

SEPTEMBER 2026 · EDITORIAL POSITION

The House Is Responsible for Its AI.

Using AI to exploit gambling vulnerability is wrong. Protections for at-risk customers must reach the systems deciding whom to pursue.



BY ALAN SHIMEL

CEO, TECHSTRONG GROUP · A FUTURUM COMPANY

THE CENTRAL POSITION

Using AI to exploit vulnerable gamblers is wrong. Protections against harmful targeting must reach the systems selecting customers.

A New York Times investigation into DraftKings brings a long-standing responsibility into sharp focus. Selected gambling regulations expressly protect vulnerable customers. Moving customer selection into AI does not remove applicable obligations. This report examines those rules and what a defensible system looks like.

03	THE OPENING POSITION Exploiting vulnerability is wrong. Oversight must reach AI.
04	THE MEDICAL FRAME The condition behind the customer profile
05	THE HOUSE STUDIES THE PLAYER From advantage play to personalized pursuit
06–07	INSIDE THE MACHINERY What the system is actually being asked to optimize
08	PERMISSION TO OPERATE What the license already requires
09–10	OVERSIGHT MUST REACH THE DECISION Examination belongs inside the regulatory perimeter
11	WHAT CREDIBLE PROTECTION LOOKS LIKE Standards for a defensible commercial system
12	THE HOUSE REMAINS RESPONSIBLE The license names who answers
13	SOURCES & EDITORIAL DISCLOSURES Citations, limits and editorial disclosures

EDITORIAL NOTE This report draws on The New York Times’ reporting, published medical research and public regulatory documents. It does not claim an independent inspection of DraftKings’ systems and does not find that the company violated any specific rule. DraftKings’ response to the Times is quoted where it appears in the newspaper’s reporting.

THE OPENING POSITION

Exploiting vulnerability is wrong. Oversight must reach AI.

Using AI to identify people showing signs of gambling disorder and press them to gamble more is wrong. Applicable protections must govern automated targeting as well as human decisions.

EDITORIAL POSITION

"Moving the decision into an AI system does not move the responsibility. The house holds the license; the house answers for what its systems do."

ALAN SHIMEL, CEO, TECHSTRONG GROUP

Reading the [New York Times investigation into DraftKings](#), I kept coming back to the assignment the company reportedly gave its technology: figure out which customers would respond to promotional offers by gambling and losing more. According to the investigation, employees raised concerns that the resulting models could identify vulnerable gamblers as attractive marketing prospects. Other employees described efforts to develop predictive protection systems that were subsequently shelved.¹

Those allegations deserve scrutiny on their own. They also make a wider point plain to the technology industry: a regulated business does not shed its obligations by moving decisions into an AI system. Regulatory oversight has to follow the decision wherever the operator puts it.

The Times investigation, by Alex Klavens, Walt Bogdanich and Jenny Vrentas, describes work by data analyst Jayden Butts on a casino promotion model in 2023. The model assigned customers scores intended to indicate how their losses would respond to incentives. Butts tested its use with an initial group of about 5,000 players before expanding the experiments. He and other former employees told the newspaper they became concerned about the implications for people struggling with addiction.¹

DraftKings disputes the characterization. In its response to the Times, the company said its promotions were directed toward customers demonstrating sustained engagement, rather than toward customers based on losses. It denied improper targeting.

Its chief responsible gaming officer, Lori Kalani, defended the company's existing monitoring and intervention approach and said predictive risk technology had not been shown to be helpful enough to adopt.¹

That disagreement matters. An investigation based on internal documents and employee accounts warrants attention, but it does not give outside readers access to the full models, their validation results or the customer records necessary to establish a particular violation. This report draws on the Times' reporting, medical research and public regulatory requirements. It does not claim an independent inspection of DraftKings' systems.

