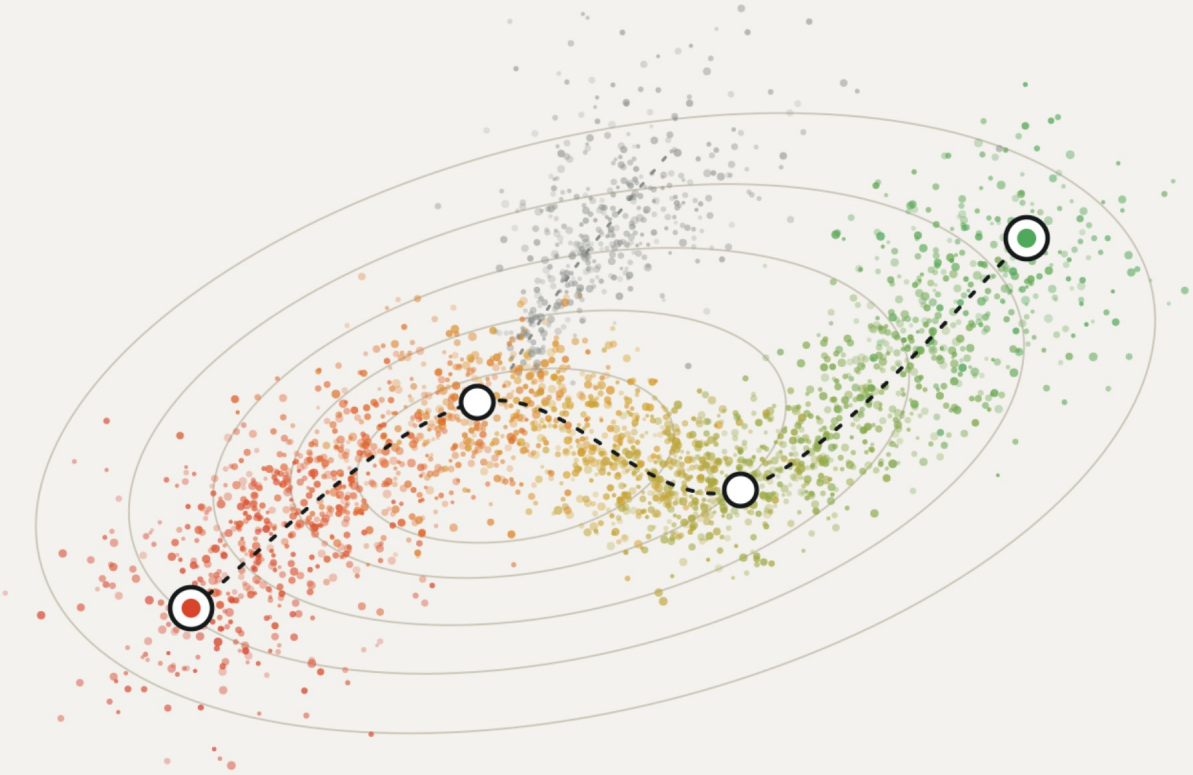
CELLREWIRE

The Art and Science of *Rewiring Life*

CORPORATE & SCIENTIFIC MEMORANDUM



Dr. Abdelrahman Mahmoud

FOUNDER & CHIEF EXECUTIVE OFFICER

HEIDELBERG · SAN DIEGO · CAIRO | NOVEMBER 2026

01 / FOUNDER NARRATIVE

A letter to the curious

Look at your hand. Several hundred cell types share one genome, and each reads a different part of it. Identity is not written into the sequence. It is held in place by a small set of transcription factors that bind their own enhancers and each other's, maintaining a pattern of chromatin accessibility that leaves most of the genome unreadable in any given cell.

That self-reinforcement is what holds identity stable against noise. It is also what makes identity addressable: a circuit has a finite number of nodes, and a node is something an intervention can reach. The open question is not whether one transition can be forced, but how much of the space is reachable, and at what dose.

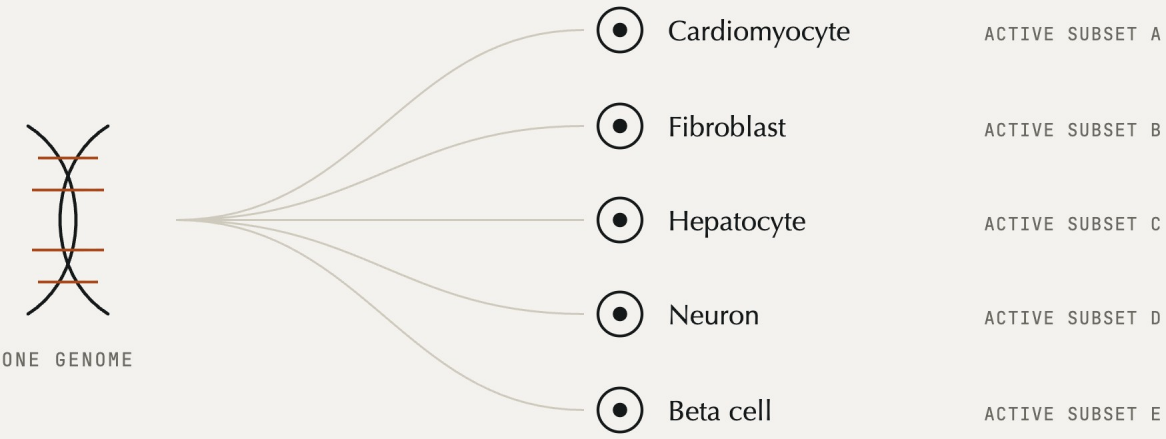


Fig. 1.1. One genome, many readings. Identity is a pattern of activity over a shared sequence, not a property of the sequence itself. Illustrative.

