IIAR JOURNALS

Innovating Knowledge. Advancing Research.

IIAR Journal of Pharmaceutical Sciences

Vol 1, No 1, 2026

Translating In Silico Predictions into Human Dose Feasibility in AI-Driven Antibacterial Discovery: A Conceptual Pharmacokinetic/Pharmacodynamic Framework

IGWE Emmanuel Sunday[1], UDEH Samuel Chibueze[2], NWOGBOJI Emmanuel Usulor[3]

[1,3] Veristzon Global Research Lab; [2] Department of Pharmaceutical Services, National Obstetric Fistula Centre

ABSTRACT

Computational discovery pipelines identify antibacterial candidates primarily on potency and, increasingly, on predicted physicochemical and ADMET properties, yet many candidates are later abandoned because achievable human exposure cannot reach pharmacodynamic targets. This theoretical paper develops a conceptual framework linking in silico predictions of potency, clearance, distribution and protein binding to early judgements of human dose feasibility through pharmacokinetic/pharmacodynamic target attainment. The study is non-empirical and uses a conceptual research design. A corpus of peer-reviewed journal articles published between 2020 and 2026 was assembled across four literatures, namely antibacterial pharmacokinetics, pharmacodynamics and target attainment, translational and population dose prediction, machine-learning prediction of pharmacokinetic and ADMET properties, and computational antibacterial discovery, with a supplementary search on resistance burden, antibiotic access and laboratory capacity in sub-Saharan Africa. Every source was verified against its Crossref publisher record for authorship, journal, volume, pagination and year and screened for retractions, corrections and expressions of concern in that record and in the corresponding PubMed record; sources failing verification were excluded rather than corrected by inference. Analysis proceeded by decomposing the translational step into its constituent links, classifying each as predicted, measured or assumed, and reasoning from the resulting structure. Three findings are reported. First, because the links combine multiplicatively, prediction errors compound rather than average and each link contributes to the result in proportion to its own variance; separately, a link whose error has never been quantified cannot be propagated at all and therefore bounds what the calculation can claim, site penetration being the link most often assumed and least often quantified. Second, because feasibility is adjudicated at a threshold rather than estimated, the informativeness of in silico reasoning varies with a candidate's distance from the feasibility boundary and is lowest precisely for difficult cases. Third, feasibility is a two-sided window bounded below by target attainment and above by toxicity, so frameworks modelling only the efficacy bound systematically overestimate how many candidates are viable. The paper derives a prioritisation rule that evaluates each candidate under favourable and unfavourable assumptions, and establishes an asymmetry: in silico reasoning can reliably exclude candidates whose window is empty under optimistic assumptions but cannot reliably rank candidates whose windows are both non-empty. Six falsifiable propositions and ten reporting principles are presented. It is recommended that uncertainty be propagated rather than point-estimated, that error in currently assumed links be quantified as a priority, that predictive pharmacokinetic models be evaluated against decision accuracy rather than endpoint accuracy alone, and that for settings without routine therapeutic drug monitoring the width of the feasibility window be treated as a design criterion.

Keywords: antibacterial drug discovery; pharmacokinetics; pharmacodynamics; target attainment; human dose prediction

Translating In Silico Predictions into Human Dose Feasibility in AI-Driven Antibacterial Discovery: A Conceptual Pharmacokinetic/Pharmacodynamic Framework

1. INTRODUCTION

Computational discovery pipelines now identify antibacterial candidates in numbers that outstrip any capacity to characterise them experimentally. Models trained on whole-cell growth inhibition have yielded compounds subsequently confirmed in the laboratory, including agents active against priority Gram-negative pathogens (Stokes et al., 2020; Liu, Catacutan, et al., 2023), and candidates are filtered primarily on predicted potency and, increasingly, on predicted physicochemical and absorption, distribution, metabolism, excretion and toxicity properties (Swanson et al., 2024; Ekins et al., 2023). The need is pressing: bacterial antimicrobial resistance was associated with an estimated 4.71 million deaths in 2021, and a drug pipeline directed at Gram-negative organisms could cumulatively avert some 11.1 million resistance-attributable deaths between 2025 and 2050 (Naghavi et al., 2024).

Yet many candidates that satisfy potency and property filters are later abandoned for a reason those filters do not address: no dose that a human can tolerate produces an exposure sufficient to meet the pharmacodynamic target. Antibacterial efficacy is governed not by potency alone but by the relation between the free drug concentration over time and the minimum inhibitory concentration of the organism, summarised in indices such as the time above the minimum inhibitory concentration or the ratio of area under the concentration-time curve to that concentration (Rao & Landersdorfer, 2021). Whether a candidate can meet such a target in people depends on its clearance, its distribution, the fraction unbound in plasma, the concentration achieved at the site of infection, and the dose that can be given without toxicity. Each of these is a separate quantity, and in early discovery each is predicted rather than measured.

This paper develops a conceptual framework linking in silico predictions of potency, clearance, distribution and protein binding to early judgements of human dose feasibility through pharmacokinetic/pharmacodynamic target attainment. Its organising claim is that the translational step is not a single inference but a chain of them, and that the reliability of the conclusion is governed by how error accumulates along that chain rather than by the accuracy of any one prediction. Because the decision at the end of the chain is a threshold judgement, whether a feasible dose exists, modest errors in the individual links can be sufficient to reverse it, and the informativeness of the whole exercise varies sharply with how close a candidate lies to that threshold.

A second claim concerns the shape of the feasibility question. Dose feasibility is usually framed as a floor: can exposure reach the pharmacodynamic target? It is in fact a window, bounded below by that target and above by toxicity, and both bounds are uncertain at the point of prediction. A framework that models only the lower bound will systematically overstate how many candidates are feasible.

