Jeremy Dietz
Toronto, Ontario, Canada
github.com/JDSource/clearmarket

April 30, 2026

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

Submitted electronically via comments.cftc.gov

Re: RIN 3038-AF65 — Public Comment on Advance Notice of Proposed Rulemaking on Prediction Markets, 91 Fed. Reg. 12516 (March 16, 2026)

Dear Secretary Kirkpatrick:

I write in personal capacity. I have spent fifteen years inside regulated market infrastructure, in operating roles at the Toronto Stock Exchange and as VP Product at Canada's first CIRO-registered cryptocurrency exchange. The substantive basis for this comment is original data drawn from the public APIs of Kalshi and Polymarket on April 27, 2026, structured against ClearMarket: an open-source schema and reference dataset I maintain at github.com/JDSource/clearmarket that catalogs prediction-market contracts and their resolution rules in machine-readable form.

I confine this comment to eight of the Commission's questions where original cross-venue data and a working reference-data implementation give me a sharper view than I would have on the broader policy questions in this rulemaking.

Summary of Recommendations

(1) A DCM's Core Principle 2 obligations for event contracts should require structured, machine-readable resolution metadata at listing, not solely narrative rules text. Specifically, the *resolution mechanism* (arbiter), *resolution source type* (data-source classification), *resolution source name and URL with explicit source-priority and tie-breaking rules*, *proposer model*, and *resolution triggers* should be discrete disclosed fields. The four prominent prediction-market resolution failures of the past 24 months were not failures of disclosure presence; they were failures of placeholder language that did not specify which source controls when sources conflict. (Q2.b)

(2) Susceptibility-to-manipulation analysis under Core Principle 3 should explicitly contemplate resolution-source manipulation, a class of attack that has no direct equivalent in commodity futures and that the existing trading-side framework does not address. (Q2.c)

(3) A regulator-mandated per-contract identifier for CFTC-regulated event contracts at listing, following the CUSIP/LEI/UPI operating model (regulator sets the standard, designated

operator runs the registry), would materially improve reporting and surveillance. Cross-venue clustering of these identifiers into canonical-event groupings is properly a market-driven layer. (Q5.a)

Responses to Specific Questions

A. Core Principles and Commission Regulations

Q2.b: Resolution criteria and dispute resolution under Core Principle 2

The most consequential and least-developed surface in the prediction-market regulatory perimeter is the resolution layer. The dominant cause of disputed outcomes in the past 24 months has not been venue conduct or trading-side manipulation. It has been a specific failure mode in the contract rules text itself, which I refer to throughout this comment as **placeholder language**: phrasing that names a source *category* rather than a specific authority. Compare two formulations of the same kind of resolution rule:

- **Specific source named:** *"as published by the Federal Reserve Board of Governors at federalreserve.gov."* The reader knows in advance which document controls.
- **Placeholder language:** *"based on a consensus of credible reporting."* The reader knows the *type* of source the rule has in mind, but not which outlets count, what counts as "consensus," or which source controls when outlets disagree.

Resolution rules using placeholder language require human judgment when sources actually conflict, and the rule itself does not bind that judgment in advance. The four prominent failure cases of the past 24 months are all instances of this pattern.

The four cases (Hall, "What to do when prediction markets fail," a16z crypto, January 24, 2026):

- A Polymarket Ukraine territorial-control market resolved against a particular online map, with no rule specifying what to do if the map were edited. The map was allegedly modified during the contract life. The vulnerable surface was the source itself, not the order book.
- A U.S. government-shutdown contract resolved against the U.S. Office of Personnel Management (OPM) website's status. Funding legislation was signed November 12; the OPM site was not updated until November 13. The rule pinned resolution to OPM's website status without anchoring to the underlying legislative event, leaving traders correct on the legislative event to lose on the cited source's lag.
- A market on Zelensky's NATO-summit attire (over $200M traded) resolved against language asking whether his outfit was a "suit," without naming a specific source to make that determination. *BBC* and *New York Post* described the outfit as a suit; the market initially resolved YES on this basis. UMA (the decentralized oracle protocol Polymarket uses to resolve disputed markets through token-holder voting) flipped the resolution to NO. The dispute was about whose interpretation controlled, not what the underlying event was.
- A Polymarket market on the 2024 Venezuelan presidential election ($6M+ traded) named two resolution sources within the same rule that pointed in opposite directions: "official information from Venezuela" as primary, with "a credible consensus of credible reporting" as alternate. On election night, the Maduro-controlled electoral authority declared Maduro the winner; the opposition published polling-station tallies showing Edmundo González had won

