Michael Izhaky CEO, PredictRails
michael@predictrails.ai

April 29, 2026

Via Electronic Submission

Christopher Kirkpatrick, Secretary of the Commission
Commodity Futures Trading Commission
Three Lafayette Centre
1155 21st Street NW
Washington, DC 20581

Re: **Advance Notice of Proposed Rulemaking — Prediction Markets (RIN 3038-AF65; 91 FR 12516)**

Dear Secretary Kirkpatrick:

PredictRails respectfully submits the following comments in response to the Commission's Advance Notice of Proposed Rulemaking on prediction markets, published March 16, 2026. We appreciate the Commission's engagement and the opportunity to contribute a practitioner perspective.

PredictRails is building infrastructure for institutional participation in prediction markets, focused on reference data, compliance, and operational risk workflows for event contracts. This submission focuses on technical and operational observations where our direct implementation experience may be useful to the Commission. We do not offer legal conclusions or policy recommendations on the permissible scope of event contracts.

Key Takeaways for Staff

- The binding constraint on institutional adoption is **infrastructure, not demand**
- The primary issue is **resolution ambiguity**, not market access or product scope
- Event contracts fall into two distinct types (**at-expiration vs. trigger-based**) that require different approaches to margin, surveillance, and risk
- **Canonical contract identifiers** are required for position limits, cross-venue aggregation, and surveillance
- **Settlement oracles introduce a distinct risk vector** related to data coverage, revisions, and manipulation susceptibility
- **Cross-margining is essential** for institutional participation across multiple venues

- The appropriate regulatory approach is **technology-neutral and infrastructure-first**

Executive Summary

The prediction market sector has experienced rapid growth, from approximately $9 billion in combined annual volume in 2024 to over $44 billion in 2025, and a $275 billion annualized run rate in early 2026. Despite this growth, institutional participation remains constrained by the absence of standardized operational infrastructure.

The current state of prediction market infrastructure resembles the credit default swap market in the late 1990s and early 2000s: bespoke contract definitions, no standardized reference entity identifiers, inconsistent credit event language, and fragmented settlement procedures. Institutional capital could not scale into CDS until industry infrastructure — beginning with ISDA's first standardized credit derivatives definitions in 1999, significantly expanded in the 2003 revision, and culminating in the formalization of determinations committees under the 2009 Big Bang Protocol — made the contracts operationally governable. We observe a similar dynamic in prediction markets today: strong institutional interest, early professional participation, growing regulatory clarity, and widespread agreement that operational infrastructure must evolve before institutional capital can scale.

A structural observation that informs much of our analysis: event contracts currently traded on prediction markets encompass at least two fundamentally different settlement types that the ANPRM does not explicitly distinguish but that have materially different implications for resolution infrastructure, margin methodology, surveillance, and operational risk. Some event contracts settle only at a scheduled expiration date based on the state of the world at that moment (analogous to European-style options). Others resolve immediately upon the occurrence of a triggering event at any point before expiration (analogous to one-touch barrier options). These two settlement types require different infrastructure, different risk models, and different surveillance approaches. We reference this distinction throughout our comments where it is relevant.

We focus our comments on ten areas where we believe operational and technical observations can meaningfully inform the Commission's deliberations:

1. Resolution criteria and settlement standards (Question 2b)
2. Impartial access and pre-trade compliance controls (Question 2a)
3. Unique product identifiers for event contracts (Question 5a)
4. Position limits and cross-venue aggregation (Question 2e)
5. Margin and cross-margin (Questions 2f, 3c)
1. Institutional participation and workflow integration (Question 4)
2. Market surveillance across venues (Question 2d)
3. Inside information and market integrity (Questions 29-32)
4. Operational risk and automated systems (Question 2g)

5. Blockchain-based prediction markets (Question 2h)

I. Resolution Criteria and Settlement Standards (Question 2b)

The Commission asks what factors should govern resolution disputes and whether CDS dispute resolution procedures offer helpful precedents. We believe this is among the most consequential questions for the institutional viability of event contracts.

The Problem

Event contracts today lack standardized resolution specifications. Each designated contract market drafts its own contract language, defines its own resolution sources, and applies its own interpretive procedures when ambiguity arises. The result is that economically equivalent contracts — two contracts referencing the same underlying event on two different exchanges — can resolve differently based on textual variations in their terms.

This creates three concrete operational problems for institutional participants:

Resolution ambiguity. Contract text frequently contains void conditions, exclusion clauses, and venue discretion provisions that are not apparent from the contract's headline description. A contract titled "Will Event X occur by Date Y?" may contain conditions under which the contract resolves to a third outcome (void, refund, or partial settlement) that is not reflected in the binary price. Institutional risk systems cannot model contracts whose resolution logic is ambiguous or unstructured.

