Tilted: Platform Exploitation, Cognitive Integrity, and the Missing Half of Sports Betting Law

Mike Sapp

California Western School of Law

v9 Working Draft — May 4, 2026

TABLE OF CONTENTS

INTRODUCTION

At 2:17 AM, Marcus's phone buzzes. He has been betting for three hours, chasing $1,800 in losses through increasingly desperate wagers. His DraftKings app knows everything: the escalating stakes, the loss-chasing pattern visible in 71% of his bets placed within thirty minutes of losing, the two previous "cooling off" periods he imposed on himself, the three deposits totaling $2,000 he has already made tonight. The platform's algorithms have flagged him.[1]

The notification reads: "EXCLUSIVE: 50% Profit Boost on Next Bet! One-Tap Claim."

What the platform calls marketing has a different name when the recipient is in crisis.[2] DraftKings could have triggered a responsible gambling intervention. Instead, it deployed a profit boost calibrated to Marcus's crisis state. The "one-tap" mechanism eliminates the friction that might have prompted reflection, and the timing targets the moment of maximum vulnerability.

Current gambling regulation has nothing to say about this. The market is functioning. The odds are accurate. No insider traded on nonpublic sports information. No game was fixed. By every metric that regulators monitor, the system worked exactly as designed.

That is the problem.

The legalization of sports betting after *Murphy v. NCAA* moved the casino from the floor to the phone.[3] Ninety-five percent of wagers are now placed on smartphones — within a regulatory framework designed for physical spaces with human oversight, geographic boundaries, and built-in friction. The smartphone eliminated every one. The casino lives in the bettor's pocket, staffed by algorithms rather than pit bosses.

Current regulation addresses market integrity: the fairness of betting markets themselves. But market integrity captures only one dimension. A fair market and an exploitative house can coexist: the betting market can function with perfect integrity while the platform systematically exploits vulnerabilities bettors cannot perceive and the platform can detect in real time. This second dimension — cognitive integrity — receives almost no regulatory attention.

---

[1] The Marcus illustrative case used throughout this Article is a composite drawn from the patterns documented in the *Sage v. DraftKings* complaint and from the public-record evidence the Article's empirical foundation rests on; the actual *Sage* plaintiffs supply the empirical foundation Marcus illustrates. Behavioral indicators attributed to Marcus (loss-chase percentages, stake escalation, late-night session concentration, multi-deposit nights) are calibrated to the patterns the *Sage* complaint alleges of named plaintiffs and to the platform-self-described detection categories documented in Part I.B.1.

[2] This Article uses "extraction" and "exploitation" in distinct senses. *Extraction* denotes the financial act — the transfer of money from bettor to platform. *Exploitation* denotes the moral and legal wrong — the use of asymmetric information to take advantage of vulnerability. The platform exploits the bettor's cognitive state in order to extract the bettor's money. Extraction without exploitation (a fair bet, honestly offered) is lawful; exploitation that enables extraction (crisis-state targeting of a vulnerable user) is the tort this Article develops.

[3] AM. GAMING ASS'N, Commercial Gaming Revenue Tracker (2025) (reporting $149.9 billion in sports betting handle for 2024 across 38 states and the District of Columbia, with approximately 95% of wagers placed via mobile apps).

The result is a regulatory gap. Market integrity protects the game. Nothing protects the player.

This Article develops a two-integrity paradigm for sports betting regulation. The framework for the missing half is the Algorithmic Toxicity Tort (ATT).[4] ATT names cognitive integrity — the interest in freedom from surreptitious, automated manipulation that substantially impairs autonomous choice — as the specific protected interest, and develops the Manipulation-Vulnerability Matrix as the administrable standard that enables courts to distinguish legitimate marketing from actionable exploitation.

The doctrinal innovation can be stated compactly. American addictive-product liability has progressed from tobacco through opioids to ultra-processed foods — products designed to produce continued return through engineered exposure to a harm-vector. Sports-betting platforms complete a particular movement within that tradition: the harm-vector has migrated inward. The trigger is digital — a notification, a microbet offer, a personalized prompt — but the dependent response it engineers is the same one the substance-addiction frameworks were built to address. The platform delivers the stimulus; the body provides the chemistry. Where the chemical lever — manufacturing standards, distribution restrictions, ingredient disclosure — does not exist, what remains regulable is the trigger architecture itself: the timing, targeting, and intensity of the stimulus that produces the dependent response. The Algorithmic Toxicity Tort is the doctrinal instrument designed for that target.

The public health movement has documented the harm with precision: Natasha Dow Schüll's ethnographic research identified the "machine zone" — the manufactured dissociative state that mobile platforms have reproduced in private, at scale, and without witnesses.[5] In parallel, the Public Health Advocacy Institute has systematized the litigation methodology from tobacco through food marketing to gambling — identifying exploitative design, documenting corporate knowledge, reframing from consumer choice to manufacturer accountability.[6] What has been missing is the conversion from empirical diagnosis and public health advocacy to a legal standard courts can apply. The ATT framework provides it.

The framework sidesteps whether compulsive betting constitutes addiction or habit. The legal question is what the platform did when it detected vulnerability. The platform's conduct is actionable not because it created addiction, but because it systematically suppressed the user's capacity to break what would otherwise be a manageable habit. When platforms possess detailed real-time knowledge of user

---

[4] While the ATT framework is developed here in the gambling context, its structure is designed for application wherever platforms deploy behavioral data to exploit detected vulnerability. Social media, fintech, and generative AI present analogous structural problems.

[5] NATASHA DOW SCHÜLL, ADDICTION BY DESIGN: MACHINE GAMBLING IN LAS VEGAS 2-12, 166-97 (2012); *see infra* Part I.B.3. The ethnographic tradition Schüll inaugurated for casino-floor gambling has been extended in the British, Australian, and Irish public-health and sociological traditions; the framework developed here adapts the tradition to the algorithmic moment in which the dissociative architecture has migrated from physical to temporal form.

[6] *See* Lissy C. Friedman et al., *Tobacco Industry Use of Personal Responsibility Rhetoric in Public Relations and Litigation: Disguising Freedom to Blame as Freedom of Choice*, 105 AM. J. PUB. HEALTH 250 (2015) (documenting PHAI's systematized litigation methodology); *see also infra* Part II.B.2 (developing the PHAI litigation genealogy from tobacco through gambling).

vulnerability — observable in their own behavioral telemetry as it unfolds — a duty arises not to exploit that knowledge.

This Article proceeds in six Parts. Part I documents the regulatory gap, locating the platform exploitation in the longer history of sport's capture by gambling commerce. Part II develops the ATT framework — cognitive integrity as the protected interest, the duty arising from information asymmetry and counterparty structure, and the Matrix as the rule of decision. Part III dissects the platform business model, applies the framework to a worked specimen case, and supplies the litigator's path through the framework's pleading, burden-shift, and discovery-stage architecture. Part IV answers the doctrinal and normative objections — § 230, the First Amendment, preemption, and arbitration. Part V extends the analysis to league secondary liability. Part VI situates the United States doctrinal frame within the comparative and global context. Three Appendices follow: a Proposed Restatement formalizing the cause of action; the Matrix Reference Tables that operationalize the breach standard; and an Evidentiary Inventory mapping the framework's claims to public records, discovery targets, and expert-testimony categories.

As of this drafting, the framework is no longer prospective: in March 2026, the Public Health Advocacy Institute filed *Sage v. DraftKings* in Philadelphia, naming DraftKings, FanDuel, the NFL, and Genius Sports as defendants. The argument that follows treats *Sage* as an analytical anchor and a live test of the framework's doctrinal reach.

PART I: THE REGULATORY GAP

Sport and gambling have not always been the same product. Through most of the twentieth century the betting shop was a separate place: a discrete physical and cultural institution that the bettor entered and, eventually, left. Sport happened on television, in stadiums, in the company of others. Gambling happened in a different room, sometimes a different building, and almost always in a register the bettor recognized as marked. The integration of the two was engineered, recently and on purpose.[7] The wager is no longer placed on the game; the game has been redesigned to produce wagers.[8]

---

[7] The integration of sport and gambling commerce in its current form is largely a post-2018 development in the United States and a post-2005 development in the United Kingdom. The Gambling Act 2005 (UK) liberalized broadcast advertising restrictions and opened the regulatory pathway for the integration of sportsbook commerce into live broadcasting; *Murphy v. NCAA*, 584 U.S. 453 (2018), accomplished the analogous shift in the United States by striking down the Professional and Amateur Sports Protection Act. The intervening industry consolidation — the Flutter / Stars Group merger (2020), the Caesars / William Hill acquisition (2021), the DraftKings public listing (2020), the NFL's exclusive data partnership with Genius Sports — created the corporate scale at which the integration could be deployed at production scale across multiple jurisdictions simultaneously. The sport-gambling integration is therefore not a long-running phenomenon but a regulatory and commercial development of the past two decades, and a US development of the past seven years.

[8] The structural inversion is observable in the platform's product design and in the league's revenue model. Microbetting — the unit of consumption that opens before the play and resolves with it — requires statistical infrastructure that did not exist before sportsbook demand created it; the live-data licensing arrangements between leagues and sportsbook data suppliers, in which leagues

The cognitive harm this Article addresses cannot be understood apart from that integration. Sport is the affective vehicle through which the platform reaches the bettor; the team allegiance, the player loyalty, the ritual of watching are the cultural substrate that makes the algorithmic targeting effective. The slot machine had to manufacture its own attractiveness through engineered reinforcement. The mobile sportsbook borrows the affective infrastructure that fandom built. The British and Australian gambling-sociology traditions converge on this point — Schüll's ethnography of casino architecture, Reith's analysis of the cultural conditions of compulsive consumption, Cassidy's political economy of regulatory capture, Livingstone's public-health critique of the responsible-gambling paradigm: the harm sits in the cultural architecture that delivers the gamble to the gambler.[9] What this Article supplies is the doctrinal completion of a diagnosis the public-health-sociological tradition has been building for two decades.

A regulatory architecture designed for casinos governs a smartphone industry.

A. MARKET INTEGRITY: WHAT CURRENT LAW PROTECTS

Market integrity concerns the fairness of betting markets themselves: accurate odds, honest settlement of bets, prevention of insider trading and match-fixing. Current regulation addresses these concerns comprehensively — leagues police material nonpublic sports information, sportsbooks monitor for suspicious wagering patterns, and federal criminal enforcement backstops the system.[10] The framework is uneven, but it exists and violations are caught. Market integrity is the success story of sports betting regulation.

Cognitive integrity is the chapter the law has not yet read.

B. COGNITIVE INTEGRITY: WHAT CURRENT LAW IGNORES

The regulatory architecture described above protects the game. It has nothing to say about how platforms exploit cognitive vulnerabilities no existing framework reaches.

---

hold equity in the data supplier and earn higher commission rates on in-play wagers than on pre-match wagers, are post-2020 phenomena. *See infra* Part V (developing the league liability theory that arises from this inversion).

[9] GERDA REITH, THE AGE OF CHANCE: GAMBLING IN WESTERN CULTURE (Routledge 1999) (developing the cultural-sociological framework for understanding gambling as a practice constituted by the social and economic conditions of late modernity rather than as an individual pathology); REBECCA CASSIDY, VICIOUS GAMES: CAPITALISM AND GAMBLING (Pluto Press 2020) (analyzing the political economy of the gambling industry's regulatory capture and the systematic failure of "responsible gambling" frameworks); *see also* Charles Livingstone & Angela Rintoul, *Moving on from Responsible Gambling: A New Discourse Is Needed to Prevent and Minimise Harm from Gambling*, 184 PUB. HEALTH 107 (2020). The British, Australian, and Irish sociological traditions — Schüll's American ethnography of casino architecture excepted only by jurisdiction — have collectively built the empirical and theoretical foundation that the framework developed here translates into doctrine.

[10] *See, e.g.*, NFL Gambling Policy (2024); NBA Lifetime Ban of Jontay Porter for prop-bet manipulation (2024); MLB Lifetime Ban of Tucupita Marcano (2024); *United States v. Clase de la Cruz*, No. 1:25-cr-00346 (E.D.N.Y. 2025); Press Release, U.S. ATT'Y'S OFF., E.D. Pa. (Jan. 15, 2026) (charging twenty-six defendants in the largest college basketball point-shaving conspiracy since 1951).

1. Platform Knowledge Asymmetry

Betting platforms know their users with a precision that would unsettle those users if they understood it. Public reporting and complaint allegations describe DraftKings as collecting 186+ data attributes on each customer: betting patterns, loss responses, session timing, deposit behavior, promotional responsiveness, and calculated "lifetime value."[11] DraftKings' own securities disclosures separately confirm an individualized customer-analytics infrastructure capable of customer-level betting-limit management.[12] The collected data is Zuboff's "behavioral surplus" — information generated far beyond what service delivery requires.[13] What the platform needs to process a wager is one dataset. What it mines to predict and modify the bettor's behavior is something else entirely.

The surplus enables predictive modeling of vulnerability. The DraftKings CEO has confirmed the targeting himself: "There's a lot of data science and customization at a player level to serve up [Same Game Parlays] to the customers that we think have proclivity to engage with that type of bet."[14] "Proclivity to engage" is the corporate translation of vulnerability to gambling harm.

The surveillance is patented methodology. Entain's U.S. Patent No. 12,555,436 B2 describes a tiered vulnerability-detection system that classifies users by behavioral signals — "amounts greater than usual," "longer time than usual" — and prescribes calibrated commercial responses at each tier.[15] The patent does not prove extraction intent. It proves capability. Real-time risk detection, tiering, and intervention are technically feasible on the same behavioral substrate platforms already collect for targeting purposes. What is tortious is the deployment asymmetry: the same behavioral intelligence is used to retain losing vulnerable users, and the protective intervention shows up only where law or regulation compels it.

---

[11] Public reporting and complaint allegations describe the 186-attribute profile: Zach Marcus, *The App Always Wins*, WASH. MONTHLY (Oct. 29, 2024), https://washingtonmonthly.com/2024/10/29/the-app-always-wins/ [*hereinafter App Always Wins*]; *City of Baltimore v. DraftKings, Inc.*, No. C-24-CV-25-002683 (Balt. City Cir. Ct. filed Apr. 3, 2025) (alleging that one Flutter entity "maintains at least 186 attributes" per bettor and uses that data to target individuals showing signs of addiction); *Sage v. DraftKings, Inc.*, No. 260303384 (Phila. Ct. Com. Pl. filed Mar. 24, 2026) [*hereinafter Sage* Complaint] ¶¶ 60-72 (describing platform-internal behavioral telemetry).

[12] DRAFTKINGS INC., Annual Report (Form 10-K) 5, 33 (fiscal year ended Dec. 31, 2025) [*hereinafter* DraftKings 10-K (FY2025)] (disclosing the customer-data infrastructure underlying the targeting capability and confirming "the industry practice of restricting and managing betting limits at the individual customer level"); *see also* DraftKings, App Privacy, Apple App Store (accessed Mar. 2026) (enumerating the categories of behavioral data collected on each user under the App Store privacy framework).

[13] SHOSHANA ZUBOFF, THE AGE OF SURVEILLANCE CAPITALISM 71-97, 202-223 (2019).

[14] *App Always Wins*, *supra* note 11 [hereinafter Robins Proclivity] (quoting DraftKings CEO Jason Robins from a 2022 earnings call: "There's a lot of data science and customization at a player level to serve up [Same Game Parlays] to the customers that we think have proclivity to engage with that type of bet.").

[15] Entain Mktg. UK Ltd, *System and Method for Encouraging Responsible Gameplay and Preventing Detrimental Gaming Activity*, U.S. Patent No. 12,555,436 B2 (filed Aug. 24, 2022, issued Feb. 17, 2026), Claims 3, 10-11. The patent's claims describe a system configured to detect risk indicators, classify users into vulnerability tiers, and recommend tools to reduce gaming risk and impose restrictions for particularly high-risk users. The patent thus establishes industry-feasibility evidence and constructive-knowledge evidence: the capability has been built, patented, and described in the public record. The framework's tortious-conduct claim does not depend on the patent itself proving extraction intent; it depends on the deployment asymmetry between platforms' use of comparable behavioral architectures abroad (where regulation requires protective deployment) and in the United States (where regulation does not).

This exploitation operates on a device the Constitution itself has recognized as intimate. *Riley* described cell phones as holding "the privacies of life"; *Carpenter* described the behavioral data they generate as enabling "near perfect surveillance."[16] Those characterizations describe the device. The doctrinal foundation for tort-law protection of what happens on it comes from the privacy-tort lineage developed below, not from the Fourth Amendment cases. Betting platforms convert the intimacy the constitutional cases recognize into a vulnerability map. The harm that follows is actionable in tort under the privacy doctrine the Restatement codifies.[17]

The industry labels its highest-loss customers "whales" and "VIPs." Targeting categories, not euphemisms. In sports betting, just 5.7% of bettors account for 80% of spending, and their dedicated hosts contact some clients up to 100 times per day.[18] The demographic vector of this targeting is consistent across jurisdictions: young adult men, principally aged 18 to 34, exhibiting the cultural and behavioral patterns the platform's targeting is calibrated to exploit. The Pennsylvania gambling-addiction hotline records that men aged 25 to 34 account for more than two-thirds of incoming calls; UK and Australian data show analogous concentrations. The VIP and hotline populations substantially overlap demographically — young adult men, high-frequency users, users showing escalating loss-chasing behavior — though demographic overlap is not the same claim as roster coextension.

The asymmetry is not merely informational. The platform's vulnerability-detection capability is, in clinical terms, diagnostic-adjacent: not a clinical diagnosis of gambling disorder, but a real-time read of behavioral risk markers that substantially overlap with the substrate clinical screening tools observe. The instrument operates continuously, on the user's own behavioral data, with sample sizes and predictive resolution no clinical screening tool can match. A licensed gambling-disorder specialist conducting a thirty-minute clinical interview operates on a fraction of the behavioral substrate the platform processes per session.

Yet where the clinical assessment triggers a duty of care — to inform, to refer, to refrain from harm — the platform's behavioral-risk assessment triggers a targeting cycle. The platform performs the risk-detection function of vulnerability assessment without inheriting any of the ethical or fiduciary

---

[16] *Riley v. California*, 573 U.S. 373, 403 (2014) ("the privacies of life"); *Carpenter v. United States*, 585 U.S. 296, 312 (2018) ("near perfect surveillance").

[17] *See* RESTATEMENT (SECOND) OF TORTS § 652B (1977) (intrusion upon seclusion); Daniel J. Solove, *A Taxonomy of Privacy*, 154 U. PA. L. REV. 477, 491-93 (2006) [*hereinafter* Solove, *Taxonomy*].

[18] *See* Complaint at 10, *Patel v. FanDuel, Inc.*, No. 1:24-cv-07402 (S.D.N.Y. filed Oct. 1, 2024) (alleging VIP host communicated with plaintiff "as much as 100 times every day"); *Sage* Complaint, *supra* note 11, ¶ 60 (alleging Pennsylvania Council on Compulsive Gambling data: men aged 25 to 34 account for more than two-thirds of incoming hotline calls); U.K. GAMBLING COMM'N, Adult Gambling Survey (2024) (analogous male / young-adult concentration in UK problem-gambling presentations); LANCET PUB. HEALTH Comm'n on Gambling, 9 LANCET PUB. HEALTH e950, e958 (2024) (finding that 5.7% of sports bettors accounted for 80% of spending); David Forrest & Ian G. McHale, *The Dependence of Online Gambling Businesses on High-Spending Customers*, 41 J. GAMBLING STUD. 693 (2025).

infrastructure that historically attaches to that function. The harm is engineered in this specific sense: the risk-detection capability has been built and is being deployed, and the directionality has been inverted.

What in any clinical setting would be the basis for intervention is, in the platform setting, the basis for extraction.

The asymmetry is not theoretical. Part III.A documents the system as it operates on the production record. The same individual-level behavioral profiling capability the platform deploys to identify and exile profitable bettors — Stake Factoring assigning a 0.01 multiplier to winning accounts, Gubbing stripping promotions from comparison-shoppers, IESnare device-fingerprinting flagging users who beat the line — is the capability the platform withholds when the bettor is in crisis. The detection system is in production. The defendants are operating it today, calibrated to the cent, against the users whose profitability they cannot tolerate. What is missing is the deployment of that same capability in the protective direction the public-health-sociological tradition has been documenting for two decades.

2. Exploitation Mechanisms

The exploitation is engineered into product design, operating through four mechanisms.

Push notification timing targets vulnerable moments. Notifications arrive constantly, triggered by detected behavior: loss-triggered promotions, late-night targeting, post-game re-engagement.[19]

Friction asymmetry traps funds. Deposits require one tap. Withdrawals require verification steps, processing delays, minimum thresholds, and "are you sure?" prompts. Part II.D below answers this asymmetry with friction parity.[20]

Same-game parlays extract maximum value. These bets return 31% to operators versus 5.5% for standard wagers, marketed to users with "proclivity to engage" — the CEO's euphemism for susceptibility.[21]

VIP programs institutionalize extraction. Users flagged as high-loss receive dedicated hosts whose compensation depends on betting volume. The overlap between high-value status and problem gambling is structural — the UK Gambling Commission found that VIP customers were problem gamblers at eleven times the general population rate.[22] That overlap is not incidental; it is the business model.

---

[19] *App Always Wins*, *supra* note 11.

