Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Co-transduction-aware digital twinning of dual-AAV prime editing for genetically stratified CNS therapy

FAIZAN AHMAD

faizmedi16@gmail.com

Jamia Hamdard https://orcid.org/0000-0002-7768-9411

Research Article

Keywords: digital twin, dual-AAV delivery, prime editing, co-transduction, LRRK2, gene therapy

Posted Date: November 17th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11747124/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Co-transduction-aware digital twinning of dual-AAV prime editing for genetically stratified CNS therapy

Faizan Ahmad*

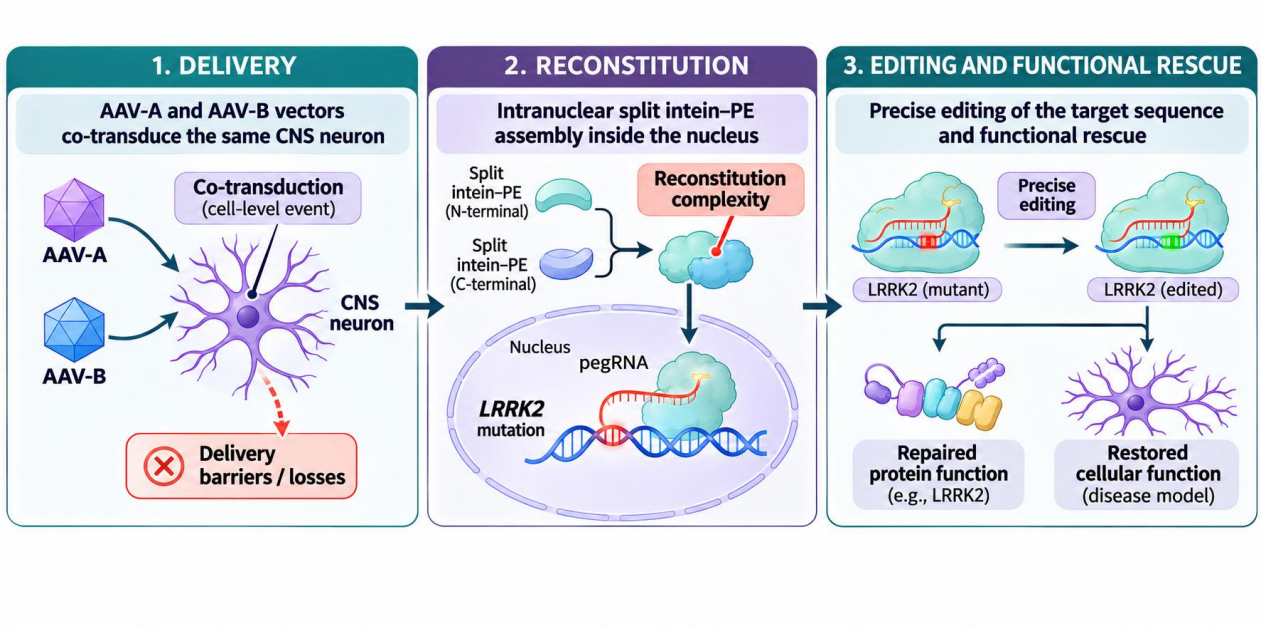
*Department of Medical Elementology and Toxicology, Jamia Hamdard University, Delhi, India

***Correspondence**
Faizan Ahmad
Department of Medical Elementology and Toxicology
Jamia Hamdard University, Delhi
India
Email: 1996faizanahmad@gmail.com; faizmedi16@gmail.com

Abstract

Dual-AAV prime editing could deliver precise gene correction to the central nervous system, but its therapeutic yield depends on a chain of events co-delivery of both vectors to the same cell, editor reconstitution, editing, and rescue that is usually measured separately in different systems. This study presents DUAL-AAV PRIME-DT, an evidence-constrained digital twin that treats AAV-A/AAV-B co-transduction as an explicit cell-level state and propagates it through the whole therapeutic chain. In a synthetic cohort of 10,000 virtual patients with LRRK2-associated Parkinson disease, mean therapeutic coverage is 1.24% (patient-level 5th\95th percentile, 0.12–3.93%), with joint delivery and editing causing the largest losses. Perfect co-transduction raises mean coverage by 74%, whereas assuming independent vectors changes it by only −5%. Stress tests, sensitivity analysis, and value-of-information analysis identify delivery, reconstitution, and editing as principal levers and show that a terminal phenotype cannot reveal which step failed. All outputs are synthetic and imply no clinical efficacy; the framework specifies the matched same-cell measurements required to replace synthetic priors and to test its predictions.

Keywords: digital twin; dual-AAV delivery; prime editing; co-transduction; LRRK2; gene therapy



Highlights

- Co-transduction is modelled as an explicit cell-level state in dual-AAV editing
- Synthetic LRRK2 cohort: 1.24% mean coverage (5^{th}-95^{th} percentile 0.12-3.93%)
- Perfect co-transduction adds 74% coverage; assuming independence changes it by 5%
- Delivery, reconstitution and editing are the principal levers in stress tests
- Editing and co-transduction measurements carry the highest information value

1. Introduction

Prime editing installs precise sequence changes without double-strand breaks or donor templates (Anzalone et al., 2019), and successive improvements to the editor, guide RNA and cellular repair environment have raised its efficiency (Chen et al., 2021; Dacquay et al., 2025; Nelson et al., 2022; Sakai et al., 2026; Yan et al., 2024). Efficiency can now be predicted from sequence and DNA-repair determinants in many cell types (Koeppel et al., 2023; Mathis et al., 2023; Yu et al., 2023), and prime editing has entered clinical use (Gori et al., 2026). In the central nervous system (CNS), however, the editor is too large for a single adeno-associated virus (AAV), so delivery relies on splitting it across two vectors that reconstitute inside the cell (Davis et al., 2024; Doman et al., 2023; Kissling et al., 2026; Levy et al., 2020; She et al., 2023). Dual-AAV prime editing has produced editing in up to roughly 40% of cells in mouse cortex (Davis et al., 2024; Doman et al., 2023) and corrected a pathogenic Scn1a allele in an epilepsy model (Kissling et al., 2026), and split-editor correction has been shown in retina (She et al., 2023).

These benchmarks conceal a structural fact. A cell can be corrected only if it receives both vectors, the two halves reconstitute, and the editor then acts on its target; therapeutic yield is therefore a product of conditional probabilities rather than a property of the editor alone. Each step has been measured, but in different species, regions, cell types and assays, so the results cannot be combined into one prediction. Single-cell quantification shows that multi-AAV success depends on whether vectors co-occur in the same neuron (Maturana et al., 2022), and regional AAV transduction predicts regional editing in Huntington disease models (Duarte et al., 2023). Yet capsid tropism and biodistribution differ across route, species and region (Bunuales et al., 2024; Goertsen et al., 2022; Okai et al., 2025; Rico et al., 2026), and pre-existing immunity and dose-dependent toxicity constrain delivery (Hinderer et al., 2018; Schulz et al., 2023). A headline editing percentage therefore cannot reveal which step limits a therapy.

Here we ask how much explicit co-transduction modelling changes predicted yield, which step dominates loss and uncertainty, and which matched measurements would most efficiently replace assumptions with data. We built DUAL-AAV PRIME-DT, a multiscale stochastic digital twin that retains AAV-A and AAV-B occupancy at the level of the cell and propagates it through reconstitution, editing, molecular rescue and functional rescue. We apply it to LRRK2-associated Parkinson disease, a genetically defined disorder with measurable pathway biomarkers (Di Maio et al., 2018; Kalogeropoulou et al., 2022; Mabrouk et al., 2026; Steger et al., 2016; Yuan et al., 2024), as a translational hypothesis rather than as evidence of efficacy: all LRRK2 outputs are synthetic, and published CNS editing studies serve only as constraints.

2 Materials and Methods

2.1 Study design and evidence hierarchy

This was a computational preclinical modelling study; no new human, animal, vector-production, genome-editing or clinical intervention experiment was performed. Published observations were used as evidence constraints and benchmark targets, and simulation-derived values are labelled synthetic. Evidence was retrieved from PubMed-indexed literature and stratified into direct dual-AAV prime-editing evidence (Davis et al., 2024; Doman et al., 2023; Kissling et al., 2026; She et al., 2023), adjacent split-editor evidence (Colella et al., 2014; Levy et al., 2020; Mittas et al., 2026; Riedmayr et al., 2023; Tang et al., 2023; Tasfaout et al., 2024; Trapani et al., 2014), CNS AAV delivery and spatial evidence (Bolon et al., 2024; Bunuales et al., 2024; Chen et al., 2023; Drouyer et al., 2024; Duarte et al., 2023; Foust et al., 2009; Goertsen et al., 2022; Hocquemiller et al., 2022; Maturana et al., 2022; Moyer et al., 2025; Okai et al., 2025; Rico et al., 2026), AAV

safety and immunogenicity evidence (Hinderer et al., 2018; Leborgne et al., 2020; Mendell et al., 2022; Rapti & Grimm, 2021; Schulz et al., 2023; Verdera et al., 2020; Weber, 2021), LRRK2 molecular and biomarker evidence (Cortés et al., 2025; Di Maio et al., 2018; Hop et al., 2024; Jennings et al., 2022; Jennings et al., 2023; Kalogeropoulou et al., 2022; Kania et al., 2023; Kluss et al., 2022; Mabrouk et al., 2026; Steger et al., 2016; Streubel-Gallasch et al., 2021; Vissers et al., 2023; Wauters et al., 2020; Yuan et al., 2024), and Parkinson disease gene-therapy and biomarker context (Brooker & Gonzalez-Latapi, 2025; Grote et al., 2024; Kirik et al., 2017; Siderowf et al., 2023; Simuni et al., 2024; Winston et al., 2025). Registry context was derived separately from ClinicalTrials.gov records (checked 28 October 2026; Table 1) and was not treated as efficacy evidence. The hierarchy prevents transitive overclaiming: cortical editing in mouse brain does not become a prior for human LRRK2 clinical response merely because both involve AAV and prime editing.

Registry	**Purpose in manuscript**	**Status / size**	**Use in Digital Twin**
NCT03866603 (ROPAD)	LRRK2 genetics and biomarker cohort	Completed; actual enrollment 20,112	Patient-state prior and genetic stratification
NCT07681219 (NEU-PD)	LRRK2 clinical/molecular/neurophysiological fingerprint	Not yet recruiting (last record retrieved); estimated enrollment 20	Multimodal patient-state schema
NCT05603312	Randomised, sham-controlled AAV-GAD (MGT-GAD-025), phase 1/2	Completed; actual 14	Clinical AAV feasibility/safety context
NCT07371338	AAV gene therapy IPS101A, phase 1, dose escalation	Recruiting	Contemporary translational context

Table 1: Clinical registry context (ClinicalTrials.gov; checked 28 October 2026). Records are contextual and do not establish efficacy

2.2 State space and delivery model

Each cell is represented by $Z_i = [G_i, C_i, X_i, T_i, A_i, B_i, N_{Ai}, N_{Bi}, AB_i, R_i, E_i, O_i, M_i, N_i, F_i]$, where G is genotype, C cell class, X anatomical position, T temporal state, A and B vector exposure, N_A and N_B vector occupancy, AB joint co-transduction, R editor reconstitution, E desired editing, O by-product or off-target editing, M molecular rescue, N network rescue and F functional rescue (Table 2). Patient-level states aggregate cell-level states through anatomical and cellular weights. Occupancy is Poisson, with marginal transduction $P(A) = 1 - \exp(-\lambda_A)$ and $P(B) = 1 - \exp(-\lambda_B)$; co-transduction is $P(AB) = P(A)P(B|A)$, with enrichment $\eta_{AB} = P(B|A)/P(B)$ and coefficient $CE = P(AB)/[P(A)P(B)]$. Spatial heterogeneity is introduced by letting λ_A, λ_B and η_{AB} vary across position, cell class and time. In the cohort implementation, $P(AB) = \min[P(A)P(B)\eta_{AB}, P(A), P(B)]$, so that enrichment is bounded by the Fréchet upper

limit; this bound was active for 8.1% of virtual patients. Across the cohort, η_{AB} had a mean of 1.10 (median 1.00), CE had a mean of 1.07, and λ_B equalled 0.92 λ_A for every patient.

Layer	State variable	Model role	Primary validation measurement
Patient	G, stage, vulnerability	Defines disease context and heterogeneity	Genotype + clinical phenotype
Anatomy	X, cell class	Controls spatial exposure and rescue weighting	Region/cell-type mapping
Vector A	A, N_A	First delivery gate	Vector-A exposure / genome abundance
Vector B	B, N_B	Second delivery gate	Vector-B exposure / genome abundance
Joint state	AB	Therapeutic co-transduction gate	Same-cell A/B measurement
Reconstitution	R	Active editor state	Reconstituted editor / transcript / protein
Editing	E, O	Desired vs unintended editing	Desired and by-product/off-target sequencing
Molecular rescue	M	Corrected molecular phenotype	Protein / pathway / biomarker
Network rescue	N	Circuit-level propagation	Functional network readout
Functional rescue	F	Final therapeutic state	Behavioural / clinical endpoint

Table 2: Core Digital Twin state variables and causal transitions

2.3 Reconstitution, editing and rescue

Reconstitution is a logistic transition, $P(R = 1|AB) = \text{sigmoid}(\beta_0 + \beta_1N_A + \beta_2N_B + \beta_3N_AN_B + \beta_4C + \beta_5t + \beta_6\delta_t)$, in which δ_t is the relative timing of the two vector components and the interaction term allows super- or sub-additive dependence on occupancy. Desired editing is $P(E = 1|R) = \text{sigmoid}(\gamma_0 + \gamma_1R + \gamma_2G + \gamma_3C + \gamma_4T + \gamma_5Q)$, where Q is a compact vector of target and edit-design descriptors; by-product and off-target editing are tracked separately so that a high total editing fraction is not mistaken for a high therapeutic fraction. Therapeutic coverage is $C = P(AB) \times P(R|AB) \times P(E|R) \times P(M|E) \times P(F|M)$, and network rescue can be inserted between M and F when a calibrated network model is available.

2.4 Hierarchical calibration of editing

Published editing observations were treated as heterogeneous study-level signals rather than as a common-effect meta-analysis, and ‘up to’ values were treated as upper-bound constraints rather than population averages. Domain-specific posteriors (prime editing in brain and retina, SCN1A, Huntington disease striatum and cortex, and base editing in brain) were estimated under an

evidence-weighted hierarchical prior. The pooled check compares one common editing prior with benchmark endpoints; the hierarchical check compares each endpoint with its own domain predictive distribution. The posterior-predictive statistic is the predictive mean for mean-type endpoints and an upper summary of the domain predictive distribution for 'up to' endpoints, and the reported quantity is the probability that a replicate meets or exceeds the reported value. Spatial calibration reproduced the regional prediction problem of Duarte et al. (2023).

2.5 Synthetic cohort

The virtual cohort comprises 10,000 patients (4,499 LRRK2 G2019S, 1,509 other LRRK2 variants and 3,992 sporadic-like) and 90,000 cells, generated with seed 20261008. Variables were sampled from literature-constrained or governance-defined distributions, and each parameter is tagged as observed, literature-derived, inferred or synthetic. Synthetic distributions are not estimates of any clinical population.

2.6 Counterfactuals, stress tests and sensitivity analysis

Counterfactual coverage replaces one transition while holding the rest of the patient state fixed. The independence counterfactual replaces P(AB) with P(A)P(B), and perfect co-transduction replaces it with min[P(A), P(B)]. Stress tests rescale a parameter class by ±25% and recompute P(A), P(B) and P(AB) before recomputing coverage. Rank sensitivity uses Spearman correlation, and conditional leverage uses standardized multivariable regression. Variance-based analysis uses Sobol total-effect indices for nine shift parameters (A_scale, B_scale, rho_AB, R_shift, E_shift, stage_shift, spatial_heterogeneity, temporal_penalty and phenotype_shift) within assumed ranges that were not equalised. Value of information is approximated by expected entropy reduction, $x^* = \text{argmax}\ E[H(\theta) - H(\theta|D_x)]$, under an assumed measurement-noise model; the highest-ranked measurement is the most informative experiment, not necessarily the most upstream variable.

2.7 Model comparison and generalisation

Six nested models were compared: M0, baseline mean (AUROC 0.5 and R^2 0 by construction; not tabulated); M1, AAV exposure; M2, plus spatial structure; M3, plus co-transduction; M4, plus reconstitution and editing; and M5, the full model with temporal and phenotype layers. Metrics were AUROC, area under the precision–recall curve, Brier score, root-mean-square error, mean absolute error and R^2. Leave-one-disease-out evaluation used LRRK2, SCN1A and Huntington disease anchor distributions. Identifiability was classified by whether parameters are directly observed, indirectly constrained or inferred only from the endpoint.

2.8 Statistics

All percentages derived from the virtual cohorts are Monte Carlo quantities. No P values are presented as evidence of biological efficacy. Bootstrap confidence intervals (5,000 resamples of the cohort) quantify simulation uncertainty only.

