



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Exposure–Response Relationships Change Meaning When Time, Tolerance, and Disease Progression Interact

Rafael Mendoza^{1*}, Isabela Cruz¹, Carlos Lopez²

¹Department of Clinical Pharmacology and PKPD Modeling, Faculty of Pharmacy, National Autonomous University of Mexico, Mexico City, Mexico.

²Department of Tolerance and Disease Progression, Faculty of Pharmacy, Monterrey Institute of Technology, Monterrey, Mexico.

*Email: rafael.mendoza@unam.mx

ABSTRACT

Exposure–response analysis is often interpreted as though exposure, pharmacodynamic responsiveness, and disease state were stable coordinates. That assumption becomes fragile when drug exposure changes during treatment, effects lag behind plasma concentrations, responsiveness adapts, disease progresses, or treatment alters the processes that determine subsequent pharmacokinetics and outcome. This Perspective argues that the clinically relevant exposure–response relationship is therefore conditional on temporal state and history rather than being a fixed mapping between a concentration summary and an endpoint. Evidence from therapeutic antibodies, turnover and effect-compartment models, tolerance and tachyphylaxis analyses, disease-progression models, and longitudinal pharmacometrics is used to distinguish forward drug effects from feedback-compatible associations and model-dependent interpretations. We propose a state-conditioned longitudinal exposure–response framework that separates measured exposure, latent pharmacodynamic state, adaptive responsiveness, disease trajectory, and clinical outcome processes. The framework is intended as an interpretive discipline, not a validated universal model. Its value depends on identifiability, prospective or external validation, appropriate temporal sampling, and decision-relevant biomarkers. Dose optimization should therefore rely on static exposure targets only when stationarity is credible; when it is not, treatment decisions should be conditioned on validated evolving states.

Keywords: Exposure–response analysis, Pharmacokinetics, Pharmacodynamics, Disease progression, Tolerance, Model-informed precision dosing

INTRODUCTION

Exposure–response analysis is central to dose selection, therapeutic drug monitoring, and quantitative clinical pharmacology, but the meaning of an exposure target depends on what that target is expected to predict. Population concentration targets can be useful, yet they do not automatically define an optimal target for an individual patient when response biomarkers, toxicity, or disease state contain additional decision-relevant information [1]. Model-informed precision dosing likewise depends not only on estimating exposure but on the adequacy, validation, and implementation context of the models used to translate patient data into dosing decisions [2]. Quantitative models already contribute to drug-development decisions, including dose and trial design, but regulatory use of modeling does not make every model structure mechanistically correct or clinically transportable [3].

The interpretive problem addressed here arises when exposure and response are not stationary coordinates. Plasma concentration may change because clearance changes during therapy; the observed effect may lag behind

measured exposure; repeated treatment may attenuate or amplify responsiveness; disease may progress independently of treatment; and response itself may alter physiology, target burden, or other determinants of later pharmacokinetics. Under these conditions, a single area under the curve, trough concentration, or average exposure can mix information generated at different biological times.

We therefore propose a state-conditioned longitudinal exposure–response interpretation. The proposal does not require every PKPD model to contain every possible state. Rather, it asks whether the exposure metric, pharmacodynamic state, adaptive response state, and disease trajectory can reasonably be treated as invariant over the decision horizon. If they cannot, the relationship of interest becomes conditional on history. The subsequent sections separate the principal sources of nonstationarity before considering how they can be represented jointly and discriminated without mistaking model fit for mechanism.

Why static exposure–response relationships mislead

A static exposure summary can be misleading for two different reasons. First, the summary itself may contain temporally contaminated information. When an exposure metric is calculated using observations accrued after treatment initiation or near an event, the metric can become associated with the event process that it is later used to predict [4]. This does not invalidate static summaries in general; it means that their causal and predictive interpretation depends on how the exposure window relates to treatment, covariate evolution, and outcome timing. Second, the observed slope can reflect patient prognosis or model structure rather than pharmacological sensitivity. Simulation work in large-molecule oncology shows that omitted confounding and exposure saturation can generate biased or apparently positive exposure–outcome relationships under defined scenarios, while methodological reviews of monoclonal-antibody development identify disease severity and changing pharmacokinetics as recurrent sources of distortion [5, 6]. The relevant contrast is therefore not “static versus dynamic” as competing statistical fashions, but a correctly specified summary versus one that erases temporal ordering needed for the question being asked.

