Pricing the Unpriceable

A Fair Value Framework for Sports Performance Index Futures

Jodie Gunzberg, CFA

May 2026

Executive Summary

The arrival of futures contracts referenced to sports performance indexes raises a question that conventional derivatives pricing was not designed to answer. There is no investable spot, no cost of carry, and no arbitrage channel through which the futures price can be tied to an underlying asset. The standard equity index and commodity futures fair value formula does not apply, yet these contracts will trade and settle, and the participants who price, clear, and trade them need a shared analytical framework for how that should be done.

This paper develops that framework using the FutureSports Performance Indexes (FSPI) for the National Hockey League as the working example. We show that fair value in this category cannot be derived from arbitrage to a tradable underlying. It must instead be derived from an expectation of the terminal index level, decomposed in a way that mirrors the structure of the underlying methodology. This framing is not novel in the broader benchmark futures landscape; it is the framework used for VIX futures, weather derivatives, and other calculated benchmarks that lack an investable spot. What is new is the application of this framework to sports performance data, and the technical work required to map the index methodology onto a tractable pricing model.

We develop the framework in stages. We first establish why the conventional cost-of-carry formula fails for FSPI and identify the appropriate analogue family. We then construct the core fair value equation and decompose the expected accumulation term into three additive components, each with its own modelling machinery. We address two structural features of the FSPI methodology that have direct pricing consequences: annual resetting and the absence of a within-season floor. We treat basis dynamics, the variance risk premium, and the role of administrator credibility as a measurable input to spread behavior. We close with a worked illustrative calculation that applies the full framework end-to-end to a representative team, with explicit placeholder inputs.

Key highlights

- Sports performance index futures occupy a category of calculated benchmark derivatives that includes VIX futures, weather derivatives, and crypto reference rate futures. The pricing framework appropriate to this category is structurally different from equity index and commodity futures pricing, which are anchored in the cost of carrying a tradable underlying.

- Fair value can be expressed as the current observable index level plus the expected accumulation of index points between the current date and contract expiry. There is no cost-of-carry term, no convenience yield, and no foregone dividend. The fair value relationship is derived from expectation rather than from arbitrage.
- Expected accumulation decomposes into three additive components: in-game statistical accumulation, end-of-game milestone probabilities, and period and competitive milestones. Each component requires a different modelling approach, drawn respectively from weather derivative pricing, insurance and digital option pricing, and joint Monte Carlo simulation.
- Two structural features of the FSPI methodology have direct mathematical implications for futures pricing: an annual reset of all index values to a defined base level, and the absence of a floor mechanism within a season. Cross-season contracts are decoupled from current-season performance. Within-season contracts on underperforming teams can have negative fair value, which is mathematically clean and operationally precedented.
- Administrator credibility, expressed through governance quality, methodology transparency, and operating track record, is a measurable input to basis behavior in calculated benchmark markets. Investment in benchmark integrity is a commercial input, not just a regulatory one, with direct effect on the spreads at which market makers will quote and the appetite of institutional end users to allocate.
- We provide a worked illustrative calculation for a representative NHL team that applies the full framework end-to-end. The example uses placeholder inputs and is intended to show the mechanics rather than to forecast any actual contract value.

1. Introduction: A New Category of Calculated Benchmark

The launch of futures contracts referenced to sports performance indexes represents a meaningful expansion of the calculated benchmark family. The FutureSports Performance Indexes are rules-based, methodology-driven aggregations of officially reported sports performance statistics, designed to align with the IOSCO Principles for Financial Benchmarks.[1] The indexes are continuously calculated, published, and licensed for use as reference values. Contracts referencing these indexes are being brought to market through a regulated derivatives exchange.

This paper is written for the practitioners who will need to price, risk-manage, govern, and clear these contracts in the absence of a traditional arbitrage anchor: market makers and structurers on derivatives desks, risk and clearing teams at exchanges and clearing houses, benchmark sponsors and product staff, and institutional end users evaluating allocation. The framework that follows is intended to provide a shared analytical starting point for that audience.

From the perspective of benchmark administration, this is a familiar undertaking. The architecture, the aggregation of real underlying activity into a synthesized published value supported by governance, oversight, and methodology change procedures, is well established across commodities, volatility, weather, credit, and digital assets. From the perspective of derivatives pricing, however, the situation is less familiar. Sports performance indexes do not have an investable spot. They are not derived from prices observed in a market. The conventional fair value formula for equity index and commodity futures, anchored in the cost of carrying a position in the underlying to expiry, does not provide a meaningful starting point because the underlying activity is not the kind of thing that can be held.

This is not unique to sports indexes. The same observation applies to VIX futures, which reference a calculated volatility index with no investable spot.[2] It applies to weather derivatives, where settlement is a cumulative integral of observed measurements over a defined period. It applies to early CME bitcoin futures, which traded for years before regulated spot bitcoin investment vehicles became available. In each of these markets, fair value is derived not from arbitrage but from expectation, and basis behavior reflects the dynamics of that expectation rather than financing costs.

This paper develops the fair value framework for sports performance index futures by drawing on these precedents and applying them to the specific structure of the FSPI methodology. The structure of the paper proceeds as follows. Section 2 establishes why the conventional cost-of-carry formula fails for FSPI and identifies the appropriate analogue family. Section 3 surveys the calculated benchmark futures markets that constitute the appropriate set of precedents and identifies what each contributes to our framework. Section 4 develops the core fair value equation, defines the notation, decomposes the expected accumulation term into its three modellable components, and identifies a whole-instrument analogue. Section 5 addresses three structural features of the FSPI methodology that have direct pricing implications: the annual reset and its interaction with the listed-contract calendar, the absence of a within-season floor, and the structural distinctions between FSPI futures and event contracts. Section 6 considers basis dynamics, the variance risk premium, and the role of administrator credibility. Section 7 presents a worked illustrative calculation that applies the full framework to a representative team. Section 8 discusses implications and open questions for the market to resolve. References, endnotes, and disclosures follow.

2. Why Conventional Fair Value Fails for FSPI

The conventional fair value relationship for equity index and commodity futures is anchored in the cost of carrying a position in the underlying. A market participant who can buy the underlying basket and hold it to expiry can replicate the futures payoff. Arbitrage forces the futures price to converge to the spot price compounded at the risk-free rate, adjusted for any income foregone. The relationship is mechanical, anchored in a tradable underlying, and pinned by an arbitrage channel that operates whenever observed prices deviate from the no-arbitrage value.

The standard expression of this relationship for an equity index future is:

$$F(t,\ T) = S(t) \cdot e^{(r-q)(T-t)} \tag{1}$$

where $F(t, T)$ is the futures price at time t for expiry at time T, $S(t)$ is the observable spot price, r is the risk-free rate, and q is the dividend yield. The relationship is enforceable through cash-and-carry arbitrage. A participant observing the futures trade above this value can short the future, buy the spot, finance the position at r, and lock in a profit at expiry; below this value, the symmetric strategy applies.

Sports performance index futures lack the inputs to this calculation. There is no spot market in NHL team performance. A market participant cannot purchase a basket of goals, saves, hits, blocks, and takeaways, hold them to expiry, and deliver them against the contract. There is no commodity to store. There is no instrument to lend. There is no convenience yield because there is nothing physical to be conveniently held. The cost-of-carry components are not just zero; they are undefined, because the conceptual framework that produces them does not apply to an underlying that is itself a calculated value rather than a tradable asset.

