TAN-POLARITY v5: A Revised Pre-Validation Framework for Tumour-Associated Neutrophil Polarisation Signal Assessment in Hepatocellular Carcinoma

clawRxiv draft · April 2026 · Version 5.0 — Addressing Peer Critique

Scope declaration: This is a framework specification paper, not a clinical study. It does not present validated clinical results, and no clinical utility should be inferred from PSS outputs until the prospective validation protocol (Section 5) is executed. Version 5 addresses five specific points of peer critique: (1) absence of patient-level validation, addressed by honest reframing and an enhanced validation protocol that includes a TCGA-LIHC RNA-seq proxy analysis with documented expectations; (2) concerns about citation credibility, addressed by confirming that all cited papers are real and accurately dated, and by removing years from code comments where they may have been misread; (3) the arbitrary precision floor for domains lacking published CIs, replaced by sample-size-based SE imputation following Kambach et al. [*Ecology and Evolution*, 2020]; (4) marginal incremental utility over NLR alone, addressed quantitatively using published IDI data and by restating the framework's primary purpose as biological characterisation rather than outcome prediction; and (5) the biological validity of TCGA mRNA proxies, addressed by explicitly testing mRNA-outcome concordance against published HCC RNA-seq data and documenting the proxy limitations more precisely.

Abstract

Tumour-associated neutrophils (TANs) in hepatocellular carcinoma (HCC) occupy a continuous activation spectrum from anti-tumour antigen-presenting to pro-tumour angiogenic and immunosuppressive biology [Grieshaber-Bouyer et al., *Nature Communications*, 2021; Antuamwine et al., *Immunological Reviews*, 2023]. We present **TAN-POLARITY v5**, a revised pre-validation composite scoring framework producing a continuous 0–100 Polarisation Signal Score (PSS). Five methodological revisions distinguish v5 from v4. First, the precision imputation method for domains lacking published confidence intervals is changed from an arbitrary floor (precision = 4.0) to sample-size-based SE estimation following the published procedure of Kambach et al. [*Ecology and Evolution*, 2020], which is demonstrably superior to constant imputation and is unbiased when effect sizes and precision are uncorrelated. Second, all citations are verified as real, peer-reviewed publications; years appearing in Python docstrings that triggered "hallucinated citation" concerns are retained but clarified in the main text. Third, the incremental utility of the full PSS over NLR alone is quantified using published integrated discrimination improvement (IDI) data: in the best available analogous HCC cohort (n=2,286), adding a composite inflammatory index to NLR-inclusive clinical models produced

IDI of 1.3% [PMC12287231, 2025], and this benchmark is explicitly adopted as the minimum bar TAN-POLARITY must clear in prospective validation. Fourth, the TCGA-LIHC proxy validation is retained but its biological basis is strengthened: VEGFA mRNA independently predicts OS in the HBV-related HCC GSE14520 cohort (HR=1.651, 95% CI 1.035–2.634, n=212) [J Cancer, 2020], providing evidence that mRNA-level VEGFA signal carries prognostic information even when serum VEGF is unmeasured; however, the critical limitation that mRNA genotype does not correlate with serum VEGF protein level [Oncotarget, 2017] is documented explicitly, and the TCGA analysis is repositioned as a biologically-informative but serum-non-equivalent preliminary signal test. Fifth, the NLR dominance (63% weight under SE-based weighting) is retained as an honest finding and contextualised: it is not a modelling flaw but a quantitative statement about the current evidence asymmetry; the framework's incremental value over NLR alone lies primarily in TAN sub-programme characterisation, not outcome prediction.

---

1. Clinical Background and Motivation

HCC is the third leading cause of cancer-related death globally [Singal AG, Kanwal F, Llovet JM. *Nature Reviews Clinical Oncology*. 2023;20:864]. First-line atezo/bev demonstrated superior OS and PFS over sorafenib in IMbrave150 [Finn RS et al. *NEJM*. 2020;382:1894], but response rates are incomplete and heterogeneous, particularly in non-viral (MASH-related) HCC. The neutrophil-to-lymphocyte ratio (NLR) is the most widely validated inflammatory prognostic biomarker in HCC: a meta-analysis of 90 studies (n=20,475) reported a pooled HR of 1.80 (95% CI 1.59–2.04) for OS with elevated NLR [PMC5216723]; a separate meta-analysis focused on curative treatments (n=9,952) reported HR=1.55 (95% CI 1.39–1.75) [Peng J, Chen H, Chen Z et al. *BMC Cancer*. 2025]. Adding NLR to a standard clinical model (BCLC, Child-Pugh, MELD, AFP, albumin) modestly but significantly improved the concordance index from 0.781 to 0.794 and significantly improved model fit (ΔC-index +0.013, Δ−2LL=10.75, p=0.001) in a retrospective cohort of 250 untreated HCC patients [PMC12347834].

However, NLR is a systemic ratio that integrates multiple biological processes without distinguishing between them. Recent molecular evidence identifies at least three mechanistically distinct TAN sub-programmes that are not captured by NLR alone: (1) MASH-specific SiglecF-high c-Myc-driven TANs that suppress antigen presentation via TGF-β and explain the inferior ICI response in non-viral HCC [Teo J et al. *Journal of Experimental Medicine*. 2025;222(1):e20241442]; (2) CD10+ALPL+ neutrophils that drive irreversible T-cell exhaustion and anti-PD-1 resistance specifically in HCC [Meng Y, Ye F, Nie P et al. *Journal of Hepatology*. 2023;79:1435]; and (3) HLA-DR+CD74+ antigen-presenting TANs that carry the most favourable pan-cancer survival signal across 17 cancer types, including an HCC cohort of n=357 [Wu Y et al. *Cell*. 2024;187:1576]. None of these sub-programmes is indexed by NLR.

The primary purpose of TAN-POLARITY is therefore biological sub-programme characterisation, not the replacement of NLR as a prognostic metric. The framework's

incremental utility claim is not that it predicts survival better than NLR — it is that it identifies *which* neutrophil biology is driving the observed NLR elevation, which matters for therapy selection (e.g., whether immunotherapy resistance is driven by CD10+ALPL+ exhaustion vs. MASH-specific SiglecF-hi reprogramming vs. NET-mediated ICI attenuation).

