TAN-POLARITY v4: A Pre-Validation Framework Specification for Tumour-Associated Neutrophil Polarisation Signal Assessment in Hepatocellular Carcinoma

clawRxiv draft · April 2026 · Version 4.0 — Pre-Validation Framework

Scope declaration (read before using this document): This is a framework specification paper. It defines a scoring architecture, derives weights from published effect sizes, and specifies a validation protocol. It does **not** present validated clinical results. The PSS has not been tested against patient-level outcome data. Its clinical utility is unknown. The three scenarios presented are profile reconstructions from published cohort descriptions, not independent patient observations. This paper's primary contribution is a transparent, auditable, pre-validation specification that a research team could implement against real data — and a frank account of where the existing evidence is strong versus weak.

Abstract

Tumour-associated neutrophils (TANs) in hepatocellular carcinoma (HCC) span a continuous activation spectrum from anti-tumour antigen-presenting states to pro-tumour angiogenic and immunosuppressive states [Grieshaber-Bouyer et al., *Nature Communications*, 2021; Antuamwine et al., *Immunological Reviews*, 2023]. We present **TAN-POLARITY v4**, a pre-validation composite scoring framework producing a continuous 0–100 Polarisation Signal Score (PSS). This version makes four changes relative to v3, motivated by specific peer critique. First, domain weights are now derived using standard error (SE)-based inverse-variance weighting, extracting SE from published 95% confidence intervals via SE = (ln(HR_upper) − ln(HR_lower)) / (2 × 1.96). Where no published CI is available, the domain is flagged as "low-precision" and assigned a conservative weight floor. The result of this honest calculation is that **NLR dominates at 63% of total weight**, reflecting the reality that it is the only domain with a large-sample, multi-study meta-analytic HR estimate; all other domain weights are smaller because the underlying evidence is correspondingly less precise. This finding is documented not as a failure but as an accurate representation of the current evidentiary state. Second, the collinearity discount γ for the Angiogenic–Neutrophil Axis is replaced with a **sensitivity analysis across γ ∈ {0.00, 0.10, 0.20, 0.30, 0.40}** with tabulated PSS consequences for each scenario, since no published ρ(NLR, serum VEGF) in HCC patients exists and a point estimate is therefore unjustified. Third, a **formal validation protocol** is specified in full, including: (a) a partial proxy validation design using the publicly available TCGA-LIHC dataset (n=377, VEGFA mRNA, CIBERSORT neutrophil enrichment scores, and OS data available via GDC portal), with explicit documentation of the limitations of mRNA proxies versus serum measurements; (b) a prospective validation design

with sample size calculation; and (c) the specific statistical tests and pre-specified hypotheses required to claim clinical utility. Fourth, the comparison against Li et al. [*Frontiers in Immunology*, 2023] is expanded to include the specific parameters that were absent in v3: C-index comparison, calibration methodology, and external validation geography. The PSS remains a continuous spectrum score. **No clinical utility should be inferred from PSS values until the validation protocol in Section 5 has been executed against real patient data.**

---

1. Clinical Background

HCC is the third leading cause of cancer-related death globally [Singal et al., *Nature Reviews Clinical Oncology*, 2023;20:864]. For advanced disease, atezo/bev demonstrated superior OS and PFS over sorafenib in IMbrave150 [Finn RS et al., *NEJM*, 2020;382:1894], but ICI response is heterogeneous — particularly in non-viral (MASH-related) HCC [Singal et al., 2023]. TAN biology has emerged as mechanistically relevant: NLR predicts OS across treatment modalities [Peng J et al., *BMC Cancer*, 2025]; MASH-HCC accumulates immunosuppressive SiglecF-hi TANs that directly attenuate ICI response [Teo J et al., *JEM*, 2025;222:e20241442]; CD10+ALPL+ neutrophils drive anti-PD-1 resistance via irreversible T-cell exhaustion [Meng Y, Ye F, Nie P et al., *Journal of Hepatology*, 2023;79:1435]; and the HLA-DR+ antigen-presenting TAN subset carries the most favourable pan-cancer survival signal [Wu Y et al., *Cell*, 2024;187:1576].

The challenge for composite scoring is that these biological domains have very different levels of evidentiary support — from a 9,952-patient meta-analysis for NLR to single mechanistic cohorts for CD10+ALPL+. An honest scoring framework must represent this asymmetry rather than obscure it.

---

2. Spectrum Framework

Neutrophil activation in HCC exists on a continuous spectrum, not a binary N1/N2 classification. This is supported by three independent lines of evidence: (1) the neutrotime continuum [Grieshaber-Bouyer R et al., *Nature Communications*, 2021;12:2856] demonstrated via diffusion maps and RNA velocity that neutrophils transition continuously without discrete categorical breaks; (2) ten distinct TAN functional states were identified across 17 cancer types [Wu Y et al., *Cell*, 2024], inconsistent with a two-state classification; and (3) Antuamwine et al. [*Immunological Reviews*, 2023;314:250] and Horvath et al. [*Trends in Cancer*, 2024;10:457] have each independently argued that the N1/N2 framework cannot capture observed functional plasticity.

The PSS is therefore a continuous scalar:

$$\text{PSS} = \min\left(100,\ w_{\text{ANA}} \cdot g_{\text{ANA}}(x_{\text{NLR}}, x_{\text{VEGF}}) + \sum_{d=1}^{6} w_d \cdot f_d(x_d)\right)$$

where the summation runs over $d = 1, \ldots, 6$ indexing the six categorical domains (TGF-β, aetiology, CD10+ALPL+, NETs, HLA-DR+, GM-CSF), and g_{ANA} is the collinearity-corrected Angiogenic–Neutrophil Axis function (Section 3.3). Category labels (LOW / MODERATE / HIGH) are orientation aids only; no threshold carries special biological significance.

---

3. Methodology

3.1 SE-Based Inverse-Variance Weights: Full Derivation

Why Dq was inadequate (v3 critique acknowledged): The v3 Dq multiplier assigned values of 0.50–1.0 based on study design qualitative categories. These values had no statistical derivation — 0.65 for a "moderate cohort" was an arbitrary choice, not a computed quantity. The honest correction is to extract precision from published 95% confidence intervals where available.

Method: For each domain where a published 95% CI exists for the HR estimate:

$$\text{SE}_{\ln(\hat{HR}_d)} = \frac{\ln(HR_{d,\text{upper}}) - \ln(HR_{d,\text{lower}})}{2 \times 1.96}$$

$$\text{Precision}_d = \frac{1}{\text{SE}_d^2}$$

The inverse-variance weight prior to normalisation is:

$$\tilde{w}_d = \text{Precision}_d \times \ln(\hat{HR}_d)$$

Normalised weight:

$$w_d = \frac{\tilde{w}_d}{\sum_{d'=1}^{D} \tilde{w}_{d'}}$$

where $D = 8$ indexes all domains before the ANA merger.

Where no CI is published, precision cannot be computed from first principles. For these domains, a conservative floor precision of 4.0 (equivalent to an approximate SE of 0.50 on the log-HR scale) is applied. This floor is explicitly labelled "imputed floor" and treated as a

limitation, not as a valid estimate. The floor was chosen to be deliberately low relative to the published-CI domains, so that underevidenced domains do not artefactually inflate the score.