2.9 Cell-level cross-disease cohort

A second, cell-level virtual cohort of 90,000 cells (30,000 each for LRRK2 Parkinson disease, SCN1A epilepsy and Huntington disease) was generated across four brain regions (cortex, hippocampus, midbrain and striatum) and six cell classes. Each cell carries binary AAV-A and AAV-B occupancy, joint co-transduction, editor reconstitution, desired, by-product and off-target editing, and functional rescue. Spatial coordinates were drawn uniformly in a unit cube and carry no anatomical information; they are shown only to illustrate how cell-level states are stored. This cohort uses its own governance-defined parameter set, so its absolute stage

probabilities are not interchangeable with those of the 10,000-patient LRRK2 cohort; it is used to examine structure across regions, cell classes and diseases. The share of cumulative loss attributed to a stage was defined as the negative logarithm of its conditional probability divided by the sum over stages. Sobol indices, held-out-disease metrics, the editing-parameter posterior (100,000 samples) and the prior and posterior standard deviations used for value of information follow Sections 2.4-2.7.

3 Results

3.1 A cell-resolved, evidence-constrained model

Each virtual cell carries a state vector spanning genotype, cell class, anatomical position, time, AAV-A and AAV-B exposure and occupancy, co-transduction, editor reconstitution, desired and by-product editing, and molecular, network and functional rescue (Figure 1a,b and Table 2; Methods). Vector occupancy follows Poisson processes, so the marginal probabilities of receiving AAV-A or AAV-B are $P(A) = 1 - \exp(-\lambda_A)$ and $P(B) = 1 - \exp(-\lambda_B)$. The probability that a cell receives both is $P(AB) = P(A)P(B|A)$, and departure from independence is captured by the enrichment $\eta_{AB} = P(B|A)/P(B)$. Joint occupancy is retained at cell and region level before aggregation to patients (Figure 1c). The primary estimand, therapeutic coverage, is the product of conditional transition probabilities, $C = P(AB) \times P(R|AB) \times P(E|R) \times P(M|E) \times P(F|M)$, in which R, E, M and F denote reconstitution, desired editing, molecular rescue and functional rescue. Coverage is a synthetic probability, not a clinical response rate.

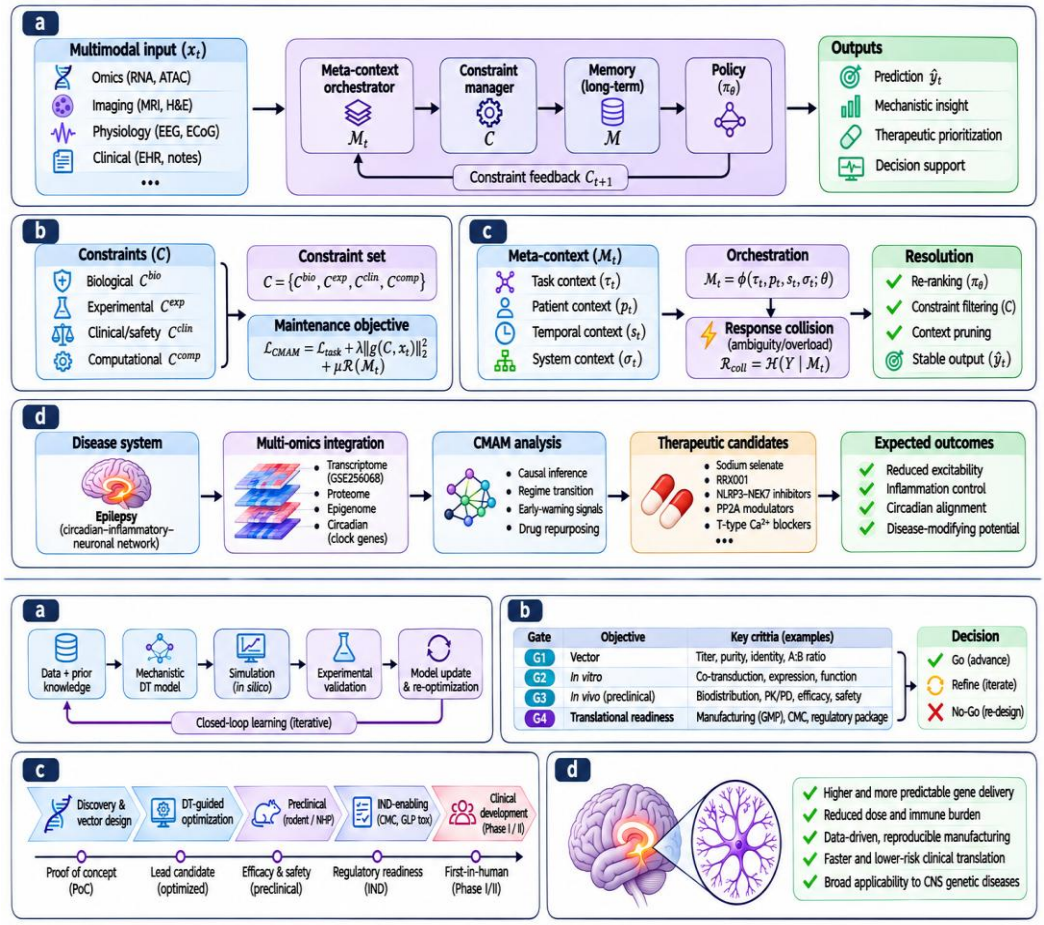


Figure 1*: Architecture of DUAL-AAV PRIME-DT. a, Multiscale therapeutic cascade linking patient state and spatial context to AAV-A and AAV-B exposure, co-transduction, editor reconstitution, prime editing, and molecular, network and functional rescue. The primary estimand is synthetic therapeutic coverage C, the product of conditional transition probabilities; it is not a clinical response rate. b, Causal state graph separating observed inputs from latent biological transitions. c, Spatial and same-cell representation, showing why joint vector occupancy is retained rather than collapsed into population averages. d, Closed-loop acquire update-predict-decide structure used to propagate uncertainty and prioritise the next informative measurement.*

Evidence enters at defined strength. Direct dual-AAV editing, spatial AAV modelling and split-editor technology constrain the delivery and editing layers (Davis et al., 2024; Doman et al., 2023; Duarte et al., 2023; Kissling et al., 2026; Maturana et al., 2022; Mittas et al., 2026; Riedmayr et al., 2023; She et al., 2023; Tasfaout et al., 2024; Trapani et al., 2014), whereas the LRRK2 rescue layers carry wide priors; published percentages are used as envelopes and are never pooled as an efficacy prior (Table 3). A closed acquire-update-predict-decide loop (Figure 1d) converts uncertainty into a ranking of the next measurement.

Evidence stream	Representative evidence	Role in Digital Twin	Interpretation
Direct dual-AAV prime editing	Mouse brain dual-AAV PE; compact PE systems (Davis et al., 2024; Doman et al., 2023)	Constrains brain editing/reconstitution envelope	Mechanistic benchmark, not LRRK2 efficacy
Disease-model prime editing	SCN1A neurological editing (Kissling et al., 2026)	Anchors genotype-to-editing-to-phenotype hierarchy	Contextual validation anchor
CNS AAV spatial modelling	Single-cell co-transduction and brain-region prediction (Duarte et al., 2023; Maturana et al., 2022)	Constrains A/B joint state and spatial heterogeneity	Directly informative to model structure
Large-cargo reconstitution	Trans-splicing, intein and DNA-recombination systems (Colella et al., 2014; Mittas et al., 2026; Riedmayr et al., 2023; Tang et al., 2023; Tasfaout et al., 2024; Trapani et al., 2014)	Constrains reconstitution uncertainty	Technology-class evidence
LRRK2 molecular biomarkers	Rab phosphorylation, lysosomal and inflammatory markers (Cortés et al., 2025; Di Maio et al., 2018; Hop et al., 2024; Jennings et al., 2022; Jennings et al., 2023; Kalogeropoulou et al., 2022; Kania et al., 2023; Kluss et al., 2022; Mabrouk et al., 2026; Steger et al., 2016; Streubel-Gallasch et al., 2021; Vissers et al., 2023; Wauters et al., 2020; Yuan et al., 2024)	Defines patient-state and pharmacodynamic layers	Disease-specific biological grounding

PD gene-therapy clinical experience	AAV programmes and recent clinical reviews (Grote et al., 2024; Kirik et al., 2017; Winston et al., 2025)	Defines translational endpoint and safety context	Clinical context; not dual-AAV PE evidence

Table 3: Evidence architecture and role of each evidence stream.

3.2 Benchmarks constrain but do not transfer

The published in vivo anchors differ in target, species, vector architecture, route and assay, and several are reported as 'up to' maxima: 42% cortical editing with dual AAVs (Davis et al., 2024), about 40% loxP insertion with compact editors (Doman et al., 2023), 34.7% mean Scn1a correction (Kissling et al., 2026) and about 16% correction in retina (She et al., 2023) (Figure 2a). We therefore calibrated a hierarchical editing model with domain-specific effects rather than a common effect. The calibrated latent editing mean was 0.304 (95% credible interval 0.108–0.581), with large between-domain heterogeneity (median $\tau = 0.342$). A single pooled prior underpredicted the high-performing Huntington disease striatum (probability that a replicate met or exceeded the reported value, 0.038), whereas the hierarchical model placed all seven benchmark endpoints within the central range of the predictive distribution (probabilities 0.220–0.565; Figure 2c, Figure 3a,c and Table 4). This argues for region- and architecture-specific calibration.

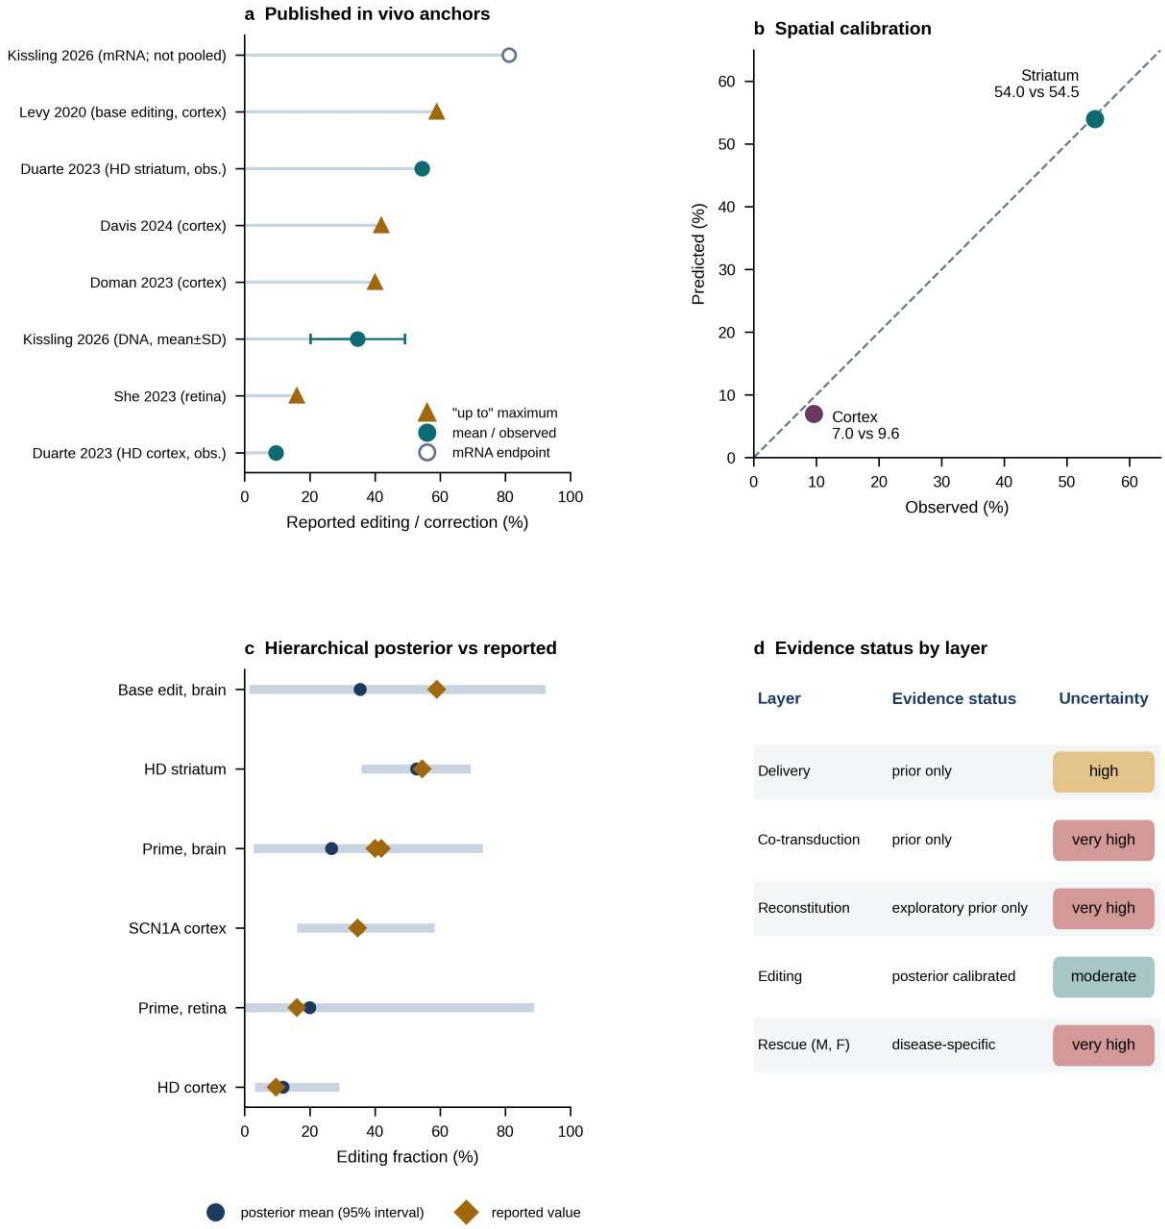


Figure 2: *Evidence anchors and calibration. a, Published in vivo editing or correction anchors, shown as heterogeneous evidence rather than pooled effects. Triangles denote 'up to' maxima, filled circles means or observed values, and the open circle the SCN1A mRNA endpoint, which is not pooled with DNA-editing endpoints. b, Predicted versus observed regional editing in the Huntington disease benchmark. c, Hierarchical posterior means with 95% intervals for each editing domain (bars) compared with reported values (diamonds). d, Evidence status and ordinal uncertainty level of each model layer, taken from the evidence-to-model mapping and uncertainty register (analyst-assigned levels, not statistical estimates). Bars in c show 95% credible intervals and circles posterior means; a shows reported values (error bar, s.d.).*

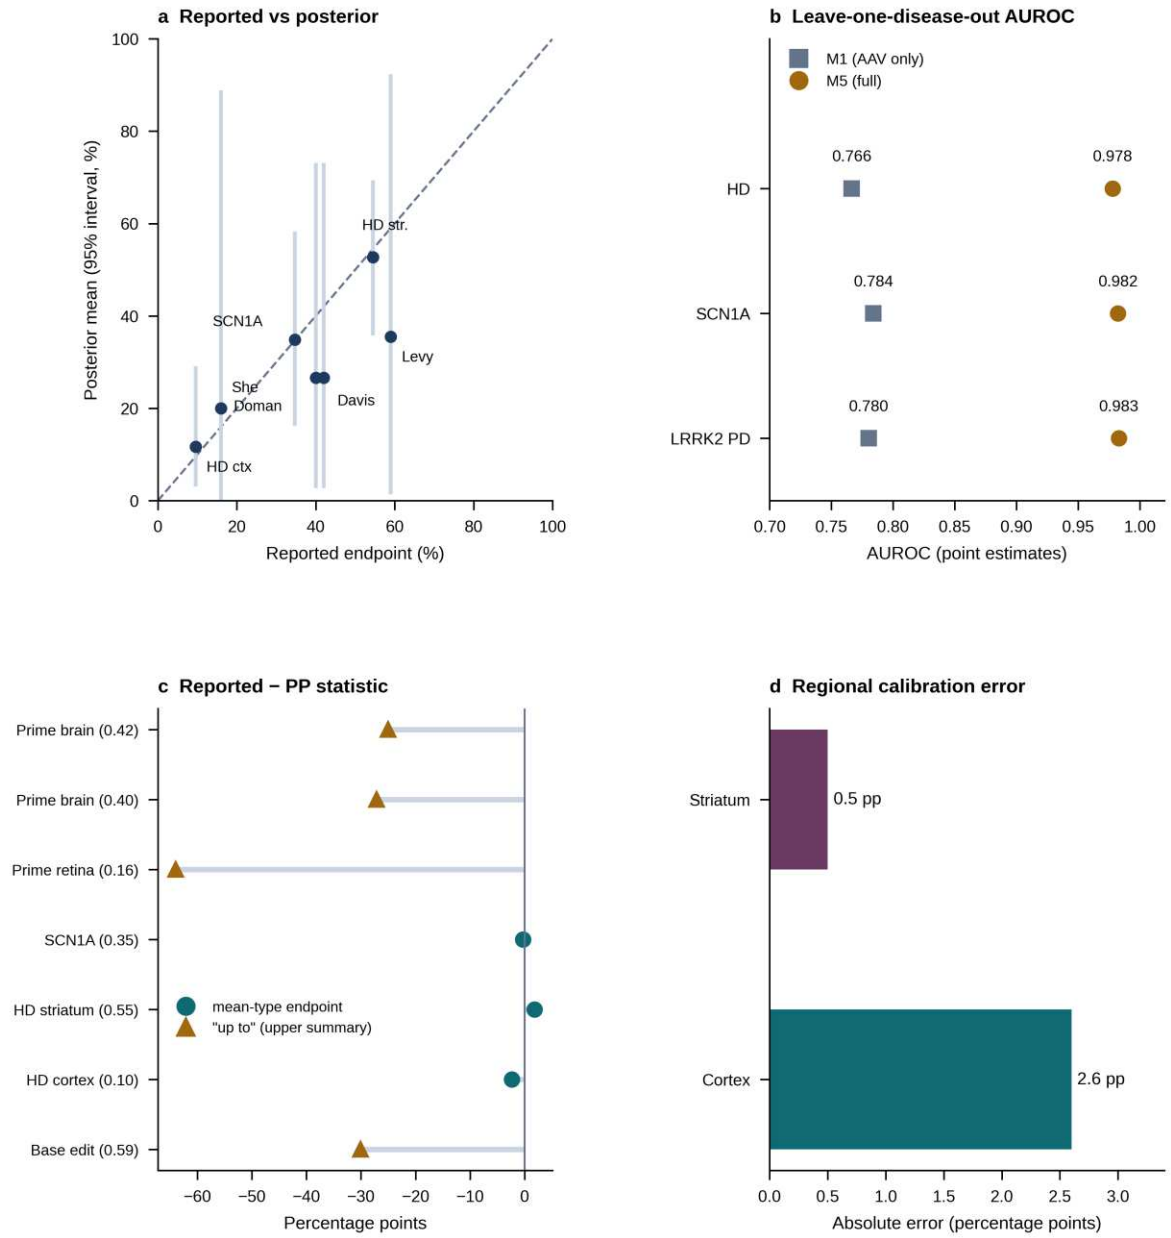


Figure 3: *Spatial calibration and cross-context generalisation. a, Reported benchmark endpoints versus hierarchical posterior means with 95% intervals; the identity line is shown for orientation. b, Leave-one-disease-out AUROC for models M1 and M5 when LRRK2 Parkinson disease, SCN1A epilepsy or Huntington disease is held out (point estimates; no confidence intervals were computed). c, Reported minus posterior-predictive statistic; triangles denote 'up to' endpoints, for which the reported value is expected to lie below the upper summary. d, Absolute regional calibration error for the striatum and cortex benchmark.*

