

Pricing the Unpriceable

Frequently Asked Questions

A companion to the FSPI Fair Value Framework

FutureSports Index Services · May 2026

Overview

This document addresses frequently asked questions about the fair value framework for futures that reference FutureSports Performance Indexes (FSPI), referred to in this document as FSPI futures for brevity. It accompanies the white paper *Pricing the Unpriceable: A Fair Value Framework for Sports Performance Index Futures* and summarizes that paper in question-and-answer form. The white paper remains the authoritative reference on matters of detail; questions of contract specification, settlement, and exchange operation are governed by the listing exchange's rules and contract terms.

FSPI futures are a calculated benchmark derivative. Because no investable spot underlies them, the conventional cost-of-carry approach to futures pricing does not apply. The framework summarized here derives fair value from an expectation of the future index level rather than from arbitrage. This is the established approach for calculated benchmarks that lack an investable spot, drawing on the pricing practice of VIX futures, weather derivatives, and crypto reference rate futures.

It is offered as a shared analytical reference, not a proprietary black box: each input and each step is intended to be transparent and reproducible by the participant. It is written for the market makers, risk and clearing teams, benchmark and product staff, and institutional users engaged with the product.

The basics

What is an FSPI futures contract?

An FSPI futures contract references a FutureSports Performance Index, a rules-based benchmark that aggregates officially reported sports performance statistics into a single published value under a defined methodology. The indexes are continuously calculated, published, and licensed as reference values, and they are designed to align with the IOSCO Principles for Financial Benchmarks. The futures contract is listed on a regulated derivatives exchange and settles against the published index level on a defined calculation date.

Why can't these be priced with the standard futures formula?

The conventional fair value formula for equity index and commodity futures is anchored in the cost of carrying a position in the underlying to expiry, a relationship enforced by arbitrage: a participant can buy

the underlying, hold it, and deliver it against the contract. FSPI futures lack every input to that calculation. There is no spot market in team performance, nothing to store, nothing to lend, and no convenience yield. The cost-of-carry components are not merely zero; they are undefined, because the underlying is a calculated value rather than a tradable asset.

How is fair value determined instead?

Fair value is derived from a mathematical expectation rather than from arbitrage. It is expressed as the current observable index level plus the expected accumulation of index points between now and contract expiry. There is no cost-of-carry term, no convenience yield, and no foregone dividend. This is the same structural form used to price VIX futures and weather derivatives, both of which reference calculated benchmarks with no investable spot, and it is well established in that landscape rather than novel.

The pricing framework

What is the core fair value equation?

Fair value equals the current index level plus the expectation, under the relevant pricing measure, of the change in the index level over the life of the contract. In words: where the index sits today, plus how much it is expected to accumulate between now and expiry. The relationship is anchored in the dynamics of that expectation rather than in any financing relationship.

How is the expected accumulation calculated?

It decomposes into three additive components that mirror the three ways the FSPI methodology awards points. The first is in-game statistical accumulation, the points earned from game events across the remaining schedule, which has the structural form of a weather derivative. The second is end-of-game milestone probabilities, fixed point values for outcomes such as a win or a shutout, priced as a book of binary or digital options. The third is period and competitive milestones, such as leading a division or winning a championship, realized only at the end of a period or on a defined competitive event and estimated through Monte Carlo simulation across the league. Each component is modelled separately and the three are summed.

Structural features

What is the annual reset and why does it matter?

Under the current methodology, all team index values are reset to a defined base level approximately one week after the end of each postseason. For a contract that expires within a season, the reset is irrelevant and the standard framework applies directly. For a contract that would span the reset, expected

accumulation begins from the post-reset base level rather than from the current index level, so current-season performance does not carry across. The reset functions as a scheduled, exogenous event that pricing must account for.

Can fair value, or a traded price, be negative?

The methodology applies no floor within a season, so a contract on a deeply underperforming team late in a season can rationally carry a negative model fair value. The framework handles that case the same way it handles positive values. Whether a negative number can also trade and settle is a separate question, governed by contract design and exchange rules; the May 2020 WTI crude oil settlement at a negative price demonstrated that the clearing and risk-management infrastructure for negative pricing already exists at major derivatives exchanges.