Cross-venue inconsistency. When two exchanges list contracts on the same event, differences in contract language can produce different settlement outcomes for the same real-world event. An institutional participant holding positions across multiple venues cannot reliably aggregate or hedge exposure without a canonical mapping of contract equivalence.

Operational fragility at settlement. Recent high-profile geopolitical event contracts have demonstrated the operational consequences of underspecified resolution criteria. When the resolution source or interpretive method is ambiguous, exchanges face post-hoc pressure to modify settlement terms, refund fees, or reinterpret contract language — all of which undermine market integrity and institutional confidence.

Settlement Type as a Resolution Variable

The resolution infrastructure required for an event contract depends critically on its settlement type. This distinction is not well-captured in current contract specifications, yet it has first-order implications for margin, surveillance, and operational risk.

At-expiration contracts (e.g., "Will inflation exceed 3% as of June 30, 2026?") resolve at a scheduled date based on a specific data release or observable condition. The resolution source is typically known in advance, the timeline is fixed, and the resolution process is a

single point-in-time determination. This is operationally straightforward, analogous to the settlement of a European-style option.

Trigger-based contracts (e.g., "Will Event X occur by December 31, 2026?") resolve immediately upon the occurrence of a triggering event at any point before expiration. This settlement type introduces substantially greater operational complexity:

- **Continuous monitoring.** The resolution source must be monitored in real time, not at a scheduled date. For contracts referencing geopolitical events, judicial decisions, or executive actions, the triggering event may occur at any hour of any day.
- **Ambiguity at the trigger boundary.** Determining whether a trigger has occurred can be substantially more difficult than evaluating a point-in-time condition. Has a political leader been "removed from office"? Does a partial withdrawal constitute a "cease-fire"? These boundary questions are inherent to trigger-based contracts and are the primary source of the resolution disputes the Commission is seeking to address.
- **Compressed settlement.** Once triggered, settlement may need to occur within hours, not days — creating operational demands that traditional settlement workflows do not accommodate.

We believe standardized resolution specification should require contracts to explicitly identify their settlement type, and that the Commission should consider whether different infrastructure requirements — including margin methodology, surveillance approach, and reporting obligations — should apply to each type.

Case Study: Resolution Risk in Practice

A recent event contract illustrates these dynamics concretely. In March 2026, a contract listed on a prominent prediction market venue asked whether shipping traffic through the Strait of Hormuz would "return to normal" by the end of April. The contract attracted attention as a proxy for energy market normalization, and the broader category of Hormuz-related contracts on the same venue generated over $25 million in combined trading volume.

The plain-language title suggested a broad real-world judgment about shipping conditions. The operative settlement rule, however, was narrower: the IMF PortWatch 7-day moving average of transit calls must reach or exceed 60. This is a trigger-based contract — it resolves YES the moment a qualifying print is published, at any point before expiration.

Structured analysis of the contract's resolution framework — performed before any independent field validation — identified six specific structural risks:

1. **Title-to-trigger mismatch.** "Returns to normal" is a semantic claim; the legal trigger is a specific threshold on a specific proxy metric.

2. **Source-model dependency.** The contract settles on one dataset (IMF PortWatch), not on all real-world traffic. Vessel activity not captured by that dataset does not contribute to settlement.
3. **Filtered-entity risk.** Only certain ship classes and AIS-reporting vessels are included. Ships operating without transponders are invisible to the settlement metric.
4. **Smoothing methodology.** Settlement depends on a 7-day moving average, not a spot reading, which may diverge from real-time conditions.
5. **Revision asymmetry.** Data revisions can cause the contract to resolve YES (qualifying a previously non-qualifying date), but cannot undo a prior qualifying print — an asymmetry that materially affects the contract's risk profile.
6. **Publication-surface ambiguity.** Settlement data appears across multiple formats (chart and downloadable files); the contract does not specify which controls in the event of a discrepancy.

Twelve days after this analysis was published, independent field research confirmed its practical implications: AIS data — the foundation of the PortWatch dataset — was found to be missing approximately half of actual vessel transits through the strait, including dozens of tankers operating without transponders. The market had been pricing the contract based on a broader economic interpretation than the dataset used for settlement could support.

This dynamic is not isolated. It illustrates a generalizable dynamic: **market participants can be economically correct about the underlying event yet contractually incorrect at settlement.** Institutional risk resides in the gap between what a contract appears to measure and what it legally settles on. Standardized, machine-readable resolution specifications would make this gap visible before settlement, not after.

Subsequent Developments

On April 7, 2026, the United States, Israel, and Iran agreed to a two-week ceasefire that included a coordinated reopening of the Strait of Hormuz. The contract price rose sharply. However, the resolution risk identified in the original analysis remains unchanged: Iran's "coordinated" reopening involves managed transit, not open passage, and the AIS-dark vessels that represent approximately half of actual strait traffic are unlikely to appear in the PortWatch dataset regardless of whether physical shipping normalizes. The ceasefire changed the real-world conditions but did not change the structural gap between what the contract measures and what is actually happening in the strait.