[20] *See* Philip Newall, *Sludge, Dark Patterns and Dark Nudges: A Taxonomy of Online Gambling Platforms' Deceptive Design Features*, 120 ADDICTION 1916 (2025).

[21] *App Always Wins*, *supra* note 11 (reporting 31% operator hold on same-game parlays versus 5.5% on standard wagers).

[22] U.K. GAMBLING COMM'N, *High Value Customer and VIP Scheme Monitoring: Executive Summary* (2020) (8% of online VIP customers were problem gamblers — eleven times the general population rate).

3. The Machine Zone

These mechanisms produce something Schüll recognized on the casino floor — a state the smartphone has reproduced at scale. Schüll's fieldwork with slot machine players revealed what she termed the "machine zone" — not a clinical condition but a manufactured dissociative state, engineered by the product's architecture, in which reflective capacity shuts down and play continues on autopilot.[23] Time dissolves; losses accumulate without registering; the next bet is placed before the last one is processed. The zone is a state the product creates and the platform can detect in real time through the same behavioral data it already collects.

Schüll's contribution was methodological as well as descriptive. Years of fieldwork on the Las Vegas casino floor demonstrated that the machine zone was the product's design objective rather than an aberration of an otherwise neutral product: the architecture of casino space was engineered to produce the dissociative state in which the bettor's reflective capacities cease to operate. The mobile platform completes a transition the casino industry could not. Casino architecture required the bettor to enter it. Platform architecture follows the bettor through the device the bettor already carries. The "play to extinction" phenomenon Schüll documented — the bettor remaining at the machine until all available funds have been wagered — reappears in the platform context as the 2 AM session, the loss-chase signature, the depleted bankroll followed by the new deposit. What the public-health and ethnographic traditions have spent twenty years documenting is, in the platform context, observable in the platform's own behavioral telemetry: the 186-attribute profile, the loss-chase trigger, the late-night session metric. The ethnographer's notebook and the platform's targeting log describe the same phenomenon from opposite sides of the screen.

The slot machine achieved this through speed: seconds between wager and outcome, no pause, no offramp. Mobile sports betting has replicated that compression through microbetting — live in-game wagers that resolve within seconds, turning every play into a new pull of the lever. A bettor watching an NFL game can wager on the next snap, learn the result before the teams line up, and bet again on the snap after that. DraftKings now offers an average of 517 different bet types during a single game — four times what it offered in 2022 — with 54% of all wagers placed live and in-game.[24] FanDuel's "Pulse" feature wraps microbets in real-time game narrative, surfacing new wagers as the action unfolds.[25] The result is a product that behaves like a slot machine dressed in the language of sports fandom — and in one Australian

---

[23] Schüll, *supra* note 5, at 12-18, 166-97.

[24] DRAFTKINGS INC., Investor Day 2026 Presentation, at 29, 38 (Mar. 2, 2026) [*hereinafter* DraftKings Investor Day] (reporting 517 average bet types per game — four times the 2022 figure — and 54% live in-game wagering share).

[25] FanDuel, *Take Live Betting to the Next Level With FanDuel's New Feature: The Pulse* (FanDuel.com, accessed Mar. 2026) (describing The Pulse as tracking "the biggest storylines in sports" with new bets added "as the action unfolds").

study of 1,813 sports bettors, 78% of those who placed bets on micro events met diagnostic criteria for problem gambling.[26]

Microbetting is not an incremental variation on traditional sports wagering. A pregame bet on the Eagles to cover the spread resolves in three hours; a microbet on the next play resolves in thirty seconds. That difference in temporal compression is the difference between a poker game and a slot machine: between a product that permits reflection and one that engineers its absence. The machine zone Schüll documented on the casino floor required physical architecture: windowless rooms, recirculated air, chairs bolted in place. Microbetting reproduces the zone through temporal architecture alone, delivered to a device the bettor carries to bed.

The harm is measurable and accelerating. A UC San Diego study documented a 23% national increase in gambling-related help-seeking since *Murphy*. Maryland reported a 42% increase in disordered gambling within three years of mobile legalization. Almost twenty million American adults — 8% of the adult population — report experiencing at least one indicator of problematic gambling behavior "many times" in the past year.[27] The harm concentrates where the mechanisms concentrate — in the states that legalized mobile betting and on the platforms that engineered the zone.

Current regulatory vocabulary has no category for this.

C. WHY THE GAP EXISTS

The gap is an artifact of gambling regulation's criminal-law origins — suppress the vice, license the activity, take the tax — which treated the problem as moral vice rather than market failure.[28] What emerged was a regulatory architecture that asks whether the house is honest but has no vocabulary for whether it is predatory.

The most visible response to these harms — "responsible gambling" programs — is structurally inadequate. These programs are designed by the same platforms that profit from vulnerable users, reliant

---

[26] Alex M.T. Russell et al., *Who Bets on Micro Events (Microbets) in Sports?*, 35 J. GAMBLING STUD. 205, 212 (2019) (finding 78% of microbetters met criteria for problem gambling). The study reports correlation between microbetting participation and problem-gambling diagnostic criteria; the framework's causation argument does not depend on the directional inference (microbetting → problem gambling) that the cross-sectional study could not test directly. The temporal-compression argument developed throughout the Article identifies the mechanism by which microbet architecture produces the harm; the Russell et al. correlation establishes that the population-level harm signature is present in the predicted location.

[27] UC San Diego Today, *Study Reveals Surge in Gambling Addiction Following Legalization of Sports Betting* (2024); *New Study Finds Uptick in Problematic Gambling Since Legalization of Online Sports Betting*, MD. MATTERS (Oct. 23, 2025); NAT'L COUNCIL ON PROBLEM GAMBLING, NGAGE 3.0, at 2 (2024).

[28] *See generally* John T. Holden & Marc Edelman, *A Short Treatise on Sports Gambling and the Law: How America Regulates Its Most Lucrative Vice*, 2020 WIS. L. REV. 907; Wire Act, 18 U.S.C. § 1084 (1961); Unlawful Internet Gambling Enforcement Act, 31 U.S.C. §§ 5361-5367 (2006); Indian Gaming Regulatory Act, 25 U.S.C. §§ 2701-2721 (1988). The securities-regulation contrast developed in the body text is sourced *infra* at note 31.

on opt-in tools used by only 2-8% of customers, measuring success by metrics the industry controls.[29] The industry's position on targeting vulnerable users is explicit. When FanDuel was criticized for targeting problem gamblers, it responded that restricting such advertising would be "analogous to a liquor store not being able to advertise to customers who 'may be' alcoholics."[30]

Securities law, built specifically because Depression-era investment markets had come to resemble casinos, developed the consumer-protective logic that gambling regulation never imported. The Securities Act of 1933 and Exchange Act of 1934, together with the regulatory and enforcement architecture they enabled, identified information asymmetry coupled with misaligned incentives as the structural problem and built a multi-layered protective response: fiduciary duty, suitability requirements, anti-churning rules, Know Your Customer obligations. Each has a direct gambling counterpart that gambling law has never developed.[31]

The translation runs three columns:

Securities Doctrine	Gambling Counterpart	ATT Application
Suitability	Product appropriateness for vulnerable users	Matrix vulnerability classification determines permissible product offerings
Anti-Churning	VIP host contact volume	Contact frequency triggers under Actionable Behavioral Conditioning
Know Your Customer	Vulnerability assessment	Platform behavioral surveillance already performs this function

PART II: THE ALGORITHMIC TOXICITY TORT (ATT)

The framework developed here to close the regulatory gap is the Algorithmic Toxicity Tort.[32] ATT recognizes cognitive integrity as a protected legal interest, completing the arc of common-law autonomy. The framework rests on a three-part architecture: contextual integrity provides the duty and breach

---

[29] *See* Michael Auer, Niklas Hopfgartner & Mark D. Griffiths, *The Effects of Voluntary Deposit Limit-Setting on Long-Term Online Gambling Expenditure*, 23 CYBERPSYCH., BEHAV. & SOC. NETWORKING 113, 114 (2020) (1.3% adoption rate among 49,560 players studied); UK Gambling Comm'n Survey (2019) (9% limit-setting, 3% timeouts, 2% self-exclusion). The public health critique of responsible gambling programs has been developed comprehensively. *See, e.g.*, Livingstone & Rintoul, *supra* note 9 (arguing that industry-designed "responsible gambling" programs individualize systemic harm, manage liability rather than prevent harm, and must be replaced by a public health framework); P. Johnstone & M. Regan, *Gambling Harm Is Everybody's Business: A Public Health Approach and Call to Action*, 184 PUB. HEALTH 63 (2020); Cassidy, *supra* note 9 (developing the political-economic analysis of how the regulatory failure persists despite the empirical evidence).

[30] Callum Jones, *Top US Betting Firm Lobbied Against Rules to Protect Young People and Problem Gamblers*, GUARDIAN (Dec. 9, 2023).

[31] Securities Act of 1933, 15 U.S.C. §§ 77a-77aa; Securities Exchange Act of 1934, 15 U.S.C. §§ 78a-78qq. *See generally* Holden & Edelman, *supra* note 28, at 920-45 (documenting how state regulatory frameworks prioritized revenue generation over consumer protection).

[32] *See supra* note 4 (discussing the framework's broader applicability beyond gambling).

standard,[33] cognitive integrity identifies the protected interest, and the tobacco-litigation causation methodology grounds the causal framework, attaching liability to the operator's response to detected vulnerability.

A note on doctrinal posture before the framework is developed. The Algorithmic Toxicity Tort is articulated as a new common-law cause of action in the lineage of the Warren-Brandeis privacy tort: a recognition by courts that an evolving technology has produced a category of harm the existing tort vocabulary does not name, and that the common-law method permits naming.[34] The Proposed Restatement in the Appendix is the formalization of that cause of action. ATT is not, however, dependent on judicial recognition of a new tort *eo nomine*. The same conduct it identifies — the platform's detection of vulnerability and deployment of behavioral-conditioning mechanics in response — is also actionable under existing causes of action, and is in fact pleaded under existing causes of action in *Sage v. DraftKings*: products liability and design defect; common-law negligence with a heightened duty of care; intentional infliction of emotional distress; aiding-and-abetting liability against the leagues; consumer-protection statutes such as Pennsylvania's UTPCPL, Massachusetts's Chapter 93A, and California's Unfair Competition Law. ATT is the doctrinal scaffolding for those claims and the descriptive name for the harm those claims target. The litigation pathway runs through the existing torts; the eventual destination is the formal recognition the Restatement Appendix proposes.

A. COGNITIVE INTEGRITY AS PROTECTED INTEREST

Cognitive integrity is the interest in freedom from surreptitious, automated manipulation of one's mental processes that substantially impairs the capacity for autonomous choice. The concept does not require that choices be "good" or even rational; it requires that choices be the decision-maker's own, not the product of another party's exploitation of vulnerabilities the decision-maker cannot perceive.

The philosophical and empirical work that cognitive integrity names has been done. Calo developed the digital market-manipulation framework that distinguishes manipulation from persuasion as a matter of mechanism rather than degree.[35] Susser, Roessler, and Nissenbaum supplied the philosophical taxonomy of online manipulation as the targeting of vulnerabilities to bypass rather than engage rational agency.[36]

---

[33] *See* HELEN NISSENBAUM, PRIVACY IN CONTEXT: TECHNOLOGY, POLICY, AND THE INTEGRITY OF SOCIAL LIFE (2010); HELEN NISSENBAUM, *Privacy as Contextual Integrity*, 79 WASH. L. REV. 119 (2004) (originating contextual integrity as a privacy-norm theory describing the violation that occurs when information flows breach the contextual norms of the relationship in which the information was shared).

[34] *See* Samuel D. Warren & Louis D. Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193 (1890). The common-law method's openness to naming new categories of harm is not a recent innovation but a structural feature of how Anglo-American tort law accommodates technological change. The framework developed here is offered in that lineage.

[35] Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995, 1003-10 (2014).

[36] Daniel Susser, Beate Roessler & HELEN NISSENBAUM, *Online Manipulation: Hidden Influences in a Digital World*, 4 GEO. L. TECH. REV. 1, 18-24 (2019).

Nussbaum's capabilities theory identified practical reason as the central human capability autonomy presupposes.[37] Frankfurt's first-order/second-order desires framework, developed long before the algorithmic moment but with unusual precision for it, distinguished the autonomous agent from the "wanton" whose first-order desires operate without second-order reflective endorsement.[38] What has been missing is the doctrinal translation: the conversion of two decades of work in philosophy, technology policy, and public health into a tort-cognizable protected interest with an administrable breach standard. Cognitive integrity is the doctrinal label for the interest those traditions have collectively named without naming.

The translation follows a pattern of common-law evolution. The nineteenth century recognized battery to protect bodily autonomy from unwanted physical contact.[39] In 1890, Warren and Brandeis argued that new technology — the instant camera — required recognition of a "Right to Privacy" to protect the "inviolate personality" from unwanted publication.[40] The mid-twentieth century built on that foundation through Prosser's four-tort schematic and the Restatement (Second) §§ 652A-E privacy torts; late-twentieth and early-twenty-first century scholarship, most prominently Daniel Solove's taxonomy work, extended privacy doctrine from concealment of facts to control over information processing.[41] The architecture of engagement-optimized gambling algorithms calls for the next step: tort-law recognition of cognitive integrity as the protected interest the manipulation literature has been describing for two decades, in the doctrinal vocabulary of common-law evolution that the privacy tort itself inaugurated.

The lineage answers a federal-court standing concern that follows from *TransUnion LLC v. Ramirez*: where plaintiffs allege intangible injury, federal courts ask whether the asserted harm has a "close historical or common-law analogue" the common law has long recognized.[42] The battery-to-privacy-tort-to-cognitive-integrity progression developed here is designed to supply that analogue. *Sage v. DraftKings* proceeds in Pennsylvania state court and the question does not arise on the *Sage* facts directly, but it returns wherever the doctrine moves into federal court — in companion ATT actions, in eventual federal RICO and consumer-protection claims, and in the appellate posture every state-court action eventually faces.

The decomposition lets each tradition do the work it does best. The manipulation literature supplies the philosophical content of what cognitive integrity protects against: the targeting of vulnerability to override rather than engage rational agency. Contextual integrity supplies the breach standard. Behavioral data generated in a consumer-entertainment relationship flows appropriately to bet processing and account

---

[37] MARTHA C. NUSSBAUM, CREATING CAPABILITIES 33-34 (2011).
[38] Harry G. Frankfurt, *Freedom of the Will and the Concept of a Person*, 68 J. PHIL. 5, 7-10 (1971).
[39] RESTATEMENT (SECOND) OF TORTS § 18 (Am. L. Inst. 1965).
[40] Warren & Brandeis, *supra* note 34, at 205-07.
[41] William L. Prosser, *Privacy*, 48 CALIF. L. REV. 383, 389 (1960); RESTATEMENT (SECOND) OF TORTS §§ 652A-652E (1977); Solove, *Taxonomy*, *supra* note 17, at 484-91.
[42] *TransUnion LLC v. Ramirez*, 594 U.S. 413, 425-27 (2021).

management; it does not flow appropriately to a vulnerability-detection-and-extraction system, and the violation occurs when the platform routes the data across that contextual boundary.[43] Capabilities theory supplies the underlying autonomy theory the protected interest presupposes. Frankfurt supplies the precision tool: the agent acts freely when first-order desires (*I want to bet*) are endorsed by second-order desires (*I want to gamble rationally, within my bankroll*); the platform that detects the loss-chase moment and deploys a profit-boost notification engineers a wanton — a person acting on first-order desires the second-order will has been engineered out of the room to endorse.[44] The philosophical concept has been available for decades. What the Article claims is narrower: the doctrinal moment has arrived for the law to name it.

Paternalism restricts choices for the chooser's benefit; protecting cognitive integrity ensures that choices are actually choices. The platform that identifies when Marcus is loss-chasing and deploys a profit-boost notification does not appeal to his reasoned judgment. It targets the vulnerability itself, using his crisis state as the condition for engagement. The Frankfurt distinction supplies the doctrinal precision that distinguishes this conduct from paternalism: liberty interests presuppose the operating capacity for second-order endorsement, and what the framework targets is the platform's engineering of the conditions under which that capacity cannot operate. The agent who has been engineered into the wanton condition is not exercising the autonomy the liberty interest protects; the agent's first-order conduct has been severed from the agent's second-order will. The conduct is actionable not because the bettor's choices were unwise. It is actionable because the platform manufactured the conditions in which the bettor's choices ceased to be the bettor's own.

Gambling exploits specific neurological vulnerabilities in ways most commercial activity does not. Variable reward schedules — the core mechanic of the slot machine — drive compulsive engagement by delivering rewards at unpredictable intervals, triggering dopamine loops that make stopping difficult even when the rational course is clear.[45] Near-miss effects trick the brain into experiencing losses as almost-wins; near-misses activate the ventral striatum — the same brain region activated by actual wins — creating a

---

[43] Nissenbaum, *Privacy as Contextual Integrity*, *supra* note 33, at 137-43 (developing the normative framework for appropriate information flows within contexts).

[44] Frankfurt's original framework described person-level autonomy. Applying it to *episodic* states — transient configurations within a person's life rather than dispositional capacities — is a doctrinal extension, developed in the autonomy literature subsequent to Frankfurt. *See, e.g.*, John Christman, *Autonomy and Personal History*, 21 CAN. J. PHIL. 1 (1991); GERALD DWORKIN, THE THEORY AND PRACTICE OF AUTONOMY 15-20 (1988). The framework's claim is structural, not metaphysical: the platform's targeting architecture engineers transient states that satisfy Frankfurt's description of the wanton condition, and what the framework regulates is the manufacture of those states. The decomposition the body performs — distributing the philosophical work across the manipulation, contextual-integrity, capabilities, and Frankfurt traditions — is, as far as the author can determine, original to this Article, though all four traditions have circulated together in the philosophical literature. *See, e.g.*, Susser, Roessler & Nissenbaum, *supra* note 36 (drawing on three of the four without integrating them in the doctrinal-translation register).

[45] Luke Clark et al., *Pathological Choice: The Neuroscience of Gambling and Gambling Addiction*, 33 J. NEUROSCIENCE 17617, 17618-20 (2013).

neurological false signal that sustains play.[46] Loss chasing follows predictable patterns: craving predicts chasing behavior independently of gambling severity, with loss-chasers experiencing withdrawal-like symptoms that mirror substance-dependence patterns.[47] Each mechanism is observable in the platform's own telemetry, and the platform's response to that telemetry is the legal question.

The machine zone documented in Part I provides the empirical anchor. The translation from Schüll's casino ethnography to the mobile context is itself an inferential move, and one the framework does not assert without support. Three observations carry the translation. First, the temporal compression of microbet placement is documented by the platforms themselves: 517 average bet types per game, a fourfold increase since 2022, with 54% of all wagers placed live and in-game.[48] Second, the loss-chase signature is documented in the platforms' behavioral telemetry: the 71% of bets placed within thirty minutes of a loss, the late-night session concentration, the depleted-bankroll-followed-by-redeposit pattern, all observable in the same 186-attribute profile the platforms maintain for targeting purposes. Third, the late-night session concentration and the demographic vector of harm are documented in independent public-health data converging across jurisdictions.[49] The platform's own observations, the ethnographic tradition's observations, and the public-health tradition's observations describe the same phenomenon from three vantage points; the machine zone is the frame that names what each is observing.

Product liability doctrine provides additional support. Design defect analysis asks whether a reasonable alternative design existed. The UK's mandatory intervention requirements answer affirmatively: the same platforms already implement crisis-state protections in jurisdictions that require them — on the same corporate infrastructure, for the same parent companies.[50] The platforms that call responsible gambling operationally difficult already operate it where regulation requires.

---

[46] Luke Clark et al., *Gambling Near-Misses Enhance Motivation to Gamble and Recruit Win-Related Brain Circuitry*, 61 NEURON 481, 486-88 (2009).

[47] Marina Cosenza et al., *Getting Even: Chasing Behavior, Decision-Making, and Craving in Habitual Gamblers*, 12 BMC PSYCH. 445 (2024).

[48] DraftKings Investor Day, *supra* note 24, at 29, 38.

[49] *See supra* Part I.B (developing the demographic vector and the public-health prevalence data); PA. COUNCIL ON COMPULSIVE GAMBLING (2024); U.K. GAMBLING COMM'N, Adult Gambling Survey (2024); UC San Diego Today (2024).

[50] RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2(b) (Am. L. Inst. 2011) (design defect where reasonable alternative design existed); *see also Tincher v. Omega Flex, Inc.*, 104 A.3d 328, 384-89 (Pa. 2014) (adopting composite consumer-expectations and risk-utility tests for products liability in Pennsylvania); *see also* U.K. GAMBLING COMM'N, Licence Conditions and Codes of Practice, Social Responsibility Code Provision 3.4.3 (Remote Customer Interaction), paras. 7-11 (effective Sept. 12, 2022) [hereinafter LCCP Provision 3.4.3] (requiring operators to flag indicators of harm, prevent marketing to users with strong harm indicators, and automate intervention processes). Flutter's Real-Time Intervention system already operates in production on the same corporate infrastructure.

