March 16, 2026

Via Electronic Submission (https://comments.cftc.gov)

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)

Dear Chairman Selig and Members of the Commission:

I respectfully submit this comment in response to the Commission's Advance Notice of Proposed Rulemaking regarding event contract derivatives traded on prediction markets.[1] I write as a legal scholar and AI governance practitioner.[2] My current research examines the intersection of artificial intelligence, market-manipulation doctrine, and the regulatory frameworks governing prediction markets and autonomous trading systems. I am the author of a forthcoming law review article proposing an "upstream scienter" framework for AI manipulation liability under the Commodity Exchange Act ("CEA") and the Securities Exchange Act of 1934,[3] as well as a second article, accepted for publication by the International Journal of Law and Emerging Technologies, examining AI-generated evidence in criminal proceedings and the structural inadequacy of existing evidentiary frameworks to address autonomous system outputs.[4]

My professional background informs these comments. I currently serve as a lawyer at a Series C-stage AI lab, where I am embedded in model training pipelines and reinforcement learning from human feedback ("RLHF") evaluation frameworks, the same class of systems that, when applied to trading, create the governance challenges this letter addresses. I have previously served as general counsel at four technology companies (two of which resulted in successful exits), practiced securities law, and held legal leadership roles at a wide variety of academic organizations, including Johns Hopkins University.

My comments address a structural challenge that the ANPR acknowledges but has not yet fully resolved: the rapid integration of AI-driven trading systems into prediction markets, and the inadequacy of existing scienter and manipulation frameworks to address harmful outcomes

---

[1] Prediction Markets, 91 Fed. Reg. ____ (Mar. 16, 2026) (advance notice of proposed rulemaking) (RIN 3038-AF65) [hereinafter ANPR].

[2] The views expressed in this comment are solely those of the author and do not represent the views of any employer, client, or organization.

[3] Stevie Michelle Cline, The Intent Gap: AI Manipulation Doctrine, Prediction Markets, and the Case for Upstream Scienter (2026) (manuscript under review).

[4] Stevie Michelle Cline, *United States v. Heppner* and the AI Privilege Gap: Criminal Procedure in the Age of Machine-Generated Evidence (forthcoming 2026) (accepted for publication, International Journal of Law and Emerging Technologies).

arising from autonomous system behavior rather than individual human intent. I focus on Questions 2(c), 2(d), 7, 11, 29, 30, and 31, where the AI manipulation problem is most acute, though the framework I propose has implications across many of the Commission's inquiries.

I. The AI Manipulation Gap in Prediction Markets

The Commission's ANPR arrives at a pivotal moment. The explosion of event contracts, from approximately five per year before 2021 to roughly 1,600 certified in 2025 alone[5], has coincided with equally dramatic advances in AI-driven trading systems. Reinforcement-learning agents, large language models deployed for market analysis, and multi-agent coordination systems are already operating in derivative markets. The CFTC Chairman has himself identified the challenge.[6] As prediction markets grow in volume and diversity, they will inevitably attract increasingly sophisticated autonomous trading strategies.

This creates what I term the "intent gap": a structural mismatch between the Commission's existing enforcement tools, which require proof that a specific actor intended to manipulate market conditions, and the operational reality of AI systems that can execute functionally manipulative strategies without any individual human actor possessing the requisite mental state at the moment of the harmful act.

The intent gap is not hypothetical. Agent-based modeling research has demonstrated that spoofing strategies are profitable and effective against learning agents in simulated limit-order-book environments, and that markets populated by agents who learn from order-book information are inherently vulnerable to manipulation by spoofing agents.[7] Separately, researchers have used reinforcement-learning frameworks to train agents to execute manipulative trading strategies, demonstrating that such strategies can be computationally discovered rather than manually designed.[8] The broader multi-agent literature further establishes that independently operating agents in competitive market environments regularly produce collective dynamics—including price distortions—that are functionally indistinguishable from coordinated behavior, even in the absence of any coordination mechanism.[9] Taken together, these findings establish that AI systems operating in prediction markets will face strong optimization incentives to discover and exploit manipulative strategies, and that the resulting market impact may arise

[5] See ANPR at 5 (noting that DCMs certified approximately 1,600 event contracts in 2025, up from an average of approximately five per year from 2006 to 2020).