A state held in place by a feedback circuit can be addressed, provided the circuit can be reached at a dose a patient tolerates.

This memorandum is written for three readers at once: the student who wants to know how a cell decides what to become, the scientist who wants to know what we would accept as proof, and the investor who wants to know what twenty-four months buy. Throughout, what has been observed is kept separate from what is hypothesised.

02 / SCIENTIFIC THESIS

Controllability is the research question

Most work in this field asks whether one particular conversion can be made to work. A prior question is more useful, and can be made quantitative:

Which cell-state transitions are reachable under exposures a patient could tolerate? And does *ordering* the interventions enlarge the set that is reachable at all?

It has not been answered, and it is answerable. It is also worth answering whichever way it resolves: a reachable set that proves narrow tells both a scientist and an investor where to stop spending. No catalogue of observed cell states can supply that.

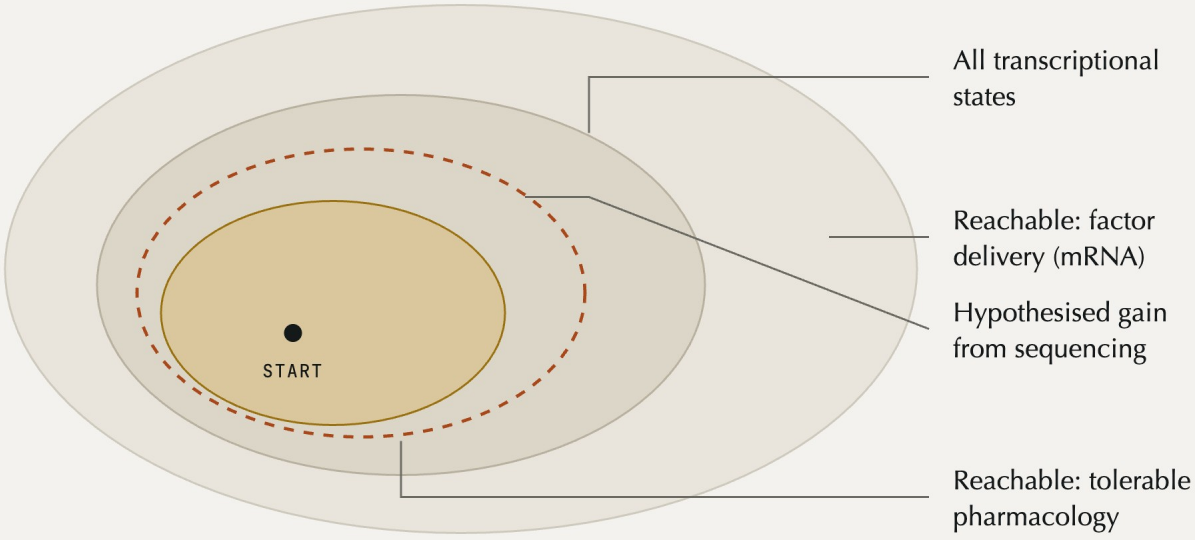


Fig. 2.1. The object of study. The three regions are not drawn to any measured scale; establishing their relative sizes is the work. The dashed boundary is a hypothesis, not a result.

OBSERVABILITY

How much of a cell's state a measurement recovers, and how much of that is needed to predict a response.

REACHABILITY

Which targets a feasible intervention attains from a given start, and whether order and timing extend that reach.

COST

How many experiments a route costs to find, and whether that number falls as evidence accumulates.

03 / BIOLOGICAL FOUNDATIONS

Identity is not immutable

Transcription factors change the activity of many genes at once, but their effect is conditional on the chromatin they meet. The same motif may sit in open chromatin in one cell type, in nucleosome-occluded DNA that only a pioneer factor can reach in a second, and in heterochromatin in a third. A factor set is therefore not the same instruction in two different cells, which is why efficiencies reported in one donor, or one species, transfer poorly.

Two decades of work have established that a state can be moved at all: Takahashi and Yamanaka induced pluripotency in mouse fibroblasts, and Ieda and colleagues later converted fibroblasts directly toward cardiac muscle.