This is not a defect of the FSPI methodology. It is a structural feature of any benchmark that synthesizes a published value from underlying activity that is not itself an asset. The same observation has been made and addressed in other calculated benchmark markets. The VIX index is calculated from the prices of S&P 500 options, but the VIX itself is not a tradable instrument; it is a computed value. CBOE bitcoin futures launched in 2017 against an underlying for which no regulated spot investment vehicle existed in the United States until 2024. Heating degree day weather contracts settle against a cumulative integral of temperature measurements. In each case, the question of how to price futures against a non-investable underlying was addressed by deriving fair value from expectation rather than from arbitrage.

The appropriate fair value relationship for FSPI is therefore not derived from a no-arbitrage condition. It is derived from a mathematical expectation. Section 4 develops the explicit form of that expectation. Before doing so, we survey the calculated benchmark markets whose pricing practice provides the relevant precedent.

3. The Calculated Benchmark Family

Calculated benchmark futures form a coherent family of derivatives products whose pricing depends on the dynamics of an expectation rather than on a financing relationship. While the underlying methodologies differ across markets, the structural form of fair value is the same. We survey the four most informative members of this family and identify what each contributes to a pricing framework appropriate for FSPI.

3.1 VIX futures

The VIX futures market provides the closest precedent for FSPI pricing. The VIX index is calculated from the implied volatilities of S&P 500 options under a published methodology.[2] There is no spot market in volatility itself. A market participant cannot purchase volatility, store it, and deliver it. The VIX futures price reflects the market's expectation of where the VIX will be at the contract's expiry, not a discounted spot price.

Two structural features of the VIX futures market are directly relevant to FSPI. First, the futures price typically differs from the level of the spot index by an amount that reflects expected mean reversion in volatility together with a variance risk premium. The VIX futures curve is upward-sloping in calm markets and inverted in stressed markets, reflecting these expectation dynamics.[2] Second, the VIX futures consistently trade above the realized level of the VIX over their contract lives, a persistent gap attributable to the variance risk premium that volatility sellers earn on average.[2] Both features are likely to have analogues in FSPI pricing.

3.2 Weather derivatives

Weather derivatives, including heating degree day and cooling degree day futures, settle against a cumulative integral of temperature measurements over a defined contract period. The settlement value is therefore the sum of contributions from each day in the period, with each day's contribution determined by a function of the observed temperature.

This structural form is directly applicable to FSPI. The FSPI value at any time is the cumulative sum of contributions from each game played up to that time, with each game's contribution determined by the methodology's points attribution table applied to the observed game statistics. The mathematical form is the same: a sum of contributions across a known schedule of contributing events. Fair value at any time before expiry is the accumulated value to date plus the expected accumulation across the remaining events. The schedule of contributing events is a known input. The per constituent contribution distribution is the modelled object.

Weather derivatives also illustrate that a benchmark futures market can clear, settle, and serve its intended purpose at relatively modest scale, provided the underlying methodology is sound and the administrator is credible. This is a useful precedent for the early stage of any new calculated benchmark market.

3.3 Crypto reference rate futures

CME bitcoin futures launched in December 2017 against the CME CF Bitcoin Reference Rate.[3] The reference rate had been published since November 2016 and was designed around the IOSCO Principles for Financial Benchmarks.[1] At the time of futures launch, spot bitcoin traded across a fragmented set of venues at materially different prices, with no consolidated reference embedded in the spot market and no regulated U.S. spot exchange-traded product. The CME CF reference rate itself was constructed precisely to address this absence of a consolidated spot benchmark.[3] Regulated U.S. spot bitcoin exchange-traded products did not appear until January 2024.[4] For more than six years, CME bitcoin futures traded against a calculated benchmark that consolidated a fragmented spot market into a single reference value.

This is the most direct precedent for the architectural question FSPI raises. The bitcoin futures market demonstrated that a regulated derivatives market can develop, clear, and grow against a calculated benchmark whose role is to consolidate a fragmented underlying market into a single reference value, provided the benchmark itself is robustly administered and the contract specifications are clear. The growth trajectory of bitcoin futures during the 2018 through 2024 period, before any regulated U.S. spot exchange-traded product was available, provides empirical evidence that the absence of a consolidated, regulated spot reference is not a barrier to a viable derivatives market, provided the calculated benchmark fills that role.

3.4 Carbon allowance markets

Carbon allowance markets contribute a distinct piece of the framework. Allowance prices reflect not only supply and demand fundamentals but also the market's assessment of administrator credibility. When the administrator's credibility is high, basis is tight and curves are stable. When credibility is questioned, basis widens and curves become volatile. The European Union Emissions Trading System provides documented evidence of this dynamic across multiple regulatory cycles.[6]

This is empirically important for FSPI. Calculated benchmark markets are systematically more sensitive to administrator credibility than markets in tradable underlyings, because there is no arbitrage channel to enforce a particular price relationship and basis is therefore more dependent on the market's confidence in the calculation. The investment that benchmark administrators make in compliance, governance, and transparent methodology change procedures is, in a calculated benchmark market, a direct input to basis stability and therefore to the spreads at which futures contracts will trade.

3.5 What FSPI inherits from the family

Taken together, these four markets provide the full set of components needed to build a pricing framework for FSPI. From VIX, we inherit the structural form of fair value as an expectation under a risk-neutral measure, together with the observation that a variance risk premium will be a meaningful component of basis. From weather derivatives, we inherit the decomposition of the expectation into a sum of per constituent contributions across a known schedule. From crypto reference rate futures, we inherit the architectural precedent that a regulated derivatives market can operate viably against a calculated benchmark that consolidates a fragmented underlying market into a single reference value.

From carbon markets, we inherit the empirical observation that administrator credibility is a measurable input to basis behavior in this category.

These are not loose analogies. Each contributes a specific, mathematically tractable element to the framework that follows. We turn now to the construction of that framework.

4. A Fair Value Framework for FSPI

We are now in a position to develop the explicit form of fair value for sports performance index futures. We proceed in three stages. First, we state the core fair value relationship and define notation. Second, we decompose the expected accumulation term into three additive components corresponding to the three types of points contribution defined by the FSPI methodology. Third, we develop each component in turn, identifying the appropriate modelling approach for each.

4.1 The core fair value equation

Consistent with the framework used for VIX futures and other calculated benchmark futures, fair value for an FSPI futures contract is the current observable index level plus the expectation of the change in index level between the current date and contract expiry, taken under the relevant pricing measure:

$$F(t, T) = I(t) + \mathbb{E}^{\mathbb{Q}}[\Delta I(t, T)] \tag{2}$$

where $F(t, T)$ is the futures fair value at time t for expiry at time T, $I(t)$ is the observable index level at time t, and $\mathbb{E}$ denotes expectation under a risk-neutral pricing measure denoted $\mathbb{Q}$. The change in index level $\Delta I(t, T) = I(T) - I(t)$ is the accumulation of index points over the contract life.

The structural form of equation (2) is identical to that used for VIX futures, weather derivatives, and other calculated benchmark contracts. There is no cost-of-carry term, no convenience yield, and no foregone dividend. The fair value relationship is anchored in the dynamics of expectation rather than in a financing arbitrage.