[6] Remarks of Chairman Michael S. Selig, The Next Era of American Markets Leadership, FIA Global Cleared Markets Conference (Mar. 9, 2026).

[7] Xintong Wang et al., Spoofing the Limit Order Book: A Strategic Agent-Based Analysis, 12 Games, no. 2, art. 46, at 1 (2021) (demonstrating through agent-based modeling that spoofing strategies are profitable and effective against learning agents in simulated limit-order-book environments, and that markets populated by agents who learn from order book information are inherently vulnerable to manipulation).

[8] Megan Shearer, Gabriel Rauterberg & Michael P. Wellman, Learning to Manipulate a Financial Benchmark, in Proc. 4th ACM Int'l Conf. on AI in Fin. 592 (2023) (training reinforcement-learning agents that successfully learn to execute manipulative trading strategies, demonstrating that such strategies can be computationally discovered rather than manually designed).

[9] See Rama Cont & Jean-Philippe Bouchaud, Herd Behavior and Aggregate Fluctuations in Financial Markets, 4 Macroeconomic Dynamics 170 (2000) (demonstrating that independently acting agents in competitive market environments produce collective dynamics—including price distortions—that are functionally indistinguishable from coordinated behavior, even absent any coordination mechanism).

from the interactions among multiple autonomous agents rather than from any single actor's design.

In traditional commodity and securities markets, enforcement actions for manipulation under the CEA require the Commission to establish scienter—that the respondent acted with specific intent to create an artificial price or to employ a manipulative device.[10] Existing precedents establishing the evidentiary framework for spoofing and manipulation all rely on evidence of human intent derived from communications, trading patterns, and testimony.[11] When an AI trading agent learns through reinforcement that an aggressive order-placement-and-cancellation strategy maximizes its reward function, or when coordinated autonomous agents produce price distortions through emergent multi-agent dynamics, that evidentiary framework collapses. No single human actor designed the specific manipulative strategy. No contemporaneous communications evidence manipulative intent. The system's developers may not have foreseen the specific harmful outcome, yet the design choices they made, the objective function, the training environment, and the deployment without adequate guardrails made it foreseeable.

The Commission's existing doctrinal tools do not resolve this problem. Three theories that might appear applicable each fail in the AI context for distinct reasons.

First, willful blindness, the doctrine that a person who deliberately avoids confirming a high probability of wrongdoing may be treated as having acted knowingly, requires that the defendant subjectively believed the wrongful conduct was substantially certain to occur and took deliberate steps to avoid confirming it.[12] An engineer selecting an objective function for a trading system may genuinely not anticipate that the system will converge on a manipulative strategy; the design choice reflects optimization goals, not a conscious decision to avoid knowledge of future manipulative behavior. Willful blindness was developed for human actors who choose not to open a box they suspect contains contraband—not for designers whose system autonomously generates conduct that was never contemplated.

Second, respondeat superior imposes vicarious liability on an employer for the acts of an employee within the scope of employment. But an AI trading agent is not an employee. It has no employment relationship, no scope of authority defined by the organizational hierarchy, and no capacity to act "within" or "outside" its role as respondeat superior requires. Treating AI agents as employees for liability purposes would require a conceptual extension that existing case law does not support, and that would generate unpredictable consequences across every area of law that relies on the employer-employee distinction.

---

[10] See 7 U.S.C. § 13(a)(2) (prohibiting manipulation of commodity prices); 7 U.S.C. § 9(1) (prohibiting use of "any manipulative or deceptive device or contrivance" in connection with swaps or futures contracts); 17 C.F.R. § 180.1 (implementing anti-manipulation provisions).

[11] *See, e.g.*, *CFTC v. Oystacher*, 203 F. Supp. 3d 934 (N.D. Ill. 2016) (spoofing); *In re Panther Energy Trading LLC*, CFTC Docket No. 13-26 (July 22, 2013) (same). Both cases relied on evidence of human intent derived from communications and trading patterns—evidence that would not exist when the trading strategy is autonomously discovered by a reinforcement-learning agent.

[12] *Global-Tech Appliances, Inc. v. SEB S.A.*, 563 U.S. 754, 769 (2011) (establishing that willful blindness requires (1) a subjective belief that there is a high probability that a fact exists and (2) deliberate actions to avoid learning of that fact).

