

July 20, 2026

Christopher Kirkpatrick

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

Re: Comment on Notice of Proposed Rulemaking, *Prediction Markets; Public Interest Determinations* (RIN 3038-AF65; 91 Fed. Reg. 35806, June 12, 2026)

Dear Commissioners,

I submit this comment in my personal capacity as an academic researcher who has studied prediction markets, forecasting, and information aggregation. The views expressed here are my own and not of any employer or other professional affiliation.

I am writing in response to the Commission's Notice of Proposed Rulemaking (NPRM) on prediction markets, which proposes amendments to the "Special Rule" from 17 CFR §40.11 governing when event contracts may be found contrary to the public interest. I previously submitted a comment on the Commission's Advance Notice of Proposed Rulemaking issued under this same RIN (3038-AF65). This letter builds directly on that submission and I [linked it here](#) for reference.

My earlier comment made three related points: (1) prediction markets are not a homogeneous category, and the Commission should codify a distinct research and information pathway separate from the large commercial pathway; (2) informational value varies substantially across categories of events, with professional and college sports representing a clear low-information limiting case; and (3) regulation of manipulation and insider-trading risk should focus on incentive-compatible market design rather than on categories of events. I am writing now to respond to the specific text of the proposed rule and to supply new empirical evidence that bears directly on it.

The proposed rule does not yet address these three points, and in several respects moves in the opposite direction. I therefore respectfully resubmit and sharpen them here in response to the specific text of the NPRM.

1. The proposed rule adopts a category-based framework where a design- and purpose-based framework is warranted.

The proposed rule sorts event contracts by the subject matter of the underlying event: activity unlawful under federal or state law (terrorism, assassination, or war) and "gaming" (proposed §40.11(a)(2), (b)), and applies the public-interest factors of proposed

§§40.11(a)(5) and (a)(6) on that basis. As I explained in my prior comment, the relevant risks the Commission identifies, manipulation and insider trading, are not intrinsic to any category of event. They are properties of a market's design and of its operator's incentives: position limits, subsidy and liquidity parameters, contract payout structures, participation rules, and whether the operator is optimizing for transaction volume and fee revenue or for research and public-interest objectives. Two contracts on the same underlying event can differ enormously in their susceptibility to manipulation depending on these design choices. A framework keyed solely to event type cannot distinguish between them, and will predictably misclassify robust, incentive-compatible markets while clearing manipulable ones that happen to fall outside an enumerated category.

For example, a small information-oriented market with verified participants, bounded stakes, and compensation or scoring rules designed around forecast accuracy presents very different incentives from an open, high-stakes exchange that takes all comers and monetizes transaction volume. The small information-oriented market can ensure an insider has neither the incentive nor the opportunity to trade on private knowledge of the result. An open exchange that takes all comers, by contrast, can give an insider a reason to trade on advance knowledge of the outcome, or even give someone with influence over the event a reason to alter the outcome itself, once the stakes are large enough. Put precisely: manipulation and insider trading occur when, accounting for risk tolerance, the expected benefits of cheating exceed the expected costs of being caught. An exchange holds many levers over the expected benefits of cheating (position and payout limits, participation rules), and, absent measures such as verified identity, comparatively few over the expected costs of being caught.

I therefore renew my recommendation that the Commission, within the public-interest analysis of proposed §§40.11(a)(5) and (a)(6), evaluate whether a contract's design and governance make the expected returns to manipulation or exploitation of private information lower than their expected costs, and treat that showing, rather than the event's category, as the operative question. This design-and-purpose orientation is also the framework least likely to entrench incumbents: both category-based federal rules and state-by-state gambling regimes can operate to protect the rents of established operators, whereas an incentive-compatibility standard evaluates a market's design rather than its operator's identity, size, or location so long as the required showing is a published, low-cost checklist a small or new entrant can satisfy, rather than a bespoke proceeding only a well-capitalized incumbent can afford. A category-based rule treats unlike markets as though they were alike, while a design-based rule treats alike risks alike regardless of category.

To be concrete without being prescriptive: the Commission has a menu of design and governance tools it could combine to make this operational. For example: published per-participant position limits and per-contract aggregate or open-interest caps, calibrated

more tightly where the underlying event is cheaply influenced and more loosely where it is robust; a verified, one-person-one-account identity requirement, which is the precondition that makes any per-participant limit enforceable, permits detection of insider trading, self-trading, and collusion, and allows interoperability with existing participant-protection regimes (for instance, honoring the problem-gambling self-exclusion registries that no pseudonymous market can enforce today); and alignment of the operator's revenue model with informational accuracy rather than transaction volume, whether through nonprofit operation or a business that predominantly monetizes data and forecasts rather than trading fees. I do not urge any one of these as a stand-alone mandate, nor do I endorse particular numerical thresholds. My submission is narrower: that some combination of such safeguards, set at levels the Commission and operators agree are adequate, is what a design-and-purpose analysis should require and reward, rather than the category of the underlying event.