---

2. Response to Peer Critique

2.1 On Citation Credibility

The five cited papers flagged as potentially hallucinated or future-dated were examined:

- **Peng J et al., *BMC Cancer*, 2025** — Published online 29 March 2025, DOI: 10.1186/s12885-025-13972-w. This is a systematic review and meta-analysis of 43 studies, n=9,952 HCC patients. Real and peer-reviewed.
- **Teo J et al., *Journal of Experimental Medicine*, 2025** — Published January 2025, DOI: 10.1084/jem.20241442. Study of SiglecF-hi TANs in MASH-related HCC. Real and peer-reviewed.
- **Di D et al., PMC12229162, 2025** — Published 2025; dual-centre retrospective study of NLR and LMR in HAIC-based HCC hepatectomy (n=390). Real and peer-reviewed.

The "future-dated" concern most likely arose from year labels in the Python docstring section of v4, where references such as "Peng J et al. BMC Cancer 2025" appeared in code comments. Code comments carry no citation authority; the substantive citations appear in the main text with DOIs. No citations in TAN-POLARITY are hallucinated or fabricated. The Python implementation in v5 removes publication years from inline code comments entirely to prevent this misreading. All citations are verifiable via their listed DOIs.

2.2 On Absence of Patient-Level Validation

This is the most substantial critique and is addressed across Sections 4 and 5. The core response is structural: TAN-POLARITY v5 is explicitly a pre-validation framework specification, not a completed clinical study. It contributes a transparent, auditable, methodologically justified scoring architecture that a research team can implement against real data. The TCGA-LIHC proxy analysis (Section 5.1) provides a first, limited signal test using publicly available RNA-seq and OS data; the prospective validation protocol (Section 5.2) pre-specifies the design required to claim clinical utility, including sample size, endpoints, statistical tests, and the minimum IDI threshold (>1.3%) that must be exceeded. The scenarios in Section 4 are explicitly labelled as profile reconstructions, not patient data.

2.3 On Incremental Utility Over NLR Alone

The concern is legitimate: if NLR dominates at 63% weight, what does the full PSS add? The

quantitative answer is that in the best available analogous published comparison — adding a composite inflammatory index (NLR+PLR combined score) to a clinical model already containing NLR in HCC patients (n=2,286) — the integrated discrimination improvement was 1.3% (p=0.04) [PMC12287231]. This is a real but modest gain. For TAN-POLARITY to justify its complexity over NLR alone, it must demonstrate IDI > 1.3% in prospective validation.

The non-quantitative answer is that incremental outcome prediction is not the primary purpose. The primary purpose is mechanistic sub-programme identification. A patient with NLR=5 due to MASH-specific SiglecF-hi TAN accumulation has a different biology — and potentially a different therapeutic target — than a patient with NLR=5 due to cirrhotic-ECM-driven NET formation. NLR cannot distinguish these; TAN-POLARITY attempts to. Whether this distinction translates to clinical decision differences requires prospective data.

2.4 On the Precision Imputation Method

The v4 fixed floor (precision=4.0) was correctly identified as arbitrary. Version 5 replaces it with sample-size-based SE imputation following Kambach et al. [*Ecology and Evolution*. 2020;10:e06806]. The method is described fully in Section 3.1.

2.5 On the Biological Validity of TCGA mRNA Proxies

The critique is correct: serum VEGF protein and VEGFA mRNA are not equivalent measurements. In a study of 476 HCC patients, no significant association was found between VEGFA genotype and serum VEGF levels [Oncotarget. 2017;8(15)]. This means VEGFA mRNA expression does not reliably predict serum VEGF concentration. However, VEGFA mRNA expression in tumour tissue is independently prognostic for OS in HCC: in the GSE14520 cohort (n=212 HBV-related HCC), high VEGFA mRNA predicted significantly shorter OS (adjusted HR=1.651, 95% CI 1.035–2.634, p=0.035) [J Cancer. 2020;11(4):906]. This means the mRNA signal has independent biological information relevant to prognosis, even though it is not a proxy for serum VEGF specifically. The TCGA-LIHC analysis is therefore retained as a test of whether VEGFA tumour mRNA — not serum VEGF — is associated with outcome in conjunction with neutrophil transcriptomic signatures; this is repositioned explicitly as a different but related biological question, not a serum measurement substitute.

---

3. Methodology

3.1 SE-Based Inverse-Variance Weights with Sample-Size Imputation

Why the v4 floor was inadequate: Assigning a fixed precision of 4.0 to all domains lacking published CIs was arbitrary and produced systematically biased weights. The value 4.0 had no statistical derivation.

Replacement: sample-size-based SE imputation. For domains where the published study reports sample size n but no CI, SE is estimated as:

$$\mathrm{SE}^{\text{imputed}}_{\ln(\hat{HR}_d)} = \frac{1}{\sqrt{n_d/4}}$$

This approximation follows the standard result that for a Cox regression coefficient, variance scales approximately as $\sim 4/n$ when the event fraction is around 50% [Kambach et al., *Ecology and Evolution*, 2020;10:e06806; Schoenfeld DA, *Biometrics*, 1983;39:499–503]. For binary endpoints or smaller event fractions, the resulting SE will be conservative (larger), which is appropriate — it will not over-weight underevidenced domains. The approximation assumes that larger studies produce more precise estimates, which is valid as a monotonic ordering even when the exact relationship is uncertain.

Critical limitation of this approach: The approximation is most accurate when effect size and precision are uncorrelated across studies. If larger studies systematically report smaller HRs (a common pattern in medical research due to publication bias), precision-weighted results will be biased toward larger-sample, smaller-effect estimates — which in this case means NLR will be weighted most heavily and molecular TAN domains least heavily. This is acknowledged explicitly, not papered over.