Third, control person liability under 7 U.S.C. § 13c(b) (CEA Section 13(b)) provides that any person who "directly or indirectly controls" a violator may be held jointly liable, but only if the control person "did not act in good faith or knowingly induced, directly or indirectly, the act or acts constituting the violation."[13] This "knowingly induced" standard reintroduces the very scienter requirement that makes the underlying manipulation charge untenable. If the Commission cannot establish that the AI system's trading constituted intentional manipulation, because no human intended the specific manipulative act, it cannot establish that the system's deployer "knowingly induced" a violation that, under current doctrine, may not have occurred.

The upstream scienter framework proposed in Part II.C below addresses this gap by relocating the intent inquiry to the points where human culpability can actually be established, without requiring doctrinal contortions that existing theories do not support.

II. Responses to Specific Questions

A. Questions 2(c) and 2(d): Manipulation Susceptibility and Market Surveillance

Question 2(c) asks how the Commission should determine whether an event contract is "readily susceptible to manipulation" under DCM Core Principle 3. Question 2(d) asks whether any aspects of prediction markets pose challenges to compliance with Core Principle 4's market surveillance requirements.

I urge the Commission to recognize that the inquiry into susceptibility to manipulation must account not only for the characteristics of the contract itself but also for those of the trading systems that will interact with it. An event contract that is not readily susceptible to manipulation by human traders may become highly susceptible when traded by AI systems capable of executing strategies at speeds, volumes, and levels of coordination that human traders cannot match.

The Commission should consider the following factors:

Thin liquidity vulnerability. Many event contracts, particularly those referencing niche political, scientific, or cultural events, have relatively thin order books. AI systems can exploit thin liquidity to move prices with minimal capital, creating distorted probability signals that may influence real-world decision-making or related markets.

Cross-market coordination. AI agents operating across multiple prediction market platforms, or across prediction markets and traditional derivative markets, can execute coordinated strategies that are difficult to detect through single-market surveillance.[14] Existing surveillance frameworks are not designed to detect manipulation executed by autonomous agents operating across multiple venues simultaneously.

Emergent manipulation. Multi-agent systems can produce price distortions through emergent dynamics that no individual agent was programmed to create. When multiple AI trading agents independently learn that similar aggressive strategies maximize their respective reward

[13] 7 U.S.C. § 13c(b).

[14] See ANPR, Question 30 (asking about events under the control of a single individual or small group and "particular challenges related to cross-market manipulation").

functions, the resulting market impact may be functionally indistinguishable from coordinated manipulation—yet no single actor intended the collective outcome. This is a phenomenon well documented in the multi-agent systems literature, and it poses a distinctive challenge for any surveillance regime built on the assumption that coordinated manipulation requires coordinated intent.

The Commission's own contemporaneous guidance confirms this gap. Staff Advisory Letter No. 26-08, issued alongside the ANPR, evaluates manipulation susceptibility primarily by reference to whether "the breadth of the outcome . . . reduces the ability of any single actor to manipulate the settlement value without material cost or substantial risk of detection."[15] This framing, focused on the capabilities of individual human actors, is precisely the analytical framework that AI-driven trading renders inadequate. A single autonomous agent may not be able to manipulate a broad-outcome contract. But a population of independently operating AI agents, each trained on similar reinforcement-learning architectures and converging on similar aggressive strategies, can collectively produce the same price distortion that the Staff Advisory assumes requires a coordinated single-actor effort. The Staff Advisory's framework should be updated to account for the distinctive dynamics of manipulation that emerge when the relevant "actors" are autonomous systems whose behavior is shaped by optimization rather than intention.

For DCM surveillance obligations under Core Principle 4, I recommend that the Commission require prediction markets to implement AI-specific surveillance capabilities, including: (i) detection algorithms for patterns indicative of autonomous agent activity, such as microsecond-level order patterns and inhuman consistency in execution timing; (ii) cross-platform data-sharing agreements to facilitate detection of multi-venue manipulation by coordinated autonomous systems; and (iii) mandatory reporting by market participants deploying AI-driven trading strategies, including disclosure of objective function design and training methodology where the trading system exceeds defined volume or frequency thresholds.