How are FSPI futures different from single-event or “event” contracts?

They differ structurally in several ways. The underlying is a published, governed benchmark with formal oversight and documented procedures, administered independently of the trading of contracts that reference it. Settlement references the benchmark level on a calculation date, not the outcome of any single game; no single event causes the benchmark to settle or terminate. The contracts are listed and centrally cleared with daily mark-to-market, margin, and position limits. Their horizons span many contributing events, so the payoff is an aggregation rather than a single binary outcome. This description is for pricing and risk-management purposes only and should not be read as a characterization of regulatory status, which is a matter for regulators, exchanges, and counsel.

Settlement and data

What does a contract settle against, and when?

Settlement references the official Index Close, which is published at 9:00 AM Central Time on the calculation day following each game day (T+1), reflecting final verified league data from the prior day. Monthly and seasonal milestones are applied on a T+1 basis at defined times after the relevant observation period concludes. Games that begin on one date but finish on the next are attributed to the date the game began. These rules are set out in the FS Index Calculation and Publication Schedule. Last-trading-day specifications and the timing of final settlement against the T+1 close are matters of contract specification addressed by the exchange.

How are later corrections to official statistics handled?

Official league amendments issued after a game are treated as discrete statistical events and incorporated into the index in accordance with the Points Attribution Table. They are reflected in the official end-of-

day index value on the calculation day following publication and do not retroactively affect historical milestone determinations. The treatment of amendments issued after a contract's final settlement is an operational matter addressed in contract specifications.

Where does the underlying data come from, and what if it is wrong or late?

The methodology draws on official league data through a defined hierarchy of inputs, with secondary sources permitted only when official data is temporarily unavailable. Three data risks are relevant to futures pricing. Data finality risk is the possibility that statistics treated as official are later corrected; in practice it can produce small ex-post jumps between a provisional close and the final close. Data latency risk is the gap between game-end and official publication; it can widen intra-day basis when the market trades ahead of the final close. Settlement-source risk is the possibility that a feed disruption requires reliance on an alternate source with different accuracy or timing; it can cause discrete basis moves if the alternate source diverges from the primary feed. Each moves basis independently of model fair value, and they are best monitored as distinct categories rather than aggregated into a single line item.

Governance and administration

Who administers the index, and how is its independence protected?

The index is administered as a published, IOSCO-aligned benchmark with a defined methodology, formal governance, and oversight committees, and is administered independently of the trading of contracts that reference it. In a calculated benchmark market there is no arbitrage channel to enforce price, so the market's confidence in the administrator feeds directly into basis. Investment in IOSCO-aligned governance, independent oversight, transparent methodology-change procedures, robust complaint handling, and input-data integrity therefore supports tighter basis and better execution, making benchmark integrity a commercial input rather than only a regulatory one.

Can the methodology change, and what protects integrity if it does?

Yes. The benchmark operates under documented procedures for methodology change, error correction, and complaint handling, and changes are subject to governance oversight with material changes communicated in advance. The framework itself accommodates change; introducing a within-season floor, for example, would be a reversible methodology change that the pricing approach can be extended to handle. The treatment of open positions when a methodology change occurs is governed by the exchange's contract specifications, not by the benchmark methodology. Because a calculated benchmark has no arbitrage anchor, the credibility of these procedures is itself reflected in basis and spreads.

Use and considerations

How would the index be used in practice?

FSPI futures provide a rules-based reference to the cumulative on-ice performance of a team across a season, expressed in a single continuously calculated index rather than through any single game outcome. The framework is built to support the participants who would engage with such a reference: liquidity providers on the trading side, and benchmark users seeking exposure or analytical use cases on the other.

The character of demand, and the correlation and capacity properties of FSPI exposure relative to other holdings, will be established as the market develops. The framework provides the pricing scaffolding; the commercial profile is a question the early market will answer.

How may a participant hedge or risk-manage exposure?