The paper therefore sets out to identify the assumptions required to carry a prediction from chemical structure to a statement about human dose; to analyse how prediction error accumulates through each link of that chain; to distinguish the decisions that early pharmacokinetic/pharmacodynamic reasoning can reasonably inform from those it cannot; to state in falsifiable form when in silico predictions are sufficient to support lead prioritisation; and to propose principles for incorporating dose-feasibility reasoning into artificial-intelligence-driven antibacterial discovery. The study is conceptual and generates no new measurements. Section 2 reviews the relevant literature, Section 3 states the design and verification protocol, Section 4 presents the framework, the propositions and the reporting principles and discusses their implications, including for antibacterial use in African health systems, and Sections 5 and 6 conclude.

2. LITERATURE REVIEW

2.1 CONCEPTUAL REVIEW

Dose feasibility is the proposition that some administrable regimen produces, in the intended patient population and at the intended site of infection, an exposure that attains the pharmacodynamic target without exceeding the

© IIARJournal. All rights reserved.

exposure at which unacceptable toxicity occurs. Four features of this definition organise the paper. It is a statement about a population rather than an individual, since exposure varies between patients; it is site-specific, since the concentration that matters is the one at the infection; it is two-sided, bounded by efficacy below and toxicity above; and it is probabilistic, properly expressed as the probability of attainment across the population rather than as attainment in a typical patient.

Each of these features has an established operational counterpart. Probability of target attainment is estimated by Monte Carlo simulation over a population pharmacokinetic model, and is the basis on which breakpoints and regimens are set (Liu, Moodley, et al., 2023). Site specificity is addressed by measuring or simulating concentrations in the relevant compartment. The two-sided character is addressed when toxicity is modelled alongside efficacy, which is done less often. The point of the present paper is that early discovery must make the same judgement with predicted rather than measured inputs, and that the properties of the judgement change when it does.

2.2 THEORETICAL REVIEW

The theoretical structure is a chain. Predicted potency against the target organism supplies the pharmacodynamic denominator. Predicted clearance and volume of distribution determine the shape and magnitude of the concentration-time curve for a given dose. Predicted fraction unbound converts total to free concentration, since only free drug is microbiologically active. Predicted or assumed tissue penetration converts plasma exposure to exposure at the site. The pharmacodynamic index and its target magnitude convert that exposure into a statement about attainment. A dose is then feasible if the regimen required for attainment lies below the toxicity bound.

Two properties of this chain drive the analysis. The first is that the links are multiplicative rather than additive in their effect on the final quantity: exposure scales inversely with clearance, free exposure scales with fraction unbound, and target attainment depends on the ratio of free exposure at the site to the minimum inhibitory concentration. Errors therefore compound rather than average, and the contribution of each link to the uncertainty of the result is weighted by its own variance rather than equalised across the chain.

The second is that the conclusion is a threshold judgement. Feasibility is not a continuous quantity to be estimated but a binary decision, advance or abandon, taken by comparing a predicted window with an empty set. Near the boundary, a modest error in predicted exposure changes the decision; far from it, even a large error does not. The informativeness of in silico pharmacokinetic/pharmacodynamic reasoning is therefore not uniform across candidates, and this non-uniformity is the basis of the prioritisation rule developed in Section 4.

2.3 EMPIRICAL REVIEW

Each link of the chain has a literature characterising what prediction can presently achieve. Machine learning for pharmacokinetic parameters has advanced substantially, with models predicting animal and human parameters directly from structure (Obrezanova, 2023) and reviews documenting the breadth of application and the persisting constraints of limited and heterogeneous training data (Paliwal et al., 2024; Huang et al., 2025). Magnitudes matter here: in a structurally diverse series of insulin analogues with abundant in-house data, machine-learning models predicted clearance with average fold errors of roughly 2.5 to 2.9, with around seventy per cent of predictions within twofold (Einarson et al., 2023). That is a creditable performance for a difficult endpoint, and it is also a severalfold uncertainty entering one link of a multiplicative chain.

Absorption, distribution, metabolism, excretion and toxicity prediction more broadly is now available at library scale through public platforms (Swanson et al., 2024) and continues to improve architecturally (Xu et al., 2024), with reviews documenting both progress and the endpoint-dependent character of accuracy (Zhang et al., 2025; Rácz et al., 2021). Reflective assessments of two decades of the field note that optimisation now addresses many properties simultaneously while the underlying data constraints remain (Ekins et al., 2023).

The clinical pharmacokinetic literature supplies the magnitudes for the remaining links. Exposure varies markedly between patient populations for the same agent and dose, as documented for gentamicin across diverse groups (Hodiamont et al., 2022) and for teicoplanin in a systematic review of attainment (Mouton et al., 2025), and in states such as continuous renal replacement therapy where

© IIARJournal. All rights reserved.

regimens must be re-derived by simulation (Charoensareerat et al., 2023). Population variation is accordingly handled by estimating attainment probabilities rather than point values, including in paediatric and neonatal populations where the relevant parameters differ again (Cafaro et al., 2024; Haynes et al., 2025).

Site specificity is well evidenced. Plasma is an imperfect surrogate for the site of infection: individual target attainment in epithelial lining fluid diverges from that in plasma for meropenem (Rohani et al., 2023), and concentrations in cerebrospinal fluid and brain tissue require their own population models (Alshaer, Barlow, et al., 2022). Attainment also carries clinical consequence rather than being a modelling abstraction, with early and cumulative beta-lactam target attainment associated with outcomes in bloodstream infection (Alshaer et al., 2023) and with outcomes in intensive care pneumonia, where machine learning was used to define the association (Alshaer, Maranchick, et al., 2022).

