can be represented as a repeated state–action problem in which the value of an intervention depends on subsequent observations rather than on a single optimization step [27]. This computational formulation is useful conceptually, although present evidence does not justify treating reinforcement learning as a clinically validated replacement for pharmacometric modeling or clinical judgment.

Prospective evidence nevertheless shows that repeated model-based adjustment can improve defined exposure endpoints. In kidney-transplant recipients, electronic-health-record-embedded MIPD of tacrolimus improved attainment of the study's repeated trough-concentration target compared with physician-guided dosing [28]. The result demonstrates prospective benefit for exposure control within that setting; it does not establish corresponding improvements in rejection, toxicity, graft survival, or other patient outcomes.

Sequential adaptation must also permit reversal of an initially plausible recommendation. In a piperacillin scenario analysis, a priori model-informed dosing performed poorly for the evaluated target, whereas incorporation of the first and subsequent therapeutic drug monitoring results markedly improved simulated target attainment [29]. Dynamic precision pharmacotherapy should therefore avoid making the initial model recommendation path-dependent. New evidence may justify continuing, intensifying, reducing, delaying, or abandoning the previous dosing trajectory.

Stopping and escalation rules

A dynamic system requires explicit rules for when not to continue escalating. Busulfan provides a particularly instructive example. External evaluation and simulation showed that optimizing the initial dose from existing population PK models did not improve target attainment over the comparator approach, whereas incorporation of post-treatment TDM substantially improved exposure targeting [30]. The manuscript-level inference is narrower than the numerical result: a credible pretreatment model can still generate a dose that should be revised once patient-specific observations contradict its assumptions.

Stopping logic should therefore respond to evidence conflict rather than to a universal exposure threshold. Escalation may become inappropriate when concentration data indicate unexpected accumulation, when toxicity evolves despite apparently acceptable exposure, when the response endpoint has an unresolved delay, when organ function changes, or when the model no longer represents the patient's current state. Conversely, failure to reach

an exposure target does not automatically justify escalation if the target itself is uncertain or if the observed response suggests adequate effect.

This conditionality is especially important in oral oncology. Practical MIPD and therapeutic-drug-monitoring frameworks for tyrosine-kinase and PARP inhibitors emphasize that exposure targets, toxicity, efficacy evidence, timing, adherence, and the maturity of drug-specific exposure–response knowledge differ across agents [31]. The proposed stopping and escalation logic is therefore bidirectional and evidence-specific: dose movement should halt whenever the uncertainty surrounding the next action becomes clinically more important than the predicted benefit of continuing in the same direction.

A dynamic PKPD updating framework

The preceding sections support a proposed architecture in which treatment proceeds through repeated state estimation rather than through independent dose calculations. Electronic-health-record-integrated clinical decision support can technically connect patient data, pharmacometric models, and dose recommendations within the clinical workflow [32]. Model-informed artificial-intelligence approaches can additionally represent sequential decisions under evolving information [33]. Neither infrastructure nor algorithmic sophistication, however, removes the need to preserve uncertainty and validation boundaries.

The proposed framework therefore separates four uncertainty components at every clinically relevant update: patient-parameter uncertainty, concerning the individual's current PK or PD parameters; model uncertainty, concerning whether the structural and population model remains appropriate; observation uncertainty, concerning measurement quality and informativeness; and clinical-state uncertainty, concerning whether disease, physiology, concomitant treatment, adherence, or response state has changed since the previous estimate. Available MIPD software already differs substantially in model implementation, clinical integration, validation features, interoperability, and usability, reinforcing that the software layer itself must be distinguished from scientific model validity [34].

Each posterior state also receives a decision horizon. The state remains actionable only while the assumptions supporting it remain credible. A new concentration can reduce parameter uncertainty without resolving structural uncertainty. A new response measurement can challenge PD assumptions without implying abnormal PK. A changing covariate can make both estimates stale. **Table 3** compresses these state transitions into a temporal logic that can support, but does not prescribe, clinical decisions.

