Available online **www.ijpras.com**

International Journal of Pharmaceutical Research & Allied Sciences, 2025, 14(3):103-112
https://doi.org/10.51847/ld93VGQWjn

Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Reconstructing the Causal Dose–Exposure–Response Chain under Adherence Variation, Disease Progression, Concomitant Therapy, and Measurement Error

Mateo Alvarez[1]*, Sofia Herrera[2], Diego Cruz[1], Andres Castro[3]

[1] Department of Causal Dose–Exposure–Response Modeling, Faculty of Pharmacy, National University of Colombia, Bogota, Colombia.
[2] Department of Adherence and Disease Progression, Faculty of Pharmacy, University of Chile, Santiago, Chile.
[3] Department of Measurement Error and Concomitant Therapy, Faculty of Pharmacy, University of Antioquia, Medellin, Colombia.

***Email:** mateo.alvarez@unal.edu.co*

ABSTRACT

Observed exposure–response associations are often interpreted as evidence that changing dose or exposure will alter clinical outcome. Such interpretation may fail when prescribed dose differs from administered dose, adherence changes over time, disease severity modifies pharmacokinetic disposition, prior response influences subsequent treatment, concomitant medicines alter exposure or pharmacodynamic state, or measurement and sampling processes determine which observations enter the analysis. This Original Causal PK/PD Architecture Article proposes the Causal Dose–Exposure–Response Reconstruction Architecture, an original conceptual synthesis that separates treatment intention, realized administration, absorbed drug input, latent systemic exposure, measured concentration, pharmacodynamic state, latent biological response, and observed outcome. The architecture represents disease progression, time-varying physiology, clinical dose modification, and concomitant therapy as temporally ordered processes whose causal roles depend on the estimand and treatment history. It also distinguishes measurement error, informative sampling, censoring, and missingness from biological variability. The contribution provides a causal graph, construct definitions, testable propositions, failure modes, and staged identification, estimation, and validation requirements. Its central proposition is that no single exposure metric, adherence measure, pharmacokinetic model, pharmacodynamic model, causal estimator, or digital method can independently reconstruct the chain. Validation must separately address dose-input accuracy, structural assumptions, latent-state recovery, observation models, external transportability, causal identification, and decision consequences. The architecture may support study design, assumption auditing, model criticism, and validation planning, but it does not constitute an empirically validated causal model, clinical dosing recommendation, regulatory standard, digital twin, or deployment-ready decision system.

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

Received: 15 May 2025; **Revised:** 04 September 2025; **Accepted:** 11 September 2025

INTRODUCTION

Dose–exposure–response reasoning connects a therapeutic intervention to drug disposition, pharmacological action, and clinical outcome. In its simplest form, the chain appears sequential: a dose generates exposure, exposure produces pharmacodynamic activity, and activity changes an endpoint. In clinical data, however, each

component may be incompletely observed, dynamically modified, or affected by processes that also influence subsequent treatment and outcome.

Model-informed precision dosing has established the value of integrating pharmacological models with patient-level information, while also demonstrating that model adequacy, validation, implementation infrastructure, and clinical context are inseparable from decision use [1-3]. The persistent challenge is therefore not only predicting concentration or selecting a nominal dose. It is determining what treatment was administered, which biological process generated the exposure, what the exposure represents causally, and whether the measured response can be attributed to an intervention rather than to evolving patient state.

Prevailing exposure–response analyses frequently compress these questions into a small number of variables. Prescription may substitute for administration, a summary exposure metric may replace a longitudinal trajectory, baseline covariates may stand in for evolving physiology, and observed endpoints may be treated as direct measurements of latent pharmacological response. This compression can preserve prediction while obscuring the causal pathways that produced the prediction.

This article develops the proposed Causal Dose–Exposure–Response Reconstruction Architecture, or CDERRA. As an original non-empirical PK/PD architecture, it integrates adherence, time-varying disposition, delayed pharmacodynamics, disease progression, concomitant therapy, treatment-confounder feedback, and observation processes within one temporally ordered framework. The contribution defines constructs and relationships, derives testable propositions, specifies validation requirements, and establishes decision boundaries without claiming empirical validation or universal applicability.