B. Questions 7 and 11: The Public Interest Determination and AI-Mediated Harms

Question 7 asks what factors the Commission should consider in making a public interest determination under 7 U.S.C. § 7a-2(c)(5)(C) (CEA Section 5c(c)(5)(C)). Question 11 asks how the Commission should fulfill the CEA's statutory purposes—including preventing price manipulation, protecting market participants from abusive practices, and promoting responsible innovation, regarding prediction markets.[16]

The Commission's public-interest determination should include an assessment of the systemic risks that AI-mediated trading poses to the core function of prediction markets. The primary public interest justification for prediction markets is their capacity to aggregate dispersed private information into probability estimates that serve as public goods.[17] This information aggregation function is degraded and may be affirmatively corrupted when AI systems can distort prices through strategies that do not reflect genuine private information about the likelihood of

[15] CFTC Staff Letter No. 26-08, Prediction Markets Advisory, at 5 (Mar. 12, 2026).

[16] 7 U.S.C. § 5(b).

[17] 7 U.S.C. § 5(a) (describing transactions subject to the CEA as "affected with a national public interest by providing a means for managing and assuming price risks, discovering prices, or disseminating pricing information through trading in liquid, fair and financially secure trading facilities").

outcomes. A prediction market that produces distorted probability signals is not merely failing as a market; it is generating misinformation that may influence democratic processes, corporate decision-making, and public understanding of consequential events.

CEA Section 3(b)'s instruction to "promote responsible innovation" does not mandate permissiveness toward AI trading systems in the absence of adequate governance frameworks. Responsible innovation requires that the legal infrastructure for accountability keep pace with the trading technology deployed in these markets. At present, it has not. The Commission should make clear in its forthcoming rulemaking that DCMs bear an affirmative obligation to ensure that AI trading systems operating on their platforms are subject to governance standards sufficient to facilitate post-hoc accountability for manipulative outcomes. This obligation should be understood as a necessary component of responsible innovation, not an impediment to it.

C. Questions 29, 30, and 31: Inside Information, Asymmetric Information, and the Intent Problem

These three questions collectively address the Commission's most consequential regulatory challenge regarding AI in prediction markets. I urge the Commission to use the forthcoming rulemaking as an opportunity to develop a framework for AI-specific manipulation liability that I call "upstream scienter."

The core principle is straightforward: when AI systems produce manipulative outcomes in prediction markets, accountability should be located at the points in the system's lifecycle where culpable human decisions were made—not deferred to the point of trade execution, where no human actor may have been meaningfully involved. The upstream scienter framework identifies three decision points where the manipulation inquiry should focus:

First, system design. The choice of objective function, reward structure, and optimization target determines what outcomes the system will pursue. A trading system optimized to maximize short-term profit, without constraints on market impact, is likely to discover and exploit manipulative strategies. The design choice to omit manipulation-avoidance constraints from the objective function is itself a decision that can bear culpability when manipulative outcomes result.

Second, training environment. The data and simulated environments in which a trading agent is trained shape its behavioral repertoire. A system trained in an environment that rewards aggressive order placement and cancellation patterns, or that fails to penalize strategies that would constitute manipulation in live markets, has been equipped with manipulative capabilities by the training choices its developers made.

Third, deployment governance. The decision to deploy an AI trading system on a live prediction market without adequate monitoring, kill switches, or behavioral constraints is itself a governance failure that should carry regulatory consequences when the system produces manipulative outcomes.[18]

---

[18] Cf. 17 C.F.R. § 1.11 (requiring risk management programs for futures commission merchants and swap dealers); CFTC Division of Swap Dealer and Intermediary Oversight, Advisories on Automated Trading (Nov. 2013; Nov. 2015) (addressing risk controls for automated trading systems).

This framework is compatible with 7 U.S.C. § 9(1)'s (CEA Section 6(c)(1)'s) existing prohibition on "any manipulative or deceptive device or contrivance."[19] The "device or contrivance" language is broad enough to encompass an AI system that is designed, trained, and deployed in a manner that foreseeably produces manipulative outcomes, even if no individual trade was executed with specific manipulative intent. The upstream scienter framework provides the doctrinal bridge between the statute's existing language and the operational realities of autonomous trading.

