

Behavioral Biases in Football Lottery Decision-Making: An Integrated Prospect Theory and Machine Learning Approach

Finn Erickson

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

finn1991@uab.edu

Devidor J. Carpenter

Department of Computer Science, University of Houston, Houston, TX, USA.

carpenter529@uh.edu

Abstract

The football lottery represents a complex socio-technical domain in which millions of participants routinely make decisions under uncertainty, often exhibiting systematic departures from rationality. This paper presents an integrated analytical framework that synthesizes prospect theory with modern machine learning architectures to examine, detect, and contextualize behavioral biases in large-scale lottery decision-making environments. Rather than proposing a single predictive model, the study adopts a systems-level perspective to interrogate the structural trade-offs, infrastructural requirements, governance mechanisms, and ethical implications inherent in deploying such integrated systems. We discuss how core prospect theory constructs such as reference dependence, loss aversion, and probability weighting can be operationalized within machine learning pipelines that combine real-time transactional data, contextual match information, and user interaction streams. The analysis is organized around a set of interconnected architectural axes, including the tension between interpretability and predictive fidelity, the design of streaming data fabrics to capture time-varying behavioral signatures, the maintenance of algorithmic fairness across heterogeneous demographic segments, the robustness of inference under adversarial contamination and concept drift, and the sustainability of computational workloads at population scale. Policy considerations regarding consumer protection, responsible gambling mandates, and transnational regulatory fragmentation are examined as constitutive elements of the system design space rather than external constraints. Through cross-domain comparisons with financial trading and insurance underwriting, we identify recurring patterns and critical divergences that inform forward-looking deployment strategies. The paper contributes a conceptual blueprint for the responsible, transparent, and resilient integration of behavioral economic theories into the automated decision-support infrastructures that increasingly mediate collective gambling behavior.

Keywords

football lottery, behavioral biases, prospect theory, machine learning, system architecture, algorithmic fairness, responsible gambling, infrastructure governance.

1. Introduction

The football lottery, encompassing a variety of pari-mutuel, fixed-odds, and hybrid wagering formats tied to match outcomes, ranks among the most globally pervasive forms of regulated gambling. Its decision-making environment is characterized by high-dimensional stochasticity, real-time information flow, and pronounced emotional engagement, all of which create fertile ground for the emergence of systematic behavioral biases. For decades, the study of such biases drew primarily upon the cognitive psychology and behavioral economics traditions, with prospect theory providing the foundational explanatory apparatus for understanding deviations from expected utility maximization. More recently, the advent of large-scale machine learning has introduced powerful new capabilities for detecting, modeling, and potentially influencing choice patterns in near-real time. However, the intersection of these two intellectual streams remains insufficiently theorized at the systemic level, where architectures, governance frameworks, and socio-technical feedback loops become central to any deployment.

This paper argues that integrating prospect theory with machine learning in the football lottery domain must be treated as a systems design problem of considerable complexity, not merely as a computational exercise in feature engineering or model selection. Viewing the lottery ecosystem as a distributed information processing architecture, we foreground the structural trade-offs that arise when behavioral constructs are embedded within automated pipelines operating at population scale. These trade-offs span multiple layers: the data infrastructure that must ingest streaming wager logs, odds movements, and match-event feeds; the modeling layer that must translate theoretical constructs like loss aversion into computable representations without succumbing to interpretability erosion; the governance layer that must enforce fairness constraints and responsible gambling obligations under heterogeneous jurisdictional regimes; and the sustainability layer that must reconcile computational demands with environmental and economic viability. The paper adopts a cross-domain comparative lens, drawing parallels with algorithmic decision systems in financial markets and insurance underwriting, where similar tensions between behavioral modeling, real-time adaptation, and regulatory compliance have been extensively catalogued.