Why observed exposure–response relationships may be noncausal

An observed association between higher exposure and improved outcome need not represent the effect of increasing exposure. In therapeutic-antibody development, disease burden, cachexia, treatment response, and time-varying clearance may jointly influence measured exposure and outcome, creating associations in which exposure partly reflects patient condition rather than an independently manipulable pharmacological cause [4-6]. The resulting gradient may therefore combine treatment effect, prognostic structure, and reverse causation.

The problem is intensified when exposure is summarized after treatment initiation. Average concentration, cumulative exposure, or model-derived clearance may incorporate information generated by earlier response, toxicity, treatment interruption, dose modification, or disease evolution. Such metrics can remain predictive while failing to represent a well-defined intervention. Prediction asks whether an observed quantity forecasts an endpoint; causal interpretation asks what would happen if a specified treatment process were changed while the relevant competing pathways were controlled.

Confounding may also arise before treatment. Disease severity, organ function, body composition, inflammation, genotype, comorbidity, access, and clinician judgement can affect dose selection, adherence, exposure, sampling, and outcome. After treatment begins, several of these variables may become consequences of prior therapy while continuing to influence later therapy and response. Their status therefore cannot be determined by variable name alone; it depends on timing, treatment history, and the causal question.

CDERRA does not assume that exposure–response evidence is generally invalid. Rather, it requires the observed relationship to be decomposed into treatment input, biological disposition, pharmacodynamic mechanism, disease evolution, and observation processes. Evidence of association remains scientifically useful, but it does not by itself identify the effect of intervening on dose, adherence, concentration, target engagement, or another component of the chain.

Proposed causal dose–exposure–response graph

Causal pharmacometric analysis requires an explicit intervention, temporally ordered variables, and assumptions connecting the observed data to the counterfactual contrast of interest. Causal graphs, examination of confounded exposure metrics, and analyses of time-dependent confounding show why longitudinal exposure–response questions cannot be resolved through an exposure summary alone [7-9]. CDERRA therefore separates treatment intention, realized administration, latent biological states, measured proxies, and clinical observation processes.

Figure 1 presents the original conceptual synthesis for causal graph for the dose–exposure–response chain, showing how the article's principal components, evidence relationships, uncertainties, and decision boundaries are connected.

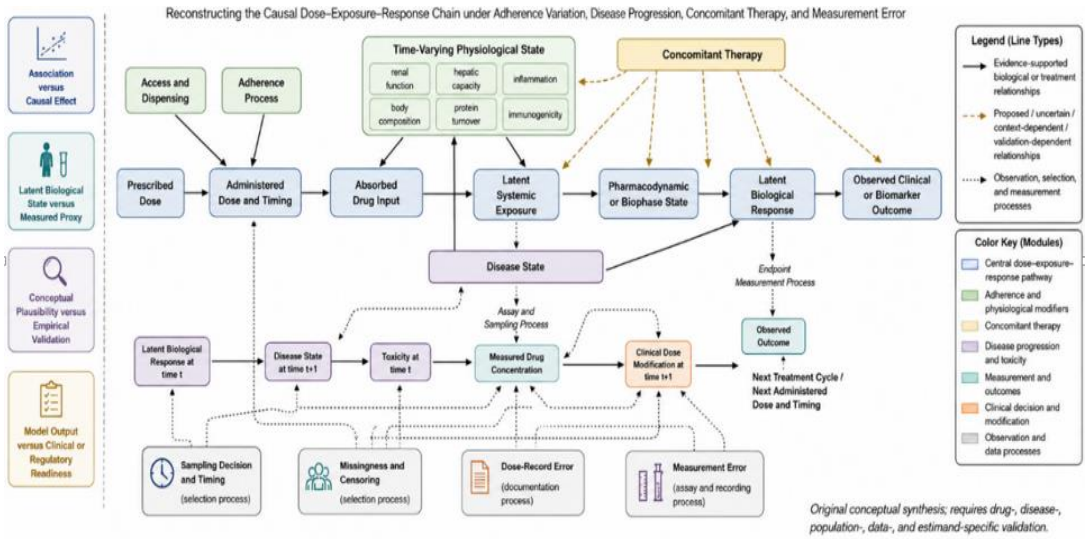

Figure 1. Causal Graph for the Dose–Exposure–Response Chain.