A note on the pricing measure is in order. Under the physical measure $\mathbb{P}$, the expectation reflects the modeller's best estimate of where the index will accumulate to under the empirical distribution of outcomes. Under the risk-neutral measure $\mathbb{Q}$, the expectation reflects the price at which a representative market participant is willing to take the other side of the contract. The two measures differ by an amount that captures the market's aversion to bearing the relevant risk, which is the variance risk premium discussed in Section 6. For most of the development that follows, we use $\mathbb{E}$ without superscript when the result holds under either measure; we use $\mathbb{E}\mathbb{Q}$ explicitly when the distinction matters.[c]

4.2 Notation

The remainder of this paper uses the following notation:

$F(t, T)$	Futures fair value at time t for expiry at time T
$I(t)$	Observable index level at time t
$\Delta\, I(t, T)$	Change in index level from time t to time T
$\mathbb{E}$	Expectation operator (under measure $\mathbb{P}$ or $\mathbb{Q}$ as specified)
$\mathbb{P}$	Physical (empirical) probability measure
$\mathbb{Q}$	Risk-neutral pricing measure
G	Set of game-level statistical categories (e.g., goals, saves, hits, blocks, takeaways)
K	Set of end-of-game milestone categories (e.g., win, shutout, overtime)
L	Set of period and competitive milestone categories (e.g., division leader, playoff berth)
SAF_g	Statistical attribution factor for game-statistic category g, from the FSPI methodology
S_g	Realized count of category-g events in a game
m_k	Index points awarded for end-of-game milestone k, from the methodology
M_k,g	Indicator that game g triggers end-of-game milestone k
p_ℓ	Index points awarded for period or competitive milestone ℓ, from the methodology
M_ℓ	Indicator that period or competitive milestone ℓ is triggered before time T
B	Annual base level to which all FSPI values are reset (currently 7,500)
τ	Time of the annual reset event (one week after the end of the postseason)

4.3 Decomposing the expected accumulation

The FSPI methodology assigns index point contributions through three distinct mechanisms. In-game statistical events generate points equal to the per-constituent multiplier from the Points Attribution Table, applied to the realized count of each category and scaled by the seasonal adjustment factor SAF_game. End-of-game milestones, such as a win or a shutout, generate fixed point values when triggered in a given game. Period and competitive milestones, such as First in Conference or Season Most Goals, generate fixed point values realized only at the end of a period or upon the achievement of a defined competitive event.

Because these three mechanisms are separately specified and additively incorporated in the methodology, subject to defined override rules that enforce mutually exclusive outcomes (Overtime Loss overrides regulation Loss; Goals Scored = 1 overrides Opponent Shutout; Goals Allowed = 1 overrides

Shutout (Team)), the expected accumulation decomposes additively into three components, one for each mechanism:[b]

$$\mathbb{E}[\Delta I(t, T)] = \mathbb{E}[\Delta I_{\mathrm{stat}}] + \mathbb{E}[\Delta I_{\text{game-mil}}] + \mathbb{E}[\Delta I_{\text{period-mil}}] \tag{3}$$

Each of the three terms requires a different modelling approach. We develop them in turn.

4.4 In-game statistical accumulation

The first component, expected statistical accumulation, is structurally analogous to the settlement of a weather derivative. The realized contribution from any future game is the sum across statistical categories of the per constituent factor times the realized event count. Aggregated across all remaining games on the team's schedule between time t and contract expiry T, the expected statistical accumulation is:

$$\mathbb{E}[\Delta I_{\mathrm{stat}}] = \textstyle\sum_{\mathrm{g}\in \mathrm{G}}\sum_{\mathrm{game}\in[\mathrm{t,T}]}(SAF_g \cdot \mathbb{E}[S_{g,game}]) \tag{4}$$

where the inner sum runs over all games scheduled to be played between time *t* and contract expiry *T*, and the outer sum runs over all in-game statistical categories defined by the methodology. The schedule of games is a known input. The expected event count $\mathbb{E}[\, S_{\mathrm{g,\ game}}\,]$ is the modelled object: it is the expectation of the per-game count of category-g events for the team in question.

Several modelling choices are available for the per-game expected event count. A historical resampling approach uses the empirical distribution of recent per-game event counts for the team. A parametric approach fits a Poisson, negative binomial, or related distribution and projects forward. A regression approach incorporates additional covariates such as opponent strength, home and away venue, days of rest, and roster availability. The choice of modelling approach is an implementation question for market participants; the framework itself is agnostic to the choice.

The seasonal adjustment factor SAF deserves explicit comment because of how it enters equation (4). Under the FSPI methodology, SAF takes the value 1.0 for regular-season games and 1.5 for postseason games. The factor does not transition smoothly as the season progresses; it activates discontinuously the moment the team qualifies for the postseason. From a derivatives pricing perspective, this is structurally equivalent to a *knock-in feature on a structured note*: the elevated multiplier knocks in upon the qualification event, and the contract from that point forward accumulates points at the elevated rate. For a within-season contract whose expiry falls after the postseason begins, the expected statistical accumulation must therefore be computed as a probability-weighted blend of the regular-season path (multiplier 1.0 throughout) and the postseason path (multiplier 1.0 through qualification, then 1.5 thereafter), with the weighting determined by the team's playoff probability and conditional expected playoff length. The same observation applies to the milestone components developed in the next two subsections.

It is worth noting that this component is the largest and most predictable contributor to expected accumulation. In-game statistical events are recurring across the schedule, and even substantial uncertainty about per-game event counts averages out across a forward window of many games. For a typical mid-season FSPI futures contract with weeks or months of remaining games, the expected statistical accumulation can be estimated with relatively narrow confidence intervals.

4.5 End-of-game milestones

The second component, expected accumulation from end-of-game milestones, has a different modelling structure. End-of-game milestones such as wins, shutouts, and game winning goals are binary or near-binary events that occur at the conclusion of each game. The contribution from each milestone is the methodology's per-milestone point value times the probability that the milestone is triggered in each remaining game, summed over milestones and games:

$$\mathbb{E}[\Delta I_{\text{game-mil}}] = \sum_{\text{k}\in\text{K}} \sum_{\text{game}\in[\text{t,T}]} m_k \cdot \mathbb{P}(M_{k,game} = 1) \quad (5)$$

where *mk* is the per-milestone point value, *Mk, game* is the indicator that milestone *k* is triggered in the specified game, and ℙ(*Mk, game* = 1) is the probability of that milestone being triggered, calculated under the relevant pricing measure.

This is the structural form of an insurance contract or a binary digital option. The expectation of a binary event is its probability times its payoff. For wins, the relevant probability is the team's win probability in the upcoming game, which can be estimated from any of several established approaches: market-implied probabilities from established sportsbook lines, model-based probabilities from team strength estimators, or hybrid approaches. For shutouts and overtime outcomes, the methodology is similar: the probability of the event is estimated, and the per-milestone point value is applied.

The actuarial and digital options precedent is more than analogical. The mathematical machinery for valuing a portfolio of binary events with independent or weakly correlated payoffs is well developed in both fields. Insurance pricing routinely values books of binary policies; equity derivatives markets routinely value books of binary digital options. The FSPI end-of-game milestone component is a direct application of the same machinery.

It is important to be precise about what this component is and is not. It is the expectation, calculated under a specified pricing measure, of the milestone-based contribution to the index. It is a derivatives pricing input, not a wagering activity. The contracts that price these expectations are exchange-listed, centrally cleared, multi-day futures on a regulated benchmark; they are not single-event binary contracts of the kind that have attracted regulatory scrutiny in adjacent markets.

4.6 Period and competitive milestones

The third component, expected accumulation from period and competitive milestones, is structurally the most complex. Milestones such as Monthly Most Goals for a hockey team or winning the Stanley Cup are not realized game-by-game. They are realized only at the end of a defined period or upon the achievement of a defined competitive event. Their realization depends on the joint outcome of the team's remaining schedule and the schedules of other teams in the relevant competitive group.

The expected contribution from period and competitive milestones is:

$$\mathbb{E}[\Delta I_{\text{period-mil}}] = \sum_{\ell \in \mathrm{L}} p_\ell \cdot \mathbb{P}(M_\ell = 1) \tag{6}$$

where p_ℓ is the per-milestone point value and $\mathbb{P}(M_\ell = 1)$ is the probability that period or competitive milestone ℓ is triggered before contract expiry. The structural form is the same as equation (5), but the probability calculation is materially more complex.

