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.

Identifiability and Out-of-Sample Limits of a Digital Shadow for Harvest Timing in Automated T-Cell Perfusion Bioreactors

FAIZAN AHMAD

faizmedi16@gmail.com

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

Research Article

Keywords: digital shadow, T-cell manufacturing, perfusion bioreactor, practical identifiability, harvest timing, process analytical technology

Posted Date: November 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11874858/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.

Abstract

Digital shadows may support harvest timing in automated T-cell manufacturing, but published models are rarely stress-tested for identifiability, out-of-sample behaviour or process consequences. We re-analysed published run-level data and parameters of a digital shadow of T-cell expansion in a glucose-controlled perfusion bioreactor (six runs), re-implemented its controller–metabolism equations, and used published smart-sensor counts. Endpoint agreement with run-specific parameters (0.57% mean error) is an internal-consistency result, because the maximum growth rate is computed from the final concentration; pooled parameters erred by 69% (leave-one-run-out). Under set-point control, growth follows from the perfusion rate through the glucose yield, so online data identify $\mu_{max}/Y_{X/S}$, $Y_{X/S}{\cdot}Y_{P/X}$ and v but not the yields separately. In pseudo-real-time tests, fixing $Y_{X/S}$ at its pooled mean reproduced the published 13% error and tracked its run-to-run mismatch (r = 0.79); one offline count reduced it to about 10%. The model implies 405–731 mL medium per 10^9 cells and 0.8–6.4 d to 0.5–2×10^9 cells across runs. An evidence gate and reverse-feasibility map govern transfer to other T-cell sources but prospective validation is still required.

Highlights

Endpoint agreement of a T-cell digital shadow is a consistency check, not validation

Pooled parameters mispredict new runs: leave-one-run-out error is 69%

Online data identify parameter combinations, not the glucose yield

Fixed glucose yield limits real-time forecasts to ~ 12% error in best-case tests

Model implies 1.8-fold media and 2.5-fold time-to-dose spread across runs

1. Introduction

Harvest timing in autologous T-cell manufacturing balances cell yield, phenotype and process cost, and manufacturing conditions shape the differentiation state of the product (Song et al., 2024). Automated perfusion platforms, including stirred-tank systems, now provide continuous online measurements and tight control over feeding (Amini et al., 2020; Costariol et al., 2019, 2020; Hood et al., 2024; Gatla et al., 2022; Springuel et al., 2025), which creates the conditions for model-based decision support (Narayanan et al., 2020; Hewitt et al., 2021). A digital shadow, a model fed one way by process data without actuating the process, is the lowest-risk first step towards that goal.

Egan et al. (2025) recently reported such a digital shadow for T-cell and CAR T-cell expansion in an automated perfusion bioreactor (Aglaris FACER), in which a controller adjusts the perfusion rate to hold glucose at a set-point. Growth is described by a Richards-type equation (Richards, 1959) with a lag, a maximum specific growth rate and a shape parameter, coupled to glucose and lactate balances.

Parameters were fitted to five of six runs. In pseudo-real-time tests, with the lag and the glucose yield fixed at pooled means because the model was otherwise not practically identifiable, the final cell concentration up to 2.5 d ahead was predicted with a mean relative error of 13% (standard deviation about 6%) once at least 2.5 d of online data were available. A companion study defined a smart-sensor framework whose consensus alerts flagged unstable behaviour in the same experimental campaign (Nettleton et al., 2023).

Four questions limit how far such a shadow can be trusted or transferred, and each can be examined with published information alone. First, how much of the agreement between model and data is independent evidence? When a parameter is computed from the quantity being predicted, endpoint agreement is a consistency check, not validation. Second, how well do parameters pooled across historical runs predict a new run before run-specific data exist? Third, which parameters can the online data actually identify, and what limits real-time forecast accuracy (Wanika et al., 2024)? Fourth, what do the fitted kinetics imply for manufacturing, namely media demand, time to a cell dose and the effect of design choices such as the glucose set-point? A fifth, transfer question follows: what must be shown before parameters from adult peripheral-blood αβ T cells are used for cord-blood, γδ or iPSC-derived T cells, for which process-level kinetic data are scarce (Hippen et al., 2008; Sutton et al., 2016; Lo Presti et al., 2018; Nianias & Themeli, 2019)?

Here we re-analyse the published run-level data and parameters of Egan et al. (2025) and the consensus counts of Nettleton et al. (2023), and re-implement the controller–metabolism equations of the shadow. We (i) separate in-sample consistency from out-of-sample prediction; (ii) quantify harvest-time sensitivity and the information in summary observables; (iii) analyse practical identifiability and pseudo-real-time forecasting; (iv) derive media-demand and time-to-dose metrics and their sensitivity to the glucose set-point; (v) propagate parameter ranges as scenarios and examine smart-sensor consensus burden; and (vi) formalise an evidence gate with a reverse-feasibility map for source transfer. We did not have the original 10-s trajectories, so we make no claim to refit or independently validate the published model. Our aim is to define what a digital shadow can and cannot claim from published information and which additional data would change that.

2. Materials and methods

2.1. Data sources and provenance

All inputs were transcribed from published tables: run-level inputs and controller settings (Table S1) and fitted parameters (Table S2) from Egan et al. (2025), and consensus-event counts from Nettleton et al. (2023). Egan et al. report six expansions in a 134 mL chamber: four donor runs and one patient-derived CAR T-cell run (Run 4) in medium A, and one donor run in medium B (Run 6). Fitted parameters are published for five runs. Run 3 had a mechanical interruption of about 0.5 d without online data and was used by the originators as an independent test run, not for fitting. Nettleton et al. report consensus counts for five batches. We assumed that their batch labels correspond to Runs 1–5 of Egan et al.

because the descriptors match (a mechanical shutdown in Run 3, patient-derived cells in Run 4), but the correspondence is not stated explicitly and should be confirmed with the originators. Raw trajectories are not public and were not used.

2.2. Growth model

We use the convention of Egan et al. (2025), in which the shape parameter $v \leq 0$. With time measured from seeding ($t_I = 0$), viable-cell concentration X after a lag t_{lag} is

$$X(t) = X_I \, [1 - \mu_{max} \, v \, (t - t_{lag})]^{-1/v} \quad (1)$$

with the exponential limit $X_I \exp[\mu_{max}(t - t_{lag})]$ as $v \to 0$. The specific growth rate is

$$\mu(t) = \mu_{max} / [1 - \mu_{max} \, v \, (t - t_{lag})] \quad (2)$$

Egan et al. do not fit μ_{max} freely in training; it is computed from the final concentration X_F reached at time t_F:

$$\mu_{max} = [1 - (X_I / X_F)^v] / [v \, (t_F - t_{lag})] \quad (3)$$

The time at which a target concentration X_T is reached is

$$t_T = t_{lag} + [1 - (X_I / X_T)^v] / (\mu_{max} \, v) \quad (4)$$

We use $X_T = 10 \times 10^6$ cells mL^{-1}, as in Egan et al., unless stated otherwise.

2.3. Controller-metabolism model

Following Egan et al., a proportional-integral controller adjusts the perfusion rate D (vessel volumes per day, VVD) to hold glucose S at the set-point S_{sp}, while glucose and lactate P obey mass balances with yields $Y_{X/S}$ (cells $pmol^{-1}$) and $Y_{P/X}$ (pmol $cell^{-1}$):

$$dD/dt = H(t - t_{sed}) \, [-K_p \, dS/dt + K_i(S_{sp} - S)] \quad (5)$$

$$dS/dt = H(t - t_{sed}) \, D \, (S_I - S) - \mu X / Y_{X/S} \quad (6)$$

$$dP/dt = Y_{P/X} \, \mu X - H(t - t_{sed}) \, D \, P \quad (7)$$

where H is the Heaviside step switching perfusion on after sedimentation (t_{sed}), S_I is the glucose concentration of fresh medium, and K_p and K_i are controller gains (derivative action omitted, as in Egan et al.). With X in 10^6 cells mL^{-1}, the glucose uptake rate is $\mu X/Y_{X/S}$ in mmol L^{-1} d^{-1} and the lactate release rate is $Y_{P/X}\mu X$, so no further unit conversion is needed. If glucose is held at its set-point ($dS/dt \approx 0$), the perfusion rate is

$D_{sp} = \mu X / [Y_{X/S} (S_I - S_{sp})]$ (8)

and therefore $dX/dt = Y_{X/S}(S_I - S_{sp})D$: cell growth is a glucose mass balance scaled by $Y_{X/S}$. Equations 5–7 were integrated by forward Euler (step 0.0005 d for single runs, 0.005 d in vectorised batches), with D floored at the published minimum perfusion rate $D_{min} = 0.067\ d^{-1}$ after sedimentation (our assumption; Egan et al. force non-negativity). Controller gains, S_I, S_{sp}, t_{sed} and vessel volume (134 mL) were taken from Egan et al. Table 1. The re-implementation was checked against the set-point relation (Table S12) and against the published final cell numbers (Table S11) but, without raw trajectories, not against measured profiles.

2.4. Consistency and out-of-sample tests

(i) **Run-specific consistency**: X_F was recomputed from Eq. 1 using the published (rounded) parameters of each fitted run. (ii) **Circularity check**: μ_{max} was recomputed from the reported X_F with Eq. 3. (iii) **Pooled prediction**: Run 3 was predicted from the arithmetic mean of the five fitted parameter sets (t_{lag}, v, μ_{max}). (iv) **Leave-one-run-out (LOO)**: each fitted run was predicted from the mean of the other four. Pooled predictions are an open-loop, prior-only shadow without run-specific online data and are not a replication of the real-time updating of Egan et al. Errors are signed relative errors; MAE is the mean absolute relative error.