by a wide margin. The market resolved for González, contradicting the named primary source. The rule itself was the failure: it specified two sources without specifying which controlled in a conflict.

The common thread across the four cases is not absence of disclosure. All four markets had verbose rules text that named, or appeared to name, a resolution source. The failure was that the rules used *placeholder language*: a map source without rules for handling edits; a website status without a controlling date; a vague trigger ("suit") without naming a source for making the determination; "official information" alongside an alternate source with no controlling rule. Such language required human judgment when sources actually conflicted. UMA token-holders or platform staff exercised that judgment; the resolution they reached was contested precisely because the rule had not pre-committed.

The data confirms the pattern is structural, not anecdotal. I pulled active-market metadata from Kalshi and Polymarket on April 27, 2026: 37,382 active Polymarket markets representing $6.0B in cumulative volume, and 10,165 active Kalshi markets representing approximately $594M in cumulative volume.

The two platforms expose resolution-rule text differently. Polymarket has a structured `resolutionSource` field designed to hold a URL plus a free-text `description` that contains the binding rules narrative. Kalshi has only free-text `rules_primary` and `rules_secondary` fields with no dedicated structured source field. The operational question on either venue is the same: does the rules text identify a specific authority in advance, or does it use placeholder language?

To measure placeholder language at scale, I scanned each market's resolution rules for phrasing that gestures at a source category without naming a specific authority ("consensus of credible reporting," "official sources," "primary source," and similar formulations).

On the table below, the per-market columns report what fraction of markets contain placeholder language. The volume-weighted columns report how much of the trading volume sits in those markets. The two can diverge sharply. For surveillance and systemic-risk purposes, the volume-weighted figure is the more diagnostic measure.

Category	Polymarket placeholder rate	Polymarket placeholder, volume-weighted	Kalshi placeholder rate	Kalshi placeholder, volume-weighted
Politics	71.7%	29.9%	7.8%	10.1%
Geopolitics	**77.7%**	**98.4%**	**68.9%**	**41.7%**
Culture (entertainment, awards)	51.2%	90.2%	3.1%	3.9%
Sports	83.1%	70.3%	0.0%	0.0%
Macro (Fed, BLS, indices)	29.2%	45.9%	0.1%	0.3%
Crypto	14.2%	11.4%	0.0%	0.0%
Tech / AI	22.2%	41.3%	0.0%	0.0%
Weather / Science	3.0%	56.9%	0.0%	0.0%
All categories	64.1%	43.9%	3.7%	6.7%

Two patterns are worth highlighting:

(i) Kalshi (CFTC-registered) has materially less placeholder language than Polymarket (unregistered) in most categories. Politics 7.8% vs 71.7%, sports 0.0% vs 83.1%, macro 0.1% vs 29.2%, culture 3.1% vs 51.2%. The data does not support a claim that prediction-market venues uniformly fail to document resolution; the regulatory perimeter is producing tighter rules-writing on Kalshi in most categories. The exception is geopolitics, where Kalshi runs 68.9% and Polymarket 77.7%. This is the only category where the regulated venue is roughly as bad as the unregistered one.

(ii) All four prominent resolution failures of the past 24 months arose at the politics-geopolitics intersection. On Polymarket specifically, the volume-weighted placeholder rate is 98.4% in geopolitics and 90.2% in culture. Almost every dollar of trading in those categories sits in markets where the resolution rule requires human judgment. Where placeholder language pairs with high volume, it pairs hard.