Nivolumab provides a clinically important illustration. Clearance decreased as disease status improved, so steady-state exposure could partly encode information about prior treatment response rather than only the forward effect of drug on outcome [7]. That finding does not establish a universal reverse-causal mechanism, nor does it exclude a true exposure effect. It shows that the same numerical exposure metric can change interpretation when the biological process generating that metric changes over treatment. For dose optimization, the key question is consequently not only whether exposure correlates with response, but whether exposure was measured before, during, or partly after the response-generating process.

Time-varying drug exposure

Therapeutic-antibody pharmacokinetics make nonstationarity especially visible. Time-dependent clearance has been reported recurrently across therapeutic-antibody programs and can confound exposure–response interpretation when changing clearance is ignored [8]. The mechanisms and magnitude are heterogeneous, so this evidence should not be generalized automatically to small molecules or treated as proof of one pathway. It does, however, establish that exposure may sometimes function as an evolving patient-state readout rather than as a fixed independent input.

The same caution applies to time-varying covariates. A postbaseline variable can be a confounder, a mediator of treatment effect, a consequence of response, or some combination across time. Adjusting for it without defining the estimand can remove part of the treatment effect one intends to estimate [9]. Covariate modeling therefore needs to be purpose-aware: pharmacometric covariates can improve prediction or explain variability without necessarily identifying a causal mechanism [10].

Clinical data reinforce this distinction. In a population pharmacokinetic analysis of avelumab, clearance declined over time and the magnitude of change varied across tumor types and response states [11]. The observation is compatible with co-evolution of disease state and pharmacokinetics, but it is not sufficient to establish a unique disease-to-clearance mechanism. Conversely, nonstationary exposure can originate from the drug itself. A semimechanistic TAS-114 model incorporated autoinduction, showing that repeated treatment can alter future exposure while pharmacodynamic inhibition follows a different trajectory [12]. Thus, time-varying exposure can arise from disease-linked physiology, treatment-induced disposition changes, target-related processes, or covariate evolution, and these sources are not inferentially interchangeable.

Because the same observed concentration trend can imply different mechanisms and decisions, **Table 1** separates the monitoring and validation conditions that determine when time-varying exposure changes the exposure–response question.

Table 1. When changing exposure becomes a different exposure–response question

Time-varying exposure scenario	Monitoring signal	Validation criterion	Clinical state to retain	Principal PK driver	State-conditioned framework endpoint
Response-associated clearance drift	Serial concentration profile together with disease-response measurements	Exposure metric remains predictive when response-related clearance evolution is represented	Tumor or disease trajectory	Time-dependent systemic clearance	Distinguish forward exposure effect from response-contaminated exposure
Postbaseline covariate evolution	Covariate trajectory aligned to dosing and outcome time	Covariate role matches the intended estimand and does not unintentionally remove mediated treatment effect	Treatment-modified patient state	Covariate-linked parameter change	Declare whether the covariate is predictor, mediator, consequence, or unresolved
Indication-dependent antibody clearance	Repeated PK across clinically distinct disease contexts	Model retains performance across the disease states for which inference is claimed	Tumor type and response status	Disease-associated disposition change	Restrict exposure–response interpretation to validated clinical states
Treatment-driven autoinduction	Repeated dose and concentration measurements over the induction period	Dynamic disposition model reproduces evolving PK without forcing pharmacodynamic change into clearance	Enzyme or metabolic adaptation state	Induction of elimination capacity	Separate changing exposure generation from changing pharmacodynamic sensitivity
Apparently stable summary exposure	Early and later exposure summaries examined against event timing	Summary is not materially altered by information accrued after response or event development	Outcome-risk state	Metric construction rather than biological PK change	Accept a static target only when temporal ordering remains defensible

Note: Entries synthesize the evidence into a proposed decision-oriented classification; they are not a validated scoring system, causal hierarchy, or universal dosing rule.