2.5. Sensitivity and summary-observable information

Sensitivity of t_T to ± 20% changes in μ_{max}, v and t_{lag} was summarised per run as the secant elasticity $E = (\Delta t_{+20\%} - \Delta t_{-20\%})/(0.4\ t_T)$. For the observable vector $[X_F, t_T]$, the log-scaled Jacobian with respect to (t_{lag}, v, μ_{max}, $Y_{X/S}$, $Y_{P/X}$) was evaluated for Run 1 by central differences and decomposed by singular-value decomposition. $X_T = 19\times10^6$ cells mL^{-1} (equal to the Run 1 X_F) was used as a contrast case.

2.6. Practical identifiability and pseudo-real-time forecasting

Fisher information

For each fitted run, the full ODE was used to compute sensitivities of log D (multiplicative noise 5%) and lactate (additive noise 0.5 mmol L^{-1}) observed from 0.3 d after perfusion starts for windows of 0.5–3.5 d, and Cramér–Rao coefficients of variation (CV) were obtained for (μ_{max}, v, $Y_{X/S}$, $Y_{P/X}$), for the combinations $\mu_{max}/Y_{X/S}$ and $Y_{X/S}{\cdot}Y_{P/X}$ (log-parameter covariance), and for the real-time set-up of Egan et al. in which $Y_{X/S}$ and t_{lag} are fixed. Sensor noise is not published and was assumed.

Forecasting: Pseudo-data (log D, lactate, glucose) were generated from the published fitted ODE of each run, with run-specific noise replicates (16 per run, window and noise level; noise on log D 5% or 15%). Three estimators predicted X_F with the run excluded from every pooled value: the oracle ($Y_{X/S}$ and t_{lag} fixed at the run's true values); the Egan-like estimator (both fixed at the mean of the other runs); and a

mass-balance estimator in which $Y_{X/S}$ is calibrated from one offline count X_c (10% CV) at the end of the window,

$$Y_{X/S} = (X_c - X_I) / \{(S_I - S_c) + \int D\,(S_I - S)\,dt\} \quad (9)$$

with t_{lag} pooled. In all cases (μ_{max}, v, $Y_{P/X}$) were estimated by least squares on log D and lactate (grid search; $Y_{P/X}$ in closed form). Because the pseudo-data come from the model itself, errors are best-case lower bounds.

2.7. Biomanufacturing metrics

Cumulative medium use (vessel volumes) is the integral of D from the end of sedimentation to t_F in the full ODE; media per 10^9 cells is 1,000 × cumulative medium/X_F (mL; vessel volume cancels). Glucose and lactate per 10^9 cells follow from $1/Y_{X/S}$ and $Y_{P/X}$ (180.16 and 90.08 g mol^{-1}). Time to N cells uses Eq. 4 with X_T = N/V. The design sensitivity of media demand to S_{sp} used Eq. 8 with each run's kinetics held fixed; the model has no glucose-limitation term, so this is a demand calculation, not a prediction of growth at other set-points.

2.8. Scenario propagation and smart-sensor consensus

To illustrate forecast width when parameters are bounded only by historical ranges, t_{lag}, v and μ_{max} were drawn independently and uniformly between the minimum and maximum of the five fitted values (N = 50,000; NumPy default_rng, seed 42) and propagated through Eqs. 1 and 4 using the initial concentration and expansion time of Run 1; tail stability was assessed over 20 seeds. Independent sampling ignores the coupling between v and μ_{max}, so the envelopes are not prediction intervals. Consensus events (three or more of six smart sensors simultaneously in alert) were taken from Nettleton et al. and divided by the final process time. A process-state rule was defined for the framework: Red, documented shutdown; Amber, at least 2 consensus events per day without a documented shutdown; Green, otherwise. The 2 events d^{-1} threshold is a design choice and was not calibrated.

2.9. Evidence gate, reverse feasibility and software

Every model input is assigned one of six evidence classes: E1, direct observation; E2, published fitted parameter; E3, derived quantity; E4, literature prior; E5, modelling assumption; E6, hypothetical scenario. Only E1–E2 may calibrate the model, E4 may inform priors, and E6 may define requirements only (Fig. 1). A targeted, non-systematic evidence set was compiled for six source classes (primary adult PBMC αβ T cells, umbilical-cord-blood T cells, γδ T cells, iPSC-derived T cells, menstrual-blood T cells and breast-milk T cells) and rated qualitatively on three dimensions (Table 3; Table S9). Reverse feasibility asks which μ_{max} a new source must show, for a given v and t_{lag}, to take a reference process (1.49×10^6 cells mL^{-1}, 4.94 d, target 19×10^6 cells mL^{-1}) to its target, by solving Eq. 3 for μ_{max}. Analyses used Python 3 (NumPy, SciPy, pandas, Matplotlib); all scripts, inputs and outputs are available in the repository listed under Data availability.

2.10. Run-level reporting and exploratory analyses

To separate what was measured, selected and excluded, we adopted the layered reporting of a large clinical transcriptomic study (Andreoletti et al., 2023): a characteristics table, an explicit inclusion and quality-control map and several analysis tiers. The analogues here are (i) a run-characteristics table (Table 4) and a data-availability matrix (Fig. 8a); (ii) unsupervised hierarchical clustering of the five fitted runs on ten z-scored features (μ_{max}, v, t_{lag}, $Y_{X/S}$, $Y_{P/X}$, X_I, $S_I - S_{sp}$, peak D, media per 10^9 cells and time to 10^9 cells), using Pearson correlation distance and average linkage, annotated by medium, cell source and seeding time t_I (Fig. 8b); (iii) a Spearman correlation analysis of the same features, ordered by hierarchical clustering of $1 - |\rho|$ and compared with a null in which each feature was permuted independently across runs (20,000 permutations, seed 2026) for the mean $|\rho|$ and the number of pairs with $|\rho| \geq 0.9$ (Fig. 8c); (iv) a comparison of runs reaching 1×10^9 cells within 3 d with the other runs, using the exact permutation distribution; (v) an attribution of the time to 1×10^9 cells (Eq. 4 with $X_T = 10^9$ cells / 134 mL = 7.46×10^6 cells mL^{-1}) to seeding density and kinetics, in which each run's X_I and (t_{lag}, v, μ_{max}) were swapped with a reference (median X_I of the fitted runs and pooled kinetics) and the two contributions were computed as Shapley values, which sum exactly to the difference from the reference (Fig. 8d; Table 5); and (vi) a per-run comparison of the three forecasting estimators of Section 2.6 at 2.5 d with 5% noise on log D (Fig. 8e). With five fitted runs these analyses are descriptive; the smallest attainable two-sided exact permutation P value for a two-versus-three split is 0.2.

Table 1

Final viable-cell concentration: run-specific and pooled-parameter predictions (10^6 cells mL^{-1}).

Run	Source	X_I	X_F observed	Run-specific parameters	Pooled parameters
1	Donor 1	1.49	19.0	18.95 (− 0.25%)	32.6 (+ 71.3%)
2	Donor 2	0.37	19.0	18.78 (− 1.14%)	7.9 (− 58.3%)
3	Donor 3	1.49	32.5	n/a	38.1 (+ 17.3%)
4	Patient 1 (CAR T)	0.62	10.9	10.86 (− 0.36%)	12.2 (+ 12.2%)
5	Donor 4	1.13	8.73	8.67 (− 0.74%)	20.0 (+ 129.2%)
6	Donor 5	1.49	18.2	18.27 (+ 0.38%)	31.4 (+ 72.7%)
MAE				0.57%	68.7% (LOO, n = 5)

Run-specific: published rounded parameters (Eq. 1). Pooled: mean of the five fitted runs for Run 3; mean of the other four fitted runs for each fitted run (leave-one-run-out). X in 10^6 cells mL^{-1}; percentages are signed relative errors. Run 3 has no published fitted parameters.

3. Results

3.1. Endpoint agreement with run-specific parameters is a consistency check, not validation

Using the published, rounded parameters, Eq. 1 reproduced the reported final concentration of the five fitted runs within 1.14% (MAE 0.57%, median 0.38%; Table 1; Fig. 2a,b). Times to target (X_T = 10×10^6 cells mL^{-1}) matched published values within 0.005 d (Fig. 2c). Both agreements are expected by construction: μ_{max} is computed from X_F through Eq. 3, and recomputing it that way differed from the published value by − 0.65% to + 0.34% (Table S2), the size of the rounding in three-significant-figure parameters. These results confirm that the equations and published values were interpreted correctly; they are not evidence of predictive validity.

3.2. Pooled parameters fail out of sample

Predicting Run 3 from pooled mean parameters (t_{lag} 0.092 d, v − 0.537, μ_{max} 1.51 d^{-1}) overshot X_F by 17.3% (38.1 versus 32.5×10^6 cells mL^{-1}). LOO predictions for the five fitted runs had errors of + 71.3%, − 58.3%, + 12.2%, + 129.2% and + 72.7% (MAE 68.7%, RMSE 78.2%; Fig. 2a,b). The pooled error for Run 3 was smaller than four of the five LOO errors, so the documented interruption did not make Run 3 exceptional for a pooled-parameter prior and its error cannot be attributed to the disturbance. Parameter heterogeneity was large (sample CVs: μ_{max} 46%, v 65%, t_{lag} 93%; Table S2) and μ_{max} and |v| were positively rank-correlated (Spearman ρ = 0.7, n = 5, P = 0.19; Fig. 3a,b), consistent with the coupling imposed by Eq. 3. Pooled parameters are therefore a prior-only baseline: at this sample size a shadow run without run-specific updating is uninformative.