The figure is an original conceptual synthesis developed for "Reconstructing the Causal Dose–Exposure–Response Chain under Adherence Variation, Disease Progression, Concomitant Therapy, and Measurement Error." 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 graph's principal treatment pathway begins with prescribed dose but does not equate prescription with intervention delivery. Prescribed dose represents clinical intention; administered dose represents the realized treatment event; absorbed input represents the fraction and timing entering the pharmacokinetic system. Latent systemic exposure is the underlying concentration trajectory, whereas measured concentration is an error-prone and selectively observed representation of that trajectory.

The response pathway is similarly separated. A pharmacodynamic state represents mechanism-proximal processes such as receptor binding, target occupancy, signalling, turnover, or adaptation. Latent biological response represents the resulting disease or physiological change, while the observed endpoint is generated through measurement, assessment, and documentation. These distinctions are proposed to prevent convenient measurements from being interpreted automatically as direct biological states.

Disease state, physiology, concomitant therapy, and adherence occupy context-dependent causal positions. A physiological measure may precede exposure, result from earlier treatment, mediate response, or influence later dose selection. CDERRA therefore represents feedback through successive time points rather than through a contemporaneous cycle. The same variable may occupy different nodes across time without acquiring a universal status as confounder, mediator, or outcome.

The graph is not an adjustment template. Conditioning on every displayed node could block part of the treatment effect, open selection pathways, or adjust for consequences of earlier treatment. Its purpose is to force explicit definition of the intervention, estimand, temporal ordering, measurement process, and unobserved assumptions before an estimator is selected. Each application requires a drug-, disease-, population-, and question-specific graph.

Dose administration and adherence processes

Adherence is proposed as a longitudinal treatment-input process rather than a stable patient attribute. It includes treatment initiation, access, timing, dose omission, interruption, compensatory administration, persistence, and discontinuation. A population pharmacokinetic adherence framework demonstrates that realized administration patterns can materially shape inferred exposure and treatment assessment [10]. Consequently, prescription alone cannot define the delivered intervention.

Dose histories may be partially observed through administration records, dispensing data, electronic monitoring, self-report, pill counts, infusion documentation, or concentration patterns. These sources measure different events and contain different errors. When dosing history is missing, population pharmacokinetic estimates may be biased

unless the uncertainty is explicitly reconstructed or integrated into estimation [11]. CDERRA therefore treats dose-record reliability as a distinct observation process.

The pharmacological consequence of nonadherence is also mechanism-dependent. Drug forgiveness reflects the capacity of a regimen to preserve pharmacological effect despite delayed or missed administration, and it depends jointly on pharmacokinetic persistence, pharmacodynamic dynamics, dosing interval, and the relevant response criterion [12]. The same adherence pattern may therefore have different consequences across drugs, formulations, disease states, and therapeutic objectives.

The proposed dose-input reconstruction layer should represent each administration event as observed, uncertain, or missing and should propagate that uncertainty into exposure and response estimates. It may combine independent evidence streams, but it cannot recover an event that leaves no informative trace without assumptions. A reconstructed history is therefore an estimated treatment process, not verified ground truth or a sufficient basis for clinical dose adjustment.

Pharmacokinetic disposition and time-varying physiology

Pharmacokinetic disposition should not automatically be represented by fixed individual parameters. Clearance, distribution, absorption, and bioavailability may change during treatment because physiology, disease burden, inflammation, body composition, immunogenicity, organ function, or concomitant interventions evolve. Time-varying pharmacokinetic models of therapeutic monoclonal antibodies illustrate the importance of representing such changes explicitly rather than assigning all longitudinal variability to residual error [13].

Disease dynamics may themselves modify disposition. In nivolumab treatment, time-varying clearance was associated with evolving disease characteristics, demonstrating that exposure may partly reflect the clinical trajectory that it is subsequently used to explain [14]. CDERRA therefore treats physiological state as a temporally indexed latent or measured construct that can influence exposure while also being influenced by prior treatment and disease evolution.

Inflammation and drug interactions may jointly alter metabolic capacity rather than acting as independent baseline covariates. Physiologically based pharmacokinetic modelling has shown how inflammatory suppression and perpetrator-drug effects may be integrated when predicting CYP3A- and CYP2C19-mediated disposition [15]. Such modelling supports mechanistic plausibility but does not establish transportability to unstudied populations, diseases, or treatment combinations.

The proposed architecture consequently distinguishes baseline determinants from treatment-induced physiological changes. Renal function, hepatic capacity, inflammatory activity, protein turnover, and disease burden may operate as confounders, mediators, consequences, or effect modifiers depending on their timing. Their causal status must be specified for each estimand rather than inferred from clinical terminology alone.

