Seeing Like a Market: Event Contracts and Market Topology

Working Paper Summary – March 2026

1 Core Question

When a derivative market and a prediction market both price the same binary event, which venue offers lower total cost to the end user, and what determines when capital migrates from one to the other? This paper provides a structural decomposition of the cost gap between these two market architectures and tests the resulting predictions across 87 event-contracts spanning 11 categories, 5 asset classes, and elections in 12 countries, using 2.9 million trade-level observations from November 2024 through February 2026.

The central mechanism is a *structural cost wedge* that separates the two venues. In the options market, replicating a binary payoff requires purchasing volatility exposure, which embeds three distinct costs: a variance risk premium, a dealer balance sheet charge, and a replication friction. In a prediction market, the same binary claim trades directly at the probability. The gap between these two costs (the "Vega Wedge") determines which venue dominates. The paper shows that this wedge varies systematically across event categories, generates a testable ordering of which categories prediction markets should capture first, and identifies a threshold mechanism (the "Liquidity Seesaw") governing capital reallocation between venues.

2 The Vega Wedge: Structural Cost Decomposition

The cost of expressing a binary view through derivatives can be decomposed into three components:

$$W = \mathrm{VRP} + B + F \tag{1}$$

where W is the Vega Wedge, VRP is the variance risk premium (the spread between implied and realized volatility), B is the dealer balance sheet cost (intermediation capital charges under post-crisis regulation), and F is the replication friction (the practical cost of constructing a binary payoff from continuous instruments). The prediction market bypass is direct: a binary contract trades at its implied probability with none of these embedded costs.

The Vega Wedge measures the structural advantage of the prediction market architecture. It is not a temporary mispricing or an arbitrage opportunity. It is an equilibrium cost difference that arises because the two venues use fundamentally different production technologies to deliver the same payoff. Options markets produce binary information as a byproduct of continuous hedging infrastructure; prediction markets produce it natively.

This decomposition connects directly to the frictions framework in intermediary-based asset pricing. The VRP component reflects the compensation demanded by volatility sellers for bearing variance risk. The balance sheet component B captures the cost of dealer capital allocated to intermediation, the same channel through which leverage constraints and margin requirements affect derivative pricing. The replication friction F reflects the practical cost of constructing a digital payoff from a continuum of strikes, a cost that is zero by construction in a prediction market.

A key feature of the decomposition is that VRP is the only component that varies substantially across event categories. The balance sheet cost B and replication friction F are relatively stable across categories (the paper estimates $B + F \approx 0.07\%$ for Tier 1 and Tier 2 categories as a constant). This means the VRP gradient is, for practical purposes, the structural wedge gradient. Categories where options are most expensive relative to realized risk are exactly those where prediction markets hold the largest cost advantage.

3 The VRP Gradient: Empirical Ordering

The variance risk premium varies systematically across event categories, generating a testable prediction: categories with higher VRP should cross the Seesaw threshold first, because the structural cost advantage of prediction markets is largest where the options market charges the most for volatility exposure.

Table 1: VRP Gradient Across Event Categories

Category	Proxy	VRP (%)	Tier	Verdict
BTC (Annual 2025)	DVOL / Spot	+8.70	Direct	PM wins
BTC (Monthly avg.)	DVOL / Spot	+4.30	Direct	PM wins
Elections (median)	Country ETF IV	+3.50	Proxy	PM wins
Gold	GC options	+2.82	Direct	PM wins (partial)
Equity Indices	ES/NQ options	+1.90	Direct	Threshold
Commodities (CPI)	ZN options	+1.42	Proxy	Threshold
FOMC	SR3/ZQ decomp.	+0.52	Constructed	PM loses

The gradient spans an order of magnitude, from BTC at approximately 4.3% (monthly average) to FOMC at approximately 0.5%. The ordering is robust to seven alternative specifications of the FOMC VRP calculation (window width, return frequency, annualization method, repo noise adjustment), with the mean FOMC estimate ranging from 0.454% to 0.561% across all specifications. The qualitative ordering is invariant: BTC always occupies the top, FOMC always occupies the bottom, and the intermediate categories maintain their relative positions.