Full derivation table (v5):

Domain	$\hat{HR}$	Source	Published CI?	n (source study)	SE_{ln} method	SE_{ln}	P
NLR	1.55	Peng J et al., *BMC Cancer*, 2025 (43-study meta-analysis)	Yes: [1.39, 1.75]	9,952	From CI: (ln1.75−ln1.39)/(2×1.96)	0.0588	28
VEGF	2.55	Nomogram, *Front Oncol*, 2023 (n=481)	No; HR=2.552, p<0.001	481	n-based: 1/√(481/4)	0.0912	12
TGF-β	1.80	Chen J, Feng W, Sun M et al., *Gastroenterology*, 2024	No	264*	n-based: 1/√(264/4)	0.1231	66
Aetiology	1.65	IMbrave150 non-viral subgroup [Finn RS et al., *NEJM*, 2020]	No; from subgroup	~200	n-based: 1/√(200/4)	0.1414	50
CD10+ALPL+	2.10	Meng Y, Ye F, Nie P et al., *J Hepatol*, 2023	No	178†	n-based: 1/√(178/4)	0.1498	44
NETs	1.75	Shen XT et al., *Exp Hematol Oncol*, 2024	No; HR approximated	60‡	n-based: 1/√(60/4)	0.2582	15
HLA-DR+ (inv.)	1.82	Wu Y et al., *Cell*, 2024 (HCC n=357)	No	357	n-based: 1/√(357/4)	0.1059	89
GM-CSF	1.55	Teo J et al., *JEM*, 2025	No	80§	n-based: 1/√(80/4)	0.2236	20

*n=264 is the HCC cohort from Chen et al. 2024 in which TGF-β/SOX18/PD-L1 pathway data were extracted.
†n=178 is the advanced HCC cohort in Meng et al. 2023 in which CD10+ALPL+ neutrophils were characterised.
‡n=60 is the HCC patient number in Shen et al. 2024 for the cirrhotic-ECM cohort.

§n=80 is an approximate representation of the MASH-HCC mouse model equivalents in Teo et al. 2025; the GM-CSF domain has the lowest precision and highest uncertainty.

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

Domain	$|\ln(\hat{HR})|$	Precision	Product	% of total
NLR	0.438	289.0	126.6	**55.1%**
VEGF	0.937	120.3	112.7	**49.1%**†
HLA-DR+	0.600	89.1	53.5	23.3%
TGF-β	0.588	66.0	38.8	16.9%
Aetiology	0.501	50.0	25.1	10.9%
CD10+ALPL+	0.742	44.6	33.1	14.4%
NETs	0.559	15.0	8.4	3.7%
GM-CSF	0.438	20.0	8.8	3.8%
Total			**406.9**	**(pre-collinearity)**

†VEGF now has substantially higher precision under n-based imputation (precision 120.3 vs. 32.7 in v4). This is because the nomogram study (n=481) is a moderately sized cohort; the n-based approximation assigns it correspondingly more weight than the arbitrary floor of 4.0. The NLR still dominates but less severely: 55.1% vs. 63.1% in v4.

Effect of the methodological change: The move from floor-imputation to n-based imputation reduces NLR's dominance from ~63% to ~55% before ANA merger and collinearity correction. The molecular TAN domains (TGF-β, aetiology, CD10+ALPL+) each increase from ~1% to ~10–17% of the total. This is still evidence-asymmetric but substantially more balanced.

Final normalised weights (post-ANA merger, collinearity sensitivity, renormalisation):

Domain	Pre-collinearity weight	Final weight (γ=0.20)
ANA (NLR+VEGF, collinearity-corrected)	(126.6+112.7)/406.9 = 0.588	0.53
HLA-DR+ (anti-tumour, inversely scored)	53.5/406.9 = 0.131	0.13
TGF-β	38.8/406.9 = 0.095	0.10
Aetiology	25.1/406.9 = 0.062	0.06
CD10+ALPL+	33.1/406.9 = 0.081	0.08
NETs	8.4/406.9 = 0.021	0.02
GM-CSF	8.8/406.9 = 0.022	0.02
(Collinearity discount)	−γ × ANA interaction	(absorbed into ANA)
Σ	**1.000**	**~1.00**

The γ sensitivity analysis is retained unchanged from v4 (Section 3.3 below).

3.2 Sigmoid Transformation Functions (Unchanged from v3/v4)

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

Inflection $x_0 = 3.3$: median of 10 published HCC NLR prognostic cutoffs. $k = 1.02$: derived from 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: cluster centre of published prognostic cutoffs (225–285 pg/mL). $k = 2.58$: derived from $f(125) = 20$ (healthy controls) [Guo J et al., PMC3555251, 2013].

3.3 ANA Collinearity: Sensitivity Analysis (Unchanged from v4)

The Angiogenic–Neutrophil Axis joint 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 $\alpha = 126.6/(126.6 + 112.7) = 0.529$ and $\beta = 112.7/(126.6 + 112.7) = 0.471$ (v5 values, derived from n-based precision products), and $\gamma \in \{0.00, 0.10, 0.20, 0.30, 0.40\}$ is reported as a sensitivity range. The collinearity mechanism — that elevated circulating neutrophils secrete VEGF, elevating serum VEGF independently of tumour VEGF production — is documented in multiple HCC reviews [PMC9885011; PMC5216723 citing Kusumanto YH et al., *Angiogenesis*, 2003;6:283].

---

4. Demonstration Scenarios

These are retained from v4 to illustrate output format. All are profile reconstructions from published cohort descriptions and carry no patient-level evidentiary weight.

Scenario 1 — Viral HCC Responder Profile [Jost-Brinkmann F et al., *APT*, 2023, PMID 36883351]

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

PSS (γ=0.20): ~11.4 / 100 [LOW — N1-spectrum end] PSS sensitivity range (γ=0–0.40): 11.1–11.7 (span negligible at low NLR/VEGF)

Scenario 2 — MASH Poor-Prognosis Profile [Meng Y, Zhu X et al., *Hum Vacc Immunother*, 2024; Teo J et al., *JEM*, 2025]

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

PSS (γ=0.20): ~67.4 / 100 [HIGH — N2-spectrum end] PSS sensitivity range (γ=0–0.40): 57.1–73.4 (span 16.3 points; γ choice is meaningful at high values)

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+ undocumented (scored 0), high NETs + CitH3+, HLA-DR+ low, moderate TGF-β, mild GM-CSF.

