International Journal of Innovative Research in Computer Science and Technology (IJIRCST)
ISSN (Online): 2347-5552, Volume-13, Issue-6, November 2025
DOI: https://doi.org/10.55524/ijircst.2025.13.6.9
Article ID IJIRE-1414, Pages 71-91
www.ijircst.org

Predictive Algorithms in Crypto Investment Platforms: Rethinking the SEC's Disclosure and Suitability Regimes

Awais Amjad[1], Mohammad Salman Iqbal[2], and Waqar Ahmad[3]

[1]Department of Faculty of Business and Law, University of Northampton, Northampton, UK
[2]Department of Penn Carey Law School, University of Pennsylvania, Philadelphia, PA, USA
[3]Department of Penn State Dickinson Law, The Pennsylvania State University, Carlisle, PA, USA

Correspondence should be addressed to Awais Amjad; awaisamjad9988@gmail.com

Received: 1 November 2025 Revised: 16 November 2025 Accepted: 29 November 2025

ABSTRACT- Crypto investment platforms rely on predictive algorithms that analyze past blockchain activity, market patterns and conditions, and user behavior to surface specific assets and encourage user engagement. These systems combine on-chain signals, market microstructure features, and user profiles to rank tokens and personalize messages. Based on the established user profile, predictive algorithms push personalized notification content during a user's habitual activity window. As a result, predictive models increasingly determine how retail investors encounter information.

As these systems increasingly rely on AI-driven ranking, personalization, and automated prompts, the investor decision-making environment has shifted away from traditional, human-mediated financial guidance. This poses potential risks for investors due to a lack of transparency related to the algorithm, conflicting with disclosure and potentially accountability. Disclosure and suitability regimes that were built for and evolved within environments for human-mediated brokerage strain when personalization and ranking shift material choices into a system layer. This paper examines the impact of predictive algorithms on the crypto investment landscape and the subsequent need for evolving disclosure, suitability, and accountability policies and regulatory frameworks. The paper develops a qualitative framework supported by market context, relevant court rulings, and cross-jurisdictional regulatory guidance on privacy, accountability, intellectual property, consumer protection, and governance. It examines how these principles apply to human oversight and system-level accountability within predictive models, to mitigate algorithmic investor risk and strengthen transparency.

The discussion situates predictive crypto tools within broader legal contexts, emphasizing how regulatory principles (disclosure, suitability, accountability) intersect with algorithmic design. Particular attention is given to judicial precedents and emerging case law that center on disclosure, investor protection, and system accountability in regards to regulatory obligations for automated systems. An example of the tools under discussion includes robo-advisors that recommend crypto portfolios to trading bots that can execute rapid trades based on market signals. The paper contributes an integrated analysis that combines a detailed risk taxonomy - including provenance, adversarial risk, privacy/leakage, IP/trade-secret tension, and the conflict between transparency and protection of IP, and human oversight failures. Further, an evaluation design that operates without external datasets using immutable input snapshots, ledger-based backtesting, shadow deployments, counterfactual testing, and adversarial red-teaming. Lastly, a governance and compliance structure focused on supervisory accountability and the institutional processes that determine how predictive algorithms are monitored and evaluated. The study delves into three short scenarios that show plausible failure modes. Health care and experimental sources are used only as analogies for validation, documentation, human oversight, and challenge procedures, and not for financial claims. The study stays within qualitative methods and avoids invented statistics. Limits include the reference set scope and the absence of empirical market tests. The report provides an overview of how predictive algorithms are being used in crypto investments, and what type of challenges are posed to the SEC, what the relevant actions and disclosures are, and comparisons to approaches in other jurisdictions.

KEYWORDS: Predictive algorithms, Crypto platforms, Disclosure, Suitability, Governance, Privacy, Trade secrets, Intellectual property, Accountability, Risk taxonomy.

I. INTRODUCTION

Predictive algorithms are increasingly embedded in crypto-investment platforms and actively guide investor attention through ranked lists, trending bars, curated tiles, and behaviorally timed alerts.[1] Rather than serving as neutral infrastructure, these systems transform raw trading feeds and on-chain events into features that are combined with profile attributes - often derived from risk-tolerance questionnaires - to generate surface-level outputs such as rankings, highlights, or prompts that investors may interpret as the platform's default options.[2] Algorithmic advising and automated decision-support, therefore, occur through system-layer interactions that materially influence visibility, salience, and investor behavior.[3]

Research on credit risk and financial equity demonstrates that algorithmic scoring can shape access and outcomes, and highlights the importance of documentation of purpose, data lineage, and model limits for meaningful

review.[4] These concerns map directly onto predictive crypto-advisory tools, which depend on intensive data collection, behavioral profiling, and the linkage of wallet activity to inferred attributes.[5] Privacy and surveillance writing warns that such practices create risks of secondary data use, inference-based manipulation, and opaque categorization, each of which becomes legally significant when the resulting outputs are positioned as investment prompts or recommendations.[6]

Scholarship on judicial automation illustrates that automated decision-support systems undermine procedural integrity when institutions lack robust oversight, the ability to replay and reconstruct system reasoning, and structured review mechanisms for evaluating algorithmic outputs.[7] These insights intersect with financial regulation when algorithmic prompts, rankings, or curated displays influence investor actions.[8] Financial disputes involving misleading impressions, suitability failures, or system-level misrepresentations may require regulators or courts to examine how the model operated, what data it relied on, and whether platform governance provided adequate documentation, validation, and supervisory controls.[9]

Although predictive crypto-advisory systems use contemporary computational methods, the regulatory issues they raise remain grounded in duties of disclosure, fairness, and investor protection.[10] Jurisdictions beyond the United States – including the European Union, the United Kingdom – have articulated parallel concerns related to transparency, suitability assessment, conflict mitigation, and algorithmic explainability in digital advisory contexts.[11] While the comparative analysis is reserved for the Discussion section, these global developments indicate that predictive systems trigger similar regulatory questions across global institutions.[12]

This paper studies the intersection between predictive ranking, personalization, and the regulatory expectation that advisory processes, whether human or through algorithmic systems, must produce interactions that are fair, comprehensible, and suitable to the investor's circumstances.[13] Consistent with the paper's methodological boundaries, parallels to health-care validation and experimental research are used only to illustrate mechanisms of documentation, stress testing, challenge procedures, and human oversight.[14] These analogies serve to clarify how structured evaluation, reproducibility, and audit-ready records can be applied to predictive systems without importing empirical claims from external domains.[15] Their conceptual function is to highlight how reproducibility, replayability, and formalized oversight – tools common in regulated technical environments – become essential when algorithmic recommendations must withstand scrutiny under disclosure, suitability, and supervisory standards.[16] This framing supports the paper's three fundamental research questions that connect the legal framing to the operational analysis developed in later sections.[17]

The first asks which risk dimensions within predictive crypto-investment systems most directly threaten disclosure clarity and suitability fit, an issue that becomes legally salient when ranking, prompting, or personalization mechanisms shape investor impressions in misleading ways.[18] The second examines how an internal review team can evaluate whether these systems satisfy regulatory expectations using only artifacts such as model documentation, logs, and code-level inspection, reflecting the supervisory and recordkeeping duties emphasized in the SEC's regulatory guidance and proposed rules.[19] The third investigates how governance structures can balance privacy constraints and trade-secret protections with accountability and auditability obligations, a tension that becomes central when platforms must reconstruct algorithmic reasoning to address disputes over disclosure accuracy, suitability, or conflict mitigation.[20] Together, these aspects build the fundamental foundation for the risk taxonomy developed in Section V, and the governance and compliance framework elaborated in Section VI, ensuring that the technical and legal analyses operate within a unified evaluation structure pertinent to predictive algorithms.