Table 1 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for definitions, components, relationships, and intended uses in reconstructing the causal dose–exposure–response chain under adherence variation.

Table 1. Definitions, components, relationships, and intended uses in Reconstructing the Causal Dose–Exposure–Response Chain under Adherence Variation.

Component or scientific dimension	Problem addressed	Inputs or determinants	Proposed mechanism or relationship	Expected contribution	Evidence required	Failure or uncertainty risk	Boundary statement
Prescribed versus administered dose	Prescription may not represent delivered treatment	Clinical order, access, timing, persistence	**Proposed:** separate intended treatment from realized administration	Improves definition of the intervention	Prospective administration records and adherence evidence [10]	Misclassification of treatment input	Prescription is not proof of administration
Uncertain dose history	Dose timing or amount may be missing	Records, self-report, dispensing, concentrations	**Proposed:** represent dose events probabilistically	Propagates administration uncertainty into PK estimates	Dose-reconstruction assessment [11]	False precision or non-identifiability	Reconstruction is not verified ground truth

Drug forgiveness	Similar adherence patterns may have different effects	Half-life, PD persistence, interval, response dynamics	Missed-dose consequences depend on integrated PK/PD properties	Links adherence to clinically relevant persistence	Mechanistic and regimen-specific evaluation [12]	Unsupported universal thresholds	Forgiveness is drug- and objective-specific
Time-varying disposition	Fixed PK parameters may obscure longitudinal change	Organ function, disease state, immunogenicity	Clearance and distribution may evolve during treatment	Improves temporal exposure representation	Repeated PK and covariate measurements [13]	Confusing parameter drift with biological change	Time variation does not establish its cause
Disease-linked clearance	Disease state may affect both exposure and outcome	Disease burden, cachexia, response	Disease evolution modifies clearance and apparent exposure–response	Supports reverse-causation assessment	Longitudinal disease and PK information [14]	Attributing prognosis to exposure	Association does not identify an exposure intervention effect
Inflammation and interaction state	Physiological and drug effects may combine	Inflammation, genotype, co-medication	Joint modulation of metabolic capacity	Supports mechanism-based interaction hypotheses	Context-specific PBPK qualification [15]	Extrapolation beyond validated conditions	A plausible pathway is not a dosing rule

Pharmacodynamic mechanisms and delayed outcomes

Systemic concentration is rarely identical to the mechanism producing response. Turnover, receptor regulation, signalling, adaptation, homeostasis, and disease-system dynamics can separate measured plasma exposure from the biologically active state. Systems pharmacodynamic models demonstrate that response may require intermediate states and regulatory processes beyond a direct instantaneous concentration–effect relationship [16]. Slow reversible binding provides one example of temporal dissociation. When association and dissociation occur over clinically relevant time scales, target occupancy and effect may persist after plasma concentration declines, producing delay or hysteresis [17]. A measured concentration at one time point may therefore be an inadequate representation of the pharmacodynamic state responsible for a later endpoint.

Modern modalities can require still more specific states. Targeted protein degraders, for example, may involve tissue distribution, ternary-complex formation, target degradation, resynthesis, and downstream response. Predictive PK/PD modelling can organize these processes, but each state requires mechanistic justification and validation before it is interpreted as biologically identified [18].

CDERRA therefore inserts a latent pharmacodynamic layer between exposure and response. The layer may contain biophase concentration, target engagement, turnover, signalling, adaptation, or another mechanism-proximal construct. It is proposed as an organizing device, not as evidence that every component can be separately estimated from sparse clinical observations.

Disease progression and treatment-confounder feedback

Disease is represented as a changing state rather than a fixed baseline label. Disease-progression models can organize trajectories, transitions, and treatment-associated changes across development and decision contexts [19]. Within CDERRA, disease state may influence symptoms, adherence, dose selection, physiology, clearance, concomitant therapy, sampling, and outcome measurement.

A central complication occurs when a time-varying variable predicts later treatment but is also affected by earlier treatment. Ordinary regression adjustment may then block part of the treatment pathway or induce bias rather than remove confounding [20]. Clinical status, toxicity, biomarker response, organ function, and disease severity can all acquire this dual role.