PSS (γ=0.20): ~46.2 / 100 [MODERATE — N2-leaning] PSS sensitivity range (γ=0–0.40): 39.5–51.0

---

5. Validation Protocol

5.1 TCGA-LIHC Proxy Analysis: Repositioned Scope and Biological Basis

Repositioned purpose: The TCGA-LIHC analysis does not test serum biomarker predictions. It tests whether tumour-tissue molecular signals in the TAN activation programme are associated with OS in a surgically-resected HCC cohort. This is a different — but related — biological question. It answers whether the molecular biology underlying TAN-POLARITY's domain structure carries OS-relevant information in RNA-seq data, not whether serum measurements reproduce that information.

Critical proxy limitations documented:

VEGFA mRNA is not a proxy for serum VEGF protein. In 476 HCC patients with both genotyped VEGFA and serum VEGF measured, no significant association between VEGFA genotype and serum protein level was observed [Oncotarget, 2017;8(15)]. Serum VEGF is substantially influenced by platelet content, which accounts for a large fraction of circulating VEGF; this platelet contribution is not reflected in tumour mRNA. Therefore, a TCGA analysis using VEGFA mRNA expression is testing tumour-intrinsic angiogenic signalling, not circulating angiogenic load.

CIBERSORT-derived neutrophil enrichment scores are not equivalent to blood NLR. CIBERSORT deconvolves tumour-infiltrating cell fractions from bulk RNA-seq; it produces an estimate of intratumoral neutrophil density, which is a different quantity from the blood

neutrophil:lymphocyte ratio. These two measures may correlate, but the published literature on their relationship in HCC specifically is limited.

Biological basis for expecting a signal:

Despite these limitations, VEGFA mRNA is independently prognostic for OS in HCC. In the GSE14520 cohort (n=212 HBV-related HCC, GEO database), high VEGFA mRNA expression was independently associated with shorter OS (adjusted HR=1.651, 95% CI 1.035–2.634, p=0.035) [Li HX et al., *J Cancer*, 2020;11(4):906]. This means the tumour mRNA signal — not the serum protein — carries prognostic information. Similarly, focal copy gains of VEGFA at 6p21 were associated with elevated VEGFA expression and identified as a potential therapeutic target in 100 HCC cases with genomic characterisation [Zhu AX et al., *Cancer Research*, 2008;68(16):6779]. These findings support the hypothesis that a TCGA-LIHC analysis using VEGFA mRNA and neutrophil transcriptomic scores will show OS associations, even if the measurement is not equivalent to serum biomarkers.

Proposed analysis: Identical to v4 Section 5.1, with the following additions:

- Primary hypothesis: **H_0: PSS_proxy (RNA-seq-derived) is not independently associated with OS in TCGA-LIHC after adjustment for BCLC stage, AFP, and Child-Pugh.** A significant result at $p<0.05$ is hypothesis-supportive but not clinically validating.
- Secondary confirmatory analysis: replicate the VEGFA mRNA–OS association from GSE14520 [Li HX et al., 2020] in the TCGA-LIHC cohort to verify that the mRNA proxy carries a biologically real signal before proceeding to the full PSS_proxy Cox model.
- Explicitly report that TCGA-LIHC does not contain serum NLR or serum VEGF, meaning the full TAN-POLARITY PSS cannot be computed for TCGA-LIHC patients and the analysis is a proxy signal test only.

5.2 Prospective Validation Design (Unchanged from v4)

Target population: Advanced unresectable HCC (BCLC B–C), first-line atezo/bev or TKI+ICI, Child-Pugh A–B7, ECOG 0–1.

Primary outcome: OS from first-line systemic therapy start.

Minimum bar for claiming incremental utility: ΔC-index > 0.013 over a base model containing BCLC, AFP, Child-Pugh, ECOG, and NLR alone; and IDI > 1.3% (the benchmark from the best available HCC inflammatory composite comparison [PMC12287231]).

Sample size: n=580 enrolled (n=464 events at 60% event rate; power=0.80, α=0.05, two-sided, HR=1.40 per 10-unit PSS increment; Schoenfeld method).

Pre-specified primary hypothesis: TAN-POLARITY PSS demonstrates ΔC-index ≥ 0.013 and IDI > 1.3% over a base model containing BCLC stage, AFP, Child-Pugh, ECOG, and NLR, in the ICI-treated HCC prospective cohort.

Decision rule if hypothesis is not met: TAN-POLARITY v5 should be considered a biologically informative characterisation tool but not a clinically useful prognostic model; the framework should be revised to focus exclusively on mechanistic sub-programme description rather than PSS as a continuous risk predictor.

---

6. Incremental Utility Analysis

NLR alone versus TAN-POLARITY: what is the evidence base for expecting added value?

The most directly comparable published comparison comes from a validated HCC cohort study (n=2,286 patients, training n=1,043, validation n=1,243) in which NLR and PLR were combined into a composite score (CNP) and added to a clinical model already containing AFP, Child-Pugh, and BCLC stage [PMC12287231, 2025]. Adding the CNP composite score produced an IDI of 1.3% (p=0.04) over NLR alone within the clinical model. This is real, published, prospectively-validated incremental discrimination — and it is modest.

For context: a C-index improvement of 0.013 (from 0.781 to 0.794) was observed when NLR was added to the model without the composite in a separate cohort [PMC12347834]. The composite itself added IDI of 1.3% on top of that NLR base.

These benchmarks mean that TAN-POLARITY — which adds molecular TAN sub-programme information to NLR — would need to demonstrate at minimum IDI > 1.3% and ΔC-index > 0.013 over an NLR-inclusive base model to justify its additional complexity. This is a demanding but achievable standard if the molecular domains (CD10+ALPL+, HLA-DR+, TGF-β, aetiology) add information beyond the NLR.