I recommend the Commission set the following expectation under Core Principle 2 for all event contracts:

A DCM's rules should require, at the time of contract listing, structured machine-readable disclosure of each of the following, not solely narrative rules text:

- *Resolution mechanism*: the arbiter that decides disputes (exchange staff, third-party oracle, named determinations committee, government agency, automated rule).
- *Resolution source type*: the data source's objectivity classification (statistical agency, central bank, regulated data vendor, court filing, issuer announcement, scheduled event, media consensus, subjective judgment).
- *Resolution source name and URL*: the specific feed, page, or release referenced; not the arbiter's admin page. *Where the rule names multiple sources, the order of priority and the rule for breaking ties when sources conflict.*

- *Proposer model*: for venues that delegate resolution proposal to third parties, whether proposers are permissionless, a managed whitelist, platform staff, or a government agency.
- *Resolution triggers*: parsed conditions, Boolean logic between them, explicit exclusions, and known ambiguity risks.

The italicized addition in the third bullet, the source-priority and tie-breaking requirement, is the specific element that would have prevented the Venezuela-style failure. It addresses the exact failure mode the data and the case studies support.

Machine-readable disclosure matters because surveillance, cross-venue comparison, and Commission oversight at the scale of approximately 1,600 new event contracts per year (the Commission's footnote 9 figure for 2025) cannot be performed through manual review of prose. A structured field is parseable by a surveillance system; placeholder prose is not. Structured disclosure makes the regulatory perimeter scalable in a way the current prose-only practice does not.

ClearMarket implements this structure across 15 canonical events and 55 markets, spanning Kalshi and Polymarket. Each of the five fields above is defined as a discrete data field with a controlled vocabulary, queryable through a public API and analyzable at scale. The schema is open-source. I do not propose ClearMarket itself as the standard; I propose the *pattern* (a structured, machine-readable resolution layer alongside the prose rules text), and offer the schema as a reference implementation that the Commission, a swap data repository, or an industry body could adopt or supersede.

The ISDA Determinations Committee precedent referenced in the Commission's Q2.b is partial. ISDA DCs adjudicate after a credit event has been alleged. The more pressing problem in event contracts is *pre-trade clarity*: specifying triggers with sufficient precision that the trigger event is observable, not negotiated. The CDS-DC parallel should inform Core Principle 14 dispute procedures but should not displace structured pre-trade disclosure under Core Principle 2.

Q2.c: Readily susceptible to manipulation under Core Principle 3

The Commission's existing Core Principle 3 framework adequately addresses trading-side manipulation: cornering, spoofing, layering, wash trades. For event contracts, a second class of manipulation exists that has no direct analogue in commodity futures: resolution-source manipulation, where an actor targets the resolution layer rather than the order book. The data shows two distinct vectors.

Source-attack vector. The named source is specific but brittle. Two recent cases illustrate the vector.

Data-editing attack: a Polymarket Ukraine territorial-control market resolved against a particular online map, which was allegedly modified during the contract life.

Physical tampering: in April 2026, a Météo-France weather sensor at Paris's Charles de Gaulle airport recorded two unexpected temperature spikes that crossed Polymarket resolution thresholds, prompting Météo-France to file a criminal complaint for physical interference with the sensor.

Different mechanisms, same failure mode: a single-source resolution feed compromised at the source rather than the order book. Mitigants for this vector are source diversity combined with an explicit aggregation rule (quorum, cross-validation, or named tie-breaking), preference for regulated or objective sources over single low-integrity ones, and explicit snapshot/versioning where the source is

mutable. Source diversity alone, without an explicit conflict-resolution rule, produces the source-ambiguity failure addressed below. The classification taxonomy in my Q2.b response is also a manipulation-susceptibility filter: a contract typed as `media_consensus` or `subjective` carries a different susceptibility profile than one typed as `gov_stat_agency` or `central_bank`.

Source-ambiguity vector. The rule itself is exploitable. Where rules use placeholder language ("consensus of credible reporting," "official sources," "primary source"), the vulnerability is not principally that bad actors will manipulate news outlets (an attack that is costly and visible). The vulnerability is that placeholder language creates a discretionary resolution moment: a window where platform staff or oracle voters must interpret the rule. That window is itself the attack surface, and it produces two specific susceptibility patterns.

