Submitted Electronically via CFTC Comments Portal

Christopher Kirkpatrick
Secretary of the Commission
Commodity Futures Trading Commission
Three Lafayette Centre
1155 21st Street, NW
Washington, D.C. 20581
April 30, 2026

Re: Prediction Markets Advanced Notice of Proposed Rulemaking (RIN 3038-AF65)

Dear Secretary Kirkpatrick,

Solidus Labs is submitting this comment in response to the Commission's Advanced Notice of Proposed Rulemaking on prediction markets. Solidus Labs is a born-in-crypto trade surveillance and market integrity firm with direct operational experience in prediction market oversight: we work with numerous prediction markets, including DCMs and on-chain platforms, DCOs, and SEFs, as well as market leaders such as Kalshi. The observations and analysis in this letter are drawn from that operational experience, from proprietary research into prediction market activity, and from our broader work monitoring a wide array of digital asset, securities, and commodity markets for regulators and trading venues globally. While the analysis below is framed primarily around DCMs as the registration pathway most commonly associated with prediction market activity to date, the structural surveillance challenges described apply to SEFs offering event contracts to institutional participants, and are relevant to DCOs to the extent that clearing infrastructure intersects with the participant identity and settlement integrity issues discussed herein.

We would like to note that members of our team have contributed to the response letter submitted by The Digital Chamber - which we strongly and wholeheartedly endorse. To avoid repetition and provide the most valuable feedback reflecting Solidus' expertise, in this response we focus on market surveillance-related issues throughout the ANPRM. These topics - and specifically the challenge of ensuring market integrity and preventing abuse and manipulation in innovative capital markets across both on- and off-chain

environments - are integral to our mission and the expertise we have developed over a decade of building and deploying our solutions in frontier markets like digital assets.

Our core submission is this: Prediction market surveillance cannot be treated simply as an extension of traditional futures market oversight. The structural differences are not incremental, they are categorical. Applying conventional surveillance frameworks as is to prediction markets produces simultaneous failure modes: false positives that exhaust compliance teams, and false negatives that allow real abuse to pass undetected. The Commission's rulemaking should consider reflecting this reality by establishing surveillance expectations calibrated to how these markets function.

This letter is structured by topic rather than question-by-question in response to the ANPRM; the questions to which each section responds are identified in the footnotes accompanying each section heading. It proceeds in four parts. We first explain why the surveillance challenge in prediction markets is structurally distinct from other derivatives markets. We then present specific documented cases of market abuse on prediction markets that illustrate these challenges in practice. Finally, we outline our views on necessary elements for effective surveillance standards of prediction markets, and specific suggestions in relation to DCM core principles.

Table of Contents for Quick Reference:

I. Why Prediction Markets Create Structurally Distinct Surveillance Challenges 3
- A. Bounded Prices Create Calibration Traps 3
- B. The Scale and Heterogeneity of the Symbol Universe 3
- C. Pseudonymous, Permissionless Access and Identity Fragmentation 6
- D. The Cross-On-and-Off-Chain Architecture of Onchain Markets 6
- E. Novel Manipulation Typologies With No Traditional Market Equivalents 7

II. Market Abuse in Practice: Documented Cases 8
- A. Profit Concentration as a Structural Indicator 9
- B. Wash Trading: $253 Million in the 2024 U.S. Election Market 9
- C. Insider Trading: The Axiom Case 10
- D. Coordinated Multi-Chain Insider Trading: Geopolitical Markets 11

III. What Effective Prediction Market Surveillance Requires 12
- A. Cluster-Based Behavioral Identity Analysis 12
- B. Dynamic Cross-Symbol Detection Logic 12
- C. Long-Horizon Behavioral Based Insider Trading Detection 12
- D. Integrated Cross-On-and-Off-Chain Surveillance 13
- E. Dynamic Category-Aware Parameterization and Bounded-Price Calibration 13
- F. 24/7 Realtime Continuous Surveillance Architecture 13

IV: Suggestions in Relation to DCM Core Principles 13

V. Conclusion 15

I. Why Prediction Markets Create Structurally Distinct Surveillance Challenges[1]

The Commission's ANPRM asks, among other things, whether existing surveillance, compliance, and enforcement practices for other types of exchanges would be useful in prediction markets. We address this directly: the answer is not without fundamental recalibration. The features that make prediction markets distinctive also make them uniquely difficult to surveil. We survey here five structural categories of challenge.