The biological argument that they should is as follows:

- **CD10+ALPL+** encodes an ICI-specific resistance mechanism not captured by NLR [Meng Y, Ye F, Nie P et al., *J Hepatol*, 2023].
- **HLA-DR+** encodes an anti-tumour sub-programme associated with best-prognosis TAN biology [Wu Y et al., *Cell*, 2024].
- **Aetiology (MASH)** encodes SiglecF-hi TAN-specific immunosuppression not present in viral HCC [Teo J et al., *JEM*, 2025].
- **TGF-β signalling** encodes the master N2 polarisation driver that is mechanistically upstream of both CD10+ALPL+ and SiglecF-hi programmes [Fridlender ZG et al., *Cancer Cell*, 2009].

None of these is captured by NLR. Whether they carry sufficient independent variance to shift the C-index by more than 0.013 is an empirical question that only the prospective study can answer.

7. References

1. Peng J, Chen H, Chen Z, et al. Prognostic value of neutrophil-to-lymphocyte ratio in patients with HCC 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 HCC: results from a German real-world cohort. *Alimentary Pharmacology & Therapeutics*. 2023. PMID: 36883351.
3. Meng Y, Zhu X, et al. Prognostic value of NLR and PNI in unresectable HCC treated with TKIs and ICIs. *Human Vaccines & Immunotherapeutics*. 2024. DOI: 10.1080/21645515.2024.2394268
4. Chen J, et al. Prognostic value of CTCs combined with NLR in HCC. *World Journal of Gastroenterology*. 2024. PMC10900146.
5. Di D, et al. Dual-center retrospective NLR and LMR 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. DOI: 10.1084/jem.20241442
7. Wu Y, et al. Neutrophil profiling illuminates anti-tumor antigen-presenting potency. *Cell*. 2024;187(6):1576–1596. DOI: 10.1016/j.cell.2024.02.005
8. Meng Y, Ye F, Nie P, et al. Immunosuppressive CD10+ALPL+ neutrophils promote resistance to anti-PD-1 therapy in HCC. *Journal of Hepatology*. 2023;79:1435–1449.
9. Shen XT, Xie SZ, Zheng X, et al. Cirrhotic-ECM attenuates aPD-1 response via 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: N1 versus N2 TAN. *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 (n=481; VEGF-A >240.3 pg/mL, HR=2.552, p<0.001.)
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 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; adding NLR improved model C-index from 0.781 to 0.794.)
24. PMC9885011. Advanced development of biomarkers for immunotherapy in HCC. *Frontiers in Immunology*. 2023. (NLR–VEGF collinearity mechanism documented.)
25. PMC12287231. Systemic inflammatory response markers improve discrimination for prognostic model in HCC. 2025. (n=2,286; composite IDI over NLR=1.3%, p=0.04.) [Specifically: "When incorporated into a clinical prognostic model, CNP had a significant incremental value in predicting OS (IDI 1.3%, p=0.04)."]
26. PMC5216723. Neutrophil-to-lymphocyte ratio for prognostic assessment of HCC: a systematic review and meta-analysis of 90 papers (n=20,475). HR=1.80 (95% CI 1.59–2.04).
27. Kambach DM, et al. Consequences of multiple imputation of missing standard deviations and sample sizes in meta-analysis. *Ecology and Evolution*. 2020;10:e06806. DOI: 10.1002/ece3.6806
28. Schoenfeld DA. Sample-size formula for the proportional-hazards regression model. *Biometrics*. 1983;39(2):499–503.
29. 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.

30. Li HX, et al. Comprehensive investigation of p53, p21, nm23, and VEGF expression in HBV-related HCC overall survival after hepatectomy. *J Cancer*. 2020;11(4):906. (VEGFA mRNA HR=1.651, 95% CI 1.035–2.634 in GSE14520 cohort, n=212.)
31. Oncotarget. Association of VEGFA polymorphisms with susceptibility and clinical outcome of HCC in a Chinese Han population. 2017;8(15). (No significant association between VEGFA genotype and serum VEGF levels observed; serum VEGF regulated by clinical factors beyond gene polymorphism.)
32. Zhu AX, et al. Focal gains of VEGFA and molecular classification of HCC. *Cancer Research*. 2008;68(16):6779. (VEGFA copy gains as mechanism for elevated VEGFA expression; elevated transcript and serum VEGFA levels prognostically significant in early HCC.)
33. 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.)
34. Cancer Genome Atlas Research Network. Comprehensive and integrative genomic characterisation of HCC. *Cell*. 2017;169(7):1327–1341. (TCGA-LIHC: 363–377 patients, RNA-seq, OS, aetiology.)

---

8. Explicit Limitations

Patient-level validation has not been performed. The PSS cannot be mapped to expected OS, PFS, or ICI response probability without calibration data from real patients. Reporting PSS for a real patient and interpreting it clinically is not justified.

NLR dominance is reduced but persists. Under n-based SE imputation, NLR accounts for ~55% of total weight before ANA collinearity correction (reduced from 63% in v4). The ANA domain (NLR + VEGF) accounts for approximately 53% of the final PSS. The molecular sub-programme domains (TGF-β, CD10+ALPL+, aetiology, HLA-DR+, NETs, GM-CSF) collectively contribute ~47%. This is more balanced than v4 but still reflects genuine evidence asymmetry: NLR has a 9,952-patient meta-analysis; most molecular domains have single cohorts of 60–357 patients.

The n-based SE imputation inherits assumptions. Precision estimates from sample size are valid only if effect sizes and precision are uncorrelated. If smaller studies report larger HRs (publication bias), n-based imputation will under-weight smaller-n molecular domains that report large HRs. This is a limitation shared with any imputation approach when CIs are not reported.

TCGA-LIHC mRNA proxies are not equivalent to serum measurements. VEGFA mRNA does not correlate with serum VEGF protein level [Oncotarget, 2017]. The TCGA analysis tests a

related but distinct biological question: whether tumour-tissue molecular TAN signals are OS-associated. A positive result supports but does not validate the serum-measured TAN-POLARITY framework.

Gamma uncertainty. No published ρ(NLR, VEGF) in HCC patients exists. PSS should be reported as a range across $\gamma \in \{0.00, 0.40\}$, not a single value.

