ZYVOLANCE THERAPEUTICS LLC

TECHNICAL NOTE · NO. 01 · MAY 2026

Why the patient is the *reference.*

A structural argument for per-patient self-reference and the abandonment of cohort means as the inferential baseline for individual care.

The average patient does not exist. Cohort means are statistical constructs that no individual occupies. Population reference ranges describe groups; they do not describe persons. For control-system applications — monitoring, dosing, prediction — the only physiologically valid baseline is the patient's own pre-perturbation trajectory. This note sets out the structural reasons, the variance mathematics, and the historical pharmacology that converge on a single claim: precision medicine begins with self-reference.

ZYVOLANCE TECHNICAL NOTES — SERIES N-OF-1
NO. 01 OF AN ONGOING SERIES

SERIES
N-of-1 Foundations

CORRESPONDENCE
contact@zyvolance.com

PATENT REFERENCE
OEPM ZYV-2026-HARL-001

SECTION 1

Cohort means describe groups. They do not describe patients.

The average patient does not exist. A cohort mean is an arithmetic summary over individuals — a number that exists in statistics, not in physiology. No patient simultaneously occupies the mean glucose, mean heart rate, mean creatinine clearance, and mean QT interval of a population. The mean is a property of the cohort. It is not a property of any patient.[1,2]

A systolic blood pressure of 90 mmHg is catastrophic in a patient whose usual BP is 130 and who is collapsing toward shock. The same 90 mmHg is excellent news in a patient recovering from coma whose BP yesterday was 65. The number is identical. The physiology is opposite. The clinical inference rests entirely on the patient's own prior state — on the trajectory, not the value.[3]

An HbA1c of 7·4% can describe a patient with stable, well-governed glucose. It can equally describe a patient with daily severe hypoglycaemia and post-prandial peaks above 300 mg/dL. The aggregate has no relationship to either clinical reality. The cohort average hides the trajectory — and HbA1c is among the most misleading single parameters in modern medicine because it is constructed precisely to suppress within-patient variation.[4,5]

The variance decomposition

For any physiological variable measured longitudinally, total variance decomposes into between-patient and within-patient components:

$$\sigma^2_{total} = \sigma^2_{between} + \sigma^2_{within}$$

Between-patient variance is large — patients differ from one another. Within-patient variance is typically tight: each patient operates around a stable set-point with a narrow band of normal fluctuation. Cohort reference ranges absorb both components into one threshold, and lose the very signal a control system needs to detect.

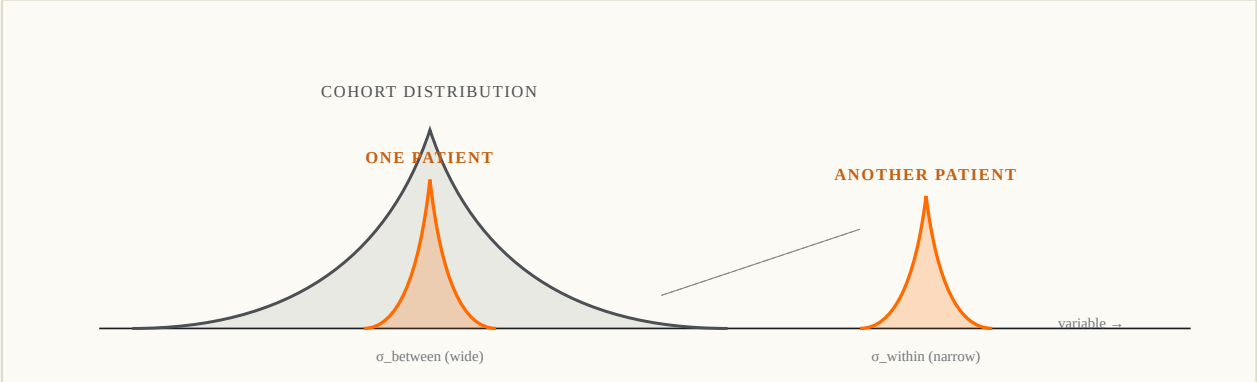

SECTION 2

The patient's own pre-perturbation trajectory is the baseline.

When the patient is their own reference, the inferential question changes form. It is no longer "where does this measurement fall in the population distribution?" It is "where does this measurement fall in this person's own historical band of operation?" The first is a classification question. The second is a control question. They require different mathematics.