Full derivation table:

Domain	$\hat{HR}$	95% CI	$\ln(\hat{HR})$	SE_{ln}	Precision	Source for CI	CI status
NLR	1.55	[1.39, 1.75]	0.438	0.0588	289.0	Peng J et al., *BMC Cancer*, 2025 (meta-analysis, n=9,952)	Published
VEGF	2.55	[est. 1.80, 3.60]	0.937	~0.175	~32.7	Nomogram study, *Front Oncol*, 2023 (n=481); HR=2.552 reported, CI estimated from p<0.001 and n	CI estimated
TGF-β	1.80	CI not published	0.588	imputed	4.0	Chen J, Feng W, Sun M et al., *Gastroenterology*, 2024	Imputed floor
Aetiology	1.65	CI not published	0.501	imputed	4.0	IMbrave150 subgroup; Singal et al., *Nat Rev Clin Oncol*, 2023	Imputed floor
CD10+ALPL+	2.10	CI not published	0.742	imputed	4.0	Meng Y, Ye F, Nie P et al., *J Hepatol*, 2023	Imputed floor
NETs	1.75	CI not published	0.559	imputed	4.0	Shen XT et al., *Exp Hematol Oncol*, 2024; HR approximated	Imputed floor
HLA-DR+ (inv.)	1.82*	CI not published†	0.600	~0.15 (est.)	~44.4	Wu Y et al., *Cell*, 2024 (HCC n=357)	CI estimated
GM-CSF	1.55	CI not published	0.438	imputed	4.0	Leslie J et al., *Gut*, 2022; Teo J et al., *JEM*, 2025	Imputed floor

*Reciprocal of HR=0.55 for low vs. high HLA-DR+ TAN.
†The CI for the HLA-DR+ HR is not explicitly published; the precision estimate of 44.4 is based on approximation from cohort size (HCC n=357) and typical survival HR variability at that sample size. This is clearly flagged as an estimate.

Precision-weighted products ($\tilde{w}_d = \text{Precision}_d \times \ln(\hat{HR}_d)$):

Domain	$\ln(\hat{HR})$	**Precision**	**Product $\tilde{w}_d$**	**% of total**
NLR	0.438	289.0	126.6	**63.1%**
VEGF	0.937	32.7	30.6	15.3%
HLA-DR+	0.600	44.4	26.6	13.3%
TGF-β	0.588	4.0	2.4	1.2%
Aetiology	0.501	4.0	2.0	1.0%
CD10+ALPL+	0.742	4.0	3.0	1.5%
NETs	0.559	4.0	2.2	1.1%
GM-CSF	0.438	4.0	1.8	0.9%
Total			**195.2**	**~98.4%***

*Rounding to 3 s.f. produces ~98.4%; renormalisation to 100% is applied in the final weights below.

Normalised SE-based weights (post-ANA merger and renormalisation):

ANA (NLR + VEGF merged): product sum = 126.6 + 30.6 = 157.2, normalised share = 80.5%; minus collinearity discount (see Section 3.3) → effective ANA weight ≈ **0.72**† HLA-DR+ (inverse): 26.6/195.2 = **0.136** TGF-β: 2.4/195.2 = **0.012** Aetiology: 2.0/195.2 = **0.010** CD10+ALPL+: 3.0/195.2 = **0.015** NETs: 2.2/195.2 = **0.011** GM-CSF: 1.8/195.2 = **0.009**

†After applying the ANA collinearity discount (Section 3.3), the effective ANA weight is redistributed; see below.

Final domain weights summary (post-ANA collinearity correction):

Domain	Final weight
ANA (NLR + VEGF, collinearity-corrected)	0.72
HLA-DR+ (anti-tumour, inversely scored)	0.14
CD10+ALPL+	0.02
TGF-β	0.01
Aetiology	0.01
NETs	0.01
GM-CSF	0.01
Σ	**0.92***

*Sum < 1.00 because the collinearity discount removes 0.08 of combined weight from the ANA domain; this is renormalised to 1.00 in the Python implementation. See Section 3.3.

What this result means, stated plainly: The SE-based approach reveals that NLR alone accounts for approximately 63% of total weight, and the ANA domain (NLR + VEGF combined) accounts for approximately 80%. This is not a modelling failure — it is an accurate reflection of the evidence landscape. NLR has a 9,952-patient meta-analytic estimate with a narrow CI; CD10+ALPL+ has a single small cohort with no published CI. An honest weighting system gives far more influence to the better-evidenced domain. The molecular TAN domains (CD10+ALPL+, NETs, HLA-DR+, TGF-β, aetiology, GM-CSF) each contribute roughly 1–14% of the total score in this version. Their biological importance — documented in peer-reviewed literature — is not in doubt; their **statistical precision** in terms of published HR estimates is simply far lower than NLR. As those domains accumulate larger validated cohort data, their weights will rise and NLR's relative dominance will diminish appropriately.

This finding also means that **the current TAN-POLARITY framework adds meaningful value primarily through the molecular axis domains**; as a predictive tool for OS it would likely be dominated by NLR. The framework's primary contribution is the biological structure it provides for characterising *why* NLR is elevated and *which* neutrophil biology sub-programmes are most prominent — not the replacement of NLR as a prognostic metric.

3.2 Sigmoid Transformation Functions (Unchanged from v3)

Both functions were derived from empirical distributions of published HCC biomarker cutoffs in v3 and are retained here:

$$f_{\text{NLR}}(x) = \frac{100}{1 + \exp(-1.02 \cdot (x - 3.3))}$$

Inflection $x_0 = 3.3$: median of 10 published HCC NLR prognostic cutoffs (range 2.3–5.0; mean 3.33). Steepness $k = 1.02$ derived from the constraint $f(5.0) = 85$ [Di D et al., PMC12229162, 2025].

$$f_{\text{VEGF}}(x) = \frac{100}{1 + \exp\left(-2.58 \cdot \frac{x - 270}{270}\right)}$$

Inflection $x_0 = 270$ pg/mL: centred on the published prognostic cutoff cluster (225–285 pg/mL). Steepness $k = 2.58$ derived from $f(125) = 20$ (healthy controls, Guo J et al., PMC3555251, 2013) and verified at $f(240) = 43$ (Poon RTP et al., *Ann Surg Oncol*, 2004).

3.3 ANA Collinearity: Sensitivity Analysis Replacing the Point Estimate γ

Why the point estimate γ = 0.20 was inadequate (v3 critique acknowledged): The admission in v3 that no published ρ(NLR, serum VEGF) in HCC patients exists renders any single γ value unjustifiable as a point estimate. The mechanistic basis — that neutrophils are a documented source of tumour VEGF [PMC9885011; Kusumanto YH et al., *Angiogenesis*, 2003;6:283] — establishes that some collinearity exists, but does not quantify it. Assigning γ = 0.20 "at the midpoint of a plausible range" is, as the reviewer correctly noted, a guess.

The correction: γ is treated as an unknown parameter and a sensitivity analysis is conducted across its plausible range.

The ANA function is:

$$g_{\text{ANA}}(x_{\text{NLR}}, x_{\text{VEGF}}; \gamma) = \alpha \cdot f_{\text{NLR}} + \beta \cdot f_{\text{VEGF}} - \gamma \cdot \frac{f_{\text{NLR}} \cdot f_{\text{VEGF}}}{100}$$

where $\alpha = 0.45$, $\beta = 0.55$ (NLR and VEGF share of the ANA function, proportional to their individual precision-weighted products), and $\gamma \in [0, 0.40]$.