Contract Construction and Market Alignment

This case raises a broader question that the Commission may wish to consider: whether event contracts that serve as proxies for real-world economic conditions should be required to align their settlement mechanics with the economic risk they purport to represent.

The contract's title — "Strait of Hormuz traffic returns to normal by end of April?" — implies a durable assessment of shipping conditions. Participants in related markets — energy traders, macro desks, insurance underwriters — may read the contract's price as a real-time probability signal about the state of one of the world's most critical shipping chokepoints. However, the contract's trigger-based settlement structure means that a transient one-week spike in AIS-visible traffic could resolve the contract YES, even if conditions subsequently deteriorate and the strait is not functionally "normal" at the end of April or thereafter.

This creates a form of headline risk that extends beyond the contract itself. Where event contract prices are used as signals or hedging inputs in adjacent commodity, currency, or insurance markets, a misleading relationship between the contract's headline and its operative settlement mechanics can propagate incorrect assumptions into markets that have no direct connection to the prediction market venue.

We respectfully suggest that the Commission consider, as part of its evaluation of listing standards under Core Principle 2 and susceptibility to manipulation under Core Principle 3, whether event contracts whose titles imply durable real-world conditions but whose settlement mechanics can be satisfied by transient or non-representative data prints present a distinct category of market integrity concern. The remedy is not to restrict contract design, but to require that the relationship between the contract's headline description and its operative settlement mechanics be explicitly disclosed and clearly analyzable.

Settlement Oracle Reliability

The Hormuz example also highlights a distinct category of resolution risk that warrants the Commission's attention: the reliability of the data source — or "settlement oracle" — on which a contract's resolution depends.

Event contracts are unique among derivatives in that their settlement is determined entirely by an external data source rather than a market price. This creates a dependency on oracle reliability that has no direct analogue in traditional derivatives:

- **Coverage scope.** Does the oracle capture all economically relevant activity, or does it measure a subset? In the Hormuz case, the oracle excluded approximately half of actual vessel traffic.
- **Known exclusions.** Are there entity classes, geographies, or conditions systematically excluded from the oracle's measurement?
- **Historical reliability.** Has the oracle experienced outages, material methodology changes, or retroactive revisions that affected prior measurements?
- **Manipulation susceptibility.** Can participants influence the inputs to the oracle — through selective reporting, spoofing, or other actions that alter the dataset on which settlement depends — rather than manipulating the contract price directly?

This is a distinct manipulation vector that existing market surveillance frameworks do not address.

DCMs listing event contracts should disclose, in structured form, the settlement oracle, its known coverage limitations, its revision behavior, and any identified manipulation susceptibility — so that market participants and the Commission can evaluate oracle reliability on a consistent basis across contracts and venues.

The CDS Precedent

The Commission specifically asks whether CDS dispute resolution procedures are helpful precedents. We believe they are directly instructive.

Before ISDA began standardizing credit event definitions — first in 1999 and more comprehensively in 2003 — credit default swaps suffered from an identical problem: each dealer drafted bespoke contract language, reference entity identification was inconsistent, and credit event determinations were frequently disputed. Institutional adoption remained limited because risk managers could not reliably model instruments whose settlement behavior was unpredictable.

ISDA's response was to standardize three things:

1. **Credit event definitions.** A closed taxonomy of triggering events (bankruptcy, failure to pay, restructuring) replaced bespoke language. Each definition had specific, machine-testable criteria.

2. **Reference entity identifiers.** Standardized identifiers — notably Markit's Reference Entity Database (RED codes), later supplemented by LEI — replaced ad hoc entity naming. Cross-dealer and cross-product equivalence became deterministic.

3. **Determinations committees.** Centralized committees adjudicated ambiguous cases using structured procedures, published precedent, and defined escalation paths.

The effect was transformative. Once contracts were identifiable, standardized, and operationally governable, institutional capital scaled rapidly. Prediction markets are at a comparable stage today.

Recommendation

The Commission should consider requiring:

- **Structured resolution specifications** for all event contracts listed on DCMs. At minimum, resolution criteria should be disclosed in structured, machine-readable form and should enumerate: (a) the settlement oracle — the specific dataset or publication governing settlement; (b) the trigger definition, including thresholds, smoothing or averaging methodology, and settlement type (at-expiration or trigger-based); (c) the treatment of data revisions, including any asymmetric treatment of

qualifying versus disqualifying revisions; (d) known coverage limitations and exclusions in the settlement oracle; (e) the authoritative publication surface in the event of discrepancies between data formats; and (f) the dispute escalation procedure.

- **A standardized taxonomy of resolution event types**, analogous to ISDA's credit event definitions. Categories might include: official data release, judicial/regulatory action, verified third-party reporting, or discretionary venue determination. Each category should carry defined interpretive standards.