Self-referenced state

Define the patient's pre-perturbation state as the rolling trajectory of the variable in question across a window prior to any putative event:

$$\sigma(t)_i = [\, X_i(t) - \mu_i(t) \,] \,/\, \sigma_i(t)$$

where $X_i(t)$ is the patient's current measurement, $\mu_i(t)$ the patient's own moving baseline, and $\sigma_i(t)$ the patient's own historical dispersion. The reference is the patient. There is no fictional comparator. There is no cohort heterogeneity to absorb. Every patient's reading is interpreted against their own physiology.[6,7]

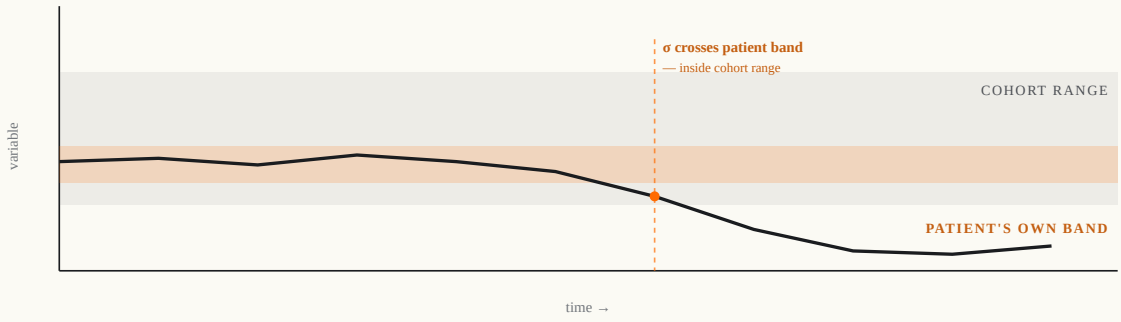

FIGURE 2 SELF-REFERENCED DETECTION. THE PATIENT'S MEASURED VARIABLE DEPARTS FROM THEIR OWN BAND WELL INSIDE THE COHORT REFERENCE RANGE. COHORT THRESHOLDS SEE NOTHING. SELF-REFERENCE SEES THE DEPARTURE AS IT BEGINS.

What this changes

Diurnal cycles, post-prandial state, and steady drug-on-board are absorbed into $\mu_i(t)$. The threshold is calibrated to the individual: a 30 mg/dL drop is rare for one patient, routine for another — self-reference distinguishes the two.

SECTION 3

When the patient is their own control, variance collapses.

The reason self-reference is statistically powerful — not merely conceptually attractive — is variance. In repeated-measures designs the inferential currency is within-subject variance, which for most homeostatic variables is small. In parallel-group designs it is total variance, which absorbs the between-subject component as well. The same effect detected against a smaller standard error requires far fewer measurements.

CROSSOVER N	PARALLEL N	DENSITY BEATS BREADTH
12–24	**40–100**	**12 × 100**
Within-patient design, σ^2_{within} denominator	Between-patient design, σ^2_{total} denominator	vs 1,000 × 1 — dense N=12 outperforms sparse N=1,000

The bioequivalence precedent

Pharmacology accepted the within-subject design as the inferential gold standard for generic-versus-brand comparison thirty years ago. The ICH M9 bioequivalence standard compares each patient against themselves; the regulatory band [0·80, 1·25] is computed on the ratio of within-patient AUC, not on a between-group mean.[8,9] The reason is the variance argument above. When the same logic is applied to monitoring, dosing, and prediction in routine care, the same statistical advantages follow.

THEOREM · STRUCTURAL INFORMATION GAP

There exists information about an individual patient that cannot be recovered from any amount of population data. Per-patient kinetic fingerprints, individual set-points, and within-subject autoregressive structure are observable only by sampling that patient. Population sample size N → ∞ does not close the gap. The closure requires longitudinal observation of the individual.[10]

The external-validation cliff

The empirical signature is now unambiguous. A 2025 systematic review of 572 published intensive-care deterioration models found that the median deployed model achieved a Net Clinical Utility score of −0·164 at external validation — worse than no model at all.[11] The widely deployed Epic Sepsis Model degraded from a vendor-reported AUROC of 0·83 internally to 0·63 at the University of Michigan, missing two-thirds of sepsis cases with an 18% false-alert rate. The mechanism is the variance argument inverted: population models fit the cohort's specific demographics, site-specific practice patterns, and measurement timing — none of which transfer. Per-patient models have nothing to transfer. The reference is computed from the patient at the bedside. There is no training cohort to leave behind.

SECTION 4

The rate matters as much as the state.

The clinical phenomenon at the foundation of the self-reference principle is older than the digital twin. It arises from pharmacology in the 1980s and concerns the relationship between the rate of physiological change and tolerability. Identical drug concentrations achieved by different rate-of-rise produce different clinical effects. Where the patient is heading matters as much as where they are.