II. BACKGROUND

The multilayered predictive systems that crypto-investment platforms rely on use mechanisms such as ranked lists, curated tiles, trending indicators, token categories, order-book mechanisms, and behaviorally timed alerts, which traditionally have been described in technical terms as data pipeline components or interface features.[21] They function in practice as investor interactions under emerging regulatory interpretations such as the 2017 Robo-Adviser Guidance.[22] Within the architecture of most platforms, data flows from ingestion to feature formation, to model scoring, and finally, to interface-level surfacing.[23] With each stage contributing to how investors encounter and interpret key asset information.[24]

The data plane typically ingests raw trading feeds, microstructure indicators, order-book movements, volatility metrics, and on-chain events such as transfers, liquidity shocks, or mempool summaries.[25] These inputs are combined with profile attributes, often derived from investment-horizon prompts or self-reported preferences, in addition to the aforementioned risk-tolerance questionnaires.[26] The resulting feature set determines the signals a model uses to construct personalized surfaces, trending lists, and categorization tiles. Because these inputs materially influence what assets appear salient or prominent, they carry direct implications for disclosure clarity.[27] Under existing regulatory expectations, platforms may be responsible for explaining which data sources shape the ranking logic, what signals are weighted, and how profile attributes affect the resulting investor interaction.[28]

The decision plane translates these features into scoring, ranking, sorting, and timing mechanisms that actively structure investor attention. Algorithms determine which tokens appear first, which are grouped into prominent categories, which receive tile placement, and which assets trigger prompts or notifications. These are not neutral presentation choices: they operate as regulated investor interactions, as was articulated by the SEC's 2023 Predictive Data Analytics Proposal,[29] where personalized rankings, prompts, timing engines, and interface-level prioritization may influence investor behavior in ways that violate suitability or require conflict-mitigation and supervisory obligations. Ranking logic has implications for suitability fit when lists or tiles imply relevance or appropriateness for a particular investor profile.[30] Timing

engines, which align notifications with behavioral patterns, raise concerns about prioritizing firm-side metrics over investor needs through these personalization mechanisms, which may create conflicts of interest if prioritization improves platform engagement while misaligning with client objectives.[31]

The interface layer operationalizes these decisions through the aforementioned interface mechanisms (curated tiles, tailored lists, prompts). These surface-level outputs carry legal relevance because they shape investor impressions and constitute representations that must be fair, balanced, and comprehensible. In this environment, the distinction between technical design and advisory communication collapses; that is to say, interface-level mechanisms become forms of financial advisory and guidance that fall under the same duties of disclosure and fairness as a human advisor would.

Underlying these technical layers are governance artifacts such as gating logic, logs of model actions, replayable traces of decision pathways, documentation of code changes, lineage records and performance summaries which are of critical importance due to their function as internal records and as the foundation for reconstructibility, which is central to disclosure and supervisory expectations under the Advisers Act and is later developed in Section VI. Logs and lineage records determine whether a firm can demonstrate how a ranking or prompt arose, what version of the model was deployed, and whether suitability constraints or conflict-mitigation controls functioned properly. Because predictive systems can produce misleading impressions even without intent, these governance elements are also essential for risk identification and compliance reviews.

This architecture parallels system-layer challenges identified in global regulatory frameworks. European and British-aligned guidelines similarly emphasize transparency of model logic, suitability assessments in automated environments, robust oversight of personalization mechanisms, and explainability of interface-level outputs.

By outlining how platforms transform data ingestion, feature formation, ranking logic, and prompt generation into system-layer investor interactions, this section clarifies the technical pathways through which disclosure clarity, suitability fit, and conflict-management obligations may later arise. These operational details establish the context necessary for examining how predictive systems intersect with regulatory duties, allowing the subsequent legal framing to evaluate how these mechanisms are treated under existing securities law. The interface layer operationalizes these decisions through the aforementioned interface mechanisms (curated tiles, tailored lists, prompts). These surface-level outputs carry legal relevance because they shape investor impressions and constitute representations that must be fair, balanced, and comprehensible. In this environment, the distinction between technical design and advisory communication collapses; that is to say, interface-level mechanisms become forms of financial advisory and guidance that fall under the same duties of disclosure and fairness as a human advisor would.[32]

Underlying these technical layers are governance artifacts such as gating logic, logs of model actions, replayable traces of decision pathways, documentation of code changes, lineage records, and performance summaries, which are of critical importance due to their function as internal records and as the foundation for reconstructibility, which is central to disclosure and supervisory expectations under the Advisers Act and is later developed in Section VI.[33] Logs and lineage records determine whether a firm can demonstrate how a ranking or prompt arose, what version of the model was deployed, and whether suitability constraints or conflict-mitigation controls functioned properly.[34] Because predictive systems can produce misleading impressions even without intent, these governance elements are also essential for risk identification and compliance reviews.[35]

This architecture parallels system-layer challenges identified in global regulatory frameworks. European and British-aligned guidelines similarly emphasize transparency of model logic, suitability assessments in automated environments, robust oversight of personalization mechanisms, and explainability of interface-level outputs.[36]

By outlining how platforms transform data ingestion, feature formation, ranking logic, and prompt generation into system-layer investor interactions, this section clarifies the technical pathways through which disclosure clarity, suitability fit, and conflict-management obligations may later arise.[37] These operational details establish the context necessary for examining how predictive systems intersect with regulatory duties, allowing the subsequent legal framing to evaluate how these mechanisms are treated under existing securities law.[38]

III. RELATED WORK

Credit risk and financial equity sources show how algorithmic assessment affects access and outcomes and raise distributional and documentation issues that transfer to investment suggestion systems.[39] [40] Privacy and surveillance work outlines data collection risks, data sharing risks, and the need for transparency and consent mechanisms.[4142] Automation in judicial administration surveys oversight structures for algorithmic tools and the importance of human responsibility, which informs reviewability and audit in platforms.[43] Legal writing on intellectual property and on AI as creator or inventor addresses ownership of outputs and the need to separate protected logic from what must be disclosed to users, which informs model card design and trade secret management.[44] [45] [46] Criminal liability analysis connects automated decision effects to responsibility lines, a useful lens for incident response and remediation.[47] Equality writing anchors fairness testing and user tiering controls.[48] Health care and clinical sources are used only as analogies for validation, transparency, safety, and human oversight. Cardiac and neurological application surveys describe validation pipelines, dataset documentation, and harm mitigation, which map to testing and review in financial recommendation settings.[49] [50] [51] [52] [53] [54] [55] Experimental pharmacology papers provide an analogy for controlled evaluation and documentation practices, not for financial claims.[56] [57] Work on ICT for migrants supports the inclusion aim and the need to document accessibility outcomes, which can inform fairness testing for novice users of investment platforms.[58] The role of these sources in this paper is analogous only.