A. Bounded Prices Create Calibration Traps[2]

Event contracts are priced between $0.00 and $1.00. This seemingly simple feature breaks the foundational logic of traditional surveillance calibration. In conventional derivatives markets, percentage-based price movement thresholds are a reasonable proxy for economic significance. In prediction markets, they are actively misleading. A move from $0.03 to $0.04 represents a 33% price change but only a one-cent shift in absolute terms, economically trivial for most positions. A move from $0.80 to $0.90 represents only a 12.5% percentage change, yet depending on position size, can represent hundreds of thousands of dollars in PNL impact. Percentage-based alert logic inherited from traditional markets will fire constantly on low-probability contracts while systematically missing meaningful moves in high-probability ones. This is not a configuration problem that can be solved by adjusting thresholds. It is a fundamental modeling problem that requires segmentation by price band and probability range, an approach with no equivalent in traditional surveillance frameworks.

B. The Scale and Heterogeneity of the Symbol Universe[3]

A major traditional futures DCM may list hundreds or low thousands of instruments with broadly similar characteristics, settlement cycles, underlying asset classes, and trading hours are relatively standardized. Prediction markets are categorically different.

A single prediction market platform may have hundreds of thousands of simultaneously active contracts spanning political elections, sporting events, geopolitical developments, economic indicators, and crypto market milestones. Each contract has its own liquidity profile, resolution timeline, and category-specific manipulation risk. A sports event market peaks sharply and resolves within hours. A political market builds slowly over weeks and

[1] Responsive to ANPRM Questions 1, 2c, 2d, 2g, and 2h.
[2] Responsive to ANPRM Question 2d.
[3] Responsive to ANPRM Questions 2d and 2g.

may be particularly susceptible to insider positioning in the days before resolution. An economic data market has entirely different manipulation characteristics tied to information release schedules.

Applying uniform detection logic across this heterogeneity produces oversight gaps. Effective surveillance requires category-aware parameterization that distinguishes legitimate volatility in live sports markets from the patient, suspicious positioning patterns characteristic of geopolitical or political markets.

A further structural distinction is that prediction market contracts must be evaluated both in isolation and as part of a broader group of related contracts. Most contracts exist within a taxonomy of explicitly related contracts (e.g., "Tom Brady Returns on June 5", "Tom Brady Returns on June 6", etc.) but also within an array of loosely to strongly correlated contracts (e.g., "Fed will increase rates by 50 bps"; "The US will enter a recession in Q2"). The former contracts are tied to the same outcome, governed by different resolution dates, whereas the latter contracts are tied to different outcomes and circumstances but could be strongly correlated and influenced by the same subset of insiders with access to privileged economic information. Thus, the event contract must be looked at in isolation, in conjunction with its explicitly related contracts, and in combination with more loosely related contracts. This has direct implications for detecting manipulation across a number of typologies: a holistic picture of an entity's activity and intentions requires a view into trading and positioning across a wide array of contract symbols and the architecture to ingest and maintain a dynamic set of related contract symbols.

The multi-outcome structure of many prediction market events compounds this challenge further. While binary outcome markets - contracts that resolve to YES or NO - represent the most common prediction market structure, a significant and growing share of prediction market events involve three or more mutually exclusive outcomes. An election market with ten candidates, for example, generates ten correlated symbols, each with its own Yes and No token pair. These symbols are not independent instruments. If one outcome's implied probability rises, others must fall. Conventional surveillance frameworks, which ingest each symbol as an independent instrument, misread this correlation structure entirely.