B. THE DUTY

1. Knowledge, Power, and Counterparty Structure

Traditional fiduciaries — lawyers, doctors, financial advisors — owe elevated duties because they possess specialized knowledge that clients rely upon and cannot readily verify.[51] Digital platforms occupy an identical position. The data architecture documented in Part I enables predictive modeling of vulnerability with a precision no traditional fiduciary possesses — continuous, behavioral, real-time, and operating on a constitutionally intimate device.[52]

If an algorithm recommending a stock portfolio owes a duty of loyalty to its user (and the SEC has confirmed that it does), then an algorithm deploying high-risk wagering products to a user it has identified as vulnerable must as well.[53]

The gambling platform's position is structurally stronger. A financial adviser earns fees regardless of client outcome; the betting platform is the bettor's counterparty — its profit is the user's loss. The zero-sum structure creates an incentive to maximize user losses that distinguishes the betting platform from virtually every other digital platform.

The objection that this argument generalizes too broadly — that every retailer profits from customer purchases, that subprime credit issuers and high-spread retail FX brokers occupy similar adverse positions — sharpens the framework rather than weakening it. Three structural features distinguish the algorithmic gambling platform. First, the platform's revenue is monotonically increasing in user impairment: as the bettor's capacity for autonomous choice degrades, the bettor's willingness to wager rises, and the platform's expected return on that wager rises with it. Ordinary commerce does not exhibit that monotonicity, because in ordinary commerce a sufficiently impaired customer eventually stops being a profitable customer. The gambling platform's profit curve runs through impairment to depletion. Second, gambling exploits acute neurological mechanisms (variable reward schedules, near-miss effects, loss-chase dopamine signaling) that have no functional analogue in ordinary retail and that have been the subject of two decades of clinical and neuroscientific literature.[54] Third, the Supreme Court's gambling-regulation jurisprudence has consistently treated wagering as a category warranting distinct regulatory architecture.[55] The framework is

---

[51] *See* RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR ECONOMIC HARM § 4 (Am. L. Inst. 2020); Jack M. Balkin, *Information Fiduciaries and the First Amendment*, 49 U.C. DAVIS L. REV. 1183, 1206-15 (2016).

[52] *Riley*, *supra* note 16, at 403; *Carpenter*, *supra* note 16, at 312.

[53] *See* Commission Interpretation Regarding Standard of Conduct for Investment Advisers, 84 FED. REG. 33,669 (July 12, 2019) (confirming fiduciary duties for "automated advisers"; noting fiduciary duty is strongest when the adviser's financial interest conflicts with the client's).

[54] *See supra* notes 45–47 and accompanying text (developing the neurological mechanisms specific to gambling); *see also* Marc N. Potenza, *The Neurobiology of Pathological Gambling and Drug Addiction: An Overview and New Findings*, 363 PHIL. TRANSACTIONS ROYAL SOC'Y B 3181 (2008).

[55] *See Greater New Orleans Broad. Ass'n v. United States*, 527 U.S. 173, 185-95 (1999) (recognizing the federal government's substantial interest in differentially regulating gambling-related advertising and analyzing such restrictions under intermediate scrutiny, while striking the federal restriction at issue as inconsistent under the third and fourth *Central Hudson* prongs); *Murphy*

portable to other counterparty-impairment contexts; it would, with appropriate doctrinal adaptation, reach a subprime credit issuer that detected acute financial-distress signatures and responded by extending additional credit. Gambling is the cleanest doctrinal vehicle because the three features converge there with unusual clarity.

The counterparty relationship activates fiduciary logic: the duty is elevated because the platform's financial interest is directly adverse to the user it serves. This is the operating principle of the system itself, as Part III.A documents — the same individual-level behavioral profiling capability the platform deploys to identify and exile profitable bettors is the capability it withholds from vulnerable losing bettors.

Massachusetts already applied this logic in the investment context. *In re Robinhood Financial LLC* sanctioned an investment platform for using gambling-design mechanics — scratch-off animations, confetti tied to trading frequency, push notifications without suitability analysis — to drive engagement without regard for client vulnerability.[56] The state applied fiduciary law to behavioral architecture structurally identical to what sports betting platforms deploy: the same gamification mechanics, the same absence of suitability analysis, the same disregard for the asymmetric information the platform has about the user. The framework extends *Robinhood*'s principle from investment to gambling, where the vulnerability detection is more precise and the counterparty structure makes the financial-interest-conflict question sharper than the broker-dealer context.

The Entain patent, properly characterized, supplies the constructive-knowledge foundation. The patent classifies users into vulnerability tiers and prescribes calibrated commercial responses at each tier; it proves industry capability and constructive knowledge of the harm-detection mechanism. It does not prove extraction intent.[57] Platforms deploying behavioral-architecture systems comparable to the patented architecture cannot credibly claim ignorance of user vulnerability. The patent establishes that the detection capability exists and has been built; the deployment-asymmetry record establishes that the capability has been deployed protectively where regulation requires and withheld where regulation does not.

---

v. NCAA, 584 U.S. 453, 481-84 (2018) (Thomas, J., concurring) (engaging the gambling-as-distinct-category question more directly than the majority opinion); *cf. Posadas de P.R. Assocs. v. Tourism Co.*, 478 U.S. 328, 345-47 (1986) (treating gambling as a regulatory category warranting heightened state authority), *abrogated on other grounds by 44 Liquormart, Inc. v. Rhode Island*, 517 U.S. 484 (1996). The proposition supported is that the doctrinal architecture of gambling regulation has historically treated wagering as a category requiring distinct regulatory treatment; *Greater New Orleans Broadcasting* carries the proposition for substantial-interest purposes (its ultimate Central Hudson holding is taken up *infra* at Part IV.B); *Murphy* is included for the concurrence's engagement with the question, not for the majority's holding, which addressed the anticommandeering doctrine.

[56] *In re Robinhood Fin., LLC*, Admin. Docket No. E-2020-0047 (Mass. Sec. Div. Jan. 18, 2024) (consent order; $7.5M settlement for gamification features deployed without suitability analysis); *see also Robinhood Fin. LLC v. Sec'y of the Commonwealth*, 492 Mass. 696 (2023) (upholding the Secretary's authority to impose fiduciary duties on broker-dealers providing investment advice).

[57] Entain Patent, *supra* note 15. Entain's UK operations were fined £17 million by the UK Gambling Commission for having behavioral data and choosing not to act; staff knew one customer was a delivery driver whose £17,000 in losses were inconsistent with their income and did not escalate. U.K. GAMBLING COMM'N, *Entain to Pay £17 Million for Regulatory Failures* (Aug. 17, 2022). The patent establishes the capability and the constructive-knowledge anchor; the enforcement action establishes that the capability was not deployed to protect users.

2. The PHAI Litigation Genealogy

The ATT's causation methodology descends from a proven lineage. The Public Health Advocacy Institute pioneered the litigation framework that dismantled the tobacco industry's "personal responsibility" defense — identifying exploitative product design, documenting corporate knowledge of harm, and reframing the legal question from consumer choice to manufacturer accountability.[58] PHAI has carried that approach forward across tobacco, food marketing, e-cigarettes, and now gambling. Each iteration applies the same framework to a new product that exploits cognitive vulnerability for profit. PHAI's Director of Gambling Policy described the current generation of betting platforms as "defectively designed and inherently dangerous AI-generated gambling model[s]" in testimony before the Senate Judiciary Committee.[59]

The ATT adapts this methodology to the algorithmic context. The "defective design" is the targeting system — the architecture that detects vulnerability and deploys extraction-maximizing stimuli — and the "corporate knowledge" is in the SOX-certified filings.

The reframing is the heart of the matter: from "the bettor chose to gamble" to "the platform chose to target a bettor it had classified as vulnerable."

3. Actionable Behavioral Conditioning

The core obligation is to refrain from Actionable Behavioral Conditioning: the systematic use of algorithmic feedback loops to identify specific cognitive vulnerabilities and deliver stimuli designed to override rational decision-making and compel engagement.[60]

Contextual integrity provides the normative standard. Behavioral data generated in a consumer entertainment relationship flows appropriately to bet processing and account management. It does not flow appropriately to a vulnerability detection and extraction system. Users consented to a gambling relationship, not a surveillance-and-exploitation relationship. Consent to one context does not authorize data flows that violate the norms of that context.[61]

---

[58] Friedman et al., *supra* note 6. The PHAI methodology traces from Richard Daynard's pioneering tobacco litigation through successive applications to food marketing and e-cigarettes. *See generally* Richard A. Daynard, *Regulating Tobacco: The Need for a Public Health Judicial Decision-Making Canon*, 30 J.L. MED. & ETHICS 281 (2002) (arguing that courts lack a public health canon of judicial decision-making and that this absence systematically produces decisions hostile to public health — a structural gap the ATT framework addresses in the algorithmic context); PUB. HEALTH ADVOC. INST., *Litigation as a Tool for Public Health* (ongoing project documentation).

[59] *Hearing on America's High-Stakes Bet on Legalized Sports Gambling Before the S. Comm. on the Judiciary*, 118th CONG. (Dec. 17, 2024) (statement of Harry Levant, Dir. of Gambling Pol'y, Pub. Health Advoc. Inst. at Northeastern Univ. Sch. of Law) (testifying that current platforms are "defectively designed and inherently dangerous AI-generated gambling models").

[60] *See* Nissenbaum, *Privacy as Contextual Integrity*, *supra* note 33, at 137-43 (supplying the contextual-norm theory the ABC standard operationalizes).

[61] *Id.* at 137-43; Susser, Roessler & Nissenbaum, *supra* note 36, at 18-24.

C. THE MANIPULATION-VULNERABILITY MATRIX

The duty not to exploit requires a line-drawing mechanism. When does marketing become manipulation? When does behavioral targeting become predation? The Manipulation-Vulnerability Matrix answers by mapping platform conduct against user susceptibility to create clear liability zones, operationalizing the breach standard of Actionable Behavioral Conditioning.[62]

The Matrix measures two variables. *Manipulation intensity* assesses the degree to which platform conduct exploits behavioral vulnerabilities — from Low (standard marketing) through Extreme (active exploitation of crisis states in real time). *User vulnerability* measures susceptibility as knowable through platform data — from Low (recreational bettor, stable stakes) through Extreme (2-4 AM sessions, multiple deposits in a single session, depleted bankroll followed by new deposits, or a user who has self-excluded and thereby communicated vulnerability through the most unambiguous channel available). The question is not whether the user is objectively vulnerable, but whether the platform can detect vulnerability through its behavioral surveillance systems. The Matrix provides a framework for courts, not a claim about the full complexity of human experience. The law requires categories, and these are anchored in what the platform itself can see.

The intersection creates three liability zones. *Green Zone* (low manipulation against ordinary susceptibility, or where the operator does not weaponize known or knowable vulnerability): no liability — normal commercial marketing to capable consumers. *Yellow Zone* (medium combinations): viable claims requiring case-by-case analysis. *Red Zone* (the high-manipulation / high-vulnerability combinations specified in Appendix B): a presumption of breach arises because the platform cannot credibly deny knowledge of vulnerability or intent to exploit.[63] The burden shifts to the platform on the breach element to demonstrate that it did not engage in Actionable Behavioral Conditioning.

The most serious objection to the framework comes from Lina Khan and David Pozen, who argue that information fiduciary theory fails to generate administrable duties.[64] The objection was legitimate when they wrote it in 2019. But the Matrix does what fiduciary theory alone cannot: it converts a principle ("don't exploit vulnerability") into a rule ("don't deploy high-manipulation conduct against users in extreme-vulnerability states without clear justification"). Their second concern — that the fiduciary label sits awkwardly on entities whose financial interests are adverse to their users' — is answered by the counterparty structure identified above. The adversity is the feature, not the bug. Where the Restatement's information-asymmetry rationale meets a directly adverse counterparty, the case for the duty is strongest.

---

[62] The Matrix's full axis definitions are set out *infra* in Appendix A § 3; the zone classification table is set out in Appendix B.

[63] *See infra* Appendix A § 3(b)-(c) (formalizing Red Zone presumption and Green Zone safe harbor with reference to the Matrix table in Appendix B).

[64] Lina Khan & David Pozen, *A Skeptical View of Information Fiduciaries*, 133 HARV. L. REV. 497, 500-520 (2019).

The Appendix formalizes this as a Proposed Restatement. With the standard set, the framework's remaining moves are mechanical: causation, remedies, and the burden-shift in Red Zone cases.

D. CAUSATION AND REMEDIES

The ATT resolves the causation challenge by adopting the substantial factor test from toxic tort litigation — the methodology PHAI systematized.[65] When a defendant's conduct materially increases the risk of a harm that eventually materializes, causation is established — even if other factors contributed. Defendants cannot simultaneously argue that their behavioral targeting justifies billions in marketing expenditure but fails to influence behavior.

The structural fit is not coincidental: algorithmic harm and toxic exposure share an evidentiary profile that the substantial-factor test was designed to address. Both involve cumulative injury produced through many small interactions, no single one of which is provably the proximate cause of the eventual harm. A user receives thousands of push notifications across months of betting; no single notification can be isolated as the cause of the gambling disorder, just as no single cigarette can be isolated as the cause of the lung cancer. What the substantial-factor test does in the toxic-tort context — relieve the plaintiff of an impossible burden of singular causation while preserving the requirement that the defendant's conduct materially contributed to the harm — is the analytical move algorithmic harm requires for the same structural reason.

The platform's own targeting determinations also supply substantial evidentiary weight on the causation question. When a platform identifies a user for VIP treatment or "proclivity" targeting, it has made its own determination, based on its proprietary behavioral data and trained against its own outcome metrics, that this user is susceptible to commercial influence. That determination is admissible evidence of the platform's belief in the targeting's effectiveness, and supports the factfinder's inference that subsequent targeting in fact contributed to the harm.[66]

The case for causation is fortified by three additional doctrinal moves that operate independently of the substantial-factor analysis. First, the lost-chance doctrine — developed in the medical-malpractice context to address harms produced through the deprivation of a chance rather than the imposition of a discrete injury — supplies the conceptual frame for the platform's specific role in the harm. The breach

---

[65] *See* RESTATEMENT (SECOND) OF TORTS § 431 (Am. L. Inst. 1965) (substantial-factor test for legal cause); *Anderson v. Minneapolis, St. Paul & Sault Ste. Marie Ry. Co.*, 179 N.W. 45, 49 (Minn. 1920) (foundational substantial-factor case in cumulative-causation contexts).

[66] A defendant's own assessment that particular conduct will be effective is generally admissible to prove effectiveness. *See* FED. R. EVID. 801(d)(2) (party-opponent admissions); 2 MCCORMICK ON EVIDENCE § 254 (8th ed. 2020). The defendant's commercial commitment to the targeting — measurable in the resources allocated to designing, deploying, and continuously refining the system — is itself evidence that the defendant believes the targeting is effective. The proposition advanced here is not that the platform is estopped from contesting causation, but that the platform's own behavioral predictions are substantial evidence on the causation question that the factfinder is entitled to weigh.

need not have caused the gambling disorder; the breach deprived the bettor of a chance to break what would otherwise have been a manageable habit. The 31% intervention-success rate at Sportsbet, the operational metric the same parent company already produces in Australia, is the lost chance quantified. (The lost-chance doctrine is recognized in roughly half of state jurisdictions, including Pennsylvania, where *Sage v. DraftKings* proceeds; where the doctrine is unavailable, the substantial-factor and burden-shifting analyses below carry the same evidentiary weight without the lost-chance label.) Second, where the defendant uniquely possesses the targeting records and the plaintiff cannot reconstruct the counterfactual, the burden-shifting logic deployed here extends *Sindell v. Abbott Laboratories* and *Summers v. Tice*, which shifted the burden of proof in cases where the plaintiff could not identify which defendant caused the harm. The principle these cases establish, that parties with exclusive control of dispositive evidence should not benefit from that control, applies *a fortiori* to the algorithmic-harm context, where the defendant possesses not merely superior knowledge but the entire causation record.[67] Third, foreseeability is established by the Entain patent itself. A defendant who has patented the targeting-of-vulnerability mechanism cannot credibly claim the resulting vulnerability-caused harm was an unforeseeable consequence; the patent recites the harm-causation chain explicitly, classifies users into vulnerability tiers, and prescribes calibrated commercial responses at each tier. Patenting the mechanism by which the harm-causation chain operates is the most thorough demonstration of foreseeability the common law recognizes.

Class certification under Rule 23, a recurring obstacle in algorithmic-harm litigation under *Wal-Mart Stores, Inc. v. Dukes* and *Comcast Corp. v. Behrend*, is more available under the framework than under most analogous theories, though not automatic. The key procedural posture is bifurcation under Rule 23(c)(4): class certification on the liability question, individualized determinations for injury and damages.

Wal-Mart's commonality requirement is satisfied at the liability stage where the class can be defined by objective platform records — for example, all users to whom the platform sent crisis-state-triggered notifications during sessions exhibiting specified loss-chase, repeated-same-session-deposit, late-night-session, or post-self-exclusion signatures. The common question is the legality of the defendant's own classification-and-targeting scheme. The individualized inquiries that defeated certification in *Wal-Mart* — whether each plaintiff suffered the same kind of discriminatory conduct from a different supervisor under a discretionary policy — do not arise. The class definition tracks targeting categories evidenced in the platform's own records, although the precise classification taxonomy the platform uses internally —

---

[67] *Sindell v. Abbott Labs.*, 26 Cal. 3d 588, 611-13 (1980) (market-share liability; shifting the burden of proof to defendants with superior knowledge in defendant-identification cases); *Summers v. Tice*, 33 Cal. 2d 80, 86 (1948) (same). The framework's application of the *Sindell-Summers* principle to causation-evidence asymmetry is a doctrinal extension of the burden-shift logic those cases established for defendant-identification asymmetry; the underlying principle — that parties should not benefit from controlling the evidence that would establish their own liability — carries from one context to the other.

whether it includes formal "Red Zone" tiers or some functional equivalent — will need to be confirmed in discovery.

Comcast's damages-capable-of-class-wide-proof requirement is the harder question, and the one the framework answers through remedy structure rather than through aggregation of individualized loss. The remedy structure prioritizes disgorgement of platform profits per affected user. Disgorgement is more tractable for class-wide proof than individualized compensatory damages, because it can be modeled from platform revenue data. The same SOX-certified filings that disclose the customer-data infrastructure and the per-customer "lifetime value" calculation establish the dataset from which per-user profit attribution can be computed — though the computation will still require expert proof to separate lawful betting losses from exploitation-period losses. *Comcast*'s damages-model requirement is satisfied because the disgorgement model maps onto the platform's own revenue accounting in a way fully individualized compensatory-damages models do not. Compensatory damages — net losses during the exploitation period, varying by plaintiff — are pursued in individualized post-certification proceedings under Rule 23(c)(4)'s bifurcation framework. The platform's behavioral telemetry records each user's exposure with timestamps and amounts, which makes the individualized damages determinations procedurally tractable even where they are not class-wide-provable in *Comcast*'s strict sense.

The bifurcated-disgorgement model is the framework's strongest claim about class certification, and a more modest claim than a more aggressive draft might assert. Class certification is not easy. What the framework offers is a procedural pathway that algorithmic-harm theories without the disgorgement structure do not have, located in the platform's own revenue accounting rather than in a doctrinal innovation a court would have to accept.

In Red Zone cases, the burden shifts to the platform on the breach element under the framework developed in Appendix A. Causation remains separately analyzed under the substantial-factor standard, with the *Sindell-Summers* records-asymmetry logic supplying additional support where the defendant uniquely possesses the targeting records that would establish or disprove the causal chain. To require plaintiffs to prove the counterfactual of what they would have done absent targeting effectively immunizes the most egregious conduct.

ATT treats sustained manipulation as a continuing tort. Rather than isolating a single bet, the law views prolonged exposure to extraction algorithms as a single ongoing wrong — each interaction reinforcing the last.[68] This extends the statute of limitations to the last act of exploitation and recognizes that cumulative harm exceeds the sum of any individual interaction.

---

[68] *See* RESTATEMENT (SECOND) OF TORTS § 834 cmt. e (Am. L. Inst. 1979) ("for continuing intrusions either by way of trespass or nuisance each repetition or continuance amounts to another wrong, giving rise to a new cause of action").

Remedies are keyed to Matrix zones: injunctive relief (prohibition of crisis-state targeting, friction parity, cooling-off mechanisms),[69] compensatory damages (net losses during the exploitation period), disgorgement (platform profits from the affected user), and punitive damages for Red Zone conduct (2-5× multipliers).[70] Friction parity requires that the cognitive and temporal barriers to depositing funds mirror those for withdrawing them — if withdrawal demands a waiting period, identity verification, and a confirmation prompt, deposit must require the same. The UK already mandates these injunctive requirements on the same platforms, the same infrastructure.[71]

E. THE AUTONOMY OBJECTION

The law protects autonomy, including the right to make self-destructive choices. Adults who choose to gamble with full knowledge of the risks are exercising a liberty interest the state has no business overriding. The objection cuts across the political spectrum, and it is the strongest argument against cognitive-integrity regulation. It also depends on a premise the framework has already rejected.