Pharmacodynamic delay and hysteresis

Even if plasma exposure were stationary, measured concentration and measured response need not be synchronized. Omalizumab modeling in chronic spontaneous urticaria required target turnover, an effect-compartment component, and placebo and time-course components to reproduce delayed biomarker and clinical responses [13]. Such structures show why plasma concentration cannot automatically be treated as the instantaneous driver of an observed endpoint, but a fitted effect compartment should not be reified as a directly measured anatomical biophase.

Hysteresis makes the same problem visible geometrically. On the ascending and descending limbs of a concentration–time profile, the same measured concentration may correspond to different responses when equilibration or downstream turnover introduces delay [14]. Importantly, a hysteresis loop identifies temporal displacement, not its unique cause. Active metabolites, tolerance, sensitization, disease evolution, or measurement timing can produce similar concentration–effect separation.

Patient and interaction states can also change the temporal mapping. In a clopidogrel proof-of-concept study, CYP2C19 phenotype altered active-metabolite exposure and its relationship to platelet reactivity within a turnover PKPD model [15]. The resulting dose simulations were not prospective prescribing validation, but they show that genotype can condition both exposure generation and downstream response kinetics. A physiologically based PKPD analysis of S-warfarin similarly combined CYP2C9-related disposition, fluconazole interaction, genetic variability, and delayed INR response [16]. These examples support a general interpretive rule: the relevant exposure for a delayed endpoint is the exposure history that drives the response process, not necessarily the concentration measured at the moment the endpoint is observed.

Tolerance and sensitization

Repeated exposure can change the response system itself. A population pharmacodynamic model of sarilumab-associated neutrophil changes used margination and a tolerance component to reproduce a transient decrease followed by recovery [17]. The model supports an adaptive-state representation, but it does not uniquely prove the underlying cellular mechanism. For exposure–response interpretation, the important point is that identical concentrations at different treatment times may act on different effective response states.

Abrocitinib platelet modeling provides an even clearer boundary between useful structure and mechanistic certainty. A tachyphylaxis component improved the description of platelet time courses, yet the analysis could not distinguish that time-dependent attenuation from underlying feedback with the available data [18]. A parameter that captures diminishing effect can therefore be useful for prediction while remaining mechanistically nonidentifiable.

Sensitization represents the opposite direction of state change. Under inflammatory conditions, experimental work in mice and dorsal-root-ganglion systems showed NGF-dependent enhancement of receptor-linked TRPV1 responsiveness and pain behavior [19]. This evidence is mechanistically informative but preclinical; it should not be used to claim that human therapeutic sensitization follows the same pathway. Taken together, tolerance, tachyphylaxis, feedback, and sensitization are better treated as competing state-dependent explanations for changing response gain than as interchangeable labels for any nonstationary concentration–effect curve.

Disease progression

Disease progression changes the baseline against which drug response is interpreted. Disease-progression modeling can separate natural history, placebo or standard-of-care effects, and drug-related components, but its utility depends on data maturity, model qualification, and the intended development question [20]. A static endpoint can therefore misattribute expected disease evolution to treatment failure or, conversely, attribute spontaneous or background change to drug effect.

Long-horizon disease time may itself be latent. In Parkinson disease, a data-driven model reconstructed a much longer disease trajectory from substantially shorter individual observation windows [21]. The resulting trajectory is model-derived, not a record of decades of direct follow-up. In Alzheimer disease, a longitudinal CDR-SB model was developed using one dataset and tested with separate clinical-trial placebo data, illustrating the value of external evaluation when a progression model is intended to transport beyond its construction sample [22]. External validation strengthens confidence within the evaluated context but does not establish universal transferability.

The analytical distinction is therefore between what is observed now and what is inferred about the disease process that would have occurred, or will occur, over the relevant decision horizon. Because dose interpretation changes differently when background progression, latent disease time, or treatment-responsive disease dynamics dominate, **Table 2** separates what is observed from what must be inferred before dose consequences are assigned.