2. Prediction market prices increasingly shape outcomes, not merely forecast them, heightening the case for design- and purpose-based safeguards.

As prediction-market prices become embedded in news coverage and public debate, the incentive to influence a price (even at a financial loss, to sustain a favored narrative) grows. This is not a mere hypothetical: all the way back in the 2012 election, Rajiv Sethi and I documented a single trader who spent millions supporting one candidate's contract prices through Election Day, apparently valuing the narrative above the expected trading loss. The risk is most acute for contracts on small, frivolous, or easily influenced outcomes, for example, whether a public figure will utter a particular phrase, where the cost of moving both the price and the underlying event is low. These are incentive-and-design problems, and they reinforce why the Commission's analysis should turn on bounded returns to manipulation rather than on event category alone.

3. New empirical evidence: sports prediction markets provide little measurable incremental price-discovery value beyond simple public-data models.

Informational value is a separate axis, and, unlike manipulation risk, it does vary systematically across domains. The proposed rule's treatment of "gaming", defined in proposed §40.11(b) to reach essentially any rule-governed activity engaged in for recreation or entertainment whose outcome depends on luck, skill, or athletic ability, and thus to encompass all professional and college sports, rests on an unstated premise: that these contracts function as the kind of price-discovery and information-aggregation mechanisms the Commission's framework was built to accommodate. My ongoing research provides evidence that, in the tested samples, pre-game sports prediction-market prices provide little or no measurable incremental forecasting value relative to deliberately simple models based only on public data.

Using only publicly available data, a deliberately minimal Elo rating model with a margin-of-victory adjustment (with parameters drawn from the published literature and trained only on prior seasons, not tuned on the test data) matches or outperforms Polymarket's pre-game closing prices across major U.S. sports leagues on standard proper scoring rules (Brier score and log-loss) and on calibration:

- In MLB (2025 regular season, $N = 557$ games), the simple model outperforms the market on both Brier score (0.2356 vs. 0.2389) and log-loss, with no statistically significant advantage for the market.
- In the NFL (2024, $N = 272$) and NHL (2024–25, $N = 491$), the model and the market are statistical ties, with no significant difference in either direction.
- In the NBA, the one league where the market had a small but statistically significant advantage over the base model, the entire gap is attributable to publicly available player-availability information (mandatory injury reports and rest patterns) that the minimal model deliberately ignores but could mechanically incorporate.

Prices on major sports prediction-market platforms are highly correlated for matched contracts, suggesting these results are unlikely to be platform-specific. The pattern is exactly what information economics predicts: sports outcomes are driven by abundant, structured, publicly available signals—team quality, home advantage, pitching matchups, rest, injuries—that a simple model incorporates directly. There is little dispersed private information for a market to aggregate, so market prices reflect the same public information plus distortions from recreational-bettor bias, spreads, and entertainment demand. This stands in sharp contrast to political and geopolitical forecasting, where information is genuinely dispersed and markets demonstrably add value.

These results do not imply that sports contracts have no private value to participants. They imply that, on the Commission's own price-discovery and information-aggregation dimension, their public informational contribution appears limited in the tested settings.¹

The regulatory implication is direct. The Commission's proposed "Price Discovery and Information Aggregation Utility" factor (proposed §40.11(a)(5)) is the correct lens—but if it is applied rigorously, it counsels that the sports contracts now accounting for a substantial volume carry negligible informational value², while the low-volume, high-information markets (elections, policy, scientific and technological milestones) are precisely where public value is concentrated. A categorical "gaming" definition inverts this: it sweeps in the low-information, high-volume contracts under a single rubric while

¹ I did not have the bandwidth to run the same test on within-game predictions, but theory would dictate the outcome would be the same (or worse) for prediction markets relative to statistical models. Traders would be more susceptible than models to biases associated with high information flow.

² "Sports has made up 80% of total trading volume on Kalshi and 39% on Polymarket since July 2024." <https://www.pewresearch.org/short-reads/2026/05/27/trading-volume-on-prediction-markets-has-soared-in-recent-months/>

doing nothing to protect the high-information, low-volume research markets that most merit accommodation. Greater informational and societal value does not eliminate manipulation concerns; rather, it creates a genuine public benefit that must be weighed against those concerns.