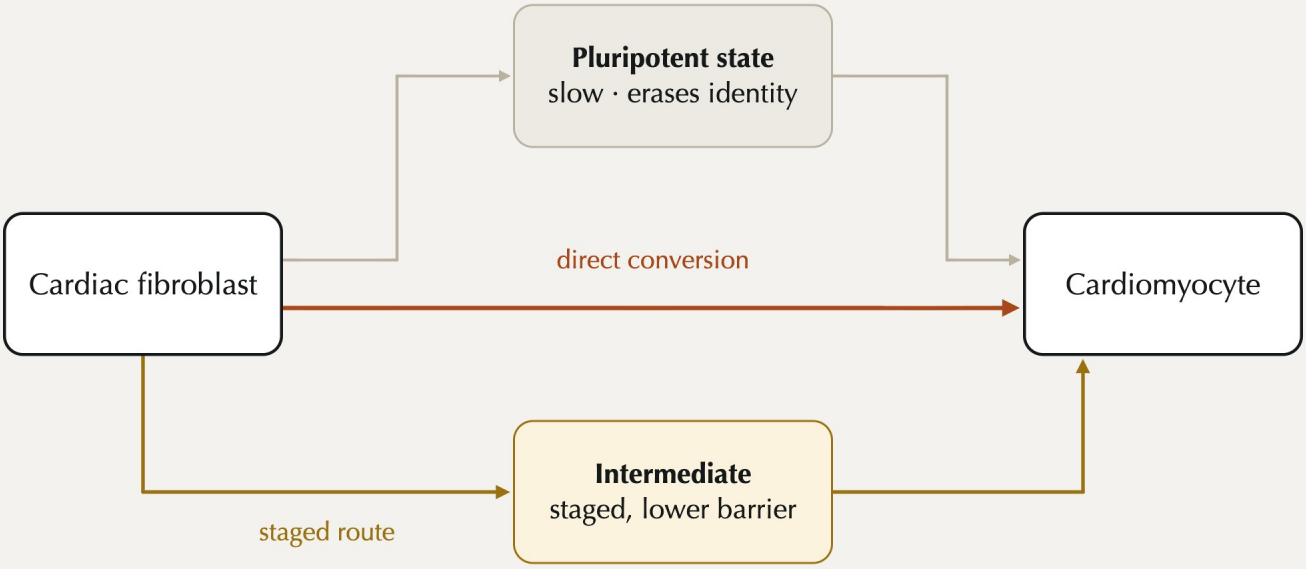


Fig. 3.1. Three routes to a new identity. We pursue direct transitions that preserve tissue function without passing through a pluripotent state; the staged route is the subject of the following chapter. Schematic.

Routing through pluripotency discards the residual chromatin context we intend to exploit and carries teratoma risk from any incompletely differentiated cells that survive. The direct route is less forgiving: there is no intermediate at which the cell is a defined, characterisable type, so a partially converted population cannot be purified away and must instead be shown to be harmless. We therefore treat three things as separate claims with separate evidence: transcriptional profile, electrical and contractile function, and persistence after withdrawal. A conversion that satisfies the first and fails the third is a transient phenotype, not a therapy.

04 / SEQUENCE AND TIMING

Why order matters

The barrier is not a metaphor: it is the cost of opening repressed target enhancers while the incumbent regulatory network is still reinforcing itself. There are two ways past it. Pushing harder means more perturbation: higher exposure, more factors, longer expression. That ceiling is set by toxicity, not biology. Going around means reaching an intermediate state in which part of the incumbent network has already relaxed, leaving a lower barrier than the one that was refused.

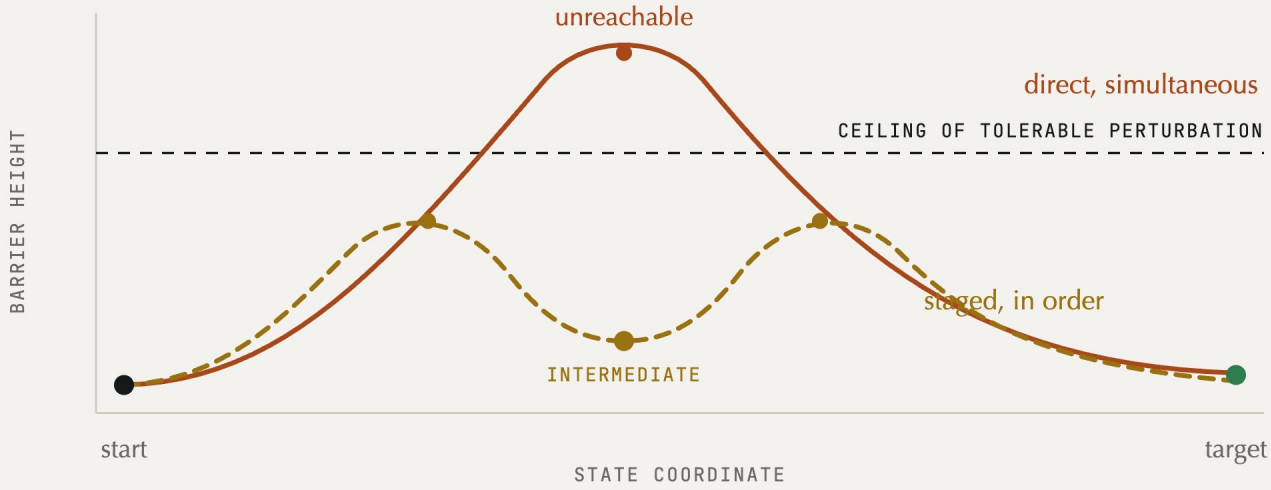


Fig. 4.1. The central hypothesis, drawn. Both routes end at the same target; only one stays under the ceiling throughout. Schematic: no axis carries measured units, and the position of the ceiling is itself an experimental quantity.

Sequential intervention makes transitions accessible that simultaneous intervention cannot reach, because an intermediate state lowers the barrier that remains.

This is the one claim here that is both novel and cheap to falsify. Four arms settle it: the sequence; the same agents together; each alone; and the additive prediction from the single-agent arms. If the sequence does not beat the other three, the claim is wrong, at the cost of one well-designed experiment.

WHY REINFORCEMENT LEARNING

The object to be learned is a policy over time, not a ranking of molecules. The formalism follows from the problem, not from fashion.

WHY THE SAME CONTROLS DEFEND THE IP

A regimen patent survives an obviousness objection on evidence of synergy that does not follow from the single agents. The arms that test the science are the arms that build that file.

05 / CAUSAL BIOLOGY

A map is not a mechanism

Single-cell measurement tells us which molecular profiles resemble one another. It does not tell us how to move between them. Two states can sit beside each other in an embedding and be separated by a barrier no tolerable intervention crosses; two far apart can be joined by a single factor.

The reason is concrete. An expression embedding is dominated by abundant transcripts (housekeeping, ribosomal, cell-cycle), while the barrier between two identities sits in a handful of lowly expressed regulators that contribute almost nothing to the distance. Similarity is correlational; reachability is causal, and the only instrument that measures it is a controlled perturbation.