The surveillance implications are material. In a binary market, an actor who buys both YES and NO tokens in equal dollar-weighted proportions is paying approximately the full $1.00 probability-space cost and taking no directional risk - a clear wash trade signal. In a multi-outcome market with ten candidates, the same logic does not apply directly. An

actor who buys two low-probability outcomes at $0.05 each spends only $0.10 in total, covers only 10% of the probability space, and retains genuine directional exposure. A detection methodology that treats this as wash trading produces false positives. One that fails to account for the full probability-space cost across all outcomes in the event will miss genuine cross-symbol wash trading entirely. Effective detection in multi-outcome markets requires grouping all outcome symbols by their parent event, computing delta-neutrality across the full outcome set in a price-aware manner, and applying detection thresholds that reflect the actual structure of the market - not assumptions imported from binary prediction markets or traditional derivatives.

A related challenge arises from markets that are functionally equivalent but structurally unlinked. Two independently listed contracts may ask materially the same question using different phrasing, or may share nearly identical resolution criteria despite carrying distinct symbols with no programmatic relationship in the data model. A contract titled "Will the Fed cut rates in June?" and one titled "Federal Reserve rate decision - cut at the June 2026 FOMC meeting?" are, for surveillance purposes, the same instrument - but conventional surveillance systems have no mechanism to recognize this. An actor who splits a large informed position across both contracts would evade any per-symbol position, volume, or PnL threshold, while simultaneously generating separate alerts on each - inflating alert counts without surfacing the coordinated behavior that matters.

Addressing this requires a semantic linkage layer applied to both market title text and resolution criteria language prior to surveillance evaluation. Resolution criteria text is often a stronger signal of functional equivalence than market titles alone: two contracts may have divergent titles but resolve based on the same underlying data release, court ruling, or official announcement. Surveillance programs that cannot identify and cluster semantically equivalent markets will systematically undercount position concentration, undercount PnL, and overcount alert volume - three simultaneous failure modes that undermine the integrity of any oversight program.

Solidus research finds that in active prediction markets, more than 90% of contracts trade for less than one day. This poses significant challenges to a foundational assumption in conventional surveillance models - that instruments have meaningful price and volume history against which current activity can be benchmarked. When many contracts have no multi-day history, reference-data-dependent detection logic has nothing to reference. Alert thresholds calibrated to rolling averages, historical volatility bands, or multi-day

volume baselines can be structurally blind to the majority of the prediction market symbol universe.

Addressing this effectively requires different detection architecture. Behavioral surveillance, which identifies anomalous patterns relative to actor conduct across resolved events rather than instrument price history, is the only methodology that functions when instrument history is measured in minutes or hours rather than months. Surveillance systems built to monitor thousands of instruments with deep historical records cannot simply be extended to cover millions of ephemeral contracts. Horizontal scaling for symbol ingestion, automated schema normalization for free-form contract titles that do not conform to standard financial instrument identifiers, and self-service threshold management across a dynamic symbol universe are not optional features. They are prerequisites for meaningful oversight.

C. Pseudonymous, Permissionless Access and Identity Fragmentation[4]

For onchain prediction markets, the participant identity model is even more fundamentally different from traditional derivatives markets. Especially in fully decentralized platforms, there is no KYC at the point of market access. A single actor can control hundreds of pseudonymous wallets, each appearing as an independent participant. The onchain funding infrastructure compounds this: participants fund their trading accounts through proxy wallets, often via Gnosis Safe or EIP-1167 smart contracts, which further obscure the relationship between the funding entity and the trading identity.

Per-address monitoring is structurally insufficient in this environment. An actor amassing $200,000 of insider trading profit across ten wallets may evade detection at the wallet level. The same activity becomes visible when the surveillance tool naturally ingests onchain data and has the ability to cluster onchain wallets based on behavioral heuristics like age of wallet, common funders, timing correlation, bridge fingerprints, and execution patterns. This is a detection methodology with no equivalent in traditional surveillance.