My position on the underlying principle is not open to negotiation. Building a system whose purpose is to find people likely to lose, or people already showing the pattern of a gambling disorder, and to press them harder to gamble is wrong. Knowingly exploiting addiction is despicable. Giving the job to software does not reduce the responsibility of the company directing it.

Gambling makes that principle concrete because selected regulations expressly protect vulnerable customers. Applicable obligations do not evaporate when a machine picks the recipient. The decision about whom to approach, when and with what offer belongs to the company. The specific legal duties depend on the product, jurisdiction and facts, regardless of whether a person or a model selected the customer.

For Techstrong's audience this is a question about model objectives, data, deployment controls and corporate authority. It is also a question about people. The customer receiving a carefully timed bonus may be placing another recreational wager. That customer may instead be struggling to stop an activity that is damaging a family, a job or a life. A licensed operator does not get to hide that difference behind a score.

The operator does not need perfect knowledge of every circumstance to owe those customers a duty. It does need a defensible account of how its systems act on evidence of vulnerability. Regulators must be able to examine that account, and the answer must be that protection was built into the decision rather than appended around it.

THE MEDICAL FRAME

The condition behind the customer profile



Conceptual illustration: an ordinary customer weighing a promotional offer. Restraint, not caricature.

Gambling disorder is a recognized behavioral addiction. The [American Psychiatric Association](#) describes a pattern of gambling that causes significant problems or distress. Its diagnostic framework includes unsuccessful efforts to control or stop gambling, returning to recover losses, and jeopardizing important relationships or opportunities. Restlessness or irritability when attempting to cut back can also be part of the condition. These are clinical concerns, not simply descriptions of someone who made an expensive mistake.²

The distinction changes how we should think about responsibility. If impaired control is part of a disorder, a protective approach that depends entirely on the customer exercising control has an obvious limitation. A person may understand the danger and still struggle to stop. The existence of a button, a disclaimer or a telephone number does not settle whether a business is doing enough to avoid making that struggle worse.

The [World Health Organization](#) describes gambling harms extending beyond financial losses to mental health, relationships and other aspects of well-being. It also warns about the limitations of voluntary approaches and the tendency to place responsibility on individuals. Harm can affect family members and can occur before a person meets the threshold for a clinical diagnosis.³

We should be careful with the language. A losing customer is not necessarily addicted. Frequent gambling does not establish a diagnosis. A model-generated risk score cannot substitute for a clinical assessment. These distinctions protect people from careless labeling and protect the credibility of the argument.

They should not become excuses for ignoring warning signs. An operator deciding whether to send an inducement faces a different question from a clinician deciding whether someone meets diagnostic criteria. A company can reasonably withhold a promotion when it sees troubling behavior without declaring that customer medically ill. The evidentiary standard for taking a modest precaution need not be identical to the standard for making a diagnosis.

THE INCENTIVE EFFECT

10%+

In a randomized experiment with 622 men under 40 interested in Euro 2024, inducements increased betting expenditure by more than 10% and nearly halved the likelihood of choosing not to bet. Participants also selected more inferior-value bets. Short experiment; endowed credits; treatment-history exclusions.⁴

The limits are substantial. Participants used endowed credits, treatment-history was excluded, and the regulator-funded research did not test DraftKings, personalized AI or clinical relapse. The narrower conclusion, that incentives can influence expenditure and decisions, is enough to make targeting consequential when an operator uses behavioral records to select recipients whose activity shows possible signs of impaired control.

THE HOUSE STUDIES THE PLAYER

From advantage play to personalized pursuit



Conceptual illustration: the contest between players and the house predates machine learning. Historical, not documentary.

The contest between gamblers and the house long predates machine learning. Players have searched for mathematical advantages, developed betting systems and tried to distinguish a useful pattern from a coincidence. Operators have had their own reasons to watch the people placing bets.