- **Settlement oracle reliability assessments.** DCMs should be encouraged to disclose a good-faith assessment of each settlement oracle's coverage scope, known exclusions, historical reliability, and susceptibility to input manipulation. No standardized framework currently exists for evaluating whether a settlement mechanism is fit for purpose across contracts and venues.

- **Structured dispute resolution procedures** for contested settlements, with defined timelines, escalation paths, and published precedent — modeled on ISDA determinations committee procedures.

II. Impartial Access and Pre-Trade Compliance Controls (Question 2a)

The Commission asks what aspects of prediction markets affect how a DCM provides impartial access and prohibits abusive trade practices, and whether there are particular barriers to impartial access.

The Access Problem for Institutional Participants

Core Principle 2 requires that a DCM provide access to its markets on terms that are "impartial, transparent, and applied in a non-discriminatory manner." For traditional derivatives, this requirement is well-understood: access criteria are based on financial qualifications, regulatory status, and operational connectivity.

Event contracts introduce a complication. Institutional participants operating under fiduciary obligations or regulatory mandates may need to restrict their own access — not because they are excluded by the DCM, but because their internal compliance policies require screening at the contract level before any order is submitted. A hedge fund that is prohibited by its limited partnership agreement from trading contracts referencing certain geopolitical events cannot rely on the DCM's listing standards alone; it needs the ability to programmatically evaluate each contract against its own policy set.

In traditional markets, institutional participants maintain restricted lists, sector exclusion policies, ESG screens, and sanctions databases that are evaluated pre-trade and enforced automatically. For event contracts, the absence of standardized contract metadata makes this type of automated screening impractical. Each institution must independently parse

natural-language contract terms to determine whether a given contract falls within its compliance boundaries.

Impartial Access Requires Standardized Metadata

Impartial access requires structured, machine-readable metadata. When contract metadata is unstandardized, only participants with the resources to build custom classification systems can effectively screen contracts against their policies. Smaller institutional participants — regional asset managers, smaller pension funds, community banks — face a disproportionate compliance burden that functions as a de facto barrier to access.

Standardized, machine-readable contract metadata (as discussed in Sections I and III) would lower this barrier by enabling third-party compliance tools to operate across all DCMs consistently. This would make institutional access to prediction markets more impartial, not less.

Recommendation

We respectfully suggest the Commission should consider requiring:

- Recognizing that institutional pre-trade compliance screening is a necessary component of impartial access — not an impediment to it — and that standardized contract metadata is a prerequisite for such screening at scale.
- Encouraging DCMs to publish contract metadata in formats that support automated compliance evaluation, enabling institutional participants to apply their own policy frameworks without requiring bespoke integration with each venue.

III. Unique Product Identifiers for Event Contracts (Question 5a)

The Commission asks whether CUSIP, ISIN, LEI, or comparable identifiers are appropriate for event contracts. We believe standardized event contract identifiers are essential infrastructure and that existing identifier schemes are insufficient for the purpose.

Why Existing Identifiers Are Insufficient

CUSIP and ISIN were designed for securities with static reference data: an issuer, a maturity, a coupon. An event contract's identity is defined by its resolution criteria, resolution source, expiration, and the specific event it references — none of which map cleanly to the fields in a CUSIP or ISIN record.

LEI identifies legal entities, not instruments, and is therefore orthogonal to the contract identification problem (though LEI may be useful for identifying resolution sources, such as the Bureau of Labor Statistics for CPI contracts or NOAA for weather contracts).

What Is Needed

Event contracts require a purpose-built identification system that encodes:

- **The underlying event** (e.g., "Federal Reserve federal funds rate decision, June 2026 FOMC meeting")
- **The contract's specific position on that event** (e.g., "rate cut of at least 25 basis points")
- **The resolution source** (e.g., "Federal Reserve press release, as published on federalreserve.gov")
- **The resolution methodology** (e.g., "binary: YES if announced rate is lower than current rate by ≥25bp")
- **The settlement type** (at-expiration or trigger-based, as discussed in Section I)
- **The expiration and settlement terms**

This identifier must be deterministic — two contracts with identical underlying events, positions, resolution criteria, and settlement type should produce the same identifier regardless of which exchange lists them. Critically, two contracts referencing the same event but with different settlement types (one trigger-based, one at-expiration) are not economically equivalent and must not share an identifier. This is the fundamental requirement for cross-venue equivalence: an institutional participant must be able to determine, programmatically, that Contract A on one exchange and Contract B on another exchange reference the same real-world event with the same resolution logic.

Industry experience suggests that this type of canonical mapping is technically feasible with existing technology. The absence of a standardized approach, however, means that each market participant must build this capability independently, creating fragmentation and inconsistency.

