



Original Article

ISSN : 2277-3657
CODEN(USA) : IJPRPM

Translational Pharmaceutical Research Should Optimize the Evidence Chain Rather Than Individual Experiments

Paolo Ricci^{1*}, Marco De Luca¹, Giulia Ferraro², Antonio Russo¹

¹Department of Translational Research and Evidence Chain Optimization, Faculty of Pharmacy, University of Naples Federico II, Naples, Italy.

²Department of Experiment Design and Translation, Faculty of Pharmacy, University of Bologna, Bologna, Italy.

***Email:** paolo.ricci@unina.it

ABSTRACT

Translational pharmaceutical research is commonly organized around experiments that are optimized locally: targets are selected for biological plausibility, models for mechanistic sophistication, doses for acceptable exposure, biomarkers for measurable change, and early trials for interpretable signals. Yet a technically strong result at one layer has limited developmental value when the inference connecting it to the next layer is weak. This Perspective argues that translational performance should therefore be evaluated at the level of the evidence chain rather than the individual experiment. Target evidence, mechanistic models, exposure, target engagement, biomarkers, preclinical-to-human bridging, and early clinical findings are treated as distinct but conditionally dependent evidentiary layers. Their relationships are neither automatically additive nor interchangeable. A proposed evidence-chain view emphasizes concordance, transfer boundaries, competing explanations, and the localization of weak links. It also reframes discordant human findings as potentially informative evidence about which assumption failed rather than as indiscriminate falsification of all preceding work. The resulting approach prioritizes experiments according to their capacity to resolve consequential uncertainty across connected decisions. Evidence-chain optimization does not imply that every layer requires maximal certainty. Instead, it requires that confidence, validation effort, and experimental investment be proportionate to the influence of each link on downstream development decisions.

Keywords: Translational pharmacology, Model-informed drug development, Target engagement, Biomarkers, Preclinical-to-human translation, Evidence integration

Received: 15 August 2026; **Revised:** 03 October 2026; **Accepted:** 07 October 2026

INTRODUCTION

Pharmaceutical development generates evidence through specialized experimental systems, each optimized for a particular scientific question. Target-biology programs strengthen causal hypotheses; mechanistic studies characterize pathways; pharmacokinetic analyses quantify exposure; pharmacodynamic measurements examine target or pathway response; biomarkers summarize biological change; and early clinical trials test whether selected relationships survive contact with human biology. The difficulty is that development decisions rarely depend on any one of these outputs independently. The multidimensional 5R experience at AstraZeneca illustrates how coordinated attention to target, tissue exposure, safety, patient selection, and other development dimensions can accompany improved portfolio performance [1]. The more general implication is not that one fixed checklist is universally sufficient, but that translational quality emerges from relationships among evidence domains.

This distinction matters because an experiment may be internally rigorous and still possess limited downstream decision value. Predictive-validity analyses in drug discovery emphasize that the usefulness of a screening or selection tool depends on how well its output relates to eventual clinical utility, not simply on analytical precision or performance within the experimental system [2]. Translational research therefore faces an optimization problem that differs from ordinary experimental optimization. The relevant unit is not only the assay, model, biomarker, or trial but also the inferential link connecting its result to the next decision.

This Perspective develops that argument across target evidence, mechanistic models, exposure and target engagement, biomarkers, bridging systems, and early clinical evidence. It proposes that pharmaceutical translation should be managed as a conditional evidence chain in which weak or untested links remain visible. The proposed framework does not assign universal scores or claim prospective validation. Its purpose is to identify what a result permits investigators to infer, what it does not permit them to infer, and which experiment would most efficiently discriminate between competing explanations.

The local-optimization problem in translation

Local optimization occurs when a development layer is improved according to its own performance criterion without adequately testing whether that improvement strengthens the downstream inference. A highly predictive computational model can fail if its prediction concerns a biological state that is not reproduced in humans. A compelling target can fail because therapeutically relevant tissue exposure is unattainable. Strong target occupancy can coexist with inadequate pathway modulation. A pharmacodynamic biomarker can change without establishing clinical relevance. These failures are not equivalent, yet they are often compressed into the broad statement that a candidate “did not translate.”

Model-informed drug development demonstrates why cross-layer integration can matter operationally. Portfolio analysis has associated MIDD activities with reductions in development cycle time and trial cost, although the magnitude is company- and program-specific rather than a universal effect [3]. The value comes partly from using accumulated information to alter decisions rather than merely generating another technically complete analysis.

The problem becomes more pronounced when disease and treatment responses are heterogeneous. Machine-learning approaches in quantitative pharmacology increasingly combine clinical, biological, longitudinal, and real-world information to characterize disease trajectories and variable outcomes [4]. Such methods do not remove causal uncertainty, but they illustrate why a single average model may conceal clinically consequential states.