Our contribution is an integrative architectural synthesis rather than a single empirical study. We articulate the conceptual building blocks of a system capable of ingesting rich behavioral traces, diagnosing bias signatures such as the favorite-longshot bias, the disposition effect, and the gambler's fallacy, and feeding these diagnoses into adaptive mechanisms that support harm minimization and market integrity. The analysis is structured so as to provide a roadmap for policymakers, platform architects, and researchers who must navigate the competing demands of innovation and consumer protection in an increasingly digitized gambling landscape.

2. Behavioral Foundations and the Applicability of Prospect Theory

Understanding decision-making in the football lottery requires a departure from the rational-agent assumptions that underpin conventional financial pricing models. Prospect theory, as originally formulated by Kahneman and Tversky, posits that individuals evaluate outcomes relative to a reference point, exhibit loss aversion such that losses loom larger than equivalent gains, and apply nonlinear probability weighting that overweights low probabilities and underweights high probabilities [1]. Subsequent generalizations, notably cumulative prospect theory, refined the treatment of probabilistic sensitivity and rank-dependent weighting, offering a more robust descriptive account of choices under risk [5]. These foundational

contributions have spawned an extensive empirical literature demonstrating their relevance across diverse gambling contexts, from racetrack betting to state lotteries [2].

In the specific ecology of the football lottery, several classes of behavioral bias can be mapped onto prospect-theoretic constructs with significant structural implications. The favorite-longshot bias, wherein bettors systematically overvalue low-probability high-return outcomes due to probability weighting distortions, has been widely documented in betting markets and aligns with the overweighting of small probabilities predicted by the theory. Similarly, the disposition effect, habitually observed in financial trading but also present in sequential lottery decisions, reflects a tendency to cash out winning positions prematurely while holding onto losing wagers in the hope of recovery, a pattern intimately tied to reference-point adaptation and loss aversion in sequences of gambles. The gambler's fallacy, the erroneous belief that a streak of one outcome makes the opposite outcome due, can be reinterpreted through the lens of representativeness heuristics interacting with probability weighting, where bettors mentally substitute stochastic processes with deterministic balancing expectations. Sentiment-driven betting on popular clubs, often amplified by media narratives, introduces another layer of bias where reference points become socially anchored to group identity and emotional affiliation rather than probabilistic assessment [6].

What makes the football lottery an especially challenging domain for behavioral modeling is the temporal structure of decision-making. Unlike a simple lottery ticket purchased once, the football lottery often unfolds across a weekend of fixtures, with opportunities for in-play wagering, cash-out mechanisms, and portfolio-style combinations. Decision-makers receive a continuous stream of partial outcomes, shifting their reference points dynamically and creating path-dependent patterns of risk appetite. Capturing these dynamics requires modeling architectures that explicitly incorporate time, state, and experience in ways that static utility functions cannot. This temporal richness simultaneously increases the empirical identifiability of prospect-theoretic parameters and amplifies the risk of ecological feedback loops, where the system's own behavioral predictions might induce behavioral change if used to nudge or intervene. A systems perspective must therefore attend not only to the accuracy of bias detection but to the recursive consequences of deploying such detection within a live socio-technical environment.

3. Machine Learning Architectures for Behavioral Diagnostics

Translating prospect-theoretic constructs into operational machine learning workflows presents a non-trivial architectural challenge. The core tension lies between the desire to preserve fidelity to the original psychological constructs and the pragmatic demands of large-scale, low-latency inference. A purely theory-driven approach would encode explicit parametric forms for value functions and weighting functions and estimate their parameters from data. However, such parametric specifications risk imposing distributional assumptions that fail to capture the heterogeneity and context-dependent nature of real wager data. Conversely, a purely data-driven deep learning model might achieve superior predictive performance on observational data but yield representations that are disconnected from the behavioral mechanisms of interest, making it difficult to certify that detected biases align with the intended constructs or to generate actionable explanations for regulators and platform operators.