The framework does not pit autonomy against regulation. It identifies the autonomy interest that regulation defends. The protected interest the framework names is precisely the autonomy the objection invokes; what the framework targets is the platform's engineering of the conditions under which that autonomy cannot operate. The Frankfurt distinction supplies the precision: the autonomy interest is the agent's capacity for second-order reflective endorsement of first-order conduct, and that capacity is exactly what the platform's targeting architecture is designed to suppress. The framework does not restrict the bettor's freedom to gamble. It restricts the platform's freedom to manufacture the conditions of suppression.[72]

The objection's strongest version is preserved in full. Nothing in the framework prevents the bettor from gambling, from losing, or from gambling and losing on the same product the framework regulates. What the framework prevents is the platform's exploitation of the moment when the bettor's reflective capacities have been engineered out of operation by the platform's own design — and on the platform's own observation that the engineering has succeeded. The right to gamble is the right of an autonomous

---

[69] *Supporting Affordability and Fairness with Every Bet Act of 2025*, H.R. 2087, 119th CONG. (2025) (introduced Mar. 11, 2025); S. 1033, 119th CONG. (2025) (companion Senate measure). H.R. 2087's substantive provisions parallel several of the ATT's design mandates: § 103(b)(2)(E) prohibits sports wagering on any sporting event once such event has commenced (eliminating in-game wagering and microbetting); § 103(b)(2)(H) bans VIP programs and delinks customer-facing employee compensation from gambling activity; § 103(b)(6)(F) imposes affordability checks; § 103(b)(6)(G) restricts artificial-intelligence-driven targeting; § 103(b)(7) restricts marketing during 8 AM–10 PM and during live sporting event broadcasts; § 202 establishes a national self-exclusion list. Critically, § 301(b) preserves state common-law claims. The friction-parity principle the framework develops is the Article's synthesis rather than the bill's express structure; the ATT framework provides the private right of action that complements such federal standards.

[70] *See infra* Appendix A § 5 (formalizing the remedies framework).

[71] U.K. GAMBLING COMM'N, LCCP Social Responsibility Code Provision 3.4.3, *supra* note 50, paras. 7-11.

[72] *See supra* Part II.A (developing the Frankfurt second-order desires framework and the protected-interest decomposition).

agent. The framework restricts only the platform's freedom to manufacture the conditions in which the bettor is no longer that agent.

F. LOAD-BEARING ARCHITECTURE

Not every move in the framework is structural. The architecture above stacks several novel claims, and a hostile reader will push to see which dominoes can be knocked over without bringing the rest down. So here is the inventory: the framework has a spine, and it has accretions. The spine has to hold for the framework to do anything. The accretions sharpen its edges and could be set aside without collapsing the structure. The distinction matters most for litigators, who are pleading the framework's substance under existing tort categories before any court has formally recognized it.

Four moves carry the structural weight.

The first is the duty arising from the platform's possession of asymmetric behavioral data and its position as the bettor's financial counterparty.[73] Without the duty, there is nothing the platform was obligated not to do, and no other move in the framework saves it.

The second is substantial-factor causation under the toxic-tort lineage.[74] A court that requires single-bet but-for proof across years of cumulative targeting collapses the algorithmic-harm theory for the same reason the asbestos and tobacco cumulative-exposure theories would have collapsed under that standard.

The third is the *Lemmon* conduct/content distinction.[75] A court reading § 230 to immunize the platform's own targeting decisions as derivative of the third-party content displayed produces a federal statute that swallows the cause of action — whatever else the framework establishes.

The fourth is the burden-shift on causation grounded in *Sindell* and *Summers*.[76] The platform alone holds the dispositive evidence: the classification logs, the targeting timestamps, the comparative deployment metadata. A regime that requires the plaintiff to prove the counterfactual without access to the records that would establish it functions as an immunity for the very conduct the framework targets.

If any of the four falls, the framework falls with it.

The remaining moves are doctrinal sharpeners — the framework operates with them but can operate without them. Cognitive integrity as a named protected interest in tort doctrine is the most prominent of these accretions, but the conduct it identifies is actionable under existing torts (products liability, design defect, IIED, UTPCPL, common-law negligence with a heightened duty), and the *Sage* complaint pleads

---

[73] *See supra* Part II.B.1 (developing the counterparty-and-information-asymmetry duty).
[74] *See supra* Part II.D (developing the substantial-factor analysis under the toxic-tort lineage).
[75] *See infra* Part IV.A (developing the conduct/content distinction); *Lemmon v. Snap, Inc.*, 995 F.3d 1085, 1091-93 (9th Cir. 2021).
[76] *See supra* Part II.D; *Sindell*, *supra* note 67; *Summers*, *supra* note 67.

exactly those.[77] The Manipulation-Vulnerability Matrix supplies the framework's administrable rule of decision and the answer to Khan and Pozen's objection that fiduciary theory cannot generate administrable duties, but a court can find the same breach without the formalized Green/Yellow/Red taxonomy; the Matrix sharpens the analysis at the doctrinal level without being required to generate it at the litigation level. The Red Zone presumption formalizes a burden-shift the *Sindell* logic supplies on its own. The *Moody* algorithmic-feed reservation is the framework's preferred First Amendment posture, but *Central Hudson* is the doctrinal fallback that does not depend on the *Moody* reservation extending further than it currently has.[78] The four-factor synthesis applied to league liability extends *Halberstam* as *Taamneh* applies it, but *Halberstam*'s three-element § 876(b) test is directly available without the four-factor extension if a court declines to draw *Taamneh*'s distinction the way the Article reads it.[79]

The practical payoff is the litigator's. The *Sage* complaint is pleaded under existing tort categories and does not require any court to recognize a new tort *eo nomine*. The framework structures and frames those pleadings, supplying the duty's doctrinal foundation, the causation methodology, and the § 230 distinction, without committing the pleader to the accretive moves a court might decline to adopt at the motion-to-dismiss stage. An attorney can plead the existing torts under the framework's lens, carry the four load-bearing moves into the briefing as doctrinal architecture, and reserve the accretive moves for the courts willing to adopt them. The framework's structural integrity does not depend on a judiciary that adopts every claim it makes. It depends on a judiciary that holds the load-bearing four.

PART III: THE EXTRACTION MACHINE

The framework developed in Part II has a concrete target: the modern sports betting app as a system — one that identifies vulnerability, cultivates it through manufactured friendship, and exiles any user who manages to win.

A. THE SYSTEM ANATOMY

Every user enters the same machine. The phases have different names — Acquisition, Grooming, Extraction, Exile — but the direction is always the same: inward, deeper, and toward the house.

---

[77] *See Sage* Complaint, *supra* note 11 (pleading strict products liability, design defect, common-law negligence with a heightened duty of care, IIED, aiding-and-abetting against the leagues, UTPCPL, and unjust enrichment counts on the same evidentiary base the framework structures).

[78] *See infra* Part IV.B (developing the *Central Hudson* analysis as the doctrinal foundation and the *Moody* algorithmic-feed reservation as the supplementary argument); *Central Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n*, 447 U.S. 557, 566 (1980).

[79] *See infra* Part V.B (developing the four-factor synthesis as the framework's reading of *Halberstam* and *Taamneh*); *Halberstam v. Welch*, 705 F.2d 472, 477 (D.C. Cir. 1983); *Twitter, Inc. v. Taamneh*, 598 U.S. 471, 488-93 (2023).

Acquisition. The entry point is a "risk-free" bet offer — the first cognitive manipulation. The promise is that the new bettor cannot lose the first wager; the implicit promise is that the platform sees the new bettor's onboarding as a costless sample to be recovered through subsequent extraction. The asymmetry is built into the offer's mechanics: the "risk-free" refund issues as site credit, not cash, and unlocks only after the credit is wagered through — converting a first loss into the structural foundation of continued engagement. Once inside, the platform deploys "the silver bullet": Same-Game Parlays returning 31% to the operator.[80] The targeting architecture is self-reported to investors, certified to regulators, and declared to Apple.[81]

Grooming. Once a user is identified as "high-value" — industry code for high-loss — the machine shifts to VIP programs contributing up to 80% of revenue.[82] The VIP host is the human face of the extraction system, deploying manufactured intimacy to override the natural "stopping logic" of a loss.[83]

The mechanics are visible in a recent complaint filed by the Public Health Advocacy Institute. A FanDuel VIP host assigned to a high-loss bettor cultivated a relationship built on the fiction of shared stakes: "NICE WIN TERRY!!!!!! I knew we were due for one!!!!!!" The "we" positioned the host as an ally; she was a revenue instrument.[84] When the bettor's losses mounted, the host suggested he "take a timeout and enjoy the holidays with the family." Three weeks later, she called with what she described as an "emergency." The emergency was a Super Bowl gift package — two tickets, hotel, party access, transportation — extended to a bettor she had just counseled to stop.[85] The cycle of concern followed by inducement is not a failure of training. It is the system operating as designed: the appearance of care calibrated to sustain the relationship through which extraction continues.[86]

A second plaintiff in the same action self-excluded from online gambling on March 15, 2025 — the most unambiguous signal of vulnerability a bettor can send. His DraftKings VIP host continued to message him with promotional offers after the self-exclusion.[87] Under the Matrix, continued targeting of a self-excluded user is Extreme Manipulation against Extreme Vulnerability. The presumption of breach does not require inference; the platform received an explicit request to stop and chose to keep selling.

---

[80] *App Always Wins*, *supra* note 11.
[81] DraftKings 10-K (FY2025), *supra* note 12, at 5, 33; DraftKings, App Privacy, Apple App Store (accessed Mar. 2026); *Sage* Complaint, *supra* note 11, ¶¶ 60-72.
[82] *Sage* Complaint, *supra* note 11, ¶ 60 (citing PA Council on Compulsive Gambling data); *Patel* Complaint, *supra* note 18, at 10.
[83] Keith Whyte, *VIP Programs and the Cultivation of Manufactured Intimacy* (NAT'L COUNCIL ON PROBLEM GAMBLING, working paper 2024).
[84] *Patel* Complaint, *supra* note 18, at 10-12 (alleging VIP-host communications quoted in body text).
[85] *Id.* at 13-14.
[86] *See* RESTATEMENT (SECOND) OF TORTS § 323 (Am. L. Inst. 1965) (liability for negligent performance of an undertaken duty); *see also Sage* Complaint, *supra* note 11, ¶¶ 88-95 (developing the "undertaken to act" theory in the VIP-host context).
[87] *Sage* Complaint, *supra* note 11, ¶¶ 96-101.

Extraction. The "Whale Trap" is sprung during moments of acute crisis, detected in real time. The platform sends push notifications at 2 AM to users chasing thousands in losses — one-tap deposits, instant, while withdrawals are buried under verification delays and "reverse withdrawal" periods.[88] The capability to intervene already exists in production: Flutter Entertainment's Real-Time Intervention system, deployed at Sportsbet in Australia, intervenes when behavior deviates from baseline. Thirty-one percent of prompted customers lower or cancel their deposits.[89] The technology exists, it works, and it is deployed selectively.

Exile. The clearest evidence of intent. Users who win are not rewarded — they are removed. Stake Factoring assigns profitable bettors a factor of 0.01, limiting their bets to pennies. Gubbing strips all promotions, leaving accounts as hollow shells. IESnare device-fingerprinting technology identifies devices to flag users who comparison-shop odds.[90] All three amount to individual-level behavioral profiling, deployed to protect revenue from winners. That profiling capability is withheld from vulnerable losers. The asymmetry is a design choice, documented in SOX-certified federal securities filings.[91]

The industry's standard defense before regulators — that problem gambling is subjective and real-time detection operationally infeasible — collapses against its own Exile infrastructure. A platform that identifies a profitable bettor with enough precision to limit wagers to pennies cannot credibly testify that comparable precision is unavailable when the user is in crisis. The detection capability is deployed to protect the house's money, not the bettor's life.

Losers are retained; winners are exiled.

B. THE CONTRADICTION

To the public: responsible gaming is "fundamental to our mission."[92] To investors: CEO compensation — 100% tied to Revenue and Adjusted EBITDA, with zero weight on any responsible gambling metric; $195 million for Simplebet — a microbetting platform promising "betting opportunities during every moment of a game" — versus approximately $2 million cumulative for responsible gambling councils.[93] Flutter's CEO, in a February 2026 shareholder letter, described FanDuel's strategy as "leveraging our scale, proprietary technology, and data advantages" to deliver "smarter personalization" and "richer live engagement,"

---

[88] Newall, *supra* note 20.

[89] FLUTTER ENTM'T PLC, *Sustainability Report* (FY2024) (describing Sportsbet Real-Time Intervention deployment and the 31% intervention-success rate); *see also Sage* Complaint, *supra* note 11, ¶¶ 110-118 (alleging that the same Flutter parent corporate infrastructure operates FanDuel without the comparable intervention deployment).

[90] *App Always Wins*, *supra* note 11 (describing Stake Factoring, Gubbing, and IESnare device-fingerprinting practices).

[91] DraftKings 10-K (FY2025), *supra* note 12, at 5 (acknowledging "the industry practice of restricting and managing betting limits at the individual customer level").

[92] DraftKings, *Responsible Gaming Mission Statement* (DraftKings.com, accessed Mar. 2026).

[93] DRAFTKINGS INC., Definitive Proxy Statement (Schedule 14A) 38-42 (FY2025) (disclosing CEO compensation tied 100% to Revenue and Adjusted EBITDA, with no responsible-gambling metric weighting); DraftKings, *DraftKings to Acquire SimpleBet* (press release, Sept. 12, 2024) (announcing $195M acquisition); NAT'L COUNCIL ON PROBLEM GAMBLING, Annual Report 2024-2025 (reporting cumulative industry contributions to responsible-gambling councils).

because "excellence in the fundamentals compounds into retention and lifetime value."[94] "Retention and lifetime value" is the investor translation of what the bettor experiences as inability to stop.

The contradiction extends to the company's own tools. DraftKings' "My Stat Sheet" gives players a retrospective behavioral dashboard — a rearview mirror showing losses after the fact. The company uses identical data in real time to identify users with "proclivity to engage" with high-hold products — a targeting dashboard. Players see what happened; DraftKings uses the same data to drive what happens next.[95]

This is cosmetic compliance: symbolic programs that create the appearance of good standing while internal incentives remain extraction-dependent.[96] The responsible gambling tools are the symbolic layer; the compensation structure and investment priorities are what govern. The contradiction establishes scienter: the companies know what they are doing because they have told two different audiences two different stories about why they do it. Federal courts have accepted exactly this evidentiary architecture in the tobacco context, sustaining scienter findings where industry public statements contradicted internal corporate knowledge.[97]

The contradiction is also evidentially decisive. The defendants operate the same behavioral-detection infrastructure on demographically and behaviorally similar populations under different regulatory regimes, with measurably different deployment outcomes — the comparative regulatory configuration developed in Part VI. The defendants' United Kingdom and Australian operating records constitute experimental controls against which the United States operations can be measured. The regulatory variation is the most legally salient explanatory variable, though not the only one; rigorous comparative analysis must account for market maturity, enforcement capacity, demographic differences, and platform-specific implementation choices. The discovery, in the deepest sense, has already been done. The defendants did it themselves, on their own books, in service of their own foreign regulators. Each disclosure the framework relies on belongs to one of several evidentiary categories addressed in detail in Appendix C: party-opponent admissions where made by the defendant entity itself; industry-feasibility evidence where made by a competitor or affiliate; public-record allegations where pleaded in related litigation. The defense will not be litigating against the framework's empirical premises. It will be litigating against its own clients' prior

---

[94] FLUTTER ENTM'T PLC, Q4 2025 Results Presentation (Feb. 26, 2026) (CEO Peter Jackson letter describing FanDuel's strategy of "leveraging our scale, proprietary technology, and data advantages" to deliver "smarter personalization" and "richer live engagement").

[95] DraftKings, *My Stat Sheet* (DraftKings.com, accessed Mar. 2026); *Robins Proclivity*, *supra* note 14 (describing data-science targeting of users with "proclivity to engage" with high-margin parlays).

[96] *See* Cassidy, *supra* note 9, at 142-58 (developing the "responsible gambling" critique as cosmetic compliance); Livingstone & Rintoul, *supra* note 9.

[97] *United States v. Philip Morris USA Inc.*, 449 F. Supp. 2d 1, 26-27 (D.D.C. 2006), *aff'd in part, vacated in part on other grounds*, 566 F.3d 1095 (D.C. Cir. 2009) (sustaining scienter findings against tobacco industry where industry public statements contradicted internal corporate knowledge).

public statements, made under the legal compulsions of two regulatory regimes more demanding than the United States's.

C. THE MATRIX APPLIED: A SPECIMEN CASE

Marcus — whose 2:17 AM notification opened this Article — opened his DraftKings account fourteen months earlier. His account file reflects: weekly recreational wagering escalating to daily betting by month seven; average bets rising from $25 to $210; 71% of bets placed within thirty minutes of a loss; sessions concentrated between 11 PM and 3 AM; four deposits in a single overnight session.[98]

At 2:17 AM, following $1,800 in losses across three hours, DraftKings pushed a notification:

"EXCLUSIVE: 50% Profit Boost on Your Next Bet. One-Tap Claim →"

Vulnerability: EXTREME. Each indicator the Matrix recognizes is present in the record and traceable to a specific source. The loss-chase signature — 71% of bets within thirty minutes of a loss — is observable in the platform's own session-and-wager telemetry, obtainable through discovery in the same data structure DraftKings has declared to Apple under the App Store privacy framework.[99] Stake escalation from $25 to $210 over fourteen months sits in the platform's wager-history table for the named user. Late-night concentration between 11 PM and 3 AM sits in the platform's session timing logs. Four same-session deposits on the night in question sit in the platform's deposit ledger. None of this requires judgment. The platform's own data-collection capabilities, certified to investors in the DraftKings 10-K and to Apple in the App Privacy declarations, make the determination algorithmic: the customer-data infrastructure and the per-customer "lifetime value" calculation already perform the classification the Matrix asks the trier of fact to make.[100] The Entain patent supplies the doctrinal capstone. The detection capability is patented, the vulnerability tiers are taxonomized, the calibrated commercial responses are prescribed.[101] Industry feasibility and constructive knowledge are on the public record.

Manipulation: EXTREME. The notification has four constituent elements, each mapping to a category of evidence already developed in this Article. The real-time behavioral trigger is the platform's own targeting decision, executable only through the detection infrastructure DraftKings has self-described to investors and to Apple. The profit framing — "50% Profit Boost" — recasts loss as opportunity. It is the literal text of the notification, preserved in the user's device record. The one-tap mechanism eliminating

---

[98] *See supra* note 1 (Marcus is a composite drawn from the patterns documented in the *Sage v. DraftKings* complaint and the public-record evidence base; the behavioral signatures attributed to Marcus are calibrated to those the platform's own self-described detection categories recognize).

[99] DraftKings, App Privacy, Apple App Store (accessed Mar. 2026) (enumerating behavioral data categories collected); DraftKings 10-K (FY2025), *supra* note 12, at 5, 33.

[100] DraftKings 10-K (FY2025), *supra* note 12, at 5 (per-customer "lifetime value" calculation).

[101] Entain Patent, *supra* note 15, Claims 3, 10-11.

deliberative friction is observable on the live product, and it stands in deliberate asymmetric relation to the friction the platform places on withdrawals — the asymmetry the SOX-certified 10-K language about "the industry practice of restricting and managing betting limits at the individual customer level" already concedes is calibrated bettor-by-bettor.[102] The timing — 2:17 AM, after $1,800 in losses across three hours — exploits the acute distress the platform's vulnerability-detection system was designed to recognize.

Zone: RED. The presumption of breach attaches under the framework developed in Appendix A § 3(b). The burden shifts to DraftKings on the breach element to demonstrate that its conduct was not Actionable Behavioral Conditioning. None of the available rebuttals is plausible on these facts. The platform either deployed a real-time intervention (the Sportsbet system documented in Part III.A is in production today on the same Flutter corporate infrastructure that operates FanDuel) or it did not; the answer is in the platform's deployment-region metadata, obtainable through targeted discovery.

Remedies. Full net losses during the exploitation period. Disgorgement of DraftKings' profits from Marcus. Punitive damages warranted — the Entain patent establishes that the detection capability was known, developed, and patented; the choice to target rather than intervene is the knowing, willful conduct punitive damages exist to deter.[103] The UK operates each injunctive requirement today, under the same parent companies.[104]

The defendants built the evidentiary foundation themselves.

D. THE LITIGATOR'S PATH

The framework structures and frames the *Sage* complaint's pleadings. It supplies the duty's doctrinal foundation, the causation methodology, and the § 230 distinction. It does not commit the pleader to accretive moves a court might decline to adopt at the motion-to-dismiss stage.[105] An attorney can plead the existing torts under the framework's lens, carry the spine into briefing as doctrinal architecture, and reserve the accretive moves for courts willing to adopt them. The path through the framework runs through three doctrinal moves and one evidentiary architecture.

Red Zone pleading. The first move is Matrix classification at the complaint stage. The Marcus specimen case in Part III.C models the form. Identify the user's vulnerability indicators by reference to the platform's own self-described detection categories: loss-chase, stake escalation, late-night session concentration, same-session redepositing, self-exclusion. Identify the manipulation indicators by reference to the platform's own real-time targeting categories: crisis-state push notifications, profit-boost framing,

---

[102] DraftKings 10-K (FY2025), *supra* note 12, at 5, 33 (acknowledging individual-customer-level betting-limit management).
[103] *See infra* Appendix A § 5 (formalizing the remedies framework keyed to Matrix zones).
[104] U.K. GAMBLING COMM'N, LCCP Social Responsibility Code Provision 3.4.3, *supra* note 50, paras. 7-11.
[105] *See supra* Part II.F (developing the load-bearing architecture and the litigator's payoff of the spine/accretion distinction).

one-tap mechanisms, VIP-host promotional contact. Plead the intersection as Red Zone, with specificity tied to dates, times, and notification text where available. At the pleading stage, Red Zone allegations create a plausible inference of breach sufficient to survive a motion to dismiss. The formal presumption-of-breach burden-shifting becomes operative at later stages, where the platform must produce the records that uniquely lie within its possession. The *Sage* complaint already does this in substance; the framework supplies the doctrinal category and the burden-shift consequence that follows from it.