Common goals for a securities regulator include investor protection, fair disclosure, conflict management, market integrity, and sales-practice duties, including suitability.[59] Moreover, any advisor - whether robo or algorithmic - is bound by the fiduciary obligations embedded in the Investment Advisers Act.[60] A disclosure regime requires that material features of a product or recommendation be stated in plain language at a time and place where the user can act.[61] A suitability regime expects that personalized suggestions fit a declared user profile and that higher-risk prompts are regulated through gating.[62] Predictive ranking and personalization shift decision-making into a system layer, but that layer must still support meaningful disclosure and suitability checks. Documentation, human review, and logging provide the bridge between algorithmic action and legal responsibility.[63] The SEC has recently required that disclosures for crypto-asset offerings be "clear, concise, and understandable," avoiding technical jargon.[64] By analogy, the same expectation applies to robo-advisory systems: a platform must offer a clear rationale for how a suggestion was generated, even though firms often struggle to explain how their models reach a particular output.[65]

IV. LEGAL FRAMING

As the preceding sections have shown, predictive crypto tools algorithmically shape investor attention by determining which tokens appear first on screens, impacting and altering visibility through personalized algorithms, tailoring the content and wording of investment prompts, and timing notifications to coincide with behavioral patterns - each of which constitutes an investor interaction that materially influences decision-making. To understand how these dynamics strain existing protections, it is necessary to situate them within the legal framework that governs disclosure and suitability obligations. These obligations in U.S. securities law were built for a human-governed environment and are grounded in well-established antifraud principles, such as Section 17(a) of the Securities Act of 1933, which prohibits obtaining capital or property "by means of any untrue statement of a material fact or any omission to state a material fact."[66] Courts have reinforced that these obligations apply even when misleading impressions arise from automated processes.[67]

The 2023 Predictive Data Analytics Proposal (Release No. 34-97990) by the SEC further extends these foundational principles into predictive-algorithmic contexts. It states that when a platform presents an investment in a manner that influences a user's decision, the firm is accountable for ensuring that the representation is fair, balanced, and consistent with applicable disclosure and suitability requirements.[68] This aligns with the SEC's earlier analysis of predictive algorithms, which notes that PDA systems may "optimize for the firm's interest over the investor's interest" through model-design choices and subtle prioritization patterns - conflicts of interest that must be neutralized under the proposed rule.[69]

Under the Securities Exchange Act of 1934 and the Investment Advisers Act of 1940, the SEC's proposed conflicts rules would require firms to "neutralize conflicts of interest associated with broker-dealers' or investment advisers' interactions with investors through these firms' use of technologies that optimize for, predict, guide, forecast, or direct investment-related behaviors or outcomes."[70]

The PDA Proposal also clarifies that personalized ranking, prompts, timing mechanisms, and interface designs constitute "investor interactions employing covered technology," thereby triggering conflict-mitigation duties and supervisory/recordkeeping obligations.[71] These places predictive models under the same legal scrutiny applied to human advisers.

These expectations mirror standards outlined in FINRA Rule 2210, which requires investor communications to be "fair and balanced" and prohibits selective or asymmetrical presentation of information.[72] The PDA Release affirms that PDA-driven interactions fall under these same principles.[73]

In addition, recordkeeping, traceability, and reconstruct ability are not optional under the proposed rule. The SEC requires firms to maintain "records of the version of any covered technology," "records of inputs and outputs associated with investor interactions," and documentation enabling supervisors to "reconstruct the basis of an interaction."[74] This parallels the paper's emphasis on replay ability, lineage, and model documentation. These artifacts are crucial to meet legal expectations for reconstruct ability - a relevant factor in assessing suitability and disclosure adequacy.

Similarly, due diligence is placed on robo-advisors. Suitability principles apply in the case of robo-advisers. The SEC's 2017 robo-adviser guidance states that advisers must "explain their algorithmic methodology, including the assumptions and risks," and must gather "sufficient client information to form a reasonable basis for recommendations."[75] It also clearly states, "To address potential gaps in a client's understanding of how a robo-adviser provides its investment advice, the robo-adviser (like all registered investment advisers) should disclose, in addition to other required information, information regarding its particular business practices and related risks."[76] These principles map to predictive crypto tools because both rely on automated processes that shape investor decision environments, whether through portfolio construction or personalized ranking and prompt timing.

The regulatory expectations articulated in the SEC's 2017 Robo-Adviser Guidance hold relevance. The guideline establishes the baseline regulatory expectations for automated investment tools. Because these platforms often interact with clients exclusively through digital interfaces, the Guidance frames disclosure and suitability as system-level obligations.[77] It covers three fundamental aspects:[78]

The presentation/transparency of disclosures to all clients about the implementation of robo-adviser services[79]. The active obligation to obtain accurate information to supplement all automated advisory tools[80]. The strict implementation of SEC-accurate compliance programs specifically tailored to automated tools[81]. The SEC emphasizes that a client's ability to make informed decisions may depend entirely on the adviser's electronic disclosures, especially when there is little or no human interaction. In such environments, robo-advisers must provide full and fair disclosure of all material facts, avoid misleading statements, and ensure that information is presented in a manner clients are reasonably likely to read and understand - prioritizing accessibility[82].

The Guidance stresses that disclosures must be sufficiently specific to allow clients to understand the nature and scope of the advisory service[83]. This expectation directly parallels the issues raised in the present paper: when predictive models determine how assets are ranked, surfaced, or timed for presentation, those processes become part of the adviser's "business practices" and therefore must be disclosed in a clear, comprehensible manner. The SEC's insistence on specificity anticipates the concerns raised by predictive ranking and personalized prompts, which likewise can influence investor decisions in ways that are material to client understanding.

To address gaps in client understanding, the Guidance requires robo-advisers to explain their business model, the algorithmic functions used to manage accounts, any assumptions and limitations embedded in those models, and the specific risks inherent in automated portfolio construction[84]. Examples include algorithms that may rebalance without regard to market conditions, models that may not account for sustained market changes, and systems that rely solely on questionnaire data. The Guidance also instructs advisers to disclose any circumstances under which human personnel override the algorithm, the degree of human oversight involved, and the extent of third-party involvement in developing or operating the algorithm - including resulting conflicts of interest[85].