3.3. Harvest time is set by growth parameters

For Run 1, ± 20% in μ_{max} shifted t_T by + 0.61 and − 0.40 d, ± 20% in v by − 0.48 and + 0.62 d and ± 20% in t_{lag} by ∓ 0.03 d, corresponding to secant elasticities of − 0.98, + 1.07 and + 0.06; other runs gave similar values (μ_{max} − 0.98 to − 1.04; v + 0.88 to + 1.07; Fig. 3c; Table S5). Because Eq. 1 depends only on t_{lag}, v and μ_{max}, $Y_{X/S}$ and $Y_{P/X}$ have zero influence on X_F and t_T: endpoint agreement cannot validate the metabolic layer. The Jacobian of [X_F, t_T] with respect to the five parameters had two non-zero singular values (2.32 and 0.19 at X_T = 10×10^6 cells mL^{-1}; condition number 12) and therefore rank 2, so two scalar summaries cannot constrain three growth parameters. When X_T equalled X_F, the two observables became nearly collinear (singular values 2.59 and 7.0×10^{-4}; condition number 3.7×10^3; Fig. 3d). This follows from counting and is not a defect of the model; continuous online data carry far more information, which we quantify next.

3.4. Data identify parameter combinations

Under set-point control, $X(t) - X_I$ equals $Y_{X/S}(S_I - S_{sp})$ times the integral of D, so the online perfusion record fixes the shape of cell growth but its absolute scale only through $Y_{X/S}$. Fisher-information analysis showed this directly (Fig. 4a; Table S13). For Run 1, with 2.5 d of data, the CVs of μ_{max}, $Y_{X/S}$ and $Y_{P/X}$ were 183%, 185% and 185% (condition number 4.3×10^8), whereas $\mu_{max}/Y_{X/S}$, $Y_{X/S}\cdot Y_{P/X}$ and v were determined to 2.1%, 0.8% and 6.5%. The degeneracy was strongest for Run 1 ($v \approx -0.92$, nearly constant perfusion rate) and weaker for the other runs (μ_{max} CV 11–59% at 2.5 d; $\mu_{max}/Y_{X/S}$ 1.7–3.7%; $Y_{X/S}\cdot Y_{P/X}$ 0.6–0.8%). With $Y_{X/S}$ fixed, as in the real-time set-up of Egan et al., the remaining parameters were tightly determined (CV at most 1.8% at 2.5 d in all runs). The cost of this choice is a forecast whose scale inherits any error in the fixed yield.

3.5. Real-time forecast error is set by the fixed glucose yield

In pseudo-real-time tests with 5% noise on log D, the Egan-like estimator had a median error in X_F of 33.6%, 20.6%, 12.6%, 12.3%, 9.9% and 9.3% with 1.0, 1.5, 2.0, 2.5, 3.0 and 3.5 d of online data (Fig. 4b; Table S14). The 2.5 d value of 12.3% agrees with the 13% (standard deviation about 6%) reported by Egan et al. for the measured data, suggesting that the published accuracy is close to what a pooled $Y_{X/S}$ allows. The oracle estimator, with the run's own $Y_{X/S}$ and t_{lag}, reached 7.5% at 2.5 d and 3.6% at 3.5 d; the Egan-like error fell little with more data, indicating a floor set by the fixed parameters rather than by noise. The signed error of the Egan-like estimator at 2.5 d tracked the mismatch between the pooled and true $Y_{X/S}$ (Pearson r = 0.79, n = 5; mean absolute mismatch 16.1%, mean absolute error 12.1%; Fig. 4c); Run 5, which grows exponentially (v = 0), was the outlier (+ 9% error for + 32% mismatch). Calibrating $Y_{X/S}$ from one offline count (Eq. 9) estimated the yield with a median error of 8.9% and reduced the forecast error to 10.4% at 2.5 d (7.7% at 3.5 d). Results were similar with 15% noise on log D (Table S14). These are best-case values, because the pseudo-data are model-generated.

3.6. Perfusion demand, media use and time to dose

The re-implemented model reproduced the published final cell numbers within 1% (for example 2.54 versus 2.54×10^9 cells in Run 1 and 1.16 versus 1.17×10^9 cells in Run 5; Table S11). The set-point relation (Eq. 8) agreed with the full ODE within 0.4–4.9% after 2 d of perfusion but differed from the full ODE in cumulative medium use by -1% to + 20% because it neglects the controller transient (Table S12; Fig. S5), so the full ODE is used below. Predicted perfusion rates rose to 2.0-4.1 VVD, and cumulative medium use ranged from 3.5 vessel volumes (Run 5) to 13.9 (Run 1), equivalent to 405–731 mL medium per 10^9 cells, a 1.8-fold spread (Fig. 5a,b; Table 2). Glucose use was 0.58–0.87 g and lactate release 0.56–0.82 g per 10^9 cells. Time to 1×10^9 cells ranged from 1.9 d (Run 1) to 4.7 d (Run 5), and to 2×10^9 cells from 3.9 to 6.4 d (Fig. 5c), so a fixed expansion duration would leave some runs well past a given dose and others short.

Media demand per cell is inversely proportional to the glucose margin $S_l - S_{sp}$ (Eq. 8). At fixed kinetics, raising S_{sp} from 20 to 22 mmol L^{-1} (margin 5 to 3) raises media use by 67%, and lowering it to 17 mmol L^{-1} (margin 8) lowers it by 38% (Fig. 5d; Table S15). The model has no glucose-limitation term, so this identifies a cost lever to be tested rather than a recommended setting. The published yields also carry a stoichiometric signal: $Y_{P/S} = Y_{X/S} \cdot Y_{P/X}$ was 2.48 and 2.49 in Runs 4 and 6, 24% above the glycolytic ceiling of 2 mol lactate per mol glucose, and 1.42 in Run 5 (Table 2). Values above 2 imply lactate from sources other than glucose or bias in the yield estimates, which matters for any glucose-balance soft sensor.

Table 2
Model-implied biomanufacturing metrics for the five fitted runs.

Run	Medium used (VV)	Peak D (VVD)	Media per 10^9 cells (mL)	Glucose per 10^9 cells (g)	Lactate per 10^9 cells (g)	$Y_{P/S}$	t to 0.5×10^9 (d)	t to 1×10^9 (d)	t to 2×10^9 (d)
1	13.9	3.8	731	0.78	0.71	1.83	0.8	1.9	3.9
2	12.0	4.1	641	0.66	0.56	1.72	2.3	3.5	5.2
4	6.4	2.4	589	0.66	0.82	2.47	2.3	3.8	6.0
5	3.5	2.1	405	0.87	0.62	1.41	2.9	4.7	6.4
6	9.6	3.0	525	0.58	0.72	2.48	0.8	2.0	4.0

3.7. Scenario envelopes are wide and heavy-tailed

Independent min–max sampling gave a Run 1 final concentration with median 23.9×10^6 cells mL^{-1} and a 2.5–97.5th percentile range of $6.1–1{,}628\times10^6$ cells mL^{-1} (267-fold; Fig. 6a). Eighteen per cent of draws exceeded 100×10^6 cells mL^{-1}, far above any concentration in the analysed runs (maximum 32.5×10^6 cells mL^{-1}). Time to $X_T = 10\times10^6$ cells mL^{-1} had median 2.6 d (2.5–97.5th percentile 1.2–8.2 d), bracketing the published 2.59 d (Fig. 6b). The 97.5th percentile of the endpoint ranged from 1,541 to $1{,}747\times10^6$ cells mL^{-1} across 20 seeds (Fig. 6d), and the largest endpoints came from high μ_{max} with small $|v|$, combinations the fitted runs do not show (Fig. S2). These envelopes are bounds under ignorance of parameter coupling, not forecasts.

3.8. Smart-sensor consensus burden does not by itself identify the documented disturbance

Consensus events per day of process time ranged from 0.59 (Run 1) to 3.02 (Run 5), a 5.1-fold spread (Fig. 7a; Table S8). The documented-shutdown run (Run 3; 2.33 events d^{-1}) ranked second of five. Run 5, without a documented shutdown, had the highest burden, and Run 4 sat exactly at the 2 events d^{-1}

threshold. Individual sensors disagreed markedly: the standard-deviation-index sensor registered 1 alert in Run 3 and none in Run 5, whereas the window-based and fuzzy sensors registered 12–14 (Run 3) and 16–21 (Run 5) (Fig. 7b). With five runs, and alert thresholds set by the originators using these batches, we draw no inference about association with endpoint error or yield. A process-state gate needs event timing relative to known disturbances, not only counts.

3.9. An evidence gate restricts calibration to sources with process-level data

Within the targeted evidence set, only primary adult PBMC αβ T cells had process-level expansion and metabolic data (Bental & Deutsch, 1993; Maciolek et al., 2014; Amini et al., 2020; Costariol et al., 2020; Egan et al., 2025) and published fitted parameters (Table 3; Fig. S4; Table S9). Cord-blood T cells (Hippen et al., 2008; Lo Presti et al., 2018), γδ T cells (Sutton et al., 2016; Xiao et al., 2018; Bold et al., 2023) and iPSC-derived T cells (Nianias & Themeli, 2019; Netsrithong & Wattanapanitch, 2021) had protocol-level expansion evidence but no glucose and lactate kinetics or matched trajectories in this set. Menstrual-blood (Sabbaj et al., 2011) and breast-milk T cells (Bertotto et al., 1990; Wirt et al., 1992) had phenotyping evidence only. The gate therefore permits calibration for PBMC αβ T cells only, permits cord-blood parameters only if matched trajectories are obtained, assigns γδ cells a partial-expansion model with metabolic priors, requires a differentiation-aware model for iPSC-derived cells, and restricts menstrual-blood and breast-milk T cells to reverse-feasibility analysis. The map is qualitative and reflects the evidence set, not the biology of these cells.

Table 3
Targeted, non-systematic evidence set for six T-cell source classes and the model treatment permitted by the evidence gate.