The architectural design space can be organized along several dimensions. The first dimension concerns the representation of behavioral state. Temporal sequences of wager amounts, odds selection, cash-out decisions, and platform navigation events can be modeled

using recurrent neural architectures, temporal convolutional networks, or Transformer-based sequence models that learn distributed representations of user trajectories. These representations can then be aligned with prospect-theoretic dimensions through multi-task learning objectives that jointly predict behavioral outcomes and reconstruct interpretable latent variables corresponding to reference point, loss domain, and gain domain. Such an approach has been explored in financial behavior modeling, where deep architectures trained on transactional data are constrained to produce factorizations that correspond to known cognitive biases. A second dimension involves the choice of inference granularity. Population-level models estimate aggregate bias parameters and can inform market-wide regulatory interventions, while individual-level models personalize diagnostics and could drive targeted responsible-gambling nudges. The trade-off between these levels is not merely statistical; it implicates privacy, fairness, and the ethical legitimacy of personalization.

A further architectural consideration is the integration of counterfactual reasoning. Detecting a behavioral bias requires not only observing what a bettor did but also estimating what a normatively calibrated agent would have done under identical information conditions. This necessitates a normative baseline model of fair odds and value. Sports forecasting models based on Poisson goal distributions, Elo ratings, or ensemble machine learning predictors [3] can serve as reference engines against which behavioral deviations are measured. The systems challenge consists in maintaining consistency between the normative model and the behavioral model, particularly when both are updated online as new match data arrive. As bets are settled and new information becomes available, the normative yardstick can shift, retroactively altering the diagnosis of past biases. This temporal entanglement between measurement and reality demands careful versioning, audit logging, and dispute resolution mechanisms within the system infrastructure.

4. Data Infrastructure and Systemic Integration

The effective integration of prospect theory and machine learning for football lottery analytics is contingent upon a data infrastructure capable of ingesting, processing, and storing heterogeneous, high-velocity streams with clearly defined provenance. The fundamental data types include wager transaction logs from lottery operators and betting platforms, real-time odds feeds from multiple market makers, match event streams capturing goals, cards, substitutions, and injury news, and increasingly, social media sentiment signals that correlate with betting volume shifts. The data fabric must reconcile the tension between the low-latency requirements of operational bias detection, which might seek to flag problematic gambling patterns within a session, and the batch-oriented analytical needs of regulatory reporting, model retraining, and longitudinal behavioral studies. Architectures built on distributed streaming platforms for log processing and in-memory cluster computing for iterative analytics have become standard templates for such cyber-physical financial systems [7].

The design of data schemas and ontologies in this context is far from a trivial technical task; it is a deeply conceptual exercise that shapes what behavioral phenomena can be observed and theorized. Representing a wager as an atomic event with attributes such as stake, odds, selection, and timestamp is necessary but insufficient for capturing the reference-dependent utility calculus. The system must model the temporal relationship between placement and settlement, the user's contemporaneous portfolio state, the sequence of intervening informational shocks, and the cash-out opportunities that were available but declined. This requires an event-sourced architecture in which all state transitions are logged as an append-only sequence, enabling retroactive reconstruction of the information set visible to the bettor

at the moment of decision. Furthermore, the integration of lottery-specific data models must account for the unique combinatorial structures of football pools, where tickets often involve selections across multiple matches, introducing dependencies that complicate the attribution of a decision to a specific match-outcome belief.

The emergence of novel gameplay formats in the lottery ecosystem adds further layers of infrastructural complexity. Recent explorations have proposed interactive guessing modes that blend traditional lottery mechanics with dynamic feedback and gamified elements, potentially altering the behavioral signature of participants in ways that existing models may not capture [11]. When such innovations are deployed at scale, the data infrastructure must be sufficiently flexible to accommodate new event types and interaction modalities without requiring disruptive schema migrations. The adoption of semi-structured, schema-on-read paradigms based on document stores and distributed file systems provides a degree of agility, but at the cost of increased governance complexity in ensuring data quality, consistency, and compliance with data protection regulations such as the General Data Protection Regulation (GDPR) [16]. The systems architect must balance these forces, designing for evolvability while maintaining the strict integrity guarantees required for financial and behavioral auditing.