Applied longitudinal research has shown that treatment-confounder feedback becomes especially difficult when follow-up is sparse or observation depends on clinical history [21]. CDERRA therefore places treatment, disease state, physiological state, and response in successive temporal slices. The graph identifies candidate feedback pathways but does not prescribe one estimator for every question.

The proposed decision rule is temporal: before adjustment, each variable must be classified relative to prior treatment, current treatment, future treatment, and the estimand. A variable should not be called a confounder solely because it predicts exposure and outcome. Its causal history determines whether adjustment is necessary, harmful, or insufficient.

Table 2 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for testable propositions, evidence requirements, and validation criteria for reconstructing the causal dose–exposure–response chain under adherence variation.

Table 2. Testable propositions, evidence requirements, and validation criteria for Reconstructing the Causal Dose–Exposure–Response Chain under Adherence Variation.

Architecture layer or process stage	Core function	Information or material flow	Interaction with other components	Validation criterion	Potential failure mode	Human or experimental responsibility	Readiness boundary
Dose-input reconstruction	Estimate realized administrations	Prescription and adherence evidence → dose events	Determines absorbed input and exposure	Concordance with prospectively observed administration [10]	Prescription treated as administration	Clinical pharmacology and data-capture teams	Exploratory until dose-input accuracy is established
Exposure reconstruction	Estimate latent concentration trajectory	Dose input and physiology → exposure	Feeds PD state and concentration measurement	Recovery under realistic dose and sampling error [11]	False precision from missing doses	Pharmacometric modeller	Not decision-ready from fit alone
PK/PD persistence	Represent consequences of interruptions	Exposure and PD memory → sustained or lost effect	Modifies adherence consequences	Regimen-specific missed-dose evaluation [12]	Universal forgiveness assumption	Experimental and clinical pharmacology teams	No generalized missed-dose recommendation
Disease-state layer	Represent evolving clinical condition	Prior disease and treatment → later disease	Affects physiology, treatment, sampling, and outcome	Longitudinal external evaluation [19]	Disease trajectory treated as fixed	Disease-area experts	Prognostic performance is not causal validity
Feedback layer	Preserve prior-treatment pathways	Earlier treatment → covariate → later treatment/outcome	Determines valid adjustment strategy	Sensitivity to feedback-aware estimation [20]	Adjustment for treatment consequences	Causal analyst and statistician	Estimator depends on the estimand
Observation layer	Represent follow-up and sampling	Clinical history → observation probability	Determines which latent states become visible	Assessment of informative follow-up [21]	Selection bias from sparse observation	Investigators and data-governance teams	Missingness assumptions must remain explicit
Integrated CDERRA proposition	Align all layers in time	Treatment, biology, disease, and observation histories	Produces an estimand-specific architecture	Prospective structural and transport evaluation	Correct prediction for the wrong causal reason	Multidisciplinary validation group	**Proposed:** not clinically or regulatorily ready

Concomitant therapy and interaction structures

Concomitant therapy should be represented as a dynamic intervention rather than a fixed nuisance covariate. A co-administered drug may change absorption, metabolism, transport, elimination, target engagement, toxicity, adherence, disease state, or clinical dose modification. Its placement in the graph must correspond to the mechanism being proposed.

Drug, gene, and disease factors may combine to produce a metabolic phenotype that differs from the effect predicted by any component alone. Published drug–drug–gene–disease interactions demonstrate the complexity of these combined pathways and the limitations of assigning metabolic status from genotype or co-medication in isolation [22].

Mechanistic prediction does not remove the need for validation. Qualification and reporting procedures for physiologically based pharmacokinetic models emphasize that model credibility must be linked to the specific interaction question and context of use [23]. A model adequate for exploratory interaction assessment may be insufficient for individual dose modification.

Phenoconversion models provide one means of representing how genotype and enzyme inhibition jointly alter functional metabolic activity [24]. CDERRA generalizes this idea as an interaction-state layer, but it does not assume that every interaction is metabolic, additive, stable over time, or identifiable from routine data.

Measurement error, informative sampling, and missingness

The observed data are generated by an observation process as well as by biology. Missing doses, incorrect times, assay error, values below quantification limits, irregular sampling, missing covariates, endpoint misclassification, and undocumented treatment changes are distinct mechanisms. Pharmacometric guidance on problematic data emphasizes that these mechanisms require different assumptions and analytic responses [25].