Table 2. Disease trajectories that can masquerade as changing pharmacodynamic sensitivity

Time-varying PKPD state	Governing disease condition	What is directly observed	What is inferred	Discriminating condition	Dose implication	State-conditioned framework endpoint
Background progression dominates	Endpoint changes despite stable treatment and exposure	Serial clinical or biomarker measurements	Natural-history contribution	Progression model predicts trajectory independently of dose change	Avoid reflexive escalation based solely on worsening endpoint	Separate disease drift from loss of drug effect
Latent disease time differs across patients	Similar chronological follow-up begins at different biological stages	Fragmented longitudinal measurements	Position on an underlying disease trajectory	Reconstructed disease-time model remains stable under alternative plausible structures	Interpret equal exposure against disease stage rather than visit number alone	Condition response on inferred disease state

Progression model transports incompletely	Model fits development data but performs differently in a new cohort	Prediction residuals in external data	Population or state mismatch	Independent evaluation identifies where trajectory assumptions fail	Do not carry dose–response conclusions across unvalidated populations	Mark transportability as a validation boundary
Treatment modifies disease trajectory	Disease markers change after therapy and later PK also changes	Joint PK and disease–response time courses	Feedback-compatible coupling	Temporal model distinguishes earlier drug exposure from later disease-mediated PK change	Prefer temporally ordered exposure metrics	Preserve directionality among treatment, disease, and subsequent PK
Endpoint contains mixed components	Clinical change reflects disease, placebo, and pharmacodynamic processes	Composite longitudinal response	Relative contribution of latent components	Competing models remain identifiable with available sampling	Use the simplest identifiable decomposition compatible with the decision	Prevent disease progression from being absorbed into a generic PD term

Note: Entries are a proposed synthesis of disease-progression and longitudinal PKPD evidence; they neither assign causal mechanisms nor constitute validated dose algorithms.

Figure 1 is needed to preserve the parallel evolution of exposure, disease trajectory, and response while making the interpretation and validation gates visible rather than burying them in sequential prose.




Figure 1. Disease-state drift can redirect the inferential path from exposure to observed response

Treatment–disease feedback

The strongest reason to avoid a one-way exposure-to-response interpretation is that treatment and disease can co-evolve. In nivolumab, response-associated changes in clearance meant that later exposure partly reflected the evolving clinical state [7]. Across therapeutic antibodies, time-dependent clearance is sufficiently recurrent to require explicit consideration when late exposure metrics are linked to outcome [8]. Avelumab provides another direct clinical example in which clearance trajectories varied with tumor and response context [11]. These observations support feedback-compatible coupling, but they do not prove that disease improvement directly causes every change in clearance or that the same mechanism applies across modalities.

The covariate problem is therefore causal as well as statistical. A postbaseline disease marker may lie on the pathway from treatment to outcome, may influence later PK, or may itself be affected by earlier response. Adjustment can then answer a different question from the total treatment effect and may remove mediated information [9]. Similarly, delayed clinical response can contain drug, placebo, and time-course components that should not automatically be merged into one concentration–effect term [13]. Disease-progression work reinforces the need to preserve natural-history and treatment components as analytically distinct when the data can support that separation [20].

The proposed framework therefore treats treatment–disease feedback as a hypothesis-bearing architecture rather than a default causal loop. Earlier exposure may influence disease response; evolving disease may alter target burden, catabolic state, organ function, or other determinants of later disposition; and both processes may affect the clinical endpoint. Each directional relation requires its own temporal evidence and validation. **Figure 2** makes these directional and inferential boundaries explicit so that late exposure cannot be read automatically as an independent upstream predictor.