Recommendation

We believe standardized event contract identifiers are required. Specifically:

- Establishing a working group, potentially within the CFTC Innovation Advisory Committee, to develop a standardized event contract identifier specification. This effort could draw on the experience of CUSIP Global Services, ANNA (the Association of National Numbering Agencies), and ISDA's reference data infrastructure.

- Encouraging DCMs to publish machine-readable contract metadata — including resolution sources, resolution methodology, and void/exclusion conditions — in a standardized format. This would enable third-party reference data providers to build canonical mappings without requiring each participant to independently parse natural-language contract terms.

IV. Position Limits and Cross-Venue Aggregation (Question 2e)

The Commission asks what factors it should consider in adopting position limitations for prediction markets, including how position limits across similar event contracts should be aggregated and whether there is the same underlying reference.

The Aggregation Problem

Core Principle 5 requires DCMs to adopt position limits or position accountability "as is necessary and appropriate to reduce the potential threat of market manipulation or congestion." For traditional commodities and derivatives, aggregation rules are well-established: positions in economically equivalent contracts are aggregated across venues using standardized product identifiers.

For event contracts, aggregation is not currently operational at scale. Consider a participant who holds a position in "Will the Fed cut rates in June 2026?" on one DCM and an economically identical contract on a second DCM. Without a canonical identifier that maps these contracts as equivalent, neither the exchanges nor the Commission can determine whether the participant's aggregate position exceeds any applicable limit.

This is a structural supervisory gap. It is a structural gap in the Commission's ability to monitor concentration risk and enforce position accountability across the growing number of DCMs listing event contracts. As the number of venues and overlapping contract listings increases, the gap widens.

Canonical Identity as a Prerequisite for Position Aggregation

Effective position limit aggregation for event contracts requires the same canonical identification infrastructure discussed in Section III. Specifically:

- **Deterministic equivalence mapping.** Two contracts on two different exchanges must be programmatically identifiable as referencing the same underlying event with the same resolution logic. Without this, aggregation is manual, inconsistent, and unscalable.

- **Event-level grouping.** Beyond exact-match equivalence, position accountability may require grouping contracts that reference related but non-identical events (e.g., "Fed cuts rates by 25bp" and "Fed cuts rates by 50bp" reference the same FOMC decision with different strike conditions). This requires structured event metadata, not just contract-level identifiers.

- **Settlement-type awareness.** As discussed in Section I, a trigger-based contract and an at-expiration contract on the same event are economically correlated but not equivalent. The aggregation framework must determine whether and how to group positions across settlement types — a question that does not arise in traditional derivatives position limits.

- **Cross-venue reporting.** Aggregation requires that positions across multiple DCMs be reportable in a common format against a common identifier. This is the event contract equivalent of the CFTC's existing large trader reporting system — but it requires an identifier infrastructure that does not yet exist.

Recommendation

Effective position limit frameworks require:

- Recognizing that cross-venue contract equivalence standards would materially strengthen the effectiveness of any position limit or accountability framework for event contracts. Without canonical identifiers, cross-venue position aggregation will remain structurally challenging.

- Encouraging DCMs to report contract listings using structured metadata sufficient to enable automated equivalence determination by the Commission and by market participants.

V. Margin and Cross-Margin for Event Contracts (Questions 2f, 3c)

The Commission asks whether event contracts should trade on margin, how margin should be calculated, and what products are eligible for cross-margining. We address each in turn.

Margin Calculation

Event contracts have a defined maximum loss (the contract price for a buyer; the complement for a seller). This bounded loss profile simplifies margin calculation relative to traditional derivatives but introduces challenges that existing models do not address.

The settlement-type distinction discussed in Section I is particularly consequential for margin. Trigger-based contracts (one-touch) carry jump risk that can materialize at any moment — their risk profile resembles a barrier option, with the potential for instantaneous and complete loss of remaining value upon triggering. At-expiration contracts exhibit time-decay dynamics more analogous to European-style binaries, where value converges toward the binary payoff as expiration approaches. These two profiles require fundamentally different margin treatment.

Additionally, correlation-based margin offsets between event contracts are poorly understood and difficult to validate empirically. Two contracts on related events (e.g., "Will the Fed cut rates in June?" and "Will inflation fall below 3% in Q2?") have real economic correlation, but the correlation is event-driven and discontinuous — it does not follow the continuous distributions assumed by standard portfolio margining models (SPAN, CME CORE). Cross-margining between a trigger-based and an at-expiration contract on the same underlying event introduces further complexity, as the two settlement types respond differently to the same information flow.

These challenges are addressable but require event-contract-specific approaches rather than retrofitting existing derivatives margin models.

Cross-Margin

Cross-margining is the most important structural question for institutional adoption. Absent cross-margining, capital requirements scale inefficiently across venues, materially limiting institutional adoption regardless of market growth.