D. The Cross-On-and-Off-Chain Architecture of Onchain Markets[5]

Major onchain prediction markets like Polymarket operate through a hybrid architecture that is not captured by either purely onchain or purely off-chain surveillance: customers fund their accounts onchain, typically through Gnosis Safe proxy wallets, but order

[4] Responsive to ANPRM Questions 2d and 2h.
[5] Responsive to ANPRM Questions 2d and 2h.

matching occurs off-chain on a central limit order book (CLOB). That CLOB then instructs a smart contract to swap USDC for YES or NO outcome tokens, settling on-chain.
This cross-on-and-off-chain structure is not a technical curiosity. It is the primary mechanism through which manipulation is executed and obscured. Funding obfuscation happens onchain, through multi-chain bridge routes spanning Polygon, Solana, Ethereum, and BSC. The coordinated trading activity happens off-chain, through the CLOB. Profits are collected and consolidated onchain. Surveillance that covers only one domain, either the CLOB or the blockchain, cannot connect the dots. Both layers must be integrated in a single detection framework.

A relevant July 2023 study by the Japan Financial Services Authority has explored the question of whether on-chain data alone is enough to assess the integrity and health of decentralized assets and ecosystems.[6] The study found that "the large amount of off-chain data in the crypto asset market, including DeFi, makes it difficult to get a picture of overall market activity from on-chain data alone."

This challenge is distinct from "cross-chain" surveillance, which refers to monitoring across multiple blockchains. Cross-on-and-off-chain surveillance refers to the integration of offchain data including orderbook activity, account deposits and withdrawals, KYC, social media, news and other non-blockchain data with onchain activity, wallet behavior and funding flows. In blockchain-enabled trading environments, both capabilities are critical and neither is sufficient without the other.

On-chain markets also trade 24/7 and are immutable, while traditional surveillance infrastructure is built around exchange hours and delayed processing. Insider-motivated positioning in prediction markets concentrates in low-activity windows precisely because surveillance coverage is weakest during those periods. An architecture designed around batch-cycle processing tied to session hours is structurally blind to the activity patterns that matter most.

E. Novel Manipulation Typologies With No Traditional Market Equivalents[7]

Several forms of market abuse in prediction markets have no direct equivalent in traditional derivatives markets, making detection using conventional surveillance logic

[6] Report of FSA's Joint Research on Analyzing Decentralized Financial System using On-Chain and Off-Chain Data (https://www.fsa.go.jp/policy/bgin/Research_Paper_QUNIE_en.pdf)
[7] Responsive to ANPRM Questions 2a, 2c, and 2d.

inefficient (specific documented examples are provided in Section II of this response). One clear example is cross-symbol wash trading. In conventional derivatives markets, wash trading involves an actor buying and selling the same instrument to create artificial volume. Prediction markets introduce a structurally distinct variant: an actor simultaneously buys YES tokens and NO tokens on mutually exclusive outcomes in the same market. Because outcomes must sum to approximately $1.00, the combined position is economically delta-neutral, no directional risk, no economic exposure, but every individual transaction registers as real volume. Detection requires evaluating position neutrality across outcome tokens within a parent event, a data model and algorithmic approach with no equivalent in traditional finance.

A second novel typology is narrative manipulation. A single misleading social media post or premature rumor can simultaneously shift prices across multiple related event contracts. This creates a combined social-media and cross-product surveillance challenge that requires correlating news and sentiment signals with order flow across related markets, a capability that traditional surveillance does not contemplate.
Insider trading in prediction markets also presents differently than the front-running model familiar from traditional markets. In futures and equities, insider trading typically manifests as concentrated activity immediately before a material announcement, a minutes-scale behavioral signature. In prediction markets tied to geopolitical, political, or event-driven outcomes, informed actors position days or weeks in advance of resolution, with position sizes that may not appear unusual in isolation but are revealing when viewed against the actor's account history, funding patterns, and the timing of market resolution. Per-address, price-trigger-based detection is structurally blind to this behavioral signature.