Domain	Posterior mean	Median	95% interval
prime_brain	0.266	0.218	0.027-0.731
prime_retina	0.200	0.080	0.001-0.888
SCN1A	0.349	0.340	0.161-0.583

HD_striatum	0.527	0.528	0.358-0.694
HD_cortex	0.117	0.102	0.031-0.291
base_edit_brain	0.355	0.287	0.014-0.923

Benchmark (domain)	**Reported**	**PP statistic (domain-specific)**	**P(replicate ≥ reported)**
prime_brain	0.420	0.670	0.220
prime_brain	0.400	0.671	0.240
prime_retina	0.160	0.800	0.402
SCN1A	0.347	0.350	0.495
HD_striatum	0.545	0.527	0.446
HD_cortex	0.096	0.119	0.565
base_edit_brain	0.590	0.891	0.236

Table 4: *Hierarchical posterior and posterior-predictive diagnostics. Upper panel, domain posteriors for editing; lower panel, probability that a replicate meets or exceeds the reported value (PP, posterior predictive).*

We tested the spatial component by reproducing the regional prediction of Duarte et al. (2023). Predicted editing of 54% in striatum and 7% in cortex compared with 54.5% and 9.6% observed (mean absolute error 0.0155; root-mean-square error 0.0187; Figure 2b and Figure 3d). The larger relative error in cortex reflects its small denominator and shows that low-abundance regions need explicit intervals. Figure 2d summarises which model layers rest on prior assumptions alone.

3.3 Joint delivery limits therapeutic coverage

Across 10,000 virtual patients (90,000 virtual cells; fixed seed 20261008), mean coverage was 1.24% (median 0.79%; bootstrap 95% interval of the mean, 1.21-1.26%), and the 5^{th}-95^{th} percentile of patients spanned 0.12–3.93% (Figure 4 and Figure 5; Table 5). Mean probabilities along the chain were 0.262 for P(AB), 0.558 for reconstitution, 0.320 for desired editing, 0.549 for molecular rescue and 0.432 for functional rescue. The delivery and co-transduction gate therefore removed the largest share of probability (73.8%), followed by editing (68.0%). Co-transduction was positively enriched on average (η_{AB} mean 1.10, median 1.00; Figure 4b).

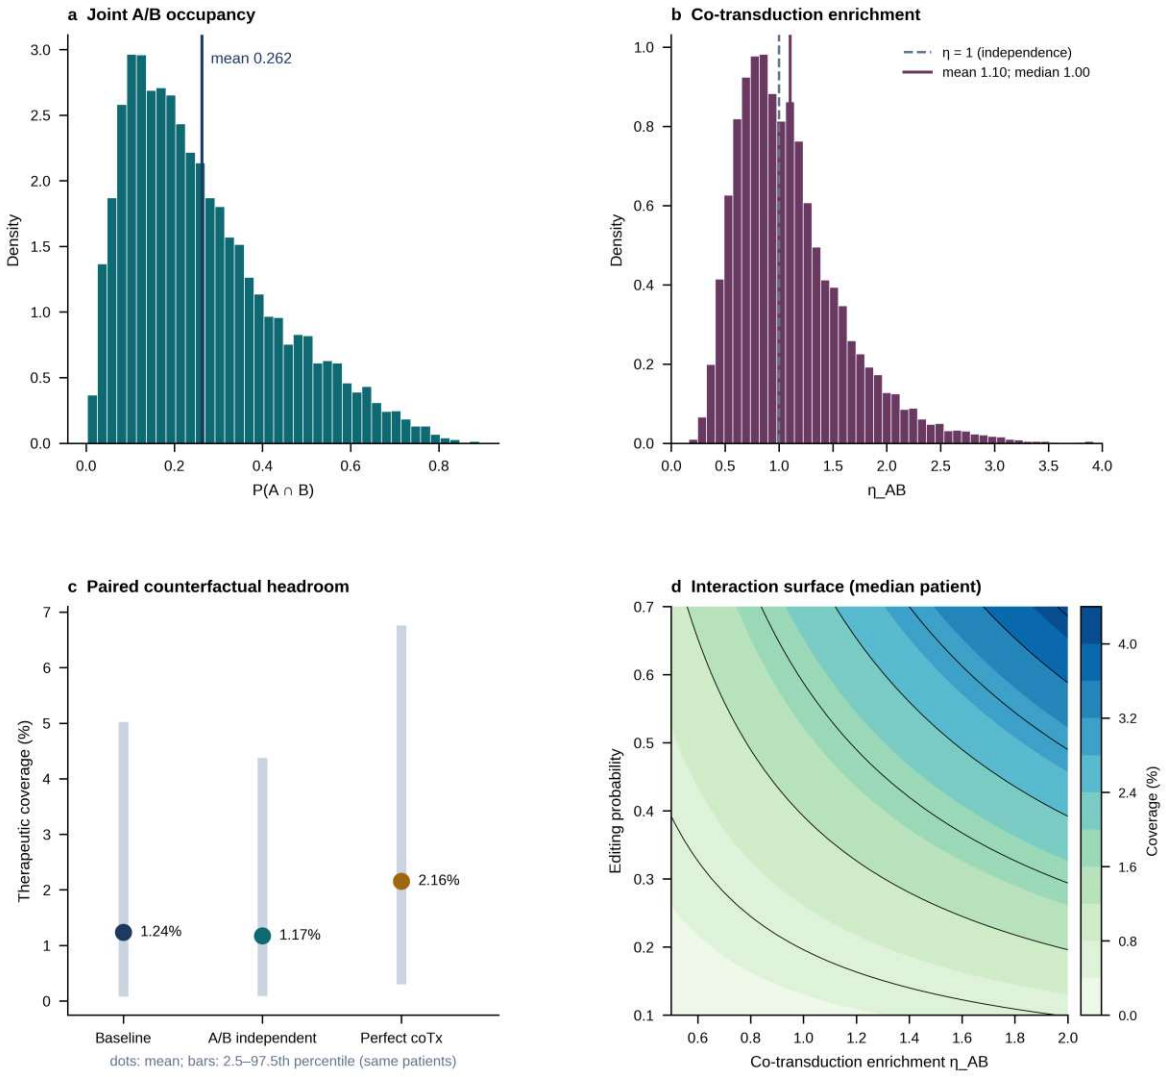


Figure 4: *Joint delivery is a quantifiable bottleneck. a, Distribution of joint AAV-A/AAV-B occupancy across 10,000 virtual patients. b, Distribution of co-transduction enrichment η_{AB} relative to the independence reference of 1 (mean 1.10, median 1.00). c, Paired comparison of baseline, A/B-independent and perfect-co-transduction scenarios in the same virtual patients; dots show means and bars the 2.5th–97.5th percentiles. d, Coverage surface over co-transduction enrichment and editing probability for the median virtual patient. Counterfactuals are mechanistic sensitivity analyses, not clinical efficacy predictions. Centre values are cohort means; no error bars are shown other than the percentile ranges stated.*

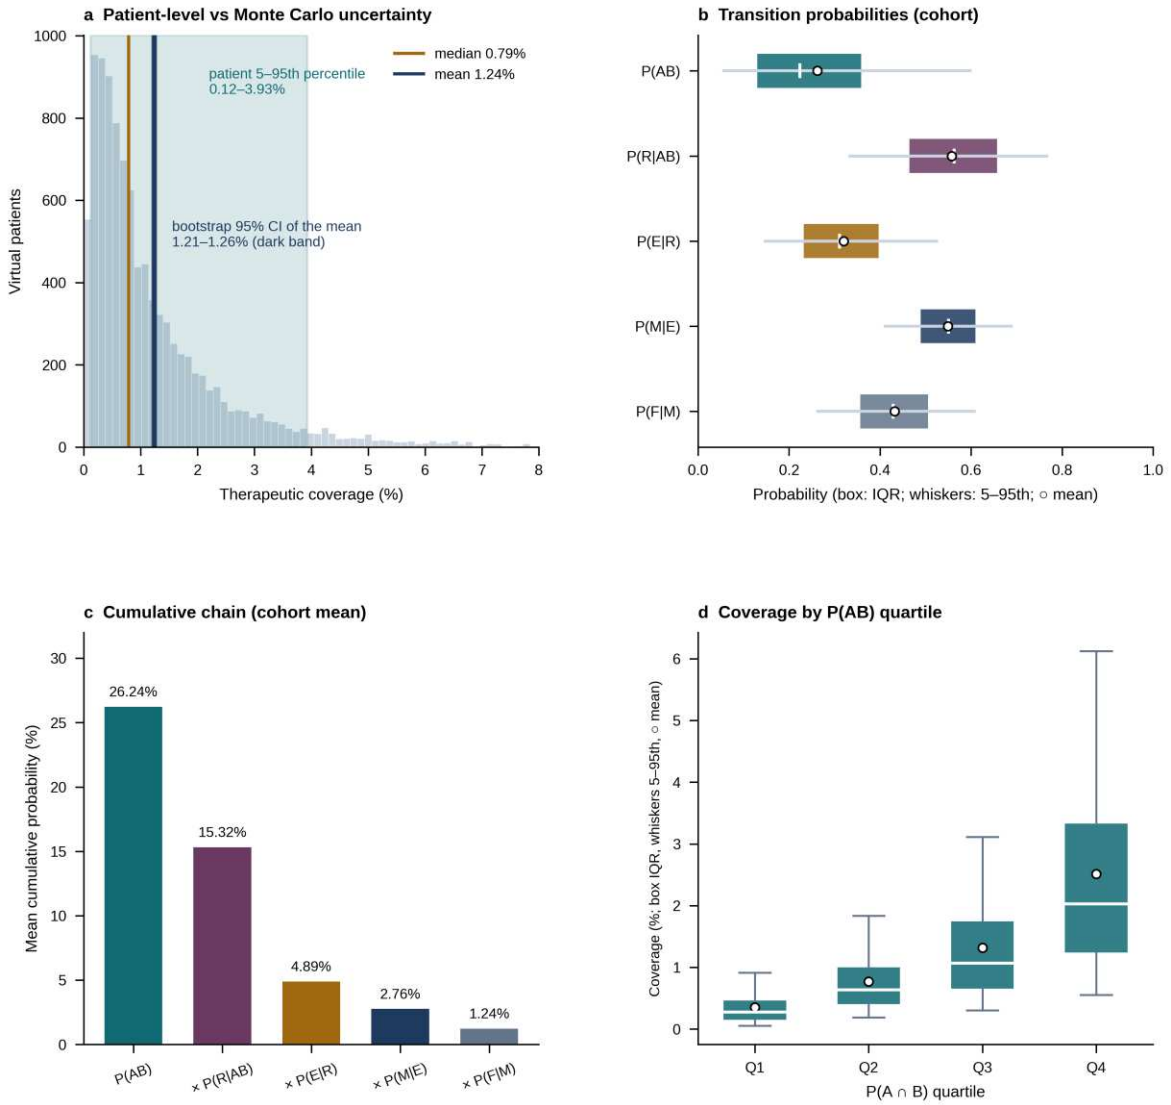


Figure 5: *Distributional uncertainty and transition-state structure. a, Coverage distribution; the shaded band is the patient-level 5th–95th percentile range and the dark band the bootstrap 95% interval of the cohort mean. b, Cohort distributions of conditional transition probabilities (box, interquartile range; whiskers, 5th–95th percentiles; circle, mean). c, Mean cumulative probability after each stage of the serial chain. d, Coverage by P(AB) quartile.*

Quantity	**Estimate**	**95% synthetic interval / context**
Mean therapeutic coverage	0.012385	0.000795-0.050243
Median therapeutic coverage	0.007872	5th-95th percentile: 0.001171-0.039296
Mean P(A and B)	0.262355	0.039084-0.667131
Mean reconstitution probability	0.557804	0.288758-0.801020
Mean desired editing	0.320112	0.122336-0.577297

probability		
Mean molecular rescue probability	0.549387	0.384534-0.715798
Mean functional rescue probability	0.431954	0.232713-0.640959
Bootstrap CI for mean coverage	0.01212-0.01264	Monte Carlo confidence interval; not biological CI

Metric	**Mean**	**Median**	**5th percentile**	**95th percentile**
Therapeutic coverage	0.012385	0.007872	0.0012	0.0393
P(A∩B)	0.2624	0.2240	0.0532	0.6015
Reconstitution	0.5578	0.5628	0.3306	0.7697
Desired editing	0.3201	0.3111	0.1447	0.5273
Molecular rescue	0.5494	0.5504	0.4077	0.6918
Functional rescue	0.4320	0.4286	0.2605	0.6102

Table 5: Virtual-cohort operating characteristics (synthetic). Upper panel, estimates with 95% synthetic intervals; lower panel, mean, median and percentiles

To ask how much this matters, we used paired counterfactuals in which each patient is held fixed and only the dependence structure changes (Table 6). Replacing the joint state with independent vectors lowered mean coverage only slightly, from 1.24% to 1.17% (−5%), whereas perfect co-transduction raised it to 2.16% (+74%; ratio of means 1.74; Figure 4c). Under independence the patient-level median ratio was 1.00 and 49.8% of patients gained, so the patient-level mean relative change (+11.8%) and the ratio of means (−5%) are different estimands. For perfect co-transduction, the patient-level fold change had a median of 2.08 and a mean of 2.49, and 91.9% of patients gained; in the remainder P(AB) already reached its upper bound. These estimands differ because gains are right-skewed and concentrated where baseline coverage is low, so they should be reported separately. Headroom was greatest where joint delivery was poor: across P(AB) quartiles mean coverage rose from 0.35% to 2.51%, whereas the perfect-co-transduction counterfactual rose only from 1.36% to 2.98% (Table 7 and Figure 6d). Co-transduction and editing probability act together, such that improving one has limited value while the other remains constraining (Figure 4d). Perfect co-transduction is an upper bound under the present downstream priors, not a realistic expectation.

Counterfactual	**Mean coverage**	**Median coverage**	**Mean absolute gain**	**Mean relative gain**
Baseline	0.012385	0.007872	Reference	Reference
A/B independence	0.011740	0.008048	-0.000645	+11.8% (patient-level mean)
Perfect co-transduction	0.021582	0.016842	+0.009197	2.49-fold (+148.6%)

				patient-level mean; median 2.08-fold
Ratio of means (perfect / baseline)	1.743× (+74.3%)	—	—	Different estimand

Table 6. Paired counterfactual and stress-test results (synthetic). Upper panel, counterfactuals; lower panel, one-at-a-time ±25% perturbations relative to baseline.