The upper bound of the window is less often modelled but is well documented where it has been. A population analysis of cefepime estimated both antibacterial target attainment and the probability of neurotoxicity as functions of exposure in critically ill patients, making explicit that the regimen must satisfy two constraints simultaneously (Bilal et al., 2023).

The gap this paper addresses is therefore not an absence of work on any single link. It is the absence of an account of how the links compose: how error accumulates along the chain, how that accumulation interacts with a threshold decision, and what follows for which candidates in silico reasoning can usefully triage.

3. MATERIALS AND METHODS

3.1 RESEARCH DESIGN

This study is non-empirical and uses a conceptual research design. The unit of analysis is the published literature rather than compounds, bacteria or patients, and the output is a framework with falsifiable propositions rather than a dataset. No compound was synthesised, no accumulation was measured and no model was trained. There is consequently no sample, no randomisation and no statistical inference; what follows describes how the literature was assembled, how claims were admitted and on what basis the paper's two recurring judgements were made.

3.2 CORPUS CONSTRUCTION

Four bodies of literature were searched: antibacterial pharmacokinetics, pharmacodynamics and target attainment; translational and population dose prediction; machine-learning prediction of pharmacokinetic and ADMET properties; and computational antibacterial discovery. A supplementary search covered antimicrobial resistance burden, antibiotic access and laboratory capacity in sub-Saharan Africa, to anchor the analysis in the setting for which the authors write. Searches were conducted in PubMed and supplemented by citation tracing from the retrieved reviews.

Sources were included where they were peer-reviewed journal articles published between 2020 and 2026, treated at least one of these domains substantively, and reported enough methodological detail for the cited claim to be attributable, in particular, for pharmacokinetic studies, the population, the matrix sampled and the exposure metric reported. Preprints, conference abstracts, theses, commentaries without original synthesis, and sources whose claims could not be traced to the specific statement they were cited for were excluded. Priority was given to primary reports stating population characteristics and exposure metrics explicitly, and to reviews aggregating quantitative performance rather than restating conclusions.

3.3 CITATION VERIFICATION PROTOCOL

Every source admitted to the reference list was verified against its publisher record before citation, under a five-check protocol. First, the digital object identifier was resolved against the Crossref registry and the record compared field by field with the citation: authors and their order, title, journal, volume, issue, pagination or article number, and year. Second, each record was screened for retractions, corrections, errata and expressions of concern through the relation and update fields of the registry record and, independently, through the publication-type fields of the corresponding PubMed record. Third, the publication year recorded in the registry was adopted throughout, which for several items in this corpus is the year of advance online publication rather than that of the print issue to which the article was later assigned. Fourth,

© IIARJournal. All rights reserved.

discrepancies between an indexing database and the publisher record were resolved in favour of the publisher record. Fifth, sources failing any check were excluded rather than corrected by inference.

No source is cited on the strength of an abstract alone, and no bibliographic field in the reference list was supplied from memory or inference. Where a cited study's population, sampling matrix or exposure metric bears on the claim it supports, that detail is stated in the text rather than left implicit, because the argument of this paper is precisely that such details determine what an exposure figure means.

3.4 BASES OF JUDGEMENT

Two judgements recur and both are interpretive rather than measured. The first is how each link of the translational chain is classified in early discovery. A link is treated as predicted where its value is obtained from a model operating on chemical structure; as measured where an experimental value for the candidate itself is available; and as assumed where a value is carried over from class precedent or set to a default, the commonest example being the assumption that concentration at the site of infection equals that in plasma. The classification is made from what a study states it did, not from what would have been preferable.

The second is what magnitude of uncertainty to attribute to a predicted link. Where a primary study reports a quantitative measure of predictive error for an endpoint (an average fold error, a proportion of predictions within twofold, or a correlation on a held-out set), that figure is used and attributed to the study and the chemical series that produced it. Where no such figure is available, the paper says that the magnitude is uncharacterised rather than substituting a value from a different endpoint or series; uncharacterised links are treated in Section 4 as a distinct category, because an unquantified uncertainty cannot be propagated.

3.5 DISCLOSURE OF ARTIFICIAL INTELLIGENCE USE

Generative artificial intelligence tools assisted with literature search, with retrieval and cross-checking of bibliographic records against the Crossref registry, and with language editing of the authors' drafted text. No text entered the manuscript without the authors' verification of its substantive claims against the cited sources. No such tool was used to generate data, to select citations without subsequent registry verification, or to formulate the framework and propositions in Section 4, which are the authors' own. The authors accept full responsibility for the content, including the accuracy of all references.

3.6 VALIDITY AND LIMITATIONS OF THE APPROACH

Construct validity concerns whether the links of the chain correspond to quantities that actually determine whether a dose is feasible; they are drawn from established pharmacokinetic/pharmacodynamic practice rather than invented. Descriptive validity concerns whether the predicted, measured and assumed classification fairly represents the cited studies; the criteria in Section 3.4 are stated so that a reader can contest a particular assignment. Theoretical validity concerns falsifiability; each proposition names the comparison that would contradict it. The principal limitation is that error magnitudes are available for some links and not others, so the propagation argument in Section 4 is structural rather than numerical; it specifies how errors compose without claiming to quantify the result for any particular candidate.

4. RESULTS AND DISCUSSION

4.1 RESULTS

4.1.1 The translational chain and its error structure

Setting the links of the chain side by side shows that they differ not only in what they contribute but in how well their uncertainty is characterised. Table 1 states the position.

Table 1. Links of the translational chain from chemical structure to human dose feasibility, and the effect of error in each.