Evidence accumulation is also temporal. Bayesian approaches provide a formal example of how prior knowledge and newly generated evidence can be integrated as development proceeds [5]. The proposed evidence-chain perspective adopts the broader learning principle without requiring every analysis to be Bayesian: evidence should remain revisable when a new result changes confidence in an upstream assumption or a downstream transfer.

Target evidence

A translational chain begins before compound exposure. The first question is whether perturbing the proposed target is sufficiently connected to the disease process to justify downstream investment. Human genetics can strengthen this inference because naturally occurring variation offers evidence linking biological perturbation to human phenotypes. Reviews of genetics-driven drug discovery document direct examples in which germline observations contributed to target identification and ultimately to approved therapies [6]. These successes, however, do not mean that genetic association by itself establishes pharmacological tractability, an appropriate direction or magnitude of perturbation, therapeutic index, or treatment benefit.

Target confidence therefore requires reconciliation across evidence types. Platforms such as Open Targets integrate genetics, genomics, functional perturbation, expression, tractability, and other evidence to prioritize target–disease hypotheses [7]. Convergence is informative because independent evidence streams can reduce reliance on one potentially misleading observation. Discordance is equally informative. A genetically supported target with weak pharmacological tractability presents a different development problem from a tractable target with uncertain disease causality.

Knowledge maturity adds another boundary. Mapping of the human proteome has shown that substantial portions remain comparatively understudied despite potential biological or chemical tractability [8]. Lack of accumulated evidence should not automatically be interpreted as negative evidence, but neither should novelty be substituted for validation. The proposed evidence-chain view therefore distinguishes target plausibility, causal support, tractability, and evidentiary maturity rather than collapsing them into a single target-confidence label.

A target-layer result becomes translationally useful when investigators can specify which uncertainty it reduces and what evidence must follow. A target supported by human genetics may justify stronger mechanistic investment; it does not justify assuming adequate exposure or engagement. Conversely, poor downstream pharmacology does not retroactively erase the biological association. This separation allows failures to be localized rather than indiscriminately assigned to the target.

The practical consequence is an exception-first view of target evidence. Before advancing a candidate because one upstream result is compelling, development teams should ask what downstream pattern would reveal that the apparent target advantage was misleading and what experiment could distinguish the competing explanations.

Table 1 organizes these failure states without imposing a numerical ranking.

Table 1. Target-Evidence Failure States That Produce Different Downstream Translational Diagnoses

Observed configuration		Mechanism and competing diagnosis	Corrective discrimination
Target-evidence condition / Exposure support	Biomarker pattern	Misleading or failure pattern / Target-engagement state	Discriminating test / Residual development consequence
Target-evidence condition: Causal support comparatively strong; perturbation strategy unresolved Exposure support: May be adequate or unknown	Downstream signal may be absent because engagement was never achieved	Misleading or failure pattern: Strong disease association but uncertain druggability Target-engagement state: May remain technically inaccessible	Discriminating test: Demonstrate feasible, selective perturbation at relevant tissue and magnitude Residual development consequence: Target remains plausible, but candidate strategy is not yet validated
Target-evidence condition: Model evidence strong; human relevance uncertain Exposure support: Exposure may appear adequate	Biomarker may move in model systems	Misleading or failure pattern: Compelling preclinical target effect with weak human causal support Target-engagement state: Engagement can occur	Discriminating test: Test target–disease relationship in human-relevant systems or stratified human data Residual development consequence: Positive pharmacology cannot by itself establish human disease relevance
Target-evidence condition: Upstream target hypothesis remains viable Exposure support: Systemic or tissue exposure supported	Biomarker response absent or ambiguous	Misleading or failure pattern: Adequate exposure with absent engagement Target-engagement state: Expected engagement not demonstrated	Discriminating test: Measure direct occupancy, proximal pharmacodynamics, or target-proximal effect Residual development consequence: Candidate, assay, tissue penetration, or target-accessibility problem must be separated
Target-evidence condition: Target is reached Exposure support: Exposure supported	Expected pathway biomarker unchanged	Misleading or failure pattern: Engagement without expected downstream biology Target-engagement state: Engagement demonstrated	Discriminating test: Test pathway dependency, compensatory biology, assay timing, and biomarker validity Residual development consequence: Mechanistic link downstream of engagement becomes the dominant uncertainty
Target-evidence condition: Target hypothesis incompletely localized Exposure support: Exposure compatible	Biological signal changes	Misleading or failure pattern: Biomarker response without secure target specificity Target-engagement state: Engagement may be indirect or uncertain	Discriminating test: Use orthogonal engagement or perturbation evidence and competing-mechanism controls Residual development consequence: Biomarker cannot yet be interpreted as target-mediated evidence