II. Market Abuse in Practice: Documented Cases[8]

The surveillance challenges described above are not theoretical. Solidus Labs has documented specific cases of each manipulation typology on active prediction markets. The following cases are drawn from Solidus' HALO Trade Surveillance system, operating on real market data.

[8] Responsive to ANPRM Questions 2c, 2d, and Part II.E.

A. Profit Concentration as a Structural Indicator[9]

Before examining individual abuse cases, it is worth noting a market-level finding from Solidus' analysis of approximately 7.9 million trades across Polymarket's Politics category between December 2025 and February 2026, spanning 94,222 maker addresses and 353,019 taker addresses.

0.55% of profitable maker addresses (201 wallets) captured 50.0% of all maker-side gains ($8.02M of $16.04M total). On the taker side, 0.26% of winning addresses (316 wallets) captured 49.7% of all taker-side gains ($8.36M of $16.83M total).

In both maker and taker populations, fewer than 1% of addresses captured approximately 50% of all profitable returns over a three-month period. Not all concentrated profit is evidence of wrongdoing, skilled quantitative traders and sophisticated market-makers can legitimately outperform retail participants. But when the concentration is this extreme, and when the same wallets recur across surveillance alerts, it demands systematic scrutiny. It is also a baseline against which individual cases of suspected abuse can be evaluated.

B. Wash Trading: $253 Million in the 2024 U.S. Election Market[10]

Solidus analyzed 7.18 million executions on Polymarket's 2024 U.S. Election market from October 10 to November 6, 2024, covering approximately $1.87 billion in notional traded over 27 days. Two distinct wash trading patterns were defined and measured.
The first pattern, intra-symbol wash trading, mirrors the classic A-B-A structure familiar from traditional markets: a single account buys and sells the same outcome token within a tight time window, ending near-flat but generating inflated gross volume. Detection uses a roundtrip score (1 minus absolute net quantity divided by gross quantity), with a VWAP spread filter to exclude legitimate market-making activity.

The second pattern, cross-symbol wash trading, is unique to prediction markets and has no equivalent in traditional finance. The actor simultaneously buys YES tokens on mutually exclusive outcomes (e.g., buying both "Trump Win" and "Harris Win" within the same 15-minute window). Because outcomes must sum to approximately $1.00, the combined position is economically delta-neutral. The metric is a delta-neutrality score computed as the sum of (quantity times price) across positions divided by maximum single-symbol

[9] Responsive to ANPRM Questions 2d and Part II.E.
[10] Responsive to ANPRM Questions 2a and 2d.

notional; a score of 0.95 or higher with a combined cost approaching $1.00 identifies the pattern. Unlike traditional markets, where taking offsetting positions across related contracts still carries residual risk, binary prediction markets offer a mechanism for genuinely riskless volume inflation with no traditional-market analogue.

Results: $253 million in gross notional flagged as wash trading across both patterns, 14% of total market volume. 18,300 accounts flagged for intra-symbol wash trading (9.4% of active accounts). 519 accounts flagged for cross-symbol wash trading, with the top 25 accounts accounting for 99.1% of flagged cross-symbol volume.

A secondary finding illustrates both the validity of the detection methodology and the nature of the behavior: the ratio of flagged to total volume dropped sharply in the days immediately preceding the election, as genuine organic interest surged and diluted the proportion of inorganic volume. Wash trading concentrated in low-organic-interest periods, exactly the pattern the algorithm is calibrated to identify.

C. Insider Trading: The Axiom Case[11]

In late February 2026, onchain investigator ZachXBT announced he would be publishing an investigation into a crypto firm called Axiom. Polymarket listed a cluster of contracts on the news, including the contract WILL-AXIOM-BE-ACCUSED-OF-INSIDER-TRADING. HALO's insider trading algorithm flagged unusual activity in this contract before the announcement.