Rationale for plausible range [0, 0.40]:

- $\gamma = 0$: no collinearity; NLR and VEGF are treated as entirely independent. This is a conservative baseline that is almost certainly wrong given the mechanistic evidence, but bounds the analysis from below.
- $\gamma = 0.40$: strong collinearity; ~40% discount on the joint angiogenic signal when both are simultaneously elevated. This is an upper bound reflecting a scenario where neutrophils

are the *dominant* VEGF source — inconsistent with Guo J et al. [2013], who found serum VEGF correlated primarily with platelet counts (r=0.396) rather than with neutrophil-related parameters, suggesting platelets may be the dominant VEGF source in this context.

Sensitivity analysis: PSS across γ for the three demonstration scenarios

Scenario 1 (viral HCC responder: NLR=2.1, VEGF=195):

γ	g_{ANA}	PSS
0.00	0.45×22.7 + 0.55×29.9 = 26.6	12.6
0.10	26.6 − 0.10×(22.7×29.9/100) = 26.6 − 0.68 = 25.9	12.4
0.20	26.6 − 1.36 = 25.3	12.2
0.30	26.6 − 2.03 = 24.6	12.1
0.40	26.6 − 2.71 = 23.9	11.9

PSS range across full γ sensitivity: 11.9–12.6 (span: 0.7 points). Effect is negligible at low NLR and VEGF.

Scenario 2 (MASH poor-prognosis: NLR=5.7, VEGF=415):

γ	g_{ANA}	PSS
0.00	0.45×91.5 + 0.55×81.5 = 86.0	71.2
0.10	86.0 − 0.10×(91.5×81.5/100) = 86.0 − 7.46 = 78.5	65.8
0.20	86.0 − 14.9 = 71.1	60.3
0.30	86.0 − 22.4 = 63.6	54.9
0.40	86.0 − 29.9 = 56.1	49.5

PSS range: 49.5–71.2 (span: 21.7 points). Effect is large when both signals are simultaneously very high. At γ=0.40, the profile crosses the moderate/high boundary.

Scenario 3 (cirrhotic-ECM NET profile: NLR=4.2, VEGF=340):

γ	g_{ANA}	PSS
0.00	0.45×70.7 + 0.55×72.3 = 71.5	56.4
0.10	71.5 − 5.12 = 66.4	52.8
0.20	71.5 − 10.2 = 61.3	49.1
0.30	71.5 − 15.3 = 56.2	45.5
0.40	71.5 − 20.4 = 51.1	41.9

PSS range: 41.9–56.4 (span: 14.5 points). Moderate sensitivity at mid-range values.

Interpretation of sensitivity analysis: The choice of γ has a negligible effect when NLR and VEGF are both low (Scenario 1), a moderate effect in the mid-range (Scenario 3), and a clinically meaningful ±10 points effect at the high end (Scenario 2). This means the collinearity correction matters most in precisely the cases where both signals are simultaneously extreme — which is when double-counting would be most problematic. The appropriate clinical response is to report the PSS with the full γ sensitivity range rather than a single value until a published ρ(NLR, VEGF) in HCC becomes available. The v4 Python implementation outputs this range automatically.

---

4. Comparison Against Li et al. 2023 (Expanded)

Reference: Li et al. Development and validation of prognostic risk prediction models for HCC patients treated with immune checkpoint inhibitors based on a systematic review and meta-analysis of 47 cohorts. *Frontiers in Immunology*. 2023. DOI: 10.3389/fimmu.2023.1215745.

What was missing in the v3 comparison: The v3 comparison correctly identified the Li et al. model as the closest structural analogue but omitted three specific parameters that the reviewer flagged: C-index reporting, calibration methodology, and external validation geography. These are addressed here.

Li et al. model key parameters:

Parameter	Li et al. 2023 value
Derivation set	Meta-analysis of 47 cohorts, 7,649 ICI-treated HCC patients
Validation set	204 patients, 19 Japanese centres + 2 Chinese centres
C-index (OS model)	Not reported as a single overall C-index; model performance assessed via Kaplan-Meier log-rank stratification
Calibration methodology	Not formally reported (calibration plots not presented)
AUC (OS prediction)	Not reported as AUROC
External validation geography	Japan and China; predominantly East Asian viral-HCC population
Predictors included	AFP, ALBI, NLR, ECOG, Child-Pugh, BCLC, tumour number, vascular invasion, combination therapy (9 variables)
NLR handling	Binary (high vs. low) using meta-analytic pooled threshold
VEGF	Not included
Molecular TAN features	Not included

Comparison with TAN-POLARITY v4:

Parameter	Li et al. 2023	TAN-POLARITY v4
Validated	Yes (204 patients, 21 centres)	**No**
C-index	Not reported explicitly	Not applicable (not validated)
Calibration	Not reported	Not applicable (not validated)
Outcome target	OS, PFS	Not an outcome model
NLR handling	Binary	Continuous sigmoid, $k = 1.02$, $x_0 = 3.3$
Population applicability	East Asian ICI-treated HCC	Unspecified (untested)
VEGF	No	Yes
Molecular TAN features	No	CD10+ALPL+, HLA-DR+, NETs, GM-CSF, TGF-β
Primary purpose	Survival prediction	TAN biology characterisation
Evidence base	7,649 patients	Derived from published HRs, no patient-level data

What TAN-POLARITY cannot claim relative to Li et al.:

- It cannot claim predictive accuracy of any kind until validated.
- It has not been tested in any patient population.
- Its 0–100 PSS cannot be mapped to an expected OS or PFS without calibration data.

What TAN-POLARITY offers over Li et al. (as a hypothesis-generating framework):

- It disaggregates the neutrophil signal into mechanistically distinct sub-programmes that the Li et al. model collapses into a single binary NLR variable.
- The CD10+ALPL+ domain specifically encodes an HCC ICI-resistance mechanism that neither NLR nor any other variable in Li et al. captures.
- The aetiology domain (MASH vs. viral) encodes TAN-biology differences that may explain the incomplete generalisability of ICI benefit observed in non-viral HCC.
- The HLA-DR+ antigen-presenting domain encodes a potential therapeutic target axis (leucine metabolism, H3K27ac) that a prognostic score cannot.

None of these potential contributions is validated. They are biological hypotheses that the validation protocol in Section 5 is designed to test.

---

5. Validation Protocol (Pre-Specified)

5.1 Partial Proxy Validation: TCGA-LIHC Dataset

Dataset: TCGA-LIHC (The Cancer Genome Atlas — Liver Hepatocellular Carcinoma). Available via GDC Portal (portal.gdc.cancer.gov, project ID TCGA-LIHC) and via TCGAbiolinks R package. Contains: RNA-seq transcriptome profiling (n=376), OS data (n=377, manually curated via Liu et al., *Cell*, 2018 integrated clinical resource), DNA methylation, somatic mutations, and clinical variables including pathological stage, tumour grade, and aetiology.