Source class	Expansion / manufacturing evidence	Metabolic (glucose/lactate) evidence	Matched process data	Model treatment (evidence gate)
Primary adult PBMC αβ T	Direct (bioreactor-scale)	Direct	Published fitted parameters only (no raw trajectories)	Calibrate
Umbilical-cord blood T	Partial (protocol-level)	Absent in set	Absent	Calibrate only if matched trajectories obtained
γδ T	Partial (protocol-level, G-Rex / serum-free)	Absent in set	Absent	Partial-expansion model with metabolic priors
iPSC-derived T	Partial (differentiation protocols)	Absent in set	Absent	Differentiation-aware model
Menstrual-blood T	Absent (phenotyping only)	Absent	Absent	Reverse feasibility only
Breast-milk T	Absent (phenotyping only)	Absent	Absent	Reverse feasibility only

Evidence levels refer to this targeted, non-systematic set only; absence is not evidence of absence. Key references per class are listed in Table S9.

3.10. Reverse feasibility defines what a new source must show

For a reference process taking 1.49×10^6 cells mL^{-1} to 19×10^6 cells mL^{-1} in 4.94 d, the required μ_{max} depends strongly on v: 0.53 d^{-1} for exponential growth (v = 0), 1.12 d^{-1} at the pooled v (− 0.537) and 2.12 d^{-1} at v = −0.922 (t_{lag} = 0.092 d). Varying t_{lag} from 0 to 0.19 d changed the requirement by 2% or less. Overlaying the fitted runs (Fig. 6c; Fig. S3), Runs 2 and 4 exceeded the requirement at their own v and t_{lag} (+ 94% and + 25%), Runs 1 and 6 lay on it (− 0.3% and 0.0%), and Run 5 fell short (-21%). A new source can be placed on this map once its (μ_{max}, v, t_{lag}) are measured. The thresholds are properties of the process and target, not of the source, so they cannot rank sources or predict which will meet them.

3.11. Run inclusion, data availability and run-level structure

Of the six published runs, five had fitted parameters and entered the pooled, leave-one-run-out, identifiability and forecasting analyses; Run 3 (documented interruption) had no fitted parameters and

served only as a held-out test; Run 6 had no consensus counts in the source study; and none had raw trajectories available (Fig. 8a; Table 4). Hierarchical clustering of the five fitted runs on ten features gave a two-cluster cut of Runs 1, 2 and 6 versus Runs 4 and 5 (Fig. 8b). The cut did not follow medium (Run 6, medium B, clustered with medium-A donors), cell source (the CAR T Run 4 clustered with donor Run 5) or the glucose set-point. It coincided with the two runs whose seeding time t_l was about 2 d (Runs 4 and 5: 2.09 and 1.92 d, versus ≤ 0.14 d in the others), an annotation we did not hypothesise in advance and cannot separate from other differences between those two runs. No correlation structure exceeded chance at this sample size: the mean |ρ| over 45 feature pairs was 0.48 against 0.42 expected under independent permutation (P = 0.07), and no pair reached |ρ| ≥ 0.9 (null mean 0.86 pairs; Fig. 8c). The visible blocks, namely μ_{max}, t_{lag}, peak D and media use varying together, and v, time to dose and the glucose margin varying together, are therefore hypotheses: the first is partly algebraic (perfusion demand follows growth through Eq. 8) and the second is driven by Run 5, the only run with a different glucose margin. Runs reaching 10^9 cells within 3 d (Runs 1 and 6) differed completely from the other three in v (more negative) and in seeding density (1.49 versus 0.37–1.13×10^6 cells mL^{-1}), but not in μ_{max}, t_{lag}, the yields or peak perfusion rate. With two versus three runs the smallest attainable exact P value is 0.2, so this is a pattern to test, not a result.

Table 4
Run and process characteristics.

Run	Source	Medium	X_l	t_l (d)	S_{sp}	μ_{max} (d^{-1})	v	$Y_{X/S}$	Used for
1	Donor 1	A	1.49	0.11	20	2.14	−0.92	0.231	F L I C
2	Donor 2	A	0.37	0.14	20	2.02	−0.48	0.275	F L I C
3	Donor 3	A	1.49	0.11	20	n/a	n/a	n/a	H C
4	Patient 1 (CAR T)	A	0.62	2.09	20	1.39	−0.52	0.272	F L I C
5	Donor 4	A	1.13	1.92	17	0.41	0.00	0.207	F L I C
6	Donor 5	B	1.49	0.08	21	1.58	−0.76	0.311	F L I

Fitted values from Egan et al. (2025). *X_l in 10^6 cells mL^{-1}; S_{sp} in mmol L^{-1}; μ_{max} in d^{-1}. Used for: F, fitted run; L, pooled and leave-one-run-out tests; I, identifiability and forecasting; C, consensus counts available; H, held-out test run. Run 3 has no published fitted parameters.*

3.12. Seeding density, not only kinetics, sets time to dose

Relative to a reference of X_l = 1.13×10^6 cells mL^{-1} with pooled kinetics (2.26 d to 10^9 cells), the time to dose differed by − 0.37 d (Run 1) to + 2.40 d (Run 5) (Table 5; Fig. 8d). Runs 1 and 6, the fastest, were faster because they were seeded densely (− 0.56 and − 0.57 d), whereas their kinetics lengthened the

time to dose (+ 0.19 and + 0.31 d): they did not grow faster than the pooled mean. Run 2 was 2.30 d slower because of its low seeding density (0.37×10^6 cells mL^{-1}) but 1.07 d faster in kinetics; Run 4 was 1.33 d slower from seeding with kinetics adding 0.22 d; and the 2.40 d delay of Run 5, seeded at the reference density, was entirely kinetic (exponential growth, μ_{max} 0.405 d^{-1}). Seeding therefore contributed more than kinetics in four of five runs, so much of the 2.5-fold spread in time to dose in Section 3.6 reflects process design rather than donor-intrinsic growth. This is a model-implied decomposition, not a causal estimate: the fitted kinetics were estimated conditional on each run's own X_I, and the pooled reference predicts poorly out of sample (Section 3.2). Per run at 2.5 d (5% noise), the median absolute error in X_F ranged from 1.2% to 26.6% for the Egan-like estimator (mean 12.3%), from 0.4% to 22.2% for the oracle (mean 7.5%) and from 6.4% to 13.6% for the single-count mass-balance estimator (mean 10.4%) (Fig. 8e). The single offline count therefore roughly halved the worst case of the Egan-like estimator without improving its best case, and the oracle was not best in every run (worst in Run 5, the exponential run).

Table 5

Attribution of time to 10^9 cells (134 mL vessel) to seeding density and kinetics.

Run	X_I	t to 10^9 cells (d)	Δ vs reference (d)	Seeding contribution (d)	Kinetics contribution (d)
1	1.49	1.89	−0.37	−0.56	+ 0.19
2	0.37	3.48	+ 1.22	+ 2.30	−1.07
4	0.62	3.81	+ 1.55	+ 1.33	+ 0.22
5	1.13	4.66	+ 2.40	0.00	+ 2.40
6	1.49	2.00	−0.26	−0.57	+ 0.31
Ref.	1.13	2.26	0	0	0

Shapley values relative to a reference with X_I equal to the median of the fitted runs and pooled kinetics (t_{lag} 0.092 d, v − 0.537, μ_{max} 1.51 d^{-1}); contributions sum to the difference from the reference. Model-implied, not causal.

Shapley values relative to a reference with X_I equal to the median of the fitted runs and pooled kinetics (t_{lag} 0.092 d, v − 0.537, μ_{max} 1.51 d^{-1}); contributions sum to the difference from the reference. Model-implied, not causal.

4. Discussion

The main finding is methodological: for a digital shadow whose growth parameter is computed from the endpoint, agreement with the endpoint is guaranteed to within rounding and carries no validation weight, while the same model driven by pooled historical parameters errs by tens of per cent (LOO MAE 69%).

Reporting these quantities separately is inexpensive and prevents a favourable in-sample statistic from being read as predictive performance. The 13% error reported by Egan et al. for real-time updating is the relevant benchmark for a shadow with online data, and our pseudo-real-time analysis suggests where it comes from.

Under glucose set-point control the shadow is, in effect, a glucose mass-balance soft sensor: the perfusion record determines growth up to the yield $Y_{X/S}$. Online data therefore identify $\mu_{max}/Y_{X/S}$, $Y_{X/S}{\cdot}Y_{P/X}$ and v, and the absolute cell concentration is only as good as the yield assumed. Fixing $Y_{X/S}$ at its pooled mean, a pragmatic choice that Egan et al. justify by its low variability, is sufficient to reproduce their reported error in our model-correct test, and the error tracks the run's yield mismatch (r = 0.79). The practical implication is that improving the real-time forecast requires run-specific information on the yield rather than more or better perfusion data. A single offline count, for example the suspension sample already taken at passage or at the end of perfusion, cut the error from 12% to 10% in our tests, and hierarchical priors on the yields, informed by more runs, are a natural next step. A resuspension count disturbs a sedimented culture, so its value must be weighed against that cost.

The biomanufacturing metrics show why run-specific forecasts matter operationally. Time to a given cell number varied up to 2.5-fold between runs (for example 1.9 to 4.7 d for $1{\times}10^9$ cells), and media use per 10^9 cells varied 1.8-fold. Because demand scales with $1/(S_l - S_{sp})$, the glucose margin is a first-order cost lever that standard set-point choices may not exploit. The model lacks glucose-limitation and lactate-inhibition terms, so such levers need experimental testing, ideally within a design-of-experiments or quality-by-design framework (Hood et al., 2024; Garcia-Aponte et al., 2021). The $Y_{P/S}$ values above 2 in two runs deserve attention in their own right, since lactate may derive partly from other substrates, and a shadow that assumes glucose-only stoichiometry would then misestimate glucose demand.