Figure 2. When later exposure carries information generated by earlier treatment response

Caption. The figure separates early-treatment and later-treatment states around a central treatment–disease feedback field. Earlier exposure can drive pharmacodynamic change, whereas evolving disease or target state can be associated with subsequent changes in disposition. Solid arrows indicate evidence-supported or manuscript-asserted directional relationships; dashed arrows indicate proposed or conditional links; bidirectional connectors identify feedback-compatible coupling without claiming mechanistic reciprocity. A target-mediated-disposition region is retained only where such biology is relevant and is visually separated from broader disease-associated clearance change. Uncertainty corridors retain alternative explanations including prognostic confounding, postbaseline mediation, adaptive pharmacodynamics, and unmeasured disposition changes. Validation gates require temporal precedence, mechanistic plausibility where mechanism is claimed, and model discrimination before later exposure is treated as an independent predictor of treatment benefit.

Alt text. Wide feedback architecture read from left to right. The left side contains early dose and exposure linked to an early pharmacodynamic state and subsequent disease response. The center is a large feedback field where evolving disease or target state can connect to later clearance and later exposure. A separate target-mediated-disposition module is nested only within the relevant biologic-drug pathway rather than being applied to all drugs. Parallel uncertainty corridors show prognostic confounding, mediation, adaptation, and unmeasured PK change as competing explanations. The right side contains clinical-outcome interpretation behind validation gates. Solid, dashed, and bidirectional lines distinguish supported direction, proposed conditionality, and feedback-compatible coupling, with no quantitative claims.

Joint longitudinal PKPD representation

The preceding sections imply that a clinically useful exposure–response model may need to represent several trajectories without assuming that they are interchangeable. Joint longitudinal–event modeling provides one practical foundation because repeated biomarker measurements and time-to-event outcomes can be analyzed within a linked framework rather than as independent snapshots [23]. Such linkage can improve temporal coherence, but statistical dependence between submodels does not establish a causal pathway.

A clinical example is provided by amcenestrant, for which tumor dynamics and progression-free survival were modeled jointly in early-phase advanced breast cancer development [24]. The example demonstrates that exposure, longitudinal disease behavior, delay or resistance-related dynamics, and event risk can coexist within one quantitative representation. It does not establish that the same structural model transports to other tumors, mechanisms, or decision horizons.

We therefore propose a state-conditioned longitudinal PKPD representation with five separable elements: observed exposure history, latent or delayed pharmacodynamic state, adaptive response state, disease trajectory, and clinical–event process. These elements should be included only when the data and decision problem justify them. The framework is deliberately modular: a stationary exposure–response model remains appropriate when temporal invariance is credible, whereas additional states are introduced only when ignoring them changes interpretation. Observed variables should remain visually and mathematically distinct from latent states, and every proposed feedback relation should be treated as a testable structural alternative rather than as assumed biology.

A practical consequence is that sampling design becomes part of the scientific interpretation. Sparse concentrations collected at one phase of treatment may identify average exposure adequately while providing little information about whether clearance, effect delay, or disease state changed between visits. Likewise, an endpoint measured infrequently can make an adaptive pharmacodynamic state appear indistinguishable from slow disease evolution. The proposed representation therefore treats observation schedules as constraints on what state transitions can legitimately be inferred. It should not create latent compartments merely because a more complex model fits; a state earns interpretive status only when its inclusion corresponds to a plausible process, improves the decision-relevant representation, and remains estimable with the available information.

Identifiability and model discrimination

A richer longitudinal model can reduce interpretive bias only if its components are distinguishable. Structural or practical nonidentifiability can make different parameterizations or mechanisms reproduce similar observations, so identifiability should be examined before fitted parameters are interpreted mechanistically [25]. This issue is central when delayed effect, tolerance, disease progression, and feedback can each generate superficially similar concentration–response patterns.

Model diagnostics also require care. Derived outcomes may show predictive discrepancies even when the model describes the directly modeled quantities adequately, and such divergence should be understood rather than treated automatically as proof of model failure [26]. Conversely, a good global fit cannot rescue an unidentifiable mechanism. The abrocitinib example is instructive because tachyphylaxis improved the platelet model while remaining difficult to distinguish from underlying feedback [18].