Scenario	Mean coverage	Relative to baseline
Delivery (lambda_A and lambda_B) −25%	0.008637	0.70×
Delivery (lambda_A and lambda_B) +25%	0.015803	1.28×
Co-transduction enrichment −25%	0.009650	0.78×
Co-transduction enrichment +25%	0.014614	1.18×
Reconstitution −25%	0.009289	0.75×
Reconstitution +25%	0.015451	1.25×
Editing −25%	0.009289	0.75×
Editing +25%	0.015481	1.25×
Molecular + functional rescue −25%	0.006967	0.56×
Molecular + functional rescue +25%	0.019351	1.56×

Stratum	N	Mean P(A and B)	Mean editing	Mean functional	Mean coverage	Median coverage
LRRK2_G2019S	4499	0.263	0.312	0.431	0.01214	0.00764
LRRK2_other	1509	0.262	0.322	0.429	0.01225	0.00779
sporadic_PD_like	3992	0.262	0.328	0.434	0.01271	0.00824

P(A and B) quartile	N	Mean P(A and B)	Mean coverage	Perfect-coTx counterfactual
Q1	2500	0.083	0.00352	0.01364
Q2	2500	0.176	0.00772	0.01930
Q3	2500	0.286	0.01317	0.02363
Q4	2500	0.505	0.02513	0.02975

Table 7: Genotype-stratified and P(AB)-quartile-stratified virtual cohort (synthetic)

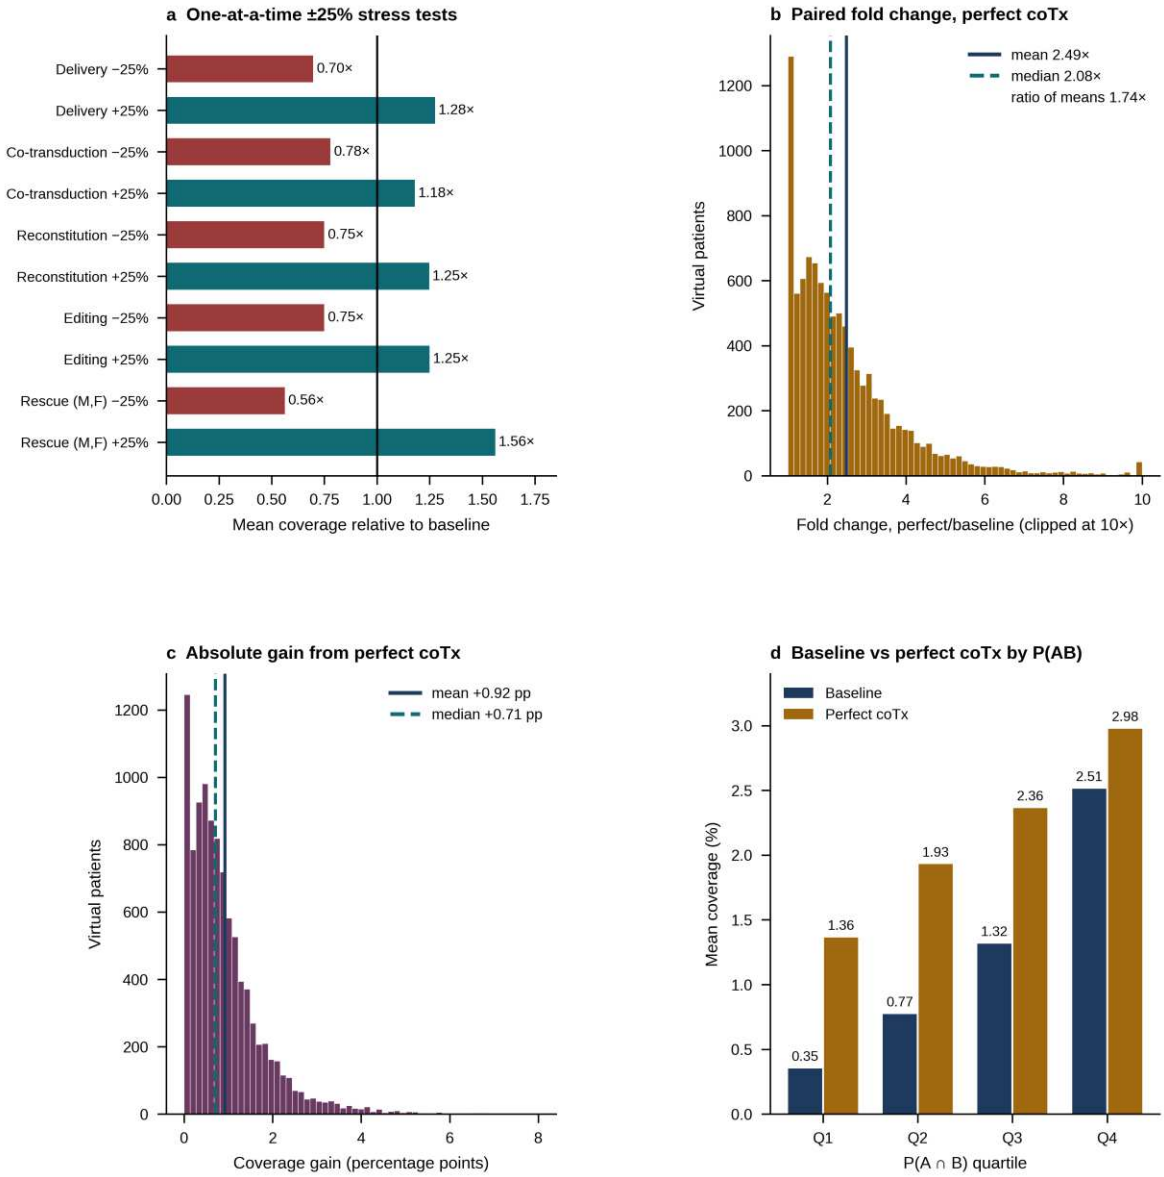


Figure 6: *Paired counterfactuals and stress tests. a, One-at-a-time ±25% parameter perturbations, expressed as mean coverage relative to baseline. b, Patient-level fold change in coverage under perfect co-transduction, in the same virtual patients (clipped at 10-fold for display). c, Absolute coverage gain from perfect co-transduction. d, Mean coverage under baseline and perfect co-transduction by P(AB) quartile.*

3.4 Heterogeneity across virtual patients

Coverage was strongly right-skewed. The probability that a virtual patient exceeded 0.5%, 1%, 2% and 5% coverage was 66.9%, 41.5%, 18.3% and 2.6% (Figure 7a,b); these are decision cut-points, not response criteria. Population spread (patient 5th–95th percentile, 0.12–3.93%) was far larger than Monte Carlo uncertainty in the cohort mean (1.21–1.26%), so intervals around the mean must not be read as biological confidence. Genotype differences were small mean coverage

was 1.21% for LRRK2 G2019S, 1.22% for other LRRK2 variants and 1.27% for a sporadic-like stratum because the current model encodes genotype through vulnerability and downstream rescue, whereas delivery dominates. This should prevent over-interpretation of genotype-specific efficacy until variant-level mechanism and patient-level delivery data are available (Figure 7c and Table 7). Mean coverage declined from 1.40% to 1.07% across disease-stage bins, and patient-level coverage rose with joint occupancy and editing probability (Figure 7c,d).

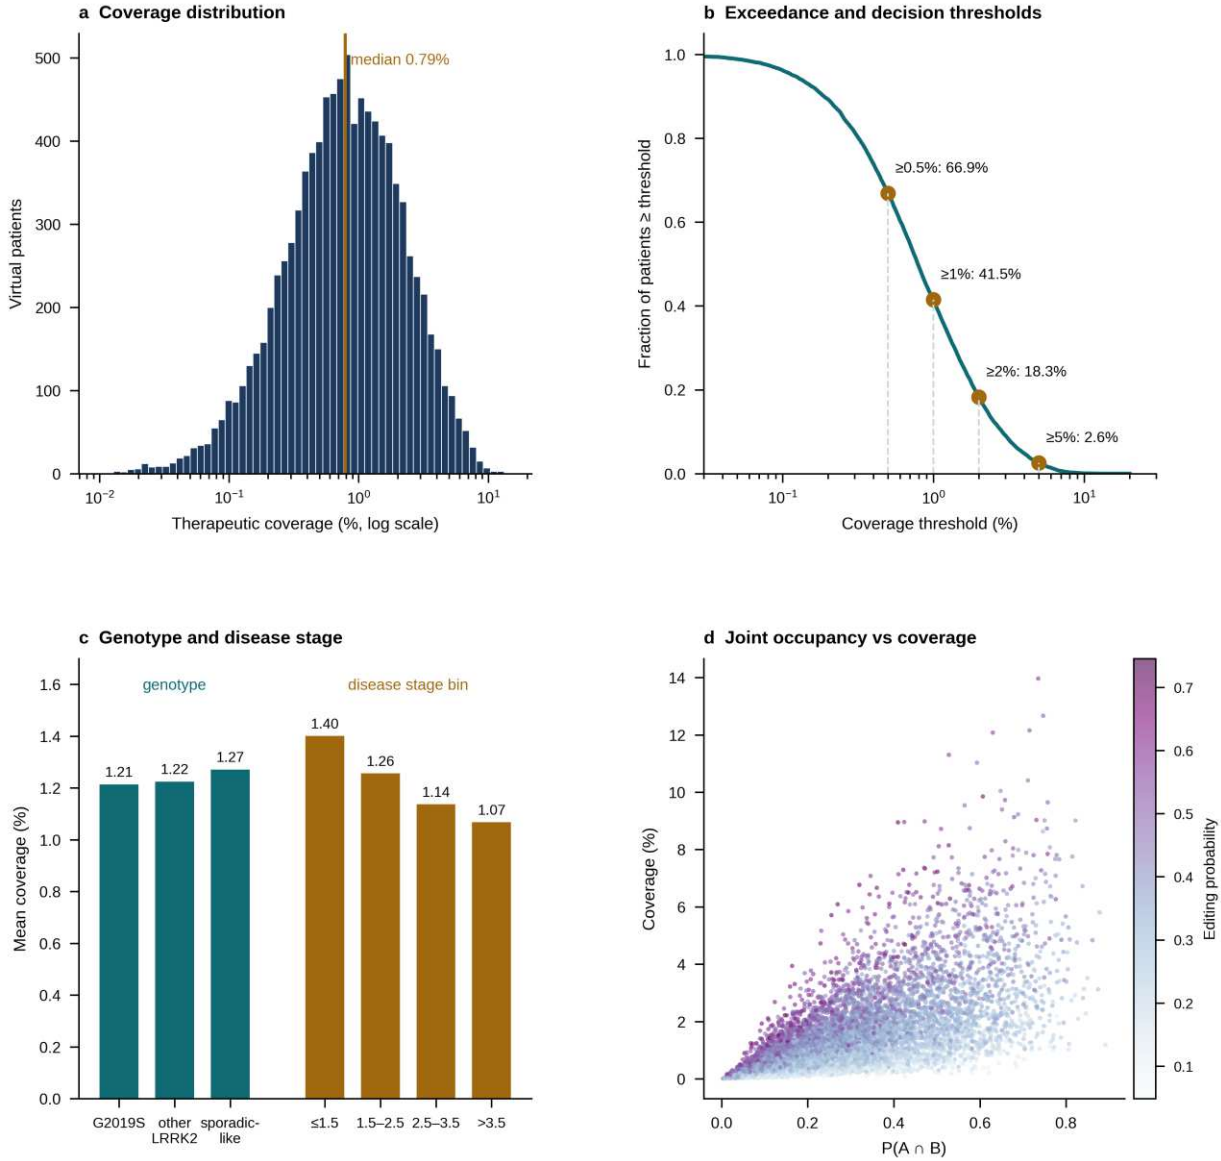


Figure 7: *Virtual-patient heterogeneity. a, Distribution of predicted therapeutic coverage on a logarithmic axis. b, Fraction of patients at or above coverage thresholds of 0.5%, 1%, 2% and 5%. c, Mean coverage by genotype stratum and by disease-stage bin. d, Patient-level joint AAV occupancy versus predicted coverage, coloured by editing probability. All outcomes are synthetic Digital Twin outputs and are not estimates of observed patient treatment response.*

3.5 Sensitivity, stress tests and information value

Rank correlation placed delivery first (λ_A and λ_B; $\rho = 0.60$ each), followed by editing (0.50), reconstitution (0.44), co-transduction enrichment (0.42) and the rescue probabilities (0.38–0.41;

Figure 8a and Table 8). In the cohort λ_B equals 0.92 λ_A for every patient, so the two exposures are collinear and their individual effects cannot be separated. After conditioning on all inputs, standardized coefficients were largest for editing probability (β = 0.34) and co-transduction enrichment (0.34), then for exposure (0.24), functional rescue (0.23), reconstitution (0.21) and molecular rescue (0.14; Figure 8b). Variance-based indices were dominated by the reconstitution shift (mean total effect 0.87), followed by the AAV-A scale (0.30) and the editing shift (0.28; Figure 8c). These indices depend on the ranges assumed for each parameter, which were not equalised, and therefore describe sensitivity to the assumed priors.

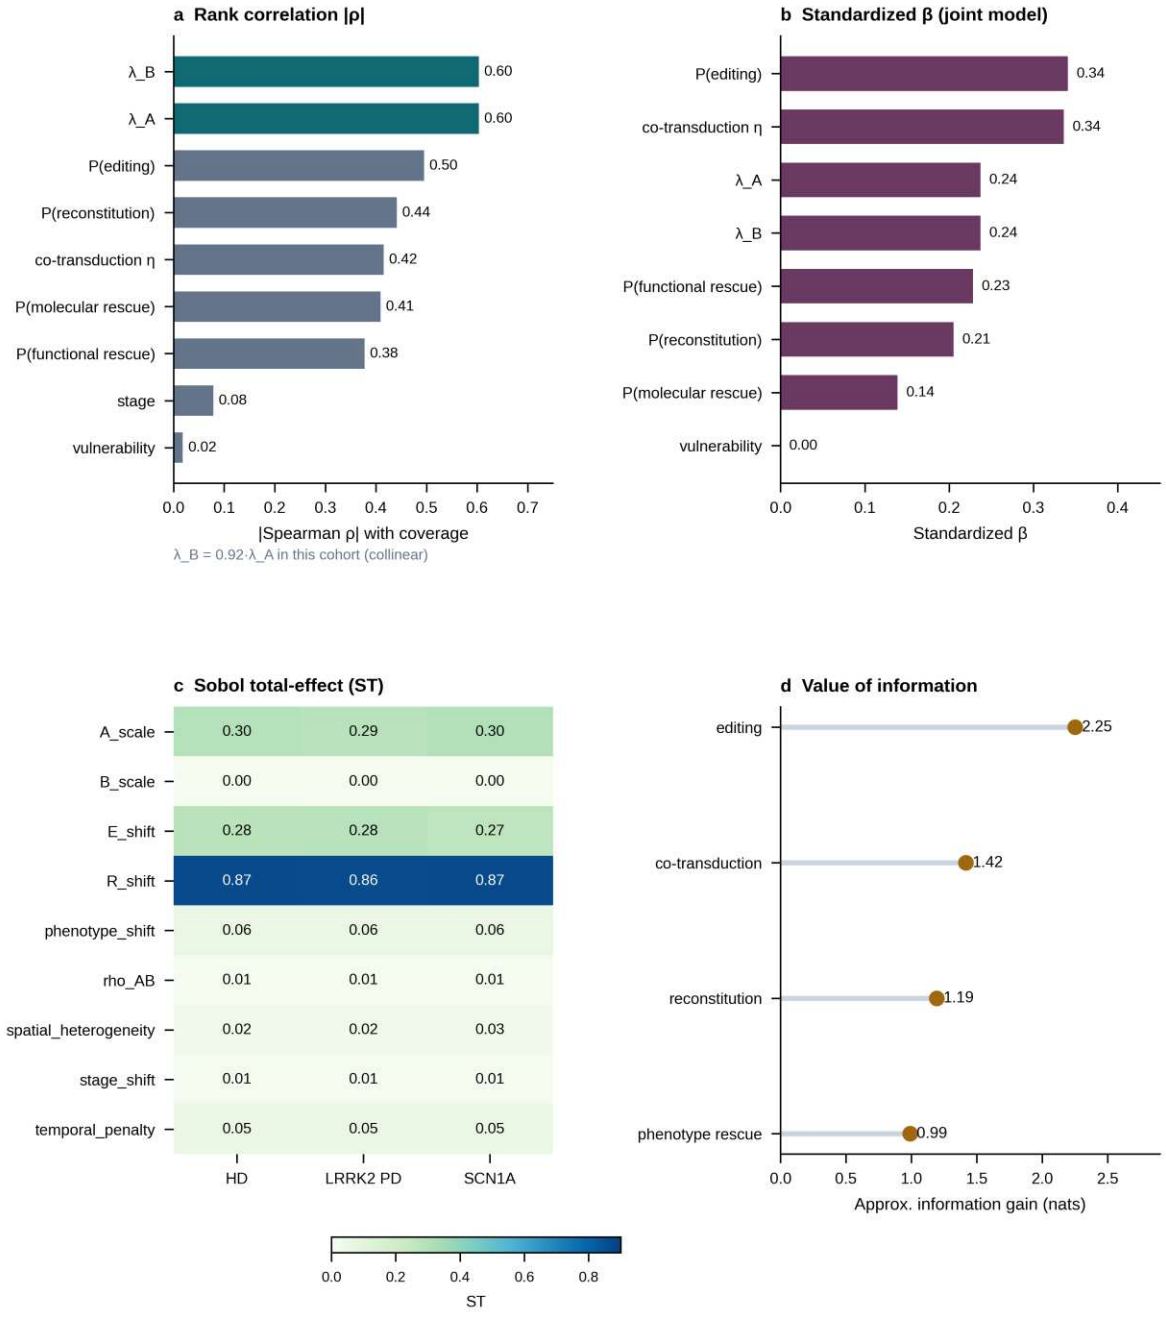


Figure 8: *Global sensitivity analyses. a, Absolute Spearman correlation with coverage. b, Standardized regression coefficients. c, Sobol total-effect indices for Huntington disease, LRRK2*

and SCN1A anchors. d, Approximate information gain. λ_B equals 0.92 λ_A in this cohort, so their values are identical; Sobol indices depend on assumed, non-equalised ranges.