Borrowing the layered reporting of a large clinical transcriptomic study (Andreoletti et al., 2023) clarified three points. First, as in that study, where most participants improved and few differences separated responders from non-responders, five runs leave little power: no clustering, correlation module or stratified contrast here exceeds what chance and model algebra allow, and we present them as descriptive. Second, an explicit inclusion map is itself a result: only five of six runs support any fitting and none supports validation against raw data. Third, the attribution of time to dose shows that process-design factors such as seeding density can masquerade as donor kinetics, so a harvest-timing model should report seeding and kinetic contributions separately. Like that study, ours lacks independent control data and mixes sampling contexts, here different donors, media and one patient-derived product.

The smart-sensor analysis supports a cautious design for the process-state gate. Consensus counts varied five-fold between runs, were not highest in the documented-shutdown run and depended on sensor choice, so a count-only gate would misclassify the runs analysed here. The 2 events d^{-1} threshold is illustrative and should be replaced by a threshold calibrated on held-out runs. Such calibration is not possible from the published counts: the thresholds were set by the originators on the same batches, and five runs with one documented disturbance cannot estimate detection or false-alarm

rates. It would require raw sensor trajectories from many independent runs, independently logged disturbances (for example shutdowns and feed or glucose excursions) to label alerts, per-sensor thresholds and the number of agreeing sensors fitted on training runs, validation on held-out runs, and re-checking whenever cell source, medium or process changes. The evidence gate and reverse-feasibility map address transfer to other T-cell sources: they state what must be measured before parameters may be reused, namely time-resolved viable-cell concentration, glucose, lactate and perfusion rate from multiple independent runs, and do not predict that any source will meet a target.

This study has limitations. It is a secondary analysis of published tables: five fitted runs, rounded parameters, no raw trajectories and an assumed correspondence between the run labels of two papers. The identifiability and forecasting experiments use pseudo-data generated by the model itself, with assumed sensor noise and a known controller, so they give best-case results; structural errors, sensor bias, bubble artefacts and unmodelled heterogeneity would make real errors larger. The re-implemented controller (including the D_{min} floor) was not validated against measured perfusion profiles. Media and time-to-dose figures are model-implied, and the set-point sensitivity ignores growth effects of glucose. Scenario envelopes ignore parameter coupling by design. The smart-sensor rule was not calibrated, and the evidence set is targeted and non-systematic. Product-quality attributes and release criteria are outside the scope of the model, and no regulatory claim is made. The run-level clustering, correlation, stratified and attribution analyses use five fitted runs and are descriptive; with two versus three runs the smallest exact permutation P value is 0.2.

Three steps would convert these findings into validated claims: obtaining the raw online trajectories from the originating groups to repeat the identifiability and forecasting analyses on measured data with blocked cross-validation; fitting a hierarchical model with run-specific yields and the $v–\mu_{max}$ coupling; and prospectively scoring harvest-window forecasts, with the process-state gate frozen in advance, in new runs including at least one non-PBMC source.

5. Conclusions

Endpoint agreement of a published T-cell digital shadow is a consistency check, and pooled parameters predicted new runs with a mean absolute error of 69%. Under glucose set-point control the online data identify parameter combinations but not the glucose yield, which sets a floor on real-time forecast error (about 12% in a model-correct test, matching the published 13%) that one offline count can partly lower. The fitted kinetics imply a 1.8-fold range in media demand and up to 2.5-fold range in time to dose, with media use governed by the glucose margin. An evidence gate and reverse-feasibility map make explicit which inputs may calibrate a model and what a new T-cell source must show. Run-specific yield information and matched source-specific datasets are the priorities for turning digital shadows into reliable harvest-timing tools.

Abbreviations

CAR
chimeric antigen receptor
CQA
critical quality attribute
CV
coefficient of variation
LOO
leave-one-run-out
MAE
mean absolute relative error
ODE
ordinary differential equation
PBMC
peripheral blood mononuclear cells
PI
proportional-integral
RMSE
root-mean-square error
VVD
vessel volumes per day

Declarations

Conflict of interest:

No conflict of interest

Ethics statement

Not applicable. This is a secondary analysis of previously published, aggregated data.

Use of generative AI

No use of AI

Funding:

No source of funding

Author contributions:

Analysis, design and writing all performed by Faizan Ahmad

Acknowledgements:

We thank the authors of Egan et al. (2025) and Nettleton et al. (2023) for openly publishing the run-level data used here.

Data availability.

All inputs are transcribed from published tables (provenance in data/README.md; files data/*.csv). Code, derived tables and figures are available at https://github.com1996faizanahmad-beep/tcell-digital-shadow-bb; Raw 10-s trajectories are held by the originating groups (requests via the corresponding author of Egan et al., 2025) and were not used.

References

1. Amini A, Wiegmann V, Patel H, Veraitch F, Baganz F (2020) Bioprocess considerations for T-cell therapy: Investigating the impact of agitation, dissolved oxygen, and pH on T-cell expansion and differentiation. Biotechnol Bioeng 117(10):3018–3028. https://doi.org/10.1002/bit.27468
2. Andreoletti G, Romano O, Chou H-J et al (2023) High-throughput transcriptome analyses from ASPIRO, a phase 1/2/3 study of gene replacement therapy for X-linked myotubular myopathy. Am J Hum Genet 110:1648–1660. https://doi.org/10.1016/j.ajhg.2023.08.008
3. Bental M, Deutsch C (1993) Metabolic changes in activated T cells: An NMR study of human peripheral blood lymphocytes. Magn Reson Med 29(3):317–326. https://doi.org/10.1002/mrm.1910290307
4. Bertotto A, Gerli R, Fabietti G et al (1990) Human breast milk T lymphocytes display the phenotype and functional characteristics of memory T cells. Eur J Immunol 20(8):1877–1880. https://doi.org/10.1002/eji.1830200838
5. Bold A, Gross H, Holzmann E et al (2023) An optimized cultivation method for future in vivo application of γδ T cells. Front Immunol 14:1185564. https://doi.org/10.3389/fimmu.2023.1185564
6. Costariol E, Rotondi M, Amini A et al (2019) Establishing the scalable manufacture of primary human T-cells in an automated stirred-tank bioreactor. Biotechnol Bioeng 116(10):2488–2502. https://doi.org/10.1002/bit.27088
7. Costariol E, Rotondi MC, Amini A et al (2020) Demonstrating the manufacture of human CAR-T cells in an automated stirred-tank bioreactor. Biotechnol J 15(9):e2000177. https://doi.org/10.1002/biot.202000177

8. Egan JR, Marí-Buyé N, Benítez-Cano V, E., et al (2025) A digital shadow of CAR T cell expansion in a perfusion bioreactor: Informing optimal harvest times for autologous cell therapy. Biotechnol Prog 41(6):e70045. https://doi.org/10.1002/btpr.70045
9. Garcia-Aponte OF, Herwig C, Kozma B (2021) Lymphocyte expansion in bioreactors: Upgrading adoptive cell therapy. J Biol Eng 15(1):13. https://doi.org/10.1186/s13036-021-00264-7
10. Gatla H, Uth N, Levinson Y et al (2022) Enabling allogeneic T cell-based therapies: Scalable stirred-tank bioreactor mediated manufacturing. Front Med Technol 4:850565. https://doi.org/10.3389/fmedt.2022.850565
11. Hewitt MM, Trainor N, Ostrout N, Abraham E (2021) Cell therapy manufacturing: Process analytic technologies needed to achieve flexible, feedback-driven automation. Curr Opin Biomedical Eng 20:100358. https://doi.org/10.1016/j.cobme.2021.100358
12. Hippen KL, Harker-Murray P, Porter SB et al (2008) Umbilical cord blood regulatory T-cell expansion and functional effects of tumor necrosis factor receptor family members OX40 and 4-1BB expressed on artificial antigen-presenting cells. Blood 112(7):2847–2857. https://doi.org/10.1182/blood-2008-01-132951
13. Hood T, Slingsby F, Sandner V et al (2024) A quality-by-design approach to improve process understanding and optimise the production and quality of CAR-T cells in automated stirred-tank bioreactors. Front Immunol 15:1335932. https://doi.org/10.3389/fimmu.2024.1335932
14. Lo Presti V, Nierkens S, Boelens JJ, van Til NP (2018) Use of cord blood derived T-cells in cancer immunotherapy: Milestones achieved and future perspectives. Expert Rev Hematol 11(3):209–218. https://doi.org/10.1080/17474086.2018.1431119
15. Maciolek JA, Pasternak JA, Wilson HL (2014) Metabolism of activated T lymphocytes. Curr Opin Immunol 27:60–74. https://doi.org/10.1016/j.coi.2014.01.006
16. Narayanan H, Luna MF, von Stosch M et al (2020) Bioprocessing in the digital age: The role of process models. Biotechnol J 15(1):e1900172. https://doi.org/10.1002/biot.201900172
17. Netsrithong R, Wattanapanitch M (2021) Advances in adoptive cell therapy using induced pluripotent stem cell-derived T cells. Front Immunol 12:759558. https://doi.org/10.3389/fimmu.2021.759558
18. Nettleton DF, Marí-Buyé N, Marti-Soler H et al (2023) Smart sensor control and monitoring of an automated cell expansion process. Sensors 23(24):9676. https://doi.org/10.3390/s23249676
19. Nianias A, Themeli M (2019) Induced pluripotent stem cell (iPSC)-derived lymphocytes for adoptive cell immunotherapy: Recent advances and challenges. Curr Hematol Malig Rep 14(4):261–268. https://doi.org/10.1007/s11899-019-00528-6
20. Richards FJ (1959) A flexible growth function for empirical use. J Exp Bot 10(2):290–301. https://doi.org/10.1093/jxb/10.2.290
21. Sabbaj S, Hel Z, Richter HE, Mestecky J, Goepfert PA (2011) Menstrual blood as a potential source of endometrial derived CD3 + T cells. PLoS ONE 6(12):e28894. https://doi.org/10.1371/journal.pone.0028894