Table 2: FOMC VRP Sensitivity Analysis

Specification	Change	Mean VRP (%)	Gradient Position
Canonical	–	0.516	Low
Alt RV window: $T-2$ to $T+2$	Wider window	0.488	Low
Alt RV window: $T-0.5$ to $T+0.5$	Announcement day	0.561	Low
Alt annualization: $\sqrt{N_{\text{bars}}}$	Bar-count scaling	0.503	Low
Alt return freq: 15-minute	Coarser sampling	0.529	Low
No repo adjustment	Raw IV $-$ RV	0.454	Low
Extended hours RV	Include Globex overnight	0.542	Low

The range across specifications (0.454% to 0.561%) is narrow relative to the gap between FOMC and the next category on the gradient (equity indices at 1.90%). This stability is important: it means the FOMC's position at the bottom of the gradient is a structural feature of the category, not an artifact of any particular methodological choice.

The measurement architecture itself reflects the VRP hierarchy. BTC VRP is computed directly: Deribit's DVOL index (a model-free 30-day implied volatility, constructed via the Carr-Madan variance swap methodology) minus realized spot volatility. No proxy chain is needed because the derivative instrument (Deribit options, summarized by DVOL) and the prediction market contracts (Polymarket BTC touch markets) reference the same underlying asset. This is Tier 1 measurement.

FOMC VRP, by contrast, requires a constructed proxy because Fed Funds options (OZQ), the native instrument for FOMC binary hedging, emit no observable market data. The CME Interest Rate Liquidity Review omits the product. CFTC Bank Participation Reports exclude it. The physical trading pit closed in May 2021. Direct VRP measurement is impossible. The paper builds an SR3/ZQ decomposition: implied volatility is extracted from Three-Month SOFR options via Black-76 inversion at the ATM strike nearest to the $T-1$ settlement price, and realized volatility is computed from Fed Funds futures (ZQ) 5-minute midpoint returns over a three-day event window ($T-1$ market open to $T+1$ market close) centered on each FOMC announcement. A repo noise adjustment removes SOFR-specific financing dynamics (quarter-end effects, repo market stress, collateral scarcity) that affect SOFR independently of Fed decisions. The adjustment is modest in magnitude (mean of approximately 10 basis points) but critical for the December 2025 meeting, where year-end financing dynamics inflated the raw SR3-ZQ spread to 28.8 basis points of repo noise.

This measurement asymmetry is itself informative: categories where VRP is hardest to observe tend to have the lowest structural wedge. The difficulty of measurement and the smallness of the premium are two expressions of the same underlying fact, that the options market does not charge very much for FOMC event risk.

4 The Liquidity Seesaw: Threshold Mechanism

The Seesaw is a binary regime classification applied at the $3M reference position size. For each event-contract, the paper compares the structural wedge (VRP + B + F) against the effective execution cost in the prediction market venue. Effective spread is computed as the mean absolute deviation of executed trade prices from the contemporaneous mid-price, inferred from the most recent buy-side and sell-side trade prices. The classification:

- **PM wins:** Structural wedge exceeds PM execution cost by more than 0.5 percentage points. Capital should migrate to the prediction market.
- **Threshold:** Gap within ±0.5 percentage points. Venues are approximately competitive.
- **PM loses:** PM execution cost exceeds structural wedge. The options market remains the lower-cost venue.

Across 87 event-contracts at the canonical ±0.5pp boundary: 30 are classified as PM wins, 12 as threshold, 1 as PM loses (marginal), and 44 as PM loses (high friction). The classification is robust to boundary perturbation, as shown in Table 3.

Table 3: Classification Sensitivity to Boundary Choice

Boundary	PM Wins	Threshold	PM Loses (Marginal)	PM Loses (High Friction)
±0.3pp	27	6	4	50
±0.5pp (canonical)	30	12	1	44
±0.7pp	33	16	0	38
±1.0pp	36	18	0	33

The qualitative story is stable: BTC dominates PM-wins at all boundaries; FOMC threshold events remain at threshold under tighter or wider bands; the binding constraint taxonomy is unchanged at all tested boundaries.

The Seesaw mechanism is not a static comparative-statics exercise. PM execution costs are declining rapidly over the sample period: FOMC effective spreads fell 78% and BTC spreads fell 57%. This trajectory implies that categories currently at threshold are approaching the tipping point, and the ordering of which categories cross next follows the VRP gradient.