The burden-shift sequence. The Red Zone presumption of breach operates with different procedural force at each stage. At Rule 12(b)(6), the presumption supplies plausibility: a complaint alleging Red Zone facts states a claim for which relief can be granted. At discovery, the presumption justifies requiring the operator to produce the classification logs, targeting timestamps, and intervention-deployment metadata uniquely within its possession. At summary judgment and trial, the presumption shifts the persuasion burden on the breach element to the operator, while causation remains separately analyzed under the substantial-factor standard.[106] In a counterparty-structure case the burden-shift carries unusual evidentiary force: the platform is in exclusive possession of the very classification data on which the Matrix assessment turns, and the same SOX-certified filings that concede individual-level betting-limit management for profitable bettors concede the same capability is available for vulnerable losing bettors. The platform that owes the rebuttal is the only party that can produce the records that would discharge it. Discovery requests drafted to that asymmetry — formulated to compel production rather than to seek information the platform has incentive to withhold — are the instrument the burden-shift exists to enable.

Discovery targets. Six categories of platform records, each identifiable in advance from public sources, supply the evidentiary spine of the Red Zone case.

The customer-attribute behavioral profile maintained on each named plaintiff, with assigned vulnerability classifications and risk scores, is the dataset DraftKings has declared to Apple under the App Store privacy framework and disclosed in its 10-K.

The vulnerability classification logs themselves — the tier assignments, the triggering criteria the Entain-patent-style system applies, and the timestamps at which each named plaintiff moved between tiers — are in the platform's own systems, though the precise internal taxonomy remains a discovery target rather than a known artifact.

The complete log of all targeted promotional communications sent to plaintiffs, with timestamps, triggering criteria, and the algorithms that selected them, is the platform's own "why-did-this-notification-fire" record.

---

[106] *See infra* Appendix A §§ 3(b), 4(b) (formalizing the breach presumption and the causation analysis); RESTATEMENT (SECOND) OF TORTS § 433B(2) (Am. L. Inst. 1965) (providing the doctrinal foundation for the burden-shifting framework where the defendant uniquely possesses the records that would establish causation).

VIP host compensation structures — the formula by which hosts are paid, the relationship of that formula to the assigned bettor's losses, and the performance metrics against which hosts are evaluated — are in the platform's HR and compensation systems.

Real-time intervention deployment metadata in non-US markets is in the production records of the Flutter Real-Time Intervention system at Sportsbet (Australia) and the LCCP 3.4.3 affirmative-customer-interaction systems at Entain UK, FanDuel UK, and the other UK subsidiaries. These are the records that establish that the same corporate infrastructure deploys the protective system the United States platforms claim is operationally infeasible.

Internal documents reflecting comparison of US versus UK or Australian operational practices, including any analyses of why detection systems deployed abroad are not deployed domestically, supply the documents that, in the tobacco litigation, produced the scienter findings on which RICO liability ultimately rested.[107]

The evidentiary foundation the platform built itself. The framework's structural advantage is that the moves it requires are documented in materials the platforms have already produced. SOX-certified 10-K language concedes individual-customer-level betting-limit management; Apple App Store privacy declarations enumerate the targeting attributes; the Entain patent recites the vulnerability-tier classification system and the calibrated commercial responses; investor-day presentations announce the microbet bet-type expansion and the live-in-game share; VIP host complaints filed in *Sage* and *Patel* document the manufactured-intimacy operational manual. These sources fall into different evidentiary categories: party admissions where made by the defendant entity itself (the DraftKings 10-K and investor presentations are admissions against DraftKings; the Flutter Sustainability Report is admission-grade against Flutter and FanDuel); industry-feasibility and constructive-knowledge evidence where made by a competitor or industry incumbent (the Entain patent supplies industry capability evidence rather than a DraftKings admission); public-record allegations where pleaded in related litigation (the *Sage* and *Patel* complaints are allegations subject to proof, not admissions); and discovery roadmaps where they identify likely internal records (Appendix C maps each public source to the platform records likely to corroborate or amplify it).[108] The litigator's task is to compel the records that complete the picture; the framework supplies the doctrinal category that makes those records relevant; and the platform's prior disclosures supply the causal-knowledge findings the substantial-factor analysis requires. Public records, discovery targets, and expert

---

[107] *Philip Morris*, *supra* note 97, 449 F. Supp. 2d at 851-908 (RICO liability premised on internal documents reflecting industry knowledge contradicting public statements).

[108] *See infra* Appendix C (formalizing the four-category evidentiary taxonomy and mapping each public source to its evidentiary category, the platform records likely to corroborate or amplify it, and the discovery vehicle most likely to obtain those records).

testimony — the categories Appendix C maps in detail — together produce a Red Zone complaint whose evidentiary spine is, in substance, already on the platform's filing record.

PART IV: LEGAL BARRIERS

The framework's reach across § 230, the First Amendment, federal preemption, arbitration, and the assumption-of-risk defense is the test of whether ATT survives the doctrinal architecture courts have built around the platforms it targets. Each barrier is genuine; each is also surmountable in this configuration. Two conceptual moves recur: the conduct/content distinction that permits courts to address what platforms *do* even where they cannot address what platforms *publish*, and the targeted/general-audience distinction that permits courts to address surreptitious manipulation directed at identified individuals even where they cannot regulate marketing to a general audience.

A. SECTION 230: CONDUCT, NOT CONTENT

Section 230 immunizes platforms for publishing third-party content. It does not immunize platforms for designing systems that target users for exploitation.[109] Recent decisions have made this distinction operative in exactly the configurations the framework concerns.

In *Lemmon v. Snap, Inc.*, the Ninth Circuit refused § 230 immunity for Snapchat's "Speed Filter," a feature design that allegedly encouraged dangerous driving. The court held that liability for the platform's own design choices does not depend on third-party content.[110] The Third Circuit followed in *Anderson v. TikTok* (2024), allowing a wrongful-death claim to proceed where a ten-year-old died attempting a "Blackout Challenge" promoted by TikTok's recommendation algorithm; the court treated the algorithm as the platform's own expressive activity rather than protected third-party content.[111] The post-*Anderson* line is not unanimous. The Ninth Circuit in *Doe v. Grindr* affirmed dismissal of state-law product-liability claims against the platform's matching design, treating those claims as necessarily implicating Grindr's role as publisher of third-party content.[112] The New York Appellate Division, Fourth Department, in *Patterson v. Meta Platforms* declined to extend *Anderson*, holding that algorithmic recommendation of third-party content remains publishing activity protected by § 230 and editorial activity protected by the

---

[109] 47 U.S.C. § 230(c)(1).

[110] *Lemmon v. Snap, Inc.*, 995 F.3d 1085, 1091-93 (9th Cir. 2021).

[111] *Anderson v. TikTok, Inc.*, 116 F.4th 180, 188-92 (3d Cir. 2024).

[112] *Doe v. Grindr Inc.*, 128 F.4th 1148, 1152-56 (9th Cir. 2025) (affirming dismissal of state-law claims for defective design, defective manufacturing, negligence, failure to warn, and negligent misrepresentation as barred by § 230 because each claim necessarily implicated the platform's role as publisher of third-party content).

First Amendment under *Moody*.[113] The split runs along a doctrinal seam ATT does not need to traverse: ATT targets the platform's own conduct, not the algorithmic curation of third-party content.

ATT exclusively targets platform conduct: real-time vulnerability detection, targeted notification deployment, friction asymmetry, VIP host compensation structures, exile mechanisms. Each is the platform's own conduct. None is third-party content the platform merely hosts. Marcus's notification was generated by DraftKings, sent by DraftKings, and timed by DraftKings to exploit the loss-chase signature DraftKings detected in DraftKings' own behavioral telemetry. There is no third-party content in the chain of conduct. § 230 is not engaged.

The Supreme Court has signaled that the question of § 230's scope for algorithmic conduct is unresolved. *Twitter, Inc. v. Taamneh* declined to resolve § 230's reach to algorithmic recommendation.[114] The framework does not depend on resolution: ATT's targets are platform design choices and operational conduct, which fall outside § 230's protected publishing function under the consistent line of conduct/content cases the courts have already developed.

B. First Amendment: Central Hudson Foundation, Moody Reservation

The First Amendment objection is the most serious. Commercial speech enjoys real protection, and a framework targeting "manipulation" risks doctrinal architecture that could be turned against legitimate marketing. The framework survives the objection on a doctrinal foundation it has held for forty-five years (the *Central Hudson* test) and on a doctrinal supplement the Supreme Court left open in 2024 (the *Moody* algorithmic-feed reservation).[115]

1. The Central Hudson Foundation

The conventional First Amendment analysis applies *Central Hudson Gas & Electric Corp. v. Public Service Commission* to commercial speech: lawful and not misleading; substantial state interest; direct advancement; reasonable tailoring.[116] Each prong yields the framework's result on the conduct it targets.

Lawful and not misleading. Lawful commercial speech directed at general audiences receives intermediate scrutiny. ATT does not target lawful commercial speech directed at general audiences. It targets surreptitious manipulation directed at users the platform has identified as vulnerable. The prong is satisfied because the speech the framework regulates is not the kind the prong was designed to protect.

---

[113] *Patterson v. Meta Platforms, Inc.*, 2025 NY Slip Op 04447 (App. Div. 4th Dep't July 25, 2025) (declining to follow *Anderson v. TikTok* and holding that § 230 and the First Amendment bar tort claims against social-media defendants for content-recommendation algorithms; aligning with *Force v. Facebook, Inc.*, 934 F.3d 53 (2d Cir. 2019), and *M.P. v. Meta Platforms, Inc.*, 124 F.4th 519 (4th Cir. 2025)).

[114] *Twitter, Inc. v. Taamneh*, 598 U.S. 471, 506 (2023) (declining to resolve § 230's reach to algorithmic recommendation).

[115] *Central Hudson Gas & Elec. Corp. v. Pub. Serv. Comm'n*, 447 U.S. 557 (1980); *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024).

[116] *Central Hudson*, *supra* note 115, at 566.

Substantial state interest. Cognitive integrity, and the addiction crisis the framework documents, supplies the state interest. *Greater New Orleans Broadcasting Ass'n v. United States* recognized the federal government's substantial interest in differentially regulating gambling-related advertising; the case cuts the framework's way on the substantial-interest prong even though it ultimately struck the federal restriction at issue under the third and fourth prongs.[117] The state interest in regulating exploitation of detected vulnerability is, if anything, more substantial than the state interest in regulating gambling advertising generally. The framework is not regulating who may say what about gambling. It is regulating what data-targeting platforms may do to users they have classified as vulnerable.

Direct advancement. Restrictions targeted at the specific exploitation mechanisms — crisis-state push notifications, friction asymmetry, VIP host structures with extraction-correlated compensation — directly advance the state interest in cognitive integrity by addressing the specific conduct that produces the harm. The framework does not regulate gambling advertising in general. It regulates the platform conduct that converts general advertising into individualized exploitation.

Reasonable tailoring. The framework's tailoring is the work *Greater New Orleans Broadcasting* and *Lorillard Tobacco Co. v. Reilly* did to drive home: where a regulation sweeps too broadly (*Lorillard* struck down a 1,000-foot-from-school outdoor tobacco-advertising restriction as too broad in geographic and audience scope), the regulation fails the fourth prong even where the state interest is compelling.[118] The framework adopts the lesson directly. The conduct it targets is *individualized* exploitation, not categorical advertising restriction. The Matrix's tailoring runs through the user's vulnerability classification, not through the platform's audience or geography. A push notification deployed at 2:17 AM to a user the platform has classified as crisis-state is regulated; the same push notification deployed at 2:17 PM to a user the platform has not classified as vulnerable is not. The tailoring is at the vulnerability/conduct intersection — the most individualized tailoring the *Greater New Orleans Broadcasting* and *Lorillard* lessons permit.

The fourth prong is also where defendants will press hardest, citing *Sorrell v. IMS Health Inc.*'s heightened scrutiny for content-based restrictions on commercial speech.[119] The framework concedes the engagement: the rules it would generate distinguish between the *content* of platform communications (a

---

[117] *Greater New Orleans Broad. Ass'n v. United States*, 527 U.S. 173 (1999). The case is sometimes cited as supporting expansive commercial-speech regulation; the more accurate reading is that *Greater New Orleans Broadcasting* recognized a substantial state interest in differential gambling-advertising regulation while striking down the federal restriction at issue under the third and fourth *Central Hudson* prongs because the federal scheme was internally inconsistent and not adequately tailored. The framework reads the case for both lessons: the substantial state interest is established (prong two), and the tailoring lesson — that regulation must be calibrated to the specific harm rather than imposed as a categorical advertising restriction — informs the framework's individualized-targeting design (prong four).

[118] *Lorillard Tobacco Co. v. Reilly*, 533 U.S. 525, 561-66 (2001) (striking down 1,000-foot-from-school outdoor tobacco-advertising restriction as too broad in geographic and audience scope under the fourth Central Hudson prong).

[119] *Sorrell v. IMS Health Inc.*, 564 U.S. 552 (2011).

profit-boost notification is regulated; a sports-news notification is not) and so are content-based in *Sorrell*'s sense. The framework does not concede the result.

Sorrell invalidated a Vermont statute that selectively suppressed pharmaceutical-detail communications because it concluded the statute was a "viewpoint discrimination": the state preferred non-pharmaceutical-detail speech to pharmaceutical-detail speech on the merits of the speech itself. The framework's distinction is not viewpoint-based. It does not prefer non-gambling speech to gambling speech, and it does not prefer non-betting messages to betting messages. The framework's distinction is *conduct-targeting*: the regulated category is messages whose deployment exploits behavioral data the platform has used to identify the recipient as vulnerable. The state's interest is in the conduct of the data-driven targeting, not in the content of the speech.

The targeting/general-audience distinction is the same distinction federal regulators draw when they regulate manipulative communications under securities and consumer-protection law without offending the First Amendment. *Sorrell* itself acknowledged that its content-based-restriction analysis was not the only First Amendment frame for commercial speech.[120] On the most demanding reading of *Sorrell*, the framework satisfies the test it imposes; on the more conventional *Central Hudson* reading, the framework satisfies the four-prong test with room to spare.

2. The Moody Algorithmic-Feed Reservation

The 2024 Supreme Court decision in *Moody v. NetChoice* signaled a doctrinal opening the framework relies on as a supplement to (not a replacement for) *Central Hudson*. The majority's central question was whether algorithmic curation enjoys First Amendment protection like a newspaper's editorial decisions. The majority's affirmative answer was qualified: Justice Barrett's controlling concurrence, joined in relevant part by Justice Jackson, expressly distinguished individualized algorithmic feeds from editorial curation, reserving the question whether such feeds receive equivalent protection.[121] The reservation does not announce a new test. It identifies an open question, and the framework's First Amendment posture leaves room for it without depending on it.

The reservation tracks the framework's own architecture. ATT distinguishes manipulation (the platform's targeting of an identified individual) from speech (commercial communications directed at audiences). *Moody*'s qualified affirmance — protecting editorial curation while reserving individualized

---

[120] *Id.* at 567 ("commercial speech is no exception" to general First Amendment principles, but recognizing that conduct-targeted regulation operates on a different doctrinal register than viewpoint-discriminatory restriction).

[121] *Moody*, *supra* note 115, at 745-50 (Barrett, J., concurring) (distinguishing individualized algorithmic feeds from editorial curation; reserving the question whether individualized feeds receive equivalent First Amendment protection); *id.* at 751 (Jackson, J., concurring in part) (joining the relevant portions of the Barrett concurrence). The reservation is the framework's preferred posture for the algorithmic-individualized-targeting register; the *Central Hudson* analysis above stands as the doctrinal foundation that does not depend on the reservation extending further than it currently has.

algorithmic feeds — maps onto exactly the same distinction. On the conduct ATT targets (crisis-state targeting of users whose vulnerability the platform has detected), the *Moody* reservation is, at minimum, available. Defendants asserting that algorithmic feeds receive editorial-curation protection must engage the reservation rather than proceeding as though it had not been written.

The framework does not require courts to read the *Moody* reservation as broadly as it could be read. The *Central Hudson* analysis above stands on its own. *Moody* is the preferred posture where the case can be made on the algorithmic-individualized-targeting register; *Central Hudson* is the doctrinal fallback where it cannot.

3. A Recent Case Worth Noting

A recent federal-court engagement with algorithmic targeting illustrates that the legal architecture the framework relies on has been developing in the courts even as the framework was being written.

In *Garcia v. Character Technologies, Inc.*, a federal district court allowed a wrongful-death claim against an AI chatbot company to survive a motion to dismiss after the platform's algorithmic targeting allegedly contributed to a teenager's suicide.[122] The court rejected a sweeping First Amendment defense at the pleading stage, treating the platform's algorithmic targeting as conduct subject to scrutiny under traditional tort principles. *Garcia* is among the early federal-court engagements with algorithmic targeting at the motion-to-dismiss stage; the framework reads it as one data point in a developing line, not as a categorical resolution.

Neither *Garcia* nor the line of post-Anderson decisions settles the framework's First Amendment question. They establish that courts are not treating algorithmic targeting as a categorical immunity zone, and they supply the doctrinal posture *Central Hudson* and *Moody* together generate when applied to specific conduct the framework targets.

C. PREEMPTION, ARBITRATION, AND ASSUMPTION OF RISK

1. Preemption

State common law claims survive UIGEA, the Wire Act, and the post-*Murphy* federal landscape. None of those statutes occupies the field of consumer-protection or tort liability arising from gambling platforms; UIGEA expressly preserves state law,[123] and the Wire Act addresses interstate transmission of wagers rather than the platform-conduct categories the framework targets. The post-*Murphy* federalism architecture is, in fact, the framework's structural ally: state common-law tort claims are the doctrinal vehicle the framework principally relies on, and the SAFE Bet Act of 2025 — were it to be enacted in its current form — would

---

[122] *Garcia v. Character Techs., Inc.*, 785 F. Supp. 3d 1157 (M.D. Fla. May 21, 2025).

[123] 31 U.S.C. § 5362(10)(B) (UIGEA's express preservation of state law).

expressly preserve state common-law claims under § 301(b).[124] The Article does not depend on enactment of the SAFE Bet Act; it identifies the bill as the federal-legislative counterpart to the framework's tort-liability architecture, and notes that the bill's preservation provisions would resolve any preemption question by statute.

State supreme courts confronting state gaming-act preemption arguments have likewise rejected categorical readings: the Michigan Supreme Court in *Davis v. BetMGM* held that the state's Lawful Internet Gaming Act did not abrogate or preempt common-law claims of fraud, conversion, and breach of contract arising from a gambling-platform dispute, reversing the Court of Appeals' preemption holding.[125]

2. Arbitration: The Three-Pathway Analysis

Arbitration is the framework's most serious procedural barrier. It deserves the depth Parts II and III could not give it.

Sportsbook user agreements universally contain mandatory binding arbitration clauses with class-action waivers. Federal Arbitration Act jurisprudence is favorable to enforcement: *AT&T Mobility LLC v. Concepcion* and *Epic Systems Corp. v. Lewis* together establish that the FAA preempts state-law restrictions on arbitration clauses in most consumer and employment contexts.[126] The arbitration architecture is the principal procedural fortification the platforms rely on against the framework's substantive moves.

The response here is not that the arbitration architecture is invalid. It is that the architecture is incomplete. Three pathways permit the framework's substantive claims to proceed despite the arbitration clauses, and the framework's strength on the merits matters in each.

Pathway one: public enforcement is unaffected. State Attorneys General, federal regulators, and state gaming commissions are not bound by private arbitration agreements between the platform and individual users. The framework's substantive analysis — the duty, the Matrix breach standard, the substantial-factor causation methodology — applies in public-enforcement actions exactly as it applies in private ones. The Baltimore action against DraftKings, filed by the Baltimore City Solicitor on behalf of the City of Baltimore in April 2025, illustrates the pathway in operation.[127] State AG actions to enforce consumer-protection statutes (Pennsylvania's UTPCPL, Massachusetts's Chapter 93A, California's UCL, New York's GBL § 349) are unaffected by the platform's private arbitration architecture. These are the venues in which the framework's substantive moves can be litigated to a public adjudicative judgment that develops the doctrine for the field.

---

[124] SAFE Bet Act, *supra* note 69, § 301(b) (preserving state common-law claims).

[125] *Davis v. BetMGM, LLC*, ___ N.W.3d ___, 2025 Mich. LEXIS 1224 (Mich. July 22, 2025) (holding the Michigan Lawful Internet Gaming Act did not abrogate common-law claims of fraud, conversion, and breach of contract arising from a gambling-platform dispute, and rejecting the operator's argument that the LIGA preempted such claims).

[126] *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333 (2011); *Epic Sys. Corp. v. Lewis*, 138 S. Ct. 1612 (2018).

[127] *City of Baltimore v. DraftKings, Inc.*, No. C-24-CV-25-002683 (Balt. City Cir. Ct. filed Apr. 3, 2025).