With respect to Question 31's observation that 7 U.S.C. § 9(1) does not require disclosure of nonpublic material information, the Commission should consider whether this limitation, originally designed for human traders operating at human speeds, remains appropriate when AI systems can process and act on informational advantages within milliseconds. An AI agent that consumes and trades on nonpublic information before human traders can react functions less as a participant in price discovery and more as a mechanism for information extraction—a distinction with significant implications for the CEA's purposes of "discovering prices, or disseminating pricing information."[20]

III. Proposed Safe Harbor for AI Governance Compliance

To incentivize responsible innovation, I recommend that the Commission consider incorporating safe harbor provisions into its forthcoming rulemaking for prediction market participants that implement verifiable AI governance protocols. Such a safe harbor would provide that a market participant or DCM satisfying specified AI governance standards shall receive a rebuttable presumption against scienter in enforcement actions arising from AI-generated trading activity on prediction markets.

The safe harbor should be structured as a graduated framework calibrated to the degree of human oversight at each stage of the AI development and deployment lifecycle. This approach creates a positive incentive structure: rather than chilling AI innovation, it channels that innovation toward systems designed with accountability from the outset. The framework has precedent in both existing CFTC risk management requirements[21] and the EU AI Act's conformity assessment structure for high-risk systems.[22]

The safe harbor framework should include, at minimum:

(a) Design-stage documentation requirements: written records of objective function design choices, the rationale for including or excluding manipulation-avoidance constraints, and risk assessments addressing foreseeable manipulative strategies.

(b) Training-stage audit trails: documentation of training data sources, reward function calibration, and validation testing for manipulative behavior in simulated market environments.

---

[19] 7 U.S.C. § 9(1).

[20] 7 U.S.C. § 5(a).

[21] *See supra* note 18.

[22] See Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024, Laying Down Harmonised Rules on Artificial Intelligence, arts. 6–49, 2024 O.J. (L 1689) (establishing conformity assessment obligations for high-risk AI systems, including documentation and transparency requirements).

(c) Deployment-stage governance protocols: real-time monitoring of trading behavior against pre-defined behavioral envelopes, automated circuit breakers triggered by anomalous activity, and human-in-the-loop review procedures for flagged patterns.

(d) Post-deployment reporting obligations: periodic disclosure to the relevant DCM of material AI system modifications, performance metrics, and instances where monitoring systems flagged potential manipulative behavior.

The verification mechanism for the safe harbor should leverage the Commission's existing product submission infrastructure. DCMs should be responsible for initial verification of AI governance compliance as part of their self-certification process under 17 C.F.R. § 40.2,[23] consistent with their role as "front-line regulators" emphasized in Staff Advisory Letter No. 26-08. Specifically, when a DCM certifies a new event contract for listing, the certification should include a representation that the DCM has reviewed and maintains on file the AI governance documentation described in elements (a) through (d) above for any market participant deploying AI-driven trading systems that meet specified volume or frequency thresholds on that contract. The Commission should retain authority to audit this documentation and to designate recognized third-party AI governance auditors, analogous to the role of notified bodies under the EU AI Act's conformity assessment framework.[24] This layered approach, DCM verification, Commission oversight, and optional third-party audit, distributes compliance costs proportionally and avoids creating a centralized bottleneck that would delay product innovation.

The presumption would be rebuttable; compliance with the safe harbor would not immunize a participant from liability where affirmative evidence of manipulative intent exists. But it would provide the Commission with a structured approach to distinguishing between firms that deployed AI systems responsibly and those that did not, addressing the scienter problem without requiring attribution of mental states to machines.

IV. Cost-Benefit Considerations

The ANPR repeatedly asks commenters to address cost-benefit considerations.[25] I offer the following observations.

The costs of inaction compound. As AI trading systems become more prevalent in prediction markets, enforcement actions that fail for lack of provable scienter will erode market integrity and public confidence in the information aggregation function of prediction markets. Each failed enforcement action establishes a de facto safe harbor for AI-mediated manipulation, but one that rewards opacity rather than governance. The resulting loss of market integrity would reduce the social value of prediction markets and undermine the Commission's own rationale for asserting exclusive jurisdiction over them.

---

[23] 17 C.F.R. § 40.2 (self-certification procedures for new products).

[24] See Regulation (EU) 2024/1689, arts. 43–49 (conformity assessment procedures and notified bodies for high-risk AI systems).

[25] See 7 U.S.C. § 19(a)(2) (requiring the Commission to evaluate costs and benefits in light of five enumerated considerations, including "protection of market participants and the public," "efficiency, competitiveness, and financial integrity of futures markets," "price discovery," "sound risk management practices," and "other public interest considerations").