Missing covariates should not automatically be replaced by a single deterministic value. Multiple-imputation workflows can propagate uncertainty when the imputation model and missingness assumptions are defensible [26]. Imputation nevertheless produces plausible completed data under assumptions; it does not recover the unknowable true value.

Censoring below the lower limit of quantification is similarly distinct from ordinary measurement error. Different handling strategies may alter parameter stability, precision, and bias, and their suitability depends on the model and information content [27]. A censored concentration should not be interpreted as a measured zero or an exact latent value.

Sampling can also be informative. Patients may be sampled more often after toxicity, poor response, dose change, clinical deterioration, or an unexpected concentration. CDERRA therefore includes nodes for sampling decision, timing, missingness, and measurement. Ignoring these processes can make the observed dataset a selected representation of the treatment and disease trajectory.

Table 3 organizes the evidence, constructs, relationships, uncertainties, and boundary conditions required for failure modes, boundary conditions, governance needs, and implementation priorities for reconstructing the causal dose–exposure–response chain under adherence variation.

Table 3. Failure modes, boundary conditions, governance needs, and implementation priorities for Reconstructing the Causal Dose–Exposure–Response Chain under Adherence Variation.

Decision scenario or design trade-off	Trigger or context	Competing objective	Evidence threshold	Uncertainty treatment	Corrective or escalation action	Expected implication	What cannot be concluded
Prescription used as delivered dose	Incomplete administration records	Simplicity versus treatment validity	Independent evidence of administration [25]	Represent dose events as uncertain	Improve record linkage or collect prospective administration data	Wider but more defensible exposure uncertainty	That prescribed treatment was received
Missing covariates replaced deterministically	Sparse clinical data	Operational convenience versus uncertainty propagation	Defensible missingness and imputation model [26]	Multiple imputation and sensitivity analysis	Escalate when results depend on missingness assumptions	More transparent parameter uncertainty	That imputed values are true
BLQ concentrations treated as zero	Limited assay sensitivity	Simplicity versus likelihood fidelity	Model-specific censoring assessment [27]	Censoring-aware likelihood or justified alternative	Compare handling strategies	Reduced risk of systematic distortion	That one BLQ method is universally optimal
Concomitant therapy entered as a fixed covariate	Changing co-medication or inflammation	Parsimony versus mechanism	Mechanistic interaction evidence [22]	Time-varying interaction state	Use pathway-specific PK or PD modelling	Clearer interpretation of interaction routes	That an observed interaction is causal

Interaction model used for dose change	High-consequence prediction	Model utility versus patient risk	Context-specific qualification [23]	Scenario and parameter uncertainty	Require external and prospective evaluation	Restricts use to supported contexts	That technical qualification ensures clinical benefit
Phenoconversion inferred from genotype alone	Inhibitors, inducers, disease effects	Genetic simplicity versus functional phenotype	Joint phenotype evidence [24]	Competing-model sensitivity analysis	Measure relevant co-therapy and physiology	More realistic metabolic-state representation	That genotype fixes phenotype over time
Informative sampling ignored	Testing triggered by clinical concern	Analytical convenience versus selection validity	Evidence on sampling determinants [25]	Joint or weighted observation modelling	Audit sampling pathways	Clarifies limits of observed concentration data	That sampled patients represent all treated patients
Integrated architecture used operationally	Transition from research to decision support	Speed versus credibility	Prospective multidisciplinary validation	Calibrated uncertainty and human oversight	Restrict use or escalate to formal impact study	Prevents premature deployment	That conceptual coherence establishes readiness

Identification, estimation, and validation requirements

Identification begins with a clearly defined estimand. The intervention may concern prescribed dose, administered dose, adherence support, exposure targeting, treatment adaptation, or concomitant-therapy modification. These interventions are not interchangeable. Model-selection guidance for precision-dosing software further shows that population relevance, model adequacy, implementation correctness, and context-specific validation must be assessed separately [28].

Credibility requirements should increase with the influence of the model and the consequence of the decision. Risk-informed credibility frameworks link verification and validation effort to context of use rather than treating model acceptance as a universal property [29]. CDERRA adopts this principle while extending it to causal structure, dose reconstruction, observation processes, and decision consequences.