(i) The resolver becomes the leak point. Anyone with foreknowledge of how the discretionary call will be made (platform staff, oracle voters) has tradeable information the market does not; influencing a small number of resolvers is also more tractable than influencing a media consensus. The Q29-32 response below addresses the legal basis for treating this as misappropriation.

(ii) The dispute itself becomes a trading window. Placeholder language extends time-to-resolution; during the dispute the market price is a forecast of the discretionary call rather than the underlying event. The Zelensky-suit market remained in dispute through the UMA voting cycle; the Venezuela market continued to trade after the Maduro-controlled authority's call until eventual resolution.

Both patterns collapse when the rule names specific sources with explicit tie-breaking. There is no discretionary moment to exploit if the rule says "Reuters and AP must agree; if they disagree, resolve NO." The Commission's susceptibility analysis under Core Principle 3 should treat placeholder language as itself a susceptibility flag, particularly in geopolitics where the placeholder-language rate exceeds two-thirds on both venues.

The Commission should expect DCMs to articulate, at listing for any contract that does not resolve against a single authoritative reference, both the source-diversity rationale and the tie-breaking rule. The two together address both vectors.

Q5.a: Standardized reporting and public identifiers

For event contracts listed as swaps on CFTC-registered venues, the SDR reporting requirement of CEA section 2(a)(13)(G) applies. The operative question is what standardization at issuance would make that reporting analytically useful, supporting the price discovery and cross-market surveillance that section 2(a)(13) contemplates. I recommend a two-tier approach to standardization: a regulator-mandated per-contract identifier at the issuance layer, and a market-driven canonical-event clustering layer above it.

Tier 1 (regulator-mandated): per-contract identifier at listing. The Commission should entrust a designated body to issue standardized identifiers for CFTC-regulated event contracts at listing, operated with low latency. The identifier is per-contract (each event-contract listing on a registered venue gets its own), following the CUSIP/LEI/UPI model: regulator sets the standard, designated operator runs the registry. Existing identifier systems are not fit for event contracts: each targets a different unit (securities for CUSIP/ISIN, legal entities for LEI, swap product types for UPI), not the individual event contract. The operating model, however, is well-precedented.

The principal friction in reporting event contracts to an SDR is venue-level fragmentation: no standardized identifier exists across event-contract venues, contracts on different platforms cannot be cross-referenced, and resolution rules are encoded as prose rather than structured fields. The identifier recommended here addresses the first two; the structured resolution metadata recommended in Q2.b addresses the third. Together they let market participants and the SDR speak a common reference language, supporting the price discovery the Commission's question contemplates and cross-market surveillance enabled by structured data.

Tier 2 (market-driven): canonical-event clustering across venues. With Tier 1 in place, market data operators can cluster per-contract identifiers into canonical-event groupings, recognizing when contracts on different venues refer to the same underlying event. This cross-venue clustering enables the direct price-discovery comparison across platforms that section 2(a)(13) contemplates. ClearMarket is one such reference implementation. This layer is properly market-driven, performed by data providers, SDRs, and reference-data utilities under competitive discipline rather than regulatory mandate.

B. Public Interest

Q12: Comparison to insurance contracts

On the public-interest question first: insurance availability for a given event should inform the Commission's analysis but not control it. Insurance markets and event contracts can be written on the same underlying event but reach different participants with different objectives. Cat bonds and event contracts on hurricane landfalls both reference catastrophe risk, but cat bonds reach institutional reinsurance investors while event contracts reach exchange traders. Both transfer risk; they just reach different participants.

Other event contracts (Fed rate decisions, election outcomes) exist where no insurance market does, because the events do not cause direct insurable loss but do create price-risk exposure participants need to hedge. A bond investor exposed to a Fed rate decision can hedge through an FOMC-decision event contract; a corporation with floating-rate debt can hedge interest-cost exposure through the same contracts. The public-interest purpose of these contracts is the one CEA section 3(a) recognizes for derivatives generally: managing and assuming price risks, discovering prices, or disseminating pricing information. Insurance availability is therefore at best a partial signal: the absence of an insurance market does not mean the event contract lacks public utility, and its presence does not make the event contract redundant.