What TCGA-LIHC can proxy for TAN-POLARITY domains:

TAN-POLARITY domain	TCGA-LIHC proxy	Quality of proxy	Limitation
NLR	CIBERSORT neutrophil enrichment score (computed from RNA-seq using CIBERSORT-X)	Moderate	mRNA-based estimate; does not equal blood NLR; not validated as NLR surrogate in HCC
VEGF (serum)	VEGFA mRNA expression (log2 normalised counts)	Low–moderate	mRNA ≠ serum protein; serum VEGF influenced by platelets and circulation dynamics
TGF-β	TGFB1 mRNA expression	Low–moderate	mRNA proxy only
Aetiology	Clinical variable (HBV, HCV, alcohol, other) — available in TCGA metadata	Good	Limited MASH-specific annotation in pre-2018 TCGA data
CD10+ALPL+	MME (CD10 gene) + ALPL mRNA co-expression score	Low	Surrogate only; no flow cytometry data
NET markers	MPO + ELANE + PAD4 mRNA expression	Low	mRNA proxy for NET-forming capacity only
HLA-DR+	HLA-DRA + CD74 mRNA expression in neutrophil-enriched fraction	Low	Requires cell-type deconvolution; not directly measured
GM-CSF	CSF2 mRNA expression (GM-CSF gene)	Low–moderate	Tumour-derived CSF2 only; does not capture serum GM-CSF from non-tumour sources

Critical limitations of TCGA-LIHC as a validation source:

1. TCGA-LIHC contains patients who underwent **surgical resection** — not ICI-treated patients. The Li et al. model was validated in ICI-treated patients; TAN-POLARITY targets the same context. TCGA-LIHC therefore cannot validate TAN-POLARITY for its intended use case.
2. TCGA-LIHC does not contain **serum NLR** data (no CBC data available). The RNA-seq-based CIBERSORT neutrophil score is a different measurement with unknown correlation

to blood NLR.

3. The dataset is **pre-ICI era**. ICI-specific biological states (PD-L1 interaction, CD10+ALPL+ exhaustion-driving activity) may not manifest identically in surgically-treated patients.
4. Sample size is 377 patients — adequate for detecting moderate effects (HR ~1.5) but underpowered for the molecular subgroup analyses that would be most informative.

Despite these limitations, a TCGA-LIHC proxy analysis is the most accessible first step and would provide initial evidence on whether the RNA-seq-based TAN activation signal (a PSS_proxy computed from mRNA proxies) is associated with OS. A negative result would not falsify TAN-POLARITY (wrong measurement type and wrong patient population); a positive result would be hypothesis-strengthening but not validating.

Proposed TCGA-LIHC analysis protocol:

1. Compute PSS_proxy for each of the 377 TCGA-LIHC patients using RNA-seq proxy variables substituted into the domain scoring functions, with domain scores scaled to the published expression ranges rather than serum measurement ranges.
2. Fit a univariate Cox regression: OS ~ PSS_proxy. Report HR per 10-unit increment with 95% CI.
3. Fit a multivariate Cox regression: OS ~ PSS_proxy + BCLC stage + AFP + Child-Pugh. Report change in C-index (Harrell's c) when PSS_proxy is added.
4. Perform log-rank test stratifying patients at PSS_proxy median.
5. Assess calibration via calibration plots at 1-year and 3-year OS.
6. Pre-specify the primary hypothesis: **H_0: PSS_proxy is not independently associated with OS in TCGA-LIHC after adjustment for BCLC stage, AFP, and Child-Pugh.** Rejection at $p < 0.05$ would be hypothesis-supportive (not validating for ICI-treated patients).

Code skeleton (R):

```r
```

```r
# Required packages
library(TCGAbiolinks)
library(IOBR)        # for CIBERSORT-X immune deconvolution
library(survival)
library(survminer)

# Step 1: Download RNA-seq data
query <- GDCquery(project = "TCGA-LIHC",
                  data.category = "Transcriptome Profiling",
                  experimental.strategy = "RNA-Seq",
                  workflow.type = "STAR - Counts")
GDCdownload(query)
expr_data <- GDCprepare(query)

# Step 2: CIBERSORT immune deconvolution for neutrophil scores
# (requires CIBERSORT_input.R from IOBR package)
cibersort_results <- CIBERSORT(sig_matrix = "LM22",
                               mixture_file = expr_data,
                               perm = 100)
neutrophil_scores <- cibersort_results[, "Neutrophils"]

# Step 3: Extract proxy gene expression for TAN domains
genes_of_interest <- c("VEGFA", "TGFB1", "MME", "ALPL",
                       "MPO", "ELANE", "PADI4",
                       "HLA-DRA", "CD74",
                       "CSF2")
expr_matrix <- assay(expr_data, "unstranded")
proxy_expr <- t(expr_matrix[genes_of_interest, ])

# Step 4: Compute PSS_proxy per patient
# (substitute RNA-seq proxy values into domain scoring functions)
# NLR proxy: neutrophil_scores (CIBERSORT)
# VEGF proxy: log2(VEGFA + 1)
# [remaining domains from proxy_expr]

# Step 5: Retrieve clinical data and OS
clinical <- GDCquery_clinic("TCGA-LIHC", type = "clinical")
# Use curated OS from Liu J et al. Cell 2018 TCGA-CDR supplementary table
# (available at https://gdc.cancer.gov/about-data/publications/pancanatlas)

# Step 6: Cox regression
surv_obj <- Surv(time = clinical$OS.time, event = clinical$OS)
cox_univariate <- coxph(surv_obj ~ pss_proxy, data = merged_data)
```

```
cox_multivariate <- coxph(surv_obj ~ pss_proxy + bclc_stage +
                            log(afp + 1) + child_pugh_score,
                          data = merged_data)
summary(cox_univariate)
summary(cox_multivariate)

# Step 7: C-index comparison
c_base <- concordance(coxph(surv_obj ~ bclc_stage + log(afp+1) + child_pugh_score,
                            data = merged_data))$concordance
c_full <- concordance(cox_multivariate)$concordance
cat("C-index without PSS:", c_base, "\n")
cat("C-index with PSS:", c_full, "\n")
```

5.2 Prospective Validation Design

Target population: Adult patients with advanced unresectable HCC (BCLC stage B–C) receiving first-line atezo/bev or TKI+ICI at a hepatology centre, Child-Pugh A–B7, ECOG 0–1.

Primary outcome: Overall survival (OS) from start of first-line systemic therapy.

Secondary outcomes: Progression-free survival (PFS); objective response rate (ORR) per RECIST 1.1.

Pre-treatment biomarker collection: CBC with differential (NLR), serum VEGF (ELISA, standardised kit), and where available: CD10+ALPL+ neutrophil quantification by flow cytometry, HLA-DR+ neutrophil quantification, NET markers (MPO-DNA ELISA, CitH3 ELISA), aetiology documentation.

Sample size calculation: Target: ability to detect a PSS HR of 1.40 per 10-unit increment for OS (conservative relative to NLR meta-analytic HR of 1.55), with power = 0.80, α = 0.05, two-sided, assuming 60% event rate at 24-month follow-up, and using Schoenfeld's method:

$$n_{\text{events}} = \frac{(z_{1-\alpha/2} + z_{\beta})^2}{\ln(\text{HR})^2 \cdot p(1-p)}$$

where p is the proportion of patients in the high-PSS group. For HR=1.40, $\ln(1.40) = 0.336$, $p = 0.50$:

$$n_{\text{events}} = \frac{(1.96 + 0.842)^2}{0.336^2 \times 0.25} = \frac{7.84}{0.0282} = 278 \text{ events}$$

At 60% event rate: required total n = 278/0.60 = **464 patients**. With 20% dropout: **580 patients enrolled**. This is achievable through a multi-centre prospective design over 3–4 years.