A newly created Polygon wallet opened a large position in the contract just before a sharp price spike. For approximately four hours, this single wallet accounted for a dominant share of total market volume. As the market resolved to YES, the wallet exited its position, realizing $354,000 in profit. The wallet was created and funded only hours before the price spike, a short-lived, single-purpose account with a funding source (Ethereum wallet 0xdb407) that connected it to a second flagged wallet.

The second wallet realized $160,000 in profit on the same contract. It was funded via six cross-chain Relay.link bridge fills spanning four blockchains (Solana to Polygon to Ethereum to Polygon), with the same Ethereum wallet funding both accounts on consecutive days. The second wallet also spread its positions across related markets,

[11] Responsive to ANPRM Questions 2d and Part II.E.

buying NO in the WILL-METEORA-BE-ACCUSED-OF-INSIDER-TRADING contract, signaling that the actor knew which firms were in scope of the investigation and which were not. Detection mechanism: HALO connected both wallets through a shared Ethereum funder cluster link, an alert that is invisible to any per-address KYT monitoring tool. Neither wallet would have been flagged in isolation; the connection emerged only through behavioral network analysis across wallet clusters.

Combined profit across the two wallets: over $514,000. Material non-public information in this case was not a state secret, it was advance knowledge of a scheduled media publication. That was sufficient to generate half a million dollars in profit across two coordinated wallets with onchain breadcrumbs that only became visible through the integration of trade surveillance and blockchain analytics.

D. Coordinated Multi-Chain Insider Trading: Geopolitical Markets[12]

On February 15, 2026, a newly created wallet opened a large position in the US-STRIKES-IRAN-BY-FEBRUARY-28 market just before a sharp price spike. The market resolved to YES. The wallet realized $27,000 in profit, modest by absolute terms, but the onchain infrastructure surrounding it was significant.

The Polymarket proxy wallet was funded days before market resolution with $4,943.97 USDC, routed through a multi-chain obfuscation chain driven by an intermediary hub address (0xc4fd...4b39). After collecting winnings, the wallet bridged its entire balance to a Solana address (4dLZ...4Fwy), which turned out to be a consolidation wallet for three other Polymarket proxy wallets. Two of those three appeared in separate HALO insider trading alerts in the same geopolitical market.

Result: Four coordinated wallets. Combined return of 500%+. Funded through multi-chain obfuscation on BSC, Polygon, and Solana. Cashed out to the same Solana consolidation address within hours of market resolution.

This case illustrates why surveillance that covers only a single blockchain or only the off-chain CLOB is insufficient. The funding chain, the trading activity, and the profit consolidation each occurred on different layers and different chains. Only integrated cross-on-and-off-chain surveillance can connect them into a single identifiable pattern.

[12] Responsive to ANPRM Questions 2d, 2h, and Part II.E.

III. What Effective Prediction Market Surveillance Requires[13]

The cases documented above, and the structural challenges that enabled them, raise a coherent set of surveillance capabilities key to effectively monitoring prediction markets. We offer the below suggestions for the Commission's consideration as it develops its approach to DCM surveillance programs as a matter of regulatory standard-setting.

A. Cluster-Based Behavioral Identity Analysis[14]

Per-wallet or per-account monitoring is not adequate for markets where a single actor may control hundreds of pseudonymous addresses. Effective surveillance must aggregate behavioral signals across wallet networks, shared funders, timing correlation, bridge fingerprints, execution patterns, to construct actor-level identity clusters. The Commission should consider expecting DCMs to demonstrate cluster-level detection capability, not merely per-address thresholds.

B. Dynamic Cross-Symbol Detection Logic[15]

Cross-symbol wash trading, simultaneous same-side trading across mutually exclusive outcomes, is a manipulation typology unique to prediction markets. Detection requires grouping symbols by parent event and evaluating position delta-neutrality across correlated outcome tokens. No traditional surveillance system can detect this pattern without purpose-built logic for the prediction market data model. The Commission should consider ensuring that DCM surveillance obligations explicitly contemplate this typology.