Link	How obtained in early discovery	State of uncertainty characterisation	Effect of error on the feasibility judgement
Potency against the target organism	Predicted, or measured in vitro against reference isolates	Assay- and strain-dependent; the relevant quantity is a distribution across isolates, not a point	Shifts the required exposure in proportion
Clearance	Predicted from chemical structure	Quantified for some series; severalfold average fold errors reported even with abundant in-house data	Inversely proportional effect on exposure

© *IIARJournal. All rights reserved.*

Volume of distribution	Predicted from chemical structure	Quantified less often than clearance	Alters the shape of the concentration-time curve and which index applies
Fraction unbound in plasma	Predicted from chemical structure	Small absolute errors become large relative errors for highly bound compounds	Multiplicative on free, microbiologically active exposure
Penetration to the site of infection	Usually assumed equal to plasma	Largely uncharacterised at the prediction stage; compound- and site-specific when measured	Can dominate the judgement for sites with restricted penetration
Pharmacodynamic index and target magnitude	Taken from class precedent	Uncertain for novel mechanisms, where neither index nor magnitude is established	Determines which feature of exposure must be optimised
Toxicity bound	Assumed, or deferred to in vivo study; rarely predicted at this stage	Commonly unknown until in vivo study	Determines whether the feasibility window is non-empty at all

Two features of the third column govern what follows. The first is that characterisation is uneven: clearance prediction has been evaluated quantitatively, with average fold errors of roughly 2.5 to 2.9 and about seventy per cent of predictions within twofold reported for a well-resourced series (Einarson et al., 2023), whereas penetration to the site of infection is ordinarily not predicted at all but assumed, and the toxicity bound is commonly unknown until animal study. The second is that these are not errors of the same kind. A quantified uncertainty can be propagated; an assumption whose error is uncharacterised cannot, and its presence bounds the credibility of the whole calculation irrespective of how well the other links are estimated. These are therefore two different deficits, and Section 4.1.2 treats them separately.

4.1.2 How error accumulates

Because exposure varies inversely with clearance, free exposure scales with the unbound fraction, and attainment depends on the ratio of free exposure at the site to the minimum inhibitory concentration, the links compose multiplicatively. Independent errors in separate links therefore do not cancel, and the uncertainty in predicted exposure is wider than that of any single link. Two distinct consequences follow, and the paper keeps them apart because they have been run together in discussions of this problem. The first is quantitative and is stated in Proposition 1: the links contribute in proportion to their variances, so a single link dominates only when its uncertainty substantially exceeds that of the others combined, and where several are comparable the chain is improved by refining any of them. The second is not a matter of magnitude at all. A link whose error has never been quantified cannot be entered into such a calculation, so its presence sets a limit on what the calculation can claim irrespective of how the quantified links behave. Site penetration is the clearest instance: the case for characterising it early rests not on a demonstration that its error is the largest, which by definition cannot be shown while it remains unquantified, but on the fact that leaving it unquantified prevents any propagated estimate from being complete.

The second consequence follows from the nature of the decision. Feasibility is not estimated but adjudicated: a candidate is advanced or abandoned according to whether a non-empty window exists. The probability that accumulated error reverses that adjudication depends on how far the candidate lies from the boundary. For a compound whose required exposure is far below any plausible toxicity bound, severalfold uncertainty changes nothing; for one whose window is narrow, the same uncertainty is decisive. In silico pharmacokinetic/pharmacodynamic reasoning is accordingly informative in inverse proportion to how difficult the case is, a property that deserves to be stated openly rather than discovered late.

4.1.3 The feasibility window

Dose feasibility is customarily discussed as a floor. The literature in which both bounds have been modelled shows why that is insufficient. A population analysis of cefepime in critically ill patients estimated antibacterial target attainment and the probability of neurotoxicity as functions of the same exposure, so that the acceptable regimen was constrained from both directions simultaneously (Bilal et al., 2023). Framed this way, feasibility is the non-emptiness of an interval, and a framework modelling only its lower edge will classify as feasible candidates whose windows are in fact empty. The error is systematic rather than random, and is largest precisely where it matters most, for potent compounds whose required exposures approach their toxic ones.

The window is also site-specific and population-specific. Target attainment in epithelial lining fluid diverges from that in plasma for the same patient

© IIARJournal. All rights reserved.

and agent (Rohani et al., 2023), and central nervous system compartments require their own models (Alshaer, Barlow, et al., 2022); exposure varies between patient populations for identical regimens (Hodiamont et al., 2022; Mouton et al., 2025), and in states such as continuous renal replacement therapy regimens must be re-derived entirely (Charoensareerat et al., 2023). A single number for feasibility therefore conceals a quantity that is properly a probability over a population at a specified site.

4.1.4 A prioritisation rule

These properties yield an operative rule with an asymmetric structure. For each candidate, the chain is evaluated twice: once under assumptions favourable to feasibility at every link, and once under unfavourable assumptions. A candidate whose window is empty even under the favourable evaluation can be excluded with confidence, because the conclusion is robust to error in the direction that would rescue it. A candidate whose window is non-empty under the unfavourable evaluation can be advanced with confidence. Candidates whose classification differs between the two evaluations are not ranked by the model at all; they are referred for the specific measurement that would resolve the discrepancy, which is identified by determining which link contributes most of the width.

This rule deliberately declines to produce a complete ordering. Its justification is Proposition 6: exclusion and selection make asymmetric evidential demands. Concluding that no feasible dose exists requires only that the optimistic case fail, which accumulated error cannot easily overturn. Concluding that one feasible candidate is better than another requires discriminating between quantities whose predicted values are separated by less than their uncertainty. Pipelines that rank candidates on predicted exposure are therefore making the second kind of claim on evidence that supports only the first.