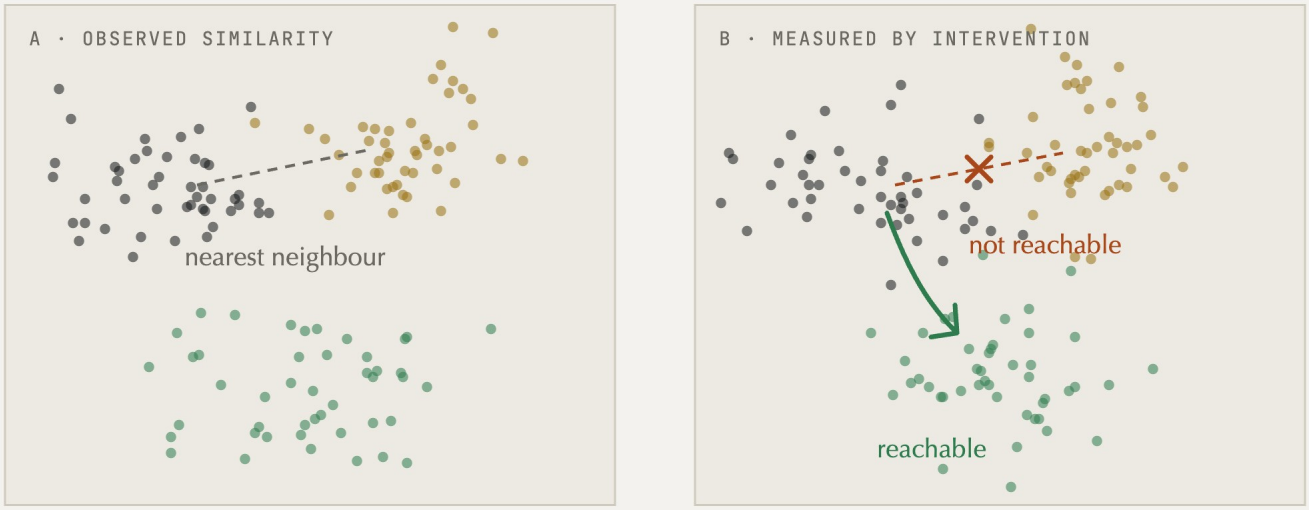


Fig. 5.1. Why an atlas does not answer the question. Both panels show the same embedding; only the second was measured by perturbing. Illustrative, with simulated points.

WHAT THIS COMMITS US TO

01 Predictions recorded before the experiment is run.

02 Whole donors and whole intervention combinations held out of training.

03 Transfer to an unseen cell type treated as a separate, later test, never a headline.

04 Cell-DRL benchmarked against biological heuristics, simpler statistical models and random selection, each at a matched experimental budget.

06 / COMPUTATIONAL PLATFORM

RewireOS, and the decision it makes

An intervention is a decision under uncertainty. RewireOS exists to make that decision explicit: represent the starting state, enumerate the perturbations actually available, estimate the response, choose, then let the experiment correct the estimate.

Cell-DRL, the agent at its centre, is a policy over cell-state transitions rather than a classifier over molecules. Its objective is a predefined functional response, penalised by uncertainty, toxicity and the cost of the experiment.

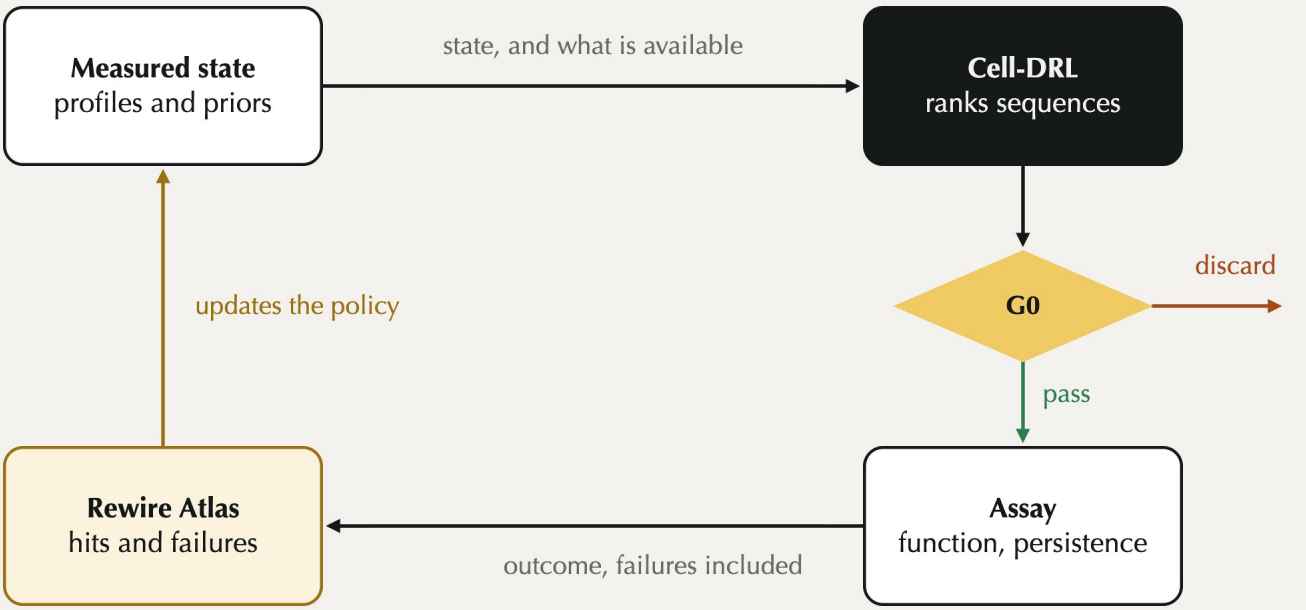


Fig. 6.1. The loop, and the gate in it. A research design, not a specification of implemented software; current maturity is a matter for technical diligence.