Several proxy categories carry real signal on FSPI exposure. Sportsbook lines on team outcomes provide continuously updated, market-aggregated probabilities for the binary milestones that drive the end-of-game component. Prediction markets on listed milestones do the same for the period and competitive component, often as a direct observable input. Instruments on correlated team indexes within the same league allow within-league offset, and news-driven inputs (rosters, lineups, goaltender starts) can be partially neutralized through other exposures. Exchange-traded funds referencing FSPI futures, once listed, would add a further mechanism.

Two caveats apply. The availability and usability of any proxy, including sportsbooks and prediction markets, depends on the legal, regulatory, and internal policy constraints applicable to each participant. And no proxy provides a complete hedge: a meaningful portion of risk is likely to be warehoused, particularly early in the market's life, consistent with the liquidity risk premium. Proxy hedging should become more effective as adjacent markets and correlation structures mature.

Will the market price always equal the model fair value?

No. Observed market price equals the model fair value plus three premia: a variance risk premium, a liquidity risk premium, and a credibility risk premium. The model fair value is the modellable part; the premia are properties of the market environment. The decomposition is useful operationally, because a widening that affects all teams uniformly points to liquidity or credibility, while a widening concentrated on one team points to a change in that team's model inputs. Basis is also expected to be wider early in the market's life, when participants are few and inventory risk is high, and to narrow as participation broadens.

What inputs feed the calculation, and is there a calculator?

The framework is agnostic to the specific modelling approach and can draw on league-licensed statistical feeds, market-implied probabilities from established lines, prediction-market prices on listed milestones, and participants' own internal models. A natural product expression is a fair value calculator that accepts these inputs and returns a model fair value, with user-selectable approaches to estimating per-game expectations so that the assumptions behind any output remain transparent.

Does the framework extend to other leagues?

Yes. The framework was developed using NHL FSPI as the working example, but it generalizes by substituting league-specific inputs: the sets of statistical categories and milestones, and their point values, are defined by each league's methodology. The structural form, the three-component decomposition, and the treatment of features such as the annual reset carry over. Whether the relative size of the three components differs across leagues is an empirical question that will be answered as the family of indexes expands.

The worked example

What does the worked example in the white paper show?

The paper applies the full framework end-to-end to a generic team referred to as Team A, at a mid-season point with 35 regular-season games remaining and a current index level of 8,400. The three components contribute approximately 1,939 points from statistical accumulation, 756 from end-of-game milestones, and 74 from period and competitive milestones, for an illustrative model fair value of about 11,169 index points, with statistical accumulation dominating at roughly 70 percent. All inputs are illustrative placeholders chosen for exposition: the statistical multipliers are simplified composites, the milestone values are drawn from the published methodology, and the figures do not represent any actual team, schedule, probability, or contract price. The example exists to show that the framework, once given specific inputs, produces a single fair value through a transparent and reproducible calculation.

If one component is so small, why does it matter?

The period and competitive milestone component contributes only about 3 percent of fair value in the example, but it carries a disproportionate share of repricing risk. A single injury, trade, or run of unexpected results can shift playoff and division probabilities sharply, with a directly proportional effect on this component, whereas the statistical and milestone components are smoothed across many remaining games and move more slowly. Exposure to news risk should therefore weight the competitive component more heavily than its share of fair value suggests.

Disclosures

This document has been prepared for informational purposes only. It does not constitute an offer to buy or sell, or a solicitation of an offer to buy or sell, any security, futures contract, derivative, or other financial instrument, nor does it constitute investment, legal, tax, or accounting advice.

All numerical inputs and outputs referenced from the worked example are illustrative placeholders chosen for exposition. They do not represent any actual team, schedule, probability assessment, or contract price, and are not intended to forecast or imply the value of any actual or contemplated FSPI futures contract.

FutureSports Performance Indexes (FSPI) are calculated benchmarks administered in accordance with the IOSCO Principles for Financial Benchmarks. The framework summarized here is a pricing framework for derivatives referencing FSPI; it is not a description of the FSPI methodology itself. Market participants seeking authoritative information on the methodology should refer to the published FSPI methodology documentation, and on the framework to the white paper this document accompanies.