Target-evidence condition: Target evidence appears coherent within selected models or populations Exposure support: Supported in tested context	Consistent within tested context	Misleading or failure pattern: Multiple upstream signals agree but remain context-limited Target-engagement state: Supported in tested context	Discriminating test: Challenge the chain across relevant human state, population, dose, or disease context Residual development consequence: Confidence remains conditional rather than universally transferable
--	----------------------------------	---	---

Note: The failure states constitute a proposed evidence-chain triage derived from the manuscript's synthesis. They are not a validated scoring system, and no row implies that every listed test is required for every program.

Mechanistic models

Mechanistic models are most valuable when they expose relationships that cannot be inferred reliably from isolated measurements. QSP models, for example, can combine disease biology, drug action, temporal dynamics, and patient heterogeneity, including in settings such as rare diseases where clinical data may be sparse. Their value, however, depends on whether the modeled mechanisms and parameterizations are credible for the decision being made. Risk-informed model-evaluation frameworks consequently tie credibility to context of use, model influence, verification, validation, and decision consequence [9, 10].

This distinction prevents mechanistic sophistication from becoming a substitute for evidence. A model can be internally coherent while containing a structurally incorrect biological assumption. It can reproduce observed data because several parameter combinations produce similar outputs. It can also be valid for interpolation within one development state yet unreliable for extrapolation to another population, dose, or endpoint.

Recent virtual-population methods illustrate both the opportunity and the boundary. In oncology QSP, calibration procedures can reconcile mechanistic tumor-dynamics outputs with progression-free-survival information and censoring [11]. Such calibration strengthens consistency between evidence layers, but goodness of fit does not by itself establish that the represented mechanism is uniquely correct or externally predictive.

The same caution applies to increasingly individualized simulations. “Digital twin” terminology can describe sophisticated updating of patient-level predictions, yet individualized modeling remains constrained by classical problems such as parameter identifiability and data informativeness [12]. In an evidence chain, model complexity therefore earns influence only when the evidence supporting its intended use is proportionate to the consequence of the decision.

Exposure and target engagement

Target plausibility and mechanistic coherence become clinically relevant only if a pharmacological intervention reaches the required biological compartment under tolerable conditions. Exposure should therefore be separated into at least systemic exposure, relevant-tissue exposure, and target engagement when the biology demands those distinctions. PET provides a clear example in neuroscience: appropriately designed imaging can establish whether a compound enters the brain and can quantify target occupancy, answering questions that plasma pharmacokinetics alone cannot resolve [13]. Even then, occupancy is evidence of engagement, not evidence of therapeutic benefit.

A second problem concerns the target value assigned to exposure itself. Precision-dosing scholarship has questioned whether population-level concentration targets necessarily define optimal exposure for individual patients and has argued for stronger individual-level response biomarkers where feasible [14]. The implication is not that population targets are generally invalid, but that exposure targets have meaning only through the response relationship and population for which they were established.

Integrated early-development analyses show how these layers can be connected without treating them as interchangeable. For the ATR inhibitor tuvusertib, population pharmacokinetic, pharmacodynamic, safety, and translational efficacy analyses were combined to inform recommended dose and schedule selection [15]. That example supports the use of convergent evidence, not the transfer of its specific dose, biomarker, or exposure relationships to other compounds.

Human exposure can sometimes be introduced earlier to reduce uncertainty in the bridge itself. Phase-0 and microdosing approaches can provide human pharmacokinetic and selected pharmacodynamic or mechanistic information before conventional therapeutic development [16]. Their usefulness depends on whether subtherapeutic behavior can inform the therapeutic setting; nonlinear pharmacokinetics, saturable processes, or dose-dependent biology can limit that extrapolation.

Thus, “adequate exposure” should not operate as a binary gate. The evidence chain asks a more specific sequence: was the relevant compartment exposed, was the intended target engaged, was the magnitude and duration

sufficient for the proposed mechanism, and was the resulting biological change interpretable in the intended patient context?

Biomarkers

Biomarkers often appear to provide the shortest path between pharmacology and clinical interpretation because they convert complex biology into measurable signals. Their inferential value, however, depends on what decision they are intended to support. Biomarkers used for diagnosis, prognosis, pharmacodynamic monitoring, prediction, safety, or potential surrogate roles do not carry interchangeable meanings. Contemporary clinical and regulatory analyses therefore emphasize context of use and fit-for-purpose analytical and clinical validation [17, 18].