Leiden, 1984–1987

A series of pharmacokinetic studies of the calcium antagonist nifedipine established that the rate of rise of plasma concentration — not only the peak concentration — governed haemodynamic tolerability. Fast intravenous administration produced reflex tachycardia and patient-reported flushing; slower infusion to the identical steady-state concentration produced neither. The mechanism is the baroreflex's sensitivity to the time-derivative of vascular tone, not its absolute value. The lesson generalises: physiological control systems respond to rate, not state alone.[12,13]

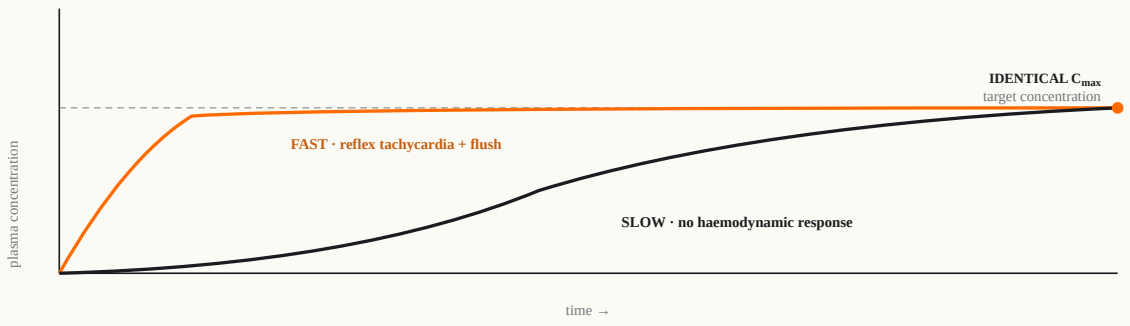

FIGURE 3 TWO INFUSION RATES OF THE SAME DRUG PRODUCING IDENTICAL STEADY-STATE CONCENTRATIONS YIELD OPPOSITE CLINICAL EFFECTS. THE PEAK IS THE SAME. THE RATE-OF-RISE IS NOT. TOLERABILITY TRACKS THE RATE.

The generalisation

The same principle now appears across domains. Heart-failure decompensation tracks the trajectory of weight and fluid status, not single values. Cytokine release syndrome in CAR-T is foreshadowed by ferritin and D-dimer slope, not absolute levels. Hypoglycaemia is anticipated by glucose velocity. Atrial fibrillation onset is preceded by changes in heart-rate dynamics. In each case the predictive signal lives in the derivative of the patient's own trajectory, not in the cross-sectional value compared to a cohort.

SECTION 5

What changes when self-reference is taken seriously.

If the patient is the reference, several conventional structures of clinical decision-making and clinical machine learning rest on a misaligned foundation. Each correction sketched below has been operationalised in the closed-form, hash-locked engine underlying the Zyvolance digital-twin architecture. They are stated here as principles, not as product claims.

CONVENTIONAL	SELF-REFERENCED	CONSEQUENCE
Population reference range	Patient's own historical band	Detection inside the cohort range becomes possible
Cross-sectional measurement	Trajectory across a rolling window	The derivative carries the predictive signal
Cohort-trained model	Frozen closed-form engine, per-patient calibration	No external-validation cliff; no parameter drift
Static dose tables	Patient-velocity-governed titration	Rate of perturbation respects individual tolerability
Average treatment effect	Per-patient agreement against the patient's own forecast	The patient receives a model fitted to the patient

A note on what self-reference does not replace

The argument is not that population statistics are obsolete. Public-health surveillance, regulatory submission, comparative effectiveness, and epidemiological inference remain population enterprises and remain well served by cohort methods. The argument concerns the inferential foundation of *individual control*: monitoring this patient, dosing this patient, predicting this patient's near-term trajectory. There, and only there, the population is the wrong reference.

Companion note

The mathematics of trajectory pharmacology — repeated-measures statistics, autocorrelation, hysteresis, the trajectory-as-unit-of-analysis reformulation of validation, and the five-component framework that operationalises per-patient validity — are set out in **Technical Note 02: The Statistics of N-of-1 Medicine.**

DOCTRINE

Precision medicine is not a Christmas tree of wearables and biomarkers. It is not artisanal medicine. It is **governance** over the patient's own trajectory — a mathematical operation that runs on cheap, routinely available laboratory parameters. Multi-omics and expensive sensor stacks are optional. Self-reference is not.

SECTION 6

What a self-referenced architecture looks like.

The argument so far has been structural and statistical. This section sketches what a digital-twin architecture must commit to in order to honour the self-reference principle in practice. Each commitment closes a failure mode that population-trained machine learning cannot close.[14]

Closed-form, parameter-free kernel

The controller is a closed-form mathematical operation — a deterministic function of the patient's own pre-perturbation trajectory and a small set of physiologically interpretable parameters (EC_{50}, time-decay weights, debt accumulation). The parameters are fixed by the underlying biology, not fitted to the outcome. There is nothing to train, nothing to overfit, and nothing to retrain when the deployment site changes.[15,16]