Finally, data sharing across organizational boundaries, between state lottery operators, commercial bookmakers, technology vendors, and academic researchers, raises profound governance challenges. Behavioral data from football lottery transactions are richly personal and can inadvertently reveal sensitive attributes such as financial distress, compulsive tendencies, and social network ties. Federated learning architectures, in which model updates are shared without exposing raw transaction records, and differential privacy mechanisms that inject calibrated noise into aggregate statistics, present promising technological pathways. However, their deployment must be embedded within a multi-stakeholder governance framework that specifies data usage purposes, retention periods, and secondary analysis constraints, aligning technical controls with evolving regulatory expectations and societal norms.

5. Governance, Fairness, and Ethical Deployment

The deployment of machine learning systems that embed behavioral economic theories into gambling platforms is not a value-neutral technical exercise; it directly engages with ethical and regulatory frameworks that govern the permissible limits of influence, exploitation, and protection. A central tension arises from the dual-use nature of behavioral diagnostics. The same infrastructure that enables an operator to identify a bettor exhibiting loss-chasing behavior and deliver a responsible-gambling message can, with minimal reconfiguration, be employed to intensify engagement by offering targeted incentives at precisely the moment of heightened vulnerability. The system architecture must therefore encode governance constraints as first-class design elements, not as after-the-fact compliance patches.

Algorithmic fairness in this domain takes on multiple intersecting dimensions. Distributional fairness concerns whether the accuracy of behavioral models and the delivery of interventions are equitable across demographic groups, geographic regions, and socioeconomic strata. Empirical audits of machine learning models in consumer finance and insurance have consistently demonstrated that models trained to maximize aggregate predictive performance can exhibit disparate performance across protected subgroups [8]. In the football lottery context, a model that systematically underestimates loss aversion for certain communities might fail to trigger timely protective interventions, while overestimating it for others could result in unjustified account restrictions that deny individuals the autonomy to participate in a

lawful activity. Procedural fairness demands that individuals have access to meaningful explanations for automated decisions that affect their accounts, a requirement that intersects with the interpretability dimension of model architecture and motivates the use of inherently interpretable models or post-hoc explanation layers where complex deep learning architectures are employed [15].

The regulatory landscape for gambling systems is highly fragmented across jurisdictions, creating significant architectural challenges for platforms that operate transnationally. The European Union's GDPR imposes stringent requirements on automated decision-making and profiling, including the right to human intervention [16]. Concurrently, gambling-specific regulations in markets such as the United Kingdom mandate affordability checks and interaction monitoring to prevent harm, while other jurisdictions emphasize consumer freedom and tax revenue generation. A well-architected system must externalize jurisdictional rule sets as parameterized policy engines that can evaluate decision logic against the relevant regulatory context without requiring divergent codebases. This is particularly critical when prospect-theoretic constructs such as reference points are operationalized, as their computational definition can directly affect whether a given pattern is classified as normal recreational behavior or a marker of emerging gambling disorder [9]. Multi-jurisdictional deployment thus demands a policy-as-code paradigm where behavioral rules are versioned, tested, and audited with the same rigor as core modeling code.

Beyond regulatory compliance, the ethical dimension of these systems encompasses the broader obligation to minimize harm without paternalistic overreach. Overly aggressive interventions, such as automated stake limits triggered by imperfect bias detections, can erode user trust and push players toward unregulated black markets, paradoxically increasing net societal harm. The governance architecture must therefore incorporate feedback loops that monitor the effectiveness of interventions in reducing clinically validated harm indicators, while simultaneously tracking displacement effects and user attrition to unregulated channels. Such socio-technical monitoring requires independent oversight mechanisms, algorithmic impact assessments conducted prior to deployment, and the routine publication of transparency reports that disclose aggregate bias statistics and intervention outcomes while preserving individual privacy.