22. Song HW, Prochazkova M, Shao L et al (2024) CAR-T cell expansion platforms yield distinct T cell differentiation states. Cytotherapy 26(7):757–768. https://doi.org/10.1016/j.jcyt.2024.03.003
23. Springuel P, Silva Couto P, Stibbs DJ et al (2025) Scalable CAR-T production in a 2-litre perfusion stirred-tank bioreactor with automated harvesting and scale-down model characterisation. Front Bioeng Biotechnol 13:1694134. https://doi.org/10.3389/fbioe.2025.1694134
24. Sutton KS, Dasgupta A, McCarty D, Doering CB, Spencer HT (2016) Bioengineering and serum free expansion of blood-derived γδ T cells. Cytotherapy 18(7):881–892. https://doi.org/10.1016/j.jcyt.2016.04.001
25. Wanika L, Egan JR, Swaminathan N et al (2024) Structural and practical identifiability analysis in bioengineering: A beginner's guide. *Journal of Biological Engineering, 18*. https://doi.org/10.1186/s13036-024-00410-x
26. Wirt DP, Adkins LT, Palkowetz KH, Schmalstieg FC, Goldman AS (1992) Activated and memory T lymphocytes in human milk. Cytometry 13(3):282–290. https://doi.org/10.1002/cyto.990130310
27. Xiao L, Chen C, Li Z et al (2018) Large-scale expansion of Vγ9Vδ2 T cells with engineered K562 feeder cells in G-Rex vessels and their use as chimeric antigen receptor-modified effector cells. Cytotherapy 20(3):420–435. https://doi.org/10.1016/j.jcyt.2017.12.014

Figures

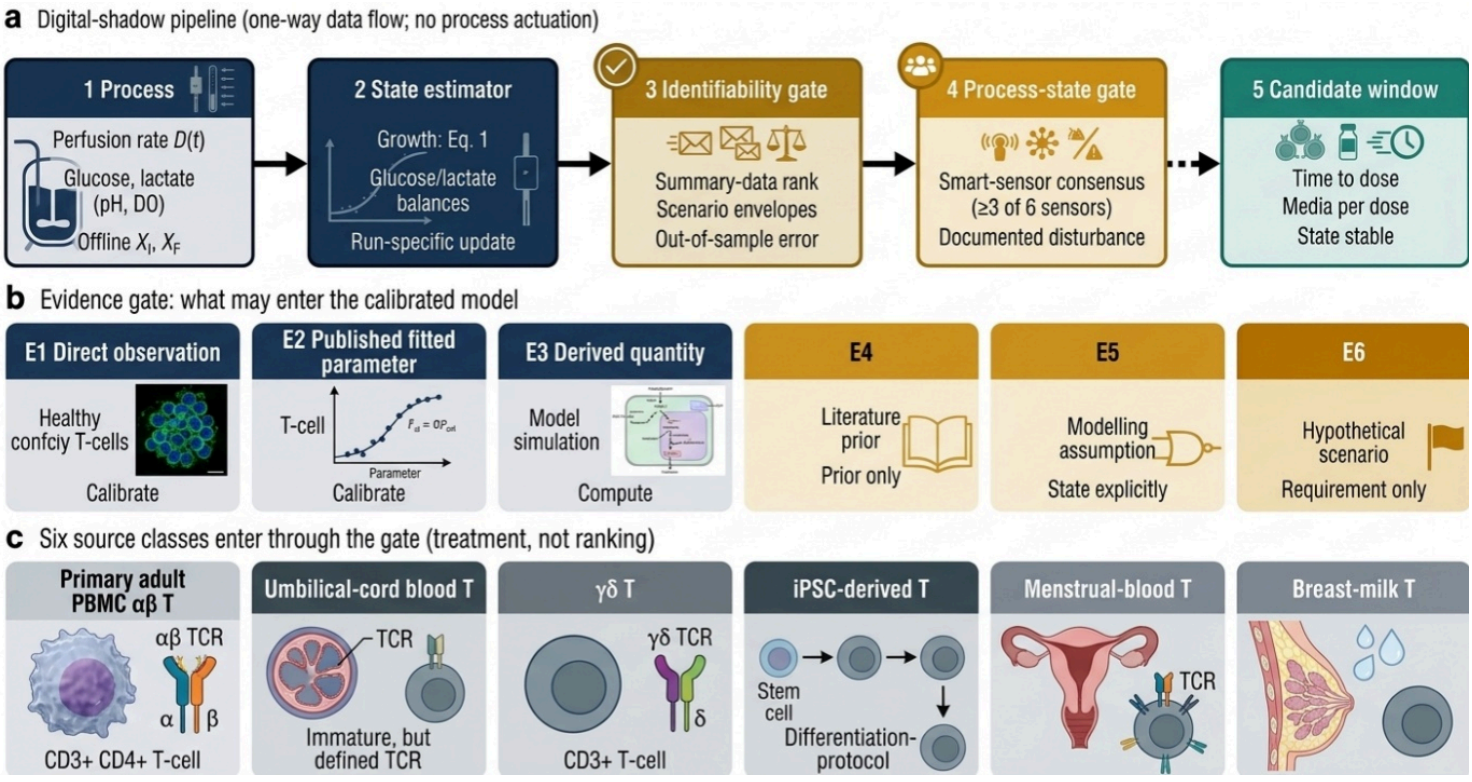


Figure 1

Evidence-gated digital-shadow framework. ***a****, One-way pipeline from process and online data to a candidate harvest window with time-to-dose and media-demand outputs; gates 3 and 4 constrain the forecast and the process state. Release decisions rely on independent critical-quality-attribute (CQA) assays and are not modelled.* ***b****, Evidence classes E1–E6 and how each may enter the model.* ***c****, Six T-cell source classes and the model treatment the gate permits for each. The schematic shows treatment rules, not a ranking of sources.*

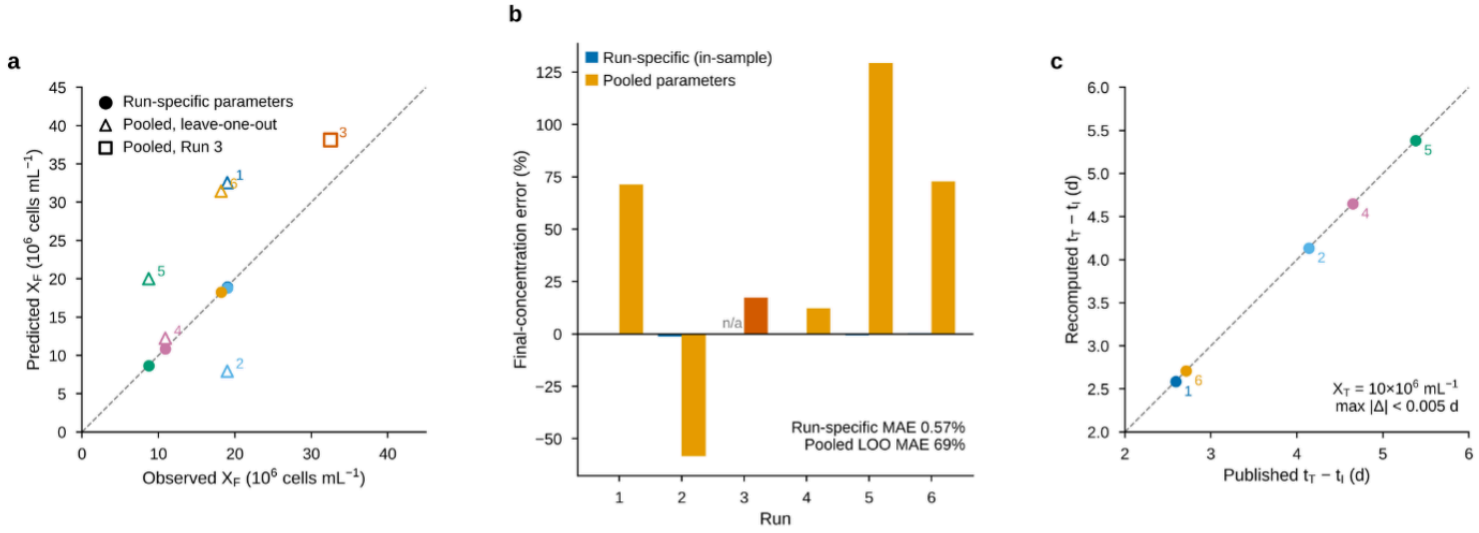


Figure 2

Run-specific agreement is a consistency check; pooled parameters fail out of sample. ***a****, Predicted versus observed final viable-cell concentration (X_F) for run-specific parameters (filled circles, five fitted runs), pooled leave-one-out parameters (open triangles) and pooled parameters for Run 3 (open square); dashed line, identity.* ***b****, Signed relative error by run for run-specific (blue) and pooled (orange; Run 3 in vermilion) predictions; run-specific MAE 0.57%, pooled LOO MAE 69%.* ***c****, Recomputed versus published time to target (X_T = 10×10^6 cells mL^{-1}); maximum absolute difference below 0.005 d. Inputs from Egan et al. (2025).*