Table 4: Execution Cost Trajectory (Effective Spread at $3M)

Category	Early Sample	Late Sample	Reduction	Current Status
BTC	1.82%	0.78%	57%	PM wins
FOMC	2.14%	0.47%	78%	Threshold
Elections	1.95%	0.92%	53%	PM wins (where VRP high)

5 Market Segmentation and the Preferred-Habitat Channel

The options market and the prediction market are segmented venues serving overlapping but distinct clienteles. This is a preferred-habitat story: institutional hedgers and volatility traders inhabit the options market because of its infrastructure (margin, clearing, prime brokerage, regulatory capital treatment); retail and information-motivated traders inhabit prediction markets because of direct binary access and lower fixed costs of participation. The structural wedge quantifies the cost of this segmentation.

The Vega Wedge framework maps onto the six frictions in the intermediary asset pricing literature as follows. Leverage constraints and margin requirements contribute to B by raising the capital cost of dealer intermediation. Asymmetric information costs appear in both venues but manifest differently: in options markets through wider bid-ask spreads on complex structures; in prediction markets through adverse selection on the binary contract itself. Participation costs differ sharply: the fixed cost of accessing CME options (clearing membership, margin infrastructure, regulatory compliance) vastly exceeds the cost of placing a prediction market trade. Search frictions are lower in the prediction market architecture, where the payoff structure is transparent and the contract specification is self-contained.

Capital migration across the segmentation boundary follows a threshold process, not a smooth rebalancing. The Seesaw captures this: when the structural wedge exceeds the execution cost friction in the prediction market venue, capital begins to flow. The flow is one-directional along the VRP gradient (high-VRP categories first) and is accelerating as prediction market liquidity deepens (lower execution costs widen the set of categories where PMs dominate).

This segmentation also explains an empirical puzzle: why do FOMC prediction markets exist at all when the structural wedge is small? The answer is the *Observability Wedge*.

Fed Funds options (OZQ), the native instrument for FOMC binary hedging, exist as listed contracts but produce no observable market data. The CME Interest Rate Liquidity Review omits the product; CFTC Bank Participation Reports exclude it; the physical trading pit closed in May 2021. Kalshi FOMC contracts, by contrast, demonstrate 100% directional forecast accuracy across six meetings in the sample (confirmed by a February 2026 Federal Reserve staff paper). The prediction market does not win on cost; it wins because the options market fails to function as an observable price-discovery mechanism for this specific binary event. The Observability Wedge is a distinct channel from the Vega Wedge: it concerns not the *cost* of producing information but the *existence* of an observable signal. For FOMC, the constructed order has failed to emit a readable price, and the spontaneous order fills the vacuum.

6 Informational Equivalence: Smile Emergence

A central empirical finding is that prediction markets spontaneously generate the same distributional features (volatility smiles, fat tails, skew) that options markets produce through the dealer-intermediated machinery of dynamic hedging, risk management, and model-based pricing. This is tested in the BTC touch market on Polymarket, where contracts at discrete price thresholds pay $1 if BTC reaches the threshold by month-end and $0 otherwise. Each contract's touch probability is inverted to implied volatility using the geometric Brownian motion barrier-option formula:

$$P_{\text{touch}} = N\left(\frac{\ln(B/S)}{\sigma\sqrt{T}}\right) + \left(\frac{B}{S}\right)^{2/\sigma^2} N\left(\frac{-\ln(B/S)}{\sigma\sqrt{T}}\right) \tag{2}$$

where B is the barrier (strike threshold), S is the spot price, σ is implied volatility (solved numerically), and T is time to expiry. The pipeline filters contracts with touch probabilities outside the 5% to 95% range, where the inversion becomes ill-conditioned.

The resulting smile exhibits the canonical features of an options-derived volatility surface: curvature (higher implied volatility at out-of-the-money strikes), skew (asymmetric treatment of upside and downside risk), and regime responsiveness (smile shape shifts with market conditions). In November 2024, when BTC rallied approximately 40%, the smile shifted in exactly the direction predicted by the framework: upside touch probabilities increased (reducing "surprise" content), shifting skew toward the downside. These features emerge without dealers, without market makers running delta-hedging algorithms, and without any participant necessarily being aware that they are collectively constructing a volatility surface. The information aggregation occurs through the price mechanism alone.

Skew is measured as the difference between OTM call-equivalent IV (upside touch) and OTM put-equivalent IV (downside touch) at matched moneyness. Positive skew indicates call-dominant (bullish) positioning; negative skew indicates put-dominant (bearish) positioning. The skew dynamics track BTC market sentiment across the sample, responding to regime changes in a manner qualitatively consistent with the Deribit options surface.