We therefore propose an exception-first discrimination rule. Competing temporal explanations should be challenged against the observation that most directly separates them: temporal precedence for feedback, rising-versus-falling concentration behavior for delay, repeated-exposure attenuation for tolerance, independent disease measurements for progression, and external or prospective data for transportability. Failure to discriminate should leave the mechanism unresolved rather than being converted into biological certainty.

Discrimination should also be asymmetric: evidence capable of falsifying an attractive mechanism is often more useful than another diagnostic confirming overall fit. A tolerance model should be challenged by measurements related to the hypothesized adaptive process where available; a disease–feedback explanation should be challenged by temporally earlier exposure metrics and independent disease-state observations; and a transport claim should be challenged outside the population used to build the model. When those data do not exist, the scientifically appropriate output is bounded uncertainty. Parsimony remains important because adding delay, adaptation, progression, and feedback simultaneously can convert a transparent uncertainty into an apparently precise but weakly identifiable state system.

Because the main danger is not merely poor fit but a convincing fit produced for the wrong reason, **Table 3** organizes the principal failure patterns around the corrective test needed before an exposure–response interpretation is carried into dosing.

Table 3. Failure patterns that can make a longitudinal exposure–response model look more certain than it is

Misleading pattern	Likely origin	Observable signature	Competing diagnosis	Corrective test	Residual consequence after mitigation
Stable fit with unstable parameters	Structural or practical nonidentifiability	Similar predictions from materially different parameter sets	Overparameterization or insufficient sampling	Identifiability analysis plus sensitivity to sampling design	Mechanistic interpretation remains limited even if prediction is adequate
Delayed response attributed to one latent compartment	Effect-site delay, turnover, active metabolite, or disease-time shift	Response lags measured plasma concentration	Adaptive response or evolving disease	Compare alternative temporal structures against informative sampling phases	Delay may be predictable without being uniquely mechanistic
Diminishing effect labeled as tolerance	Adaptive attenuation represented phenomenologically	Similar exposure produces smaller later response	Homeostatic feedback or disease evolution	Seek state measurements that separate adaptation from feedback	Retain a descriptive adaptive state if mechanisms remain unresolved
Late exposure predicts favorable outcome unusually strongly	Disease-linked clearance change or prognostic confounding	Exposure association strengthens after response has evolved	True forward exposure effect	Reanalyze with temporally ordered exposure and disease-state information	Late exposure may be unsuitable as an independent causal predictor
Progression component absorbs drug effect	Inadequate separation of natural history and treatment response	Similar trajectory under alternative disease/drug decompositions	Misspecified baseline progression	External disease data or alternative progression structures	Dose inference remains conditional on disease-model validity
Derived endpoint fails a predictive check while direct data fit	Transformation-sensitive diagnostics	Disagreement appears mainly after outcome derivation	Underlying model misspecification	Diagnose direct and derived quantities separately	Do not discard or accept the model on one diagnostic alone

Note: The table is a proposed triage framework, not a validated diagnostic hierarchy or pass–fail standard.

Consequences for dose optimization

If exposure–response meaning changes with state, dose optimization becomes a sequential inference problem. Model-informed reinforcement-learning approaches illustrate how dosing can be formulated around an evolving state, where the next action depends on accumulated information rather than on a fixed concentration target alone [27]. This establishes methodological feasibility, not clinical superiority, and implementation requires trustworthy state estimation, constrained decision rules, and prospective evaluation.

Contemporary model-informed precision-dosing work similarly emphasizes model qualification, appropriate biomarkers, and clinically meaningful endpoints [28]. A model that accurately predicts concentration but misrepresents delayed response, adaptation, or disease trajectory may still be inadequate for the decision being made. Population transport is another boundary: a pediatric vancomycin case showed how local data accumulation and model updating can address mismatch between an existing population model and a new clinical setting [29]. That strategy should not be interpreted as proof that continuous learning universally improves patient outcomes. The clinical decision horizon also matters: a short-term safety target may remain valid while long-term disease progression changes, so state-conditioned dosing should modify only decisions for which longitudinal information has been validated.

