

STRATEGY 5 - VERSION 2.0

Situational Context Betting

A dual-score football intelligence framework separating contextual research quality from market-specific prediction confidence.

Purpose. Improve the original Strategy 5 after the 20 September 2026 review without overfitting a four-match sample. The update preserves the disciplined FIS framework and adds a separate Prediction Confidence Score (PCS), a Market Fit gate, and a stricter post-match audit.

UPDATED	STATUS	PRIMARY CHANGE
20 Sep 2026	Version 2.0	FIS is no longer allowed to stand in for outcome confidence.

Research-led framework. For analytical use; no score removes football uncertainty or price risk.

Executive upgrade

What changed after the 20 September review

The existing methodology already makes an important distinction: FIS measures contextual evidence, not talent, expected goals, win probability, or expected return. The problem in practical reporting was that high FIS values could still be read as if they were strong betting predictions. Version 2.0 removes that ambiguity.

Issue observed	Version 2.0 control	Reason
High FIS looked like high outcome confidence	Publish FIS and PCS separately	A strong context dossier can coexist with an uncertain match outcome.
One match narrative could dominate	Mandatory scenario stress test for FIS 80+	Reduces story-driven overconfidence.
Directional view and market choice were blended	Add Market Fit gate	A correct team lean can still be a poor BTTS, totals, or 1X2 bet.
Post-match review focused too much on result	Separate process, probability, and market grading	Avoids rewarding luck or punishing sound analysis for variance.
Small sample tempted immediate weight changes	Freeze FIS weights until minimum sample is met	Prevents overfitting to four fixtures.

Core rule: FIS answers, "How strong and well-supported is the situational context?" PCS answers, "How confident is the model in this specific market outcome?" Market Fit answers, "Is that confidence usable at the available price?"

20 September 2026 audit

What the four Premier League matches taught us

Fixture	Pre-match signal	Final score	Lesson for v2.0
Man City v Sunderland	Very high context / City-positive	5-3	Context and direction aligned, but the eight-goal scoreline still shows outcome variance.
Leeds v Crystal Palace	High context / Leeds-positive	0-0	Key warning: high FIS did not justify high certainty on scoring or outright outcome.
Bournemouth v Liverpool	Liverpool-positive; goals angle also attractive	0-1	Directional read worked; goals angle failed. Market-specific confidence must be separate.
Fulham v Man United	Low 1X2 conviction; BTTS comparatively stronger	1-1	The market-selection layer added more value than a simple winner call.

These four matches are diagnostic, not a calibration sample. Version 2.0 therefore changes workflow and interpretation, not the underlying category weights. Formal reweighting remains locked until the paper-test database is large enough for out-of-sample evaluation.

The new five-layer architecture

Layer	Question	Output
1. Baseline	What does the stable model / no-margin market say before context?	1X2 and relevant market probabilities
2. Context dossier	What verified circumstances differ from normal?	Sourced situational evidence
3. FIS	How strong, current, and corroborated is the contextual position?	Team FIS, Delta FIS, research priority
4. PCS + Market Fit	How confident are we in this particular market, and is it robust enough?	PCS, scenario range, market status
5. Audit	Did the process, forecast, and market choice perform as expected?	Calibration, CLV, error category, version log

Non-negotiable: no single layer may substitute for another. In particular, FIS cannot be converted directly into a win probability.

Why this matters

The earlier strategy already states that a large context differential is a research queue rather than an outcome call. Version 2.0 operationalises that statement in every published report by requiring two visible numbers: FIS and PCS. This keeps the intelligence score useful while preventing readers from treating 88 FIS as "88% likely".

Football Intelligence Score - retained

FIS remains a context-quality score

Category	Weight	v2.0 treatment
Stakes & incentives	18%	Unchanged; require objective + mechanism, not "must win" wording.
Availability & lineup leverage	22%	Unchanged; confirmed absences/returns and replacement quality matter.
Recovery, travel & rotation	15%	Unchanged; use rest, minutes, travel and depth rather than generic fatigue.
Tactical & managerial fit	15%	Unchanged; score observed mechanisms, not narrative labels.
Club environment	10%	Unchanged; extreme scores require strong sourcing.
Information integrity	10%	Unchanged; must reach at least 50 for context-led use.
Market intelligence	10%	Unchanged; measures what appears priced, not whether a bet is automatically valuable.