This creates a central bottleneck. A measurable change can indicate that a biological process responded without demonstrating that the response is target-mediated, sufficiently durable, relevant to disease modification, or predictive of clinical benefit. Conversely, failure to change a distal biomarker may reflect poor engagement, incorrect timing, insensitive measurement, pathway redundancy, or an invalid biomarker rather than failure of the target itself.

Target engagement and biomarker response must therefore remain separate evidentiary states. PET occupancy evidence illustrates how engagement can be demonstrated directly while leaving the clinical consequence unresolved [13]. The proposed evidence-chain perspective adds an explicit interpretation boundary between engagement and biomarker meaning, and another between biomarker meaning and clinical relevance.

Figure 1 is needed because these distinctions are relational rather than sequential: several evidence lanes can be individually positive while the translation bottleneck persists at the boundary connecting them.




Figure 1. Biomarker Meaning at the Translational Bottleneck Between Engagement and Clinical Relevance

Preclinical-to-human bridging

The preclinical-to-human bridge is often described as though one model must predict the patient. A more defensible view treats bridging as reconciliation among models that each preserve different aspects of human biology. New approach methodologies can add human-relevant safety evidence when conventional species models are limited; organoids can reproduce selected features of human tissue architecture and response; and model-informed methods can integrate sparse nonclinical and clinical information in settings such as rare diseases [19-21]. None represents a complete human organism, and none eliminates the need to establish which relationship is actually transferable.

Bridging confidence should therefore depend on concordance across evidence that fails for different reasons. An animal model may preserve whole-organism disposition but express a different target or pathway state. An organoid may preserve human cellular architecture but lack systemic exposure, immune, endocrine, or multicompartment interactions. A computational model may integrate multiple processes but remain sensitive to structure, parameter identifiability, and training data.

The same restriction applies to in silico discovery more broadly. Computational methods can improve candidate generation and prioritization, yet reviews of preclinical translation continue to identify validation, data quality, workflow integration, and domain-shift problems between computational performance and downstream biological utility [22]. A high model metric therefore answers a different question from successful perturbation in human-relevant tissue or interpretable response in patients.

The credibility principle developed for mechanistic models can be extended cautiously across the bridge: evidence requirements should increase when an uncertain model or assay has greater influence over a consequential downstream decision [10]. This extension is part of the proposed synthesis rather than an established universal regulatory rule.

Table 2 reconciles bridge patterns by asking what concordant and discordant combinations actually permit investigators to infer.

Table 2. Reconciliation States for Preclinical Evidence Before Human Translation Is Claimed

Evidence configuration		Interpretive alternatives	Adjudication requirement
Governing bridge condition / Model or patient context	Exposure support / Engagement evidence / Biomarker readout	Inference not justified	Inference justified / Discriminating evidence still needed
Governing bridge condition: Model response is positive but human exposure state is unresolved Model or patient context: Animal, organoid, or computational system	Exposure support: Supported only in model Engagement evidence: Model-specific Biomarker readout: Positive in model	Human mechanism is established	Inference justified: Mechanism is testable under model conditions Discriminating evidence still needed: Human PK/tissue exposure and appropriate engagement evidence
Governing bridge condition: Human-relevant model and animal evidence converge Model or patient context: Complementary systems with different failure modes	Exposure support: Consistent preclinically Engagement evidence: Consistent preclinically Biomarker readout: Directionally concordant	Clinical efficacy is established	Inference justified: Bridge confidence is strengthened Discriminating evidence still needed: Human exposure, engagement, and early-response evidence
Governing bridge condition: Exposure translates but biomarker does not Model or patient context: Early human setting	Exposure support: Human exposure supported Engagement evidence: Present or uncertain Biomarker readout: Expected signal absent	Target or disease mechanism is disproven	Inference justified: Drug reaches relevant compartment under tested conditions Discriminating evidence still needed: Direct engagement, timing, assay sensitivity, pathway redundancy
Governing bridge condition: Engagement is demonstrated but distal response differs across systems Model or patient context: Model and human evidence discordant	Exposure support: Supported Engagement evidence: Supported Biomarker readout: Discordant	Downstream mechanism necessarily transfers	Inference justified: Target is pharmacologically reached Discriminating evidence still needed: Proximal-to-distal pathway testing and human-context reassessment
Governing bridge condition: Human-derived system predicts response but whole-body determinants are incomplete	Exposure support: Local exposure modeled Engagement evidence: Potentially demonstrable	Patient-level efficacy or safety is established	Inference justified: Human tissue can respond under defined conditions Discriminating evidence still needed: Systemic disposition,

Model or patient context: Organoid/NAM	Biomarker readout: Positive	multicompartment physiology, clinical evidence	
Governing bridge condition: Computational and experimental evidence agree Model or patient context: In silico plus experimental system	Exposure support: Depends on modeled context	External	Inference justified: Independent evidence supports the hypothesis
	Engagement evidence: Depends on assay	transportability is proven	Discriminating evidence still needed: Prospective validation in a new biological or clinical context
	Biomarker readout: Concordant		

Note: The reconciliation states are proposed analytical categories. They specify permissible inference rather than numerical evidence grades.