Table 3. Temporal states that determine when a PKPD estimate remains actionable

State and entry conditions	Evidence through time		Transition and reassessment
Dynamic PKPD state / Dose implication / unresolved article question	Entry condition / Covariate trajectory	What remains observable / Transition or persistence criterion	Model assumption under test / Dominant unresolved uncertainty
Dynamic PKPD state: Prior-concordant state Dose implication: Maintain current strategy while monitoring Proposed unresolved article question: How long can concordance justify carrying the same posterior forward?	Entry condition: Early observations agree with the intended-use prior Covariate trajectory: Stable or clinically unchanged	What remains observable: Concentration and emerging response Transition or persistence criterion: Concordance persists through the next decision-relevant observation	Model assumption under test: Population model remains transportable Dominant unresolved uncertainty: Individual parameter uncertainty
Dynamic PKPD state: Exposure-discordant state Dose implication: Re-estimate before further directional dose change Proposed unresolved article question: When should persistent PK discordance trigger model	Entry condition: Measured concentration conflicts with prior prediction Covariate trajectory: Stable or insufficiently explanatory	What remains observable: Repeated or strategically timed concentration Transition or persistence criterion: Discordance resolves after parameter update or persists across observations	Model assumption under test: Individual parameters or model structure may be wrong Dominant unresolved uncertainty: Parameter versus model uncertainty

replacement rather than further parameter updating?			
Dynamic PKPD state: Response-discordant state Dose implication: Avoid exposure-only escalation Proposed unresolved article question: Which response signals are sufficiently validated to overrule an exposure target?	Entry condition: Exposure is plausible but effect or toxicity is unexpected Covariate trajectory: May be stable	What remains observable: Longitudinal response or toxicity signal Transition or persistence criterion: Response becomes interpretable at an appropriate lag or alternative cause emerges	Model assumption under test: PK model may remain valid while PD state changes Dominant unresolved uncertainty: Response-lag, sensitivity, endpoint, or disease-state uncertainty
Dynamic PKPD state: State-transition condition Dose implication: Reopen the model-informed dose decision Proposed unresolved article question: When has a previously valid posterior effectively expired?	Entry condition: Organ function, disease activity, interacting therapy, support modality, or another material clinical variable changes Covariate trajectory: Nonstationary	What remains observable: Updated covariate plus PK or PD evidence Transition or persistence criterion: New observations establish a stable current state	Model assumption under test: Previous posterior may no longer represent the current patient Dominant unresolved uncertainty: Clinical-state and transportability uncertainty
Dynamic PKPD state: Recurrent model-mismatch state Dose implication: Protect the individual decision while model recalibration is evaluated separately Proposed unresolved article question: How should local learning improve a model without encoding local practice bias?	Entry condition: Systematic prediction error persists across patients or treatment episodes Covariate trajectory: Population distribution may drift	What remains observable: Cross-patient residual and validation patterns Transition or persistence criterion: Revised model passes independent fit-for-purpose validation	Model assumption under test: Population prior itself may need revision Dominant unresolved uncertainty: Model and learning-process uncertainty

Note: The temporal states, transition logic, and unresolved questions constitute the proposed synthesis of this Perspective. They are not validated dosing categories, clinical thresholds, or regulatory decision rules.

Clinical workflow and validation

Dynamic PKPD updating should be validated as a chain rather than as a single model. A clinically deployable system must establish that the prior model is appropriate for its intended population, measurements are sufficiently accurate and timely, the Bayesian or other updating method behaves as expected, dose recommendations are interpretable, software transfers data correctly, and clinicians can act on the recommendation without introducing new error. Prospective evidence such as the tacrolimus trial demonstrates that this full pathway can be tested for defined exposure endpoints, but similar evidence remains uneven across therapeutic areas.

Validation must also account for the fact that dynamic dosing changes the data subsequently available for model evaluation. Standard pharmacometric diagnostics often assume that treatment and sampling are sufficiently independent of the evolving patient state. Simulation work examining visual predictive checks in real-world data showed that adaptive dose changes, interval changes, dropout, or measurement processes can make conventional diagnostics appear inconsistent with the data even when the generating model itself is correct [35]. A failed diagnostic in an adaptive data set can therefore reflect the observation process rather than structural model failure. We consequently propose four separate validation gates. Model validation asks whether the pharmacometric representation predicts its intended quantities. Transport validation asks whether that representation remains suitable for the target population and current state. Workflow validation asks whether data acquisition, software, timing, communication, and dose execution preserve the intended inference. Clinical-utility validation asks whether the complete strategy improves a prospectively defined patient-relevant or treatment-relevant outcome. Passing an earlier gate cannot be used as evidence that a later gate has also been passed.