4.1.5 Propositions

The framework yields six propositions, each stated so that a specified comparison could contradict it.

Proposition 1 (variance-weighted contribution). Because the links of the translational chain combine multiplicatively, the uncertainty in predicted exposure is set by the combined relative variances of the links, so each contributes in proportion to its own. A single link dominates only where its variance substantially exceeds that of the rest taken together; where several carry comparable uncertainty, refining any one of them improves decision quality materially. Contradicted if decision accuracy is insensitive to refinement of a link whose variance is comparable to the largest in the chain.

Proposition 2 (threshold amplification). Because feasibility is a threshold judgement, the probability that a given prediction error reverses the decision is a function of the candidate's distance from the feasibility boundary. In silico pharmacokinetic/pharmacodynamic reasoning is therefore highly informative for candidates far from the boundary and weakly informative for those near it. Contradicted if decision reversal rates are independent of distance from the boundary at constant prediction error.

Proposition 3 (the window is two-sided). Dose feasibility is bounded below by target attainment and above by toxicity, and both bounds are uncertain at the point of prediction. Frameworks that model only the efficacy bound will systematically overestimate the proportion of candidates that are feasible, and the overestimate will be largest for potent compounds with narrow therapeutic margins. Contradicted if feasibility judgements made on the efficacy bound alone agree with those made on both bounds across a representative candidate set.

Proposition 4 (plasma is not a sufficient statistic). Attainment computed from plasma exposure does not determine attainment at the site of infection, and the discrepancy is compound- and site-specific rather than a fixed factor. Site penetration must therefore be treated as a distinct link with its own uncertainty rather than folded into a general margin. Contradicted if a single compound-independent scaling factor reconciles plasma and site attainment across agents and sites.

Proposition 5 (the decision quantity is a distribution). The quantity that determines feasibility is the probability of target attainment across the intended patient population, not attainment in a typical patient. Point predictions of exposure are therefore category errors rather than approximations, and a candidate may be feasible for one population and not for another on identical chemistry. Contradicted if population-level attainment probabilities are recoverable from typical-patient point predictions without information about between-patient variability.

© IIARJournal. All rights reserved.

Proposition 6 (exclusion and selection are asymmetric). In silico pharmacokinetic/pharmacodynamic reasoning can reliably exclude candidates whose feasibility window is empty even under optimistic assumptions, because that conclusion is robust to error in the optimistic direction. It cannot reliably select among candidates whose windows are non-empty, because ranking requires the accuracy that the accumulated error denies. Contradicted if predicted rankings among feasible candidates correspond to observed rankings of achieved human exposure.

4.1.6 Reporting principles

The propositions presuppose that studies record enough for a reader to reconstruct the chain. Table 2 states the minimum, and each item corresponds to information whose absence, in assembling this paper's corpus, made it impossible to determine what a reported prioritisation claim rested on.

Table 2. Principles for reporting pharmacokinetic/pharmacodynamic-informed prioritisation in computational antibacterial discovery.

Principle	What should be stated or done
Declare the product profile	The intended route, population and site of infection, stated before candidates are ranked
Classify every link	Which inputs to the chain are predicted, which measured for the candidate, and which assumed from precedent
Propagate rather than point-estimate	Carry uncertainty through the chain instead of combining point estimates into a single exposure figure
Report both bounds	The efficacy bound and the toxicity bound, and the resulting window, not the efficacy bound alone
Report the decision under both extremes	Whether the candidate is feasible under optimistic and under pessimistic assumptions, not only at the central estimate
State the pharmacodynamic index and its provenance	Which index is used, what target magnitude, and from which organism, model and class precedent it was taken
Express attainment as a probability	Attainment across the intended population rather than in a typical patient
Flag proximity to the boundary	Whether the candidate sits near the feasibility threshold, where prediction is least informative
Separate exclusion from selection	Whether the claim is that a candidate should be dropped or that it should be preferred to another, since these require different evidence
State uncharacterised links	Which links carry uncertainty of unknown magnitude, since these cannot be propagated and bound the whole analysis

4.2 DISCUSSION

The usual response to candidates failing on exposure is to add predicted properties to the filter stack: having screened on potency, screen also on clearance, on protein binding, on solubility. The analysis above suggests this misdiagnoses the problem. Coverage is not the binding constraint; composition is. Adding a seventh well-predicted property to a chain that still contains an unquantified assumption about site penetration does not improve the feasibility judgement, because the judgement is weighted by the variances of all the links and is bounded by any link that remains unquantified, and because it is adjudicated at a threshold. What is needed is not more predicted properties but an explicit account of how the predicted, measured and assumed quantities combine, and of how far the resulting uncertainty sits from the decision boundary.

Clinical pharmacokinetics has long handled exactly this structure correctly. Target attainment is routinely expressed as a probability over a simulated population rather than as a value for a typical patient, and regimens and breakpoints are set on that basis (Liu, Moodley, et al., 2023). Discovery has imported the vocabulary of pharmacokinetics and pharmacodynamics without importing this machinery, and the framework proposed here is in large part an argument for importing it earlier: probabilistic rather than point-valued, population-level rather than typical-patient, and two-sided rather than efficacy-only. The obstacle is not conceptual novelty but the fact that in early discovery the inputs are predicted, which changes the magnitude of the uncertainty but not the form of the calculation.

This has a corollary for how machine-learning pharmacokinetic models should be evaluated. The literature reports accuracy per endpoint, which is the right metric for model development (Obrezanova, 2023; Huang et al., 2025). It is not the right metric for the use considered here, where the quantity of interest is the proportion of advance-or-abandon decisions the chain gets right. Those are different objectives: a model that improves clearance prediction may leave decision accuracy untouched, while a crude but quantified estimate of site penetration might improve it substantially. Evaluating predictive pharmacokinetics against decisions rather than against endpoints would make

© IIARJournal. All rights reserved.

such trade-offs visible, and nothing in the present state of the models prevents it.