Early clinical evidence

Early clinical development is the point at which several assumptions that appeared individually credible can be tested together in humans. Its value is therefore greater than simply confirming a preclinical prediction. A human result can identify which relationship in the preceding chain does not survive translation.

The development of the glucose-responsive insulin MK-2640 illustrates this diagnostic function. Model-informed analyses supported rapid progression and interpretation, yet the preclinical glucose-dependent systemic-clearance behavior was not reproduced in people with type 1 diabetes, while a glucose-dependent pharmacodynamic signal remained detectable [23]. The result was neither complete confirmation nor indiscriminate failure. Instead, it separated a failed pharmacokinetic mechanistic expectation from retained evidence at another pharmacodynamic layer and supported reverse translation toward the underlying mechanism.

Dose selection presents the same problem at a different scale. Contemporary oncology development increasingly recognizes that an optimal dose cannot be inferred from tolerability alone and that the evidence required depends on molecular class, pharmacology, exposure–response relationships, biological activity, and safety [24]. A dose can therefore be locally acceptable while the evidence chain supporting its long-term use remains incomplete.

Regulatory experience demonstrates that these uncertainties can propagate downstream. Retrospective analysis of oncology approvals found that factors including maximum-tolerated-dose labeling, discontinuation burden, and exposure–safety relationships were associated with greater likelihood of postmarketing dose-optimization requirements or commitments [25]. These associations do not establish causation, but they show that unresolved early dose questions can remain developmentally consequential after approval.

Transfer across treatment contexts introduces another boundary. Exposure–response and dose–response relationships may change when a drug enters a new combination, tumor type, population, or treatment setting [26]. Prior evidence may remain informative, but portability itself becomes a hypothesis requiring assessment rather than an assumption.

The proposed evidence-chain perspective therefore treats early clinical evidence as both a forward decision input and a backward diagnostic signal. Concordance increases confidence only in the links actually tested. Discordance should trigger localization: exposure failure, inadequate engagement, incorrect mechanistic expectation, invalid biomarker interpretation, context mismatch, or an alternative biological explanation. **Figure 2** makes this diagnostic topology explicit because a simple success/failure pipeline would conceal the information contained in partial translation.




Figure 2. Early Human Discordance as a Diagnostic Test of the Translational Evidence Chain

Dependency between evidence layers

Evidence layers in translation are not independent checkpoints. Their interpretability depends on assumptions carried forward from earlier stages. A biomarker is useful only if the measured process is connected to the intended mechanism and decision; a mechanistic model is influential only to the extent that its structure and parameters are credible for its context of use; and a safety prediction matters differently when it guides exploratory prioritization than when it constrains a clinical dose. Industry experience with quantitative systems toxicology demonstrates this heterogeneity: models are already used across preclinical and clinical decision contexts, but validation practices, uncertainty analyses, resources, and model influence vary substantially [27].

Dependency also means that discordance should be localized rather than averaged away. In the MK-2640 program, failure of a predicted glucose-responsive systemic-clearance relationship in humans did not imply that every preclinical or pharmacodynamic observation was false [23]. Similarly, a biomarker's evidentiary meaning remains conditional on its context of use rather than on the simple fact that it changes [17]. The proposed evidence-chain interpretation therefore treats each connection as an inference that can be supported, weakened, or reopened independently.

Weak-link and bottleneck analysis

A translational chain can contain many technically strong components while remaining decisionally fragile because one critical connection is weak. Predictive-validity theory provides the underlying logic: a tool contributes little to candidate selection when its output is poorly related to the downstream outcome that ultimately matters, regardless of its internal technical performance [2]. The proposed weak-link principle extends this reasoning to connected translational evidence. The bottleneck is not necessarily the least precise experiment; it is the unresolved link whose failure would most alter the interpretation of downstream evidence.

Weak links can persist across development stages. In oncology, unresolved early dose relationships have been associated with later postmarketing dose-optimization requirements or commitments [25]. Likewise, widespread use of a quantitative model does not itself establish that uncertainty, identifiability, or validation have been addressed sufficiently for every intended decision [27]. Bottleneck analysis should therefore identify the evidence dependency with the greatest combination of uncertainty and decision influence, then ask whether additional experimentation can discriminate among the plausible failure explanations.