The proposed dosing consequence is therefore conditional. When exposure generation and response sensitivity are sufficiently stationary over the decision horizon, a conventional validated target may remain entirely appropriate. When they are not, future doses should be evaluated against the best validated representation of the future response state rather than an unqualified static exposure target. Clinically meaningful target definition, model validation, and decision-relevant response information remain joint prerequisites for that transition [1, 2, 28]. The framework does not prescribe a universal adaptive algorithm; it specifies when the assumption behind a static target requires explicit testing.

CONCLUSION

Exposure–response relationships are not necessarily fixed properties of a drug. Their interpretation can change when exposure evolves during treatment, pharmacodynamic effects are delayed, responsiveness adapts, disease progresses, or treatment and disease become statistically coupled over time. The proposed state-conditioned longitudinal framework separates these processes so that association is not mistaken for mechanism and model output is not mistaken for validated clinical utility. Its central practical implication is selective rather than universal: retain simple exposure targets when temporal stationarity is defensible, but require longitudinal state representation, discrimination of competing mechanisms, and appropriate validation when it is not.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Polasek TM, Peck RW. Beyond population-level targets for drug concentrations: Precision dosing needs individual-level targets that include superior biomarkers of drug responses. *Clin Pharmacol Ther.* 2024;116(3):602-12. doi:10.1002/cpt.3197
2. 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
3. Wang Y, Zhu H, Madabushi R, Liu Q, Huang SM, Zineh I. Model-informed drug development: Current US regulatory practice and future considerations. *Clin Pharmacol Ther.* 2019;105(4):899-911. doi:10.1002/cpt.1363
4. Wiens MR, French JL, Rogers JA. Confounded exposure metrics. *CPT Pharmacometrics Syst Pharmacol.* 2024;13(2):187-91. doi:10.1002/psp4.13074
5. Poon V, Lu D. Performance of Cox proportional hazard models on recovering the ground truth of confounded exposure-response relationships for large-molecule oncology drugs. *CPT Pharmacometrics Syst Pharmacol.* 2022;11(11):1511-26. doi:10.1002/psp4.12859
6. Kawakatsu S, Bruno R, Kågedal M, Li C, Girish S, Joshi A, et al. Confounding factors in exposure-response analyses and mitigation strategies for monoclonal antibodies in oncology. *Br J Clin Pharmacol.* 2021;87(6):2493-501. doi:10.1111/bcp.14662
7. Liu C, Yu J, Li H, Liu J, Xu Y, Song P, et al. Association of time-varying clearance of nivolumab with disease dynamics and its implications on exposure response analysis. *Clin Pharmacol Ther.* 2017;101(5):657-66. doi:10.1002/cpt.656
8. Proctor JR, Wong H. Time-dependent clearance can confound exposure-response analysis of therapeutic antibodies: A comprehensive review of the current literature. *Clin Transl Sci.* 2024;17(1):e13676. doi:10.1111/cts.13676
9. 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
10. Sanghavi K, Ribbing J, Rogers JA, Ahmed MA, Karlsson MO, Holford N, et al. Covariate modeling in pharmacometrics: General points for consideration. *CPT Pharmacometrics Syst Pharmacol.* 2024;13(5):710-28. doi:10.1002/psp4.13115
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. Araki H, Takenaka T, Takahashi K, Yamashita F, Matsuoka K, Yoshisue K, et al. A semimechanistic population pharmacokinetic and pharmacodynamic model incorporating autoinduction for the dose