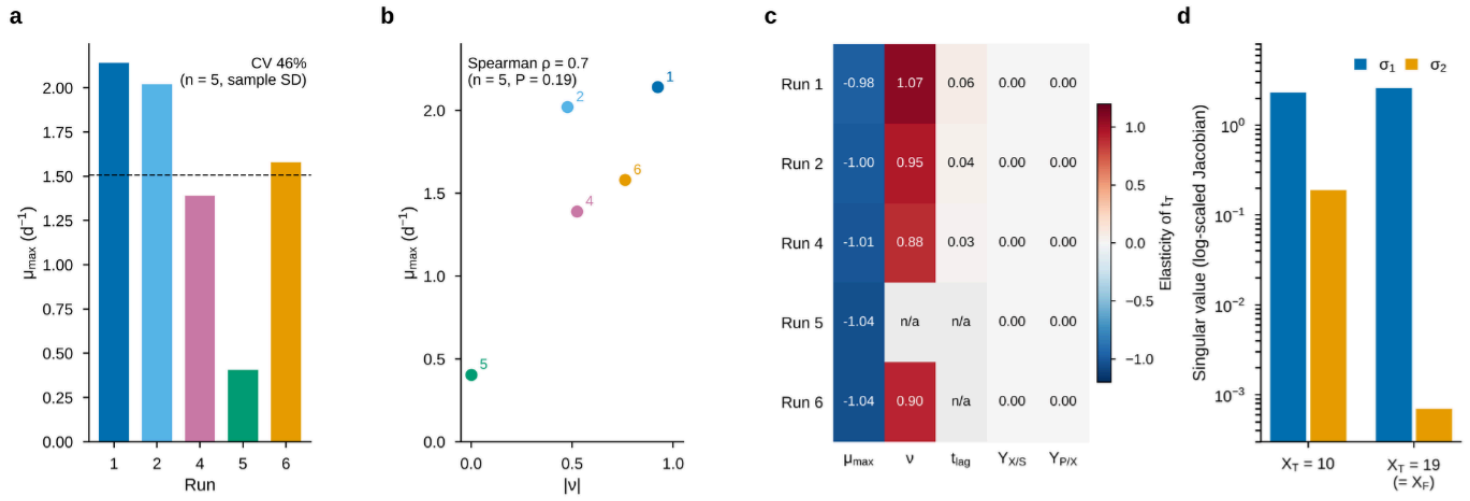


Figure 3

Parameter heterogeneity, harvest-time sensitivity and information content of summary observables. ***a***, *μ_{max} by run (dashed line, mean; CV from sample SD, n = 5).* ***b***, *μ_{max} versus |ν| (Spearman ρ = 0.7, P = 0.19, n = 5).* ***c***, *Secant elasticity of the time to target for ±20% parameter changes; grey, undefined because the parameter is zero; yields have zero elasticity by model structure.* ***d***, *Singular values of the log-scaled Jacobian of [X_F, t_T] for Run 1 at X_T = 10 and 19×10^6 cells mL^{-1} (the latter equals X_F).*

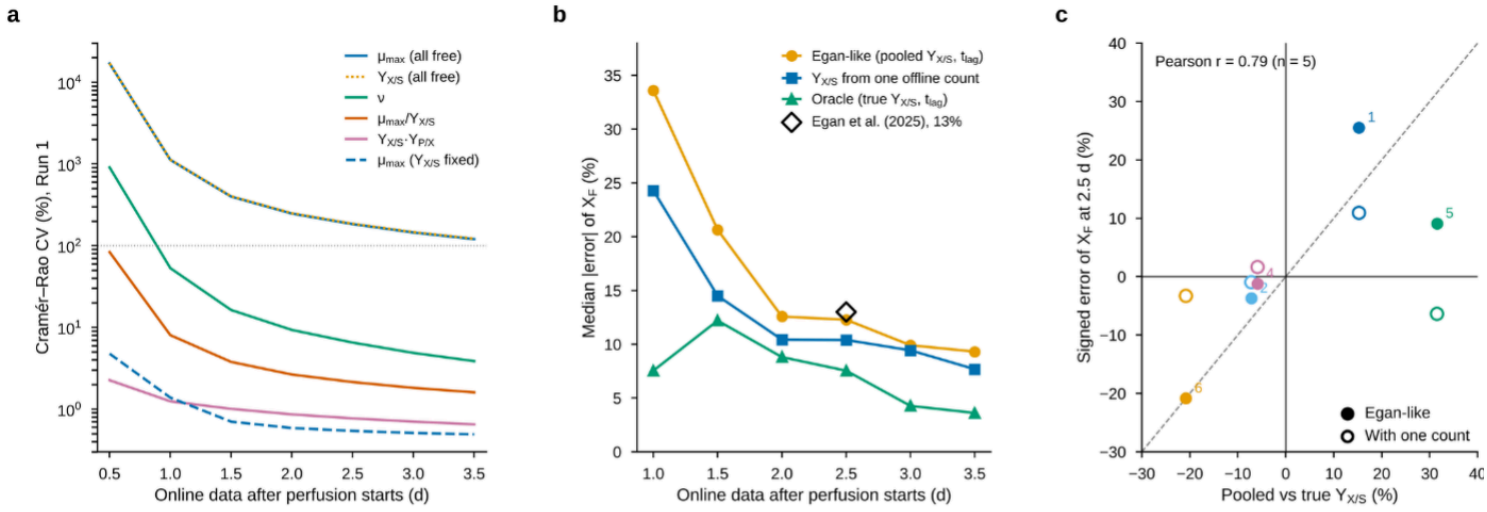


Figure 4

Practical identifiability and pseudo-real-time forecasting. ***a***, *Cramér–Rao CV of parameters and parameter combinations versus length of online data for Run 1 (log scale; 5% noise on log perfusion rate, 0.5 mmol L^{-1} on lactate; model assumed correct). The combinations $\mu_{max}/Y_{X/S}$ and $Y_{X/S}{\cdot}Y_{P/X}$ are determined far more tightly than the individual parameters; fixing $Y_{X/S}$ (dashed) removes the degeneracy.* ***b***, *Median absolute error of predicted final concentration versus online data (mean of five held-out runs; 16*

noise replicates): Egan-like estimator with pooled $Y_{X/S}$ and t_{lag} (orange), $Y_{X/S}$ calibrated from one offline count (blue) and oracle with true values (green); the diamond marks the 13% reported by Egan et al. (2025). ***c****, Signed error at 2.5 d versus the mismatch between pooled and true $Y_{X/S}$ (filled, Egan-like; open, with one count; dashed, identity). Pseudo-data were generated by the published fitted model, so results are best-case.*

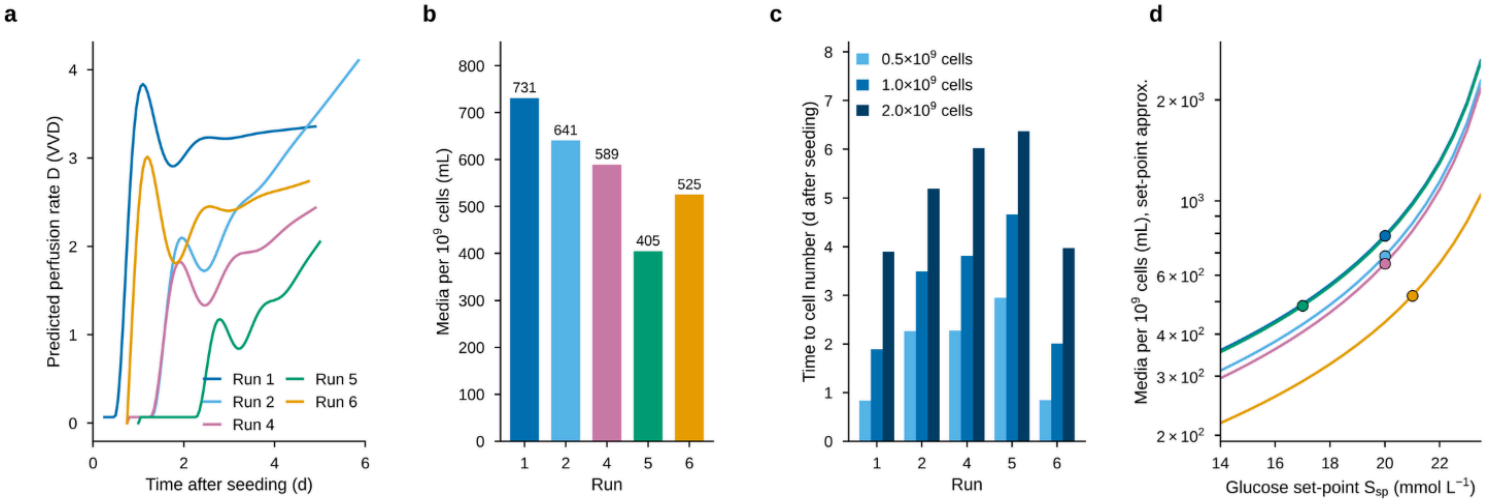


Figure 5

Model-implied perfusion demand, media use and time to dose. ***a****, Predicted perfusion rate (vessel volumes per day, VVD) from the re-implemented PI-controlled model.* ***b****, Media per 10^9 cells (full ODE).* ***c****, Time after seeding to 0.5, 1 and 2×10^9 cells (134 mL vessel).* ***d****, Media per 10^9 cells versus glucose set-point at fixed kinetics (set-point approximation; dots, published set-points; dotted line, glucose in fresh medium for Runs 1–5). Model-implied values, not measurements.*