It should be emphasised that target attainment is not a modelling abstraction with uncertain clinical meaning. Early and cumulative beta-lactam target attainment has been associated with outcomes in patients with bloodstream infections (Alshaer et al., 2023), and machine learning has been used to characterise the same association in intensive care pneumonia (Alshaer, Maranchick, et al., 2022). The quantity that the chain is trying to predict is therefore one that matters to patients, which is the justification for treating its uncertainty seriously rather than as a modelling nicety.

4.2.1 Implications for antibacterial use in African health systems

The population and the monitoring context for which a candidate is intended are part of the feasibility question rather than considerations that follow it. Three features of many African health systems bear directly on the framework. First, therapeutic drug monitoring, which elsewhere compensates for between-patient variability by correcting the dose after the fact, is not routinely available; the laboratory capacity constraints that limit identification and susceptibility testing limit concentration measurement as well (Hermans et al., 2024). A regimen that attains its target only with monitoring is, in that setting, a regimen that does not attain it.

Second, the relevant populations differ from those in which most population pharmacokinetic models are built, and exposure for a given regimen varies substantially between populations (Hodiamont et al., 2022; Mouton et al., 2025). A large share of the burden falls on children and neonates, for whom attainment must be derived separately (Cafaro et al., 2024; Haynes et al., 2025). Third, access rather than discovery is frequently the binding constraint on what reaches patients (Lubanga et al., 2024), and resistance mortality in the region is projected to rise steeply without intervention (Totaro et al., 2025).

Taken together these imply an additional criterion, which may be called robustness without monitoring: a candidate intended for such settings should have a feasibility window wide enough that a fixed regimen attains its target across the relevant population without individualised adjustment. This is a stricter requirement than feasibility as ordinarily assessed, and it can be evaluated at the same early stage, since it depends on the width of the predicted window rather than on any additional measurement. Candidates with narrow windows are not thereby disqualified, but they are identified as requiring an infrastructure that the intended setting may not have.

4.2.2 Limitations and directions for further work

As Section 3.6 notes, the framework is structural rather than numerical: the error magnitudes needed to quantify a particular candidate are available for some links and absent for others. Supplying them is the obvious next step, and the most valuable single addition would be a characterisation of prediction error for site penetration, which the analysis identifies as the link most often assumed and least often quantified. The framework also addresses dose feasibility only: it says nothing about resistance emergence during therapy, about combination regimens, or about the possibility that a pharmacodynamic index established for one mechanism does not transfer to a novel one, which for genuinely new classes is a substantial additional uncertainty.

5. CONCLUSION AND RECOMMENDATION

Carrying a prediction from chemical structure to a statement about human dose is not one inference but a chain of them, and the chain has two properties that determine what the exercise can deliver. Its links combine multiplicatively, so errors compound and each link weighs on the result in proportion to its variance, while any link left unquantified bounds what the result can claim. And its conclusion is a threshold judgement, so the informativeness of the whole calculation depends on how far a candidate lies from the feasibility boundary. Dose feasibility is further a two-sided window, bounded below by target attainment and above by toxicity, and a quantity properly expressed as a probability across the intended population at the intended site rather than as a value for a typical patient.

The framework proposed here makes these properties explicit, classifies each link as predicted, measured or assumed, and derives a prioritisation rule that evaluates each candidate under favourable and unfavourable assumptions. Its central asymmetry is that in silico reasoning can exclude candidates whose window is empty even under optimism, but cannot rank candidates whose windows are both non-empty; the first conclusion is robust to accumulated error and the second is not.

© IIARJournal. All rights reserved.

Four recommendations follow. Discovery programmes should propagate uncertainty through the chain rather than combining point estimates, and should report the feasibility decision under both optimistic and pessimistic assumptions rather than at the central estimate alone. The field should prioritise quantifying error in the links that are currently assumed rather than predicted, site penetration foremost, since an unquantified link bounds the credibility of the whole calculation. Predictive pharmacokinetic models intended to support prioritisation should be evaluated against decision accuracy and not only against per-endpoint accuracy. And for candidates intended for settings without routine therapeutic drug monitoring, the width of the feasibility window across the relevant population should be treated as a design criterion in its own right.

6. CONTRIBUTION TO EXISTING KNOWLEDGE

This paper makes four contributions. First, it sets out the translational chain from predicted structure-derived properties to human dose feasibility as an explicit sequence of links, classifying each by whether it is predicted, measured or assumed in early discovery and by whether its uncertainty has been quantified at all. Second, it shows that because the links compose multiplicatively and the conclusion is adjudicated at a threshold, decision reliability is weighted by the variance of every link, bounded by those that remain unquantified, and conditioned on proximity to the feasibility boundary, which reframes the improvement of predictive pharmacokinetics as a decision problem rather than an accuracy problem. Third, it argues that dose feasibility is a two-sided window rather than an efficacy floor, and derives a prioritisation rule whose asymmetry between exclusion and selection follows from the error structure rather than from caution. Fourth, it extends the construct with a robustness-without-monitoring criterion for health systems in which therapeutic drug monitoring is not routinely available, a condition that applies across much of sub-Saharan Africa and that the discovery literature does not presently represent.