Variable	Spearman rho	Rank	Mean Sobol ST	VOI layer
lambda_A	0.604	1	0.295	—
lambda_B	0.604	2	0.003	—
p_editing	0.495	3	0.277	editing
p_reconstitution	0.442	4	0.867	reconstitution
co_transduction_enrichment	0.415	5	0.007	co_transduction
p_molecular_rescue	0.409	6	0.060	phenotype_rescue
p_functional_rescue	0.378	7	0.060	phenotype_rescue
stage	-0.079	8	0.006	—
vulnerability	-0.018	9	—	—

Priority measurement	Information gain (nats)	Interpretation
editing	2.252	Priority 1
co_transduction	1.417	Priority 2
reconstitution	1.193	Priority 3
phenotype_rescue	0.991	Priority 4

Variable	Standardized β	Interpretation
Desired editing probability	0.341	Largest conditional model leverage
Co-transduction enrichment	0.336	Largest joint-delivery leverage after adjustment
AAV-A exposure λA	0.237	Strong upstream determinant
AAV-B exposure λB	0.237	Strong upstream determinant
Functional rescue	0.228	Important downstream translation determinant
Reconstitution probability	0.205	Major split-editor loss point
Molecular rescue	0.139	Downstream but comparatively lower conditional leverage

Table 8: *Global sensitivity, Sobol and value-of-information results (synthetic). lambda_B equals 0.92 lambda_A in this cohort; Sobol values are those of the nearest shift parameter, averaged across disease anchors, with ranges that were not equalised.*

To compare levers on an equal footing, we rescaled each parameter class by ±25% and recomputed all dependent quantities. Mean coverage changed to 0.70× and 1.28× for delivery, 0.78× and 1.18× for co-transduction enrichment, 0.75× and 1.25× for reconstitution, and 0.75× and 1.25× for editing (Figure 6a and Table 6). Because coverage is a product, stages that enter linearly respond proportionally, whereas delivery and co-transduction respond sub-

proportionally owing to saturating Poisson occupancy and the bound on P(AB); perturbing molecular and functional rescue together gave 0.56× and 1.56× and reflects two stages. Delivery, co-transduction, reconstitution and editing are therefore all first-order levers, and their exact ordering depends on the perturbation scheme.

Finally, expected information gain ranked a direct editing measurement first (2.25 nats), followed by co-transduction (1.42), reconstitution (1.19) and phenotype rescue (0.99; Figure 8d). The ranking reflects assumed measurement noise rather than biological order, but it implies that a small matched panel across the delivery-to-editing chain may be more informative than a large phenotype-only dataset.

3.6 Model depth and generalisation

Adding mechanistic layers improved performance within the synthetic benchmark. Across the nested model ladder, AUROC rose from 0.781 (AAV exposure only) to 0.981 (full model), area under the precision–recall curve from 0.225 to 0.749 and R^2 from 0.067 to 0.683, while the Brier score fell from 0.0788 to 0.0267; the largest step followed the addition of reconstitution and editing (Figure 9 and Table 9). When each disease anchor was held out in turn, the full model retained AUROC of 0.983 (LRRK2), 0.982 (SCN1A) and 0.978 (Huntington disease; Figure 3b). Because outcomes are generated by the same model family, these results show internal consistency and numerical stability, not external predictive validity.

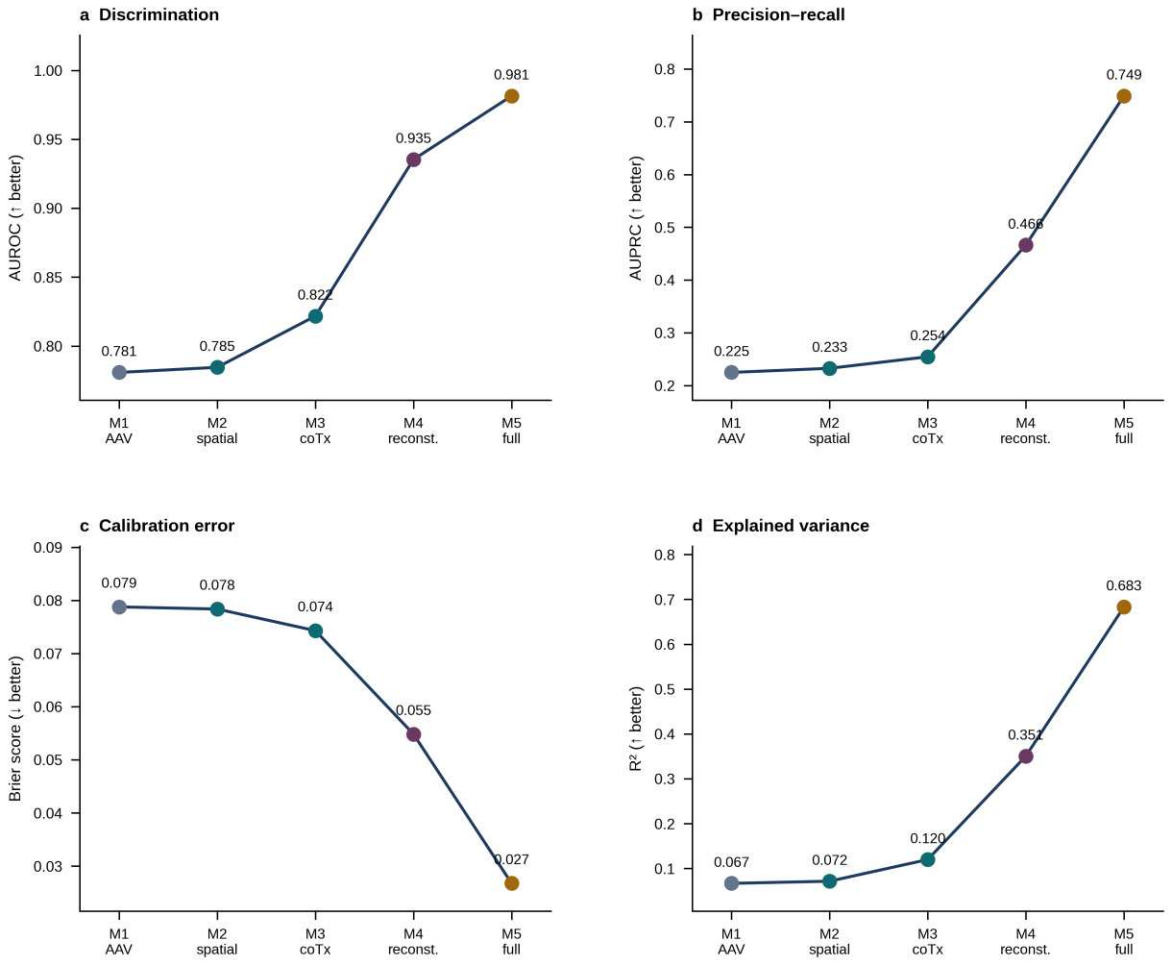


Figure 9: *Mechanistic depth improves performance within the synthetic benchmark. a–d, AUROC, area under the precision–recall curve, Brier score and R² across nested models M1–M5 (M1, AAV exposure; M2, plus spatial structure; M3, plus co-transduction; M4, plus reconstitution and editing; M5, full model). M0 is the reference and is not shown. Values are internal to the synthetic benchmark and are not external validation.*

Model	AUROC	AUPRC	Brier	RMSE	MAE	R2
M1_AAV	0.781	0.225	0.0788	0.281	0.155	0.067
M2_spatial	0.785	0.233	0.0784	0.280	0.154	0.072
M3_coTx	0.822	0.254	0.0743	0.273	0.149	0.120
M4_reconstitution	0.935	0.466	0.0548	0.234	0.110	0.351
M5_full	0.981	0.749	0.0267	0.164	0.053	0.683

Held-out disease	AUROC	AUPRC	Brier	RMSE	MAE	R2
LRRK2_PD	0.983	0.744	0.0320	0.179	0.037	0.597
SCN1A	0.982	0.759	0.0330	0.182	0.039	0.608
HD	0.978	0.722	0.0394	0.199	0.046	0.561

Table 9: Nested model comparison and leave-one-disease-out performance. Values are internal to the synthetic benchmark.

3.7 Identifiability and a falsifiable programme

A terminal endpoint cannot identify the internal failure mode. The same low response can arise from poor AAV-A exposure, poor AAV-B exposure, negative dependence, failed reconstitution, low editing or weak rescue. Exposure and desired editing are moderately to highly identifiable under current evidence, co-transduction and reconstitution only low-to-moderately, phenotype weakly and human efficacy not at all (Table 10). Adding measurements in sequence—joint A/B occupancy, then P(R|AB), P(E|R), molecular rescue and finally the functional endpoint—progressively separates these failure modes, and defines a minimum evidence package interrogating the same cellular or spatial unit (Figure 10c). The framework distinguishes what present evidence supports (ranking of bottlenecks, uncertainty and measurement priority), what is experimentally testable (paired delivery, co-transduction, reconstitution, editing and rescue) and what is not established (human LRRK2 dual-AAV prime-editing efficacy; Figure 10d).

.

Parameter / layer	Identifiability	Why	Required measurement
AAV-A exposure	high	Observed to some degree in current evidence	paired cell-level vector genome/expression data
AAV-B exposure	high	Observed to some degree in current evidence	paired cell-level vector genome/expression data
AAV-A/B co-transduction	low-moderate	Not directly identifiable from a final endpoint alone	joint single-cell measurement

reconstitution	low-moderate	Not directly identifiable from a final endpoint alone	same-cell reconstituted editor measurement
desired DNA editing	moderate-high	Observed to some degree in current evidence	standardized editing assay
by-product/off-target	moderate	Observed to some degree in current evidence	orthogonal sequencing
molecular rescue	moderate	Observed to some degree in current evidence	RNA/protein
phenotype	low	Not directly identifiable from a final endpoint alone	longitudinal disease-specific readouts
human clinical efficacy	not identifiable	Not directly identifiable from a final endpoint alone	requires human interventional evidence

Measurement set	**Purpose**	**Interpretation**
AB only	Estimate joint delivery	Separates vector exposure from downstream states
AB + reconstitution	Estimate P(R|AB)	Quantifies split-editor assembly loss
AB + reconstitution + editing	Estimate P(E|R)	Separates editor activity from delivery/reconstitution
Add molecular rescue	Estimate P(M|E)	Tests whether sequence correction changes pathway biology
Full chain	Estimate P(F|M) and total coverage	Supports end-to-end validation and counterfactual testing

Table 10: Identifiability matrix and measurement sets required for a falsifiable twin

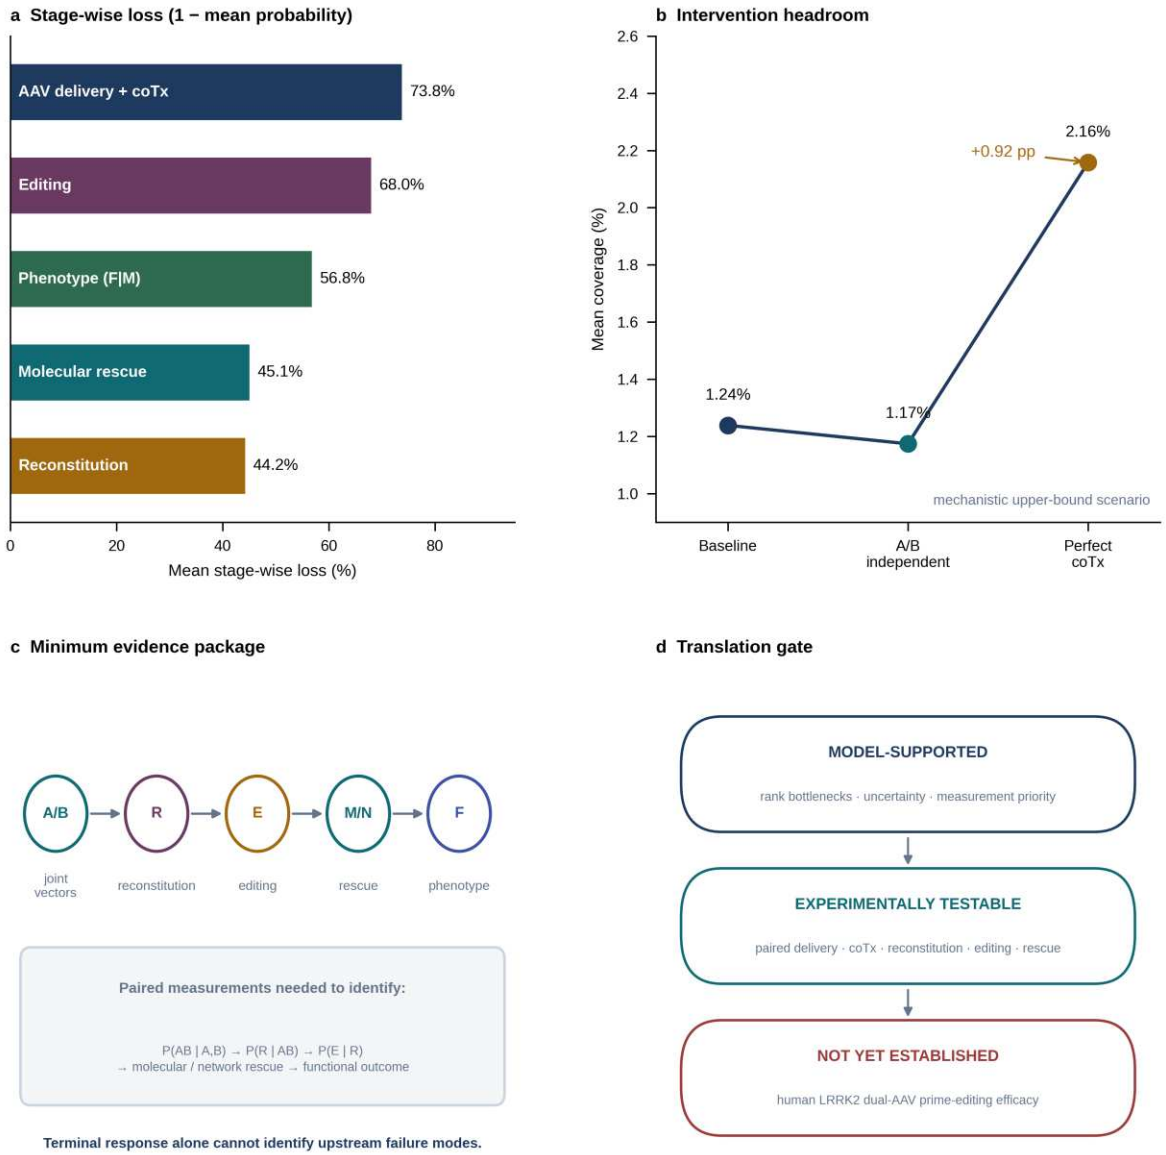


Figure 10*: From virtual patient to prospective validation. a, Mean stage-wise loss, defined as one minus the mean conditional transition probability of each stage; stage losses compound multiplicatively and are not additive. b, Mean coverage under mechanistic counterfactuals, showing the headroom from perfect co-transduction. c, Minimum evidence package required to distinguish joint delivery, reconstitution, editing and downstream rescue, shown as a linked measurement chain. d, Translation gates separating conclusions supported by the current computational and evidence-based from experimentally testable requirements and the still-unestablished claim of human LRRK2 dual-AAV prime-editing efficacy. Values in a and b are cohort means.*

Safety is carried as an explicit competing state. Anti-AAV antibodies, route- and species-dependent biodistribution and neuropathology, dose-linked toxicity and by-product editing each map to a prospective measurement, so a parameter set that raises coverage while increasing predicted adverse exposure should not be ranked on coverage alone (Hinderer et al., 2018; Okai et al., 2025; Rico et al., 2026; Schulz et al., 2023) (Table 11).

Constraint	Evidence basis	Digital Twin representation	Prospective measurement
Pre-existing anti-AAV immunity	Human AAV immunology literature (Leborgne et al., 2020; Mendell et al., 2022; Rapti & Grimm, 2021; Schulz et al., 2023; Verdera et al., 2020; Weber, 2021)	Exposure attenuation / eligibility modifier	Neutralising and binding antibody status
Biodistribution	CNS and NHP studies (Bolon et al., 2024; Bunuales et al., 2024; Chen et al., 2023; Drouyer et al., 2024; Foust et al., 2009; Goertsen et al., 2022; Hocquemiller et al., 2022; Moyer et al., 2025; Okai et al., 2025; Rico et al., 2026)	Region/cell-type exposure distribution	Vector genomes/expression across CNS and peripheral tissues
Neuropathology	Nonclinical pathology reviews and NHP studies (Bolon et al., 2024; Hocquemiller et al., 2022; Rico et al., 2026)	Safety penalty coupled to exposure	Histopathology + clinical observations
Dose-linked cytotoxicity	AAV toxicity literature (Hinderer et al., 2018)	Nonlinear safety cost	Exposure-response safety curve
Off-target/by-product editing	Prime-editing optimisation and specificity studies (Dacquay et al., 2025; Kissling et al., 2026; Koeppel et al., 2023; Mathis et al., 2023; Yu et al., 2023)	Separate safety state O	Orthogonal genomic sequencing
Longitudinal patient state	LRRK2/PD biomarker literature (Brooker & Gonzalez-Latapi,	Time-varying covariates	Serial molecular, imaging and clinical measurements

	2025; Cortés et al., 2025; Grote et al., 2024; Jennings et al., 2022; Jennings et al., 2023; Kirik et al., 2017; Mabrouk et al., 2026; Siderowf et al., 2023; Simuni et al., 2024; Vissers et al., 2023; Winston et al., 2025; Yuan et al., 2024)		

Table 11: Translational safety constraints to be carried into prospective updates.

3.8 Cell-level structure, outcome classes and cross-disease behaviour

To examine how co-transduction varies across regions and cell classes, we analysed a cell-level cohort (Section 2.9), whose spatial coordinates are illustrative only (Figure 11a,b). Across 90,000 cells, 68.8% carried AAV-A, 63.8% carried AAV-B and 51.0% carried both, a co-transduction coefficient of 1.16. Joint occupancy ranged from 42.1% in hippocampus to 56.6% in striatum and from 32.4% in microglia to 58.7% in medium spiny neurons, and in every region it exceeded the fraction expected under independence (Figure 11c,d).