This framework has direct implications for predictive crypto-asset systems. When a platform uses models to rank tokens, interpret on-chain metrics, weigh behavioral inputs, or time investor prompts, each of those mechanisms constitutes a core algorithmic function with assumptions, limitations, and risks requiring disclosure[86]. The SEC's expectations around algorithmic transparency, third-party involvement, and conflict disclosure align closely with the concerns raised by predictive analytics in crypto interfaces, which often optimize visibility or recommendations based on opaque model criteria.

The Guidance further warns that robo-advisers must avoid creating a false implication about the scope or nature of their services. For example, a robo-adviser may not suggest that it provides comprehensive financial planning or year-round monitoring of client accounts (as we will see in the enforcement cases below) if it lacks information about clients' tax or debt obligations, nor may it imply that a tax-loss harvesting feature constitutes tax advice. Similarly, advisors may not imply that data outside their questionnaire or operational scope is used in generating recommendations if it is not actually implemented[87]. This principle is critical for predictive systems: if platforms design interfaces or prompts that give the appearance of comprehensive risk analysis, holistic market scanning, or behavioral insight beyond their actual capabilities, such implications can constitute misleading statements under the Advisers Act.

The Guidance thus establishes that interface design and communication structures are part of the advisory representation itself. For predictive crypto tools that depend heavily on interface-level interactions - prominent rankings, curated lists, push notifications, prompts timed to user behavior - this means the visual and interaction design choices are not merely UX decisions, but legally regulated communications that must not mislead clients about the scope, accuracy, or robustness of the underlying model.

It is because robo-advisers rely entirely on digital communication that the 2017 Guidance places immense importance on effective and transparent presentation[88]. The SEC warns that disclosure must not be "buried" or rendered incomprehensible and encourages the use of design features such as pop-ups, tooltips, FAQs, and mobile-optimized formats to ensure that clients actually understand the information provided[89]. Advisers must also present key disclosures before account sign-ups so clients have the information necessary to make an informed investment decision[90].

This focus on design is in line with the core argument presented by this paper. When predictive algorithms operate through UI-layer mechanisms such as personalized notifications and ordered lists, the way information is presented becomes inseparable from the core function of financial advisory services. The SEC's emphasis on accessibility, clarity, and timing of disclosures predicts concerns raised by automated advisory tools and predictive analytics, and the concern that they may influence investor behavior through subtle or optimized interaction design.

The Guidance makes clear that robo-advisers are subject to the same fiduciary, disclosure, and suitability obligations as traditional advisers under the Advisers Act[91]. This understanding is foundational to the broader argument of this paper. Predictive systems operate within the legal regime. Their assumptions, limitations, risks, and advisory capacities must be disclosed with full transparency. Investors must be equipped with actionable information in order to make suitable investment decisions.

A core element of the 2017 Robo-Adviser Guidance is the requirement that automated platforms provide suitable investment advice consistent with the client's best interests as per fiduciary duty. Suitability is complex in automated environments. Information-gathering is primarily carried out via questionnaires. The SEC makes clear that, despite these technological constraints, robo-advisors are obliged to "make a reasonable determination" that the investment advice is context-appropriate, in line with the investment objectives and needs of the individual client[92]. This requirement extends logically to predictive and algorithmic systems that shape investor interactions on crypto platforms: the adequacy of client profiles becomes central to suitability.

Questionnaires offer limited information, and each varies in length, clarity, and informational scope. The SEC rightly points out that the questionnaires used by automated advisory tools do not provide clients the ability to supply context or nuances tailored to their specific situations. The lack of mechanisms for follow-up information is also discussed in the 2017 Guidance[93]. This structural rigidity creates the risk that the advice presented may be crafted using incomplete or internally inconsistent information.

To mitigate these risks and limitations, the Guidance instructs robo-advisers to evaluate their questionnaires on the sufficiency of all information gathered to meet suitability obligations[94]. This includes assessing the clarity of all information gathered, providing measures for additional explanations through implementing design features such as pop-up boxes or tooltips, and ensuring the

system itself is capable of detecting consistent client responses. Examples include alerting clients when answers conflict or implementing internal systems that automatically flag inconsistencies for review. The SEC's suitability process emphasizes active validation and interpretive checks, especially in automated environments[95].

These suitability principles integrate directly with the earlier disclosure obligations described in the Guidance. Clients must understand the scope, limitations, and risks of the algorithmic service, and robo-advisers must avoid creating false implications about their capabilities. The suitability requirements supplement this by aligning input-gathering mechanisms with the scope and transparency of the service, ensuring that all advice generated is not only suitable but grounded in factual, client-specific information that serves their investment needs and goals[96].

The 2017 Guidance clearly establishes that predictive and algorithmic systems require stricter suitability obligations. When an automated interface ranks tokens, highlights asset categories, or behaviorally timed prompts, the system is effectively generating advice tied to a user's inferred preferences, risk tolerance, or behavioral patterns. If the underlying profile is built on inadequately validated, insufficient, or inconsistent information, the resulting interactions may fail the suitability criteria[97]. This is in line with the argument of the paper, which claims that the adequacy of the fundamental design of predictive tools, such as questionnaire design, client profiling, and input validation, is central to legal compliance.

The 2017 Robo-Advice Guidance also clarifies that automated advisers are subject to the full requirements of Rule 206(4)-7, which mandates that every registered investment adviser adopt, implement, and annually review written compliance policies designed to prevent Advisers Act violations[98]. The Guidance stresses that these policies must reflect not only traditional advisory operations but also the unique risk exposures created by automation, including reliance on algorithmic code, limited human interaction, internet-based service delivery, and the scale at which predictive systems can influence client decisions[99].

The regulatory framework presented in the guidance also extends to predictive systems used in crypto investment interfaces. Compliance obligations extend into the operational logic of system design. The Guidance states that automated advisory tools must establish controls addressing the development, testing, backtesting, and post-deployment monitoring of algorithmic code. Algorithmic code must perform as represented. Updates must not create adverse effects on client accounts. Any changes to model logic must be disclosed when they materially affect portfolio outcomes[100].

These requirements map directly onto the governance and evaluative structures developed later in this paper. The call for periodic model review, performance monitoring, and version control in Section VI echoes the Guidance's expectation that advisers maintain replayable and reviewable records capable of reconstructing recommendation logic for full transparency. Similarly, Rule 206(4)-7's emphasis on written procedures around questionnaires, sufficiency, suitability processes, and data inputs supports the risk taxonomy outlined in Section V, which identifies provenance, proxy risk, personalization logic, and suitability gates as key domains requiring continuous evaluation[101].