Estimation requires assumptions appropriate to the selected estimand. These may include consistency, positivity, conditional exchangeability, correct temporal ordering, adequate representation of dose input, and defensible measurement and missingness models. Possible approaches include joint longitudinal models, g-methods, mechanistic PK/PD models, Bayesian latent-state models, or combinations of these methods. No estimator repairs an unidentified question automatically.

Validation should proceed through distinct levels: computational verification, observation-model assessment, simulation-based trajectory recovery, internal predictive evaluation, external transport evaluation, causal sensitivity analysis, and prospective decision-impact assessment. Goodness of fit is not causal validation; external prediction is not proof of the graph; and technically correct software does not establish beneficial clinical use.

Limitations and clinical implications

CDERRA is limited by the quality and scope of the available data and assumptions. Unmeasured adherence, poorly recorded co-therapy, latent disease processes, structural misspecification, limited positivity, and selective follow-up may remain unresolved. Future foundation-model components could assist data representation or extraction, but general medical representation learning does not remove the need for task-specific calibration, causal assumptions, and governance [30].

The architecture is also not a digital twin. Health digital twins remain heterogeneous and require longitudinal updating, interoperability, verification, validation, and governance tied to a specific purpose [31]. CDERRA is a conceptual causal architecture that could inform such development but does not simulate a verified individual patient or support autonomous treatment decisions.

Implementation would require multidisciplinary governance involving clinicians, pharmacometricians, causal analysts, laboratory scientists, data engineers, statisticians, and patients. Model outputs should remain accompanied by uncertainty, provenance, applicability limits, and escalation rules. Any transition from research use to clinical decision support would require prospective evaluation and continuing oversight.

CONCLUSION

The proposed Causal Dose–Exposure–Response Reconstruction Architecture reframes dose–exposure–response analysis as reconstruction of a temporally ordered treatment, biological, disease, and observation process. Its central contribution is the separation of prescribed and administered dose, latent and measured exposure, pharmacodynamic state and observed outcome, and biological variability and observation error. By integrating adherence, changing physiology, delayed pharmacodynamics, disease progression, concomitant therapy, treatment-confounder feedback, and missingness, the architecture provides a basis for clearer causal questions and validation planning. It remains a proposed non-empirical synthesis whose usefulness depends on context-specific assumptions, data, external validation, causal sensitivity analysis, and prospective evaluation. It is not a validated dosing model, clinical recommendation, regulatory determination, digital twin, or deployment-ready system.

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. 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. 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
4. Turner DC, Kondic AG, Anderson KM, Robinson AG, Garon EB, Riess JW, et al. Pembrolizumab exposure-response assessments challenged by association of cancer cachexia and catabolic clearance. Clin Cancer Res. 2018;24(23):5841-9. doi:10.1158/1078-0432.CCR-18-0415
5. 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
6. 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). doi:10.1111/cts.13676
7. Rogers JA, Maas H, Pitarch AP. An introduction to causal inference for pharmacometricians. CPT Pharmacometrics Syst Pharmacol. 2023;12(1):27-40. doi:10.1002/psp4.12894
8. Wiens MR, French JL, Rogers JA. Confounded exposure metrics. CPT Pharmacometrics Syst Pharmacol. 2024;13(2):187-91. doi:10.1002/psp4.13074
9. Yin X, Xiong Y, Bi Y, Wei X, Zhao H, Fletcher EP, et al. Evaluation and mitigation of time-dependent confounding effects in conventional exposure-response analyses for oncology drugs. CPT Pharmacometrics Syst Pharmacol. 2025;14(12):2118-27. doi:10.1002/psp4.70119
10. Ding J, Hoglund RM, Tarning J. Medication adherence framework: A population-based pharmacokinetic approach and its application in antimalarial treatment assessments. CPT Pharmacometrics Syst Pharmacol. 2024;13(5):795-811. doi:10.1002/psp4.13119
11. Wang Y, Liu X. Handling missing dosing history in population pharmacokinetic modeling: An extension to MDM method. CPT Pharmacometrics Syst Pharmacol. 2019;8(1):39-49. doi:10.1002/psp4.12374