6. Robustness, Sustainability, and System-Level Trade-offs

The robustness of an integrated prospect-theoretic and machine learning system in the football lottery domain spans multiple layers: statistical robustness to distributional shifts, adversarial robustness to strategic manipulation, and operational robustness to infrastructure failures. Distributional shifts are endemic in sports betting markets because the generative process of football matches is non-stationary, affected by rule changes, tactical evolution, seasonal variations, and exogenous shocks such as pandemic-related disruptions. A behavioral model that captures biases accurately during one league season may degrade substantially in the next if the underlying population of bettors changes composition or adapts their strategic heuristics in response to platform feedback. Continuous monitoring of model performance on temporally stratified holdout sets, coupled with automated retraining pipelines that incorporate gradual forgetting mechanisms, is necessary to maintain diagnostic validity. Concept drift detection algorithms must operate not only on prediction error metrics but also on distributional measures of the behavioral features themselves, as stable predictive performance on aggregate wagering outcomes can mask shifts in the underlying bias signatures that are the system's primary analytical target.

Adversarial robustness arises from the strategic incentives that sophisticated bettors and syndicates have to reverse-engineer and exploit behavioral models deployed by lottery operators and risk management platforms. If a system implements bet acceptance rules based on prospect-theoretic classifications, malicious actors may probe the decision boundary, deliberately submitting sequences of wagers designed to manipulate their behavioral profile and circumvent limits. This mirrors adversarial dynamics extensively studied in financial fraud detection and anti-money laundering systems, where classifiers are engaged in a continual arms race with adversaries who adapt their obfuscation strategies [10]. Architecturally, resilience to such manipulation can be enhanced through ensemble diversity, where multiple models trained on disjoint feature sets or with differing theoretical priors contribute to decisions, raising the cost of adversarial exploration. Randomized perturbation of decision thresholds, within ethically and legally permissible bounds, can further reduce the signal available to adversaries seeking to model the system's internals through queries.

Sustainability considerations are often marginalized in the design of entertainment analytics systems, yet the computational intensity of training large-scale behavioral models on population-scale transactional data carries a non-trivial environmental footprint. The push toward ever-larger language models in consumer technology has stimulated a critical discourse on the carbon implications of artificial intelligence [20], and the principles of Green AI are increasingly relevant to the sports analytics community. System designers must evaluate whether the incremental diagnostic value of a deep Transformer architecture with hundreds of millions of parameters justifies its energy cost relative to a carefully engineered gradient-boosted tree model that can be trained and served on modest hardware. Operational sustainability equally concerns the economic viability of the platform. The deployment of sophisticated real-time bias detection and intervention mechanisms represents a significant operational expenditure that must be weighed against the public health benefits and potential regulatory compliance savings. To avoid a future where only the largest oligopolistic operators can afford state-of-the-art player protection systems, industry consortia and open-source shared infrastructure initiatives may be necessary to commoditize the baseline behavioral monitoring capabilities, much as shared utilities have been developed for identity verification and anti-fraud screening in the financial sector [19].

The interplay between robustness and sustainability reveals a fundamental system-level trade-off. Highly robust, adversarially resilient, and continuously adapting systems tend to be more complex, resource-intensive, and difficult to audit for fairness and compliance. Conversely, simpler, more static systems are easier to govern and sustain but may fail catastrophically under novel behavioral regimes or deliberate attacks. Navigating this trade-off requires an architectural philosophy that embraces modularity and layered defense, where a persistent, interpretable core model grounded in domain theory provides a stable behavioral baseline, and an ensemble of adaptive, high-capacity models operates at the periphery for anomaly detection and adversarial response, subject to periodic review and decommissioning.