FIS formula: weighted average of category sub-scores on the 0/25/50/75/100 scale. Team scores remain independent; Delta FIS = FIS home - FIS away.

Interpretation bands

Band	Meaning	Required response
0-44	Adverse / unclear	Do not construct a positive narrative.
45-59	Neutral / mixed	Retain baseline; no context-led thesis.
60-69	Research priority	Challenge evidence and check market timing.
70-79	Strong documented context	Second-source review plus price comparison.
80-100	Exceptional context	Mandatory adversarial review; identify at least two failure paths.

Prediction Confidence Score (PCS)

New in version 2.0: market-specific confidence

PCS is a 0-100 score attached to a **specific market proposition**, such as Home Draw No Bet, Away +0.5, Over 1.5, BTTS Yes, or Under 3.5. It is not derived from FIS alone. It summarises how stable the forecast is after combining the baseline model, market benchmark, context challenge, lineup certainty, and scenario robustness.

Component	Weight	What earns a high score
Baseline model agreement	30%	Independent model materially supports the proposition.
Market benchmark agreement	20%	No-margin market is broadly consistent, or a documented residual edge exists.
Lineup / information certainty	15%	Key inputs are confirmed rather than assumed.
Scenario robustness	20%	Proposition survives plausible game-state alternatives.
Context support	10%	FIS evidence supports the same mechanism without double counting.
Price stability / execution quality	5%	Price has not moved beyond the acceptable entry band.

PCS discipline: until the strategy has sufficient historical records, PCS is a structured confidence grade rather than a calibrated probability. Do not print a percentage sign beside it.

PCS bands

PCS	Interpretation	Use
0-49	Weak / unstable	No selection; analysis only.
50-59	Marginal	Watchlist or paper test.
60-69	Moderate	Requires favourable price and no unresolved lineup issue.
70-79	Strong	Eligible for shortlist if Market Fit also passes.
80-100	Exceptional	Rare; requires model, market, lineup and scenario agreement.

Market Fit gate

A correct football view can still be the wrong bet

Gate	Pass condition	Fail response
Price	Offered odds remain above the model-defined minimum after margin/uncertainty buffer.	No bet / reprice.
Lineup	No critical unresolved dependency close to kickoff.	Mark Pending.
Scenario robustness	At least two plausible alternate game scripts still leave the proposition viable.	Downgrade PCS or choose a more robust market.
Correlation / duplication	The same evidence is not being counted in baseline, FIS, and PCS multiple times.	Remove duplicate support.
Market specificity	The proposition matches the evidence mechanism.	Do not infer totals/BTTS from a team-direction view without separate support.
Information quality	Both teams meet information floor; key claim is timestamped.	Insufficient Intelligence.

Lesson from Bournemouth 0-1 Liverpool: a Liverpool-positive direction and a goals/BTTS proposition are different forecasts. Version 2.0 scores them separately.

Decision matrix

How FIS and PCS interact

FIS	PCS	Status	Meaning
<60	Any	Baseline-led	Context is not strong enough to drive the case.
60-79	<60	Research only	Interesting context, weak proposition confidence.
60-79	60-69	Watchlist	Potential case; price and lineup must be favourable.
60-79	70+	Shortlist	Context and specific market are aligned.
80+	<70	Challenge hard	Exceptional story but insufficient market confidence - classic overconfidence risk.
80+	70+	High-priority review	Requires adversarial check, price validation, and failure scenarios before any action.

This table directly prevents a repeat of the Leeds-Palace problem: a very high contextual score can remain a research priority without being treated as a high-confidence outcome.

Scenario stress test

Every PCS of 70 or higher must survive three explicit game scripts. The analyst must state how the proposition behaves under each script rather than assuming one expected match narrative.

Scenario	Example question	Required note
Expected script	What happens if the favourite establishes control as projected?	Why the market should still land.
Adverse script	What if the opponent scores first, presses differently, or the game becomes low tempo?	Whether the proposition survives or fails.
Lineup / role surprise	What if one key player is benched or a tactical role changes?	Which probability or PCS component changes.