- justification of TAS-114. CPT Pharmacometrics Syst Pharmacol. 2022;11(5):604-15. doi:10.1002/psp4.12747
13. Oh E, Wada R, Le K, Zheng Y, Jin J, Poon V, et al. PK/PD modeling to characterize placebo and treatment effect of omalizumab for chronic spontaneous urticaria. CPT Pharmacometrics Syst Pharmacol. 2023;12(6):795-807. doi:10.1002/psp4.12953
 14. Standing JF. Understanding and applying pharmacometric modelling and simulation in clinical practice and research. Br J Clin Pharmacol. 2017;83(2):247-54. doi:10.1111/bcp.13119
 15. Jung YS, Jin BH, Park MS, Kim CO, Chae D. Population pharmacokinetic-pharmacodynamic modeling of clopidogrel for dose regimen optimization based on CYP2C19 phenotypes: A proof of concept study. CPT Pharmacometrics Syst Pharmacol. 2024;13(1):29-40. doi:10.1002/psp4.13053
 16. Geng K, Shen C, Wang X, Wang X, Shao W, Wang W, et al. A physiologically-based pharmacokinetic/pharmacodynamic modeling approach for drug-drug-gene interaction evaluation of S-warfarin with fluconazole. CPT Pharmacometrics Syst Pharmacol. 2024;13(5):853-69. doi:10.1002/psp4.13123
 17. Kovalenko P, Paccaly A, Boyapati A, Xu C, St John G, Nivens MC, et al. Population pharmacodynamic model of neutrophil margination and tolerance to describe effect of sarilumab on absolute neutrophil count in patients with rheumatoid arthritis. CPT Pharmacometrics Syst Pharmacol. 2020;9(7):405-16. doi:10.1002/psp4.12534
 18. Wojciechowski J, Malhotra BK, Wang X, Fostvedt L, Valdez H, Nicholas T. Population pharmacokinetic-pharmacodynamic modelling of platelet time-courses following administration of abrocitinib. Br J Clin Pharmacol. 2022;88(8):3856-71. doi:10.1111/bcp.15334
 19. Masuoka T, Yamashita Y, Yoshida J, Nakano K, Tawa M, Nishio M, et al. Sensitization of glutamate receptor-mediated pain behaviour via nerve growth factor-dependent phosphorylation of transient receptor potential V1 under inflammatory conditions. Br J Pharmacol. 2020;177(18):4223-41. doi:10.1111/bph.15176
 20. Goteti K, Hanan N, Magee M, Wojciechowski J, Mensing S, Lalovic B, et al. Opportunities and challenges of disease progression modeling in drug development—an IQ perspective. Clin Pharmacol Ther. 2023;114(2):266-74. doi:10.1002/cpt.2873
 21. Jin R, Yoshioka H, Sato H, Hisaka A. Data-driven disease progression model of Parkinson's disease and effect of sex and genetic variants. CPT Pharmacometrics Syst Pharmacol. 2024;13(4):649-59. doi:10.1002/psp4.13112
 22. Jamalian S, Dolton M, Chanu P, Ramakrishnan V, Franco Y, Wildsmith K, et al. Modeling Alzheimer's disease progression utilizing clinical trial and ADNI data to predict longitudinal trajectory of CDR-SB. CPT Pharmacometrics Syst Pharmacol. 2023;12(7):1029-42. doi:10.1002/psp4.12974
 23. Zhudnikov K, Gavrilov S, Sofronova A, Stepanov O, Kudryashova N, Helmlinger G, et al. A workflow for the joint modeling of longitudinal and event data in the development of therapeutics: Tools, statistical methods, and diagnostics. CPT Pharmacometrics Syst Pharmacol. 2022;11(4):425-37. doi:10.1002/psp4.12763
 24. Cerou M, Thai HT, Deyme L, Fliscounakis-Huynh S, Comets E, Cohen P, et al. Joint modeling of tumor dynamics and progression-free survival in advanced breast cancer: Leveraging data from amcenestrantr early phase I-II trials. CPT Pharmacometrics Syst Pharmacol. 2024;13(6):941-53. doi:10.1002/psp4.13128
 25. van Noort M, Ruppert M, DeJongh J, Marostica E, Bosch R, Mešić E, et al. Navigating the landscape of parameter identifiability methods: A workflow recommendation for model development. CPT Pharmacometrics Syst Pharmacol. 2024;13(7):1170-9. doi:10.1002/psp4.13148
 26. Prybylski JP. Indirect modeling of derived outcomes: Are minor prediction discrepancies a cause for concern? CPT Pharmacometrics Syst Pharmacol. 2024;13(10):1762-70. doi:10.1002/psp4.13219
 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. 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
 29. 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