An evidence-chain optimization framework

The central contribution of this Perspective is a proposed Evidence-Chain Optimization Framework. Its objective is not to maximize certainty at every stage. Instead, it seeks sufficient confidence across the dependencies required for the next consequential decision. The multidimensional logic resembles earlier evidence that coordinated development criteria can outperform isolated attention to individual attributes [1], but the framework developed here focuses specifically on inferential continuity between evidence layers.

Three principles govern the framework. First, confidence is conditional: model, assay, biomarker, and bridge evidence should be evaluated for the context in which they influence a decision, consistent with risk-informed model-credibility approaches [10]. Second, evidence is sequential: new information should update confidence in both downstream predictions and upstream assumptions rather than simply being appended to an accumulating dossier [5]. Third, discordance is diagnostic: a failed transition should trigger targeted reassessment of the specific connection rather than automatic rejection of the entire chain.

Because these transitions are easier to compare as a sequence than as prose, **Table 3** records the proposed decision logic from precursor evidence to a development endpoint.

Table 3. Sequence-Aware Decision Ledger for Evidence-Chain Optimization

State and entry conditions	Evidence through time	Transition and reassessment	
Precursor evidentiary state	Transition trigger / Observable readout	Failure criterion / Downstream decision consequence	Transfer boundary or interruption condition / framework endpoint
Target hypothesis considered sufficiently plausible	Transition trigger: Selection of a pharmacological perturbation strategy Observable readout: Evidence that the candidate can alter the intended target under relevant conditions	Failure criterion: Target cannot be selectively or adequately perturbed Downstream decision consequence: Reopen modality, chemistry, or target strategy	Transfer boundary or interruption condition: Disease context, tissue accessibility, direction or magnitude of perturbation differs Proposed framework endpoint: Proceed only when target evidence and perturbability are jointly adequate
Mechanistic representation supports a predicted response	Transition trigger: Candidate advances toward quantitative translation Observable readout: Model predictions remain consistent with independent biological or pharmacological observations	Failure criterion: Fit depends on poorly identified or context-specific assumptions Downstream decision consequence: Reduce model influence or obtain discriminating data	Transfer boundary or interruption condition: Population, dose, tissue, endpoint, or mechanism differs from calibration context Proposed framework endpoint: Use model according to demonstrated credibility for the decision
Exposure is predicted to support pharmacology	Transition trigger: Human or human-relevant exposure becomes measurable Observable readout: Relevant-compartment exposure and/or direct engagement	Failure criterion: Exposure achieved without engagement, or engagement cannot be demonstrated Downstream decision consequence: Diagnose penetration, accessibility, potency, assay, or mechanism	Transfer boundary or interruption condition: Systemic exposure is not equivalent to tissue exposure or occupancy Proposed framework endpoint: Advance the chain only across the exposure–engagement relation actually demonstrated
Engagement is established	Transition trigger: Biological response is evaluated Observable readout: Target-proximal or pathway biomarker behaves as expected	Failure criterion: Engagement occurs without interpretable downstream response Downstream decision consequence: Reassess pathway dependency, timing, redundancy, or biomarker validity	Transfer boundary or interruption condition: Biomarker context of use does not support the intended inference Proposed framework endpoint: Preserve engagement evidence while reopening the engagement-to-response link

Preclinical layers appear concordant	Transition trigger: Early clinical testing becomes informative Observable readout: Human PK, PD, safety, biomarker, or activity evidence	Failure criterion: Human pattern contradicts a material bridge assumption Downstream decision consequence: Localize discordance and perform reverse translation	Transfer boundary or interruption condition: Species, model, population, dose, disease state, or treatment setting changes Proposed framework endpoint: Update only the links directly tested by human evidence
Early clinical evidence supports continued development	Transition trigger: Program moves to a new regimen, population, indication, or combination Observable readout: Evidence remains directionally and quantitatively interpretable in the new setting	Failure criterion: Previous exposure–response or benefit–risk relationship no longer transfers Downstream decision consequence: Reopen dose, biomarker, or mechanistic assumptions	Transfer boundary or interruption condition: Material clinical or biological context change Proposed framework endpoint: Re-establish sufficient evidence-chain continuity for the new decision

Note: The sequence, transition states, and framework endpoints are proposed conceptual synthesis. They are not validated evidence grades, regulatory criteria, or numerical decision thresholds.

Development decision rules

The framework changes the purpose of the next experiment. Additional evidence is most valuable when it can change a consequential decision by resolving uncertainty at an influential link. Portfolio experience with MIDD supports the broader proposition that integrated quantitative evidence can alter development efficiency [3], but the proposed rule is more specific: experimental value should be judged partly by whether a result can discriminate between plausible explanations for a chain bottleneck.