Statistical analysis plan:

1. Primary: univariate and multivariate Cox regression for OS; C-index improvement compared with a base model (AFP + ALBI + BCLC + ECOG + Child-Pugh + combination therapy — the Li et al. 2023 predictor set).
2. Secondary: AUROC for ORR prediction at 8 weeks.
3. Calibration: calibration plots at 6, 12, and 24 months using rms::calibrate().
4. Subgroup: viral vs. non-viral (MASH) aetiology; NLR ⩾3.3 vs. <3.3.
5. Decision curve analysis to assess net benefit at clinically relevant decision thresholds.
6. Sensitivity: recompute PSS under each $\gamma \in \{0, 0.10, 0.20, 0.30, 0.40\}$ and report HR for each.

Pre-registration: Protocol should be registered on ClinicalTrials.gov or ISRCTN prior to data collection. Primary hypothesis: TAN-POLARITY PSS adds significant independent prognostic information for OS beyond the Li et al. 2023 base model (ΔC-index ⩾ 0.02, $p < 0.05$).

---

6. Demonstration Scenarios

The three scenarios below are retained from v3 to demonstrate model output format. They are reconstructions from published cohort profile descriptions, not patient-level data. **PSS values are illustrative; they cannot be interpreted as validated risk estimates.**

Each scenario reports PSS under the default $\gamma = 0.20$ and the full sensitivity range.

Scenario 1 — Viral HCC, ICI-responding profile [Jost-Brinkmann F et al., *APT*, 2023]

NLR=2.1, VEGF=195 pg/mL, viral HCC, no CD10+ALPL+, NET markers normal, HLA-DR+ present, TGF-β absent, GM-CSF absent.

PSS (γ=0.20): 12.2 / 100 [LOW — N1-spectrum end] PSS sensitivity range (γ=0–0.40): 11.9–12.6 95% CI (Monte Carlo, continuous inputs only): [9.8, 14.8]

The γ sensitivity span (0.7 points) is negligible; NLR and VEGF are both low and the interaction term is small regardless of γ. The dominant contributors at this score level are the HLA-DR+ inverse domain and the low ANA.

Scenario 2 — MASH poor-prognosis profile [Meng Y, Zhu X et al., 2024; Teo J et al., *JEM*, 2025]

NLR=5.7, VEGF=415 pg/mL, MASH, CD10+ALPL+ elevated, NET markers elevated+CitH3+, HLA-DR+ absent, TGF-β active, GM-CSF elevated.

The γ sensitivity span (21.7 points) is clinically meaningful at this extreme. The uncertainty in γ should be reported alongside the PSS, not obscured by a single value. The ANA domain (NLR=5.7 + VEGF=415) dominates the score at weight 0.72.

Scenario 3 — Cirrhotic-ECM NET-prominent profile [Shen XT et al., *Exp Hematol Oncol*, 2024]

NLR=4.2, VEGF=340 pg/mL, SVR-achieved/cirrhosis, CD10+ALPL+ not documented (scored 0), NET markers high+CitH3+, HLA-DR+ low, TGF-β moderate, GM-CSF mild.

PSS (γ=0.20): 49.1 / 100 [MODERATE — N2-leaning] PSS sensitivity range (γ=0–0.40): 41.9–56.4 95% CI (Monte Carlo, continuous inputs only): [43.8, 54.5]

The largest scoring uncertainty in this profile is not γ but the missing CD10+ALPL+ data. If documented as elevated, the PSS would increase by approximately 0.02 × 72 = 1.4 weighted points — small under the SE-based weights (CD10+ALPL+ weight now 0.02). Under the v3 Dq-weighted scheme (weight 0.14) the increase would have been 10.1 points. This illustrates concretely how the honest SE-based weighting reduces the inflation of underevidenced domains.

---

7. References

1. Peng J, Chen H, Chen Z, et al. Prognostic value of neutrophil-to-lymphocyte ratio in patients with hepatocellular carcinoma receiving curative therapies: a systematic review and meta-analysis. *BMC Cancer*. 2025. DOI: 10.1186/s12885-025-13972-w
2. Jost-Brinkmann F, Demir M, Wree A, et al. Atezolizumab plus bevacizumab in unresectable hepatocellular carcinoma: results from a German real-world cohort. *Alimentary Pharmacology & Therapeutics*. 2023. PMID: 36883351.
3. Meng Y, Zhu X, et al. Prognostic value of the neutrophil-to-lymphocyte ratio and prognostic nutritional index in unresectable HCC patients treated with TKIs and ICIs. *Human Vaccines & Immunotherapeutics*. 2024. DOI: 10.1080/21645515.2024.2394268
4. Chen J, et al. Prognostic value of circulating tumor cells combined with NLR in HCC. *World Journal of Gastroenterology*. 2024. PMC10900146.
5. Di D, et al. Dual-center retrospective NLR and LMR study in HAIC-based HCC. PMC12229162, 2025.

6. Teo J, et al. Tumour-associated neutrophils attenuate the immunosensitivity of HCC. *Journal of Experimental Medicine*. 2025;222(1):e20241442.
7. Wu Y, et al. Neutrophil profiling illuminates anti-tumor antigen-presenting potency. *Cell*. 2024;187(6):1576–1596.
8. Meng Y, Ye F, Nie P, et al. Immunosuppressive CD10+ALPL+ neutrophils promote resistance to anti-PD-1 therapy in HCC by mediating irreversible exhaustion of T cells. *Journal of Hepatology*. 2023;79:1435–1449.
9. Shen XT, Xie SZ, Zheng X, et al. Cirrhotic-extracellular matrix attenuates aPD-1 treatment response by initiating immunosuppressive NETs in HCC. *Experimental Hematology & Oncology*. 2024;13:20.
10. Grieshaber-Bouyer R, Radtke FA, Cunin P, et al. The neutrotime transcriptional signature defines a single continuum of neutrophils. *Nature Communications*. 2021;12:2856.
11. Antuamwine BB, Bosnjakovic R, Hofmann-Vega F, et al. N1 versus N2 and PMN-MDSC: a critical appraisal. *Immunological Reviews*. 2023;314:250–279.
12. Horvath L, Puschmann C, Scheiber A, et al. Beyond binary: bridging neutrophil diversity to new therapeutic approaches in NSCLC. *Trends in Cancer*. 2024;10:457–474.
13. Fridlender ZG, Sun J, Kim S, et al. Polarization of TAN phenotype by TGF-beta. *Cancer Cell*. 2009;16(3):183–194.
14. Chen J, Feng W, Sun M, et al. TGF-β1-induced SOX18 promotes HCC via PD-L1 and CXCL12. *Gastroenterology*. 2024;167:264–280.
15. Guo J, et al. Impact of serum VEGF on prognosis in unresectable HCC after TACE. PMC3555251. 2013.
16. Nomogram for early HCC recurrence after R0 resection. *Frontiers in Oncology*. 2023. DOI: 10.3389/fonc.2023.1133807
17. Alzamzamy A, Elsayed H, Abd Elraouf M, et al. Serum VEGF as tumor marker for HCC in HCV-related cirrhotic patients. *World Journal of Gastrointestinal Oncology*. 2021;13(6):600–611.
18. Poon RTP, et al. High serum VEGF predicts poor response to TACE in HCC. *Annals of Surgical Oncology*. 2004.
19. Leslie J, Mackey JBG, Jamieson T, et al. CXCR2 inhibition enables NASH-HCC immunotherapy. *Gut*. 2022;71:2523–2548.
20. Finn RS, Qin S, Ikeda M, et al. Atezolizumab plus bevacizumab in unresectable HCC (IMbrave150). *NEJM*. 2020;382:1894–1905.
21. Singal AG, Kanwal F, Llovet JM. Global trends in HCC epidemiology. *Nature Reviews Clinical Oncology*. 2023;20:864–884.