The probability of winning a division, for example, depends not only on the team's own remaining schedule but on the schedules and outcomes of every other team in the division. The mathematics of this is naturally addressed through joint Monte Carlo simulation across all teams in the league. A typical implementation simulates the full remaining schedule of all 32 NHL teams across many trials, applies the standings rules to each simulated outcome, and counts the proportion of trials in which each team triggers each competitive milestone. The proportion is the probability estimate that enters equation (6).

This component is also the most sensitive to information shocks during the season. A single significant injury to a key player, a trade, or a sustained run of unexpected results can materially shift the probability of winning a division or qualifying for the playoffs, and therefore can materially shift this component of fair value. The other two components are also affected by such shocks, but the effect is typically smaller because in-game statistical accumulation and end-of-game milestone probabilities are smoothed across many remaining events. Period and competitive milestones, by contrast, are concentrated single events that can swing materially on a single piece of new information.

4.7 Putting the framework together

Combining equations (2), (3), (4), (5), and (6), the complete fair value expression for an FSPI futures contract on a single team, with expiry T occurring within a single season, is:[a]

$$F(t, T) = I(t) + \sum_{\text{g,game}} (SAF_g \cdot \mathbb{E}[S_{g,game}]) \tag{7a}$$

$$+ \sum_{\text{k,game}} m_k \cdot \mathbb{P}(M_{k,game}) \tag{7b}$$

$$+ \sum_\ell p_\ell \cdot \mathbb{P}(M_\ell) \tag{7c}$$

Each component is independently modellable and the components are additive. Each component admits a clean economic interpretation. The first is a weather-style accumulation across the known schedule. The second is an insurance-style book of binary expectations across upcoming games. The third is a Monte Carlo problem in joint outcomes across the league. The total is the fair value of the contract.

Equation (7) describes the framework as it applies to within-season contracts. Section 5 addresses the modifications required for cross-season contracts and the pricing implications of two structural features of the FSPI methodology that have not yet entered the framework explicitly: the annual reset and the absence of a within-season floor.

4.8 The instrument as a whole: a convertible-style analogue

The per-component analogues developed above are useful for instantiating each piece of the framework, but the instrument also admits a clean whole-instrument analogue worth surfacing for derivatives audiences. The structure of equation (7) — a steady, predictable accumulation component plus a portfolio of contingent payoffs whose values depend on threshold and competitive events — is structurally analogous to a convertible instrument. The statistical accumulation component is the yield-bearing piece: a steady stream of expected per-game points that accrues across the schedule, predictable in expectation, modest in single-period variance. The end-of-game and period and competitive milestone components together form an embedded portfolio of contingent payoffs: binary digital options on individual game outcomes, longer-dated path-dependent options on competitive milestones.

Within this whole-instrument framing, the SAF playoff multiplier discussed in §4.4 takes its appropriate place as a knock-in feature on the yield component. Both observations are useful. The convertible analogy gives the audience a single recognizable structured-product reference for the entire instrument; the knock-in observation captures the specific mechanic by which the yield component's accrual rate steps up upon a defined qualifying event. Neither is a substitute for the per-component framework that does the actual pricing work, but both are economically faithful to what the instrument actually is, and either can be drawn on as appropriate to the audience.

5. Structural Features and Their Pricing Implications

Three structural features of the FSPI methodology have direct implications for futures pricing that go beyond the within-season framework developed in Section 4. The first is the annual reset of all index values to a defined base level shortly after the end of the postseason, which interacts with the listed-contract calendar in a specific way. The second is the absence of a floor mechanism within a season, which permits index values to fall to zero or below in extreme circumstances. The third is the structural distinction between FSPI futures and single-event contracts, which is relevant to how the contracts are positioned with regulators, end users, and counsel.

5.1 Annual resetting and the within-season contract calendar

Under the current FSPI methodology, all team index values are reset to a defined base level B (currently 7,500) approximately one week after the conclusion of each season's postseason. We denote the time of the reset event by τ. The index level at time τ is therefore exogenously fixed:

$$I(\tau) = B \tag{8}$$

Listed contracts include expiries that fall both within a single season and across the annual reset. Within-season contracts are priced directly by equation (7). Cross-season contracts require explicit treatment of the reset, developed below.

5.1.1 Cross-season pricing across the reset

Listed contracts that span the reset event need an explicit pricing treatment, because their expected accumulation begins from the post-reset base level B rather than from the current observable index level. This subsection sets out that treatment.

Consider a futures contract with expiry T > τ. The change in index level over the life of such a contract decomposes into two pieces: the change from time t to time τ, which is determined by the team's remaining current-season performance and ends with the reset; and the change from τ to T, which begins from the base level B and accumulates across the new season. The first piece is immediately replaced by the reset; it does not survive into the contract life beyond τ. The second piece begins from B regardless of where the index sat just before the reset. The fair value of such a contract would therefore be:

$$F(t, T) = B + \mathbb{E}[\Delta I(\tau, T)] \quad \text{for } T > \tau \tag{9}$$

The current-season index level I(t) does not appear in equation (9). The team's current-season trajectory does not enter the fair value of a cross-season contract except indirectly, through whatever information about underlying team strength can be extracted from the current season and used in modelling next-season performance. Two teams whose current-season indexes were at very different levels — one at 12,000 and one at 4,000, for example — would have cross-season contracts whose fair values are determined entirely by expectations about next season, not by the current spread between their index levels. The reset event functions as an exogenous shock that wipes out the carry-over of current-season levels into next-season pricing.

Cross-season contracts are an analytically distinct product from within-season contracts and warrant their own market-making treatment, their own participant communication, and their own consideration of any risk premia attaching to next-season expectations. The framework above is sufficient to price them; the operational complexity sits in modelling next-season distributions at a horizon where the inputs are substantially less observable than within-season inputs.

5.2 The absence of a within-season floor

The current FSPI methodology does not impose a floor on within-season index values. Teams that perform poorly continue to lose points through the methodology's net-of-deductions calculation, and the index level can in principle fall to zero or below. While such an outcome is unlikely under realistic distributions of NHL team performance, it is mathematically possible, and the framework must address how futures should be priced when this possibility is non-negligible.

Two features of equation (7) make the no-floor design straightforward to handle in pricing. First, the decomposition is additive across components, and each component can independently take any value, positive or negative. The expected statistical accumulation can be negative if the methodology's deductions for unfavorable events (such as goals against) are projected to outweigh the contributions from favorable events. The expected end-of-game milestone contribution can be negative if penalty milestones (such as most giveaways in a given month) outweigh reward milestones. The expected period and competitive milestone contribution can be negative if the team's expected end-of-season position triggers methodology-defined competitive penalties.

Second, the fair value relationship in equation (2) inherits the sign properties of the expectation. There is no mathematical asymmetry, no nonlinearity around zero, and no breakdown of the framework when fair value is calculated to be negative. A futures contract on a deeply underperforming team late in a season may rationally be priced at a negative fair value if the expectation of the change in index from the current level to expiry is more negative than the current level itself. The framework handles this case in the same way it handles positive fair values.

Operationally, negative fair values are well precedented in regulated futures markets. The April 20, 2020 settlement of the May 2020 WTI crude oil futures contract at -$37.63 per barrel is the most prominent recent example.[5] Earlier in April 2020, CME advised the Commodity Futures Trading Commission that it was taking operational steps to support negative pricing in energy contracts, and the necessary clearing and risk management infrastructure exists at major derivatives exchanges.[5] The same infrastructure can support negative pricing in FSPI futures should it arise.

The argument for retaining the no-floor design rests on the cleanliness of the mathematical framework that results. A floor mechanism would introduce a nonlinearity at the floor level and would require modification of equation (7) to handle the censored distribution. Such modifications are tractable but introduce complexity that is not necessary if negative values are operationally supported. The methodology change to introduce a floor is reversible if market experience suggests it is needed; the framework can be extended to handle it. For now, however, the no-floor design produces a cleaner mathematical structure and aligns with the operational precedent in other regulated futures markets.