Automatic downgrade: if one plausible adverse script destroys the proposition and the price does not compensate, cap PCS at 69 until stronger evidence arrives.

New PremierBettingTips report format

What every published match row should show

Field	Example display	Why it is required
FIS	82 / 100	Context strength only.
PCS	68 / 100 - Liverpool +0.0 AH	Confidence attached to a specific proposition.
Baseline	H 31% / D 27% / A 42%	Pre-context anchor.
Market benchmark	No-margin A 40%	Shows what price already implies.
Primary evidence	2-4 timestamped bullets	Makes the case auditable.
Failure case	"If Liverpool cannot progress centrally..."	Forces a disconfirming view.
Market Fit	PASS / WATCH / FAIL	Prevents FIS-only selections.
Uncertainty	Low / Medium / High	Makes missing information visible.

Ranking rule: reports may still be ordered highest-to-lowest by FIS for research priority, but every row must also show PCS and Market Fit so FIS cannot be mistaken for predicted win chance.

Evidence protocol - tightened

Tier	Source type	v2.0 scoring rule
Tier 1	Official club / league / competition / press conference / verified squad list	May support extreme FIS ratings when the football mechanism is clear.
Tier 2	Named specialist journalist / strong local reporting with direct quotes	Supports moderate ratings; extreme rating needs corroboration.
Tier 3	General media / sourced aggregator / informed commentary	Lead only; cannot justify extreme score.
Tier 4	Anonymous posts / unsupported social claims / rumour accounts	Do not score.

Additional v2.0 requirement

For PCS 70+, the analyst must record the publication time, access time, affected team, claimed mechanism, and whether the information was already reflected in the price. Confirmed, probable and speculative player information remain separate categories.

Pre-match execution checklist

Step	Required action
1	Capture no-margin market probabilities from at least three reputable prices before reading narratives.
2	Generate baseline probabilities from the fixed model specification.
3	Build and source the context dossier for both teams.
4	Score FIS independently for home and away; calculate Delta FIS and research priority.
5	Write at least one disconfirming case for each strong narrative.
6	Choose the exact market proposition to test.
7	Calculate PCS for that proposition; do not reuse FIS as the confidence number.
8	Run the three-script stress test.
9	Apply Market Fit: price, lineup, information floor, execution band.
10	Timestamp and lock the pre-match record. Do not rewrite it after kickoff.

Near-kickoff refresh: repeat odds and lineup checks after material news. Record the new snapshot as a separate update; never overwrite the original analysis.

Post-match audit

Grade the process, not only the final score

Dimension	Question	Grade
Evidence accuracy	Were the key situational facts actually correct at kickoff?	Pass / Partial / Fail
Baseline quality	Was the pre-context probability reasonable relative to closing market?	Good / Neutral / Weak
FIS discipline	Did the score reflect context rather than outcome certainty?	Pass / Fail
PCS quality	Did confidence match how robust the proposition really was?	Well calibrated / Too high / Too low
Market choice	Did the selected market match the evidence mechanism?	Strong / Mixed / Poor
Price execution	Was the entry better or worse than close?	Positive / Flat / Negative CLV
Outcome	Did the proposition land?	Win / Push / Loss
Error class	If wrong, why?	Variance / bad data / model miss / tactical miss / price miss / overconfidence

A losing bet with sound evidence, reasonable calibration and positive closing-line value should not automatically trigger a model change. A winning bet produced by weak evidence or bad pricing should not be treated as proof that the process worked.

Performance tracker - fields to add

Group	Fields
Fixture	Date, league, teams, kickoff, analyst timestamp
Baseline	Home/Draw/Away probabilities, model version, price source, overround method
Context	Each FIS component, source tier, FIS home, FIS away, Delta FIS, research priority
Prediction	Exact proposition, PCS, scenario notes, acceptable price band, Market Fit status
Execution	Opening observed price, selected price, closing price, CLV
Outcome	Score, market result, return in units if paper-testing
Audit	Evidence correctness, PCS calibration band, error class, lesson, version tag

Minimum review cadence

Review descriptive performance every 25 records, but do not alter weights merely because a short run looks good or bad. Consider formal calibration only after at least 100 timestamped records across different leagues, price bands and FIS ranges, then validate changes on a hold-out sample.