Edward Thorp's *Beat the Dealer*, published in 1962, helped demonstrate how mathematics could produce an advantage in blackjack. [MIT Libraries' historical account](#) describes the later success of MIT blackjack teams and the casino efforts to identify their members. It also makes a distinction worth preserving: the teams' mathematical advantage play was not the same thing as illegal cheating.⁵

That distinction is sometimes lost when the entire relationship is reduced to a cat-and-mouse game. A skilled player exploiting favorable rules, someone colluding to corrupt a game, and a customer losing control of gambling are three different cases. Their conduct may all attract attention, but they call for different responses. Protecting the integrity of a game cannot be used as a blanket description for every effort to protect the operator's margin.

Analytics can serve legitimate purposes. Detecting unusual transactions, investigating possible collusion, examining suspicious betting and checking game integrity are appropriate reasons to analyze activity. The same general capacity to identify

patterns can also support customer protection. What matters is the purpose of the analysis, the quality of the evidence and the action taken as a result.

Digital gambling expands the range of decisions that can be made from customer records. An account can connect a deposit, a session, a sequence of bets, an offer and the response to that offer. In principle, a model can use that sequence to estimate whether another incentive will change the customer's behavior. A marketing system can then act on the estimate repeatedly.

This is more specific than the familiar claim that the house knows the odds. Knowing the mathematics of a game and estimating the circumstances under which a particular person will return to it are different capabilities. The latter brings the individual's behavior into the business model more directly.

The history helps us ask the right question about AI. We do not need to pretend that customer analysis began with modern models, or that every analytical advantage is improper. We need to examine how a more detailed and automated form of analysis is being used against people whose ability to control their gambling may be compromised. Calling that activity personalization provides little reassurance by itself. The underlying technology may be similar; the responsibility depends on the decision it supports.

INSIDE THE MACHINERY

Prediction, objective and the model-governance questions

Responsibility begins by naming what the system is actually being asked to optimize. Predicting a customer's behavior is one decision. Choosing what to do about it is another, and that one belongs to the license holder.

A model might predict that a customer will stop using an app. Another might predict future expenditure. A third might estimate whether a bonus will produce additional wagering beyond what would have happened without it. Those predictions answer different questions. They can lead to very different treatment of the same customer.

The Times describes DraftKings' casino work in terms of customer "elasticity," a score related to the expected response to promotional investment. According to the investigation, the inputs included recent activity, account balances, loss patterns and a separate estimate of the likelihood that a customer would stop gambling. The account raises a specific concern about using those predictions to allocate inducements. It should not be stretched into a claim that every model used by the company has the same purpose.¹

Consider a general targeting workflow, offered here as an illustration rather than a reconstruction of DraftKings' architecture. First, the operator assembles account activity. Data preparation converts events into features: deposits per session, changes in betting frequency, recent expenditure or response to an earlier promotion. A model then produces a score. Business rules translate that score into eligibility for an offer, its value and its timing. A delivery system sends the message, and subsequent behavior becomes part of the record used to evaluate or refine the process.

The objective is crucial. Predicting that a customer is likely to spend money does not establish that a promotion will cause additional spending. An operator seeking to measure the latter needs a credible comparison with what would happen without the offer. Controlled experiments can help estimate that incremental effect, although the results depend on the design, the population and the period measured.

Even a well-designed commercial experiment leaves other questions unanswered. A measured increase in wagering is not a diagnosis. A profitable campaign is not evidence that its recipients were unharmed. If the evaluation only asks whether the promotion increased operator revenue, it cannot answer whether customers experiencing impaired control were disproportionately affected.

That omission can be consequential even without an explicit addiction label in the dataset. Imagine, hypothetically, that a

combination of repeated deposits and escalating play predicts a strong response to bonuses. A model trained to reward that response could favor the pattern without classifying anyone as addicted. Whether the pattern actually overlaps with vulnerability, and how strongly, requires evidence. The absence of a medical label does not settle the question either way.