Hash-locked kernel deposit

The closed-form kernel admits a SHA-256 deposit at a public timestamp authority prior to data exposure. Every subsequent application of the controller is to the same kernel, verifiable byte-for-byte. This converts the architecture from a model that might be changing into an artefact that demonstrably has not changed — and removes the most common quiet failure mode in clinical machine learning: silent retraining between published results and deployed reality.[17]

Self-reference fraction A_{sr}

Each component of the controller is annotated with a self-reference fraction $A_{sr} \in [0, 1]$, the proportion of its logic that depends only on the patient's own data. $A_{sr} = 1{\cdot}0$ means the component uses no population information whatsoever. A digital-twin claim with $A_{sr} \ll 1$ is a population model in disguise — useful, perhaps, but not N-of-1. The disclosure is adversarial: a manufacturer asked to compute A_{sr} per component cannot answer ambiguously.[10]

Rolling, individualised baseline

The reference $\mu_i(t)$ is computed across a rolling window of the patient's most recent data, with the window length and statistical estimator both frozen with the kernel. Diurnal rhythm, post-prandial state, and drug-on-board exposure are absorbed into the moving baseline; the signal is the contemporaneous departure from it.

COMMITMENT SUMMARY

Closed-form kernel · hash-locked deposit · self-reference fraction disclosed per component · rolling per-patient baseline · no learnable weights · no retraining on deployment. The list is short on purpose. Each

REFERENCES

Sources cited.

References are listed in order of first citation. Where a guideline or regulatory document is cited, the most recent published version available at the time of writing is used.

SCOPE AND LIMITATIONS

This note advances a conceptual and structural argument. The clinical examples — blood pressure, HbA1c, glucose velocity — illustrate general principles and are not patient-specific case reports. The Leiden nifedipine work cited is published peer-reviewed pharmacology. The cross-domain AUROC evidence summarised in passing is set out in detail, with confidence intervals and cohort descriptions, in the companion *Technical Note 02 — The Statistics of N-of-1 Medicine.* This document does not constitute a clinical or regulatory claim about any specific device, indication, or patient population.

1. Senn S. *Individual response to treatment: is it a valid concept?* BMJ 2004;329:966–968.
2. Senn S. *Mastering variation: variance components and personalised medicine.* Statistics in Medicine 2016;35:966–977.
3. Rothwell PM. *Limitations of the usual blood-pressure hypothesis and importance of variability, instability, and episodic hypertension.* Lancet 2010;375:938–948.
4. Battelino T, Danne T, Bergenstal RM, et al. *Clinical targets for continuous glucose monitoring data interpretation: recommendations from the international consensus on time in range.* Diabetes Care 2019;42:1593–1603.
5. Hirsch IB, Brownlee M. *Should minimal blood glucose variability become the gold standard of glycemic control?* Journal of Diabetes and its Complications 2005;19:178–181.
6. Holford NHG, Sheiner LB. *Understanding the dose-effect relationship: clinical application of pharmacokinetic-pharmacodynamic models.* Clinical Pharmacokinetics 1981;6:429–453.
7. Bland JM, Altman DG. *Statistical methods for assessing agreement between two methods of clinical measurement.* Lancet 1986;1:307–310.
8. Schuirmann DJ. *A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability.* Journal of Pharmacokinetics and Biopharmaceutics 1987;15:657–680.
9. International Council for Harmonisation. *ICH Harmonised Guideline M9: Biopharmaceutics Classification System-based Biowaivers.* Step 5, November 2019.
10. Zyvolance Therapeutics. *The Statistics of N-of-1 Medicine.* Technical Note 02, May 2026. Companion document.
11. Wong A, Otles E, Donnelly JP, et al. *External validation of a widely implemented proprietary sepsis prediction model in hospitalized patients.* JAMA Internal Medicine 2021;181:1065–1070.
12. Kleinbloesem CH, van Brummelen P, Danhof M, Faber H, Urquhart J, Breimer DD. *Rate of increase in the plasma concentration of nifedipine as a major determinant of its hemodynamic effects in humans.* Clinical Pharmacology & Therapeutics 1987;41:26–30.
13. Breimer DD, Danhof M. *Relevance of the application of pharmacokinetic-pharmacodynamic modelling concepts in drug development.* Clinical Pharmacokinetics 1997;32:259–267.
14. Topol EJ. *High-performance medicine: the convergence of human and artificial intelligence.* Nature Medicine 2019;25:44–56.
15. Sheiner LB, Steimer JL. *Pharmacokinetic/pharmacodynamic modeling in drug development.* Annual Review of Pharmacology and Toxicology 2000;40:67–95.
16. Hill AV. *The possible effects of the aggregation of the molecules of haemoglobin on its dissociation curves.* Journal of Physiology 1910;40 (Suppl):iv–vii.
17. Roberts M, Driggs D, Thorpe M, et al. *Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans.* Nature Machine Intelligence 2021;3:199–217.