7. REFERENCES

Alshaer, M. H., Barlow, B., Maranchick, N., Moser, M., Gramss, L., Burgmann, H., Jalali, V. A., Wölfl-Duchek, M., Jäger, W., Poschner, S., Plöchl, W., Reinprecht, A., Rössler, K., Gruber, A., Zeitlinger, M., Peloquin, C. A., & Hosmann, A. (2022). Meropenem population pharmacokinetics and simulations in plasma, cerebrospinal fluid, and brain tissue. *Antimicrobial Agents and Chemotherapy*, *66*(8), Article e00438-22. https://doi.org/10.1128/aac.00438-22

Alshaer, M. H., Maranchick, N., Alexander, K. M., Manigaba, K., Shoulders, B. R., Felton, T. W., Mathew, S. K., & Peloquin, C. A. (2023). Beta-lactam target attainment and associated outcomes in patients with bloodstream infections. *International Journal of Antimicrobial Agents*, *61*(3), Article 106727. https://doi.org/10.1016/j.ijantimicag.2023.106727

Alshaer, M. H., Maranchick, N., Bai, C., Maguigan, K. L., Shoulders, B., Felton, T. W., Mathew, S. K., Mardini, M. T., & Peloquin, C. A. (2022). Using machine learning to define the impact of beta-lactam early and cumulative target attainment on outcomes in intensive care unit patients with hospital-acquired and ventilator-associated pneumonia. *Antimicrobial Agents and Chemotherapy*, *66*(7), Article e00563-22. https://doi.org/10.1128/aac.00563-22

Bilal, M., Zoller, M., Fuhr, U., Jaehde, U., Ullah, S., Liebchen, U., Büsker, S., Zander, J., Babouee Flury, B., & Taubert, M. (2023). Cefepime population pharmacokinetics, antibacterial target attainment, and estimated probability of neurotoxicity in critically ill patients. *Antimicrobial Agents and Chemotherapy*, *67*(7), Article e00309-23. https://doi.org/10.1128/aac.00309-23

Cafaro, A., Stella, M., Mesini, A., Castagnola, E., Cangemi, G., Mattioli, F., & Baiardi, G. (2024). Dose optimization and target attainment of vancomycin in children. *Clinical Biochemistry*, *125*, Article 110728. https://doi.org/10.1016/j.clinbiochem.2024.110728

Charoensareerat, T., Chaijamorn, W., Kerdnimith, P., Kosumwisaisakul, N., Teeranaew, P., Rungkitwattanakul, D., Boonpeng, A., Srisawat, N., & Pattharachayakul, S. (2023). Optimal meropenem dosing regimens in patients undergoing continuous renal replacement therapy: Systematic review and Monte Carlo simulations. *Blood Purification*, *52*(6), 503–515. https://doi.org/10.1159/000529694

© IIARJournal. All rights reserved.

Einarson, K. A., Bendtsen, K. M., Li, K., Thomsen, M., Kristensen, N. R., Winther, O., Fulle, S., Clemmensen, L., & Refsgaard, H. H. F. (2023). Molecular representations in machine-learning-based prediction of PK parameters for insulin analogs. *ACS Omega*, *8*(26), 23566–23578. https://doi.org/10.1021/acsomega.3c01218

Ekins, S., Lane, T. R., Urbina, F., & Puhl, A. C. (2023). In silico ADME/tox comes of age: Twenty years later. *Xenobiotica*, *54*(7), 352–358. https://doi.org/10.1080/00498254.2023.2245049

Haynes, A. S., Wei, Z., Scheetz, M. H., Gonzalez, D., Messacar, K., Tang Girdwood, S., Peloquin, C. A., Fish, D. N., & Anderson, P. (2025). Oral cephalexin population pharmacokinetics and target attainment analysis in infants 7-60 days old. *Journal of the Pediatric Infectious Diseases Society*, *14*(10), Article piaf088. https://doi.org/10.1093/jpids/piaf088

Hermans, L. E., Centner, C. M., Morel, C. M., Mbamalu, O., Bonaconsa, C., Ferreyra, C., Lindahl, O., & Mendelson, M. (2024). Point-of-care diagnostics for infection and antimicrobial resistance in sub-Saharan Africa: A narrative review. *International Journal of Infectious Diseases*, *142*, Article 106907. https://doi.org/10.1016/j.ijid.2023.11.027

Hodiamont, C. J., van den Broek, A. K., de Vroom, S. L., Prins, J. M., Mathôt, R. A. A., & van Hest, R. M. (2022). Clinical pharmacokinetics of gentamicin in various patient populations and consequences for optimal dosing for Gram-negative infections: An updated review. *Clinical Pharmacokinetics*, *61*(8), 1075–1094. https://doi.org/10.1007/s40262-022-01143-0

Huang, S., Xu, Q., Yang, G., Ding, J., & Pei, Q. (2025). Machine learning for prediction of drug concentrations: Application and challenges. *Clinical Pharmacology & Therapeutics*, *117*(5), 1236–1247. https://doi.org/10.1002/cpt.3577

Liu, G., Catacutan, D. B., Rathod, K., Swanson, K., Jin, W., Mohammed, J. C., Chiappino-Pepe, A., Syed, S. A., Fragis, M., Rachwalski, K., Magolan, J., Surette, M. G., Coombes, B. K., Jaakkola, T., Barzilay, R., Collins, J. J., & Stokes, J. M. (2023). Deep learning-guided discovery of an antibiotic targeting Acinetobacter baumannii. *Nature Chemical Biology*, *19*(11), 1342–1350. https://doi.org/10.1038/s41589-023-01349-8

Liu, Y., Moodley, M., Pasipanodya, J. G., & Gumbo, T. (2023). Determining the delamanid pharmacokinetics/pharmacodynamics susceptibility breakpoint using Monte Carlo experiments. *Antimicrobial Agents and Chemotherapy*, *67*(4), Article e01401-22. https://doi.org/10.1128/aac.01401-22