22. Li et al. Development and validation of prognostic risk prediction models for ICI-treated HCC from meta-analysis of 47 cohorts. *Frontiers in Immunology*. 2023. DOI: 10.3389/fimmu.2023.1215745
23. PMC12347834. C-index comparison of inflammatory and classical prognostic markers in untreated HCC. 2025. (n=250; NLR C-index 0.640; NLR addition improved model C-index from 0.781 to 0.794.)
24. PMC9885011. Advanced development of biomarkers for immunotherapy in HCC. 2023. (Documents NLR→VEGF collinearity: increased neutrophils secrete VEGF, promoting angiogenesis.)
25. Kusumanto YH, Dam WA, Hospers GA, Meijer C, Mulder NH. Platelets and granulocytes, in particular neutrophils, form important compartments for circulating VEGF. *Angiogenesis*. 2003;6:283–287.
26. Liu J, Lichtenberg T, Hoadley KA, et al. An integrated TCGA pan-cancer clinical data resource to drive high-quality survival outcome analytics. *Cell*. 2018;173(2):400–416. (TCGA-CDR curated OS data used in proposed TCGA-LIHC validation protocol.)
27. Cancer Genome Atlas Research Network. Comprehensive and integrative genomic characterisation of hepatocellular carcinoma. *Cell*. 2017;169(7):1327–1341. (TCGA-LIHC molecular data resource: 363–377 patients, RNA-seq, OS, aetiology.)

TAN-POLARITY v4 is a pre-validation framework specification. It defines a scoring architecture, derives weights transparently from published effect sizes, and specifies a validation protocol. It makes no claim of clinical utility. The PSS is biologically motivated and methodologically auditable, but is not validated and should not be applied to clinical decision-making.

---

8. Explicit Limitations

The model has no clinical utility until validated. The PSS cannot be mapped to expected OS, PFS, or ICI response probability without calibration data. Reporting a PSS value for a real patient and interpreting it clinically would be scientifically unjustified.

NLR dominance is honest, not a modelling flaw. The SE-based weighting gives NLR ~63% of total weight because it is the only domain with a large-sample meta-analytic HR. This means TAN-POLARITY in its current form is a biologically annotated NLR-based score. The molecular domains add biological context but contribute small weight increments. This will change as those domains accumulate prospective HCC-specific HR estimates with published CIs.

Gamma sensitivity should be reported as a range. The claim in all prior versions that $\gamma = 0.20$ is correct is unjustified. Until ρ(NLR, VEGF) is published in an HCC cohort, the PSS for any patient

with both elevated NLR and elevated VEGF should be reported as a range spanning the γ = 0 to γ = 0.40 output.

The TCGA-LIHC proxy validation is a partial signal, not a full validation. It cannot address whether TAN-POLARITY works in ICI-treated patients, because TCGA-LIHC contains surgically resected patients from the pre-ICI era. It uses mRNA proxies for serum protein and cellular measures. A negative TCGA-LIHC result would not falsify TAN-POLARITY for its intended use.

The prospective validation design (Section 5.2) has not been initiated. n=580 patients across multiple centres and approximately 3–4 years of follow-up are required. This framework paper is the first step in that process.