Pathway two: state-law carve-outs and pre-arbitration plaintiffs. Several state-law doctrines preserve subsets of private claims from the arbitration architecture. The most prominent is California's *McGill* rule: a contractual waiver of the right to seek public injunctive relief in any forum is unenforceable as a matter of California public policy.[128] Where a plaintiff seeks injunctive relief that would benefit the public — an injunction against crisis-state targeting that would protect not only the plaintiff but other users similarly exposed — *McGill* preserves the action even where the underlying contract compels arbitration of damages. Other state-law doctrines supply analogous carve-outs in the relevant jurisdictions. New Jersey's *Atalese* requirement that arbitration waivers of substantive rights be made with clear understanding is the most useful in the *Sage* configuration.[129] The *Sage* litigation also involves plaintiffs whose first interactions with the platforms predated the current generation of arbitration clauses, and plaintiffs whose particular contractual architecture is subject to other state-law challenges. Identifying which named plaintiffs fall outside the arbitration architecture is itself a discovery-and-pleading question that operates independently of the substantive analysis.

Pathway three: mass arbitration. The third pathway is the most procedurally aggressive and the most consequential. Where individual users have been subject to common targeting practices, plaintiffs' counsel can file thousands of individual arbitration demands simultaneously, triggering the per-case filing fees the platforms structurally agreed to bear when they imposed the arbitration architecture.[130] The Postmates and DoorDash mass-arbitration cases illustrate the mechanic: when faced with tens of thousands of simultaneous individual demands, the platform either pays the aggregated filing fees (which can run into the tens of millions of dollars before any case is heard) or returns to the litigation table to negotiate. The framework's substantive analysis — Red Zone classification, the burden-shift on breach, the substantial-factor causation methodology, the four-category evidentiary taxonomy — operates in arbitration with the same force it operates in court. AAA arbitrators apply state common law. The doctrines the framework develops are state common law. Arbitral awards develop the doctrine.

Why this still matters when the arbitration wall holds. Even where none of the three pathways is available, the substantive analysis is still doing work. Arbitral awards are increasingly published and increasingly cited. Mass-arbitration settlements are increasingly disclosed in regulatory filings, and the doctrines that drive those settlements are the ones the framework develops. Public enforcement runs on its own track. Federal-legislative reform — the SAFE Bet Act's substantive provisions parallel several of the

---

[128] *McGill v. Citibank, N.A.*, 2 Cal. 5th 945 (2017); *see also Blair v. Rent-A-Center, Inc.*, 928 F.3d 819, 825-31 (9th Cir. 2019) (applying *McGill* to FAA-arbitration context).

[129] *Atalese v. U.S. Legal Servs. Grp., L.P.*, 219 N.J. 430 (2014).

[130] *See* Maria Glover, *Mass Arbitration*, 74 STAN. L. REV. 1283 (2022) (developing the mass-arbitration mechanic and its consequences for the FAA-arbitration architecture); *Abernathy v. DoorDash, Inc.*, 438 F. Supp. 3d 1062 (N.D. Cal. 2020) (illustrating the mass-arbitration filing-fee mechanic in the gig-economy context).

framework's design mandates — has the analytical architecture available as the doctrinal-policy foundation. Pre-arbitration-clause plaintiffs whose claims the platforms' first-generation user agreements did not anticipate carry the analysis into court directly. State gaming commissions are now beginning to develop the regulatory architecture the documented harms will require. The arbitration wall channels reach but does not extinguish it. The analysis is forced into venues other than ordinary civil litigation — public enforcement, mass arbitration, regulatory rule-making, legislative reform — where its substantive contributions develop the law governing the platforms, whether the platforms litigate or arbitrate.

3. Assumption of Risk

Bettors agreed to gamble. They did not agree to algorithmic exploitation. Assumption of risk requires informed consent to the *specific* risks one assumes. The architecture of platform exploitation — vulnerability detection, real-time targeting, friction asymmetry, VIP-host compensation structures, exile mechanisms — is not disclosed to bettors and could not be inferred from the gambling activity itself.[131] The defense fails on the conventional informed-consent analysis: a bettor cannot assume a risk the bettor has not been told exists, and cannot assume a risk produced by conduct the platform itself does not disclose.[132]

PART V: LEAGUE LIABILITY

The framework's reach extends to the leagues whose data and equity arrangements with the sportsbooks make the platform-conduct architecture economically viable. The doctrinal pathway for that extension runs through aiding-and-abetting under the Restatement (Second) of Torts § 876(b) — the doctrinal home of the *Halberstam v. Welch* / *Twitter v. Taamneh* line — and is the most novel doctrinal move the framework attempts.

A. STRUCTURAL EXPOSURE

The leagues are not bystanders to the conduct the framework targets. The economic architecture binds them to it.

The NFL holds approximately 6.0% of Genius Sports's outstanding shares — equity acquired through warrants priced at $0.01 per share, on terms calibrated to the volume of in-play betting transactions Genius Sports's data products enable.[133] In fiscal year 2025, the league earned approximately $126.1 million

---

[131] RESTATEMENT (SECOND) OF TORTS §§ 496A-496G (Am. L. Inst. 1965) (informed consent and assumption-of-risk doctrine; requiring informed consent to the specific risk assumed).

[132] RESTATEMENT (SECOND) OF TORTS § 496D & cmt. b (Am. L. Inst. 1965) (the plaintiff "does not assume a risk of harm . . . unless he then knows of the existence of the risk and appreciates its unreasonable character"; specific knowledge of the risk and an opportunity to evaluate it are required for assumption of risk to attach).

[133] GENIUS SPORTS LTD., Annual Report (Form 20-F) (Mar. 17, 2026) [*hereinafter* Genius Sports 20-F] (disclosing NFL warrant-priced equity at $0.01 per share, representing approximately 6.0% of outstanding shares).

in commissions on the GGR/NGR generated through the in-play data feed Genius Sports supplies to sportsbooks under the league's exclusive data partnership.[134] The league's revenue is monotonically increasing in the volume of in-play microbetting — precisely the product category the framework's empirical analysis identifies as the most directly correlated with problem gambling, and the product category the SAFE Bet Act § 103(b)(2)(E) would prohibit.

The league's equity-and-commissions exposure is the structural foundation on which the aiding-and-abetting analysis builds. The league is not merely the upstream content supplier whose product the sportsbooks happen to wager on; the league is an equity holder in the data-supply chain whose revenue rises with the in-play wagering volume the framework identifies as the highest-harm product category.

The further structural fact the league cannot disclaim is its institutional silence on the wager-type restrictions the Matrix would generate. The New York State Gaming Commission's letter of February 4, 2026 confirmed that no league has requested any wager-type restriction in New York: no request to restrict prop bets on individual player performance, no request to restrict in-play microbetting, no request to restrict the bet types most directly correlated with problem gambling.[135] The proposition is jurisdiction-specific (the New York Gaming Commission has confirmed the position only in New York), and the Article does not overclaim it as a categorical denial that any league has ever requested any restriction in any jurisdiction. The narrow proposition is sufficient: in a major regulatory market in which the leagues hold the plenary regulatory tool of requesting wager-type restriction, the leagues have not exercised it. The leagues' subsequent regulatory advocacy reflects a consistent pattern of institutional silence on the wager-type restrictions the framework would identify as Red Zone product categories.

B. THE DOCTRINAL FRAMEWORK

The doctrinal test is *Halberstam v. Welch*'s three-element formulation under Restatement § 876(b): the principal performed a wrongful act causing injury; the secondary actor was generally aware of its role in the overall tortious activity at the time the assistance was provided; and the secondary actor knowingly and substantially assisted the principal violation.[136] The *Sage* aiding-and-abetting count is pleaded under that test, and the elements can be established on the leagues' record without doctrinal extension.

---

[134] Genius Sports 20-F, *supra* note 133 (disclosing $126.1M in NFL GGR/NGR commissions for fiscal year 2025); *see also Sage* Complaint, *supra* note 11, ¶¶ 145-58 (developing the league-data-supply-and-equity theory).

[135] Letter from Brian O'Dwyer, Chair, N.Y. STATE GAMING COMM'N (Feb. 4, 2026) (confirming that no league has requested wager-type restriction in New York; noting specifically that no league has requested restriction of prop bets on individual player performance, no league has requested restriction of in-play microbetting, and no league has requested restriction of the bet types most directly correlated with problem gambling). The proposition the framework supports on the basis of this letter is jurisdiction-specific: the New York Gaming Commission has confirmed the position only in New York, and the framework does not represent that the position has been confirmed across all state regulatory jurisdictions; the proposition is sufficient for the framework's purposes because in a major regulatory market in which the leagues hold the plenary regulatory tool of requesting wager-type restriction, the leagues have not exercised it.

[136] *Halberstam v. Welch*, 705 F.2d 472, 477 (D.C. Cir. 1983) (three-element test under RESTATEMENT (SECOND) OF TORTS § 876(b)).

Twitter v. Taamneh applied *Halberstam* to the algorithmic-platform context, holding that mere algorithmic provision of services to bad actors did not, on the facts of that case, supply the substantial-assistance element required for aiding-and-abetting under the Anti-Terrorism Act.[137] The leagues' position is materially distinguishable from the *Taamneh* defendants' broad-spectrum service relationship: they hold equity in the data supplier whose product enables the in-play microbet category, earn commissions tied to in-play wagering volume, and have declined to exercise the regulatory tools that would limit the highest-harm product categories.

To organize the *Halberstam* analysis on the *Sage* facts, four factors carry the analytical weight: designed exposure, knowledge, substantial assistance, and structural enmeshment. The framework offers them as an analytical lens — an ordering principle for the *Halberstam* elements as they apply to the leagues — not as a separate four-element test a court must independently adopt. A court applying *Halberstam* directly can find aiding-and-abetting liability without adopting the four-factor lens.[138]

Designed exposure. The first factor asks whether the secondary actor's commercial architecture is built to benefit specifically from the principal's tortious conduct rather than incidentally from a broader product or service relationship. The leagues' equity-and-commissions exposure to Genius Sports satisfies the factor: the warrant-priced equity, the GGR/NGR commission structure, and the in-play data partnership architecture together create a commercial alignment in which the league's economic return rises specifically with the in-play microbetting volume the framework identifies as the highest-harm product category. The leagues' commercial relationship with the sportsbooks is no longer the broad-spectrum sponsorship-and-licensing arrangement that historically characterized league-licensee relationships. The post-2018 architecture is a specifically calibrated equity-and-data partnership in which the league's return is monotonically increasing in the metric the framework identifies as the harm vector.

Knowledge. The second factor — the knowledge element of *Halberstam*'s general-awareness test — is satisfied by the public record the framework has built throughout this Article. The leagues have access to the same public-health evidence the framework draws on. The leagues have the same regulatory exposure to the UK and Australian operating records as the sportsbooks themselves, through their parallel league-and-licensee architecture in those jurisdictions. The leagues have received the regulatory and public-advocacy communications about wager-type restriction that the New York and other state gaming commissions have facilitated, and have responded with the institutional silence the New York Gaming Commission letter documents. The general-awareness element does not require the leagues to know each

---

[137] *Twitter, Inc. v. Taamneh*, 598 U.S. 471, 488-93 (2023).

[138] The four-factor synthesis developed in this Section — designed exposure, knowledge, substantial assistance, structural enmeshment — is the Article's reading of how the *Halberstam v. Welch* / *Twitter v. Taamneh* line applies to the leagues in the *Sage* configuration. It is not the doctrinal test either case independently announces, and a court applying *Halberstam*'s three-element test directly need not adopt the four-factor framework to reach the same liability conclusion the framework reads as supported.

named plaintiff's particular circumstances; it requires the leagues to be generally aware of the tortious conduct's character. The Article's empirical record establishes that awareness comprehensively.

Substantial assistance. The third factor — the substantial-assistance element of *Halberstam* — is the factor on which *Taamneh* turned in the Anti-Terrorism Act context, and the factor on which the Article's reading of *Halberstam* and *Taamneh* together does the most work. *Taamneh* held that mere algorithmic provision of services to bad actors did not supply substantial assistance where the platform's relationship to the third-party content was at the broadest level of generality. The leagues' relationship to the sportsbooks' tortious conduct is materially distinguishable: the league supplies the in-play data feed without which the microbet product category cannot operate, holds equity in the data supplier, and earns commissions tied to the in-play wagering volume the data supplier facilitates. What satisfies the substantial-assistance element here is not the broad sponsorship-and-licensing architecture. It is the specific data-supply-and-equity architecture that makes the highest-harm product category economically viable. A court could read *Taamneh* more broadly and decline to find substantial assistance on these facts. The claim is that the *Halberstam* / *Taamneh* line, on the reading offered here, permits the substantial-assistance finding on the leagues' specifically calibrated data-supply-and-equity architecture, not that any court must find it.

Structural enmeshment. The fourth factor, the most novel of the four, asks whether the secondary actor's commercial relationship to the principal is so structurally integrated that the principal's tortious conduct cannot be isolated from the secondary actor's commercial participation. The leagues' equity-and-commissions exposure to the data-supply chain that makes the in-play microbet product category economically viable is the structural enmeshment the factor identifies. The league does not stand outside the tortious commercial architecture as a third-party service provider. It stands inside the architecture as an equity holder whose return rises with the volume of the tortious conduct.

The four factors together carry the Article's reading of how *Halberstam* and *Taamneh* combine in the *Sage* configuration. The claim is more modest than a categorical doctrinal extension: the four-factor synthesis identifies the configuration in which aiding-and-abetting liability appropriately attaches to the leagues, and the underlying *Halberstam* three-element test is directly available without the four-factor synthesis if a court declines to draw *Taamneh*'s distinction the way this Article reads it. The accretion is the four-factor synthesis. The spine is the underlying *Halberstam* test as it has long been understood and the *Sage* complaint's aiding-and-abetting count as it is pleaded under that test.

PART VI: THE FEDERAL AND COMPARATIVE FRAME

The framework's empirical premise is that platform exploitation of detected vulnerability produces measurable harm, and that the same platforms operate the same detection-and-targeting architecture under

different regulatory regimes with measurably different deployment outcomes. The comparative architecture is the framework's strongest single empirical claim, and what follows develops it with the care the claim deserves.

A. THE COMPARATIVE CONFIGURATION

Three regulatory regimes operate on the same multinational platform corporate infrastructure. The United Kingdom, through the Gambling Commission's Licence Conditions and Codes of Practice (LCCP), requires platforms to flag indicators of harm, prevent marketing to users showing strong harm signals, and automate intervention processes. Australia, through the Northern Territory Racing and Wagering Commission's licensing architecture and the parallel state-level frameworks, requires comparable affirmative customer-interaction practices. The United States — that is, the post-*Murphy* state-by-state regulatory landscape — requires neither the LCCP-style affirmative intervention nor the Australian-style customer-interaction protocols.

The same parent companies operate in each regulatory regime. Flutter operates Sportsbet under the Australian regulatory regime and FanDuel under the United States regime; Entain operates Ladbrokes Coral under the UK regime and BetMGM (through the joint venture with MGM) under the United States regime. The comparative configuration runs on the same corporate infrastructure, the same behavioral-detection architecture, and the same data-collection systems. The variable that differs across the configurations is the regulatory regime.

The comparative outcomes are documented in the platforms' own operational records and corporate filings. Flutter's Real-Time Intervention system operates in production at Sportsbet under the Australian regime; the same parent company has not deployed comparable intervention architecture at FanDuel under the United States regime.[139] Entain's UK operations have been the subject of regulatory enforcement actions in which the company was found to possess vulnerability data and to have failed to act on it; the company's United States operations through BetMGM operate under the same parent's behavioral-detection capability without the comparable affirmative-intervention requirement.

B. THE EMPIRICAL INFERENCE

The framework's empirical claim from the comparative configuration is more carefully calibrated than a categorical attribution of all harm-difference to the regulatory variable. The methodologically appropriate inference is a difference-in-differences inference: the same platforms, operating the same detection

---

[139] *See* FLUTTER ENTM'T PLC, *Sustainability Report* (FY2024) (documenting Sportsbet Real-Time Intervention deployment and the 31% intervention-success rate); *Sage* Complaint, *supra* note 11, ¶¶ 110-118 (alleging that the same Flutter parent corporate infrastructure operates FanDuel without the comparable intervention deployment).

architecture on demographically similar populations, deploy measurably different intervention practices under different regulatory regimes, with measurably different harm outcomes. The regulatory variable is the most legally salient explanatory variable, though not the only one. Market maturity, enforcement capacity, demographic differences, platform-specific implementation choices, and cultural variables in attitudes toward gambling all contribute to the cross-jurisdictional harm-difference, and a rigorous comparative analysis must account for them.

The claim is correspondingly modest. The comparative configuration establishes three things. First, that the affirmative-intervention practices the framework would generate are operationally feasible: the same platforms operate them today, in production, under different regulatory regimes. Second, that the platforms' standard defense before United States regulators — that real-time intervention is operationally infeasible, vulnerability detection unreliable, affirmative customer interaction too costly — is contradicted by their own operating records in jurisdictions where the regulatory regime requires the protective deployment. Third, that the regulatory variable is a strong causal candidate for the cross-jurisdictional harm-difference, even if it is not the only possible explanatory variable. The normative claim — that the United States regulatory regime should require the protective deployment the platforms already operate elsewhere — does not require the comparative configuration to be a pure natural experiment in which all other variables are held constant. It requires only that the comparative configuration establish that the protective deployment is operationally feasible and that the platforms have chosen not to deploy it where regulation does not require it.

C. THE FEDERAL-LEGISLATIVE FRAME

The SAFE Bet Act of 2025, introduced by Senator Blumenthal and Representative Tonko in March 2025, supplies the federal-legislative counterpart to the framework's tort-liability architecture. The Article does not depend on enactment of the SAFE Bet Act, and the framework's analytical architecture operates with full force under the post-*Murphy* state-by-state regulatory landscape regardless of federal-legislative developments. The bill's substantive provisions — § 103(b)(2)(E)'s prohibition of in-play wagering, § 103(b)(2)(H)'s ban on VIP programs, § 103(b)(6)(F)'s affordability checks, § 103(b)(6)(G)'s restriction on AI-driven targeting, § 103(b)(7)'s marketing-time restrictions, § 202's national self-exclusion list, and § 301(b)'s preservation of state common-law claims — parallel several of the framework's design mandates and create the federal-statutory backdrop against which the framework's tort architecture would operate.[140] The framework supplies the doctrinal-policy foundation for the bill's substantive moves; the bill, were it to

[140] SAFE Bet Act, *supra* note 69 (federal legislative counterpart to the framework's tort-liability architecture).

be enacted, would provide the federal-statutory architecture within which the framework's state common-law claims would proceed without preemption complications.

CONCLUSION

At 2:17 AM, Marcus's phone buzzed. The notification was generated by an algorithm that knew, in the precise sense the platform's behavioral-data architecture makes available, that he had been losing for three hours, that he was deep in the loss-chase signature, that he was in the late-night session window, that he had just made his fourth deposit of the night. The algorithm chose that moment to deploy the profit-boost notification. The "one-tap" mechanism eliminated the friction that might have prompted reflection. The 50% framing recast loss as opportunity. The notification was the platform's own conduct, executed against a user the platform had classified as vulnerable, on the platform's own data, in pursuit of the platform's own profit.

What current regulation cannot reach is what the framework names. Cognitive integrity is the protected interest the bettor's autonomy presupposes. Actionable Behavioral Conditioning is the breach standard the platform's targeting architecture violated. Substantial-factor causation supplies the analytical frame the cumulative-targeting harm requires. The disgorgement-and-injunction remedy structure tracks the relief the platform's own revenue-accounting makes computable. The Manipulation-Vulnerability Matrix is the rule of decision. The Red Zone presumption invokes the burden-shift the *Sindell-Summers* records-asymmetry logic supports.

The framework's empirical foundation rests on materials the platforms produced themselves. The DraftKings 10-K's concession of individual-customer-level betting-limit management is a SOX-certified admission of the targeting capability that grounds the duty. The Entain patent classifies users into vulnerability tiers and prescribes calibrated commercial responses at each: the industry's own admission of the harm-detection mechanism that supplies constructive knowledge. The Flutter Sustainability Report documents the Real-Time Intervention system at Sportsbet, conceding that the protective deployment is operationally feasible. The Genius Sports 20-F discloses the NFL warrant equity and the GGR/NGR commission structure on which the aiding-and-abetting analysis rests. Together they are the evidentiary spine of the *Sage* litigation and of the successor litigation already implicit in the platforms' own filings.

The framework's doctrinal architecture rests on common-law moves the courts have already made. *Lemmon* and *Anderson* have established the conduct/content distinction the § 230 analysis turns on, and the contrary post-*Anderson* line in *Doe v. Grindr* and *Patterson* reaches algorithmic curation of third-party content, which is not the conduct ATT targets. The *Central Hudson* / *Greater New Orleans Broadcasting* / *Lorillard* line supplies the tailoring lessons the First Amendment analysis incorporates. *Moody*'s

algorithmic-feed reservation identifies the doctrinal opening on which the algorithmic-individualized-targeting frame can be made. *Halberstam* and *Taamneh* supply the aiding-and-abetting framework applied to the leagues. *Sindell* and *Summers* supply the burden-shifting logic extended to the algorithmic-harm context. *In re Robinhood* establishes the fiduciary architecture extended from investment to gambling. The courts are not being asked to make doctrinal moves they have not made. They are being asked to apply moves they have already made to specific platform conduct the framework targets.