Nor does removing one obvious feature necessarily remove the concern. Other variables can carry related information. A review that checks only whether the training data contain the words "problem gambler" would miss the substance of the decision. The relevant examination concerns what the system selects, how those selections relate to risk and what controls constrain its actions.

Research shows why behavioral records deserve attention. A [study published online in 2024](#) examined 1,743 online casino customers in Britain, Canada and Spain who completed a problem-gambling questionnaire. Features including frequent deposits within sessions and account depletion were informative in models of gambling risk. However, voluntary participation introduced selection concerns, and the study did not establish clinical diagnosis or demonstrate that deploying its models reduced harm. Several authors disclosed industry relationships, including affiliations with a responsible-gambling technology supplier.⁶

The label being predicted also matters. Another [study used subsequent self-exclusion as an outcome](#) in data covering 25,720 players. Self-exclusion is a useful event to study, but it is an imperfect stand-in for gambling disorder: some people experiencing harm never exclude themselves, while others may exclude for different reasons. Accuracy against that label does not automatically establish accuracy against a medical condition.⁷


For a technology team, these are familiar model-governance questions. What is the target? How representative are the training examples? What errors are being counted? Which outcomes are invisible to the evaluation? In gambling, the answers bear on whether a person receives another invitation to spend.

Most of this discussion concerns predictive machine learning, not a chatbot writing text. A fixed threshold that blocks an offer is not necessarily AI, and a language model drafting copy does not necessarily decide its audience. Oversight has to follow the whole decision process, ordinary software, thresholds, and human approvals included, not just the most fashionable component.

EXHIBIT 1

Where oversight must follow a targeting decision

Five sequential stages of a general targeting workflow, with protective checks placed before selection and again before delivery. Management approval and regulator examination span the workflow. Observed responses feed evaluation, not an unqualified assumption of improved performance.



Illustrative workflow, not a reconstruction of DraftKings' systems. Effective oversight must examine the deployed decision process, including protective controls.

EXHIBIT 2

Four questions an examination should answer

Q1

What outcome is the system rewarded for?

Revenue, additional wagers, incremental losses and engagement are different measures with different consequences.

Q2

Which customers does it select?

Selection must be examined against defensible risk assessment and comparison groups, not asserted case by case.

Q3

Can a risk flag stop the offer?

Protection is only meaningful if it acts close enough to delivery to reflect current information.

Q4

Can the decision be reconstructed?

Model version, feature definitions, campaign rules and timestamps must connect a prediction to an actual message.

PERMISSION TO OPERATE

What the license already requires

Selected jurisdictions already impose safeguards on marketing and protective action. Their scope varies by product and location. Using AI to select recipients does not remove the obligations that apply to an operator.

Jurisdiction and scope	Selected requirement	Relevance to AI targeting
Massachusetts Sports wagering	Restricts vulnerable-bettor targeting and addictive AI-based promotions; requires targeting records.	Explicitly reaches AI-based selection and responsible-gaming analysis.
Virginia Sports betting	Prohibits intentionally using at-risk or problem-bettor characteristics to target potentially vulnerable bettors.	Makes the selection process and evidence of intent consequential.
Great Britain Remote licences covered by SR Code 3.4.3	Requires identification, action and evaluation; strong harm indicators must stop marketing and new bonus uptake.	Connects protection signals to limits on commercial action.

Sources: Massachusetts, 205 CMR 256 and 257; Virginia, 11VAC5-80-150; Great Britain, SR Code 3.4.3. Selected comparison; scope varies by jurisdiction and product.