Decision sufficiency must also remain context dependent. Oncology dose-optimization work illustrates why mechanism, molecular class, exposure–response, activity, and safety can require different evidentiary emphasis [24]. Evidence portability likewise should be reassessed when a drug moves into a materially different combination, indication, patient population, or treatment setting because dose–response and exposure–response relationships may change [26].

Accordingly, three proposed rules follow. Proceed when uncertainty in remaining links is unlikely to change the next decision materially. Interrogate when a discriminating experiment could identify which high-influence link is responsible for discordance. Reopen an earlier assumption when new human or contextual evidence makes its original transfer condition no longer credible. These are development-management rules, not claims of prospective predictive performance or regulatory acceptance.

CONCLUSION

Translational pharmaceutical research is weakened when experiments are optimized as isolated achievements rather than as components of a connected inferential system. Target validity, mechanistic coherence, exposure, engagement, biomarker response, preclinical bridging, and early clinical evidence answer different questions and should retain their evidentiary boundaries.

The proposed evidence-chain perspective shifts attention from whether each experiment was technically successful to whether the relationships required for the next development decision are sufficiently supported. It treats discordance as information that can localize a weak link, directs validation effort toward influential uncertainties, and requires transferability to be reconsidered when biological or clinical context changes. The framework remains conceptual and requires prospective evaluation, but it provides a disciplined basis for designing translational programs around the integrity of the evidence chain rather than the local performance of its individual experiments.

ACKNOWLEDGMENTS: None

CONFLICT OF INTEREST: None

FINANCIAL SUPPORT: None

ETHICS STATEMENT: None

REFERENCES

1. Morgan P, Brown DG, Lennard S, Anderton MJ, Barrett JC, Eriksson U, et al. Impact of a five-dimensional framework on R&D productivity at AstraZeneca. *Nat Rev Drug Discov.* 2018;17(3):167-81. doi:10.1038/nrd.2017.244
2. Scannell JW, Bosley J, Hickman JA, Dawson GR, Truebel H, Ferreira GS, et al. Predictive validity in drug discovery: What it is, why it matters and how to improve it. *Nat Rev Drug Discov.* 2022;21(12):915-31. doi:10.1038/s41573-022-00552-x
3. Sahasrabudhe V, Nicholas T, Nucci G, Musante CJ, Corrigan B. Impact of model-informed drug development on drug development cycle times and clinical trial cost. *Clin Pharmacol Ther.* 2025;118(2):378-85. doi:10.1002/cpt.3636
4. Terranova N, Venkatakrishnan K. Machine learning in modeling disease trajectory and treatment outcomes: An emerging enabler for model-informed precision medicine. *Clin Pharmacol Ther.* 2024;115(4):720-6. doi:10.1002/cpt.3153
5. Ruberg SJ, Beckers F, Hemmings R, Honig P, Irony T, LaVange L, et al. Application of Bayesian approaches in drug development: Starting a virtuous cycle. *Nat Rev Drug Discov.* 2023;22(3):235-50. doi:10.1038/s41573-023-00638-0
6. Trajanoska K, Bhérer C, Taliun D, Zhou S, Richards JB, Mooser V. From target discovery to clinical drug development with human genetics. *Nature.* 2023;620(7975):737-45. doi:10.1038/s41586-023-06388-8
7. McDonagh EM, Trynka G, McCarthy M, Holzinger ER, Khader S, Nakic N, et al. Human genetics and genomics for drug target identification and prioritization: Open Targets' perspective. *Annu Rev Biomed Data Sci.* 2024;7(1):59-81. doi:10.1146/annurev-biodatasci-102523-103838
8. Oprea TI, Bologa CG, Brunak S, Campbell A, Gan GN, Gaulton A, et al. Unexplored therapeutic opportunities in the human genome. *Nat Rev Drug Discov.* 2018;17(5):317-32. doi:10.1038/nrd.2018.14
9. Neves-Zaph S, Kaddi C. Quantitative systems pharmacology models: Potential tools for advancing drug development for rare diseases. *Clin Pharmacol Ther.* 2024;116(6):1442-51. doi:10.1002/cpt.3451
10. Musuamba FT, Skottheim Rusten I, Lesage R, Russo G, Bursi R, Emili L, et al. Scientific and regulatory evaluation of mechanistic in silico drug and disease models in drug development: Building model credibility. *CPT Pharmacometrics Syst Pharmacol.* 2021;10(8):804-25. doi:10.1002/psp4.12669
11. Braniff N, Joshi T, Cassidy T, Trogon M, Kumar R, Poels K, et al. An integrated quantitative systems pharmacology virtual population approach for calibration with oncology efficacy endpoints. *CPT Pharmacometrics Syst Pharmacol.* 2025;14(2):268-78. doi:10.1002/psp4.13270
12. Houk BE. From classical PKPD to model-informed drug development: Are "digital twins" a new paradigm? *Clin Pharmacol Ther.* 2026;120(3):590-1. doi:10.1002/cpt.70340
13. Varrone A, Bundgaard C, Bang-Andersen B. PET as a translational tool in drug development for neuroscience compounds. *Clin Pharmacol Ther.* 2022;111(4):774-85. doi:10.1002/cpt.2548
14. Polasek TM, Peck RW. Beyond population-level targets for drug concentrations: Precision dosing needs individual-level targets that include superior biomarkers of drug responses. *Clin Pharmacol Ther.* 2024;116(3):602-12. doi:10.1002/cpt.3197
15. Mukker JK, Diderichsen PM, Hellmann F, Yap TA, Plummer R, Tolcher AW, et al. Integrated population pharmacokinetic, pharmacodynamic, and safety analyses to inform dosage selection in the clinical development of the ATR inhibitor tivantinib. *Clin Pharmacol Ther.* 2026;119(3):618-28. doi:10.1002/cpt.70029
16. Burt T, Young G, Lee W, Kusuvara H, Langer O, Rowland M, et al. Phase 0/microdosing approaches: Time for mainstream application in drug development? *Nat Rev Drug Discov.* 2020;19(11):801-18. doi:10.1038/s41573-020-0080-x
17. Cummings JL, Teunissen CE, Fiske BK, Le Ber I, Wildsmith KR, Schöll M, et al. Biomarker-guided decision making in clinical drug development for neurodegenerative disorders. *Nat Rev Drug Discov.* 2025;24(8):589-609. doi:10.1038/s41573-025-01165-w
18. Schuck RN, Sekar V. Use of biomarkers in drug development for regulatory purposes. *Clin Transl Sci.* 2025;18(10):e70377. doi:10.1111/cts.70377