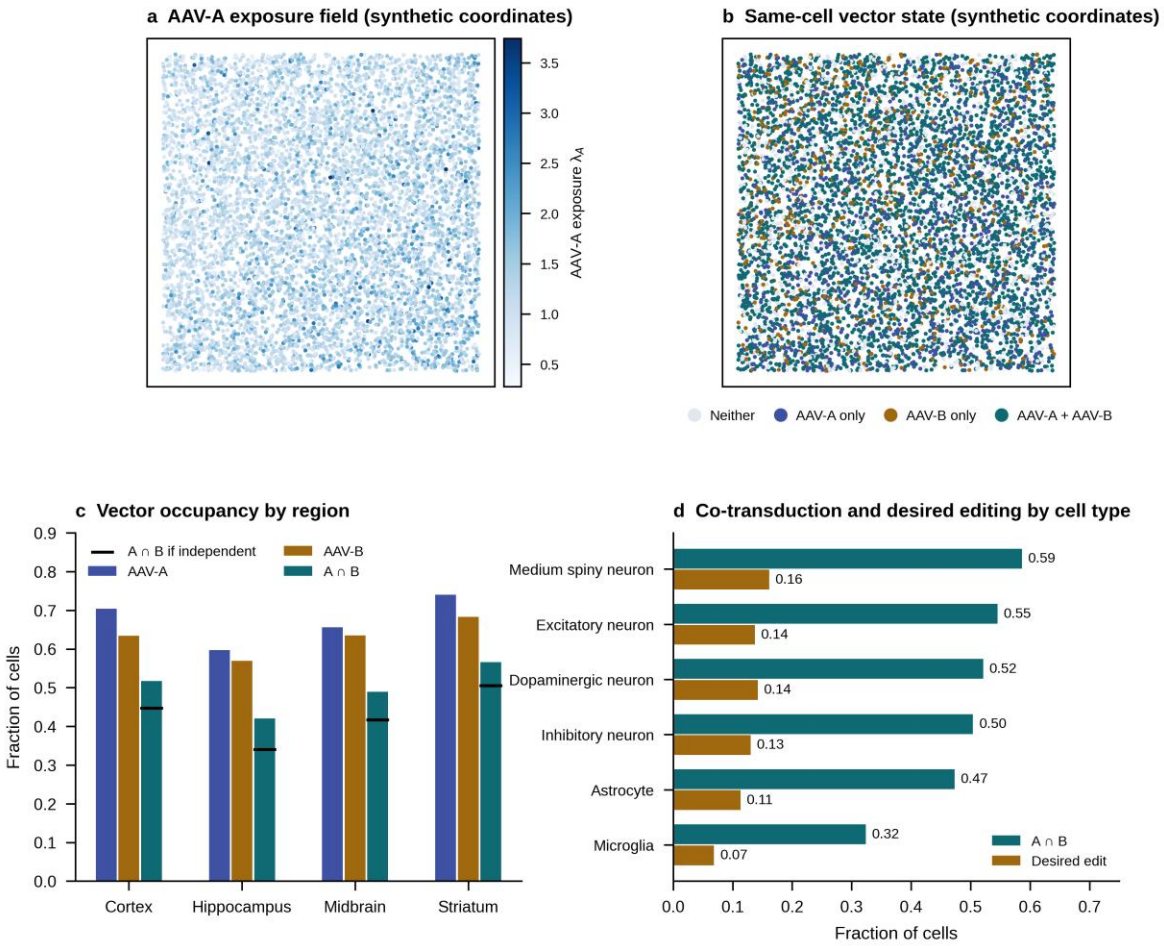


Figure 11. *Spatial delivery and same-cell vector state in the cell-level cohort. a, AAV-A exposure for 7,000 LRRK2 cells plotted at synthetic coordinates, which carry no anatomical information. b, Same cells coloured by vector state (neither, AAV-A only, AAV-B only, AAV-A + AAV-B). c, Fraction of cells carrying AAV-A, AAV-B or both in each region (all 90,000 cells); black ticks mark the joint fraction expected if vectors were independent. d, Co-transduction and desired-editing fractions by cell class. All values are synthetic.*

Reconstitution was the next limiting step. Across the three diseases, 22.2-25.0% of cells reconstituted the editor and 12.2-14.3% carried the desired edit, whereas by-product and off-target edits were found in 1.3-1.4% and 1.1-1.3% of cells (Figure 12a). Desired edits made up 83-84% of editing events, a ratio of 4.9:1 to 5.4:1 against by-product and off-target events combined (Figure 12b). Conditional probabilities were similar across diseases: 0.45–0.47 for reconstitution given co-transduction, 0.54–0.58 for desired editing given reconstitution and 0.70–0.72 for functional rescue given desired editing (Figure 12c). Desired editing varied from 0.05 in hippocampal microglia to 0.19 in striatal medium spiny neurons (Figure 12d).

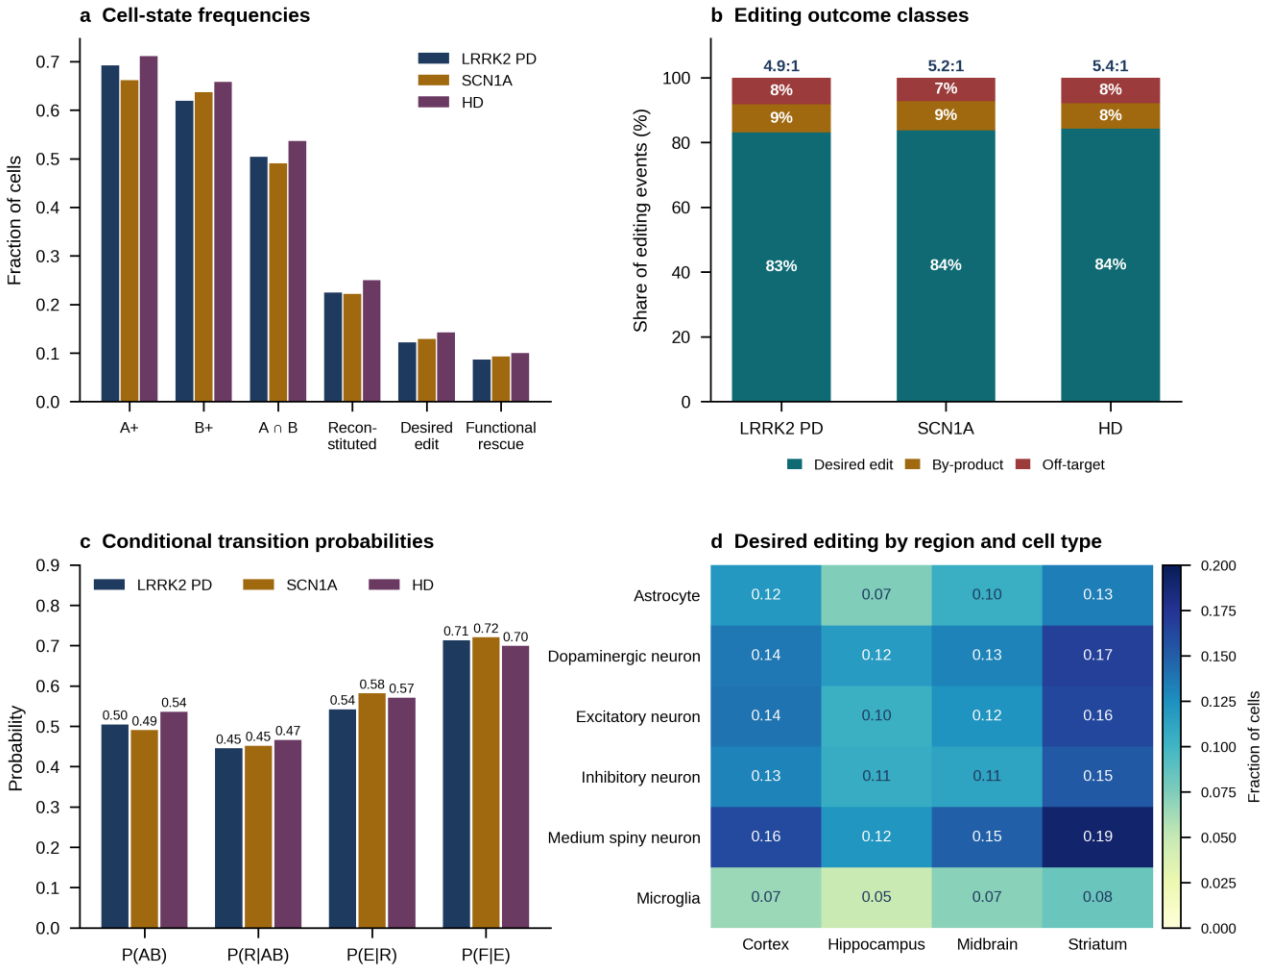


Figure 12: *Reconstitution and editing outcome classes in the cell-level cohort. a, Fraction of cells in each state for each disease (30,000 cells per disease). b, Share of editing events that are desired edits, by-products or off-target events; the ratio above each bar is desired to undesired events. c, Conditional transition probabilities in each disease. d, Desired-editing fraction by region and cell class. All values are synthetic.*

In this cohort reconstitution carried the largest share of cumulative loss in every disease (33.1–33.5% of the negative log coverage), followed by co-transduction (27.0–30.0%), editing (22.8–25.0%) and functional rescue (13.8–15.5%; Figure 13a), and 8.7–10.0% of cells were retained to functional rescue (Figure 13b). This ordering differs from the primary patient cohort, in which joint delivery removes the most probability. The two cohorts parameterise the chain differently: in the LRRK2 subset of the cell-level cohort co-transduction (0.50 versus 0.26) and editing (0.54 versus 0.32) are higher and reconstitution is lower (0.45 versus 0.56; Figure 13c). The identity of the dominant stage therefore depends on the prior set and should be treated as an assumption to be tested, which is the purpose of the matched measurements in Section 3.7. The co-transduction coefficient ranged from 1.10 to 1.25 across regions and diseases and was highest in the hippocampus (Figure 13d).

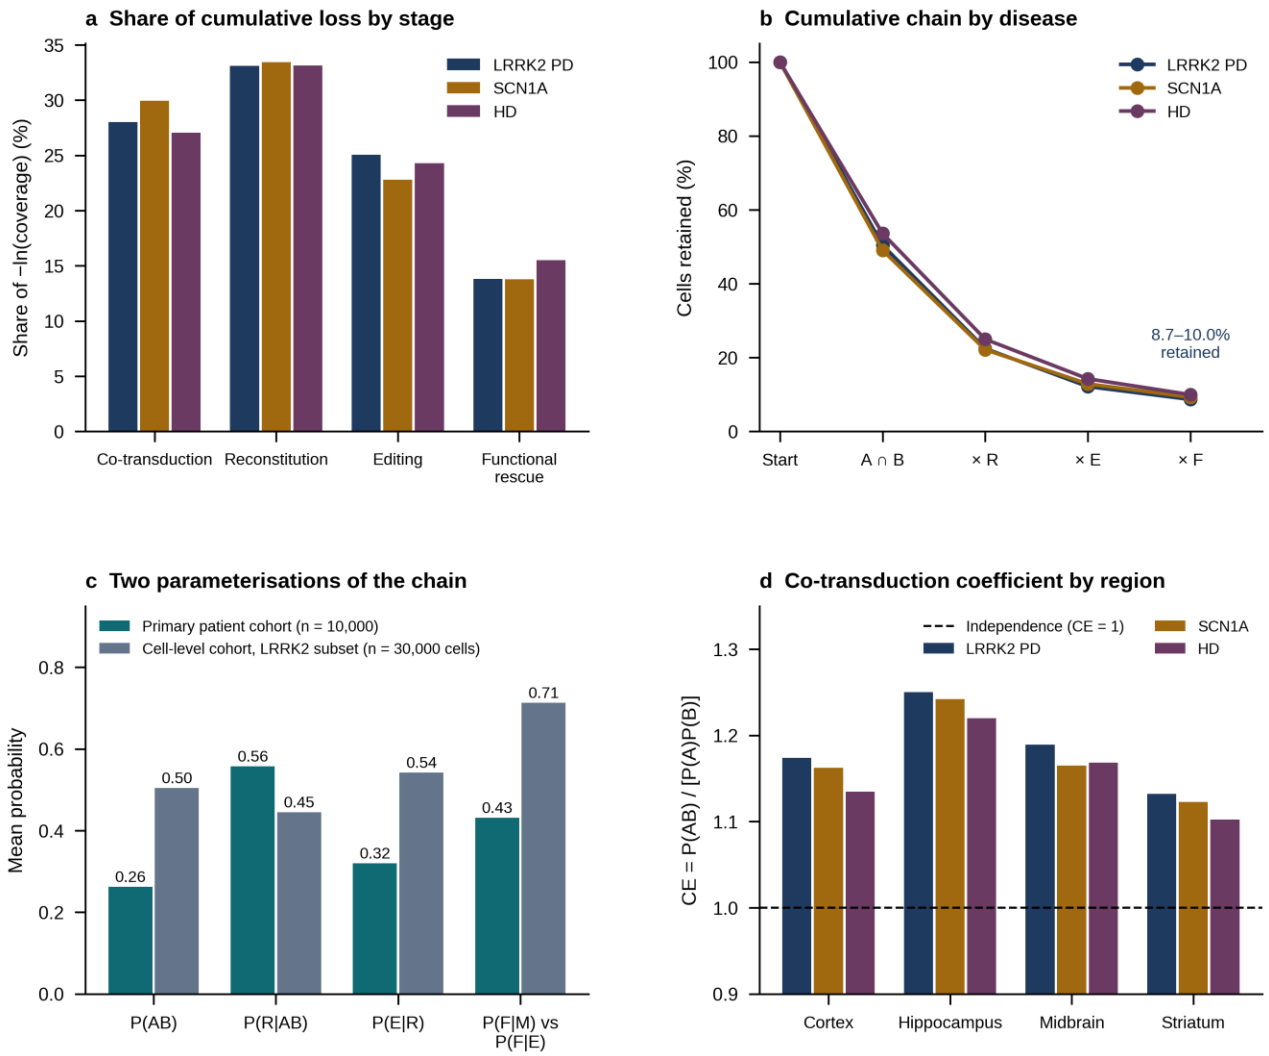


Figure 13: *Bottlenecks across diseases and across parameterisations. a, Share of cumulative loss, defined as the negative logarithm of each conditional stage probability divided by the sum. b, Fraction of cells retained along the chain. c, Mean stage probabilities in the primary patient cohort (n = 10,000) and in the LRRK2 subset of the cell-level cohort (n = 30,000 cells); the final stage is conditioned on molecular rescue in the former and on desired editing in the latter, so the last pair is not directly comparable. d, Co-transduction coefficient by region; the dashed line marks independence. All values are synthetic.*

Variance-based sensitivity in the three disease anchors agreed with the patient-level analysis. The reconstitution shift had a first-order index of 0.75 and a total-effect index of 0.86–0.87, and interactions added 0.12 to its total effect, compared with 0.07–0.08 for the editing shift and 0.07 for the AAV-A scale (Figure 14). These indices depend on the assumed parameter ranges.

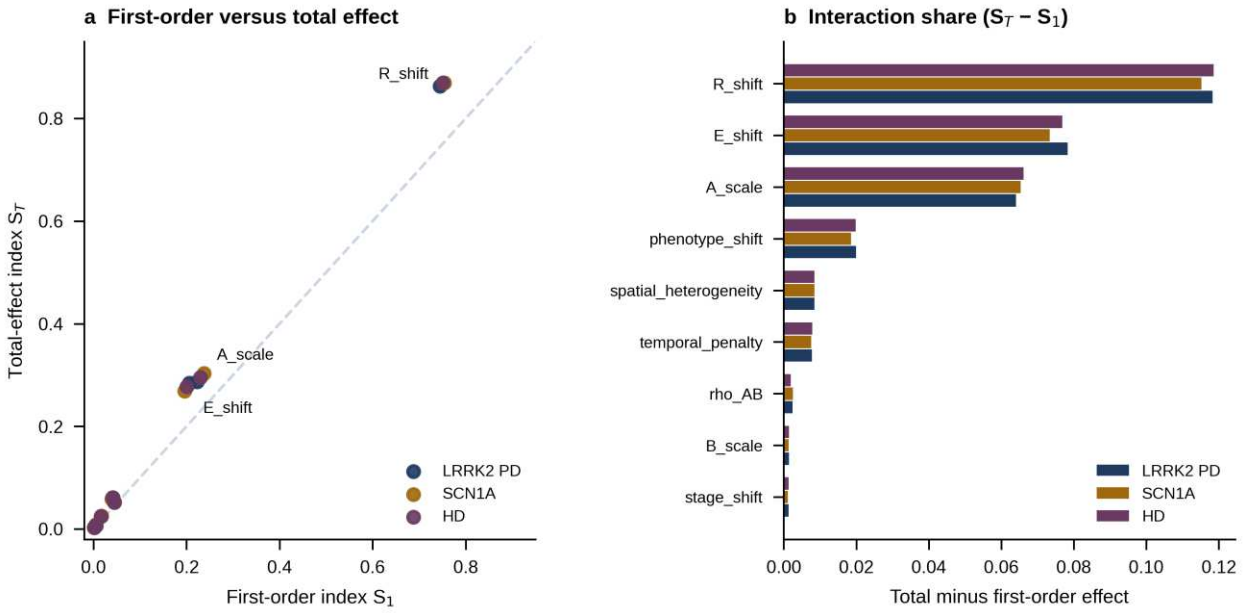


Figure 14: *Variance-based sensitivity and interactions across diseases. a, First-order (S1) against total-effect (ST) Sobol indices for nine shift parameters in three disease anchors; points above the dashed identity line indicate interactions. b, Interaction share (ST minus S1) by parameter. Indices depend on the assumed, non-equalised parameter ranges.*

When each disease was held out in turn, the full model retained an AUROC of 0.98, a Brier score of 0.032–0.039 and an R^2 of 0.56–0.61, whereas the simpler models had AUROC values of 0.77–0.94, Brier scores of 0.080–0.124 and R^2 values of at most 0.02; adding co-transduction alone (M3) gave negative R^2 values (−0.30 to −0.43; Figure 15). Held-out calibration therefore improved only once reconstitution and editing were modelled, although all outcomes come from the same model family and this does not constitute external validation.