7. Cross-Domain Comparisons and Forward-Looking Perspectives

The structural challenges of integrating behavioral economic theory with machine learning in the football lottery domain find instructive parallels in other large-scale socio-technical decision systems, particularly algorithmic trading and insurance underwriting. In high-frequency financial markets, models that attempt to incorporate behavioral biases such as overreaction and herding operate under extreme latency constraints and face adversarial conditions where other algorithmic agents are simultaneously acting upon correlated signals

[13]. The market microstructure has evolved a rich set of circuit breakers, kill switches, and post-trade transparency mechanisms to manage systemic risk, providing a reference architecture for thinking about guardrails in automated lottery intervention systems. The concept of a “behavioral circuit breaker,” activated when aggregated loss-chasing indicators breach a pre-defined threshold and triggering a temporary cooling-off period or mandatory stake reduction across the platform, would translate these financial stability mechanisms directly into the gambling-harm context.

The insurance domain similarly grapples with the responsible use of behavioral profiling. When underwriters employ machine learning to assess risk propensity, they must navigate the tension between actuarial fairness, which prices risk accurately based on all available signals, and social fairness, which seeks to prevent discriminatory outcomes against protected groups. Insurers are increasingly experimenting with “fairness-aware” machine learning pipelines that impose statistical parity or equalized odds constraints during model training, a technical approach that can be adapted to the behavioral modeling of lottery players [8]. However, the analogy also highlights a critical difference: in insurance, behavioral profiling that identifies excessive risk is used to adjust premiums or deny coverage, whereas in gambling, the ethical mandate is to use such profiling to protect the vulnerable individual, not to extract discriminatory surplus. Ensuring that the system’s incentive structure remains aligned with this protective mission, rather than drifting toward revenue maximization, requires both regulatory oversight and internal governance mechanisms such as ethics boards and independent algorithmic auditors.

Multi-agent modeling and simulation present a powerful forward-looking methodology for stress-testing integrated behavioral-ML systems before they are released into the wild. Agent-based simulations calibrated to empirical betting data can generate synthetic populations where individual agents are endowed with heterogeneous prospect-theoretic parameters and learning rules, interacting with one another and with the automated platform mechanisms [19]. Such simulations allow researchers and regulators to explore counterfactual scenarios: What happens to the distribution of biases if a personalized nudge system is deployed? Does it genuinely reduce harm, or does it merely displace it to unsimulated shadow markets? Under what conditions do behavioral feedback loops lead to emergent market volatility or synchronized panic selling of bets? While agent-based models cannot replace empirical field trials, they provide a cost-effective ethical sandbox for eliminating the most dangerous failure modes before real consumers are exposed to them.

Looking forward, the trajectory of football lottery systems appears inexorably toward deeper integration of wearable biometry, immersive virtual reality environments, and continuous affective computing. If future platforms can monitor heart rate variability, galvanic skin response, and gaze patterns during the unfolding of a live match, the granularity of behavioral state estimation will increase by orders of magnitude, enabling ever more precise real-time detection of emotional arousal and impaired decision-making. The systems community must begin now to design the legal, ethical, and architectural frameworks that will govern such intimate data streams, explicitly defining the boundaries between protective intervention and manipulative exploitation. International standards bodies and professional engineering societies have a crucial convening role to play in articulating design principles and certification criteria for behavioral AI systems in gambling, analogous to the functional safety standards that govern autonomous vehicles and medical devices.

8. Conclusion

The confluence of prospect theory and machine learning in football lottery decision-making represents a paradigmatic case of behavioral AI integration, where the stakes involve public health, consumer protection, and the integrity of a major global entertainment industry. This paper has argued that the technical challenges of bias detection and intervention cannot be meaningfully separated from their systemic context: the data architectures that determine observability, the governance frameworks that constrain permissible action, the regulatory landscapes that fragment deployment environments, and the sustainability imperatives that bound long-term viability. A system that optimizes any single dimension to the neglect of others is not merely suboptimal but potentially harmful.