It is worth being explicit about a related point. Negative fair values do not imply negative payments by long position holders at settlement. The settlement mechanics are determined by the contract specifications, which can include features such as clearinghouse-set price limits, daily mark-to-market with margin calls, and circuit breakers, all of which are standard in regulated futures markets. The framework above addresses the fair value question; the operational handling of extreme values is a contract specification question that is appropriately addressed by the exchange.

5.3 Structural distinctions from event contracts

A natural question raised by the milestone components of equation (7) is the relationship between FSPI futures and the broader category of event contracts that has attracted regulatory attention in adjacent markets. The framework developed in this paper addresses pricing only and does not opine on regulatory classification, which is a matter for the relevant regulators, exchanges, and counsel. It is nonetheless useful to identify the structural features that distinguish FSPI futures from single-event binary contracts, because those features are direct inputs to the framework and are themselves the source of much of its tractability.

Five structural distinctions are worth noting. First, the underlying is a published, IOSCO-aligned benchmark with a defined methodology, formal governance, oversight committees, and documented procedures for methodology change, error correction, and complaint handling. The benchmark is administered independently of the trading of contracts that reference it. Second, settlement references the benchmark level on a defined calculation date rather than the outcome of any single underlying event. The benchmark is itself a continuously calculated weighted aggregation of many underlying events across a defined observation period; no single game, milestone, or competitive outcome causes the benchmark to settle, terminate, or otherwise become finally determined.

Third, the contracts are listed and centrally cleared on a regulated derivatives exchange, with daily mark-to-market, initial and variation margin requirements, position limits, and the full operational and risk management infrastructure of regulated futures markets. Fourth, contract horizons span days, weeks, or months across many contributing events; the payoff is therefore an aggregation across a known schedule rather than a function of a single binary outcome. Fifth, the framework prices the expected accumulation under the methodology's defined points attribution rules; the contracts do not pay a fixed amount on a binary win-or-lose basis, and the milestone-component probabilities that enter equations (5) and (6) are inputs to a continuous fair value calculation rather than the basis for a binary payout.

Each of these features is structural rather than rhetorical, and each can be verified directly from the benchmark methodology, the contract specifications, and the operating rules of the listing exchange. They are also directly relevant to the framework: the multi-event aggregation property is what makes the weather-derivative analogue applicable to the statistical accumulation component, the IOSCO-aligned governance is what supports the credibility-driven basis dynamics discussed in Section 6, and the central clearing and margin infrastructure is what makes negative fair values operationally supportable as discussed in Section 5.2. The pricing framework and the structural design of the contract are mutually reinforcing.

6. Basis Dynamics, Risk Premia, and Administrator Credibility

The framework developed in Sections 4 and 5 produces a fair value for an FSPI futures contract under a specified pricing measure. In practice, observed market prices will not equal the model fair value at every moment. The difference between observed price and modelled fair value, the basis, will fluctuate over the life of the contract in ways that reflect three forces: the relationship between the modeller's expectation and the market's expectation, the variance risk premium that compensates the marginal market maker for warehousing risk, and the market's confidence in the underlying benchmark itself. Each is treated in turn.

6.1 The variance risk premium

The pricing measure $\mathbb{Q}$ used in equation (2) is, in general, distinct from the physical (empirical) measure $\mathbb{P}$ under which a modeller would naturally project per-game event counts and milestone probabilities. The two measures differ by an amount that reflects the marginal market participant's aversion to bearing the relevant risk. In markets where this risk is asymmetrically held — typically by liquidity providers who must commit capital to support a market in which natural longs and natural shorts arrive at different times — the difference manifests as a systematic gap between the model fair value computed under $\mathbb{P}$ and the price at which the market clears.

This gap is the variance risk premium. It is well documented in the VIX futures market, where it appears as a persistent tendency for VIX futures to settle below their entry price more often than above, even after adjusting for the upward slope of the futures curve in calm regimes. The variance risk premium represents the compensation that volatility sellers earn on average for bearing the risk that volatility could spike. Analogous risk premia appear in commodity futures markets, in credit derivatives, and in most other futures markets where one side of the trade is structurally short an option or short a tail.

In FSPI futures, the variance risk premium is most likely to manifest in the period and competitive milestone component, where outcomes are concentrated single events with potentially large index point swings. Market makers who short futures on teams approaching playoff eligibility are effectively short an option on the team's playoff outcome; they will demand compensation for this exposure in the form of a premium relative to the model fair value computed under $\mathbb{P}$. The size and sign of this premium will depend on the structural balance of natural longs and natural shorts in the market, which is empirically unknown ex ante and will emerge as the market develops.

It is worth noting that the existence of a variance risk premium does not invalidate the framework. The framework prices the expectation; the variance risk premium is an additional component that adjusts the expectation to a market-clearing price. For most operational purposes — pricing a contract for an end user, comparing two contracts, calculating delta with respect to changes in expectations — the model fair value computed under $\mathbb{P}$ is the right input. For purposes of forecasting where contracts are likely to settle relative to where they are quoted, the variance risk premium is the relevant adjustment.

6.2 Liquidity and the cost of warehousing risk

In the early stages of any new derivatives market, the cost of warehousing risk is high. There are few participants, position limits are typically modest, and the marginal market maker must be compensated not only for the variance risk premium discussed above but also for the inventory risk of holding positions in an instrument with limited secondary liquidity. The result is a wider basis between model fair value and observed price than will typically prevail in a more developed market.

This is not a defect of the framework or of the contract design. It is a feature of derivatives market development. The same observation applies to VIX futures in their early years (volume took years to develop after the 2004 launch); to bitcoin futures in 2018 and 2019; and to weather derivatives, which remain a relatively small market by notional volume despite decades of operation. The basis tightens as participation broadens, position limits expand, and the marginal market maker's inventory risk falls.

From a framework-design perspective, the practical implication is that early FSPI futures pricing should not be expected to track model fair value tightly. Wider basis is consistent with normal early-market dynamics and does not indicate a problem with the framework. Over time, as participation grows, basis dynamics should be expected to converge toward the variance risk premium plus a small liquidity component, with the liquidity component diminishing as the market matures.

6.3 Administrator credibility as a basis input

The third force shaping basis is the market's confidence in the underlying benchmark itself. In a tradable underlying market, this question does not arise: the underlying is what it is, and arbitrage forces basis to converge regardless of what anyone thinks of any administrator. In a calculated benchmark market, by contrast, the question is fundamental. The benchmark is whatever the methodology and the administrator say it is. If the market doubts the methodology — if there is concern about future changes, about the integrity of the calculation, about the independence of the oversight, about the robustness of the input data — that doubt manifests directly in basis.

Carbon allowance markets provide the cleanest empirical evidence of this dynamic. Documented episodes in the European Union Emissions Trading System show that basis widens and curves become unstable in periods when administrator credibility is in question, and tightens and stabilizes in periods of confidence.[6] The mechanism is intuitive: market participants who are uncertain about the future trajectory of the methodology demand a wider spread to compensate for the uncertainty, and the spread tightens as confidence is established.

For FSPI, this has direct implications for benchmark administration. The investment that the administrator makes in IOSCO compliance, in transparent governance, in robust complaint-handling, in independent oversight, in published methodology change procedures, and in input data integrity is not just a regulatory checkbox. It is a direct input to basis behavior in the futures market that references the benchmark. Tighter governance translates to tighter basis, which translates to better execution for end users, which translates to deeper market participation.