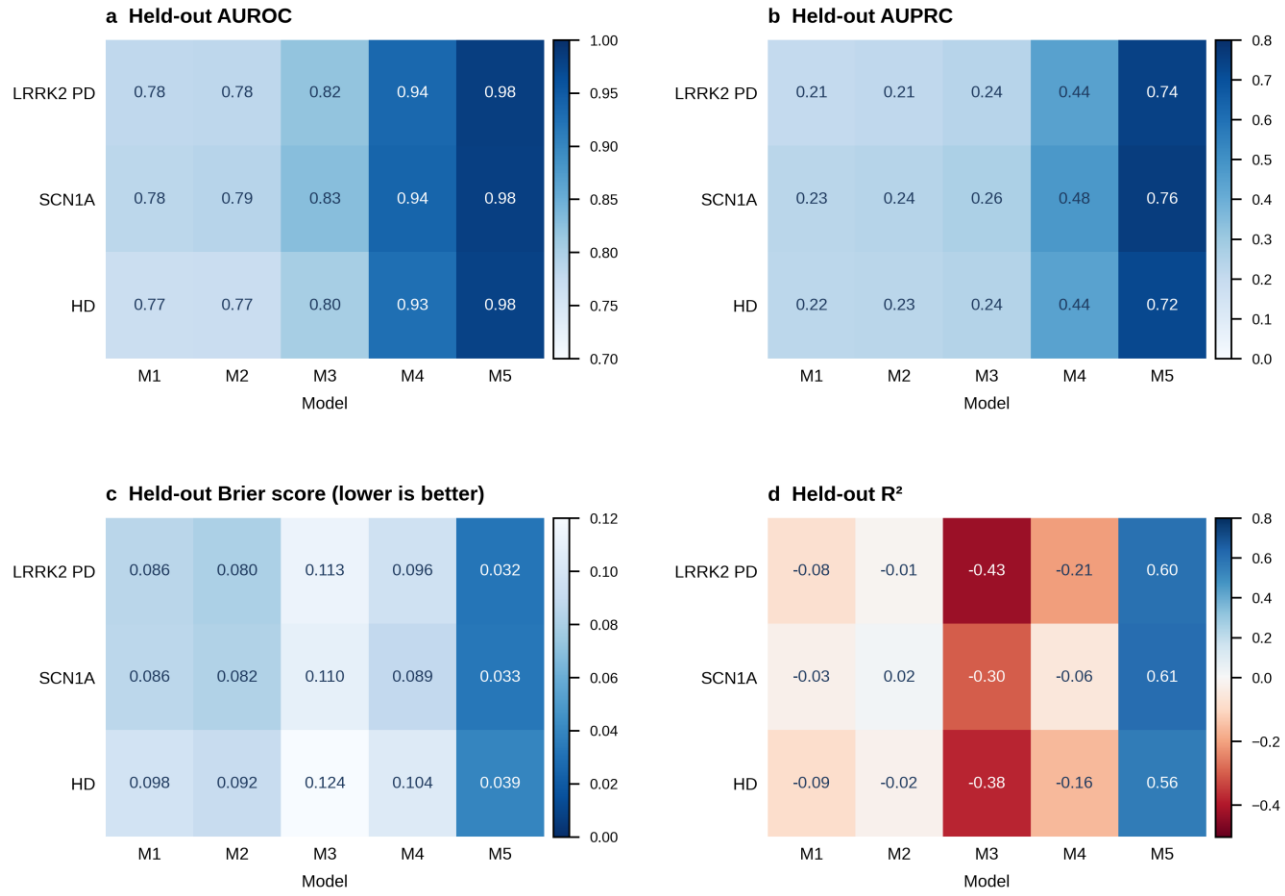


Figure 15: *Held-out-disease performance of the nested models. a, AUROC. b, Area under the precision-recall curve. c, Brier score (lower is better). d, R², centred at zero. Rows give the disease held out and columns the model (M1, AAV exposure; M2, plus spatial structure; M3, plus co-transduction; M4, plus reconstitution and editing; M5, full model). Point estimates; outcomes are generated by the same model family.*

The editing parameter had a posterior mean of 0.352 (95% interval 0.27–0.44; Figure 16a). Identifiability was high only for vector exposure, low to moderate for co-transduction and reconstitution, low for phenotype and absent for human efficacy, and uncertainty was rated very high for co-transduction, reconstitution, phenotype and model structure (Figure 16b,c). Under the assumed measurement noise, one measurement of each layer is expected to reduce its standard deviation from 0.47 to 0.05 for editing, 0.20 to 0.05 for co-transduction, 0.22 to 0.07 for reconstitution and 0.25 to 0.09 for phenotype rescue (Figure 16d).

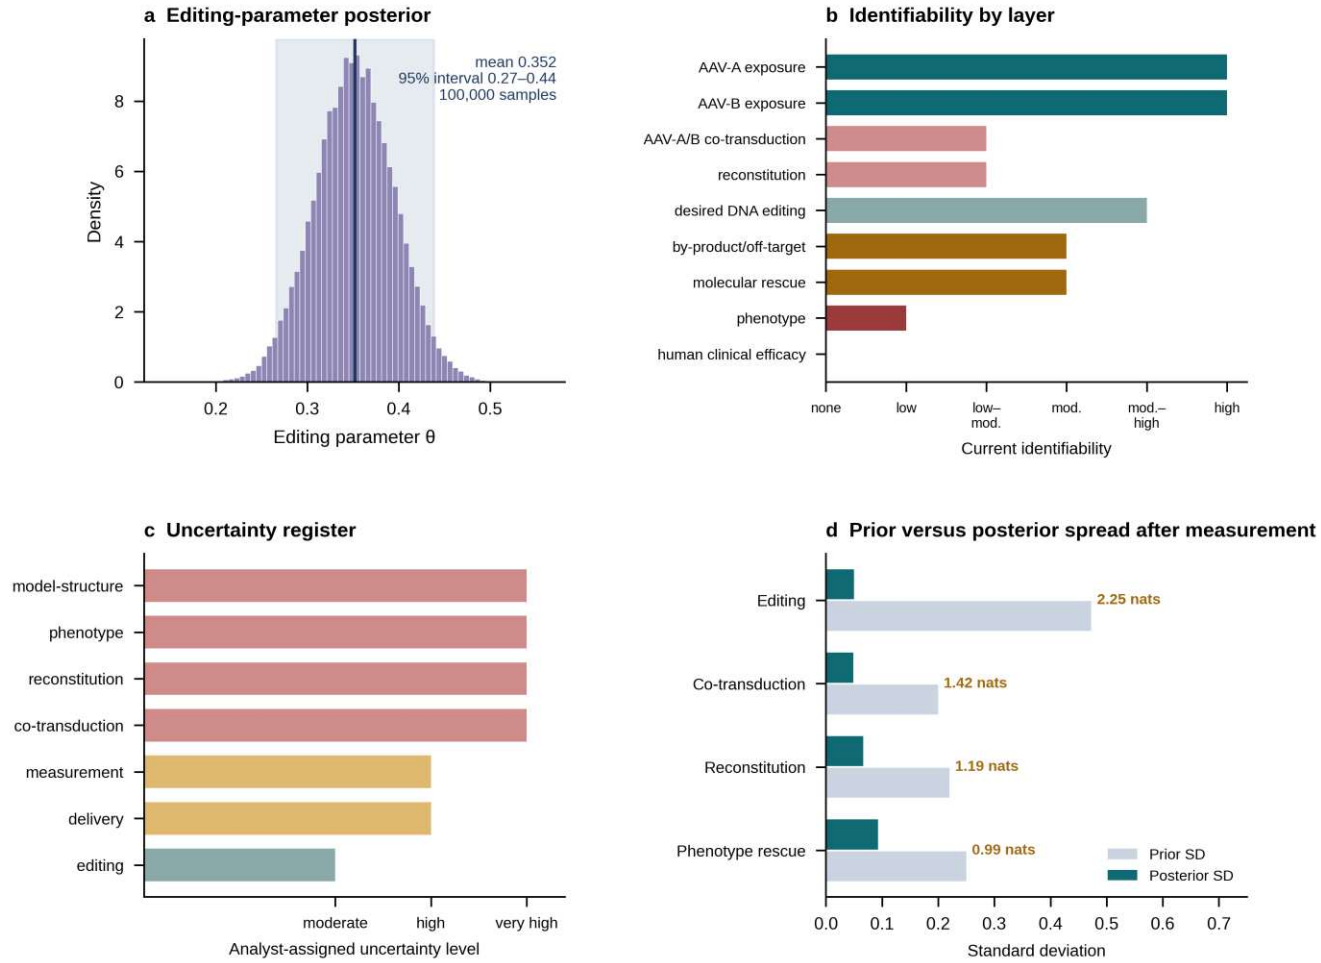


Figure 16: *Uncertainty and identifiability. a, Posterior distribution of the editing parameter (100,000 samples; shaded band, 95% interval; line, mean). b, Current identifiability of each model layer. c, Analyst-assigned uncertainty level of each model component (ordinal, not statistical estimates). d, Prior and posterior standard deviation for each value-of-information layer, with expected information gain in nats.*

Figure 17 summarises the architecture that links delivery, co-transduction, reconstitution, editing and rescue, and Figure 18 places the closed-loop workflow, the validation gates and the proposed translational pathway in sequence; the translational stages are proposals and none has been completed.

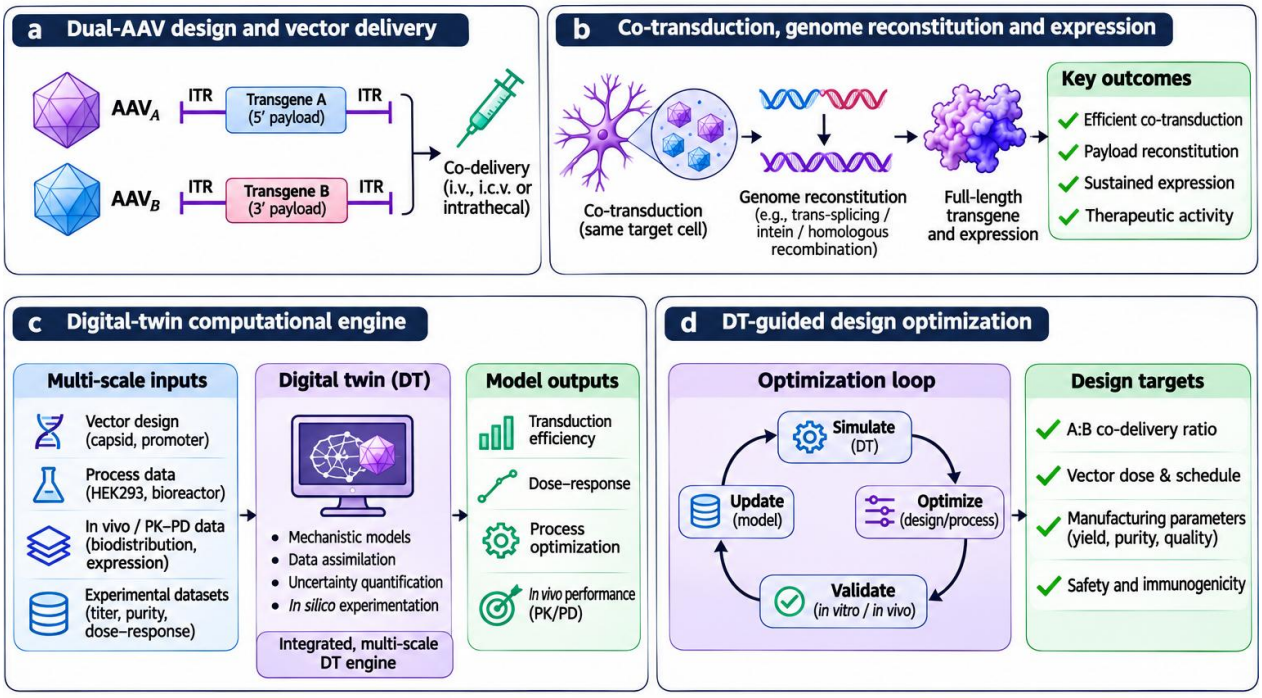


Figure 17: Multiscale architecture of DUAL-AAV PRIME-DT. a, Delivery of AAV-A (N-terminal editor half) and AAV-B (C-terminal reverse transcriptase and pegRNA), co-transduction, split-intein reconstitution and prime editing. b, Editing outcome classes and the desired-edit response chain from molecular rescue to network function. c, Computational engine and evidence status of each layer. Conceptual architecture; SCN1A and Huntington disease are proposed transferability tests and no disease-specific experimental validation is implied.

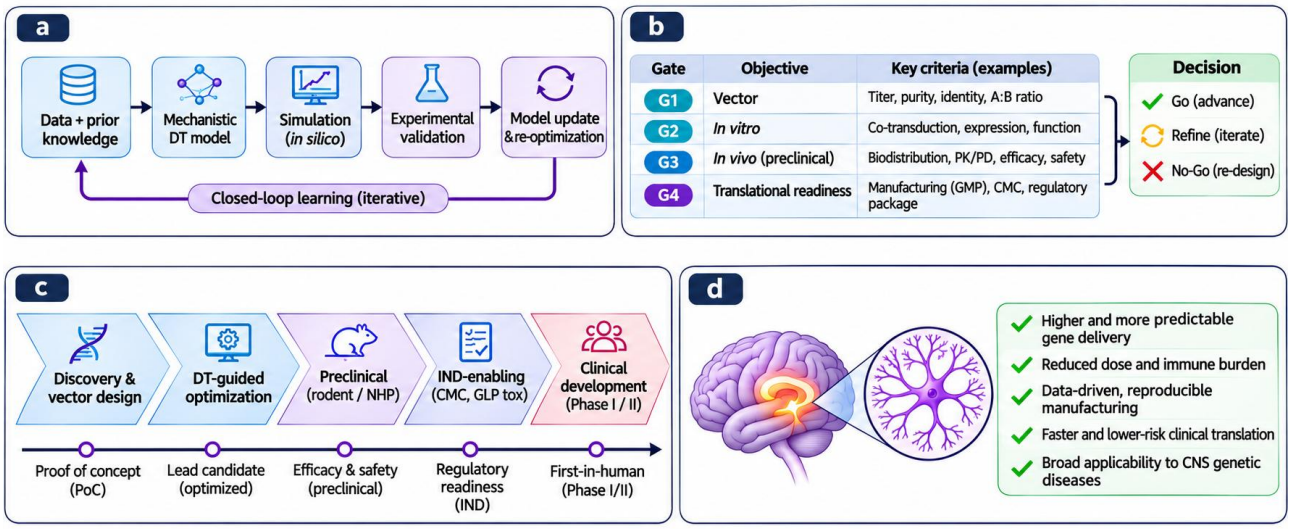


Figure 18*: From model construction to evidence-gated experimental translation. a, Closed-loop workflow from observation to evidence assimilation; the model is updated only when new evidence is available. b, Validation gates, each requiring its own measurements. c, Proposed translational pathway with go/no-go decisions; none of these stages has been completed. Decision-support framework only; model outputs are not evidence of therapeutic efficacy.*

4. Discussion

DUAL-AAV PRIME-DT recasts dual-AAV prime editing as a reliability system: because success requires every stage in series, improving a late stage cannot compensate indefinitely for a low probability of reaching it. In the present synthetic application, joint delivery removes the largest share of probability, followed by editing, and the model makes the cost of ignoring co-transduction explicit a cell that receives AAV-A but not AAV-B is a different biological object from one that receives both vectors but fails to reconstitute. This reframes the question for gene-therapy programmes from 'what percentage can be edited?' to 'what percentage of relevant cells can progress through the whole chain, with what uncertainty, and which experiment most efficiently changes the answer?'

Two results deserve emphasis. First, treating the vectors as independent changed predicted coverage by only about 5%, whereas perfect co-transduction increased it by about 74%. The relevant question is therefore not whether each vector transduces tissue but whether both do so in the same cells, consistent with single-cell and spatial studies (Duarte et al., 2023; Maturana et al., 2022). Second, a terminal endpoint cannot say why a therapy failed. This argues for matched hierarchical measurements and gives the value-of-information analysis a practical consequence: same-cell vector, reconstitution and editing data should be collected before a phenotype experiment is scaled.