On the comparison itself, insurance is heterogeneous. Event contracts are not like the indemnity insurance most readers know (auto, home), which pays out actual measured losses. They most resemble parametric insurance, which pays a fixed or capped amount when a defined external trigger is met (a wind speed, a temperature, a flight delay).

Until recently, this similarity was structural rather than functional. Event contracts and parametric insurance shared design features but served distinct risk-transfer needs. The line is starting to blur at the institutional layer. On April 27, 2026, Kalshi completed its first bespoke institutional block trade: a Houston-based environmental hedge fund hedged its exposure to California carbon allowance auction prices through a custom Kalshi contract structured by Greenlight Commodities, a CFTC-registered

Introducing Broker (Emily Nicolle and Bernard Goyder, "Kalshi Completes First Block Trade, Backed by Jump Trading," Bloomberg, April 27, 2026). The use case, an institutional risk-taker offloading event-specific exposure through a DCM-listed contract, sits in the functional space parametric insurance has historically occupied.

These features add up to a single conclusion: event contracts and parametric insurance are adjacent risk-transfer tools. They sometimes substitute, sometimes complement, and sometimes serve audiences the other does not reach. This is why insurance availability informs the Commission's public-interest analysis without controlling it.

E. Inside Information

Q29, Q30, Q31, Q32: Inside information and federal-employee restrictions

I confine my comment on these questions to the observation my data and case studies most directly support. Placeholder language in resolution rules creates a class of insider information distinct from advance knowledge of the underlying event: advance knowledge of how a discretionary rule will be interpreted at resolution. Where a contract resolves through human judgment by exchange staff or oracle voters, the resolver position is itself the exposed surface. The Polymarket geopolitics figure of 98.4% volume-weighted placeholder language, and the Kalshi geopolitics figure of 68.9% per-market, concentrate this exposure precisely where placeholder rates are highest.

A publicly articulated analytical framework relevant to this surface is the four-element misappropriation test Commissioner Caroline D. Pham described in her September 27, 2023 dissenting statement in *CFTC v. Xie* (No. 23-cv-1947, N.D. Ill.), grounded in CEA section 6(c)(1) and Regulation 180.1. The practical consequence of the absence of a CFTC-level standard is observable in the enforcement record: Matt Levine's coverage of Kalshi's enforcement actions ("Kalshi Found Some Insider Traders," *Bloomberg Money Stuff*, February 25, 2026) documents that Kalshi's venue-level rules ban any trading on insider knowledge, while a Pham-style CFTC standard would reach only misappropriation in breach of duty, and DOJ wire-fraud authority operates on yet a third standard. Whatever the Commission's resolution of that gap, the upstream fix recommended in Q2.b: structured pre-trade disclosure of source priority and tie-breaking rules, narrows the discretionary resolution moment the insider-information vector exploits, and in many cases eliminates it.

The federal-government-employee restrictions in CEA sections 4c(a)(3) and (4) extend cleanly to event contracts referencing federal-policy outcomes; interpretive guidance confirming application is sufficient.

Closing

The questions in this ANPRM cover wide ground. I have confined this comment to areas where original cross-venue data and a working schema implementation give me a sharper view than I would have on, for example, the scope of "gaming" in CEA section 5c(c)(5)(C), the contours of cost-benefit analysis under section 15(a), or the perimeter question of significantly similar instruments not listed on a CFTC-registered DCM or SEF, on which others are better placed to comment.

The April 27, 2026 data pulls and analysis scripts referenced in this comment are publicly available at github.com/JDSource/clearmarket. Readers can reproduce the segmented resolution-source figures by

running the same scripts against the public APIs.

I appreciate the Commission's consideration and am available to provide additional information.

Respectfully submitted,

Jeremy Dietz
Toronto, Ontario, Canada
github.com/JDSource/clearmarket