This argument applies with particular force in the early stages of the market. End users who are evaluating whether to take exposure to FSPI futures will do so in part by assessing the credibility of the underlying benchmark. A robust governance posture, transparently maintained and visibly enforced, lowers the credibility risk premium that end users would otherwise build into their allocation decisions. The investment in benchmark integrity is therefore not a cost center; it is an input to commercial outcomes.

6.4 Putting basis dynamics together

The full decomposition of observed market price for an FSPI futures contract is:

$$F_{\text{market}}(t, T) = F_{\text{model}}(t, T) + VRP(t, T) + LRP(t, T) + CRP(t, T) \tag{10}$$

where F_{market} is the observed market price, F_{model} is the model fair value computed under $\mathbb{P}$ from equation (7), VRP is the variance risk premium discussed in 6.1, LRP is the liquidity risk premium discussed in 6.2, and CRP is the credibility risk premium discussed in 6.3. Each component is in general non-zero and time-varying, and each can be estimated and tracked separately as the market develops.

This decomposition is operationally useful. It separates the modellable component (model fair value) from the components that are properties of the market environment (the three risk premia). It allows market participants to identify the source of unexpected basis movements: a sudden widening that affects all teams uniformly is more likely a liquidity or credibility shock than a model-driven repricing; a widening concentrated on a single team is more likely a model-input change. It provides a framework for thinking about which components should narrow as the market matures (LRP and likely CRP) and which may be more persistent features of the market (VRP).

7. A Worked Illustrative Calculation

To make the framework concrete, this section applies it end-to-end to a generic representative team. We construct a hypothetical NHL team referred to as Team A, define a set of placeholder inputs that broadly resemble those of a competitive mid-table franchise at a representative point in a season, and compute the fair value of an FSPI futures contract on Team A under each of the three components of equation (7). We close with the resulting illustrative model fair value.

All inputs in this example are deliberately generic placeholders. They do not represent any actual team, any actual game schedule, or any actual probability assessment. The example is intended to demonstrate how the framework is applied; the numerical inputs are chosen for illustrative clarity rather than for realism.

7.1 Setup

The example uses the following setup. The current date is mid-season, with 35 games remaining on Team A's regular season schedule before the postseason begins. Team A's current observable index level is $I(t) = 8{,}400$. The futures contract under consideration has expiry T at the end of the regular season, which we treat for this illustration as occurring after all 35 remaining regular season games have been played but before the start of the postseason. The contract is therefore a within-season contract, and equation (7) applies directly without the cross-season modification of equation (9).

Parameter	Value
Current date	Mid-season (35 regular season games remaining)
Current index level $I(t)$	8,400 points
Contract expiry T	End of regular season
Contract type	Within-season ($T < \tau$)

Exhibit 1 — *Setup parameters for the worked illustrative calculation.*

7.2 Component 1: Expected statistical accumulation

Following equation (4), we estimate the expected statistical accumulation as the sum across statistical categories and remaining games of the per-constituent multiplier times the expected per-game event count. For pedagogical clarity, this illustration condenses the methodology's granular structure into five composite categories with illustrative per-constituent multipliers. The actual NHL FSPI methodology defines a more granular set of multipliers (Even Strength Goal +12, Power Play Goal +10, Short Handed Goal +15, Empty Net Goal −2 for goal types alone; Saves +0.75; Hit +0.1; and so on) that a production calculation would use in full. The composite values used here are chosen to keep the exposition tractable, not to represent the methodology's actual row-level multipliers.

Category g	$\mathbb{E}[S\_g]$ per game	SAF_g (points)	Contribution
Goals scored	3.1	+25	+77.5
Goals against	2.9	−25	−72.5
Saves	27.5	+1.5	+41.3

Category g	$\mathbb{E}$[S_g] per game	SAF_g (points)	Contribution
Hits	44.0	+0.4	+17.6
Minor penalties	8.5	−1.0	−8.5
Per-game total	—	—	**+55.4**

Exhibit 2 — *Per-game expected statistical contribution for Team A.*

The per-game expected statistical contribution is therefore approximately +55.4 index points. Multiplied across the 35 remaining games on Team A's schedule, the expected statistical accumulation component is:

$$\mathbb{E}[\Delta I_{\text{stat}}] = 55.4 \times 35 = 1{,}939 \tag{11}$$

Several observations are worth making. First, the per-game contribution is the net of positive and negative items, and either the methodology's deductions or particularly poor performance can produce negative per-game contributions. Second, the result is the expectation, not a forecast of the realized value. Realized accumulation will deviate from this expectation; the realized variance is itself a quantity that can be estimated and that is relevant for risk management but does not enter the fair value calculation directly. Third, the linearity of the calculation across games means that expected statistical accumulation can be computed quickly for any team and any horizon, given the per-game expected event counts.

7.3 Component 2: Expected end-of-game milestone accumulation

Following equation (5), we estimate expected end-of-game milestone accumulation as the sum across milestones and games of the per-milestone point value times the per-game probability of triggering the milestone. For this illustration, we use four common end-of-game milestones with placeholder probabilities.

Milestone k	$\mathbb{P}$(M_k) per game	m_k (points)	Contribution
Win	0.52	+10	+182
Game Winning Goal	0.52	+30	+546
Loss	0.32	−10	−112
Shutout (Team)	0.08	+50	+140
Total milestone contribution	—	—	**+756**

Exhibit 3 — *Expected end-of-game milestone contribution for Team A across 35 remaining games.*

The milestone-component calculation reflects Team A's projected win-rate distribution across the remaining schedule. The placeholder probabilities used here imply a 52 percent win rate (the Game Winning Goal milestone shares this probability, since each win records exactly one), a 32 percent regulation-loss rate, and an 8 percent Team Shutout rate, with the balance falling into overtime outcomes. The per-milestone point values shown are drawn from the published NHL FSPI Methodology Guide. Practitioners building their own calculation would substitute their preferred probability estimates while leaving the per-milestone values fixed by the methodology.

$$\mathbb{E}[\Delta I_{\text{game-mil}}] = 756 \tag{12}$$

7.4 Component 3: Expected period and competitive milestone accumulation

Following equation (6), we estimate expected period and competitive milestone accumulation as the sum across competitive milestones of the per-milestone point value times the probability of achieving the milestone before contract expiry. The probabilities here are typically estimated through joint Monte Carlo simulation across all teams in the league. For this illustration, we use placeholder probabilities that reflect Team A as a competitive mid-table team. For simplicity, the table below assumes a single monthly period remains within the contract life; in a multi-period contract the monthly milestone probabilities would be aggregated accordingly, with a correspondingly larger contribution to expected accumulation.

Milestone ℓ	ℙ(M_ℓ)	p_ℓ (points)	Contribution
Monthly Most Goals	0.30	+125	+38
Monthly Most Blocks	0.25	+125	+31
First in conference	0.05	+250	+13
Last in conference	0.03	−250	−8
Total competitive contribution	—	—	**+74**

Exhibit 4 — *Expected period and competitive milestone contribution for Team A.*

$$\mathbb{E}[\Delta I_{\text{period-mil}}] = 74 \tag{13}$$

This component is the most sensitive to information shocks during the season. A significant injury to Team A's top goaltender in week 22, for example, could materially shift each of the four probabilities shown above, with cascading effects on the First in Conference milestone that involve joint outcomes with other teams in the division. A practical implementation would update these probabilities frequently, ideally as part of an automated end-of-day model run.