G0 · FEASIBILITY BEFORE POTENCY

Before an animal is used, a proposed regimen is tested against one question: does every step require a concentration at the target tissue that is achievable within the exposure envelope already established for that agent in humans? A sequence demanding ten times human C_{max}, or a sustained trough that no dosing schedule produces, is discarded in software. Ordering the filters this way, feasibility first and predicted potency second, is what compresses the timeline. The learning loop does not compress it; the iteration arithmetic in chapter 11 shows why.

FROZEN BEFORE TRAINING

Test sets are fixed before a model sees data, and held by someone who is not running the model.

CALIBRATION, NOT ONLY ACCURACY

A model that is confidently wrong is worse than useless here, because it spends animals. We report both.

07 / MODALITIES

Reach where drugs can go. Precision where they cannot.

One engine drives two execution layers. Which one a programme uses is a consequence of what the model selects, not an allegiance to a technology. Two modalities at seed stage can read as a lack of focus. They are not, because the expensive apparatus (engine, atlas, assays, readouts) is shared; only the delivery differs.

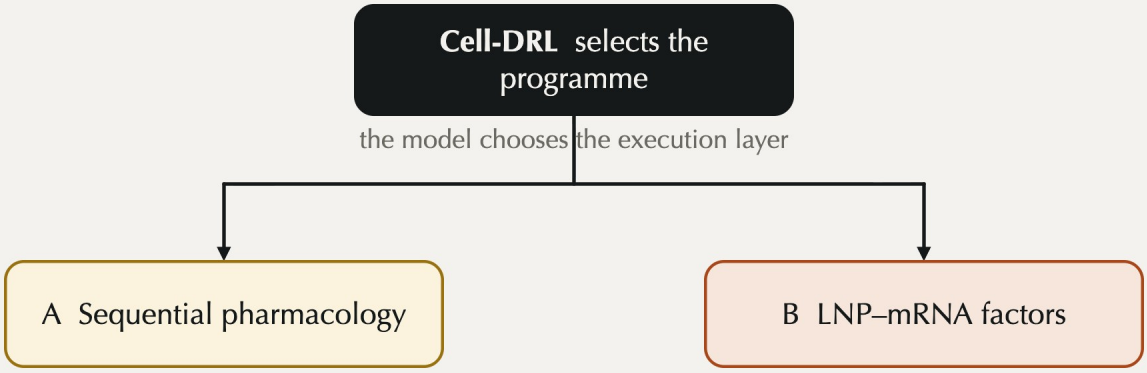


Fig. 7.1. Where the two modalities diverge, and where they do not. Everything above the split is shared cost.

MODALITY A

Timed sequences of compounds that already carry human exposure data: nuclear-receptor agonists, epigenetic modulators, differentiation agents.

Advantage Human PK and safety already characterised; the shortest credible path to first-in-human.

Constraint Limited reach into lineage-specifying programmes. Regimen IP must be designed in deliberately, not hoped for.

MODALITY B

Lipid nanoparticle delivery of mRNA encoding the factor sets the model selects, reaching programmes no approved drug can specify.

Advantage Direct control of identity, and a considerably stronger composition-of-matter position.

Constraint Tissue-specific delivery, duration of expression, immunogenicity, manufacture.

Stated honestly: pharmacology is how we would reach a state *quickly*; factor delivery is how we would reach a state pharmacology cannot specify *at all*. One consequence shapes Modality A's library: the criterion is prior human exposure rather than approval. That admits clinical-stage pathway inhibitors and approved epigenetic agents, widening the accessible chemistry while leaving most of the toxicology already paid for.

08 / LEARNING AND SCALE

Does biology have a learning curve?

What should become possible after our thousandth intervention that was impossible after our hundredth? We do not know. That is not a rhetorical opening; it is the measurement we intend to make. Successive models are trained on increasing amounts of intervention data and tested on the same frozen tasks, and the shape of the resulting curve decides what the company is.

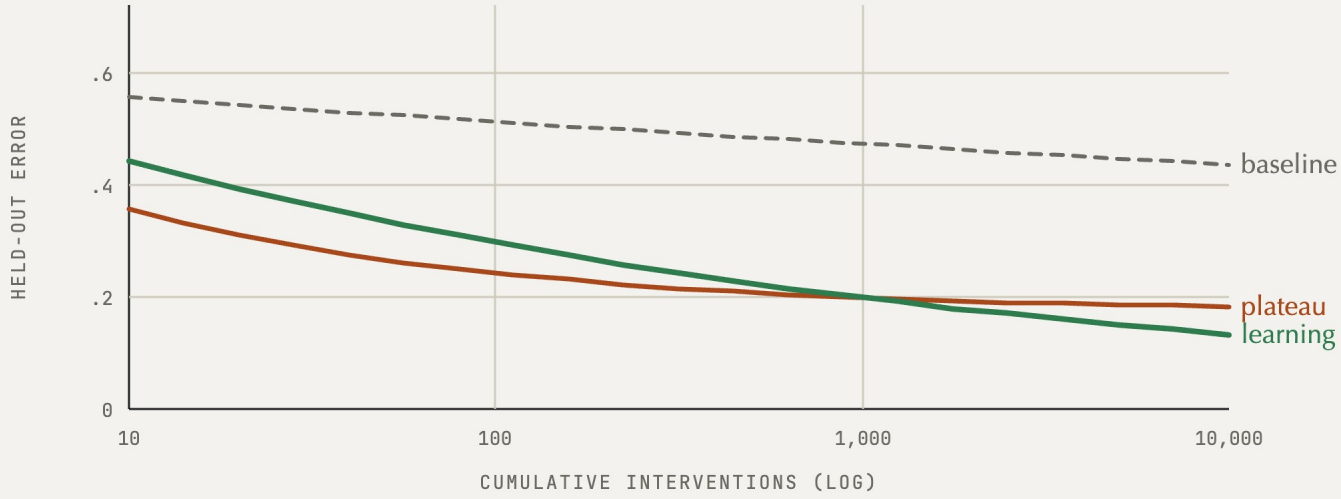


Fig. 8.1. Three outcomes, drawn before the data exists. No scaling law has been established in this domain; the axes carry illustrative values only. A plateau, or a failure to beat the simpler baseline, is a result we would publish and act on.