Massachusetts sports wagering rules prohibit using characteristics of problem or at-risk bettors to target potentially vulnerable bettors with advertisements. They require records of targeting parameters and compliance efforts, with commission access. Virginia separately prohibits intentional targeting of this kind.^{8, 12}

Massachusetts also directly addresses AI. Under 205 CMR 257.02(3)(e), operators may not use customer information to promote wagers or offers through an algorithm, machine learning or similar system known or reasonably expected by the operator or its vendor to make the platform more addictive. Section 257.02(4) requires analysis to develop responsible-gaming interventions and reports to the commission at least every six months.¹³

For covered remote licensees in Great Britain, the Gambling Commission requires monitoring from account opening, action on harm indicators and evaluation of the protective approach. Strong indicators trigger prevention of marketing and new bonus uptake. Automated protective decisions require individual manual review and an opportunity to contest them. Compliance is a licence condition.⁹

Massachusetts also addresses self-excluded customers and disciplinary action. Virginia places assistance-information requirements alongside its targeting restriction. Displaying a help number does not discharge the separate obligation concerning whom a business targets.^{8, 12}

These examples must stay within their boundaries. The American provisions discussed here concern sports wagering. The Times investigation also describes casino activity. Establishing a violation requires matching the product, location, period and facts to the applicable law. Britain provides a useful comparison; its requirements do not govern American operations simply because we find them instructive.

Those limits do not diminish the policy significance. Vulnerability is already inside the framework governing permissible marketing and protective action. The technical translation requires examining features, correlated signals and deployed behavior, not simply pointing to a compliance department.

Statute, agency guidance, voluntary commitments and this report’s recommendations carry different authority. An operator does not establish compliance by saying it did not intend to diagnose addiction: the decision under examination is whether to target an offer. What must be proven depends on the governing rule, including Virginia’s explicit intent language. Regulators need the evidence to conduct that analysis, and operators must cooperate as applicable law requires.

13. <https://www.mass.gov/doc/205-cmr-257-sports-wagering-data-privacy/download>

WHAT THE LICENSE DEMANDS

Oversight must reach the system making the decision.



Conceptual illustration: operators should be able to demonstrate to the relevant regulator how their systems comply with applicable requirements.

The obligations set out in the previous section are not aspirational. They are the conditions of doing business. Public sources establish what the license requires. What they do not establish is how thoroughly regulators have inspected the individual targeting models implicated in the Times investigation. That leaves an unanswered question about the depth of oversight. Where a licensed operator moves consequential customer decisions into an AI system, examination of that system belongs inside the regulatory perimeter, not outside it.

On September 24, the Massachusetts Gaming Commission said it would examine AI use by DraftKings and other licensed betting operators following the Times investigation. DraftKings said it complies with Massachusetts rules and denied using AI to target people based on losses or marketing based on indicators of problem gambling. The inquiry is a concrete oversight step, not a finding of wrongdoing; its scope and results still need examination.¹⁴

A useful examination would begin with the business objective. What does the organization reward the system for producing? A model description saying “engagement” may be accurate but incomplete. The reviewer needs to understand the measured outcome, the time horizon and the decisions attached to the score. Increased visits, additional wagers and incremental losses are different measures, even when they sometimes move together.

Next comes the population the model selects. A reviewer should be able to examine whether customers displaying credible signs of vulnerability receive more frequent or more valuable offers, and why. That requires a defensible method of assessing risk, with its limitations acknowledged. It also requires comparison groups. A high percentage of vulnerable customers in a campaign could

mean something different depending on the underlying customer population and the selection process.

This is where evidence of knowledge becomes relevant. Suppose a company tested a model and discovered that its most profitable audience substantially overlapped with a validated risk measure. Suppose staff then proposed a restriction and management rejected it. Such evidence would raise different questions from a model whose harmful pattern had not been identified. Those are hypothetical distinctions, not findings about DraftKings. They illustrate why records of testing, warnings and approvals can matter as much as a technical description of the model.

The next issue is the connection between protection and delivery. A risk system might flag a customer while a separate marketing platform continues sending previously scheduled offers. A suppression list might be current when a campaign is assembled but stale when it is delivered. An outside agency might receive an audience before an intervention occurs. These are potential failure modes a reviewer should test, rather than assume occurred at any particular company.