This finding is significant for market microstructure theory. It demonstrates that the informational content of the volatility surface (the market's assessment of the probability distribution of future outcomes) is not an artifact of the specific institutional architecture that produces it. The surface is a property of the underlying information, not the intermediation technology. Prediction markets and options markets are informationally equivalent venues: they produce the same distributional signature through different institutional channels. The caveats are important: touch options are path-dependent while European options depend only on terminal price, so the PM-derived smile produces *qualitatively similar* distributional features, not numerically identical implied volatilities.

7 Empirical Scope and Data Architecture

The empirical analysis covers 87 event-contracts across 11 categories and 5 asset classes: interest rates (FOMC, ECB, BOJ), cryptocurrencies (BTC touch markets), equity indices (S&P 500, Nasdaq 100 barrier events), commodities (gold, silver, retail gas), macroeconomic releases (CPI, GDP), and elections in 12 countries. The combined sample contains 2,889,424 trade-level observations from January 2024 through February 2026.

Prediction market data is sourced from Polymarket (47 events, 1.6M rows via the Goldsky subgraph, all on-chain on the Polygon blockchain) and Kalshi (23 events, 49K rows via their API, CFTC-regulated with legally standardized resolution). A critical data processing step normalizes both YES and NO token trades to the YES-equivalent probability space before computing execution metrics. Without this normalization, preliminary analysis showed spurious apparent spreads of 20 to 27% that collapsed to 1 to 5% after correction.

Derivatives data comes from Databento (10 CME/CBOT symbols across market-by-price, settlement, and OHLCV schemas), Deribit (DVOL index for BTC), and OPRA (Options Price Reporting Authority data for election ETF proxies covering 10 country ETFs across the U.S., Germany, South Korea, Canada, Chile, Argentina, Poland, UK, Australia, Taiwan, and Mexico).

The VRP measurement operates at three tiers of directness. Tier 1 (direct measurement) covers BTC, where the derivative and the prediction market reference the same underlying.

Table 5: Data Architecture Summary

Source	Events	Trade Rows	Coverage
Polymarket (Goldsky)	47	1,625,309	FOMC, BTC, ECB, BOJ, Gold, Silver, Elections
Kalshi (API)	23	49,473	CPI, GDP, Equity Indices, Retail Gas
Databento (CME)	–	–	10 derivatives symbols (IV, RV, settlement)
Deribit	–	–	DVOL index (BTC implied vol)
OPRA	–	–	10 country equity ETF option chains

Tier 2 (direct with basis) covers equity indices, gold, and silver, where ATM IV from futures options is compared against realized volatility of the underlying futures. Tier 3 (proxy measurement) covers FOMC (SR3/ZQ decomposition), central banks (FX futures options), macro releases (Treasury options), and retail gas (gasoline futures options). The paper is transparent about the declining measurement precision across tiers: proxy VRP overstates event-specific risk because the proxy derivative absorbs non-event volatility. For ECB, FX volatility around ECB meetings is estimated to be only 25 to 40% attributable to the rate decision itself, placing the rate-specific VRP close to FOMC territory (0.3 to 0.6%).

8 Binding Constraint Taxonomy

For each event-contract, the paper identifies the binding constraint preventing (or permitting) prediction market dominance. This taxonomy reveals that execution cost friction is not the only, or even the primary, obstacle. The constraints fall into four categories:

Low structural wedge. For FOMC events, the VRP is approximately 0.5%, generating a structural wedge near 0.6%. Even with 78% execution cost reductions, PM execution costs must fall to approximately 10 basis points (roughly institutional FX spreads) before PMs definitively dominate. The wedge, not the friction, is the binding constraint.

Execution cost friction. For gold events, VRP ranges from 2.8% (near-ATM) to over 20% (deep OTM). The structural wedge is large, but PM execution costs at the $3M reference size remain above the wedge for some contracts. The deep-OTM gold contracts ($5,500 target with spot near $2,600) use ATM IV, which understates the true structural wedge; wing VRP at the relevant moneyness would be 1.5 to 2.5 times the ATM figure. Gradient positions for these events are lower bounds.

Volume capacity. For some event categories (macro releases, retail gas), total PM volume is small relative to the reference position size. The depth curve must be extrapolated beyond observed data, introducing estimation uncertainty. Retail gas events carry additional quality warnings: the pipeline flagged stale or one-sided option quotes on every observation.