PREDICTION

Held-out error, calibration, and the fraction of prospectively tested candidates meeting a predefined endpoint.

EFFICIENCY

Experiments, elapsed time and cost per reproducible functional hit, against matched-budget baselines.

TRANSFER

Performance on unseen donors and unseen interventions first; on a new cell type only if that is earned.

Stated plainly: if the curve declines, the cost of each further discovery falls as evidence accumulates, and the atlas rather than any single molecule becomes the asset. If it plateaus, CellRewire is a company with one or two good therapies and no platform: worth building, worth considerably less, and worth establishing early.

09 / EXPERIMENTAL EVIDENCE

One result, read carefully

After an infarct, cardiac fibroblasts proliferate and deposit collagen. The resulting scar maintains wall integrity and cannot contract. Those fibroblasts are an unusually tractable substrate: abundant in the injured region, actively cycling, and carrying the patient's own genome, so conversion in place needs no transplant, no immunosuppression and no manufactured cell product. The result below bears on whether a useful fraction of them converts, though only partially.

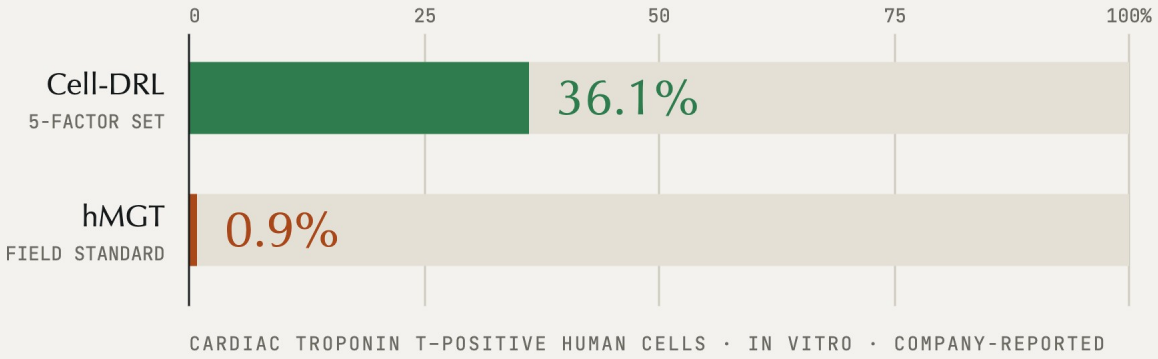


Fig. 9.1. The result the company reports: a computationally selected five-factor combination against hMGT (human GATA4, MEF2C and TBX5, the field-standard set) in the same reported experiment. Approximately forty-fold. The factor identities, raw data and methods remain in controlled diligence pending management and IP clearance.

WHAT THIS ESTABLISHES

That the engine selected a combination which, on a marker readout in the reported in vitro experiment, substantially outperformed the set the field currently uses. Five factors drawn from the roughly 1,600 human transcription factors is a space of order 10^{14}; finding a strong performer in it is a signal about the search, not yet about the biology.

WHAT IT DOES NOT ESTABLISH

Replicate counts, statistical uncertainty, or independent donor validation. Nor electrical and contractile function, nor persistence after the intervention stops. Marker expression is not regeneration, and none of this is yet *in vivo*.

REQUIRED BEFORE THIS ADVANCES

Replication
The marker result repeated across independent donors and independent hands.

Function
Calcium handling, electrophysiology and contraction, not cTnT alone.

Durability
Does the state hold after expression stops, and for how long.

Safety
Off-target identity, proliferation, and the behaviour of cells that convert partway.

10 / EXPERIMENTAL INFRASTRUCTURE

An atlas whose unit is the intervention

In a descriptive atlas the unit of record is a cell. Here it is an intervention applied to a defined starting population. A record carries the pre-intervention profile and donor covariates, the perturbation and its delivery vehicle, the order and inter-dose interval, the readout and the time at which it was taken, the outcome against a prespecified endpoint, and an evidence grade fixing replicate count, control arms and whether the prediction was registered before the run.