The framework's load-bearing architecture is simpler than its surface complexity might suggest. The duty arises from the platform's possession of asymmetric behavioral data and its position as the bettor's financial counterparty. The breach occurs when the platform deploys that data to target a user it has classified as vulnerable. Causation runs through the substantial-factor analysis the toxic-tort lineage establishes, supplemented by the *Sindell-Summers* burden-shift the records-asymmetry justifies, and grounded by the foreseeability the Entain patent's recitation of the harm-causation chain establishes. The remedies key to the disgorgement-and-injunction structure the platform's own revenue-accounting makes computable. The accretive moves — cognitive integrity as a named protected interest, the formalized Matrix taxonomy, the four-factor synthesis applied to league liability — sharpen the architecture without bearing its structural weight.

The litigation pathway runs through the existing torts. The *Sage* complaint pleads strict products liability, design defect, common-law negligence with a heightened duty of care, IIED, aiding-and-abetting against the leagues, UTPCPL, and unjust enrichment. The framework structures and frames those pleadings without committing the pleader to doctrinal moves a court might decline at the motion-to-dismiss stage. The eventual destination is the formal recognition the Restatement Appendix proposes. The pathway in the meantime is the existing torts the *Sage* complaint already pleads.

The platforms built the evidentiary foundation themselves. They certified the targeting capability to investors and to Apple. They patented the vulnerability-detection mechanism. They deployed the protective infrastructure abroad while withholding it at home. They compensated their executives on metrics that exclude any responsible-gambling weighting. They institutionalized the manufactured intimacy of the VIP host architecture. They exiled the bettors who win and retained the bettors who lose, and they did all of it in a manner the SOX-certified filings document with the precision the securities laws require.

What the framework names is conduct the platforms have already disclosed. The doctrinal architecture the *Sage* litigation and its successors require is in place; the work of applying existing doctrines to specific conduct the platforms have built belongs to the courts. The platforms will contest the reading. What they cannot easily contest is the record they themselves produced — the reading the public-health-empirical tradition has been building for two decades, the reading PHAI has been refining since the tobacco settlements, the reading the *Sage* plaintiffs are now testing in the Philadelphia Court of Common Pleas.

The casino has moved from the floor to the phone. The regulatory architecture that governed the floor does not govern the phone. Cognitive integrity is the protected interest the new architecture requires; Actionable Behavioral Conditioning is the breach standard the platforms' conduct violates; the existing torts are the doctrinal pathway through which courts can reach the conduct the platforms have built. The work this asks of courts is the work the common-law method has done before: name the new harm, apply the existing doctrines, produce the rule of decision the next case requires.

The new harm has a name now. The doctrinal pathway is open. The evidentiary foundation is on the platforms' own filing record. The litigation has begun.

APPENDIX A: THE ALGORITHMIC TOXICITY TORT — PROPOSED RESTATEMENT

What follows is the doctrinal scaffolding the framework would generate as a Restatement-style codification of the Algorithmic Toxicity Tort. It is presented in the structural form Restatements customarily use. It is "Proposed" in the literal sense: the cause of action this Article develops, formalized for the kind of work briefs and discovery memoranda will need it to do. It is not a Restatement promulgated by the American Law Institute. The litigation pathway does not depend on it. The *Sage* complaint pleads the underlying conduct under existing tort categories, and the eventual destination of the framework is the formal recognition the form below proposes. What the form below supplies in the meantime is the doctrinal architecture briefing, discovery design, and case-law engagement can pull from.

§ 1. DEFINITIONS

(a) *"Behavioral telemetry"* means data generated by a user's interactions with a digital platform, including patterns of engagement, timing, frequency, intensity, and response to platform stimuli, that the platform collects, processes, or analyzes to inform commercial decisions about the user.

(b) *"Cognitive integrity"* means the interest in freedom from surreptitious, automated manipulation of one's mental processes that substantially impairs the capacity for autonomous choice. The interest does not require that choices be wise or rational; it requires that choices be the decision-maker's own, not the product of another party's exploitation of vulnerabilities the decision-maker cannot perceive and the platform can detect.

(c) *"Vulnerability"* means a condition of impaired or compromised capacity for autonomous choice that the operator can detect through behavioral telemetry, including without limitation: loss-chase signatures (a meaningful proportion of bets placed within thirty minutes of a loss); stake escalation (substantial increase in average wager over time without offsetting account growth); late-night session concentration (substantial proportion of activity during overnight hours); same-session redepositing (multiple deposits in a single session); post-self-exclusion contact attempts; and other behavioral signatures the operator's own detection systems classify as risk indicators.

(d) *"Manipulation intensity"* means the degree to which platform conduct exploits behavioral vulnerabilities, ranging from Low (standard commercial marketing not calibrated to detected vulnerability) through Medium (general targeting using behavioral data to identify likely engagement) through High (targeting that uses behavioral data to identify users in detected vulnerability states) through Extreme (targeting that exploits crisis states in real time, including without limitation crisis-state push notifications,

one-tap wagering mechanisms deployed during crisis-state sessions, and VIP-host promotional contact directed at users showing strong vulnerability signatures).

(e) *"Actionable Behavioral Conditioning"* means the systematic use of algorithmic feedback loops to identify specific cognitive vulnerabilities and deliver stimuli designed to override rational decision-making and compel engagement. Actionable Behavioral Conditioning is the breach standard for the cause of action established by this Restatement.

(f) *"Operator"* means a digital platform, including without limitation a sports betting platform, that collects behavioral telemetry from users and uses that telemetry to inform commercial decisions about those users.

§ 2. Cognitive Integrity as Protected Interest

(a) Cognitive integrity, defined in § 1(b), is recognized as a legally protected interest, in the lineage of the common-law privacy torts the Restatement (Second) of Torts §§ 652A-E codifies and extending the privacy-tort tradition to encompass the interest in freedom from surreptitious algorithmic manipulation that the digital-platform context has produced.

(b) The interest in cognitive integrity is held by all natural persons against operators that collect behavioral telemetry from those persons and use that telemetry to inform commercial decisions about those persons.

(c) Recognition of cognitive integrity as a protected interest does not require recognition of a new tort *eo nomine*. The same conduct that violates cognitive integrity is actionable under existing tort categories, including without limitation: products liability and design defect; common-law negligence with a heightened duty of care arising from information asymmetry and counterparty structure; intentional infliction of emotional distress; aiding-and-abetting under Restatement § 876(b); state consumer-protection statutes; and breach of fiduciary duty where the operator's relationship to the user satisfies the fiduciary-relationship requirements of the relevant jurisdiction.

§ 3. Breach: The Manipulation-Vulnerability Matrix

(a) *General principle.* An operator breaches its duty to a user when the operator engages in Actionable Behavioral Conditioning toward that user. The Manipulation-Vulnerability Matrix, set out in Appendix B, provides the rule of decision the breach analysis requires.

(b) *Red Zone presumption of breach.* Where the operator's conduct toward the user falls within the Red Zone of the Matrix as defined in Appendix B's classification table — that is, the high-manipulation / high-vulnerability combinations the table identifies — a presumption of breach arises. The presumption operates with different procedural force at different stages of litigation. At the pleading stage, allegations

supporting Red Zone classification create a plausible inference of breach sufficient to survive a motion to dismiss under Rule 12(b)(6). At the discovery stage, the presumption justifies requiring the operator to produce the behavioral telemetry, classification logs, targeting timestamps, and intervention-deployment metadata uniquely within the operator's possession. At the summary-judgment and trial stages, the presumption shifts the burden of persuasion on the breach element to the operator, who must demonstrate by a preponderance of the evidence that the operator's conduct did not constitute Actionable Behavioral Conditioning.

(c) *Green Zone safe harbor.* Where the operator's conduct toward the user falls within the Green Zone of the Matrix as defined in Appendix B's classification table — including without limitation conduct toward users not in detected vulnerability states, conduct toward users in vulnerability states that does not weaponize the detected vulnerability, and conduct that the operator can demonstrate was deployed without reference to the user's behavioral telemetry — no breach is established and no presumption of breach arises. Conduct toward a user the Matrix classifies as Extreme Vulnerability falls within the Green Zone where the operator can demonstrate that its conduct was Low Manipulation as defined in § 1(d), or that the operator deployed protective intervention rather than commercial targeting in response to the detected vulnerability.

(d) *Yellow Zone case-by-case analysis.* Where the operator's conduct toward the user falls within the Yellow Zone of the Matrix as defined in Appendix B's classification table — the medium-manipulation / medium-vulnerability combinations that the table identifies — the breach analysis proceeds without the presumption that operates in the Red Zone, on a case-by-case examination of the operator's conduct and the user's vulnerability under the general principle stated in § 3(a).

§ 4. CAUSATION

(a) *General principle.* Causation in actions under this Restatement is analyzed under the substantial-factor test the Restatement (Second) of Torts § 431 establishes, recognizing that the algorithmic-targeting context produces cumulative harm through many small interactions, no single one of which is provably the proximate cause of the eventual harm — a structural configuration the substantial-factor test was designed to address in the toxic-tort context.

(b) *Substantial-factor causation.* Where the operator's Actionable Behavioral Conditioning materially increased the risk of the harm the user suffered, and the harm of that kind subsequently materialized, causation is established under § 4(a) without requiring the user to prove that any specific instance of the operator's conduct was the but-for cause of the harm. The operator may rebut the substantial-factor showing by demonstrating that its conduct did not materially contribute to the harm or that an intervening cause supplied the proximate causation.

(c) *Records-asymmetry burden-shifting.* Where the operator uniquely possesses the records that would establish or disprove the causal chain — including without limitation the behavioral telemetry, classification logs, targeting timestamps, and intervention-deployment metadata uniquely within the operator's possession — the burden of producing those records shifts to the operator under the *Sindell* / *Summers* records-asymmetry logic. The records-asymmetry burden-shift operates independently of the Red Zone presumption of breach established in § 3(b); the breach presumption operates on the breach element, while the records-asymmetry burden-shift operates on the causation element. Where the operator fails to produce the records uniquely within its possession, the trier of fact may draw an adverse inference under the conventional spoliation-and-records-control doctrines.

(d) *Foreseeability.* Foreseeability of the harm is established where the operator has patented or otherwise documented the targeting mechanism by which the harm-causation chain operates, recited the harm-causation chain in regulatory or commercial filings, or deployed the mechanism in jurisdictions where the regulatory regime requires the protective deployment the operator has chosen to withhold in the jurisdiction at issue.

(e) *Continuing tort.* Sustained Actionable Behavioral Conditioning is treated as a continuing tort. Each instance of the operator's targeting conduct constitutes a renewed wrong; the limitations period runs from the last instance of the operator's conduct, not from the first.

§ 5. REMEDIES

(a) *Injunctive relief.* A court may enjoin the operator from continued Actionable Behavioral Conditioning. Injunctive relief includes without limitation: prohibition of crisis-state targeting; friction parity (cognitive and temporal barriers to depositing funds mirror those for withdrawing them); cooling-off mechanisms; affirmative customer-interaction protocols comparable to those the operator deploys in regulatory regimes that require them; and prohibition of the targeting practices the operator has deployed in the jurisdiction at issue.

(b) *Compensatory damages.* The user may recover net losses during the exploitation period, including all amounts wagered through the operator's platform during the period in which the operator engaged in Actionable Behavioral Conditioning toward the user, minus any amounts the user received from the operator during that period.

(c) *Disgorgement.* The user may recover the operator's profits attributable to the user during the exploitation period. The disgorgement remedy is calculated by reference to the operator's own revenue accounting, including the customer-data infrastructure and per-customer "lifetime value" calculation the operator's regulatory and investor disclosures document.

(d) *Punitive damages.* Punitive damages are available where the operator's conduct falls within the Red Zone of the Matrix as defined in Appendix B's classification table. The factfinder may award punitive damages in an amount that reflects the operator's wealth, the egregiousness of the conduct, the foreseeability of the harm under § 4(d), and the deterrent effect required to prevent the operator's continued deployment of the conduct.

§ 6. DEFENSES

(a) *Operator's affirmative defenses.* The operator may assert in defense, in addition to the conventional defenses available under the relevant tort categories: (1) that the operator's conduct fell within the Green Zone of the Matrix as defined in Appendix B's classification table and § 3(c) above; (2) that the operator deployed protective intervention in response to the user's detected vulnerability rather than commercial targeting; (3) that the user's harm was caused by an intervening cause that supplied the proximate causation; and (4) the conventional informed-consent defense under Restatement (Second) of Torts §§ 496A-G, recognizing that informed consent requires consent to the specific risk assumed and that the operator's targeting architecture is not a risk a user can assume without disclosure.

(b) *Defenses unavailable.* The following defenses are unavailable to the operator: (1) the defense that the operator's conduct constituted protected commercial speech under the First Amendment, where the operator's conduct constituted Actionable Behavioral Conditioning directed at users the operator had identified as vulnerable through behavioral telemetry; (2) the defense that the operator's conduct was immunized under § 230 of the Communications Decency Act, where the operator's conduct was the operator's own design and targeting conduct rather than the publication of third-party content; and (3) the defense that the user assumed the risk of the operator's targeting conduct, where the operator did not disclose to the user the targeting architecture under which the operator operates.

APPENDIX B: THE MANIPULATION-VULNERABILITY MATRIX — REFERENCE TABLES

The Manipulation-Vulnerability Matrix supplies the rule of decision that Appendix A § 3 incorporates by reference. The tables below set out (1) the zone classification across the manipulation-intensity and vulnerability-severity axes; (2) the indicator inventories that operationalize each axis; and (3) the rebuttal architecture for each zone classification.

TABLE B.1 — ZONE CLASSIFICATION

The Matrix is a 4 × 4 grid keyed to the manipulation-intensity definitions in Appendix A § 1(d) and the vulnerability indicators in Appendix A § 1(c). Zone classifications are read row-by-column.

	Low Vulnerability	**Medium Vulnerability**	**High Vulnerability**	**Extreme Vulnerability**
Low Manipulation	Green	Green	Green	Green
Medium Manipulation	Green	Yellow	Yellow	Red
High Manipulation	Yellow	Yellow	Red	Red
Extreme Manipulation	Yellow	Red	Red	Red

The Green Zone classification on the Low Manipulation row is a structural feature of the Matrix: where the operator's conduct does not exploit detected vulnerability, the user's vulnerability state does not by itself generate liability. The Green Zone classification at High Vulnerability × Low Manipulation reflects the framework's autonomy-respecting posture: a user who chooses to gamble while exhibiting vulnerability indicators does not generate operator liability where the operator deploys ordinary commercial conduct that does not weaponize the detected vulnerability. The Red Zone is concentrated in the high-manipulation / high-vulnerability quadrant where the operator's conduct calibrates targeting intensity to detected vulnerability — the configuration the framework identifies as Actionable Behavioral Conditioning.

TABLE B.2 — VULNERABILITY INDICATORS

The vulnerability axis is operationalized by indicators observable in the operator's own behavioral telemetry. The Matrix does not require the operator to deploy any indicator the operator does not already collect; the indicators below are the categories the operator's own systems already classify, the operator's own filings already disclose, and the operator's own product architecture already targets.

Severity	Indicators
Low	Recreational engagement; stable wager amounts; activity within ordinary day-and-evening windows; no loss-chase signature; no same-session redepositing; no platform-classified risk indicators present.
Medium	Above-average frequency or stake; presence of one or more loss-chase episodes without sustained pattern; intermittent late-night activity; occasional same-session redepositing; platform-classified VIP-prospect or high-engagement indicators present without crisis indicators.
High	Sustained loss-chase signature (meaningful proportion of bets within thirty minutes of a loss); stake escalation without offsetting account growth; pattern of late-night session concentration; recurrent same-session redepositing; platform-classified high-value or VIP-active status; declining win rate sustained over time.
Extreme	Crisis-state indicators: loss-chase signature combined with same-session redepositing in the late-night window; post-self-exclusion contact attempts or successful re-onboarding within a self-exclusion period; platform's own VIP-host engagement protocols invoked in response to a session the operator's systems flagged as elevated-risk; sustained sessions exceeding the operator's own session-length thresholds for risk classification.

The indicators are illustrative rather than exhaustive. The operator's own classification taxonomy — including the specific tier names, threshold values, and triggering criteria the operator's behavioral-detection systems apply — is the controlling source for any individual case. The Matrix's indicators track the categories the platforms have publicly described in regulatory filings, in App Store privacy declarations, and in the patent-and-product literature the framework draws on; the operator's internal taxonomy will be obtained through discovery and aligned with the Matrix's categories at the merits stage.

TABLE B.3 — MANIPULATION INDICATORS

The manipulation axis is operationalized by indicators observable in the operator's own product architecture and targeting infrastructure. As with the vulnerability axis, the Matrix does not require the operator to deploy any practice the operator does not already deploy; the indicators below are practices the operator's own product is built around, the operator's own investor presentations announce, and the operator's own competitor patents recite.

Intensity	Indicators
Low	Standard commercial communications not calibrated to the recipient's behavioral telemetry; promotional offers deployed on a general-audience basis without behavioral-data-driven selection; UI elements that do not exploit deliberative-friction asymmetries.
Medium	Targeted communications using behavioral data to identify likely-engagement recipients without targeting detected vulnerability states; promotional offers selected by general engagement-prediction algorithms; gamification mechanics deployed on a general-audience basis.

Intensity	Indicators
High	Targeted communications that use behavioral telemetry to identify users in detected vulnerability states; promotional offers deployed to users the operator's classification systems have flagged as risk-elevated; VIP-host promotional contact directed at users showing strong vulnerability signatures; friction asymmetries that calibrate withdrawal friction to detected risk-elevation.
Extreme	Crisis-state targeting: real-time push notifications deployed during sessions the operator's systems have flagged as crisis-state; one-tap wagering mechanisms deployed during crisis-state sessions; profit-boost or loss-chase-encouraging communications deployed during crisis-state sessions; VIP-host promotional contact deployed in response to crisis-state classification; reverse-withdrawal mechanisms triggered on accounts the operator has flagged as elevated-risk; continued targeting of users who have invoked self-exclusion.

TABLE B.4 — REBUTTAL ARCHITECTURE

The Red Zone presumption of breach in Appendix A § 3(b) is rebuttable. The rebuttal architecture identifies the categories of evidence the operator may offer to discharge the presumption.

Rebuttal Category	Description
Protective deployment	Evidence that the operator deployed protective intervention (real-time intervention system, affirmative customer-interaction protocol, friction parity, cooling-off mechanism, or comparable protective architecture) in response to the user's detected vulnerability rather than commercial targeting.
Non-targeting	Evidence that the operator's conduct was deployed without reference to the user's behavioral telemetry — that the communication, offer, or product feature was selected by a process that did not use the user's vulnerability-classification data.
Misclassification	Evidence that the user did not in fact occupy the vulnerability state the operator's systems classified the user as occupying — for example, that the loss-chase signature reflected a one-off pattern rather than a sustained behavioral configuration.
Conduct outside Matrix	Evidence that the operator's conduct, properly classified, falls outside the Red Zone — for example, that the manipulation intensity is properly classified as Medium rather than High, or that the user's vulnerability is properly classified as Medium rather than Extreme.

Each rebuttal category requires the operator to produce records that lie within the operator's exclusive possession. The records-asymmetry burden-shift in Appendix A § 4(c) operates here: where the operator declines to produce the records that would discharge the presumption, the trier of fact may draw an adverse inference. The Red Zone presumption is, in operation, a discovery-forcing mechanism as much as a substantive presumption: it requires the operator to put the records on the table, where the framework's Matrix can be applied to them.

Figure 1: The Manipulation-Vulnerability Matrix

Mapping Platform Conduct Against User Susceptibility to Create Liability Zones

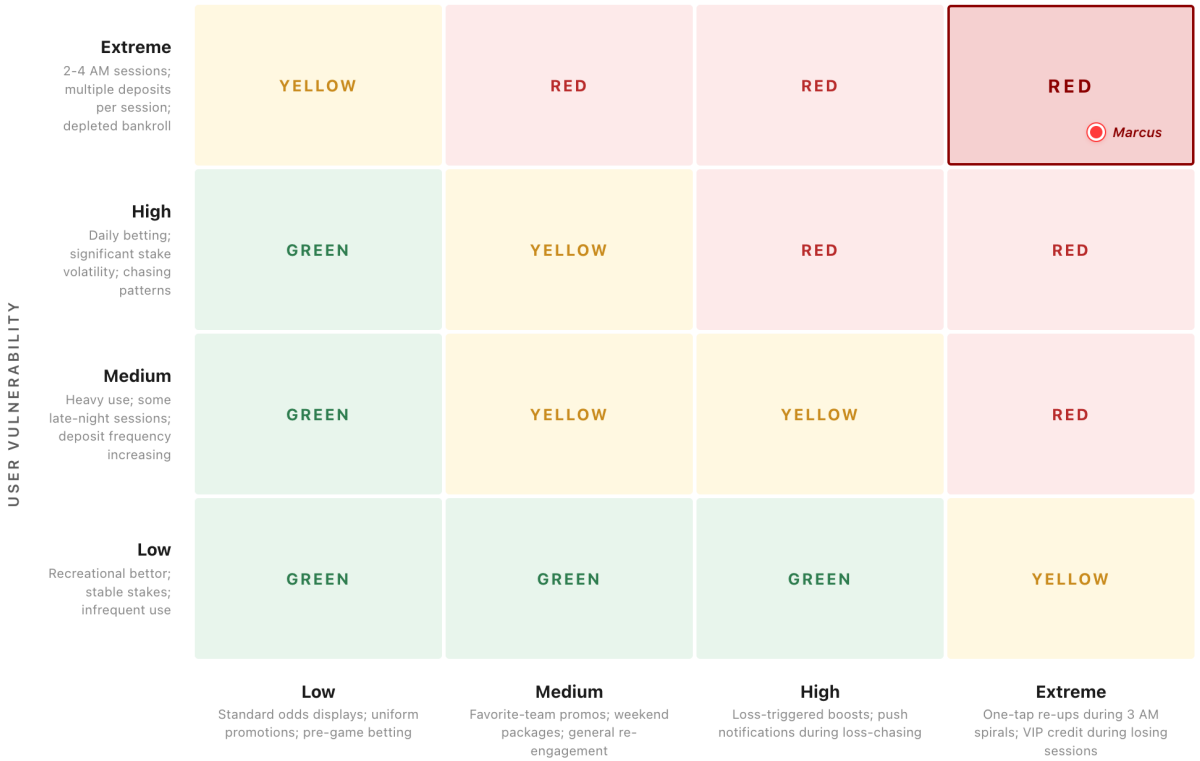

Green Zone — No liability. Standard marketing to capable consumers. **Yellow Zone** — Viable claims. Case-by-case analysis. **Red Zone** — Presumption of breach. Burden shifts to platform.