Our analysis has illuminated several architectural principles for the responsible design of such systems. Behavioral constructs from prospect theory must be operationalized through modeling pipelines that preserve interpretable links to domain theory while leveraging the pattern-recognition power of modern deep learning, a balance best achieved through multi-task learning and constrained latent representations. Data infrastructure must be designed for evolvability and temporal verisimilitude, capturing reference-point dynamics and counterfactual opportunity sets through event-sourced architectures with strong provenance. Governance must be encoded as a first-class architectural layer, externalizing jurisdictional policy rules and enabling independent algorithmic auditing without requiring access to raw personal data, potentially through federated and privacy-preserving analytics. Robustness to concept drift, adversarial gaming, and infrastructure failure is not a luxury feature but a core requirement, demanding modular architectures, layered ensemble designs, and continuous automated monitoring. The cross-domain lessons from algorithmic finance and insurance underscore both the transferability of technical mechanisms and the irreducible distinctiveness of the ethical mandate to protect rather than segment.

The integration of behavioral economics and machine learning in the lottery domain will continue to accelerate, driven by competitive pressures and genuine advances in computational social science. Without a concurrent maturation of the systems-level discourse, however, such acceleration risks embedding subtle harms and inequities into infrastructure that becomes ever harder to revise or dismantle. The interdisciplinary research community must therefore invest in building shared testbeds, open behavioral datasets with appropriate privacy safeguards, and reference architectures that embody the principles of safety, fairness, and transparency. Only through such collective engineering and governance efforts can the promise of behavioral artificial intelligence be realized in a manner that respects the autonomy, dignity, and well-being of the millions of individuals who participate in the football lottery each week.

References

1. Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.
2. Thaler, R. H., & Ziemba, W. T. (1988). Anomalies: Parimutuel betting markets: Racetracks and lotteries. *Journal of Economic Perspectives*, 2(2), 161–174.
3. Baboota, R., & Kaur, H. (2019). Predictive analysis and modelling football results using machine learning approach for English Premier League. *International Journal of Forecasting*, 35(2), 741–755.
4. Bunker, R. P., & Thabtah, F. (2019). A machine learning framework for sport result prediction. *Applied Computing and Informatics*, 15(1), 27–33.

5. Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4), 297–323.
6. Avery, C., & Chevalier, J. (1999). Identifying investor sentiment from price paths: The case of football betting. *Journal of Business*, 72(4), 493–521.
7. Kreps, J., Narkhede, N., & Rao, J. (2011). Kafka: A distributed messaging system for log processing. *Proceedings of the NetDB 2011 Workshop*.
8. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6), Article 115.
9. Abbott, M. W. (2020). The changing epidemiology of gambling disorder and gambling-related harm: Public health implications. *Public Health*, 184, 41–45.
10. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J. F., & Dennison, D. (2015). Hidden technical debt in machine learning systems. *Advances in Neural Information Processing Systems*, 28, 2503–2511.
11. Liu, S., Wang, Y., & He, H. (2020, March). A New Playing Method of the Guessing Football Lottery. In *IOP Conference Series: Materials Science and Engineering* (Vol. 790, No. 1, p. 012100). IOP Publishing.
12. Hubáček, O., Šír, G., & Železný, F. (2019). Exploiting sports-betting market using machine learning. *International Journal of Forecasting*, 35(2), 773–783.
13. Sterbenz, J. P., Hutchison, D., Çetinkaya, E. K., Jabbar, A., Rohrer, J. P., Schöller, M., & Smith, P. (2010). Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, 54(8), 1245–1265.
14. Gainsbury, S. M. (2015). Online gambling addiction: The relationship between Internet gambling and disordered gambling. *Current Addiction Reports*, 2(2), 185–193.
15. Lipton, Z. C. (2018). The mythos of model interpretability. *Queue*, 16(3), 31–57.
16. Voigt, P., & Von dem Bussche, A. (2017). *The EU General Data Protection Regulation (GDPR): A practical guide*. Springer.
17. Paul, R. J., & Weinbach, A. P. (2011). Investigating the “wisdom of crowds” in football betting markets. *Applied Economics Letters*, 18(8), 733–736.
18. Clotfelter, C. T., & Cook, P. J. (1990). *Selling hope: State lotteries in America*. Harvard University Press.
19. Bonabeau, E. (2002). Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(Suppl 3), 7280–7287.
20. Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63.