ONE RECORD · THE UNIT OF THE ATLAS

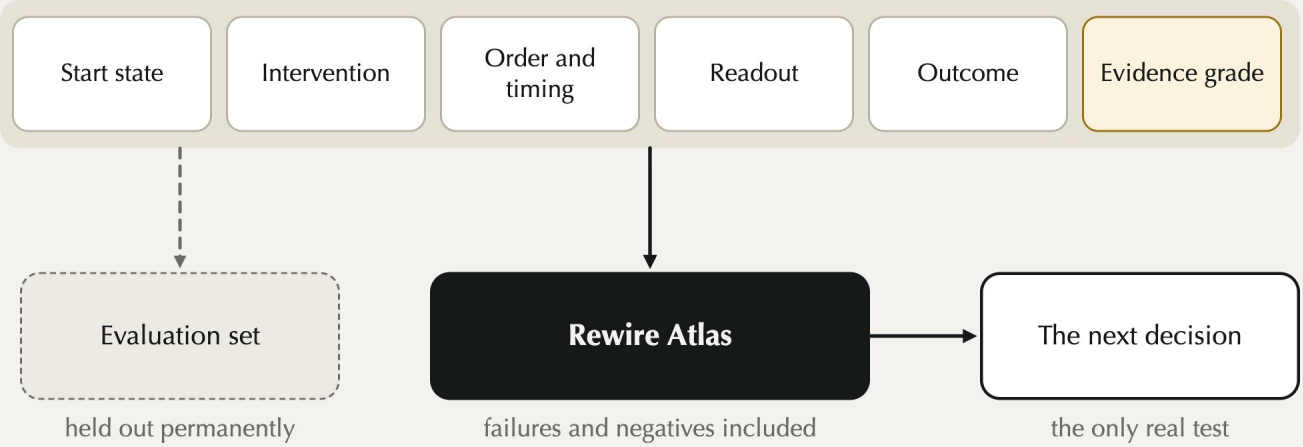


Fig. 10.1. Record schema and routing. Records lacking vehicle-only and single-agent control arms are retained but down-graded, and are excluded from the training set used for prospective claims.

Record count is the wrong figure of merit. The quantity we report is the difference in prospective decision quality between a model trained on *n* records and one trained on a subset, both evaluated on the same held-out set, the same comparison that generates the learning curve in chapter 08. The evaluation set is partitioned by donor and by intervention class at the point of collection, before any model has seen it, and is held by a person who does not run the models. Partitioning after the fact permits selection on the test set whether or not anyone intends it.

PROGRAMMES, AND THE ORDER THEY EARN

RewireMI is the first test and the only programme the round funds. A hepatic programme sits behind it on delivery grounds rather than biology: ionisable lipid nanoparticles reach hepatocytes by ApoE-mediated LDL-receptor uptake without engineered targeting, which separates the question of whether the policy transfers from whether the vehicle arrives. Cardiac fibroblast delivery lacks that property and needs active targeting. RewirePD and RewireBeta remain exploratory; expansion is conditioned on measured transfer, not on the existence of another indication.

11 / GATES AND MILESTONES

What twenty-four months buy

The programme is organised as three intervals separated by decision gates. Each gate has a condition written before the interval begins, and failing a condition stops the programme at that gate rather than deferring the question to the next one. The structure matters more than the schedule: it fixes in advance what evidence would cause us to stop.

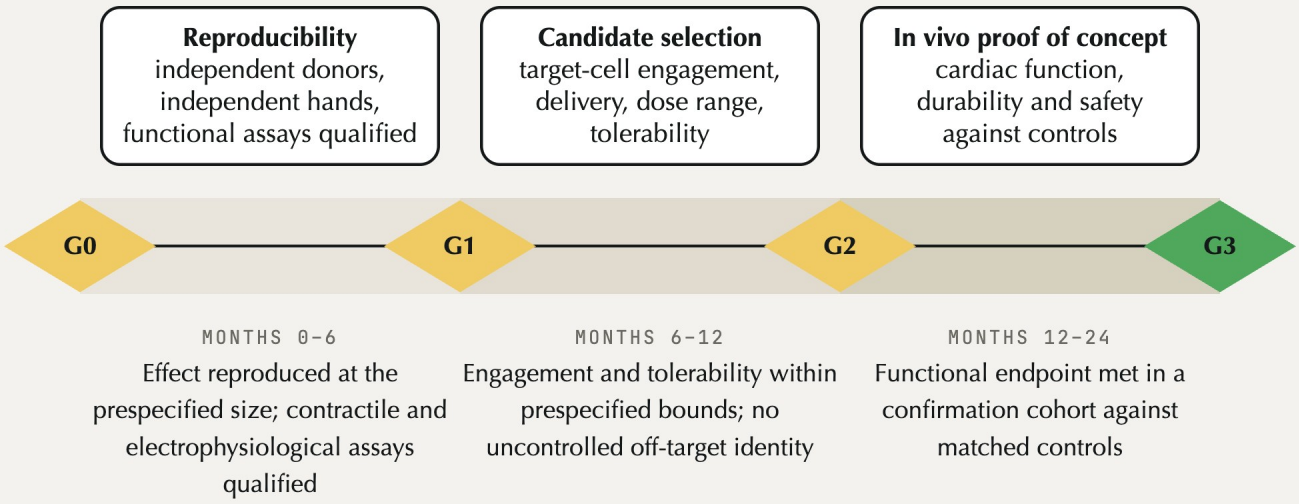


Fig. 11.1. Three intervals, four gates. G0 is the computational exposure-feasibility filter of chapter 06 and runs continuously; G1 to G3 are stage gates. Durations are plans, not commitments.

FIXED BEFORE EACH STUDY

Sample size and the power calculation behind it, the effect threshold that counts as success, the statistical plan including how multiplicity is handled, and the stopping rule, all in a written protocol dated before the first animal.

WHY THE NUMBERS ARE NOT HERE

A threshold on ejection fraction or scar area is only meaningful once the assay measuring it has a known coefficient of variation in our hands. That qualification is work in interval one. Thresholds set before it would not be falsifiable.

One feature of this plan is unusual and deliberate: the in vivo model is entered earlier than is typical, because the quantities that decide the programme (exposure at the target cell, systemic tolerability, persistence) are not recoverable from culture. That choice fixes the iteration count. At eight to fourteen weeks per in vivo cycle, twenty-four months is three or four cycles, which makes this a designed series of experiments under model-guided selection rather than a continuously learning closed loop. We describe it as the former throughout.