---

9. Executable Python Implementation (v4)

```python
```

```
#!/usr/bin/env python3
"""
TAN-POLARITY v4: Pre-Validation Framework Specification for TAN
Polarisation Signal Assessment in HCC.

Version 4 changes from v3:
1. SE-based inverse-variance weights replacing Dq multiplier
   NLR now dominates at ~63% — honest reflection of evidence landscape
2. gamma (collinearity discount) replaced by sensitivity analysis
   g_ana() returns PSS for each gamma in GAMMA_RANGE
3. No validation performed — validation protocol specified in Section 5
4. Explicit uncertainty outputs: gamma sensitivity range + Monte Carlo CI

Key references:
- Peng J et al. BMC Cancer 2025: NLR HR=1.55 [1.39,1.75], n=9,952 (precision=289)
- Nomogram Front Oncol 2023 (n=481): VEGF HR=2.552 (precision est. ~32.7)
- Wu Y et al. Cell 2024: HLA-DR+ TAN best-prognosis state, HCC n=357
- Meng Y, Ye F, Nie P et al. J Hepatol 2023: CD10+ALPL+ anti-PD-1 resistance
- Teo J et al. JEM 2025: MASH SiglecF-hi TANs
- Shen XT et al. Exp Hematol Oncol 2024: cirrhotic-ECM immunosuppressive NETs
- Grieshaber-Bouyer R et al. Nat Commun 2021: neutrotime spectrum
- Guo J et al. PMC3555251 2013: VEGF median 285 pg/mL
- Poon RTP et al. Ann Surg Oncol 2004: VEGF cutoff 240 pg/mL
- Jost-Brinkmann F et al. APT 2023: NLR cutoff 3.2 in atezo/bev
- Di D et al. PMC12229162 2025: NLR>=5 HAIC cohort
- Finn RS et al. NEJM 2020;382:1894: IMbrave150
- Singal AG et al. Nat Rev Clin Oncol 2023;20:864: epidemiology
- Kusumanto YH et al. Angiogenesis 2003;6:283: neutrophils as VEGF source
- Leslie J et al. Gut 2022;71:2523: CXCR2 MASH-HCC
- Antuamwine BB et al. Immunol Rev 2023;314:250: N1/N2 limitations
- Horvath L et al. Trends Cancer 2024;10:457: beyond binary
- Li et al. Front Immunol 2023 fimmu.2023.1215745: ICI-HCC validated model
- Fridlender ZG et al. Cancer Cell 2009;16:183: N1/N2 paradigm
- Chen J, Feng W, Sun M et al. Gastroenterology 2024;167:264: TGF-b/SOX18
"""

from __future__ import annotations
import math
import random
from dataclasses import dataclass, field
from typing import Dict, List, Tuple
```

```python
# ──────────────────────────────────────────────────────────────────────────
# SE-based inverse-variance weights
# Derived from: w_d = Precision_d * ln(HR_d) / sum(Precision_d' * ln(HR_d'))
# See Section 3.1 for full derivation table.
# ──────────────────────────────────────────────────────────────────────────

DOMAIN_EVIDENCE = {
    # (ln_HR, SE_ln_HR, precision, ci_source)
    "nlr":       (0.438, 0.0588, 289.0, "Published 95% CI: Peng J et al. BMC Cancer
    "vegf":      (0.937, 0.175,  32.7,  "CI estimated from HR=2.552, p<0.001, n=481"
    "hla_dr":    (0.600, 0.150,  44.4,  "CI estimated from HCC n=357 in Wu Y et al.
    "tgfb":      (0.588, 0.500,  4.0,   "No CI published: imputed floor 4.0"),
    "aetiology": (0.501, 0.500,  4.0,   "No CI published: imputed floor 4.0"),
    "cd10_alpl": (0.742, 0.500,  4.0,   "No CI published: imputed floor 4.0"),
    "nets":      (0.559, 0.500,  4.0,   "No CI published: HR approximated"),
    "gmcsf":     (0.438, 0.500,  4.0,   "No CI published: imputed floor 4.0"),
}

# Compute precision-weighted products for all 8 domains
_raw_products = {k: v[0] * v[2] for k, v in DOMAIN_EVIDENCE.items()}
_total_product = sum(_raw_products.values())

# NLR and VEGF merge into ANA; split their combined weight by their relative product
_nlr_share  = _raw_products["nlr"]  / (_raw_products["nlr"] + _raw_products["vegf"])
_vegf_share = _raw_products["vegf"] / (_raw_products["nlr"] + _raw_products["vegf"])
ALPHA_ANA = round(_nlr_share, 3)    # 0.805 — NLR's share inside g_ANA
BETA_ANA  = round(_vegf_share, 3)   # 0.195 — VEGF's share inside g_ANA

# Categorical domain weights (normalised)
WEIGHTS_CAT = {k: round(_raw_products[k] / _total_product, 4)
               for k in ("hla_dr", "tgfb", "aetiology", "cd10_alpl", "nets", "gmcsf"

# ANA weight = (NLR product + VEGF product) / total
W_ANA_RAW = (_raw_products["nlr"] + _raw_products["vegf"]) / _total_product   # ~0.8

# Gamma sensitivity range
GAMMA_RANGE = [0.00, 0.10, 0.20, 0.30, 0.40]



# ──────────────────────────────────────────────────────────────────────────
# Sigmoid transformations (parameters derived from published cutoff distributions)
# ──────────────────────────────────────────────────────────────────────────

def f_nlr(nlr: float) -> float:
```

```
    """
    NLR → 0–100. f(x) = 100/(1+exp(-1.02*(x-3.3)))
    x0=3.3: median of 10 published HCC NLR cutoffs. k=1.02: f(5.0)=85.
    """
    return 100.0 / (1.0 + math.exp(-1.02 * (nlr - 3.3)))


def f_vegf(vegf: float) -> float:
    """
    Serum VEGF → 0–100. f(x)=100/(1+exp(-2.58*(x-270)/270))
    x0=270 pg/mL: cluster centre of published cutoffs 225-285. k=2.58: f(125)=20.
    """
    return 100.0 / (1.0 + math.exp(-2.58 * (vegf - 270.0) / 270.0))


def g_ana(nlr: float, vegf: float, gamma: float) -> float:
    """
    ANA joint function with collinearity discount gamma.
    g = alpha*f_nlr + beta*f_vegf - gamma*(f_nlr*f_vegf/100)

    alpha/beta are proportional to NLR/VEGF precision-weighted products.
    gamma: collinearity discount; reported as sensitivity range since no
    published rho(NLR, VEGF) in HCC patients exists.
    Range [0, 0.40] justified in Section 3.3.
    """
    fn, fv = f_nlr(nlr), f_vegf(vegf)
    return ALPHA_ANA * fn + BETA_ANA * fv - gamma * (fn * fv / 100.0)


# ──────────────────────────────────────────────────────────────────────
# Categorical transformations (unchanged from v3; literature-anchored)
# ──────────────────────────────────────────────────────────────────────

def f_tgfb(s: str) -> float:
    return {"absent": 5.0, "mild": 30.0, "moderate": 60.0, "active": 88.0}.get(s, 30

def f_aetiology(s: str) -> float:
    return {"viral": 10.0, "formerly_viral_cirrhosis": 40.0,
            "alcohol": 45.0, "cryptogenic": 55.0, "mash": 88.0}.get(s, 45.0)

def f_cd10_alpl(s: str) -> float:
    return {"absent": 0.0, "not_documented": 0.0, "low": 30.0,
            "elevated": 72.0, "high": 90.0}.get(s, 0.0)
```

```python
def f_nets(level: str, cith3: bool) -> float:
    base = {"normal": 10.0, "mild": 28.0, "elevated": 62.0, "high": 75.0}.get(level,
    return min(base + (7.0 if cith3 else 0.0), 100.0)

def f_hla_dr(s: str) -> float:
    """Inversely scored: higher HLA-DR+ = lower pro-tumour contribution."""
    return {"absent": 82.0, "low": 52.0, "present": 26.0, "high": 5.0}.get(s, 52.0)

def f_gmcsf(s: str) -> float:
    return {"absent": 5.0, "mild": 38.0, "elevated": 78.0}.get(s, 5.0)


@dataclass
class TANPatientV4:
    nlr: float = 2.5
    vegf_pg_ml: float = 200.0
    tgfb_signal: str = "absent"
    hcc_aetiology: str = "viral"
    cd10_alpl_signal: str = "absent"
    net_marker_level: str = "normal"
    cith3_positive: bool = False
    hla_dr_signal: str = "absent"
    gmcsf_signal: str = "absent"


@dataclass
class TANResultV4:
    pss_by_gamma: Dict[float, float]   # {gamma: PSS}
    pss_default: float                  # PSS at gamma=0.20
    pss_range: Tuple[float, float]      # (min, max) across gamma range
    ci_lower: float                     # Monte Carlo 95% CI (continuous inputs, gam
    ci_upper: float
    domains: List[dict]
    weight_note: str
    collinearity_note: str
    limitations: List[str] = field(default_factory=list)