RB (RBOB gasoline) option liquidity at the relevant tenors is thin, and the July 2025 reading of 31.2% is almost certainly inflated by stale quotes rather than genuine risk premium.

Negative VRP (silver control case). Silver options show a VRP of -10.19% (IV of 42.9% vs. RV of 53.1%). When options are cheap relative to realized risk, the structural wedge is negative and prediction markets cannot win regardless of execution quality. This serves as a structural control: the framework correctly predicts PM disadvantage where the cost differential favors options.

9 Institutional Context and Regulatory Trajectory

The paper is situated at a specific institutional moment. On March 12, 2026, the CFTC issued a Staff Advisory Letter effectively classifying prediction market contracts as a financial asset class. Nasdaq has filed to list "Outcome Related Options," which would bring binary event contracts onto a traditional exchange infrastructure. A February 2026 Federal Reserve staff paper documented Kalshi FOMC contracts' 100% directional forecast accuracy. Kalshi reached $1.9 billion in weekly notional volume. These developments validate the paper's core premise: prediction markets are not a curiosity but a competing market architecture that is achieving institutional scale.

The convergence between the two architectures is proceeding from both directions. Prediction markets are gaining regulatory legitimacy and institutional infrastructure. Options exchanges are exploring binary payoff structures. The Vega Wedge framework provides a principled basis for predicting where this convergence will occur first (high-VRP categories) and where options markets will retain their structural advantage (low-VRP categories where the cost of the intermediation infrastructure is small relative to the precision it provides).

The election coverage illustrates the global dimension of this convergence. The paper analyzes prediction market contracts for elections in 12 countries, using country equity ETF options as the VRP proxy. The proxy chain (election outcome $\rightarrow$ country equity market reaction $\rightarrow$ ETF $\rightarrow$ ETF options chain $\rightarrow$ ATM IV $\rightarrow$ VRP) is the longest in the paper, with the heaviest contamination from non-election dynamics. To isolate the election-specific signal, the paper computes an election event premium ($EP_e = IV_e^+ - IV_e^-$) and requires that the election account for at least 10% of pre-election IV for inclusion. Eight of seventeen event-periods pass this filter. A standardized VRP metric (SVEP) normalizes by baseline IV to allow cross-country comparison despite wide variation in underlying equity volatility (Poland's EPOL at 102.5% baseline IV versus South Korea's EWY at 35.7%).

10 Theoretical Architecture: Spontaneous vs. Constructed Order

The paper employs Hayek's distinction between spontaneous order (*cosmos*) and constructed order (*taxis*) as the organizing framework. Options markets are a constructed order: they require dealers, clearinghouses, margin systems, and model-based pricing to function. Prediction markets are a spontaneous order: binary prices emerge from decentralized trading with no intermediation infrastructure beyond the matching engine and settlement mechanism.

The contribution of this framework is not philosophical but empirical. The smile emergence finding (Section 6) demonstrates that the informational output of the two architectures is equivalent despite their radically different institutional structures. Both produce volatility smiles; both encode fat tails and skew; both respond to regime changes. The difference lies entirely in the cost of production: the constructed order embeds the Vega Wedge, while the spontaneous order does not. The VRP gradient then determines which cost structure dominates for each category of binary event.

The Hayekian lens also illuminates the Observability Wedge. Fed Funds options are a constructed-order product that has failed to achieve functional price discovery: the institutional machinery exists but produces no observable signal. Kalshi FOMC contracts, a spontaneous-order alternative, fill the void not by being cheaper but by being observable. This is a market failure in the traditional sense: the existing institutional infrastructure does not deliver its intended informational function, and an alternative architecture has emerged to serve the unmet demand.

The framework generates a clear set of predictions for the path ahead. As PM execution costs continue to decline along the trajectory documented in the sample, categories will cross the Seesaw threshold in VRP-gradient order. Equity index events (VRP $\approx 1.9\%$) are next. FOMC (VRP $\approx 0.5\%$) will cross last, if ever, because the structural wedge is small and the execution cost floor is nonzero. Where negative VRP exists (silver), prediction markets will not displace options regardless of liquidity improvements. And where the binding constraint is observability rather than cost (Fed Funds options), the prediction market's advantage is categorical rather than marginal.