C. Long-Horizon Behavioral Based Insider Trading Detection[16]

Informed actors in prediction markets position days or weeks before resolution, not seconds before a price move. Effective detection requires behavioral pattern analysis calibrated to days-to-weeks lead times, tracking prediction accuracy across resolved events, wallet lifecycle signals, position sizing relative to account history, and timing patterns. Price/volume triggers calibrated to minutes-scale front-running are not sufficient. Critically, behavioral detection is equally necessary at the other end of the time spectrum. For the more than 90% of prediction market contracts that trade for less than one day, conventional surveillance approaches dependent on multi-day reference data have no historical baseline to work from. A system that cannot detect manipulation in a contract that lists and resolves within hours is blind to the majority of the prediction

[13] Responsive to ANPRM Questions 1, 2a, 2c, 2d, 2g, and 2h.
[14] Responsive to ANPRM Questions 2d and 2h.
[15] Responsive to ANPRM Questions 2a and 2d.
[16] Responsive to ANPRM Questions 2d, 2g, and Part II.E

market symbol universe. Behavioral surveillance - which evaluates actor conduct across resolved events rather than instrument price history - is the only detection methodology that functions across both long-horizon positioning and short-lived ephemeral contracts.

D. Integrated Cross-On-and-Off-Chain Surveillance[17]

For onchain prediction markets, the manipulation chain spans onchain funding activity, offchain order book execution, and onchain settlement and profit consolidation. Surveillance that covers only one domain cannot reconstruct a complete picture of coordinated activity. The Commission should consider establishing as a regulatory expectation that DCMs operating onchain prediction markets maintain integrated surveillance that correlates offchain order book data with onchain funding flows, wallet activity, and multichain bridge transactions.

E. Dynamic Category-Aware Parameterization and Bounded-Price Calibration[18]

Surveillance parameters must be calibrated to the specific characteristics of each event category. Political and geopolitical markets exhibit slow-building positioning patterns that require long-horizon detection windows. Sports markets exhibit sharp, legitimate volatility that requires different thresholds to avoid false positives. Price movement alerts must be calibrated in absolute terms and probability bands, not percentage terms, to account for the $0.00 to $1.00 pricing structure of event contracts.

F. 24/7 Realtime Continuous Surveillance Architecture

Prediction markets operate continuously. Insider activity concentrates in low-activity windows precisely when batch-cycle surveillance systems are offline. The Commission should not accept surveillance programs built on session-based processing architectures as adequate for prediction market oversight. The Commission should consider mandating realtime surveillance systems that can detect abuse in single day instruments before the markets resolve. With over 90% of prediction markets contracts lasting less than 1 day, T+1 batch processing is no longer sufficient.

IV: Suggestions in Relation to DCM Core Principles[19]

The surveillance capabilities described above are not aspirational standards external to the Commission's existing framework. They are the practical content of obligations the

[17] Responsive to ANPRM Questions 2d and 2h.
[18] Responsive to ANPRM Questions 2d and 2g.
[19] Responsive to ANPRM Questions 1, 2a, 2c, 2d, 2g, 2h, 3a, 3b, and 4.

Commission has already imposed on DCMs through the Core Principles - obligations that have yet to be extensively interpreted in the context of prediction markets and require clarification to have meaningful effect.

Core Principle 2 requires that a DCM establish, monitor, and enforce rules prohibiting abusive trade practices. Cross-symbol wash trading - the simultaneous purchase of YES and NO tokens on mutually exclusive outcomes to generate artificial volume - is an abusive trade practice with no equivalent in traditional markets and no analogue in standard surveillance tooling. The Commission should consider clarifying that CP2's prohibition on abusive trade practices encompasses manipulation typologies native to prediction markets, and that a DCM's enforcement program must be capable of detecting them.[20]

Core Principle 3 requires that a DCM list only contracts not readily susceptible to manipulation. The Commission should consider recognizing that this determination cannot be made in a meaningful way at the contract listing stage without surveillance infrastructure capable of detecting the manipulation typologies specific to that contract's category. A political market, a geopolitical market, and a sports market each present distinct manipulation risks and require distinct detection parameters. The CP3 listing determination and the CP4 surveillance obligation are therefore interdependent for prediction markets in a way they are not for traditional futures contracts.[21]