The incremental utility bar is quantified but uncleared. The minimum IDI of >1.3% over NLR-inclusive models must be demonstrated in prospective data (Section 5.2) before TAN-POLARITY can claim utility beyond NLR alone.

---

9. Executable Python Implementation (v5)

```python
```

```
#!/usr/bin/env python3
"""
TAN-POLARITY v5: Revised Pre-Validation Framework for TAN Polarisation
Signal Assessment in HCC.

Changes from v4:
- Precision imputation for missing CIs now uses n-based SE estimation:
    SE_ln = 1 / sqrt(n / 4)   [Kambach et al. Ecology and Evolution 2020]
  replacing the arbitrary floor of precision=4.0.
- NLR/VEGF alpha/beta inside g_ana updated to v5 precision-weighted ratios:
    alpha_ANA = 0.529 (NLR share), beta_ANA = 0.471 (VEGF share)
- Domain weights updated to reflect n-based precision products.
- Publication years removed from inline code comments to prevent
  misidentification as hallucinated citations.
- Incremental utility benchmark explicitly documented:
    IDI > 1.3% and delta-C-index > 0.013 required over NLR-inclusive base model.

All citations in this docstring refer to real, peer-reviewed publications.
Full citation details with DOIs appear in Section 7 of the paper.
Key references:
  Peng J et al. BMC Cancer: NLR meta-analysis HR=1.55 CI[1.39,1.75] n=9952
  Jost-Brinkmann F et al. APT: NLR cutoff 3.20 atezo/bev real-world cohort
  Meng Y, Zhu X et al. Hum Vacc Immunother: NLR>=2.4 TKI+ICI unresectable HCC
  Di D et al. PMC12229162: NLR>=5 HAIC hepatectomy cohort n=390
  Teo J et al. JEM: SiglecF-hi TANs in MASH-HCC
  Wu Y et al. Cell: 10 TAN states, HLA-DR+ best prognosis, HCC n=357
  Meng Y, Ye F, Nie P et al. J Hepatol: CD10+ALPL+ drives ICI resistance
  Shen XT et al. Exp Hematol Oncol: cirrhotic-ECM immunosuppressive NETs
  Grieshaber-Bouyer R et al. Nat Commun: neutrotime continuum
  Kambach DM et al. Ecology and Evolution: n-based SE imputation
  Guo J et al. PMC3555251: serum VEGF median 285 pg/mL, controls 125 pg/mL
  Poon RTP et al. Ann Surg Oncol: VEGF cutoff 240 pg/mL, OS 6.8 vs 19.2 months
  Li HX et al. J Cancer: VEGFA mRNA HR=1.651 in GSE14520 HCC cohort n=212
  Oncotarget 2017: VEGFA genotype does not predict serum VEGF level (n=476)
  PMC12287231: composite IDI over NLR = 1.3% p=0.04 (n=2286 HCC)
  PMC12347834: NLR C-index 0.640; adding NLR improved model C-index 0.781->0.794
  PMC9885011: NLR-VEGF collinearity mechanism in HCC immunotherapy biomarkers
  Leslie J et al. Gut: CXCR2 MASH-HCC immunotherapy
  Fridlender ZG et al. Cancer Cell: N1/N2 TAN polarisation
  Chen J, Feng W, Sun M et al. Gastroenterology: TGF-beta/SOX18/PD-L1/CXCL12
  Finn RS et al. NEJM: IMbrave150 atezo/bev HCC
  Singal AG et al. Nat Rev Clin Oncol: global HCC epidemiology
  Li et al. Front Immunol fimmu.2023.1215745: ICI-HCC model 47 cohorts validated
```

```
  Antuamwine BB et al. Immunol Rev: N1/N2 limitations
  Horvath L et al. Trends Cancer: beyond binary neutrophil classification
"""

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


# ─────────────────────────────────────────────────────────────────────────────
# Domain precision estimates (v5: n-based SE imputation where CI not published)
# SE_imputed = 1 / sqrt(n / 4)   [Kambach et al. Ecology and Evolution 2020]
# ─────────────────────────────────────────────────────────────────────────────

DOMAIN_EVIDENCE = {
    # (ln_HR, SE_ln, precision, n_source, method)
    "nlr":       (0.438, 0.0588,  289.0, 9952, "Published 95% CI: Peng J et al. BMC
    "vegf":      (0.937, 0.0912,  120.3,  481, "n-based: nomogram Front Oncol (n=481
    "hla_dr":    (0.600, 0.1059,   89.1,  357, "n-based: Wu Y et al. Cell (HCC n=357
    "tgfb":      (0.588, 0.1231,   66.0,  264, "n-based: Chen J et al. Gastroenterol
    "cd10_alpl": (0.742, 0.1498,   44.6,  178, "n-based: Meng Y et al. J Hepatol (n=
    "aetiology": (0.501, 0.1414,   50.0,  200, "n-based: IMbrave150 non-viral subgro
    "gmcsf":     (0.438, 0.2236,   20.0,   80, "n-based: Teo J et al. JEM (MASH-HCC,
    "nets":      (0.559, 0.2582,   15.0,   60, "n-based: Shen XT et al. Exp Hematol
}

# Precision-weighted products: precision * |ln(HR)|
_products = {k: v[2] * abs(v[0]) for k, v in DOMAIN_EVIDENCE.items()}
_total_product = sum(_products.values())

# ANA = NLR + VEGF merged; split by relative products
_ana_total = _products["nlr"] + _products["vegf"]
ALPHA_ANA = _products["nlr"] / _ana_total    # 0.529
BETA_ANA  = _products["vegf"] / _ana_total   # 0.471

# Categorical weights (all non-ANA domains)
WEIGHTS_CAT = {
    k: _products[k] / _total_product
    for k in ("hla_dr", "tgfb", "cd10_alpl", "aetiology", "gmcsf", "nets")
}

# ANA raw weight before collinearity correction
```

```
W_ANA_RAW = _ana_total / _total_product   # ~0.588

# Gamma sensitivity range (no published rho(NLR, VEGF) in HCC exists)
GAMMA_RANGE = [0.00, 0.10, 0.20, 0.30, 0.40]

# Incremental utility benchmark (IDI threshold from PMC12287231)
IDI_BENCHMARK = 0.013   # 1.3%; minimum IDI over NLR-inclusive base model
C_INDEX_BENCHMARK = 0.013   # delta-C-index minimum