19. Beilmann M, Adkins K, Boonen HCM, Hewitt P, Hu W, Mader R, et al. Application of new approach methodologies for nonclinical safety assessment of drug candidates. *Nat Rev Drug Discov.* 2025;24(9):705-25. doi:10.1038/s41573-025-01182-9
20. Wang D, Villenave R, Stokar-Regenscheit N, Clevers H. Human organoids as 3D in vitro platforms for drug discovery: Opportunities and challenges. *Nat Rev Drug Discov.* 2026;25(3):204-26. doi:10.1038/s41573-025-01317-y
21. Mitra A, Tania N, Ahmed MA, Rayad N, Krishna R, Albusaysi S, et al. New horizons of model informed drug development in rare diseases drug development. *Clin Pharmacol Ther.* 2024;116(6):1398-411. doi:10.1002/cpt.3366
22. Markossian S, Baskir HM, Grossman AC, Chandran S, Rele S, Hall MD, et al. In silico drug discovery: Bridging the gaps in preclinical translation. *Drug Discov Today.* 2026;31(1):104567. doi:10.1016/j.drudis.2025.104567
23. Visser SAG, Kandala B, Fancourt C, Krug AW, Cho CR. A model-informed drug discovery and development strategy for the novel glucose-responsive insulin MK-2640 enabled rapid decision making. *Clin Pharmacol Ther.* 2020;107(6):1296-311. doi:10.1002/cpt.1729
24. Zhu J, Schroeder A, Frank S, Boetsch C, Jamois C, Kassir N, et al. Oncology dose optimization: Tailored approaches to different molecular classes. *Clin Pharmacol Ther.* 2025;118(1):74-9. doi:10.1002/cpt.3658
25. Kitagaki H, Takeda K, Murai K, Maeda H. Dose optimization in oncology drug development: Risk factors for postmarketing requirements and commitments. *Clin Pharmacol Ther.* 2025;118(4):823-30. doi:10.1002/cpt.3722
26. Gao W, Liu J, Musante CJ, Zhu H, Thompson M, Shah M, et al. Using quantitative approaches to optimize dosages for new combinations and subsequent indications for oncology drugs. *Clin Pharmacol Ther.* 2025;118(4):772-7. doi:10.1002/cpt.3725
27. Beattie KA, Verma M, Brennan RJ, Clausznitzer D, Damian V, Leishman D, et al. Quantitative systems toxicology modeling in pharmaceutical research and development: An industry-wide survey and selected case study examples. *CPT Pharmacometrics Syst Pharmacol.* 2024;13(12):2036-51. doi:10.1002/psp4.13227