The same hierarchy limits interpretation of apparently successful precision dosing. Improved concentration prediction does not automatically establish improved response; improved target attainment does not automatically establish improved clinical outcome; and successful implementation at one institution does not establish

transportability to another. Dynamic updating is therefore most defensible when every new estimate carries both an uncertainty statement and an explicit statement of the decision for which that estimate has been validated.

CONCLUSION

Precision pharmacotherapy should not end when the first individualized dose is prescribed. Pretreatment models provide an essential prior, but treatment generates concentrations, responses, toxicity patterns, physiological changes, and other information capable of challenging that prior. The central proposal of this Perspective is therefore a dynamic PKPD updating architecture in which patient parameters, model structure, observations, and clinical state remain separate uncertainty domains; concentration and response information update different parts of that state; and sequential dose decisions are interrupted by explicit stopping, transportability, and validation boundaries.

The framework does not imply that every patient needs continuous monitoring or that every new measurement should change the dose. Updating is useful only when the observation is informative, the model is fit for purpose, the clinical state is interpretable, and the decision horizon makes revision consequential. Its value must ultimately be demonstrated through prospective validation of the entire data-to-decision pathway rather than inferred from model accuracy alone.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. 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
2. 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
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. Dibbets AC, Koldewij 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
5. van Os W, O'Jeanson A, Troisi C, Liu C, Brooks JT, Hughes JH, et al. Machine learning-based model selection and averaging outperform single-model approaches for a priori vancomycin precision dosing. *CPT Pharmacometrics Syst Pharmacol.* 2025;14(10):1650-60. doi:10.1002/psp4.70084
6. Broeker A, Nardecchia M, Klinker KP, Derendorf H, Day RO, Marriott DJ, et al. Towards precision dosing of vancomycin: A systematic evaluation of pharmacometric models for Bayesian forecasting. *Clin Microbiol Infect.* 2019;25(10):1286.e1-1286.e7. doi:10.1016/j.cmi.2019.02.029
7. Polasek TM, Rostami-Hodjegan A. Virtual twins: Understanding the data required for model-informed precision dosing. *Clin Pharmacol Ther.* 2020;107(4):742-5. doi:10.1002/cpt.1778
8. Kim Y, Baek SH, Jang JJ, Chung JY. Model-informed precision dosing of infliximab in Korean inflammatory bowel disease patients: External validation of population pharmacokinetic models. *CPT Pharmacometrics Syst Pharmacol.* 2025;14(10):1682-94. doi:10.1002/psp4.70089
9. Hoffert Y, Dia N, Vanuytsel T, Vos R, Kuypers D, Van Cleemput J, et al. Model-informed precision dosing of tacrolimus: A systematic review of population pharmacokinetic models and a benchmark study of software tools. *Clin Pharmacokinet.* 2024;63(10):1407-21. doi:10.1007/s40262-024-01414-y