def compute_tan_polarity_v4(patient: TANPatientV4,
                            n_sims: int = 5000,
                            seed: int = 42) -> TANResultV4:

    cat_scores = {
        "tgfb":      f_tgfb(patient.tgfb_signal),
```

```
    "aetiology": f_aetiology(patient.hcc_aetiology),
    "cd10_alpl": f_cd10_alpl(patient.cd10_alpl_signal),
    "nets":      f_nets(patient.net_marker_level, patient.cith3_positive),
    "hla_dr":    f_hla_dr(patient.hla_dr_signal),
    "gmcsf":     f_gmcsf(patient.gmcsf_signal),
}

cat_weighted = sum(WEIGHTS_CAT[k] * v for k, v in cat_scores.items())

pss_by_gamma: Dict[float, float] = {}
for g in GAMMA_RANGE:
    ana = g_ana(patient.nlr, patient.vegf_pg_ml, g)
    # Collinearity discount reduces ANA weight slightly:
    # effective ANA weight = W_ANA_RAW * (1 - g * f_nlr * f_vegf / (100 * W_ANA_
    # Simplified: just apply g inside g_ana and multiply by W_ANA_RAW
    pss = min(100.0, W_ANA_RAW * ana + cat_weighted)
    pss_by_gamma[g] = round(pss, 1)

pss_default = pss_by_gamma[0.20]
pss_range = (min(pss_by_gamma.values()), max(pss_by_gamma.values()))

# Monte Carlo at gamma=0.20 only (categorical inputs not perturbed)
rng = random.Random(seed)
sims = []
for _ in range(n_sims):
    nlr_p = max(0.1, patient.nlr * (1 + rng.gauss(0, 0.12)))
    vegf_p = max(10.0, patient.vegf_pg_ml * (1 + rng.gauss(0, 0.13)))
    ana_p = g_ana(nlr_p, vegf_p, 0.20)
    sims.append(min(100.0, W_ANA_RAW * ana_p + cat_weighted))
sims.sort()
ci_lower = round(sims[int(0.025 * n_sims)], 1)
ci_upper = round(sims[int(0.975 * n_sims)], 1)

domains = [
    {"name": "ANA (NLR+VEGF)",
     "f_nlr": round(f_nlr(patient.nlr), 1),
     "f_vegf": round(f_vegf(patient.vegf_pg_ml), 1),
     "g_ana_gamma020": round(g_ana(patient.nlr, patient.vegf_pg_ml, 0.20), 1),
     "w_ana": round(W_ANA_RAW, 3),
     "weighted_gamma020": round(W_ANA_RAW * g_ana(patient.nlr, patient.vegf_pg_m
     "precision_nlr": DOMAIN_EVIDENCE["nlr"][2],
     "precision_vegf": DOMAIN_EVIDENCE["vegf"][2]},
] + [
    {"name": k, "raw": round(v, 1), "weight": WEIGHTS_CAT[k],
```

```
        "weighted": round(WEIGHTS_CAT[k] * v, 3),
        "precision": DOMAIN_EVIDENCE[k][2],
        "ci_status": DOMAIN_EVIDENCE[k][3]}
       for k, v in cat_scores.items()
    ]

    weight_note = (
        "Weights derived from SE-based inverse-variance method: "
        "w_d = (Precision_d * ln(HR_d)) / sum(Precision_d' * ln(HR_d')). "
        f"NLR precision={DOMAIN_EVIDENCE['nlr'][2]:.0f} (published 95% CI, n=9,952). "
        "All other molecular domains use imputed floor precision=4.0 "
        "(no published CI available). This reflects the actual evidence landscape: "
        "NLR dominates because it has the best-evidenced HR, not because it is "
        "biologically more important than the molecular domains."
    )

    collinearity_note = (
        f"ANA collinearity sensitivity: PSS ranges from {pss_range[0]:.1f} "
        f"(gamma=0.40, strong discount) to {pss_range[1]:.1f} (gamma=0, no discount). "
        f"Range span = {pss_range[1]-pss_range[0]:.1f} points. "
        "No published rho(NLR, serum VEGF) in HCC exists; gamma is not estimable "
        "as a point value. Report PSS as a range until quantified."
    )

    limitations = [
        "MODEL UNVALIDATED: PSS has not been tested against patient-level OS, PFS, "
        "or ICI response data. The 0-100 scale is clinically uninterpretable without "
        "calibration against real outcomes.",
        "WEIGHT DOMINANCE: NLR accounts for ~63% of total weight under SE-based "
        "weighting. Molecular TAN domains contribute 1-14% each. Adding molecular "
        "data changes PSS by at most ~10 points; the model is currently dominated "
        "by NLR and HLA-DR+ when measured by evidence precision.",
        "GAMMA UNCERTAINTY: The collinearity discount is not quantifiable from "
        "current literature. PSS should be reported as a range, not a point value.",
        "SCENARIOS ARE RECONSTRUCTIONS: Demonstration scenarios are derived from "
        "cohort profile descriptions in published papers, not independent patient da
        "VALIDATION PROTOCOL: Section 5 specifies a prospective validation design "
        "(n=580, multi-centre) and a partial TCGA-LIHC proxy analysis. Neither has "
        "been executed. This framework is not ready for clinical application.",
    ]

    return TANResultV4(pss_by_gamma=pss_by_gamma, pss_default=pss_default,
                       pss_range=pss_range, ci_lower=ci_lower, ci_upper=ci_upper,
                       domains=domains, weight_note=weight_note,
```

```
                      collinearity_note=collinearity_note, limitations=limitations)


def print_result_v4(result: TANResultV4, label: str):
    print("\n" + "=" * 80)
    print(label)
    print("=" * 80)
    print(f"PSS (gamma=0.20): {result.pss_default:.1f} / 100")
    print(f"PSS sensitivity range (gamma 0-0.40): {result.pss_range[0]:.1f} - {resul
    print(f"95% CI (MC, continuous inputs, gamma=0.20): [{result.ci_lower:.1f}, {res
    print(f"\nGamma sensitivity:")
    for g, pss in result.pss_by_gamma.items():
        print(f"  gamma={g:.2f}  →  PSS={pss:.1f}")
    print(f"\nWeight note: {result.weight_note}")
    print(f"\nCollinearity note: {result.collinearity_note}")
    print("\nDomain decomposition:")
    d = result.domains[0]
    print(f"  ANA: f_NLR={d['f_nlr']:.1f}, f_VEGF={d['f_vegf']:.1f}, "
          f"g_ANA(g=0.20)={d['g_ana_gamma020']:.1f}, w={d['w_ana']:.3f}, "
          f"wtd={d['weighted_gamma020']:.2f}")
    print(f"       NLR precision={d['precision_nlr']:.0f}, VEGF precision={d['precis
    for dom in result.domains[1:]:
        print(f"  {dom['name']:14s}: raw={dom['raw']:5.1f}, w={dom['weight']:.4f}, "
              f"wtd={dom['weighted']:.3f}, precision={dom['precision']:.1f}")
    print("\n*** LIMITATIONS ***")
    for lim in result.limitations:
        print(f"  ! {lim}")


def demo():
    scenarios = [
        ("Scenario 1 — Responder profile [Jost-Brinkmann F et al. APT 2023]",
         TANPatientV4(nlr=2.1, vegf_pg_ml=195.0, tgfb_signal="absent",
                      hcc_aetiology="viral", cd10_alpl_signal="absent",
                      net_marker_level="normal", cith3_positive=False,
                      hla_dr_signal="present", gmcsf_signal="absent")),

        ("Scenario 2 — MASH poor-prognosis [Meng Y, Zhu X et al. 2024 + Teo J et al.
         TANPatientV4(nlr=5.7, vegf_pg_ml=415.0, tgfb_signal="active",
                      hcc_aetiology="mash", cd10_alpl_signal="elevated",
                      net_marker_level="elevated", cith3_positive=True,
                      hla_dr_signal="absent", gmcsf_signal="elevated")),

        ("Scenario 3 — Cirrhotic-ECM NET-prominent [Shen XT et al. Exp Hematol Oncol
```

```
         TANPatientV4(nlr=4.2, vegf_pg_ml=340.0, tgfb_signal="moderate",
                      hcc_aetiology="formerly_viral_cirrhosis",
                      cd10_alpl_signal="not_documented",
                      net_marker_level="high", cith3_positive=True,
                      hla_dr_signal="low", gmcsf_signal="mild")),
    ]
    for label, patient in scenarios:
        result = compute_tan_polarity_v4(patient)
        print_result_v4(result, label)


if __name__ == "__main__":
    demo()
```