Lubanga, A. F., Bwanali, A. N., Kambiri, F., Harawa, G., Mudenda, S., Mpinganjira, S. L., Singano, N., Makole, T., Kapatsa, T., Kamayani, M., & Ssebibubbu, S. (2024). Tackling antimicrobial resistance in sub-Saharan Africa: Challenges and opportunities for implementing the new people-centered WHO guidelines. *Expert Review of Anti-infective Therapy*, *22*(6), 379–386. https://doi.org/10.1080/14787210.2024.2362270

Mouton, J. W. A., De Clercq, A., De Paepe, P., Petrovic, M., Desmet, T., Brüggemann, R. J., Schouten, J. A., Jager, N. G. L., & De Cock, P. A. (2025). Pharmacokinetics and target attainment of teicoplanin: A systematic review. *Clinical Pharmacokinetics*, *64*(4), 467–509. https://doi.org/10.1007/s40262-025-01483-7

Naghavi, M., Vollset, S. E., Ikuta, K. S., Swetschinski, L. R., Gray, A. P., Wool, E. E., Robles Aguilar, G., Mestrovic, T., Smith, G., Han, C., Hsu, R. L., Chalek, J., Araki, D. T., Chung, E., Raggi, C., Gershberg Hayoon, A., Davis Weaver, N., Lindstedt, P. A., Smith, A. E., . . . Murray, C. J. L. (2024). Global burden of bacterial antimicrobial resistance 1990–2021: A systematic analysis with forecasts to 2050. *The Lancet*, *404*(10459), 1199–1226. https://doi.org/10.1016/S0140-6736(24)01867-1

Obrezanova, O. (2023). Artificial intelligence for compound pharmacokinetics prediction. *Current Opinion in Structural Biology*, *79*, Article 102546. https://doi.org/10.1016/j.sbi.2023.102546

Paliwal, A., Jain, S., Kumar, S., Wal, P., Khandai, M., Khandige, P. S., Sadananda, V., Anwer, Md. K., Gulati, M., Behl, T., & Srivastava, S. (2024). Predictive modelling in pharmacokinetics: From in-silico simulations to personalized medicine. *Expert Opinion on Drug Metabolism & Toxicology*, *20*(4), 181–195. https://doi.org/10.1080/17425255.2024.2330666

Rácz, A., Bajusz, D., Miranda-Quintana, R. A., & Héberger, K. (2021). Machine learning models for classification tasks related to drug safety.

© IIARJournal. All rights reserved.

Molecular Diversity, *25*(3), 1409–1424. https://doi.org/10.1007/s11030-021-10239-x

Rao, G. G., & Landersdorfer, C. B. (2021). Antibiotic pharmacokinetic/pharmacodynamic modelling: MIC, pharmacodynamic indices and beyond. *International Journal of Antimicrobial Agents*, *58*(2), Article 106368. https://doi.org/10.1016/j.ijantimicag.2021.106368

Rohani, R., Yarnold, P. R., Scheetz, M. H., Neely, M. N., Kang, M., Donnelly, H. K., Dedicatoria, K., Nozick, S. H., Medernach, R. L., Hauser, A. R., Ozer, E. A., Diaz, E., Misharin, A. V., Wunderink, R. G., & Rhodes, N. J. (2023). Individual meropenem epithelial lining fluid and plasma PK/PD target attainment. *Antimicrobial Agents and Chemotherapy*, *67*(12), Article e00727-23. https://doi.org/10.1128/aac.00727-23

Stokes, J. M., Yang, K., Swanson, K., Jin, W., Cubillos-Ruiz, A., Donghia, N. M., MacNair, C. R., French, S., Carfrae, L. A., Bloom-Ackermann, Z., Tran, V. M., Chiappino-Pepe, A., Badran, A. H., Andrews, I. W., Chory, E. J., Church, G. M., Brown, E. D., Jaakkola, T. S., Barzilay, R., & Collins, J. J. (2020). A deep learning approach to antibiotic discovery. *Cell*, *180*(4), 688–702.e13. https://doi.org/10.1016/j.cell.2020.01.021

Swanson, K., Walther, P., Leitz, J., Mukherjee, S., Wu, J. C., Shivnaraine, R. V., & Zou, J. (2024). ADMET-AI: A machine learning ADMET platform for evaluation of large-scale chemical libraries. *Bioinformatics*, *40*(7), Article btae416. https://doi.org/10.1093/bioinformatics/btae416

Totaro, V., Guido, G., Cotugno, S., De Vita, E., Asaduzzaman, M., Patti, G., Segala, F. V., Putoto, G., Frallonardo, L., Farkas, F. B., Lakatos, B., Veronese, N., Locantore, P., Di Gennaro, F., & Saracino, A. (2025). Antimicrobial resistance in sub-Saharan Africa: A comprehensive landscape review. *The American Journal of Tropical Medicine and Hygiene*, *113*(2), 253–263. https://doi.org/10.4269/ajtmh.25-0035

Xu, Y., Liu, X., Xia, W., Ge, J., Ju, C.-W., Zhang, H., & Zhang, J. Z. H. (2024). ChemXTree: A feature-enhanced graph neural network-neural decision tree framework for ADMET prediction. *Journal of Chemical Information and Modeling*, *64*(22), 8440–8452. https://doi.org/10.1021/acs.jcim.4c01186

Zhang, J., Li, H., Zhang, Y., Huang, J., Ren, L., Zhang, C., Zou, Q., & Zhang, Y. (2025). Computational toxicology in drug discovery: Applications of artificial intelligence in ADMET and toxicity prediction. *Briefings in Bioinformatics*, *26*(5), Article bbaf533. https://doi.org/10.1093/bib/bbaf533

© IIARJournal. All rights reserved.