An institutional participant holding a portfolio of event contracts across multiple DCMs cannot currently net margin obligations across venues. This means the capital required to maintain equivalent positions across two exchanges is roughly double the economic risk. For traditional derivatives, cross-margining arrangements between clearing organizations (e.g., CME-OCC) address this problem. No comparable arrangements exist for event contracts.

The absence of cross-margining creates a capital efficiency barrier that disproportionately affects institutional participants. Retail participants typically trade on a single venue; institutional participants operating across multiple DCMs face multiplicative capital requirements.

Recommendation

We respectfully suggest the Commission should prioritize:

- Encouraging clearing organizations to develop event-contract-specific margin methodologies that account for binary payoff structures and event-driven correlations, rather than requiring retrofit of existing derivatives models.
- Facilitating cross-margining arrangements between clearing organizations that clear event contracts, with particular attention to the interoperability requirements for contract identification and resolution-criteria standardization discussed in Sections I and III above.

VI. Institutional Participation and Workflow Integration (Question 4)

The Commission asks how institutional trading on DCMs and SEFs differs from retail trading and how these differences impact liquidity.

Infrastructure Requirements

Institutional participation in any asset class requires a minimum set of operational infrastructure:

1. **Canonical instrument reference data** — a single, reliable source of what the instrument is, how it settles, and how it maps to other instruments. (See Section III.)

2. **Pre-trade compliance controls** — the ability to automatically screen instruments against firm-specific policies, regulatory restrictions, and risk limits before an order is submitted. In traditional markets, these controls are well-established (e.g., restricted lists, position limits, ESG screens). For event contracts, no standardized compliance framework exists.

3. **Post-trade processing** — trade confirmation, allocation, reconciliation, and integration with existing portfolio management and risk systems.

4. **Auditability** — the ability to reconstruct, for any trade, the compliance decision that permitted it and the risk parameters under which it was executed.

The Compliance Gap

Institutional participants require the ability to define and enforce policies at the contract level. For example:

- An asset manager may need to exclude all contracts referencing military or geopolitical violence, regardless of whether such contracts are listed on a DCM.
- A pension fund may need to restrict trading to contracts with verified, independent resolution sources (excluding contracts resolved at venue discretion).
- A bank may need to screen all event contracts against internal restricted lists and sanctions databases.

These use cases require structured contract metadata (as discussed in Section III), standardized classification schemas, and programmable compliance controls that can evaluate each contract against a firm's policy set before permitting a trade.

The technical approaches to this type of pre-trade policy evaluation are well-understood from analogous workflows in traditional markets. The challenge specific to event contracts is the absence of standardized contract metadata on which to operate.

Recommendation

Institutional participation requires:

- Encouraging DCMs listing event contracts to publish structured, machine-readable contract metadata sufficient for institutional pre-trade compliance workflows.

- Encouraging industry working groups to develop standardized event contract classification schemas that enable programmatic compliance screening — similar to GICS codes for equities or product classification in the ISDA taxonomy.

VII. Market Surveillance Across Venues (Question 2d)

The Commission asks whether there are market surveillance practices for detecting suspicious activity in other types of exchanges and platforms that would be useful in prediction markets, and what practices the Commission should focus on.

The Cross-Venue Surveillance Challenge

Traditional market surveillance operates within a single venue: an exchange monitors order flow, detects wash trading, flags unusual price movements, and reports suspicious activity. This model assumes that a single venue has visibility into the relevant trading activity.

Prediction markets break this assumption. The same economic event can be traded on multiple DCMs simultaneously, and the same participant can hold positions across venues. Single-venue surveillance cannot detect cross-venue manipulation — for example, a participant who accumulates a large position on one exchange and then manipulates the price on another exchange referencing the same event.

The Commission is aware of instances that have been the subject of public and congressional concern, where trading patterns on one platform appeared to correlate with geopolitical events while related contracts on other platforms exhibited different dynamics. Cross-venue surveillance requires the ability to identify equivalent contracts across exchanges — which, as discussed in Sections III and IV, requires canonical contract identifiers and cross-venue equivalence mapping that do not yet exist.

The settlement-type distinction discussed in Section I further complicates surveillance. Informed trading in trigger-based contracts produces different signatures than informed trading in at-expiration contracts: a trader with advance knowledge of a triggering event will act before the trigger, while a trader with a view on an at-expiration outcome may trade gradually as expiry approaches. Surveillance systems that do not account for settlement type will generate false positives on one pattern and potentially miss the other.

Existing Precedents

Consolidated surveillance is well-established in other markets:

- **Equities.** FINRA's Cross-Market Surveillance program monitors trading across multiple exchanges and dark pools using a consolidated audit trail (CAT). This is feasible because equities have standardized identifiers (CUSIP, ticker symbols) that make cross-venue matching trivial.
- **Listed options.** The Intermarket Surveillance Group (ISG) coordinates surveillance across options and equities exchanges, relying on standardized product identifiers and shared surveillance agreements.