7.5 The full fair value calculation

Combining the three components in equation (7), the model fair value of the within-season FSPI futures contract on Team A is:

Component	Value (index points)
Current index level $I(t)$	**8,400**
Expected statistical accumulation $\mathbb{E}[\Delta I\_stat]$	+1,939
Expected end-of-game milestone $\mathbb{E}[\Delta I\_game\text{-}mil]$	+756
Expected period and competitive milestone $\mathbb{E}[\Delta I\_period\text{-}mil]$	+74
Model fair value $F(t, T)$	**11,169**

Exhibit 5 — *Full model fair value calculation for the Team A within-season FSPI futures contract.*

The illustrative model fair value is approximately 11,169 index points, against a current index level of 8,400. The contract prices a positive expected accumulation of about 2,769 index points over the 35 remaining games of the regular season. Statistical accumulation dominates: about 1,939 points, or 70 percent of expected accumulation. End-of-game milestones contribute roughly 756, or 27 percent. Period and competitive milestones contribute only 74, or 3 percent: small in expected value but operationally important, as discussed below.

This is the model fair value, computed under the empirical measure $\mathbb{P}$. The price at which the contract would clear in a market would also reflect the variance risk premium, the liquidity risk premium, and the credibility risk premium discussed in Section 6. In a developed market with mature liquidity and high benchmark credibility, observed market prices should track model fair value closely. In an early market with thin liquidity and an unproven benchmark, observed prices could differ from model fair value by several percent. Both outcomes are consistent with the framework.

The relative magnitudes of the three components are also informative for risk management. The statistical and end-of-game milestone components together contribute approximately 97 percent of the expected accumulation in this example, but they are smoothed across many remaining games and individually move slowly in response to single information events. The period and competitive milestone component contributes only about 3 percent of fair value at this horizon, but it can have disproportionate contribution to repricing risk when the contract is near a standings, monthly, or seasonal threshold: a single significant injury, trade, or run of unexpected results can shift playoff and division-winning probabilities discontinuously, with directly proportional effect on this component. A market maker calculating exposure to news risk on this position should therefore weight the period and competitive milestone component more heavily than its proportional contribution to fair value would suggest. The same observation applies to the design of any hedge.

It bears emphasizing once more that all numerical inputs in this section are illustrative placeholders. The actual FSPI methodology specifies its own per constituent factors and per-milestone point values; the actual modelling of per-game expected event counts and probabilities is an implementation question that market participants would address using their preferred methods. The contribution of this section is to demonstrate that the framework, once instantiated with specific inputs, produces a single illustrative fair value through a transparent, decomposable calculation. Each input is observable or modellable from publicly available data; each step in the calculation is reproducible. This is the operational property that the framework is designed to provide.

8. Implications and Open Questions

The framework developed in this paper provides a complete approach to fair value pricing for FSPI futures contracts. Several questions remain open, however, and are worth identifying explicitly. These questions are not deficiencies of the framework; they are aspects of the framework whose resolution will depend on choices made by market participants and on empirical data that will only become available as the market develops.

8.1 The choice of modelling approach for per-game expectations

The framework is agnostic to the specific approach used to estimate the per-game expected event counts and milestone probabilities that enter equations (4), (5), and (6). Three broad approaches are available. A historical resampling approach uses the empirical distribution of recent per-game outcomes for the team. A parametric approach fits a distribution to the historical data and projects forward. A regression or hierarchical model approach incorporates additional covariates such as opponent strength, home and away venue, days of rest, and roster availability.

Each approach has advantages and limitations. Historical resampling is transparent and free of distributional assumptions but is slow to update for structural changes in team composition. Parametric approaches are quick to update and easy to communicate but assume a specific distributional form that may not hold across all categories. Regression and hierarchical models can capture more of the relevant structure but are more complex to maintain and more sensitive to specification choices. The market will likely settle on a small number of standard approaches over time, and the relative merits of each will be a productive area for further research.

Equally important to the modelling choice is the choice of inputs. Four categories of input are relevant in practice. First, league-licensed statistical feeds, which provide the historical event counts and milestone realizations that ground both historical resampling and parametric fitting. Second, market-implied probabilities derived from established sportsbook lines, which provide a continuously updated, market-aggregated estimate of per-game win probabilities and other binary milestone probabilities and which can serve as a calibration anchor or as a direct input. Third, milestone-specific prediction-market prices, where listed prediction markets quote contracts on the same competitive milestones that drive the period and competitive milestone component of equation (6); where such markets exist and are sufficiently liquid, their prices are the most direct available input to that component, and the framework can consume them as observable probability inputs without intermediate modelling. Fourth, internal models maintained by individual market participants, which can incorporate proprietary signals such as roster availability, line combinations, and goaltender starts that are not always reflected in published lines or in prediction-market prices. Each category has its own latency, cost, and reliability profile, and each carries its own implications for the governance of the inputs to the fair value calculation. Market participants should treat the input pipeline with the same rigor as the modelling pipeline; the framework's outputs are only as defensible as the inputs that feed it.

A natural product expression of the framework is a fair value calculator that accepts the inputs above and returns a model fair value for any specified contract. We anticipate that such a calculator will be made available, with three user-selectable approaches to estimating the per-game expected event counts that enter equation (4): a current-season approach using realized per-game averages from the season to date; a schedule-weighted approach using historical opponent yields adjusted for the team's remaining schedule; and a multi-year historical approach using rolling averages over a configurable lookback window. Each approach has its own strengths and the appropriate choice will depend on the user's purpose and on the maturity of the current season's data. Making the choice user-selectable rather than baked into the calculator preserves transparency about which assumptions are driving any particular fair value output.

8.2 Cross-season contracts and the longer-horizon problem

The framework treats cross-season contracts through equation (9), which expresses cross-season fair value as the post-reset base level B plus the expected accumulation across the relevant portion of the new season. This is mathematically clean but requires next-season expected accumulation to be modelled at a horizon where the inputs are substantially less observable than within-season inputs. Roster turnover, coaching changes, and the underlying year-over-year persistence of team performance are all material drivers that need to be addressed.

A natural starting point is to model next-season per-game event counts and probabilities as a regression of next-season outcomes on current-season inputs, fitted to historical league data. This is a tractable problem but a different one from the within-season problem. Cross-season contracts may warrant their own modelling pipeline, separate from but consistent with the within-season pipeline.

8.3 Information shocks and intra-season repricing

The framework prices the contract under information available at time t. As new information arrives — injury news, trade announcements, lineup changes, results — each component of equation (7) is updated and the model fair value reprices. The expected statistical accumulation component is relatively stable to single-game shocks because it averages across many remaining games. The end-of-game milestone component is moderately sensitive because each remaining game's win probability may shift. The period and competitive milestone component can be highly sensitive because injury or trade news can shift the entire trajectory of playoff or division probability.

This creates an asymmetry in the way the three components contribute to realized basis volatility. A practical implication is that a market maker calculating exposure to news risk on an FSPI position should probably weight the period and competitive milestone component more heavily than its proportional contribution to fair value would suggest. This is a familiar pattern from other markets with concentrated event risk.

8.4 Multi-league extension

The framework was developed using NHL FSPI as the working example, but it generalizes naturally to other leagues. Each component of equation (7) admits a league-specific instantiation. The set of in-game statistical categories G, the set of end-of-game milestones K, and the set of period and competitive milestones L are all defined by the league-specific FSPI methodology. The per constituent factors and per-milestone point values are league-specific inputs to the calculation. The choice of modelling approach for per-game expectations is league-specific in its instantiation but identical in its framework.

As FSPI expands to additional leagues such as Major League Baseball or NASCAR, the framework should require only the substitution of league-specific inputs. The structural form of the fair value relationship, the additive decomposition into three components, and the treatment of structural features such as the annual reset and the no-floor design should carry over. Whether the relative magnitudes of the three components are similar across leagues, or whether (for example) statistical accumulation dominates more heavily in MLB than in NHL given the substantially larger number of games per season, is an empirical question whose answer will emerge as the family of indexes is extended.