The Guidance further requires policies for third-party oversight, cybersecurity safeguards, and controls on social and electronic media. These areas overlap with the operational governance tools in Section VI. For example, the need to supervise third-party developers or model vendors aligns with the governance structure's requirements for access control, secret handling, and audit trails. Cybersecurity expectations correspond to the risk taxonomy's treatment of system robustness, adversarial pathways, and data lineage protections. Even the requirement to supervise digital marketing channels aligns with the paper's analysis of ranking prompts and personalized nudges as regulated investor interactions rather than neutral interface elements[102].

Rule 206(4)-7's focus on testing, monitoring, documentation, change control, and supervisory oversight provides the statutory grounding for the governance checklist, lineage requirements, suitability gating, replayability mechanisms, and incident-response structures developed in Sections V and VI. Rather than treating governance as an additive layer, the Guidance positions these controls as the mechanisms through which an adviser satisfies its fiduciary duties in automated environments.

As the following enforcement cases against Wealthfront[103] & Hedgeable[104] demonstrate, these concerns are grounded in reality. They have been applied in concrete enforcement actions against automated advisory systems, demonstrating how longstanding disclosure, suitability, and compliance obligations extend directly to algorithmic processes. These cases illustrate that when firms deploy predictive or automated decision tools, representations about the functionality, accuracy, and supervision of such systems become material facts within the meaning of federal securities law. The obligations outlined earlier in this section - fair disclosure, suitability, and system-level supervision - are therefore directly implicated when platforms rely on predictive algorithms that shape how investors encounter information, rank assets, or receive prompts.

The 2018 administrative proceeding against Wealthfront Advisers LLC provides a significant example of these principles in practice[105]. Wealthfront, a software-based robo-adviser, offered a proprietary Tax-Loss Harvesting algorithm designed to reduce client tax liability. It would theoretically have achieved this through strict monitoring of accounts for transactions that could trigger potential wash sales. In client-facing whitepapers, Wealthfront stated that this obligation was met - it claimed Wealthfront "monitors all the accounts it manages for each client to avoid any transactions that might trigger a wash sale." In reality, from 2012 through mid-2016, Wealthfront did not monitor client accounts for wash-sale-triggering transactions. In fact, wash sales did occur. These statements constituted untrue or misleading representations under Section 206(4) and Rule 206 (4)-1, which prohibit registered advisers from publishing advertisements containing false or misleading statements regardless of intent[106].

Wealthfront also retweeted client testimonials on Twitter, in violation of Rule 206(4)-1's prohibition on advertisements referring "directly or indirectly" to

testimonials, and paid bloggers for client referrals without providing the disclosures required under Rule 206(4)-3[107]. Additionally, Wealthfront also failed to adopt and implement written policies and procedures reasonably designed to prevent such violations, breaching Rule 206(4)-7[108]. The SEC emphasized that a violation of Section 206(4) does not require scienter, meaning that negligent algorithmic misrepresentations or failures of automated systems such as Robo-Advisers are sufficient to trigger liability.

Another example of the SEC bypassing intent can be found within the Aaron v. SEC, 446 U.S. 680, 695 (1980) trial[109]. The Supreme Court held that liability under Sections 17 (a)(2) and (3) does not require intent, establishing that firms are responsible for misleading outcomes even if said outcomes arise from non-intentional or structural features of a system rather than deliberate misconduct.

The enforcement action against Hedgeable Inc. further underscores the applicability of these standards to automated-advisory models[110]. Hedgeable marketed a "Robo-Index" from 2016 to 2017, which allowed prospective & current clients to compare the performances from 2014 and 2015 against two other automated digital investment advisory programs, namely, "Robo-Adviser 1" and "Robo-Adviser 2". These comparisons were compiled in a Hedgeable Composite; however, the index was misleading in three significant respects[111].

First, the Hedgeable Composite used in the comparison included less than four percent of all Hedgeable client accounts for the relevant period, rendering the performance portrayal unrepresentative[112]. Second, Hedgeable's methodology for constructing competitor performance was inaccurate, as it was not based on those actual trading models but on approximations drawn from publicly available website information[113].

Third, even under its own chosen methodology, Hedgeable miscalculated the annualized returns for both the Robo-Index and its internal composite. Hedgeable also disseminated misleading fact sheets that overstated ETF returns relative to benchmark indices and failed to maintain documentation substantiating any of the performance data it published[114]. These misrepresentations were enabled by an ineffective compliance program; Hedgeable's policies did not require supervisory review or approval of marketing materials or performance data before publication, a clear violation of the requirements set forth under Rule 206(4)-7 for written policies reasonably designed to prevent Advisers Act violations[115].

This case, paired with Wealthfront, demonstrates that the SEC treats algorithm-driven performance claims, client-facing analytic and promotional tools as material representations subject to the full scope of advertising, documentation, and compliance obligations - regardless of whether the underlying outputs are driven by intent, generated by humans, or automated systems[116].

These cases also substantiate the need for the governance, documentation, and replayability mechanisms discussed in the subsequent section. Wealthfront and Hedgeable demonstrate how failures in algorithmic monitoring, disclosure accuracy, and compliance documentation can arise within automated systems. Hedgeable's performance misrepresentations and Wealthfront's inaccurate descriptions of its tax-loss harvesting algorithm both illustrate how failures in system design, client-input management, and disclosure accuracy can result in grievous violations of the Advisers Act[117]. Suitability, disclosure, and presentation of information thus operate together, forming a unified regulatory framework.

As predictive tools proliferate across crypto investment platforms - surfacing assets, ordering information, and steering investor attention - the same legal expectations apply. The enforcement record clarifies that obligations surrounding investor protection are a matter of the firm designing all automated systems with compliance as a foremost priority, along with due monitoring & disclosure[118].

V. ANALYSIS AND RISK TAXONOMY WITH EMBEDDED EVALUATION

Section V develops the paper's second major contribution: a theoretical evaluation framework for predictive systems that operationalizes the disclosure, suitability, and supervisory duties established in Section IV. The aim of this framework is not to prescribe an operational standard for industry adoption, but to offer a conceptual structure that illustrates how the legal obligations governing automated investment interactions can be interpreted through the internal mechanics of predictive models.

The framework arises from the central insight that the risks described in earlier sections - proxy-driven behavior, opacity in personalization logic, suitability violations, optimization conflicts, and model drift - manifest within a system's internal decision pathways rather than solely in observable outcomes[119]. As we have established, existing regulatory frameworks do not fully address these system-level behaviors. Accordingly, a theoretical model is needed to clarify how such risks can be rendered analyzable, traceable, and conceptually accountable within the confines of the securities-law obligations examined in Section IV.