- **CDS.** DTCC's Trade Information Warehouse, established in 2006 and expanded under post-crisis reporting mandates, centralized CDS trade records and enabled cross-dealer position visibility for the first time.

In each case, the surveillance capability was enabled by a prior investment in standardized identification infrastructure. Prediction markets currently lack this infrastructure, which means comprehensive cross-venue surveillance remains difficult to achieve at scale.

Recommendation

Effective surveillance requires:

- Recognizing that effective market surveillance for prediction markets requires cross-venue contract equivalence infrastructure. Single-venue surveillance is necessary but insufficient for a multi-DCM market.

- Encouraging the development of consolidated reporting mechanisms for event contracts — analogous to the consolidated audit trail for equities — that enable the Commission and self-regulatory organizations to monitor position concentration and trading patterns across venues.

VIII. Inside Information and Market Integrity (Questions 29-32)

The Commission asks about the tension between informed traders improving price accuracy and the fairness and manipulation concerns that information asymmetry creates.

The Operational Challenge

The Commission's questions in this area raise important issues of both policy and implementation. We confine our observations to the implementation dimension.

Event contracts present a spectrum of information-sensitivity profiles. At one end, contracts referencing publicly observable events (e.g., weather outcomes, published economic data) have relatively limited exposure to inside information. At the other end, contracts referencing government policy decisions, military operations, or actions by a small number of individuals present heightened risks of misuse of protected nonpublic information — risks that the Commission, Members of Congress, and the public have rightly identified as requiring serious attention.

The Commission already possesses authority under the Commodity Exchange Act to address misappropriation of nonpublic information.

Resolution Opacity as a Distinct Form of Asymmetry

The Commission's questions focus on the traditional information-asymmetry concern: participants with advance knowledge of the underlying event. We believe that prediction

markets exhibit a second, distinct form of asymmetry that is equally consequential: **resolution opacity.**

In many event contracts, the key informational advantage is not knowledge of whether the underlying event will occur, but understanding of how the contract resolves — specifically, the relationship between the plain-language title and the operative settlement mechanics. As the case study in Section I illustrates, a participant who understands the limitations of a settlement oracle, the implications of revision asymmetry, or the gap between a proxy metric and the economic condition it purports to measure can price the contract more accurately than participants who rely on the title alone.

This form of asymmetry — which arises from the structure of the contract rather than from access to nonpublic information — is not adequately addressed by traditional insider-trading frameworks. It is, however, directly addressable through standardized resolution disclosure: if all participants have equal access to structured, analyzable resolution data, the informational advantage from understanding settlement mechanics is eliminated.

An Infrastructure Approach

Effective market integrity requires infrastructure-level controls that make the Commission's existing authority operationally enforceable:

- **Contract classification and sensitivity labeling.** Event contracts can be systematically classified by their susceptibility to inside information along at least two dimensions: the nature of the underlying event (government policy decisions, where a defined set of individuals have advance access, versus publicly observable events) and the settlement type (trigger-based contracts, which are particularly vulnerable to advance knowledge of when a triggering event will occur, versus at-expiration contracts, where even insiders may face genuine uncertainty about the outcome at the scheduled date). A standardized classification schema incorporating both dimensions would allow intermediaries to apply differentiated surveillance and compliance controls based on each contract's specific risk profile.

- **Programmable eligibility controls.** Institutional intermediaries can implement pre-trade screening that restricts trading in information-sensitive categories based on the trader's role, affiliations, or access. This is analogous to the restricted-list and wall-crossing procedures used in securities markets.

- **Auditability.** Reference data and compliance infrastructure that maintains a contemporaneous record of what policies were applied, what contracts were screened, and what eligibility determinations were made. This audit trail is essential both for institutional compliance and for Commission enforcement.

Sensitivity to Recent Events

We note the serious concerns raised by Members of Congress and the public regarding event contracts tied to military operations, including concerns about potential misuse of

nonpublic information in connection with recent geopolitical events. These concerns implicate national security, public trust, and market integrity — and we believe they deserve the Commission's careful consideration.

Nothing in this comment should be read as urging the Commission to permit any particular category of sensitive event contracts, or as minimizing the gravity of the public interest and national security issues involved. Our observation is narrower: whatever substantive boundaries the Commission ultimately establishes for sensitive contract categories, those boundaries will only be effective if they can be operationalized through standardized event classification, pre-trade eligibility screening, and auditable compliance infrastructure. The policy determination and the technical implementation are inseparable.

IX. Operational Risk and Automated Systems (Question 2g)

The Commission asks what sources of operational risk are related to prediction markets, what operational risk analysis and other measures prediction markets currently employ, and whether there are challenges to the reliability, security, or scalable capacity of the systems used by prediction markets.