The framework transports mechanisms, not percentages. Benchmarks differ in vector configuration, locus, cell composition, species, route and assay, and CNS translation remains constrained by vector biology: capsid performance and tropism are not conserved across species, and the route of administration shifts exposure between brain and spinal cord or dorsal root ganglia (Bunuales et al., 2024; Goertsen et al., 2022; Okai et al., 2025; Rico et al., 2026). LRRK2 biology supplies an increasingly observable patient-state layer kinase activity, Rab phosphorylation, and pharmacodynamic responses to LRRK2-targeting interventions in blood and cerebrospinal fluid (Di Maio et al., 2018; Mabrouk et al., 2026; Yuan et al., 2024) that a future twin could ingest, although these studies do not validate LRRK2 editing. Earlier AAV programmes in Parkinson disease showed that local delivery is feasible but that targeting, endpoint sensitivity and open-label bias limit interpretation (Grote et al., 2024; Winston et al., 2025), and current studies are increasingly multimodal (Siderowf et al., 2023). A cell-level cohort spanning three diseases showed that co-transduction varies across regions and cell classes and that the dominant stage depends on the prior set, which is a further reason to measure rather than assume.

The study has important limitations. All LRRK2 outputs are synthetic; no study has shown dual-AAV prime editing of a pathogenic LRRK2 allele in the human CNS, and the molecular and functional rescue probabilities are inference layers. In the cohort, AAV-B exposure is a fixed multiple of AAV-A exposure, so their separate effects cannot be resolved. Sobol indices reflect unequal assumed ranges, and all performance metrics are evaluated on outcomes from the same model family. Immune neutralisation, manufacturing variability, dose-toxicity relationships,

persistence and rare genomic outcomes are not fully parameterised (Hinderer et al., 2018; Schulz et al., 2023). The central test is prospective: matched AAV-A, AAV-B, co-transduction, reconstitution and editing measurements in the same cellular or spatial units should replace synthetic priors and show whether the predicted bottleneck hierarchy is biologically correct. We offer the twin as an experimental-design and decision-support tool, not as evidence of clinical efficacy.

5. Conclusions

Dual-AAV prime editing behaves as a reliability system in which co-transduction is a distinct causal state rather than a descriptive assay. In a synthetic LRRK2 cohort, joint delivery and editing dominate predicted loss, perfect co-transduction offers an upper-bound gain of about 74%, and no terminal endpoint can identify the failing step. The digital twin is therefore best used as an experimental-design and decision-support tool: matched same-cell measurements of vector occupancy, reconstitution and editing should replace synthetic priors and test whether the predicted bottleneck hierarchy is correct.

Acknowledgments

The author would like to thank Mr Aamish Ahmad for his continued support and guidance throughout this work.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of Interest

The author declares no conflicts of interest.

Ethics Statement

This study was computational in nature and did not involve human participants, animal experiments, or identifiable clinical data. Therefore, ethical approval was not required.

Data Availability Statement

The synthetic cohort tables, parameter dictionaries, available provenance records, and computational analysis outputs associated with this study are provided in the accompanying repository (repository URL/DOI: to be inserted following deposition). Analysis scripts and figure-generation code are included where available. The repository documents the computational workflow, data provenance, fixed-seed configurations, and current reproducibility limitations in docs/REPRODUCIBILITY.md. The deposited materials distinguish reproducible outputs from legacy or unverified analyses. Full regeneration of every manuscript figure and table should not be assumed unless explicitly confirmed by the corresponding version-matched source files and execution records.

Author Affiliation and Disclosure

The author's updated institutional affiliation will be provided in the final publication record following confirmation of the applicable institutional details. Any affiliation changes arising during the manuscript's review and publication process will be communicated to the journal in accordance with its publication policies.

References

Anzalone, A. V., et al. (2019). Search-and-replace genome editing without double-strand breaks or donor DNA. *Nature*, *576*, 149–157. https://doi.org/10.1038/s41586-019-1711-4

Bolon, B., et al. (2024). Neuropathological findings in nonclinical species following administration of adeno-associated virus-based gene therapy vectors. *Toxicologic Pathology*, *52*, 489–505. https://doi.org/10.1177/01926233241300314

Brooker, S. M., & Gonzalez-Latapi, P. (2025). Biomarkers in Parkinson's disease. *Neurologic Clinics*, *43*, 229–248. https://doi.org/10.1016/j.ncl.2024.12.005

Bunuales, M., et al. (2024). Characterization of brain transduction capability of a BBB-penetrant AAV vector in mice, rats and macaques reveals differences in expression profiles. *Gene Therapy*, *31*, 455–466. https://doi.org/10.1038/s41434-024-00466-w

Chen, P. J., et al. (2021). Enhanced prime editing systems by manipulating cellular determinants of editing outcomes. *Cell*, *184*, 5635–5652.e29. https://doi.org/10.1016/j.cell.2021.09.018

Chen, X., et al. (2023). Functional gene delivery to and across brain vasculature of systemic AAVs with endothelial-specific tropism in rodents and broad tropism in primates. *Nature Communications*, *14*, 3345. https://doi.org/10.1038/s41467-023-38582-7

Colella, P., et al. (2014). Efficient gene delivery to the cone-enriched pig retina by dual AAV vectors. *Gene Therapy*, *21*, 450–456. https://doi.org/10.1038/gt.2014.8

Cortés, A., et al. (2025). In-depth mass-spectrometry reveals phospho-RAB12 as a blood biomarker of G2019S LRRK2-driven Parkinson's disease. *Brain*, *148*, 2075–2092. https://doi.org/10.1093/brain/awae404

Dacquay, L. C., et al. (2025). Dual inhibition of DNA-PK and PolTheta boosts precision of diverse prime editing systems. *Nature Communications*, *16*, 4290. https://doi.org/10.1038/s41467-025-59708-z

Davis, J. R., et al. (2024). Efficient prime editing in mouse brain, liver and heart with dual AAVs. *Nature Biotechnology*, *42*, 253–264. https://doi.org/10.1038/s41587-023-01758-z

Di Maio, R., et al. (2018). LRRK2 activation in idiopathic Parkinson's disease. *Science Translational Medicine*, *10*, eaar5429. https://doi.org/10.1126/scitranslmed.aar5429

Doman, J. L., et al. (2023). Phage-assisted evolution and protein engineering yield compact, efficient prime editors. *Cell*, *186*, 3983–4002.e26. https://doi.org/10.1016/j.cell.2023.07.039

Drouyer, M., et al. (2024). Enhanced AAV transduction across preclinical CNS models: a comparative study in human brain organoids with cross-species evaluations. *Molecular Therapy - Nucleic Acids*, *35*, 102264. https://doi.org/10.1016/j.omtn.2024.102264

Duarte, F., et al. (2023). Semi-automated workflows to quantify AAV transduction in various brain areas and predict gene editing outcome for neurological disorders. *Molecular*

Therapy - Methods & Clinical Development, *29*, 254–270. https://doi.org/10.1016/j.omtm.2023.03.013

Foust, K. D., et al. (2009). Intravascular AAV9 preferentially targets neonatal neurons and adult astrocytes. *Nature Biotechnology*, *27*, 59–65. https://doi.org/10.1038/nbt.1515

Goertsen, D., et al. (2022). AAV capsid variants with brain-wide transgene expression and decreased liver targeting after intravenous delivery in mouse and marmoset. *Nature Neuroscience*, *25*, 106–115. https://doi.org/10.1038/s41593-021-00969-4

Gori, J. L., et al. (2026). Prime Editing for p47phox-Deficient Chronic Granulomatous Disease. *The New England Journal of Medicine*, *394*, 1195–1203. https://doi.org/10.1056/NEJMoa2509807

Grote, J., et al. (2024). From lab bench to hope: a review of gene therapies in clinical trials for Parkinson's disease and challenges. *Neurological Sciences*, *45*, 4699–4710. https://doi.org/10.1007/s10072-024-07599-1

Hinderer, C., et al. (2018). Severe toxicity in nonhuman primates and piglets following high-dose intravenous administration of an adeno-associated virus vector expressing human SMN. *Human Gene Therapy*, *29*, 285–298. https://doi.org/10.1089/hum.2018.015

Hocquemiller, M., et al. (2022). AAVrh10 vector corrects pathology in animal models of GM1 gangliosidosis and achieves widespread distribution in the CNS of nonhuman primates. *Molecular Therapy - Methods & Clinical Development*, *27*, 281–292. https://doi.org/10.1016/j.omtm.2022.10.004

Hop, P. J., et al. (2024). Systematic rare variant analyses identify RAB32 as a susceptibility gene for familial Parkinson's disease. *Nature Genetics*, *56*, 1371–1376. https://doi.org/10.1038/s41588-024-01787-7

Jennings, D., et al. (2022). Preclinical and clinical evaluation of the LRRK2 inhibitor DNL201 for Parkinson's disease. *Science Translational Medicine*, *14*, eabj2658. https://doi.org/10.1126/scitranslmed.abj2658

Jennings, D., et al. (2023). LRRK2 inhibition by BIIB122 in healthy participants and patients with Parkinson's disease. *Movement Disorders*, *38*, 386–398. https://doi.org/10.1002/mds.29297

Kalogeropoulou, A. F., et al. (2022). Impact of 100 LRRK2 variants linked to Parkinson's disease on kinase activity and microtubule binding. *Biochemical Journal*, *479*, 1759–1783. https://doi.org/10.1042/BCJ20220161

Kania, E., et al. (2023). LRRK2 phosphorylation status and kinase activity regulate macroautophagy in a Rab8a/Rab10-dependent manner. *Cell Death & Disease*, *14*, 436. https://doi.org/10.1038/s41419-023-05964-0

Kirik, D., Cederfjäll, E., Halliday, G., & Petersén, Å. (2017). Gene therapy for Parkinson's disease: disease modification by GDNF family of ligands. *Neurobiology of Disease*, *97*, 179–188. https://doi.org/10.1016/j.nbd.2016.09.008

Kissling, L., et al. (2026). Prime editing of a pathogenic Scn1a allele ameliorates seizure phenotypes in a GEFS+ mouse model. *Science Translational Medicine*, *18*, eadz2557. https://doi.org/10.1126/scitranslmed.adz2557

Kluss, J. H., et al. (2022). Lysosomal positioning regulates Rab10 phosphorylation at LRRK2+ lysosomes. *Proceedings of the National Academy of Sciences*, *119*, e2205492119. https://doi.org/10.1073/pnas.2205492119

Koeppel, J., et al. (2023). Prediction of prime editing insertion efficiencies using sequence features and DNA repair determinants. *Nature Biotechnology*, *41*, 1446–1456. https://doi.org/10.1038/s41587-023-01678-y

Leborgne, C., et al. (2020). IgG-cleaving endopeptidase enables in vivo gene therapy in the presence of anti-AAV neutralizing antibodies. *Nature Medicine*, *26*, 1096–1101. https://doi.org/10.1038/s41591-020-0911-7

Levy, J. M., et al. (2020). Cytosine and adenine base editing of the brain, liver, retina, heart and skeletal muscle of mice via adeno-associated viruses. *Nature Biomedical Engineering*, *4*, 97–110. https://doi.org/10.1038/s41551-019-0501-5

Mabrouk, O. S., et al. (2026). LRRK2-targeting antisense oligonucleotide in Parkinson's disease: a phase 1 randomized controlled trial. *Nature Medicine*, *32*, 1421–1431. https://doi.org/10.1038/s41591-026-04262-4

Mathis, N., et al. (2023). Predicting prime editing efficiency and product purity by deep learning. *Nature Biotechnology*, *41*, 1151–1159. https://doi.org/10.1038/s41587-022-01613-7

Maturana, C. J., Verpeut, J. L., & Engel, E. A. (2022). Single-cell quantification of triple-AAV vector genomes coexpressed in neurons. *Current Protocols*, *2*, e430. https://doi.org/10.1002/cpz1.430

Mendell, J. R., et al. (2022). Testing preexisting antibodies prior to AAV gene transfer therapy: rationale, lessons and future considerations. *Molecular Therapy - Methods & Clinical Development*, *25*, 74–83. https://doi.org/10.1016/j.omtm.2022.02.011

Mittas, D. M., et al. (2026). Dual AAV vectors for efficient delivery of large transgenes. *Nature Protocols*, *21*, 1466–1522. https://doi.org/10.1038/s41596-025-01243-8

Moyer, T. C., et al. (2025). Highly conserved brain vascular receptor ALPL mediates transport of engineered AAV vectors across the blood-brain barrier. *Molecular Therapy*, *33*, 3902–3916. https://doi.org/10.1016/j.ymthe.2025.04.046

Nelson, J. W., et al. (2022). Engineered pegRNAs improve prime editing efficiency. *Nature Biotechnology*, *40*, 402–410. https://doi.org/10.1038/s41587-021-01039-7

Okai, T., et al. (2025). Biodistribution of AAV1, AAV5, AAV9, and AAVDJ serotypes after intra-cisterna magna delivery in non-human primates. *Molecular Therapy - Methods & Clinical Development*, *33*, 101593. https://doi.org/10.1016/j.omtm.2025.101593

Rapti, K., & Grimm, D. (2021). Adeno-associated viruses and host immunity - a race between the hare and the hedgehog. *Frontiers in Immunology*, *12*, 753467. https://doi.org/10.3389/fimmu.2021.753467

Rico, A. J., et al. (2026). Biodistribution of AAV-TT and AAV9 in the nonhuman primate brain. *Human Gene Therapy*, *37*, 744–755. https://doi.org/10.1177/10430342261427733

Riedmayr, L. M., et al. (2023). mRNA trans-splicing dual AAV vectors for (epi)genome editing and gene therapy. *Nature Communications*, *14*, 6578. https://doi.org/10.1038/s41467-023-42386-0

Sakai, H. A., et al. (2026). Directed evolution of small RNA-stabilizing motifs that improve prime-editing efficiency. *Nature Biotechnology*. https://doi.org/10.1038/s41587-026-03123-2

Schulz, M., et al. (2023). Binding and neutralizing anti-AAV antibodies: detection and implications for rAAV-mediated gene therapy. *Molecular Therapy*, *31*, 616–630. https://doi.org/10.1016/j.ymthe.2023.01.010

She, K., et al. (2023). Dual-AAV split prime editor corrects the mutation and phenotype in mice with inherited retinal degeneration. *Signal Transduction and Targeted Therapy*, *8*, 57. https://doi.org/10.1038/s41392-022-01234-1

Siderowf, A., et al. (2023). Assessment of heterogeneity among participants in the Parkinson's Progression Markers Initiative cohort using alpha-synuclein seed amplification. *The Lancet Neurology*, *22*, 407–417. https://doi.org/10.1016/S1474-4422(23)00109-6

Simuni, T., et al. (2024). A biological definition of neuronal alpha-synuclein disease: towards an integrated staging system for research. *The Lancet Neurology*, *23*, 178–190. https://doi.org/10.1016/S1474-4422(23)00405-2

Steger, M., et al. (2016). Phosphoproteomics reveals that Parkinson's disease kinase LRRK2 regulates a subset of Rab GTPases. *eLife*, *5*, e12813. https://doi.org/10.7554/eLife.12813

Streubel-Gallasch, L., et al. (2021). Parkinson's disease-associated LRRK2 interferes with astrocyte-mediated alpha-synuclein clearance. *Molecular Neurobiology*, *58*, 3119–3140. https://doi.org/10.1007/s12035-021-02327-8

Tang, H., et al. (2023). Hearing of Otof-deficient mice restored by trans-splicing of N- and C-terminal otoferlin. *Human Genetics*, *142*, 289–304. https://doi.org/10.1007/s00439-022-02504-2

Tasfaout, H., et al. (2024). Split intein-mediated protein trans-splicing to express large dystrophins. *Nature*, *632*, 192–200. https://doi.org/10.1038/s41586-024-07710-8

Trapani, I., et al. (2014). Effective delivery of large genes to the retina by dual AAV vectors. *EMBO Molecular Medicine*, *6*, 194–211. https://doi.org/10.1002/emmm.201302948

Verdera, H. C., Kuranda, K., & Mingozzi, F. (2020). AAV vector immunogenicity in humans: a long journey to successful gene transfer. *Molecular Therapy*, *28*, 723–746. https://doi.org/10.1016/j.ymthe.2019.12.010

Vissers, M. F. J. M., et al. (2023). A LRRK2 pathway biomarker characterization study in patients with Parkinson's disease with and without LRRK2 mutations and healthy controls. *Clinical and Translational Science*, *16*, 1408–1420. https://doi.org/10.1111/cts.13541

Wauters, F., et al. (2020). LRRK2 mutations impair depolarization-induced mitophagy through inhibition of mitochondrial accumulation of RAB10. *Autophagy*, *16*, 203–222. https://doi.org/10.1080/15548627.2019.1603548

Weber, T. (2021). Anti-AAV antibodies in AAV gene therapy: current challenges and possible solutions. *Frontiers in Immunology*, *12*, 658399. https://doi.org/10.3389/fimmu.2021.658399

Winston, G., et al. (2025). Gene therapy for Parkinson's disease: trials and technical advances. *The Lancet Neurology*, *24*, 548–556. https://doi.org/10.1016/S1474-4422(25)00125-5

Yan, J., et al. (2024). Improving prime editing with an endogenous small RNA-binding protein. *Nature*, *628*, 639–647. https://doi.org/10.1038/s41586-024-07259-6

Yu, G., et al. (2023). Prediction of efficiencies for diverse prime editing systems in multiple cell types. *Cell*, *186*, 2256–2272.e23. https://doi.org/10.1016/j.cell.2023.03.034

Yuan, Y., et al. (2024). Single molecule array measures of LRRK2 kinase activity in serum link Parkinson's disease severity to peripheral inflammation. *Molecular Neurodegeneration*, *19*, 47. https://doi.org/10.1186/s13024-024-00738-4