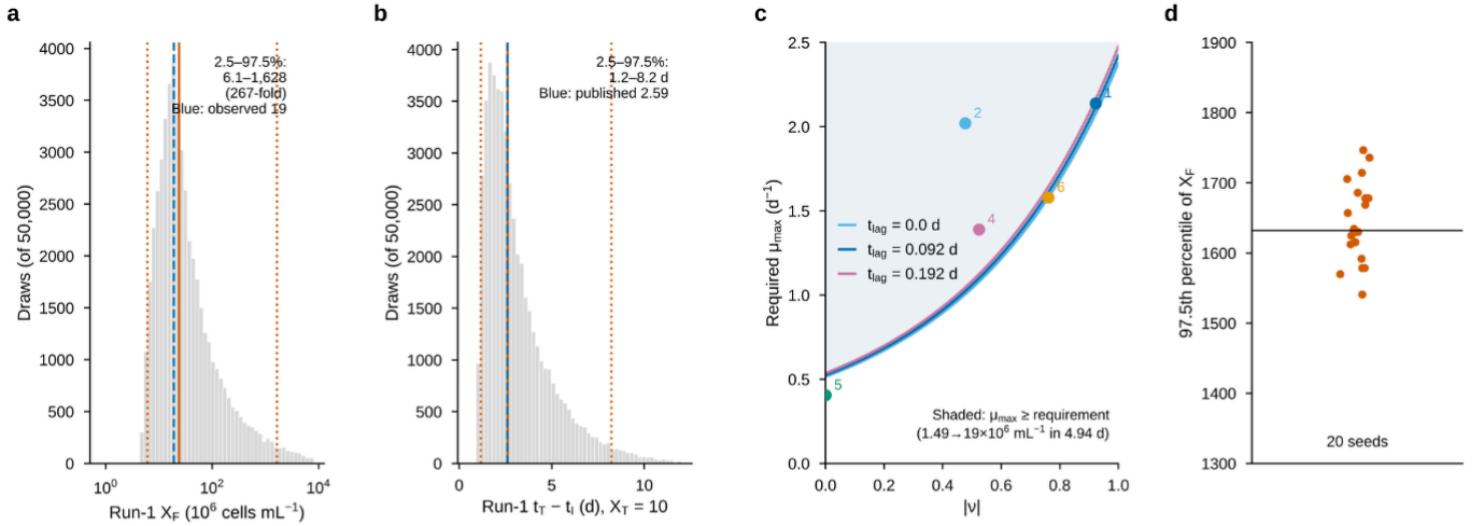


Figure 6

Exploratory scenario envelopes and reverse feasibility. a,**b**, Distributions of Run 1 final concentration (**a**) and time to target (**b**) under independent uniform sampling of t_{lag}, v and μ_{max} between fitted extremes (N = 50,000, seed 42; dotted lines, 2.5th and 97.5th percentiles; solid, median; blue dashed, observed or published value). These are not prediction intervals. **c**, μ_{max} required to take 1.49 to 19×10^6 cells mL^{-1} in 4.94 d as a function of |v| for three lag values; shading, μ_{max} at or above the requirement for t_{lag} = 0.092 d; points, fitted runs. **d**, 97.5th percentile of the endpoint across 20 random seeds.

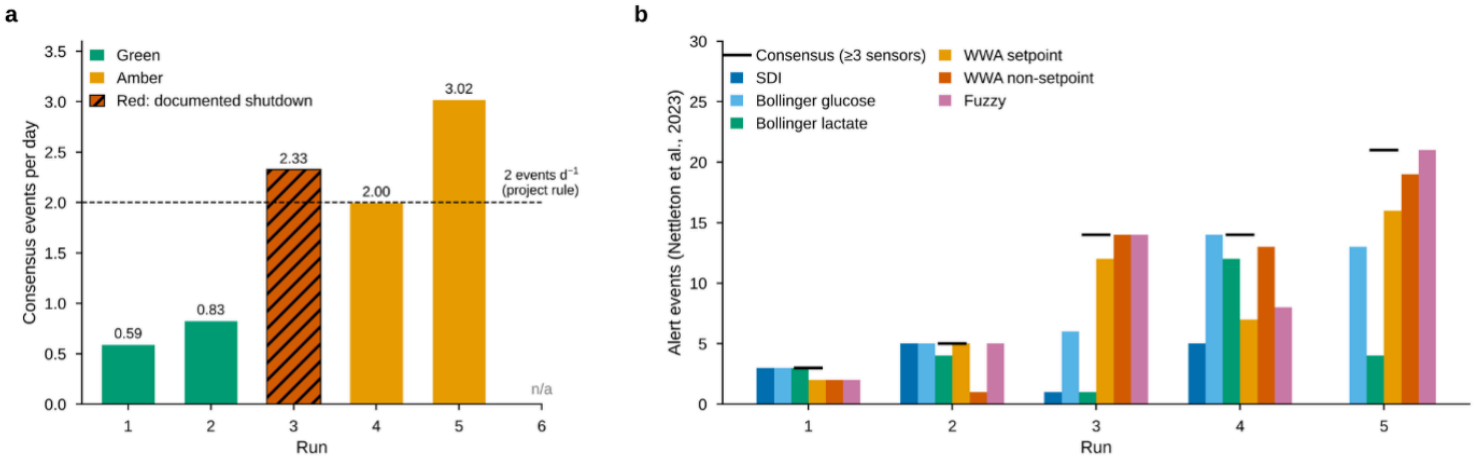


Figure 7

Smart-sensor consensus burden by run. ***a****, Consensus events (three or more of six sensors) per day of process time. Red, documented mechanical shutdown (Run 3); amber, at least 2 events* d^{-1} *without a documented shutdown; green, otherwise. Run 6 was not part of the source study.* ***b****, Alert events by individual sensor and consensus count per run. Data from Nettleton et al. (2023); run labels assumed to correspond to Egan et al. (2025).*

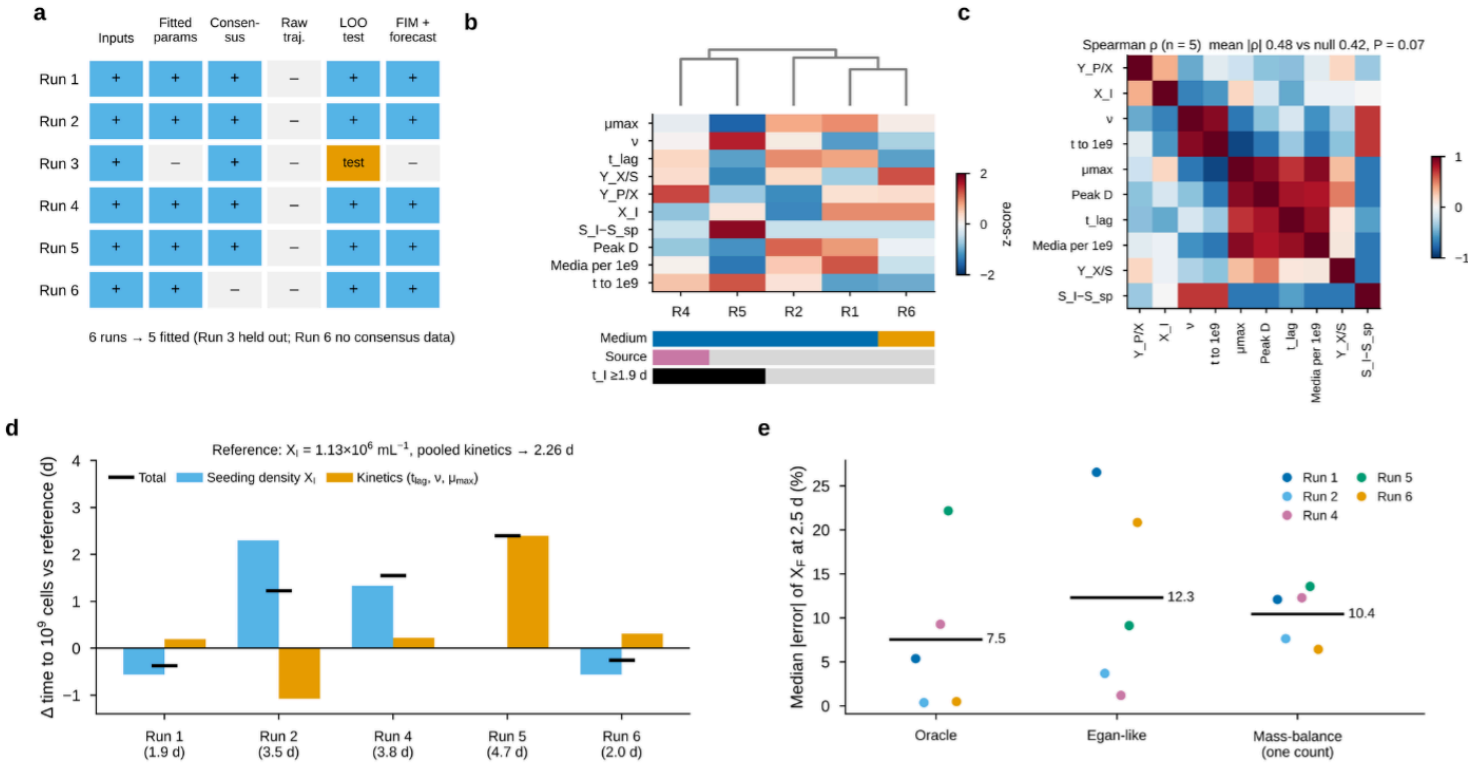


Figure 8

Run-level reporting and exploratory analyses. *a, Data-availability matrix for the six published runs (+, available; test, held-out test run; –, not available). b, Hierarchical clustering (Pearson correlation distance, average linkage) of the five fitted runs on ten z-scored features, annotated by medium, cell source and seeding time. c, Spearman correlation of the same features (n = 5) with the mean |ρ| and the P value against a permutation null. d, Attribution of the time to 10^9 cells, relative to a reference with median X_I and pooled kinetics, to seeding density and kinetics (Shapley values); bars sum to the total (black marks). e, Median absolute error in X_F at 2.5 d (5% noise) per run for three estimators; lines are means. Model-implied from published tables; five fitted runs, so analyses are descriptive.*

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.docx
- floatimage1.jpeg