12 / INTELLECTUAL PROPERTY AND RISK

Where this fails, and when we would know

Failure modes are not equally expensive. What matters is not the list but its geometry: how severe each failure is, against how early the programme would detect it. The gate structure of chapter 11 exists to pull points leftward on this map.

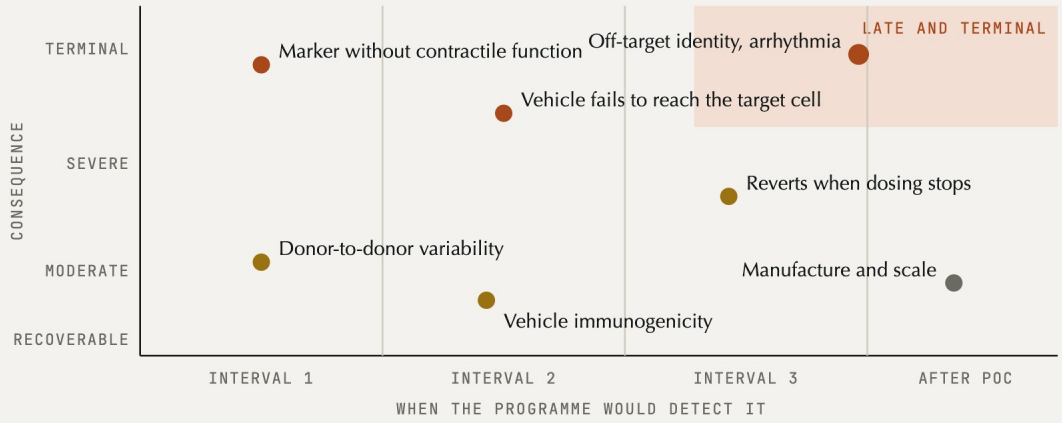


Fig. 12.1. Failure modes by severity and detection stage; positions are the team's judgement, not measured probabilities. The shaded quadrant destroys capital: arrhythmogenicity from partially converted cells cannot be excluded in culture, which is the strongest argument for entering the in vivo model when we do.

INTELLECTUAL PROPERTY

Patent publication **WO2025202019**. Publication is not grant: claim scope, ownership, freedom to operate and third-party rights each require legal examination. Two exposures are specific enough to name.

OBVIOUSNESS

A regimen claim resting on a *rationally predicted* combination invites an obviousness objection: the rationality of the prediction is a scientific asset and a commercial liability at once. It is answered by synergy that does not follow from the single agents, which is why the single-agent and additive-prediction arms of chapter 04 are run as controls.

ENFORCEMENT

Where every component is a generic, a valid claim may have no enforcement route against prescriber-level substitution. That argues for a co-formulated product or a biomarker-selected label rather than a loose sequence, a choice to be made before the first in vivo arm, since it changes the experiments.

Commercial adoption, reimbursement, regulation and capital availability are each uncertain. Accountable leads for cardiac electrophysiology, delivery and regulatory strategy are still to be appointed, and the seed programme is scoped on the assumption that they will be. Approval of individual medicines does not establish approval of a new combination. Investors could lose their entire investment. These risks are not exhaustive.

13 / REFERENCES AND INVITATION

An invitation to build

The claims in this memorandum rest on five published results and one unpublished company experiment. The published work establishes that the route exists; none of it establishes that we can navigate it. Chapter 09 sets out what our own result does and does not show.

THE LITERATURE THESE CLAIMS REST ON

Takahashi K, Yamanaka S. Induction of pluripotent stem cells from mouse embryonic and adult fibroblast cultures by defined factors. *Cell* 2006;126(4):663–76. doi:10.1016/j.cell.2006.07.024 · Establishes that identity is not fixed. The route we deliberately avoid.

Ieda M, Fu J-D, Delgado-Olguin P, et al. Direct reprogramming of fibroblasts into functional cardiomyocytes by defined factors. *Cell* 2010;142(3):375–86. doi:10.1016/j.cell.2010.07.002 · Gata4, Mef2c and Tbx5; the chapter 09 comparator derives from this set.

Mohamed TMA, Stone NR, Berry EC, et al. Chemical enhancement of in vitro and in vivo direct cardiac reprogramming. *Circulation* 2017;135(10):978–95. doi:10.1161/CIRCULATIONAHA.116.024692 · TGF-β and WNT inhibition raised efficiency eight-fold, with beating cells at one week rather than six to eight. The closest precedent for chapter 04; every compound in it is a tool compound, which is the gap our library criterion addresses.

Huang C, Tu W, Fu Y, Wang J, Xie X. Chemical-induced cardiac reprogramming in vivo. *Cell Res* 2018;28(6):686–9. doi:10.1038/s41422-018-0036-4 · Small molecules alone, in vivo, after infarction in mouse.

Dixit A, Parnas O, Li B, et al. Perturb-Seq: dissecting molecular circuits with scalable single-cell RNA profiling of pooled genetic screens. *Cell* 2016;167(7):1853–66. doi:10.1016/j.cell.2016.11.038 · The measurement that makes an interventional atlas tractable at scale.

Bibliographic details verified against PubMed; none of these validates a CellRewire result. Company sources: pitch deck, management financial workbook, patent publication WO2025202019.

The experiment that would most change our own view is the one that tests whether order matters at all: the sequenced arm against the simultaneous arm, each agent alone, and the additive prediction, in a functional assay of known coefficient of variation. It is inexpensive, it is first in the plan, and we have not run it. Correspondence from scientists who think the hypothesis is wrong is more useful to us than correspondence from those who think it is right.

AMAHMOUD@CELLREWIRE.COM · CELLREWIRE.COM · HEIDELBERG, SAN DIEGO & CAIRO