An effective control must work at the point where the action can still be stopped. Detecting a risk and recording that someone was notified are intermediate steps. The practical question is whether the relevant offer was withheld, whether the restriction persisted and whether someone could override it. If overrides are allowed, the reason and responsible person should be recorded.

14. <https://www.nytimes.com/2026/09/24/business/massachusetts-draftkings-ai-gambling.html>

Regulatory access must also produce usable evidence. A large export of scores may reveal very little without the model version, feature definitions, campaign rules and timestamps connecting a prediction to an actual message. Conversely, access to source code alone cannot show how customers were treated in production. Oversight needs enough information to trace decisions through the operating system and evaluate their consequences.

That need not mean publishing sensitive customer records or exposing a company's entire proprietary system to competitors. Confidential examination, controlled access and qualified independent assessment are possible approaches. Their suitability depends on the legal framework and the risks involved. The principle is that confidentiality should shape how scrutiny occurs, rather than become a reason it cannot occur.

Changes after deployment deserve attention too. A model approved for one population may behave differently after a new product is introduced or a campaign objective changes. A protective rule may no longer cover the behavior a new model favors. Review should therefore address material changes and ongoing performance, not assume that an initial assessment remains sufficient indefinitely.

The organizational questions are equally important. Who can stop a campaign? Who decides that a protection threshold is too restrictive? Does the person responsible for customer safety have

authority over the system producing the offers? A technically capable team cannot compensate for a management structure that leaves its findings without consequences.

The Times reports that former DraftKings employees raised concerns about safeguards and were told protection would be handled elsewhere in the company. DraftKings disputed the suggestion of tension between its business and responsible-gaming objectives. That disagreement points to a question worth asking across the sector: what evidence shows that the handoff from concern to action actually works?¹

The answer should be more substantial than a diagram assigning responsibility to another department. It should include defined authority, traceable decisions, tested controls and outcomes that can be examined. These are proposed standards for meaningful oversight, not a claim that every jurisdiction currently mandates the same architecture.

A regulator could review the language of every advertisement and still learn little about why one customer received forty offers while another received none. Content review remains useful. Customer selection, sequencing and timing add another dimension. If those decisions determine who is exposed to inducements, they belong within the scope of oversight aimed at preventing harmful targeting.

"A warning that leaves the promotional pipeline untouched is a weak response to a problem inside that pipeline."

WHAT CREDIBLE PROTECTION LOOKS LIKE

Standards for a defensible commercial system

Gambling companies already offer tools that deserve fair consideration. Their availability alone does not establish how widely they are used, how reliably they operate or how much harm they prevent.

Conceptual illustration: examination must reach the deployed decision process, not stop at the label the vendor puts on the box.



[FanDuel's published sportsbook materials](#) describe deposit, wager, time and loss limits, timeouts and self-exclusion. They distinguish alerts from limits that restrict activity and explain delays before some restrictions can be relaxed. Those are concrete functions. Their availability alone does not establish how widely they are used, how reliably they operate or how much harm they prevent.¹⁰

DraftKings told the Times that it monitors more than two dozen indicators and uses escalating interventions, including messages, educational material and, in some cases, account closure. Kalani said the company does not send promotions to customers who have been flagged or taken themselves off the platform. The operational effectiveness of those protections is precisely the kind of question that deserves independent examination.¹

There is evidence supporting proactive intervention. An [analysis of a randomized Norsk Tipping trial](#) involving 3,009 high-expenditure customers compared telephone contact, letters and no contact. Results varied by gambling subtype, with telephone contact performing better for several groups. The outcome concerned theoretical losses recorded by the operator, rather than clinical recovery or a complete picture of gambling elsewhere. Operator involvement and researcher industry relationships also warrant disclosure. The study supports investigating active outreach; it does not validate a particular AI system.¹¹

Predictive protection deserves the same seriousness about evidence that we should demand of marketing. A vendor's accuracy claim is insufficient without knowing the outcome being predicted, the population tested and the consequences of errors. A system that misses many vulnerable customers may offer false reassurance. A system that flags too broadly can burden customers, overwhelm staff and undermine confidence in legitimate interventions.