12. McAllister NP, Lawley SD. A pharmacokinetic and pharmacodynamic analysis of drug forgiveness. J Pharmacokinet Pharmacodyn. 2022;49(3):363-79. doi:10.1007/s10928-022-09808-w
13. Petitcollin A, Bensalem A, Verdier MC, Tron C, Lemaitre F, Paintaud G, et al. Modelling of the time-varying pharmacokinetics of therapeutic monoclonal antibodies: A literature review. Clin Pharmacokinet. 2020;59(1):37-49. doi:10.1007/s40262-019-00816-7
14. 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
15. Lenoir C, Niederer A, Rollason V, Desmeules JA, Daali Y, Samer CF. Prediction of cytochromes P450 3A and 2C19 modulation by both inflammation and drug interactions using physiologically based pharmacokinetics. CPT Pharmacometrics Syst Pharmacol. 2022;11(1):30-43. doi:10.1002/psp4.12730
16. Ayyar VS, Jusko WJ. Transitioning from basic toward systems pharmacodynamic models: Lessons from corticosteroids. Pharmacol Rev. 2020;72(2):414-38. doi:10.1124/pr.119.018101
17. Ren T, Zhu X, Jusko NM, Krzyzanski W, Jusko WJ. Pharmacodynamic model of slow reversible binding and its applications in pharmacokinetic/pharmacodynamic modeling: Review and tutorial. J Pharmacokinet Pharmacodyn. 2022;49(5):493-510. doi:10.1007/s10928-022-09822-y
18. Haid RTU, Reichel A. Transforming the discovery of targeted protein degraders: The translational power of predictive PK/PD modeling. Clin Pharmacol Ther. 2024;116(3):770-81. doi:10.1002/cpt.3273
19. 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
20. Mansournia MA, Etminan M, Danaei G, Kaufman JS, Collins G. Handling time varying confounding in observational research. BMJ. 2017;359. doi:10.1136/bmj.j4587
21. Karim ME, Tremlett H, Zhu F, Petkau J, Kingwell E. Dealing with treatment-confounder feedback and sparse follow-up in longitudinal studies: Application of a marginal structural model in a multiple sclerosis cohort. Am J Epidemiol. 2021;190(5):908-17. doi:10.1093/aje/kwaa243
22. Storelli F, Samer CF, Reny JL, Desmeules JA, Daali Y. Complex drug-drug-gene-disease interactions involving cytochromes P450: Systematic review of published case reports and clinical perspectives. Clin Pharmacokinet. 2018;57(10):1267-93. doi:10.1007/s40262-018-0650-9
23. Shebley M, Sandhu P, Emami Riedmaier A, Jamei M, Narayanan R, Patel A, et al. Physiologically based pharmacokinetic model qualification and reporting procedures for regulatory submissions: A consortium perspective. Clin Pharmacol Ther. 2018;104(1):88-110. doi:10.1002/cpt.1013
24. Viviani R, Berres J, Stingl JC. Phenotypic models of drug-drug-gene interactions mediated by cytochrome drug-metabolizing enzymes. Clin Pharmacol Ther. 2024;116(3):592-601. doi:10.1002/cpt.3188
25. 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
26. Vuong ML, Verbeke G, Dreesen E. A multiple imputation workflow for handling missing covariate data in pharmacometrics modeling. CPT Pharmacometrics Syst Pharmacol. 2025;14(6):991-1005. doi:10.1002/psp4.70039
27. Wijk M, Wasmann RE, Jacobson KR, Svensson EM, Denti P. A pragmatic approach to handling censored data below the lower limit of quantification in pharmacokinetic modeling. CPT Pharmacometrics Syst Pharmacol. 2025;14(6):1042-9. doi:10.1002/psp4.70015
28. Taylor ZL, Poweleit EA, Paice K, Somers KM, Pavia K, Vinks AA, et al. Tutorial on model selection and validation of model input into precision dosing software for model-informed precision dosing. CPT Pharmacometrics Syst Pharmacol. 2023;12(12):1827-45. doi:10.1002/psp4.13056
29. Kuemmel C, Yang Y, Zhang X, Florian J, Zhu H, Tegenge M, et al. Consideration of a credibility assessment framework in model-informed drug development: Potential application to physiologically-based pharmacokinetic modeling and simulation. CPT Pharmacometrics Syst Pharmacol. 2020;9(1):21-28. doi:10.1002/psp4.12479
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
31. Katsoulakis E, Wang Q, Wu H, Shahriyari L, Fletcher R, Liu J, et al. Digital twins for health: A scoping review. NPJ Digit Med. 2024;7(1):77. doi:10.1038/s41746-024-01073-0