4. Where markets provide limited informational value, participant-protection concerns become more important.

This finding also bears directly on the Commission’s “protection of market participants and the public” obligation under CEA section 15(a). If sports contracts aggregate little or no dispersed private information, then persistent profits in such markets are less likely to reflect socially valuable information production and more likely to reflect speed, capitalization, access, or sophistication. Thus, a meaningful concern is a regressive transfer from a large population of recreational, lower-information participants to a small number of well-resourced, well-informed operators and traders: mirroring the concentration of returns already observed in adjacent gambling markets. I offer this as a hypothesis consistent with the calibration evidence above and with the favorite-longshot and recreational-bias literature, rather than as a distributional result I have measured directly; the Commission is well positioned to examine trader-level profit-and-loss data to test it. Either way, a public-interest analysis that considers only whether an event category is permissible, and not who bears the losses in markets that produce no offsetting informational benefit, is incomplete.

5. The proposal leaves nonprofit or smaller information-oriented markets exactly where my prior comment warned they should not remain.

The proposed rule addresses academic and research markets only in its historical background—noting past discretionary staff no-action relief for the Iowa Electronic Markets (1992), Victoria University of Wellington (2014), and PredictIt. It does not codify a distinct pathway for nonprofit or smaller information-oriented prediction markets (like Hypermind and Manifold). The result is precisely the structural asymmetry I described: large commercial pathways receive clarity and scalability through an explicit factor-based framework, while nonprofit and smaller information-oriented markets remain dependent on discretionary relief: uncertain, non-scalable, and administratively constrained. The broad “gaming” definition may make matters worse, as it contains no carve-out for nonprofit or smaller information-oriented markets and could, on its face, impact these markets on competitive or rule-governed events.

I therefore urge the Commission to use this rulemaking to codify a distinct, lightweight regulatory pathway for bona fide nonprofit or smaller information-oriented prediction markets, defined by some tractable subset that may include their nonprofit operation, research or educational purpose, verified-participant identity, bounded position and

participation limits, and public data-sharing, consistent with the Commission's longstanding, if ad hoc, accommodation of such markets. Formalizing what has until now been discretionary relief would provide the certainty these public-interest markets need to survive and scale.

This would not exempt such markets from oversight. It would instead create a transparent path for markets whose design, scale, identity controls, and public-interest purpose make them materially different from high-volume commercial betting-like platforms.

Conclusion and Recommendations

Prediction markets are not monolithic. Their social value depends on design (incentive-compatible vs. manipulable), domain (low- vs. high-informational events), and purpose (large commercial vs. smaller research and information). The proposed rule's category-based architecture flattens these distinctions at exactly the moment the evidence counsels drawing them more sharply. *A category-based rule treats unlike markets as though they were alike, while a design- and purpose-based rule treats alike risks alike regardless of category.* Thus, I respectfully encourage the Commission to:

1. Center the public-interest analysis of proposed §§40.11(a)(5) and (a)(6) on incentive-compatible market design, requiring operators to demonstrate that manipulation and exploitation of private information are bounded through some combination of verified identity, position and open-interest limits, contract-design safeguards, surveillance, and revenue-model alignment, without mandating a single one-size-fits-all design.
2. Apply heightened scrutiny to contracts whose underlying outcomes are cheap to influence, including contracts where the same actor can plausibly move both the market price and the real-world event.
3. Apply the proposed "Price Discovery and Information Aggregation Utility" factor rigorously, recognizing that sports contracts may provide little measurable incremental forecasting value relative to simple public-data models, while elections, policy, scientific, and technological milestones may provide much greater public informational value.
4. Incorporate participant-protection and distributional evidence into the public-interest analysis, including trader-level profit-and-loss data where available, especially for markets that appear to generate limited informational benefit.
5. Codify, or otherwise formalize through clear rule text or durable guidance, a distinct lightweight pathway for bona fide nonprofit, academic, educational, and smaller information-oriented prediction markets.

Markets are tools for aggregating dispersed knowledge, but they are also social, political, and economic institutions that create incentives of their own. A framework that preserves

space for experimentation and public-interest applications—and that regulates on the basis of design and purpose rather than event category—will ensure that prediction markets continue to serve their informational function where they genuinely can. I would be honored if the Commission’s staff wished to talk further on the data or other analysis mentioned in this letter.

Sincerely,

David M. Rothschild, PhD