Responses should be proportionate. Withholding a bonus, requesting a human review and closing an account have different consequences. A company should explain why a particular signal justifies a particular action. Protection can include privacy

safeguards and meaningful review without requiring every uncertain signal to be ignored until a crisis becomes unmistakable.

My recommendations start with evaluating the commercial system itself. Operators should test whether their targeting disproportionately selects people showing credible signs of harm. They should investigate that overlap before treating strong promotional response as an uncomplicated business success. Tests should include the deployed workflow, because campaign rules can change the practical effect of a model.

Protective decisions should then constrain delivery. An operator should be able to demonstrate that relevant risk flags suppress offers across the channels and providers it controls, with checks close enough to delivery to reflect current information. Exceptions should have clear ownership and a recorded justification. A warning that leaves the promotional pipeline untouched is a weak response to a problem inside that pipeline.

Evaluation should follow the customer's behavior over an appropriate period. Messages sent, videos viewed and questionnaires completed are process measures; they do not establish that harm declined. Outcome evaluation is harder, especially when activity moves elsewhere, but that difficulty calls for careful methods rather than substituting easy metrics for meaningful ones.

Protection also needs authority and resources. Management should be accountable for material deployment decisions, unresolved warnings and the evidence used to accept risk. Outside suppliers should provide enough documentation and access for the operator to meet its obligations, because responsibility cannot disappear into a contract or a departmental boundary. None of this requires a regulator to declare one machine-learning method superior: the standard should be whether the overall system provides credible protection while remaining open to examination and correction.

THE CONCLUSION

The house remains responsible for its AI

The next promotion arrives as an email, a notification or an offer inside an app. Behind it sits a chain of data collection, prediction and commercial judgment. The customer sees the invitation. The operator chose the system. Applicable protections must follow that decision into the technology.

That is why this belongs in a technology publication. Decisions about target variables, evaluation metrics, suppression rules and deployment authority determine whether software honors a protective obligation or works against it. Those are choices made by organizations. They cannot be laundered as an algorithm doing what algorithms do.

Knowingly exploiting gambling addiction is wrong. Using signs of impaired control to press someone to gamble more deserves scrutiny. The cited rules address vulnerable bettors, not every customer expected to lose. A particular violation depends on the rule and the facts. Operators should demonstrate that their targeting systems honor applicable protections, with controls that can stop harmful pursuit.

None of this requires proving in advance that every recipient meets diagnostic criteria, and none of it authorizes declaring a specific company guilty of a specific violation without the facts. Serious scrutiny can be forceful and precise at the same time. The Times

investigation gives regulators, investors, technology leaders and the public substantial reasons to ask how these systems select customers, what the operator does with the answers and how a regulator would ever know.

The expectation should be clear to anyone building, buying or overseeing these systems. A licensed business must explain how its technology fits within its obligations to customers. Where applicable rules restrict targeting vulnerable bettors, the systems selecting recipients must be open to meaningful examination.

AI makes a business more capable of finding the people it wants to reach. That capability does not narrow accountability; it widens it. Permission to operate a gambling business carries responsibilities that reach every layer of the technology deciding whom to pursue. Those responsibilities were there before the model was trained. They remain in force after it is deployed.

THE CLOSING POSITION

The house remains responsible for its AI.

ALAN SHIMEL · TECHSTRONG GROUP

SOURCES

Citations, limits and editorial disclosures

Live hyperlinks are preserved throughout the report. Each numbered citation ties to the corresponding claim; the medical, research and regulatory sources are separated from company statements and the Times investigation. This report does not find that DraftKings violated any specific rule and does not include original inspection of the company's systems.