Unique Operational Risk Characteristics

Event contracts present operational risk characteristics that differ from traditional derivatives in several important respects:

Resolution dependency on external data sources. Unlike traditional derivatives that settle against a market price or benchmark rate, event contracts resolve based on external, often non-financial data: government announcements, judicial decisions, geopolitical events, or third-party reporting. This creates a dependency on data sources that may be ambiguous, delayed, disputed, or retroactively revised. Automated settlement systems must account for this uncertainty in ways that mark-to-market settlement does not require.

High contract proliferation rate. Prediction markets can list hundreds or thousands of contracts on a single event category (e.g., every possible outcome of every state in a national election). This proliferation rate exceeds what traditional derivatives exchanges typically manage and creates scalability challenges for risk systems, surveillance infrastructure, and compliance controls. Major prediction market venues currently list tens of thousands to hundreds of thousands of active contracts simultaneously.

Rapid time-to-resolution for trigger-based contracts. As discussed in Section I, trigger-based (one-touch) contracts resolve immediately upon the occurrence of a triggering event, which may happen at any time — including outside standard business hours. This creates operational demands for continuous resolution-source monitoring, near-real-time settlement processing, and reconciliation workflows that traditional T+1 or T+2 settlement cycles do not accommodate. At-expiration contracts, by contrast, have scheduled resolution dates and are operationally more predictable. The Commission may wish to

consider whether different operational risk requirements are appropriate for each settlement type.

Recommendation

Operational risk frameworks must address:

- Resolution-source reliability, high contract proliferation, and compressed settlement timelines as distinct risk categories that require purpose-built automated systems.
- Encouraging DCMs operating prediction markets to demonstrate that their automated risk and surveillance systems are designed for the specific operational characteristics of event contracts, rather than adapted from traditional derivatives infrastructure without modification.

X. Blockchain-Based Prediction Markets (Question 2h)

The Commission asks about challenges in applying existing regulations to blockchain-based prediction markets. We address one specific operational observation.

Technology Neutrality in Infrastructure Requirements

The infrastructure requirements outlined in the preceding sections — standardized contract identifiers, structured resolution criteria, pre-trade compliance controls, and auditable eligibility screening — apply equally to contracts traded on traditional DCMs and contracts traded on blockchain-based platforms.

The technical architecture differs (smart contracts vs. centralized matching engines, on-chain vs. off-chain settlement), but the operational requirements for institutional participation are identical. While institutional compliance and legal frameworks necessarily distinguish between regulated and unregulated venues, the underlying operational requirements for pre-trade risk management are functionally similar: both require canonical identification, resolution-risk analysis, and compliance screening before an order is submitted.

The regulatory framework should therefore specify functional requirements (what information must be available, what controls must be in place) rather than architectural requirements (how those controls are implemented). This technology-neutral approach would avoid creating regulatory arbitrage between on-chain and off-chain venues while ensuring that institutional participants can apply consistent compliance workflows regardless of the execution venue.

Conclusion

The prediction market sector is approaching institutional scale, but infrastructure, not demand, is the binding constraint.

The Commission's rulemaking has the opportunity to establish the foundational standards that enabled prior markets, such as CDS, to transition from bespoke instruments to institutional asset classes.

The core requirements are clear: standardized resolution, canonical identity, cross-venue interoperability, and operationally enforceable compliance frameworks.

These are not theoretical challenges—they are engineering problems with known solutions.

The CDS market's evolution from bespoke bilateral contracts to a standardized, institutionally accessible asset class demonstrates that coordinated infrastructure investment — by industry participants, standard-setting bodies, and regulators together — can unlock institutional adoption at scale. The prediction market sector is at a comparable stage today.

In practice, while this submission focuses on institutional infrastructure, the Commission's mandate extends equally to retail participant protection. We would encourage the Commission to consider how standardized contract metadata and structured resolution specifications could also benefit retail participants — by improving the transparency and comparability of event contracts available to individual traders and by enabling more effective disclosure of resolution risks.

PredictRails is available to provide additional technical detail or further observations and would welcome the opportunity to engage with the Commission and its staff.

Respectfully submitted,

Michael Izhaky
CEO, PredictRails
michael@predictrails.ai

[Remainder of page intentionally left blank.]

Appendix: ANPRM Question Cross-Reference

ANPRM Question	Topic	Comment Section
2a	Impartial access, abusive trade practices	II
2b	Resolution criteria, CDS dispute resolution precedents	I
2d	Market surveillance, suspicious activity detection	VII
2e	Position limitations, aggregation across similar contracts	IV
2f	Clearing and settlement, margin, cross-margin	V
2g	Operational risk, automated systems	IX
2h	Blockchain-based prediction markets	X
3c	DCO margin, cross-margin with securities	V
4	Institutional vs. retail trading differences	VI
5a	Product identifiers (CUSIP/ISIN/LEI)	III
29-32	Inside information, market integrity	VIII