10. El Hassani M, Marsot A. Guidance for external evaluation and selection of population pharmacokinetic models for precision dosing. *Clin Pharmacokinet.* 2026;65(1):1-9. doi:10.1007/s40262-025-01602-4
11. 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
12. Alrahahleh D, Thoma Y, Van Daele R, Nguyen T, Halena S, Luig M, et al. Bayesian vancomycin model selection for therapeutic drug monitoring in neonates. *Clin Pharmacokinet.* 2024;63(3):367-80. doi:10.1007/s40262-024-01353-8
13. Duong A, Le Blanc J, Projean D, Marsot A. Identifying the optimal sampling strategy for the Bayesian estimation of vancomycin AUC₀₋₂₄ in adult hematologic cancer patients. *Clin Pharmacokinet.* 2025;64(2):297-305. doi:10.1007/s40262-025-01478-4
14. Mimram L, Magreault S, Lemaitre F, Jauregui F, Mechaï F, Doukhi K, et al. Levofloxacin in bone and joint infections: Development of an application for model-informed precision dosing. *Clin Pharmacol Ther.* 2026;119(1):158-67. doi:10.1002/cpt.70044
15. Tan Z, Sancho-Araiz A, Völler S, Peeters SLJ, Manten T, Deenen MJ, et al. Advancing precision dosing of 5-FU: Population PK model development, limited sampling strategies, and fit-for-use application. *Clin Pharmacokinet.* 2026;65(7):1027-42. doi:10.1007/s40262-026-01645-1
16. Le Louedec F, Morvan L, Mourey L, Maillard M, Vachoux C, Yakoubi M, et al. Model-informed precision dosing of protein kinase inhibitors: Evaluation of model-averaging and flattened priors methods. *Clin Pharmacokinet.* 2026;65(5):681-92. doi:10.1007/s40262-026-01628-2
17. 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
18. De Carlo A, Tosca EM, Crul M, Schutte T, van Zuijlen L, Bahce I, et al. Model-informed precision dosing of carboplatin in cancer patients by leveraging myelosuppression data from electronic health records. *Br J Clin Pharmacol.* 2026;92(5):1457-72. doi:10.1002/bcp.70413
19. Ter Avest M, Langemeijer SMC, van den Heuvel LPWJ, Baas LM, van de Kar NCAJ, Ter Heine R. Model-informed precision dosing of ecilizumab in patients with paroxysmal nocturnal hemoglobinuria. *Clin Pharmacokinet.* 2025;64(8):1139-47. doi:10.1007/s40262-025-01536-x
20. 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
21. 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
22. 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
23. Liu F, Mao J, Cheng Z, Xu L, Huang S. Towards model-informed precision dosing of cyclosporine in adult renal transplantation: Assessing population pharmacokinetic models and multi-model strategies. *Eur J Pharm Sci.* 2025;214:107241. doi:10.1016/j.ejps.2025.107241
24. Gouloze 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
25. Wicha SG, Mårtson AG, Nielsen EI, Koch BCP, Friberg LE, Alffenaar JW, 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
26. Shi L, Hu N, Ling J, Yang X, Chen R. Model-informed vancomycin precision dosing by population pharmacokinetics combined with machine learning algorithms. *Br J Clin Pharmacol.* 2026;92(9):3169-79. doi:10.1002/bcp.70621
27. 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

28. Kuypers DRJ, Annaert P, Jallah B, Naesens M, Teuns M, Laenen A, et al. Electronic patient file-embedded model-informed precision dosing compared with physician dosing of tacrolimus in kidney transplantation. *Clin Pharmacol Ther.* 2026;119(2):381-92. doi:10.1002/cpt.70090
29. Haefliger D, Mina L, Guidi M, Marzolini C, Thoueille P, Rothuizen LE, et al. Individualization of piperacillin dosage based on therapeutic drug monitoring with or without model-informed precision dosing: A scenario analysis. *J Antimicrob Chemother.* 2025;80(3):840-7. doi:10.1093/jac/dkaf007
30. Mahomedradja RF, Bognàr T, van Schie KE, Kuipers IT, de Witte MA, Kuball JHE, et al. Model-informed precision dosing of busulfan for children and adults undergoing allogeneic hematopoietic stem cell transplantation: A critical evaluation of current PK models and dose recommendations. *Clin Pharmacol Ther.* 2025;118(3):723-34. doi:10.1002/cpt.3748
31. Torrent-Rodríguez A, Bastida C, Carcelero San Martín E, Monge Escartín I, Puszkiel A, Blanchet B, et al. Therapeutic drug monitoring and model-informed precision dosing of oral TKIs and PARP inhibitors: A practical framework for clinical implementation. *Clin Pharmacokinet.* 2026;65(6):773-93. doi:10.1007/s40262-026-01660-2
32. 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
33. Ribba B, Dudal S, Lavé T, Peck RW. Model-informed artificial intelligence: Reinforcement learning for precision dosing. *Clin Pharmacol Ther.* 2020;107(4):853-7. doi:10.1002/cpt.1777
34. Kantasiripitak W, Van Daele R, Gijssen 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
35. Hughes JH, Bergstrand M, Keizer RJ. The visual predictive check and real-world data. *Clin Pharmacol Ther.* 2026;120(2):394-8. doi:10.1002/cpt.70306