REFERENCES

1. Klavens, Bogdanich and Vrentas. How DraftKings Uses A.I. to Target the Gamblers Likeliest to Lose. The New York Times, September 19, 2026.
<https://www.nytimes.com/2026/09/19/business/draftkings-ai.html>
2. American Psychiatric Association. What Is Gambling Disorder?
<https://www.psychiatry.org/patients-families/gambling-disorder/what-is-gambling-disorder>
3. World Health Organization. Gambling: fact sheet.
<https://www.who.int/news-room/fact-sheets/detail/gambling>
4. O Ceallaigh et al. Effects of inducements on sports gambling and decision errors: an experimental study. Journal of Behavioral Addictions (2025).
<https://doi.org/10.1556/2006.2025.00056>
5. MIT Libraries. Year 124 - 1984: The Mathematics of Gambling by Edward O. Thorp.
<https://libraries.mit.edu/150books/2011/05/10/1984/>
6. Hopfgartner et al. Using Artificial Intelligence Algorithms to Predict Self-Reported Problem Gambling Among Online Casino Gamblers from Different Countries Using Account-Based Player Data. Online 2024.
<https://link.springer.com/article/10.1007/s11469-024-01312-1>
7. Hopfgartner et al. Predicting self-exclusion among online gamblers: An empirical real-world study. Journal of Gambling Studies. Online 2022.
<https://link.springer.com/article/10.1007/s10899-022-10149-z>
8. Massachusetts Gaming Commission. 205 CMR 256: Sports Wagering Advertising. March 27, 2023 clean copy.
<https://massgaming.com/wp-content/uploads/205-CMR-256-3.27.23-clean-copy.pdf>
9. Gambling Commission (Great Britain). Social Responsibility Code 3.4.3: Remote customer interaction.
<https://www.gamblingcommission.gov.uk/licensees-and-businesses/lccp/condition/3-4-3-remote-customer-interaction>
10. FanDuel. Play with a Plan: Sportsbook tools, limits, timeouts and self-exclusion.
<https://www.fanduel.com/playwithaplan/sportsbook>
11. Jonsson et al. Reaching Out to Big Losers: How Different Types of Gamblers are Affected by a Brief Motivational Contact Initiated by the Gambling Provider.
<https://pubmed.ncbi.nlm.nih.gov/32955693/>
12. Virginia Administrative Code. 11VAC5-80-150: Advertising in general.
<https://law.lis.virginia.gov/admincode/title11/agency5/chapter80/section150/>
13. Massachusetts. 205 CMR 257: Sports Wagering Data Privacy, sections 257.02(3)(e) and 257.02(4).
<https://www.mass.gov/doc/205-cmr-257-sports-wagering-data-privacy/download>
14. The New York Times. Massachusetts Is Investigating Gambling Companies' Use of A.I. September 24, 2026.
<https://www.nytimes.com/2026/09/24/business/massachusetts-draftkings-ai-gambling.html>

EDITORIAL DISCLOSURES

Reporting and company response. Company-specific reporting relies on the September 19 and September 24, 2026 New York Times articles. Responses are those published by the Times; no fresh outreach or independent inspection of DraftKings' systems was conducted. DraftKings disputes that its promotions target customers based on losses and defends its responsible-gaming approach.

Regulations and inquiry. Duties vary by jurisdiction, product and period. The cited American rules concern sports wagering, not a universal ban on targeting every customer expected to lose. Massachusetts announced an inquiry on September 24; the announcement is not a finding of wrongdoing. Its eventual scope, results and any enforcement conclusions remain to be examined.

Medical and research limits. Questionnaire-based risk, self-exclusion, expenditure and clinical recovery are distinct outcomes. Study populations, incentive designs and industry relationships are disclosed where material.

Illustrations and editorial status. All illustrations are AI-generated conceptual imagery, not documentary evidence, and do not depict any real facility, employee or executive. This Techstrong special report is an editorial analysis, not a finding of legal liability. Updated September 25, 2026.