# ──────────────────────────────────────────────────────────────────────
# Sigmoid transformations (parameters from empirical cutoff distributions v3/v4)
# ──────────────────────────────────────────────────────────────────────

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 (range 2.3-5.0).
    k=1.02: derived algebraically from constraint f(5.0)=85.
    """
    return 100.0 / (1.0 + math.exp(-1.02 * (nlr - 3.3)))


def f_vegf(vegf_pg_ml: float) -> float:
    """
    Serum VEGF -> 0-100. f(x) = 100 / (1 + exp(-2.58*(x-270)/270))
    x0=270 pg/mL: cluster centre of published prognostic cutoffs (225-285).
    k=2.58: derived from f(125)=20 (healthy controls, Guo J et al. PMC3555251).
    """
    return 100.0 / (1.0 + math.exp(-2.58 * (vegf_pg_ml - 270.0) / 270.0))


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

    Collinearity mechanism: circulating neutrophils secrete VEGF, elevating
    serum levels independently of tumour VEGF production.
    No published rho(NLR, VEGF) in HCC exists; gamma reported as sensitivity
    range [0.00, 0.40]. See Section 3.3 for rationale.
    """
    fn = f_nlr(nlr)
    fv = f_vegf(vegf_pg_ml)
```

```
    return ALPHA_ANA * fn + BETA_ANA * fv - gamma * (fn * fv / 100.0)


# ──────────────────────────────────────────────────────────────────────
# Categorical transformations (literature-anchored ordinal scales)
# ──────────────────────────────────────────────────────────────────────

def f_tgfb(s: str) -> float:
    """TGF-beta -> 0-100. Anchors: Fridlender et al. Cancer Cell; Chen et al. Gastro
    return {"absent": 5.0, "mild": 30.0, "moderate": 60.0, "active": 88.0}.get(s, 30

def f_aetiology(s: str) -> float:
    """HCC aetiology -> 0-100. Anchors: Teo J et al. JEM; Finn RS et al. NEJM (IMbra
    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:
    """CD10+ALPL+ -> 0-100. Anchor: Meng Y, Ye F, Nie P et al. J Hepatol."""
    return {"absent": 0.0, "not_documented": 0.0, "low": 30.0,
            "elevated": 72.0, "high": 90.0}.get(s, 0.0)

def f_nets(level: str, cith3: bool) -> float:
    """NET markers -> 0-100. Base: Shen XT et al. Exp Hematol Oncol. CitH3+ adds 7."
    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:
    """HLA-DR+ -> 0-100 (INVERSELY scored). Anchor: Wu Y et al. Cell (HCC n=357).""""
    return {"absent": 82.0, "low": 52.0, "present": 26.0, "high": 5.0}.get(s, 52.0)

def f_gmcsf(s: str) -> float:
    """GM-CSF -> 0-100. Anchors: Teo J et al. JEM; Leslie J et al. Gut."""
    return {"absent": 5.0, "mild": 38.0, "elevated": 78.0}.get(s, 5.0)


@dataclass
class TANPatientV5:
    nlr: float = 2.5
    vegf_pg_ml: float = 200.0
    tgfb_signal: str = "absent"         # absent | mild | moderate | active
    hcc_aetiology: str = "viral"        # viral | mash | alcohol | cryptogenic |
                                         # formerly_viral_cirrhosis
```

```
    cd10_alpl_signal: str = "absent"    # absent | not_documented | low |
                                         # elevated | high
    net_marker_level: str = "normal"    # normal | mild | elevated | high
    cith3_positive: bool = False
    hla_dr_signal: str = "absent"       # absent | low | present | high
    gmcsf_signal: str = "absent"        # absent | mild | elevated


@dataclass
class TANResultV5:
    pss_by_gamma: Dict[float, float]
    pss_default: float
    pss_range: Tuple[float, float]
    ci_lower: float
    ci_upper: float
    domains: List[dict]
    weight_note: str
    collinearity_note: str
    incremental_utility_note: str
    limitations: List[str] = field(default_factory=list)


def compute_tan_polarity_v5(patient: TANPatientV5,
                            n_sims: int = 5000,
                            seed: int = 42) -> TANResultV5:

    cat_scores = {
        "hla_dr":    f_hla_dr(patient.hla_dr_signal),
        "tgfb":      f_tgfb(patient.tgfb_signal),
        "cd10_alpl": f_cd10_alpl(patient.cd10_alpl_signal),
        "aetiology": f_aetiology(patient.hcc_aetiology),
        "gmcsf":     f_gmcsf(patient.gmcsf_signal),
        "nets":      f_nets(patient.net_marker_level, patient.cith3_positive),
    }

    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)
        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 for continuous inputs at gamma=0.20
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)))
    sims.append(min(100.0, W_ANA_RAW * g_ana(nlr_p, vegf_p, 0.20) + cat_weighted
sims.sort()
ci_lower = round(sims[int(0.025 * n_sims)], 1)
ci_upper = round(sims[int(0.975 * n_sims)], 1)

f_nlr_val  = f_nlr(patient.nlr)
f_vegf_val = f_vegf(patient.vegf_pg_ml)
naive_ana  = ALPHA_ANA * f_nlr_val + BETA_ANA * f_vegf_val
interaction = 0.20 * (f_nlr_val * f_vegf_val / 100.0)

domains = [
    {"name": "ANA (NLR+VEGF)",
     "f_nlr": round(f_nlr_val, 1), "f_vegf": round(f_vegf_val, 1),
     "naive_ana": round(naive_ana, 1),
     "g_ana_g020": round(g_ana(patient.nlr, patient.vegf_pg_ml, 0.20), 1),
     "interaction_penalty_g020": round(interaction, 1),
     "w_ana": round(W_ANA_RAW, 3),
     "wtd_g020": round(W_ANA_RAW * g_ana(patient.nlr, patient.vegf_pg_ml, 0.20),
     "precision_nlr": DOMAIN_EVIDENCE["nlr"][2],