To accomplish this, Section V introduces a set of evaluation criteria wherein each criterion can be measured using internal artifacts - such as model cards, input snapshots, audit logs, or replay traces - because the risks identified above arise from internal mechanisms that external performance metrics cannot diagnose or reveal[120]. These artifacts enable structured examination of the model's decision logic, data dependencies, and operational constraints, allowing the criteria to assess whether the system's internal operations conform to disclosed methodologies, align with suitability obligations, and avoid optimized patterns that disadvantage investors.

The risks posed by predictive systems operate across multiple layers of the data-to-decision pipeline, and these layers correspond directly to the disclosure, suitability, and supervisory duties outlined in Section IV. At the data and feature level, issues of provenance, representativeness, lineage, and proxy leakage can distort model outputs and propagate distributional effects that materially influence which crypto-assets are ranked, surfaced, or highlighted to users[121]. Prior literature in financial equity and credit modeling demonstrates that biased or narrow datasets generate systemic misrepresentation in downstream outputs[122].

These concerns are not only technical: under U.S. securities law, distorted or unrepresentative data that shapes rankings, surfacing, or prompts may create misleading impressions relevant to Section 17(a), and the antifraud provisions of the Advisers Act, which apply regardless of intent[119]. At the personalization and decision-making layer, feature-driven inferences about user preferences or risk tolerance function as de facto suitability mechanisms. When the personalization logic is built on incomplete, inconsistent, or profiles inferred from behavioral signals that are outdated or do not align with client goals, the resulting interactions can fail the suitability expectations embedded in regimes such as the U.S adviser-suitability obligation to base recommendations on a reasonable understanding of the client's financial profile.

Additionally, it mirrors the MiFID II Article 25 requirement that firms obtain sufficient, accurate, and relevant information to recommend suitable products and avoid constructing environments that implicitly mischaracterize client risk[120].

Privacy literature emphasizes the need for data minimization, purpose limits, and role-based access controls to reduce secondary use & confidentiality risks[121]. Equality writing demonstrates how neutral features can operate as proxies for protected attributes, necessitating periodic proxy audits in both feature-selection and personalization outputs[122]. Trade secret and intellectual property writing explains the need to protect model features, internal dashboards, and proprietary logic families while still supplying reviewers and supervisors with sufficient information to understand purpose, inputs, limits, and risks[123]. This balance parallels the SEC's expectations under the Advisers Act and the 2023 Predictive Data Analytics Proposal, which require firms to provide transparent descriptions of system behavior even when the underlying logic is proprietary[124]. Work on automation in judicial administration reinforces that replayable traces and formal audit trails are essential for oversight in automated systems[125], a point that directly connects to securities-law requirements for maintaining reviewable records of how investor-facing outputs were generated.

Figure 1 decomposes a predictive crypto system into eight sequential stages, identifying where each category of risk originates. The figure presents a conceptual pipeline that locates where specific controls, safeguards, and supervisory checks must act across the data-to-decision pathway.

Each stage in this pipeline corresponds to points where the modeling process can create, amplify, or mitigate disclosure and suitability risks. By locating controls along this pathway, the paper identifies where legally relevant decision-making occurs within a predictive system[126].

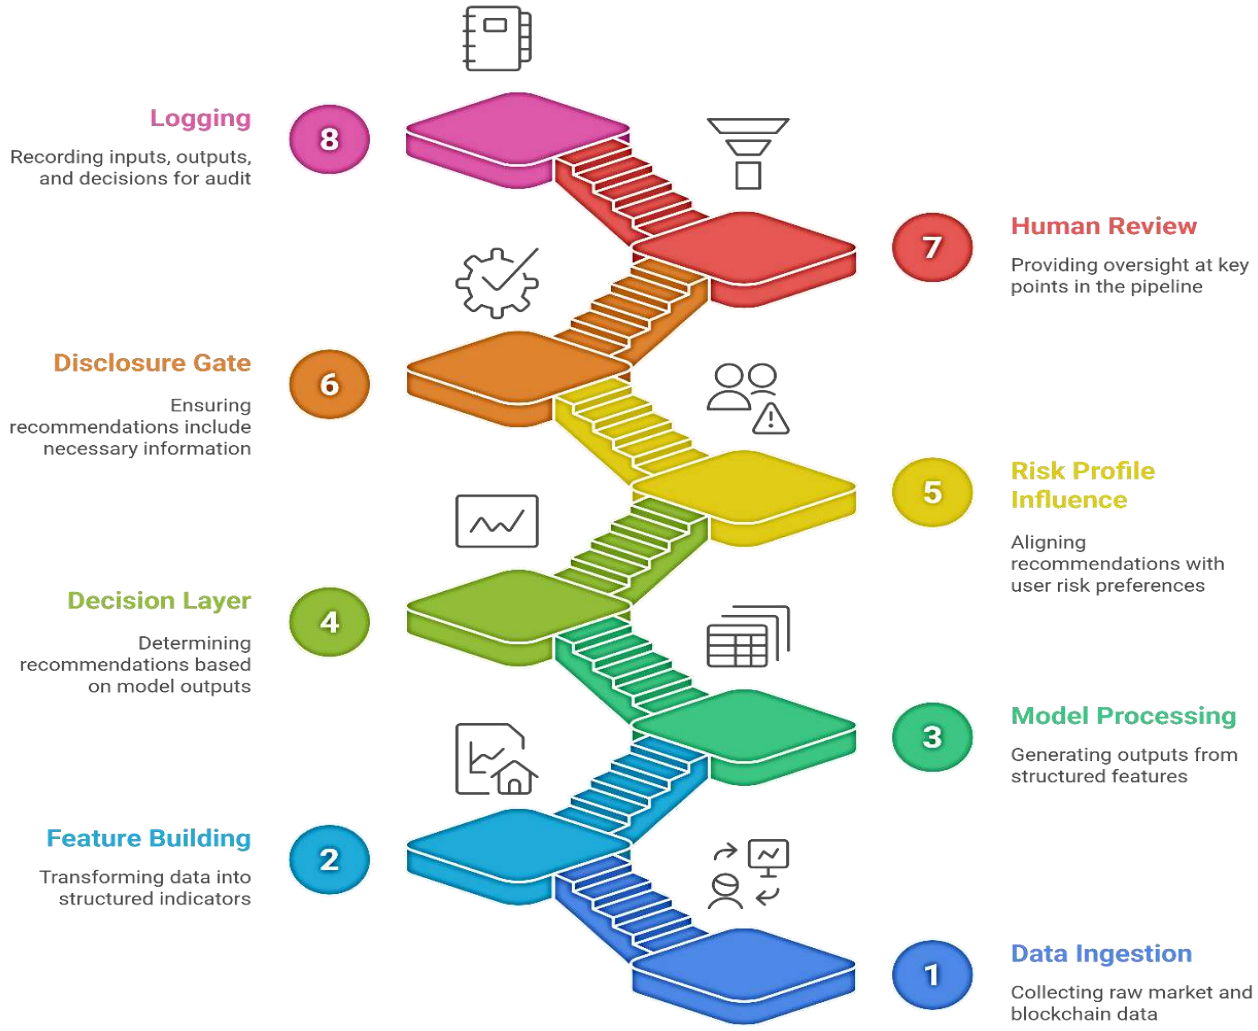

Figure 1: Conceptual pipeline for predictive algorithms in crypto platforms

Each risk dimension maps onto specific phases of the model lifecycle. Dimensions requiring early-stage structural validation (e.g., data provenance, proxy risk, personalization logic) anchor the Research and Validation gates. Mid-lifecycle gates (Staging and Go-Live) address risks related to human oversight, version control, and suitability gating. Post-deployment gates (Monitoring and Audit) address drift, logging integrity, misuse pathways, and privacy constraints. Thus, the control-gate framework operationalizes the risk taxonomy by assigning each dimension to the supervisory stage in which it can be most effectively mitigated (See the below Table 1).