8.5 Settlement mechanics, hedgeability, and data risk

The pricing framework above takes the methodology's defined settlement convention as given. Several operational features of that convention warrant explicit attention because they bear on final settlement risk, last-trading-day design, and intra-day basis behavior. The settlement rules summarized below are drawn from the FS Index Calculation and Publication Schedule, which governs how official statistics, milestones, and amendments are translated into Index values. The official Index Close for each Calculation Day is published at 9:00 AM Central Time on the following day (T+1), reflecting final verified league data from the prior game day. Monthly and seasonal milestones are applied on a T+1 basis at defined times following the conclusion of the relevant observation period. Games that begin on one date but conclude on the following date are attributed to the date on which the game began. Subsequent official league amendments to game statistics are treated by the Calculation and Publication Schedule as discrete statistical events, incorporated into Index calculations in accordance with the Points Attribution Table; they are reflected in the Daily Official EOD Index Value on the calculation day following publication and do not retroactively affect historical milestone determinations. For contracts expiring near a game day, last-trading-day specifications, the timing of final settlement against the T+1 published close, and the treatment of amendments issued after final settlement are not pricing questions but they affect realized settlement and should be addressed in contract specifications and operational documentation.

Hedgeability is the second operational question worth flagging. There is no arbitrage channel to a tradable underlying, but a market maker holding an FSPI position has several proxy instruments available: sportsbook lines on team outcomes (where legally accessible and operationally usable), prediction markets on milestones, futures or other instruments on correlated team indexes within the same league, and player or news-driven inputs that can be partially neutralized through other exposures. Exchange-traded funds referencing the FSPI futures, once listed, would provide an additional hedging mechanism, a development we view as plausible given ongoing engagement with prospective ETF sponsors interested in building products around the FSPI concept. The extent to which FSPI risk can be proxy-hedged versus warehoused will depend on the maturity of these adjacent markets and on the correlation structure between FSPI and the available proxies. We expect a meaningful portion of the risk to be warehoused in the early stages of the market, which is consistent with the liquidity risk premium discussed in §6.2. As adjacent markets develop and correlation structures become better understood, proxy hedging should become more effective and the liquidity risk premium should narrow.

Data risk deserves a distinct treatment within the credibility framework of §6.3. The methodology depends on official league data through a defined hierarchy of inputs, with secondary sources permitted only when official data is temporarily unavailable. Three sub-categories matter for futures pricing. Data finality risk reflects the possibility that statistics treated as official may be subsequently corrected. Data latency risk reflects the gap between game-end and official publication, which can leave intra-day index values in flux. Settlement-source risk reflects the possibility that a feed disruption requires reliance on alternate data sources whose accuracy or timing differs from the primary feed. Each of these can contribute to basis movements that are not driven by the model fair value, and clearing teams should treat them as distinct categories within the broader credibility risk premium rather than aggregating them into a single line item.

8.6 The path forward

The framework above is offered as a starting point for market practice. We expect it to evolve as the market develops, as empirical data on basis behavior accumulates, and as participants refine the modelling approaches that work best in practice. The principal contributions of this paper are: an explicit framing of fair value as an expectation rather than as a no-arbitrage relationship; a clean three-way decomposition that separates the three economically distinct contributions to expected accumulation; an explicit treatment of the structural features of the FSPI methodology that have direct pricing implications; and a worked example that demonstrates how the framework is applied in practice.

None of these contributions is novel in the broader benchmark futures literature. What is new, and what we hope is useful, is their synthesis and application to the specific structural form of sports performance index futures. As the market develops, we expect to refine and extend each component of the framework based on what works in practice. The publication of this framework is intended to provide a shared starting point for those refinements.

9. References, Endnotes, and Disclosures

References

1. International Organization of Securities Commissions, Principles for Financial Benchmarks: Final Report, July 2013.

2. Cboe Global Markets, VIX Index Methodology, 2003 (with subsequent revisions). VIX futures launched on the Chicago Futures Exchange on March 26, 2004; VIX options launched on Cboe Options Exchange on February 24, 2006.

3. CME Group, CME CF Bitcoin Reference Rate Methodology, 2016. CME bitcoin futures launched on December 18, 2017.

4. U.S. Securities and Exchange Commission, Order approving spot bitcoin exchange-traded products, January 10, 2024. Trading commenced January 11, 2024.

5. CME Group, communication to the U.S. Commodity Futures Trading Commission regarding operational support for negative pricing in energy futures, April 2020. The May 2020 WTI crude oil futures contract settled at −$37.63 per barrel on April 20, 2020.

6. European Union Emissions Trading System, periodic reports on allowance market dynamics. See European Commission documentation on Phase III and Phase IV reform processes for documented basis behavior across credibility cycles.

Endnotes

a. The framework as presented uses a single-team contract structure throughout. Multi-team contracts (such as basket contracts on division winners or playoff qualifiers) are a natural extension and would be priced through additive aggregation of the relevant single-team expectations, with appropriate treatment of joint outcome probabilities for milestones whose payoffs depend on more than one team.

b. The decomposition in equation (3) treats the three components as additive and independent. In practice, there may be modest correlation between components — for example, a team that overperforms on statistical accumulation in a given period may also have elevated milestone trigger probabilities in the same period. The additive form remains correct in expectation; the independence assumption affects the modelling of variance, not of the expectation itself.

c. The pricing measure $\mathbb{Q}$ in equation (2) is left unspecified beyond noting its existence and its difference from the empirical measure $\mathbb{P}$. In practice, the choice of pricing measure for FSPI futures will be determined empirically as the market develops, by observing the difference between model fair values calculated under $\mathbb{P}$ and observed market clearing prices. Theoretical work on the existence and uniqueness of an equivalent martingale measure for benchmark futures with non-investable underlyings is a natural extension of this paper.

Disclosures

This paper 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 advice, legal advice, tax advice, or accounting advice.

The information in this paper is based on data and methodology descriptions believed to be reliable, but no representation or warranty is made as to the accuracy, completeness, or timeliness of any information herein. The framework and analysis represent the views of the author as of the date noted, and these views may change without notice in response to changing circumstances, market conditions, or the development of the FSPI product family.

All numerical inputs and outputs in the worked example in Section 7 are illustrative placeholders chosen for exposition. They do not represent any actual team, any actual game schedule, any actual probability assessment, or any actual contract price. They are not intended to forecast or to imply any view about the value of any actual or contemplated FSPI futures contract. The example is included to demonstrate the application of the framework, not to provide a pricing recommendation.

The framework presented here may not be suitable for all market participants. Its application requires specialized knowledge of derivatives pricing, statistical modelling, and the underlying FSPI methodology. Market participants should consult their own advisors, perform their own due diligence, and reach their own independent judgments before acting on any analysis derived from the framework presented here.

Forward-looking statements contained in this paper, including statements about the future development of the FSPI futures market, basis dynamics, market participant behavior, and the extension of the framework to additional leagues, are based on assumptions that may not be realized. Actual outcomes may differ materially from those discussed.

FutureSports Performance Indexes (FSPI) are calculated benchmarks administered in accordance with the IOSCO Principles for Financial Benchmarks.[1] The framework presented 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 FSPI methodology should refer to the published FSPI methodology documentation.

Past performance is not indicative of future results. The framework presented here addresses fair value, which is an expectation under a specified pricing measure; it does not predict realized outcomes. Realized accumulation of any FSPI value will deviate from any expectation calculated using this or any other framework, and the magnitude of those deviations is itself a quantity subject to estimation uncertainty.

This paper has not been reviewed or approved by any exchange, regulatory authority, benchmark administrator, or other third party. The views expressed are those of the author and do not necessarily represent the views of any organization with which the author is or has been affiliated.