The Matrix assigns greater weight to manipulation intensity than to vulnerability alone. A highly vulnerable user receiving only standard marketing (Extreme × Low = Yellow) is not exploited — the platform has not weaponized its knowledge. Liability follows from what the platform did with what it knew.

APPENDIX C: EVIDENTIARY INVENTORY AND DISCOVERY ROADMAP

This Appendix maps the framework's empirical and doctrinal claims to the documentary record from which they are drawn and to the records likely to corroborate or amplify them. It is organized around three categories: public records already obtainable without discovery; discovery targets within the platform's exclusive possession; and expert-testimony categories the framework's substantive moves require. The Appendix is calibrated to the *Sage v. DraftKings* posture in the Philadelphia Court of Common Pleas and to the structurally similar successor actions the framework anticipates.

A note on evidentiary taxonomy. The materials catalogued below fall into four categories that should be tracked separately in pleading and briefing.

Party admissions are statements by a defendant entity itself. The DraftKings 10-K and definitive proxy statements are admissible as party admissions against DraftKings; the Flutter Q4 results presentation and Sustainability Report are party admissions against Flutter and FanDuel; the Genius Sports 20-F is a party admission against Genius Sports and supplies admissible foundation evidence on the league commercial architecture.

Industry-feasibility and constructive-knowledge evidence describes materials produced by an industry incumbent or competitor that establish the existence and operational availability of a capability. The Entain patent is the principal example. It does not constitute an admission by DraftKings or FanDuel; it supplies industry capability evidence rather than direct extraction-intent evidence on those defendants.

Public-record allegations are pleaded but unproven. The *Sage*, *Patel*, and *Baltimore* complaints supply allegations subject to proof, not findings.

Discovery roadmaps are the framework's forward-looking identification of internal records likely to corroborate or amplify the public materials. The framework identifies the records and the custodians where it can, marks them as research targets where it cannot, and supplies the doctrinal authority for compelling their production.

C.1 — PUBLIC RECORDS: THE EVIDENTIARY SPINE ALREADY ON THE FILING RECORD

The framework's empirical foundation is principally drawn from materials the platforms have already produced. The list that follows is the evidentiary spine of the framework's substantive claims and the evidentiary spine of the *Sage* complaint as the framework reads the operative pleading.

C.1.1 — Securities Filings and Investor Communications

DraftKings, Inc. Annual Report on Form 10-K for fiscal year ended December 31, 2025 (filed Feb. 2026) (the "DraftKings 10-K"). The filing supplies the SOX-certified concession of the targeting capability the

framework reads as the duty's structural foundation. Two passages in the filing carry particular evidentiary weight: the description at page 5, which (i) confirms "the industry practice of restricting and managing betting limits at the individual customer level" — the SOX-certified admission that the platform operates individual-level behavioral profiling for revenue-protection purposes and the concession that the same capability is structurally available for protective deployment — and (ii) discloses the per-customer "lifetime value" calculation; and the description at page 33 of the customer-data infrastructure underlying both, which together supply the disgorgement-remedy foundation on which the *Comcast*-compliant damages model the framework develops in Part II.D depends.

DraftKings, Inc. Definitive Proxy Statement (Schedule 14A) for fiscal year 2025 (filed 2026). Pages 38-42 disclose the CEO compensation structure: 100% tied to Revenue and Adjusted EBITDA, with no responsible-gambling metric weighting. The disclosure supplies the cosmetic-compliance evidence the framework develops in Part III.B and the scienter foundation the *Philip Morris* tobacco-litigation lineage reads as supporting RICO findings where industry public statements contradict internal corporate compensation structures.

DraftKings, Inc. Investor Day 2026 Presentation, dated March 2, 2026. Pages 29 and 38 disclose the bet-type expansion and live-in-game share metrics: 517 average bet types per game (a fourfold increase since 2022) and 54% live in-game wagering share. The presentation supplies the foundation for the framework's empirical claim about microbet temporal compression and the documentary record of the platform's own announcement of the targeting-architecture metrics the framework's analysis turns on.

Flutter Entertainment plc. Sustainability Report 2024 (Apr. 2025) and Q4 2025 Results Presentation (Feb. 26, 2026). Two passages carry particular evidentiary weight: the Sustainability Report's documentation of the Sportsbet Real-Time Intervention deployment and the 31% intervention-success rate, which is Flutter's own admission that the protective deployment is operationally feasible on the same parent corporate infrastructure that operates FanDuel; and the CEO Peter Jackson letter in the Q4 results presentation describing FanDuel's strategy of "leveraging our scale, proprietary technology, and data advantages" to deliver "smarter personalization" and "richer live engagement," which supplies the cosmetic-compliance contradiction the framework reads as scienter-supporting evidence.

Genius Sports Limited. Annual Report on Form 20-F, filed March 17, 2026. The filing discloses the NFL warrant equity at $0.01 per share representing approximately 6.0% of outstanding shares, and approximately $126.1 million in NFL GGR/NGR commissions for fiscal year 2025. The filing supplies the structural foundation for the league-liability analysis the framework develops in Part V.

C.1.2 — Patent and Industry-Capability Evidence

Entain Mktg. UK Ltd. Patent No. 12,555,436 B2, "System and Method for Encouraging Responsible Gameplay and Preventing Detrimental Gaming Activity," filed August 24, 2022, issued February 17, 2026. Claim 3 recites the four-tier vulnerability classification (low / medium / high / particularly high); Claims 10 and 11 recite the detection signals (anomalous deposits and prolonged sessions); Claims 4-8 recite the calibrated commercial responses at each tier; Claim 14 recites account closure as a remedial tool. The patent supplies industry-feasibility and constructive-knowledge evidence: the detection capability exists, has been built, and has been patented by a competitor in the industry the framework targets. It does not supply party-admission evidence against DraftKings or FanDuel; the patent is held by Entain Mktg. UK Ltd (a subsidiary of Entain plc, parent of BetMGM), not by the *Sage* defendants. What it supplies is the constructive-knowledge foundation underlying the duty analysis in Part II.B: the platforms' inability to credibly claim ignorance of the harm-detection mechanism. The patent also supplies the foreseeability foundation under Appendix A § 4(d): a defendant deploying a behavioral-architecture system comparable to the patented architecture cannot credibly claim the resulting vulnerability-caused harm was an unforeseeable consequence.

C.1.3 — Regulatory Records

United Kingdom Gambling Commission, Licence Conditions and Codes of Practice (LCCP), Social Responsibility Code Provision 3.4.3 (Remote Customer Interaction), paragraphs 7-11 (effective Sept. 12, 2022). The provision requires affirmative customer-interaction protocols on the licensee's UK operations. The licensee's compliance records with the LCCP — including the affirmative-intervention deployment metadata, the customer-interaction logs, and the vulnerability-classification records the LCCP requires the licensee to maintain — are the operating record that establishes the framework's comparative configuration claim: the same parent companies operate the protective deployment in the UK that they do not operate in the United States.

New York State Gaming Commission, Letter from Brian O'Dwyer, Chair, dated February 4, 2026. The letter confirms that no league has requested wager-type restriction in New York, including no request to restrict prop bets on individual player performance, no request to restrict in-play microbetting, and no request to restrict the bet types most directly correlated with problem gambling. The proposition the letter supports is jurisdiction-specific to New York; the framework reads it as supporting the narrow proposition that in a major regulatory market in which the leagues hold the plenary regulatory tool of requesting wager-type restriction, the leagues have not exercised it.

Massachusetts, *In re Robinhood Financial LLC*, Admin. Docket No. E-2020-0047 (Mass. Sec. Div. Jan. 18, 2024) (consent order; $7.5M settlement). The consent order applied fiduciary law to gamification features deployed without suitability analysis — behavioral architecture structurally identical

to what sports betting platforms deploy. The order is the regulatory-record foundation for the framework's extension of *Robinhood*'s principle from investment to gambling. *See also Robinhood Fin. LLC v. Sec'y of the Commonwealth*, 492 Mass. 696, 214 N.E.3d 1058 (2023) (upholding the Secretary's authority to impose fiduciary duties on broker-dealers providing investment advice).

C.1.4 — Pleaded Allegations in Related Litigation

Sage v. DraftKings, Inc., No. 260303384 (Phila. Ct. Com. Pl. filed Mar. 24, 2026). The framework's specimen-case analysis in Part III.C is calibrated to the *Sage* configuration; the framework's substantive moves are pleaded under existing tort categories (strict products liability, design defect, common-law negligence with a heightened duty of care, IIED, aiding-and-abetting against the leagues, UTPCPL, and unjust enrichment). The complaint's allegations are subject to proof at the merits stage; the framework reads the complaint as the operative pleading the framework's analytical architecture supports rather than as a source of established facts.

Patel v. FanDuel, Inc., No. 1:24-cv-07402 (S.D.N.Y. filed Oct. 1, 2024). Page 10 alleges VIP-host communications quoted in Part III.A; pages 10-12 develop the manufactured-intimacy operational manual the framework reads as the VIP-economy evidentiary foundation; pages 13-14 detail the loss-chase escalation pattern the framework reads as documenting the targeting architecture's effect on a known-vulnerable user. Allegations subject to proof.

City of Baltimore v. DraftKings, Inc., No. C-24-CV-25-002683 (Balt. City Cir. Ct. filed Apr. 3, 2025), *removed to* No. 1:25-cv-01487 (D. Md.), *remanded* (Nov. 10, 2025). The complaint — the first public-entity suit against sports-betting platforms post-*Murphy*, brought under the Baltimore Consumer Protection Ordinance — illustrates the public-enforcement pathway the framework develops in Part IV.C as one of the three pathways through which the framework's substantive claims can proceed despite the arbitration architecture. Allegations subject to proof.

C.1.5 — Public Health and Empirical Literature

The framework's public-health-empirical foundation rests on materials that are not party admissions or pleaded allegations but supply the empirical record on which the substantial-factor causation analysis and the design-defect alternative-design analysis depend. The principal sources are catalogued in the body's footnotes; the most heavily relied on for *Sage*-stage briefing are the Lancet Public Health Commission on Gambling (2024); the Public Health Advocacy Institute's litigation-genealogy work descended from the tobacco settlements; the Schüll machine-zone ethnographic tradition; the Clark, Cosenza, and Russell et al. neuroscientific and behavioral-research traditions on near-miss effects, loss-chasing, and microbetting; and the UK Gambling Commission's published research on VIP scheme harm. The empirical literature is the foundation for the expert-testimony categories Section C.3 develops.

C.2 — Discovery Targets: Records Within the Platform's Exclusive Possession

The framework's load-bearing claims turn on records the platform alone possesses. What follows is the discovery roadmap the framework's substantive moves generate, calibrated to the records most likely to exist on the basis of the platform's public disclosures. The precise internal taxonomy and system-name conventions remain discovery targets rather than known artifacts where they cannot be identified with public-record specificity.

C.2.1 — The Customer-Attribute Behavioral Profile

The platform maintains a behavioral profile for each registered user that includes the categories the platform has declared to Apple under the App Store privacy framework, alongside the 186-attribute profile alleged in the *Sage* and *Baltimore* complaints, and that the DraftKings 10-K's customer-data infrastructure and per-customer lifetime-value calculation document at the architectural level. The profile is the dataset the Matrix's vulnerability classification operates on. Discovery targets:

The complete behavioral profile maintained on each named plaintiff, with all attributes the platform's systems compute or store, including session-level wager history, deposit and withdrawal records, push-notification history with timestamps and triggering criteria, VIP-host engagement records, and any vulnerability-classification scores or tier assignments the platform's systems have applied to the named plaintiff.

The data dictionary for the customer-attribute system — the document or schema that defines each attribute, the algorithm that computes it, the data sources from which it is drawn, and the downstream systems that consume it.

The retention policy for the behavioral profile, including the platform's retention practices for raw session telemetry, classification scores, and historical tier assignments. The retention policy is itself a research target where the framework cannot identify it from public sources.

The custodians most likely to hold these records are the platform's data-engineering, customer-analytics, and product-data teams. The framework cannot identify the specific custodians by name from public sources; the *Sage* litigation's discovery process will identify them through the platform's standard custodian-identification protocols.

C.2.2 — Vulnerability Classification Logs

The platform's classification systems assign users to tiers — VIP, high-value, high-engagement, problem-gambling-adjacent, self-excluded, exiled — and the assignments are recorded with timestamps and triggering criteria. The classification logs are the operational record of the Matrix's vulnerability axis as the platform's own systems compute it. Discovery targets:

The complete tier-assignment history for each named plaintiff, with timestamps for each tier change, the triggering criteria the platform's systems applied at each transition, and the operational consequences that followed from each tier classification (promotional eligibility, VIP-host engagement, betting-limit changes, friction-architecture changes).

The classification taxonomy — the platform's internal nomenclature for the tiers and the criteria the platform's systems apply to each tier classification. The taxonomy is a research target where the framework cannot identify it from public sources; the Entain patent's tier nomenclature is industry-

feasibility evidence of the structure such taxonomies typically take, but the *Sage* defendants' specific taxonomy will be identified through discovery.

The relationship between the classification taxonomy and the platform's compensation and incentive structures — the documents and policies that govern how the platform's employees, contractors, and automated systems are evaluated and compensated by reference to the classification outcomes. The compensation documents are the foundation for the scienter argument the framework develops in Part III.B.

C.2.3 — Real-Time Targeting and Notification Records

The platform's targeting architecture generates push notifications, promotional offers, and VIP-host communications calibrated to the user's classification and behavioral state. The records of those communications — including the algorithms that selected them, the criteria that triggered them, and the outcomes the platform's systems recorded for each — are the operational record of the manipulation axis as the platform's own systems compute it. Discovery targets:

The complete log of all targeted communications sent to each named plaintiff, with timestamps, content (text, offer terms, promotional value), the algorithm or rule that selected the communication, the triggering behavioral signature, and the platform's recorded outcome (open rate, click-through, deposit, wager).

The "why-did-this-notification-fire" record — the platform's internal documentation of the rules engines, classification systems, and selection algorithms that produced each notification. Modern platform architectures typically maintain this record at the production-systems level for debugging and analytics purposes.

The A/B testing records — the experiments the platform has run on its targeting architecture, including the experimental designs, the populations targeted, the outcomes measured, and the changes the platform implemented in response to the experimental results. The A/B testing records are the platform's own admission, in operational form, of the experimental refinement of the targeting architecture.

C.2.4 — VIP Host Compensation and Performance Records

The VIP-host architecture is the human face of the targeting system. The compensation and performance-evaluation records for VIP hosts (the formula by which hosts are paid, the relationship of the formula to the assigned bettors' losses, the performance metrics against which hosts are evaluated, the bonus structures and quota systems) are the documentary record of the VIP economy the framework develops in Part III.A. Discovery targets:

- The compensation policies and formulas for VIP hosts, including base salary, commission structures, bonus structures, and any other incentive components.
- The performance-evaluation records for the VIP hosts assigned to each named plaintiff, including the metrics the platform measures, the targets and quotas the platform sets, and the reviews the platform has produced for the assigned hosts.
- The communications training materials and policy documents that govern VIP-host conduct, including the manufactured-intimacy operational guidance the *Patel* complaint's allegations describe.

The custodians most likely to hold these records are the platform's VIP-program management, human-resources, and compensation-administration teams.

C.2.5 — Real-Time Intervention Deployment Metadata in Non-US Markets

The framework's comparative-regulatory configuration claim turns on the platforms' deployment of real-time intervention systems abroad and the absence of comparable deployment domestically. The deployment metadata for those systems (the operational records of the Sportsbet Real-Time Intervention system at Flutter's Australian operations and the LCCP 3.4.3 affirmative-customer-interaction systems at Entain UK, FanDuel UK, and the other UK subsidiaries) is the documentary record of the comparative configuration. Discovery targets:

- The operational architecture of the Sportsbet Real-Time Intervention system, including the detection algorithms, the intervention triggers, the customer-interaction protocols, the outcome metrics, and the deployment history.
- The operational architecture of the LCCP 3.4.3 affirmative-customer-interaction systems at the UK subsidiaries, including the same categories of records.
- The internal documents and decision-records that explain why the protective deployment that operates in Australia and the UK has not been deployed in the United States. The "why-not-here" records are the most evidentiarily significant category in the comparative configuration: the platforms' own internal explanation for the comparative gap is the documentary foundation of the framework's claim that the protective deployment is operationally feasible and that the platforms have chosen not to deploy it where regulation does not require it.

The custodians most likely to hold these records are the platforms' compliance, regulatory-affairs, and corporate-development teams. The records are within the corporate group's possession even where the operating subsidiary is non-US; standard discovery practice in multinational corporate litigation reaches such records through the parent entity's control of the subsidiary documents.

C.2.6 — Internal Comparative-Practice Documents

The framework reads the platforms' own internal analysis of their cross-jurisdictional practices as the documentary record most likely to produce the scienter findings the *Philip Morris* tobacco-litigation lineage produced in the RICO context. Discovery targets:

- Internal memoranda, presentations, or analyses comparing the platforms' US operations to their UK and Australian operations.
- Internal correspondence with regulators, legislators, or industry associations addressing the operational feasibility of real-time intervention or vulnerability-detection systems in the United States.
- Internal discussions of the SAFE Bet Act, the *Sage* litigation, the *Baltimore* litigation, and the regulatory-and-litigation environment the framework documents.

The "what-the-platforms-knew-and-when-they-knew-it" records the *Philip Morris* litigation produced in the tobacco context are the analogous category here. The framework's claim is that the documents likely exist on the platforms' filing systems for the same reasons they existed in the tobacco industry: the platforms' compliance, regulatory-affairs, and corporate-strategy functions cannot operate without producing them.

C.3 — EXPERT TESTIMONY CATEGORIES

The framework's substantive moves are supported by expert testimony in three principal categories. The list below identifies the category, the doctrinal move it supports, and the body of literature on which the testimony will draw.

C.3.1 — Public Health and Behavioral-Research Expert Testimony

Expert testimony from the public-health-research and behavioral-research traditions supplies the foundation for the substantial-factor causation methodology the framework develops in Part II.D and Appendix A § 4(b). The expert categories include: clinical psychologists and addiction researchers qualified to testify about variable-reward schedules, near-miss effects, loss-chase patterns, and the neurological mechanisms of compulsive engagement; epidemiologists qualified to testify about the population-level harm patterns the cross-jurisdictional public-health data document; and machine-zone-tradition researchers qualified to testify about the translation of Schüll's casino-ethnography work to the mobile-platform context.

The testimony supports the framework's substantive moves on causation (the substantial-factor methodology), foreseeability (the harm-causation chain the platforms' detection capability operates on), and the alternative-design analysis (the operational feasibility of the protective deployment).

C.3.2 — Platform-Architecture and Behavioral-Targeting Expert Testimony

Expert testimony from computer scientists, machine-learning researchers, and platform-architecture experts supplies the foundation for the framework's claims about the operational architecture of the targeting systems. The expert categories include: data scientists and machine-learning engineers qualified to testify about the vulnerability-classification algorithms, the targeting-selection systems, and the A/B testing infrastructure the platforms operate; UX researchers qualified to testify about the friction-asymmetry and dark-patterns architecture the platforms deploy; and platform-architecture experts qualified to testify about the operational feasibility and standard industry practice of real-time-intervention systems.

The testimony supports the framework's substantive moves on duty (the platform's operational capacity to detect vulnerability), breach (the platform's targeting of detected vulnerability), and alternative design (the protective deployment the platforms operate elsewhere).

C.3.3 — Comparative-Regulatory and Industry-Practice Expert Testimony

Expert testimony from regulatory and industry-practice experts supplies the foundation for the framework's comparative-regulatory configuration claim. The expert categories include: UK and Australian gambling-regulation experts qualified to testify about the LCCP 3.4.3 and Northern Territory regulatory regimes and the operational practices the regimes require; international comparative-regulatory experts qualified to testify about the cross-jurisdictional differences in the platforms' deployment of protective architecture;

and industry-practice experts qualified to testify about the operational feasibility and standard industry practice in the protective-deployment context.

The testimony supports the framework's substantive moves on duty (the operational feasibility of the protective deployment), foreseeability (the platforms' constructive knowledge of the harm the protective deployment is designed to prevent), and alternative design (the protective deployment the platforms operate elsewhere as the alternative-design exemplar).