Table 1: Risk taxonomy matrix with qualitative levels

Dimension	Level	Notes
Data provenance and quality	Medium	Requires lineage checks [1]
Feature selection and proxy risk	Medium	Proxy audits advised [11]
Personalization and suitability logic	High	Gate high-risk prompts [8]
Explainability and reviewability	High	Replay and trace [4]
Robustness and overfitting	Medium	Stress tests [1]
Leakage and look-ahead bias	Medium	Time splits required
Privacy and confidentiality	High	Minimization [3] [9]
Trade secret protection	Medium	Limit public detail [6]
Intellectual property constraints	Medium	Ownership clarity [2] [7]
Human oversight and accountability	High	Roles and sign-off [4] [10]
Logging and auditability	High	Immutable logs [4]
Governance and change control	High	Versioning required [4] [6]
Misuse and manipulation pathways	High	Disclosure parity [11]

Explainability ensures that predictive systems disclose the logic underlying ranked lists, prompts, and timed interactions - each of which constitutes a regulated "investor interaction" under the PDA proposal and falls within the antifraud and disclosure duties discussed in Section IV[127]. Providing input families alongside statements in plain language of sensitivity and known limits operationalizes the 2017 Robo-Adviser Guidance's requirement to disclose an algorithm's "methodology, assumptions, and limitations" and prevents omissions that could create misleading impressions under Section 17(a) and Rule 206(4)-1[128]. Input explainability, therefore, functions as a compliance control by establishing a documented basis for reconstructability when investor-facing outputs are later reviewed, something that oversight and validation analogies reinforce through their emphasis on replayability and articulated model boundaries[129].

Model cards, the second criterion, serve as the core artifact through which advisers meet their obligations under the Advisers Act, Rule 206(4)-7, and the PDA Proposal's recordkeeping requirement[130]. By recording purpose, scope, data sources, lineage, known limits, retraining cadence, and gating rules, they supply the evidentiary infrastructure necessary for supervisory review and replayability[131]. Enforcement actions such as Wealthfront's & Hedgeable's cases prove that failures in documentation constitute violations even without intent. Documentation sufficiently draws a boundary between the goal of the system and its actual impact on investor interactions[132].

The third is reviewability, which requires predictive systems to reproduce exact past outputs by rerunning the same inputs, configurations, and model versions that generated a given suggestion for a given user context[133]. This capability directly supports the reconstructability, audit, and post-incident review obligations emphasized in Section IV under the PDA Proposal's recordkeeping guidelines[134]. A replay tool provides this evidentiary chain due to its ability to reconstruct specific inputs using specific data/logic within the automated advisory tool at a given point in time. Without such reproducibility, firms cannot substantiate that a predictive prompt, ranking, or suggestion adhered to disclosed assumptions or suitability constraints, which is crucial for enabling accurate audits, internal investigations, and regulatory responses[135].

Reproducibility alone does not safeguard against violations. Therefore, the fourth criterion is human-in-the-loop checkpoints, which introduce role-based review for prompts or outputs that carry elevated risk/misleading information[136], such as those affecting asset visibility or targeting users with limited profiles. As established in Section IV, disclosure and suitability failures often arise when automated systems operate on incomplete information or generate impressions that exceed their actual capabilities. Human review mitigates these risks by subjecting predictive algorithms to pre-deployment assessment against the system's disclosed methodology, articulated limitations, and suitability parameters, ensuring that outputs with heightened behavioral or financial significance are externally validated before user exposure. The correction mechanism enabled by human oversight in domains allows human review to operate as a structural safeguard against misleading impressions and improperly personalized outputs[137].

The fifth criterion concerns audit trails and access-controlled logging, which establish the core framework of a predictive system's governance architecture. Immutable audit logs provide a tamper-resistant record of the inputs, model version, configuration state, and outputs associated with each interaction. Unlike replay tools, which re-run the model to validate functional behavior, audit trails supply the necessary record to verify that the system state at the time of inference was accurately preserved and has not undergone modification. This integrity is essential for evaluating whether disclosure suitability obligations were satisfied at the moment the interaction occurred[138].

Access control further ensures that the integrity of these logs is preserved. Least-privilege permissions limit visibility of sensitive data inputs, internal configurations, and proprietary model details to only those roles with legitimate supervisory functions. This constraint protects investor privacy by restricting exposure of identifiable or behavioral data, while safeguarding trade secrets by preventing unnecessary dissemination of model architectures, feature families, or configuration states[139].

The sixth criterion is privacy-by-design, which requires that data practices be constrained from the outset through documented minimization, explicit purpose limits, and defined retention bounds. The purpose of this is to avoid unnecessary collection of behavioral or on-chain data. Any linkage between wallet activity and behavioral attributes must be justified by a clear advisory need. The SEC's PDA Proposal warns that excessive data profiling can enable systems to optimize in ways that disadvantage investors[140]. Purpose-limitation principles further require that information gathered for one function - such as wallet-activity monitoring - cannot be repurposed for unrelated tasks, e.g., targeted training prompts. This is also consistent with Regulation S-P's restrictions on the permissible use and safeguarding of customer information[141]. Retention periods must also be aligned with the minimum necessary for system functionality or legal obligations, as overly long retention increases confidentiality risks and can cause stale or outdated profile data to drive unsuitable recommendations or prompts. Privacy controls must be tested at each release gate to ensure these protections remain viable throughout the system life cycle, including any updates to models, inputs, or features that do not expand data collection, alter purpose boundaries, or weaken retention safeguards without review[142].