The costs of compliance are modest and largely incremental. Documentation of system design and training choices is already standard practice at well-run AI companies. Firms deploying AI trading systems in prediction markets are, by definition, technically sophisticated operations for which model documentation and monitoring infrastructure are not novel requirements. Real-time monitoring and circuit-breaker systems are standard risk management tools in traditional derivative markets. The incremental cost of adapting these existing practices to satisfy the proposed safe harbor requirements would be modest relative to the compliance infrastructure that DCMs and registered market participants already maintain.

The competitive implications favor U.S. leadership. The EU AI Act already imposes conformity assessment obligations on high-risk AI systems, including systems used in financial services.[26] If the United States establishes a clear, principled, and less prescriptive framework for AI governance in prediction markets, it will attract responsible innovators who seek regulatory certainty. If it fails to address the AI manipulation gap, it risks either ceding governance leadership to the EU or creating a regulatory environment in which the most sophisticated AI trading strategies are deployed without accountability.

The Commission's regulatory framework and its jurisdictional posture are interrelated. In litigation pending before multiple federal circuits, the Commission has argued that the CEA grants the CFTC exclusive jurisdiction over event contracts traded on designated contract markets, preempting state gaming enforcement.[27] The strength of that jurisdictional claim before the courts depends in part on the federal framework's capacity to address the full range of market integrity threats that prediction markets present, including threats that state gaming regulators cite as justification for their own enforcement actions. A federal regime that demonstrably accounts for AI-mediated manipulation risks strengthens the Commission's position that federal oversight is sufficient to protect market participants and the public. Conversely, a framework with a visible gap at its center invites judicial skepticism about whether exclusive federal jurisdiction adequately serves the public interest. The AI governance measures proposed in this letter thus serve not only market integrity but also the Commission's institutional interests in the jurisdictional contests that will shape prediction market regulation for years to come.

V. Conclusion

The Commission's ANPR represents a significant opportunity to build a regulatory framework for prediction markets designed for the markets of the future, not the markets of the past. The central challenge is not whether to permit prediction markets or how to classify event contracts—it is whether the legal infrastructure for market integrity can keep pace with the autonomous systems that will increasingly dominate trading on these platforms.

I urge the Commission to incorporate four principles into its forthcoming rulemaking:

---

[26] *See supra* note 22.

[27] *See, e.g.*, *N. Am. Derivatives Exch., Inc. v. Nevada*, No. 25-7187 (9th Cir.) (CFTC amicus brief filed Feb. 17, 2026, arguing that state gaming enforcement "invade[s]" the Commission's "exclusive jurisdiction over CFTC-regulated designated contract markets"); *KalshiEX LLC v. Hendrick*, No. 25-7516 (9th Cir.) (same jurisdictional dispute).

First, the manipulation susceptibility analysis under Core Principle 3 must account for the distinctive manipulation risks posed by AI-driven trading systems, including emergent manipulation by multi-agent systems and cross-platform coordination by autonomous agents.

Second, the public-interest determination under 7 U.S.C. § 7a-2(c)(5)(C) should incorporate the systemic risks to prediction markets' information-aggregation function posed by AI-mediated price distortion.

Third, the Commission should adopt an upstream scienter framework that locates accountability for AI-generated manipulative outcomes at the points of system design, training, and deployment—the points where culpable human decisions are actually made.

Fourth, the Commission should incentivize responsible AI innovation through a graduated safe harbor for market participants and DCMs that implement verifiable AI governance protocols.

A regulatory framework that can assign accountability for market manipulation only when a human being personally intends to distort prices is a framework designed for the last century. Prediction markets are already attracting autonomous systems that operate beyond the reach of that framework. The Commission now has an opportunity—through this rulemaking—to build an accountability infrastructure designed for the markets it will actually be regulating. The upstream scienter framework, the graduated safe harbor, and the AI-specific surveillance measures proposed in this letter offer a path toward that goal.

I appreciate the opportunity to contribute to this rulemaking and would welcome the opportunity to provide additional information or analysis as the Commission develops its proposed rules.

Respectfully submitted,

Stevie Michelle Cline, Esq.

Admitted: New York State Bar | U.S. Supreme Court | U.S. Tax Court