Core Principle 4 requires that a DCM have the capacity and responsibility to prevent manipulation through market surveillance, compliance, and enforcement practices and procedures. For prediction market DCMs, the Commission should consider establishing through guidance that CP4 compliance will be evaluated against the manipulation typologies native to these markets, not against standards developed for traditional futures exchanges. At minimum, a prediction market DCM's surveillance program should demonstrate the capacity to detect manipulation patterns operating across long time horizons, across correlated outcome contracts within the same parent event, and across pseudonymous participant identities at the cluster level rather than the individual address level. A surveillance program that applies percentage-based price thresholds to $0.00-$1.00 contracts, or that evaluates each outcome token independently without

[20] Responsive to ANPRM Questions 2a and 2d.
[21] Responsive to ANPRM Questions 1 and 2c.

reference to related contracts in the same market, does not satisfy CP4 as applied to prediction markets regardless of whether it satisfies CP4 as applied to traditional futures.[22]

For DCMs operating on blockchain infrastructure, CP4 compliance additionally requires that surveillance capacity extend across both onchain and offchain activity as an integrated whole. The manipulation patterns documented in this letter - coordinated funding obfuscation onchain, order execution offchain, profit consolidation back onchain across multiple blockchains - are invisible to any surveillance program that covers only one domain. A DCM whose surveillance program cannot follow a manipulation pattern across the full lifecycle of a trade, from the funding wallet through the offchain order book to onchain settlement, does not have the capacity CP4 requires. Hence, the Commission should consider placing a requirement for surveillance programs to include cross-on-and-off-chain coverage.[23]

Core Principle 20 requires that a DCM establish and maintain automated systems with reliable, secure, and adequate scalable capacity. For prediction market DCMs, the Commission should consider interpreting this to require surveillance systems capable of continuous, real-time monitoring across the full symbol universe of the platform. A system built on batch-cycle processing tied to exchange session hours cannot satisfy CP20 for a market that operates 24 hours a day, seven days a week, across hundreds of thousands of simultaneously active contracts. Insider-motivated positioning in prediction markets concentrates precisely in low-activity windows - the periods when batch-cycle systems are most blind. CP20's scalable capacity requirement should be understood to encompass both the volume of simultaneous contracts and the continuity of surveillance coverage across all hours of operation.[24]

V. Conclusion

Prediction markets represent a structurally distinct derivatives category. They are not futures contracts with an unusual underlying. They are a distinct derivatives category with manipulation typologies, participant identity models, and market microstructure that require surveillance frameworks built specifically for these characteristics for effective monitoring.

[22] Responsive to ANPRM Question 2d.
[23] Responsive to ANPRM Questions 2d and 2h.
[24] Responsive to ANPRM Questions 2d and 2g.

The Commission has the opportunity, through this rulemaking, to establish surveillance expectations that reflect the actual operational reality of these markets. Doing so will accomplish two objectives simultaneously: it will require DCMs to implement oversight programs that actually work, and it will create a principled framework that distinguishes legitimate prediction market activity from abuse.

The data presented in this letter, $253 million in wash trading in a single market, coordinated insider trading networks spanning four blockchains, profit concentration approaching 50% in fewer than 1% of accounts, demonstrates that these are not hypothetical risks. They are present, active, and detectable when the right surveillance infrastructure and framework is in place.

Solidus Labs deeply appreciates the Commission's commitment to innovation and industry-regulator dialogue, and the opportunity to provide comments as part of this rulemaking process. We would welcome further engagement on any of the issues raised in this letter, and stand ready to provide additional documentation, data, research or anything else that might be helpful to the Commission's rulemaking effort.

Respectfully submitted,

Chen Arad
Co-Founder and Chief External Affairs Officer
Solidus Labs