Version-control rules

Rule	Control
Freeze FIS weights	No category-weight changes before the minimum audit sample unless a definition is clearly defective.
Separate score versions	Store FIS rubric version and PCS rubric version with every record.
No hindsight edits	Original pre-match score remains immutable; later information is appended.
One change at a time	If calibration supports a change, alter one major element per version so the effect can be tested.
Hold-out validation	Any new probability adjustment must outperform the unchanged baseline out of sample.
Retire narratives, not evidence	If a category repeatedly adds no information, reduce its operational influence only after sufficient data.

Version 2.0 does not claim the FIS weights are wrong. It fixes an interpretation and workflow problem exposed by the match review while preserving the original strategy's anti-overfitting discipline.

Reusable match intelligence sheet - v2.0

Section	Entry
Fixture / timestamp	_____
Baseline 1X2	Home ____% Draw ____% Away ____% Model/version: _____
Market benchmark	No-margin probabilities / exact price snapshot: _____
Home FIS / Away FIS	_____ / _____ Delta FIS: _____ Research priority: _____
Primary context thesis	_____
Disconfirming case	_____
Exact market proposition	_____
PCS	_____ / 100 Scenario robustness: Low / Medium / High
Market Fit	PASS / WATCH / FAIL Minimum acceptable odds: _____
Lineup status	Confirmed / Probable / Pending
Decision status	Baseline only / Further research / Paper-test shortlist / Abstain
Post-match audit	Result: _____ CLV: _____ Error class: _____

Implementation roadmap

Stage	Objective	Deliverable
1. Deploy v2.0	Add PCS and Market Fit to all new FIS reports.	Dual-score HTML/PDF template.
2. Paper test	Log at least 100 pre-match records with locked timestamps.	Clean performance dataset.
3. Audit	Analyse calibration, CLV and errors by PCS/FIS band and market type.	Error and reliability report.
4. Calibrate	Only if supported, estimate bounded relationships between indicators and residual forecast error.	Version 2.1 with documented evidence.
5. Validate	Test the change on data not used to design it.	Hold-out comparison vs v2.0 baseline.
6. Maintain	Repeat source, price and post-match controls every cycle.	Stable intelligence process.

Bottom line. Version 2.0 keeps the strongest part of Strategy 5 - disciplined, auditable situational research - while preventing a high FIS from being mistaken for a high-probability bet. The workflow is now: establish the baseline, verify context, score FIS, score the exact market with PCS, pass Market Fit, lock the record, and audit the process.

References and source basis

Original research retained

1. Page, R. M., Field, A., Langley, B., Harper, L. D., & Julian, R. (2023). The Effects of Fixture Congestion on Injury in Professional Male Soccer. *Sports Medicine*, 53, 667-685.
2. Julian, R., Page, R. M., & Harper, L. D. (2021). The Effect of Fixture Congestion on Performance During Professional Male Soccer Match-Play. *Sports Medicine*, 51, 255-273.
3. Koning, R. H., & Zijm, R. (2023). Betting market efficiency and prediction in binary choice models. *Annals of Operations Research*, 325, 135-148.
4. Constantinou, A. C., & Fenton, N. E. (2012). Solving the problem of inadequate scoring rules for assessing probabilistic football forecast models. *Journal of Quantitative Analysis in Sports*, 8(1).
5. Dixon, M. J., & Coles, S. G. (1997). Modelling association football scores and inefficiencies in the football betting market. *Journal of the Royal Statistical Society: Series C*, 46(2), 265-280.

Version 2.0 review basis

The workflow changes were prompted by the Premier League review on 20 September 2026: Manchester City 5-3 Sunderland; Leeds United 0-0 Crystal Palace; AFC Bournemouth 0-1 Liverpool; Fulham 1-1 Manchester United. The sample is used diagnostically only and is not treated as sufficient evidence for statistical reweighting.

Original Strategy 5 source principle: FIS is a research-priority and abstention tool until enough pre-match records exist for calibration. Version 2.0 preserves that principle and makes it explicit in the published reporting format.