The AI Drug Discovery Audit

Recursion acquired Exscientia in 2024. An AI-designed drug completed Phase II. FDA reviewers are asking about AI methodology in IND meetings. The Phase II allocation decision your R&D committee is making has three governance questions embedded in it that the current process is not structured to surface.

01

Background

How AI drug discovery moved from speculative to operational — and what the platforms actually do

01 / BACKGROUND

The AI Drug Discovery Market

- Named partnerships at scale: Recursion/Roche ($150M), Recursion/Bayer ($100M+), Sanofi/Exscientia ($1.2B — now Recursion post-2024 acquisition)
- Key platforms: Recursion (phenotypic screening + ML), Insilico Medicine (generative chemistry + pathway AI), Schrödinger (physics-based ML), BenevolentAI (AstraZeneca partnership)
- Milestone: INS018_055 by Insilico Medicine completed Phase II (2024) — first AI-designed drug with Phase II human trial data
- BenevolentAI correctly identified baricitinib as a COVID-19 treatment candidate before clinical validation — prediction proved correct, FDA-approved
- Platform economics: AI screens millions of compound-target interactions per week vs. thousands via traditional HTS — pace advantage is the investment thesis, not proven clinical superiority

01 / BACKGROUND

What AI Drug Discovery Actually Does

- Recursion: phenotypic screening at scale — photographs cells treated with compounds, CNNs identify biological effect patterns; processes millions of readouts/week
- Insilico: generative chemistry designs novel molecular structures + biological pathway analysis from published literature and proprietary assay data
- Schrödinger: physics-based free energy perturbation calculations accelerated by ML — predicts binding affinity to reduce synthesis cycles
- Common mechanism: pattern recognition against training data — the AI predicts; wet lab biology confirms or refutes
- Critical gap: platforms do not run the biology — confidence scores are statistical outputs from pattern matching, not validated biological hypotheses

01 / BACKGROUND

The Training Data Problem

- Training data sources: ChEMBL, PubMed, PubChem, Protein Data Bank + proprietary experimental datasets
- Publication bias: journals publish positive results at 3–4x the rate of negative results — the AI learns from a systematically biased record
- Rare disease gap: training data scales with historical research attention — novel targets and rare disease biology have thin training signals but confidence scores do not visibly reflect this
- Confidence score interpretation: reflects pattern matching against a biased dataset, not a validated biological hypothesis — calibration depends on training data coverage for the specific target class
- Governance implication: R&D leaders committing Phase II budgets to AI confidence scores are trusting a number calibrated against a dataset they cannot inspect

02

Three Governance Questions in Every Phase II Allocation Decision

1. What validation does your organization require between an AI confidence score and a Phase II budget authorization — and is that validation designed to challenge the AI's hypothesis or confirm it?

2. What does your IND say about how the candidate was identified — and has your regulatory team prepared that documentation before the pre-IND meeting, not after a clinical hold query?

3. What are your rights to model documentation and data access if the platform partner is acquired or restructured — as Exscientia was by Recursion in 2024?

03

Four Governance Postures

Option A

Standard biology validation — apply existing Phase II criteria unchanged

Status quo. Does not address the FDA documentation gap or the challenge-vs-confirm validation problem.

Option B

Recommended

AI validation protocol — require mechanistic hypothesis + training data documentation before Phase II

Highest-leverage governance action. Addresses all three decision-point questions.

Option C

Build internal AI capability — own the model, control the training data

Governance-complete but capital-intensive. Right for top-5 pharma with internal data assets at scale.

Option D

Hybrid: external platforms + internal AI hypothesis review team

Balances platform economics with internal oversight. Most credible FDA documentation posture for mid-size pharma.

04

Develop the AI Validation Protocol Before the Next Phase II Meeting

- Require an independent mechanistic hypothesis from your pharmacology team before validation biology begins — not a restatement of the confidence score, a falsifiable biological claim
- Require the platform partner to document training data coverage depth for each candidate's target class before Phase II authorization — build this into the partnership SLA at renewal
- Assign a regulatory lead to prepare AI methodology documentation for every IND involving an AI-generated candidate — before the pre-IND meeting, not after the clinical hold query
- Review partnership agreements for data portability, model documentation entitlements, and change-of-control provisions before next renewal

05

Five Material Risks

1.

Training data bias — confidence scores inherit publication bias and rare-target data gaps

Negative results are underrepresented 3–4x in training data. Novel targets have thin training signals the interface does not surface. Confidence scores are calibrated against a dataset with known systematic gaps.

2.

FDA documentation gap — no established standard for AI-generated candidate IND methodology

Reviewers are asking. No formal requirement exists yet. First sponsors with credible documentation set the informal industry standard. Unprepared sponsors get clinical holds at the highest-cost point in development.

3.

Platform acquisition risk — Exscientia acquired by Recursion 2024; partner restructuring is not hypothetical

Standard software licensing does not cover data portability, model documentation, or change-of-control protections. AI drug discovery partnerships need bespoke agreement language for these scenarios.

4.

Confirmation bias in validation biology — designed to advance promising candidates, not falsify AI hypotheses

Validation programs run by teams incentivized to advance candidates systematically under-test disconfirming conditions. Phase II failures result from biology the AI assumed but the validation did not challenge.

5.

Phase II attrition for reasons outside the AI's predictive scope

AI predicts compound-target biology. It does not predict clinical outcomes in heterogeneous patient populations. High AI conviction increases risk that early stopping signals are rationalized rather than recognized.

06

Six Questions Before the Next Phase II Allocation Meeting

1. Does your organization have a documented protocol specifying what validation is required between an AI confidence score and a Phase II budget authorization — and does it require an independent mechanistic hypothesis from your pharmacology team?
2. What does your current IND template include about AI methodology when a candidate was identified through a platform — and has your regulatory team prepared that section before a pre-IND meeting rather than after a clinical hold query?
3. Do your AI drug discovery partnership agreements include data portability clauses, model documentation entitlements, and change-of-control provisions — and when were they last reviewed?
4. What is the training data coverage depth for the specific target classes in your current AI-assisted pipeline — and has your platform partner documented known gaps that affect confidence score calibration?
5. How does your validation biology for AI-generated candidates differ from traditionally-identified candidates — specifically, does it include challenge experiments designed to falsify the AI's hypothesis?
6. If your AI platform partner was acquired tomorrow, what would your organization receive in terms of data export and model documentation — and is any of that specified in your current agreement?

AI INSIGHT LAB

The AI Drug Discovery Audit

The platform economics are compelling and the investment is already committed. The governance architecture — validation protocol, FDA documentation, partner risk — has not been built at the same pace as the pipeline. Building it now is cheaper than building it after a Phase II failure or a pre-IND clinical hold.

Read the full memo at aiinsightlab.cloud/memos/pharma-ai-drug-discovery