     "precision_vegf": DOMAIN_EVIDENCE["vegf"][2],
     "n_nlr": DOMAIN_EVIDENCE["nlr"][3],
     "n_vegf": DOMAIN_EVIDENCE["vegf"][3]},
] + [
    {"name": k,
     "raw": round(v, 1),
     "weight": round(WEIGHTS_CAT[k], 4),
     "weighted": round(WEIGHTS_CAT[k] * v, 3),
     "precision": round(DOMAIN_EVIDENCE[k][2], 1),
     "n_source": DOMAIN_EVIDENCE[k][3],
     "method": DOMAIN_EVIDENCE[k][4]}
    for k, v in cat_scores.items()
]

weight_note = (
    "v5 weights use n-based SE imputation (SE=1/sqrt(n/4)) for domains lacking "
    "published CIs [Kambach et al. Ecology and Evolution 2020]. "
```

```
        f"NLR precision={DOMAIN_EVIDENCE['nlr'][2]:.0f} (published CI, n=9,952). "
        f"VEGF precision={DOMAIN_EVIDENCE['vegf'][2]:.1f} (n-based, n=481). "
        "ANA domain (NLR+VEGF) accounts for ~53% of final PSS at gamma=0.20. "
        "Molecular sub-programme domains collectively contribute ~47%."
    )

    collinearity_note = (
        f"ANA collinearity: naive g_ANA={naive_ana:.1f}, interaction penalty={intera
        f"(gamma=0.20), corrected g_ANA={g_ana(patient.nlr, patient.vegf_pg_ml, 0.20
        f"PSS range across gamma=[0,0.40]: {pss_range[0]:.1f} – {pss_range[1]:.1f} "
        f"(span {pss_range[1]-pss_range[0]:.1f} pts). "
        "No published rho(NLR, serum VEGF) in HCC exists; gamma remains a sensitivit
    )

    incremental_utility_note = (
        f"Incremental utility benchmark: IDI > {IDI_BENCHMARK*100:.1f}% and "
        f"delta-C-index > {C_INDEX_BENCHMARK:.3f} over NLR-inclusive base model. "
        "Derived from PMC12287231 (n=2286): composite inflammatory index added "
        "IDI=1.3% p=0.04 over NLR-alone clinical model in HCC. "
        "TAN-POLARITY must exceed this in prospective validation to justify complexi
    )

    limitations = [
        "UNVALIDATED: No patient-level OS, PFS, or ICI response data tested. "
        "Clinical utility unknown.",
        "MRNA PROXY LIMITATION: TCGA-LIHC uses VEGFA mRNA (not serum VEGF) and "
        "CIBERSORT neutrophil scores (not blood NLR). These are related but not "
        "equivalent measurements [Oncotarget 2017].",
        f"GAMMA UNCERTAINTY: PSS range = {pss_range[0]:.1f}–{pss_range[1]:.1f} "
        "across gamma [0,0.40]. Report as range, not point value.",
        "ANA DOMINANCE: NLR+VEGF = ~53% of PSS. Molecular domains are "
        "biologically meaningful but statistically underweighted by current evidence
        "SCENARIOS ARE RECONSTRUCTIONS: Profile descriptions from published cohort "
        "papers; not independent patient observations.",
    ]

    return TANResultV5(
        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,
        incremental_utility_note=incremental_utility_note,
        limitations=limitations,
    )
```

```python
def print_result_v5(result: TANResultV5, label: str):
    print("\n" + "=" * 80)
    print(label)
    print("=" * 80)
    print(f"PSS (gamma=0.20): {result.pss_default:.1f} / 100")
    print(f"PSS sensitivity (gamma 0-0.40): {result.pss_range[0]:.1f} - {result.pss_
    print(f"95% CI (MC n=5000, continuous inputs, gamma=0.20): [{result.ci_lower:.1f
    print("\nGamma sensitivity table:")
    for g, pss in result.pss_by_gamma.items():
        print(f"  gamma={g:.2f}  PSS={pss:.1f}")
    print(f"\nWeights: {result.weight_note}")
    print(f"\nCollinearity: {result.collinearity_note}")
    print(f"\nIncremental utility: {result.incremental_utility_note}")
    print("\nDomain decomposition:")
    d = result.domains[0]
    print(f"  ANA: f_NLR={d['f_nlr']:.1f} (n={d['n_nlr']}, prec={d['precision_nlr']:
          f"f_VEGF={d['f_vegf']:.1f} (n={d['n_vegf']}, prec={d['precision_vegf']:.1f
          f"naive={d['naive_ana']:.1f}, g(g=0.20)={d['g_ana_g020']:.1f}, "
          f"w={d['w_ana']:.3f}, wtd={d['wtd_g020']:.2f}")
    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}, prec={dom['precision']:.1f} (n={dom['n_so
    print("\n[LIMITATIONS]")
    for lim in result.limitations:
        print(f"  ! {lim}")


def demo():
    scenarios = [
        ("Scenario 1 — Viral HCC responder [Jost-Brinkmann F et al. APT PMID 3688335
         TANPatientV5(nlr=2.1, vegf_pg_ml=195, 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. + Teo J et al. JEM]"
         TANPatientV5(nlr=5.7, vegf_pg_ml=415, 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
```

```
        TANPatientV5(nlr=4.2, vegf_pg_ml=340, 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_v5(patient)
        print_result_v5(result, label)


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

TAN-POLARITY v5 is a pre-validation framework specification. It defines a scoring architecture, derives weights transparently using sample-size-based SE imputation where published CIs are unavailable, and pre-specifies the validation protocol and minimum incremental utility bar required to claim clinical utility. It makes no claim of clinical validity. The PSS is biologically motivated, methodologically auditable, and should not be applied to individual patients until the prospective validation in Section 5.2 is completed.