The seventh is the red team and challenge procedures to probe proxy risk leakage and over-personalization through frozen time splits, shadow deployments, and recorded outcomes. These items draw on validation analogies from adjacent domains and are strictly used as analogies only[143]. In the U.S regulatory context, these procedures support an adviser's obligations under Rule 206(4)-7 to maintain compliance programs capable of identifying and responding to system behaviors that may create misleading or unsuitable interactions[144]. The SEC's 2017 Robo Adviser Guidance similarly emphasizes the importance of a proper understanding of system limitations and failure modes of the algorithms[145]. Red team procedures can diagnose these limitations before they surface and impact investors. The 2023 PDA Proposal further emphasizes that firms must evaluate whether technology used in investor interactions could optimize in ways that disadvantage clients, making challenge testing an important supervisory tool for identifying conflicts of interest within personalization logic[146]. Red-team and challenge exercises serve as preventative mechanisms, allowing firms to identify harmful model pathways before they materialize in investor interactions.

The final criterion concerns incident-response readiness and decommission plans, which specify indicators necessitating rollback, suspension, or retirement and how users/regulators should be informed. Regulation SCI provides a relevant analogue by requiring SCI entities to maintain policies ensuring the "capacity, integrity, resiliency, availability, and security" of critical systems, illustrating the SEC's broader expectations for managing technological failures[147]. The NIST Cybersecurity Framework similarly emphasizes structured response and recovery - requiring organizations to take action upon detecting an incident and restore impaired services[148]. This aligns with the need for defined rollback pathways in predictive models. For a global counterpart, the EU's Digital Operational Resilience Act mandates incident-management and decommissioning procedures for financial ICT systems, offering a comparative model for safe retirement[149]. Together, these frameworks inform a lifecycle approach in which rollback and retirement serve as essential controls for maintaining operational and investor protection[150].

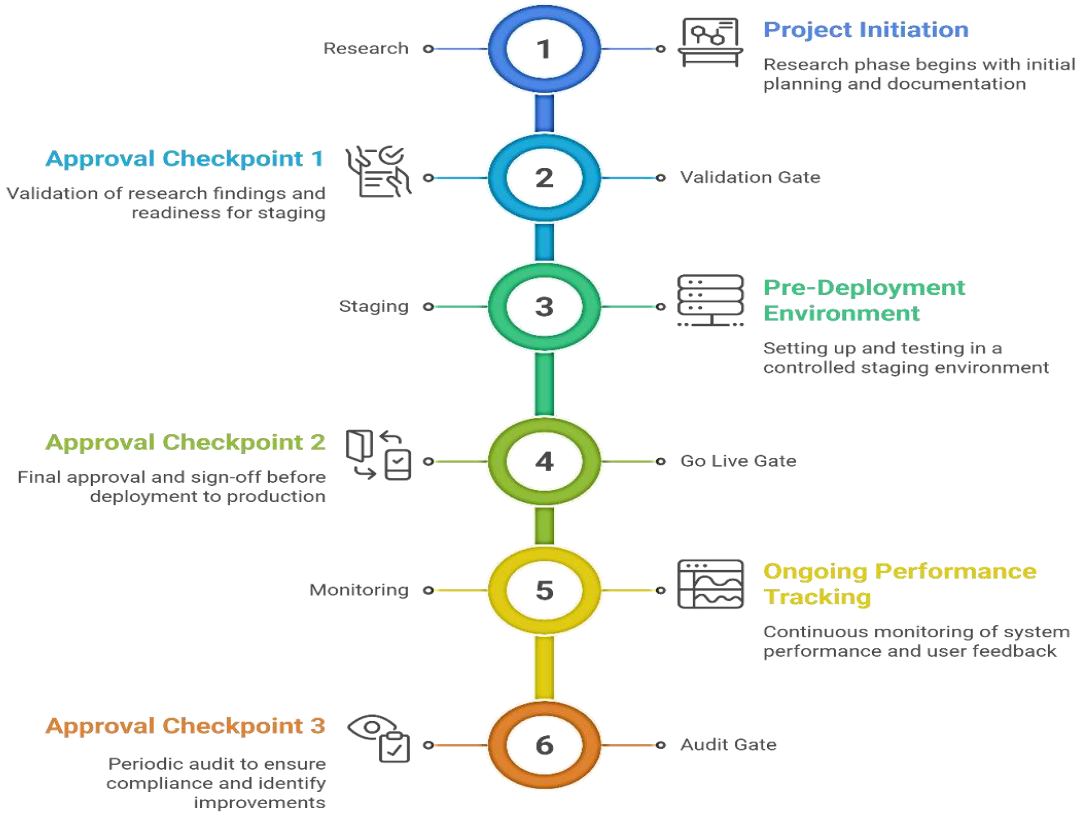

Figure 2: Control gates from research to deployment to monitoring and audit

In the above Figure 2 contributes a structured lifecycle governance model that translates disclosure, suitability, and supervisory obligations into concrete procedural gates. Each gate operationalizes one or more of the paper's evaluation criteria - such as model documentation, replayability, human checkpoints, and audit-log integrity, ensuring these requirements are validated before models progress through staging, deployment, or monitoring. The control gate structure thereby provides a proposed guideline for supervisory accountability and continuous SEC compliance across the model lifecycle (see the below Table 2).

Table 2: Mapping controls to risk dimensions with citations

Control	Risk Dimension	Source
Model card	Explainability, governance	[3] [4] [6] [9]
Version control and rollback	Governance, change control	[4] [6]
Drift and health alerts	Robustness, logging	[4]
Kill switch	Governance, human oversight	[4]
Separation of duties	Accountability, trade secret	[4] [6]
Data lineage and consent	Provenance, privacy	[3] [9]
Disclosure review	Misuse, disclosure parity	[2] [3] [4] [9]
Fairness testing	Proxy risk, equality	[11]
Access control	Trade secret, privacy	[3] [6] [9]
Record retention	Auditability, accountability	[4] [6]

VI. GOVERNANCE AND COMPLIANCE STRUCTURES

Governance must align trade secret and intellectual property constraints with disclosure clarity and suitability fit, because the obligations in Section IV require advisers to explain system behavior even when the underlying logic is proprietary. Under the Advisers Act, responsibility for misleading impressions or unsuitable outputs cannot be shifted to an automated tool; equality and liability sources similarly stress the relationship between accountability, role design, incident response, and documentation structures[151].

Privacy sources add requirements of minimization, purpose limits, user control, and access boundaries, which shape both data engineering and product flows[152]. These boundaries parallel Regulation S-P and mitigate the suitability, proxy-leakage, and drift risks identified in Section V. Trade secret and intellectual property writing supports controlled transparency that describes purpose, feature families, logic categories, and known limits without exposing protected feature formulas or parameter values[153]. This approach mirrors the SEC's demand for disclosures that are sufficiently specific for supervisory understanding, even when models are proprietary. Research on automation in judicial administration further illustrates why replayable traces, audit trails, and staged approvals are essential for oversight[154]. These structures parallel securities-law requirements for reviewable records and compliance-program supervision, showing that accountability depends on lifecycle governance - how version changes are approved, logged, and reconstructed after deployment.

Because personalization functions as a suitability mechanism, governance must ensure that user tiering gates, prompts, and ranking logic remain within suitability boundaries rather than emergent model behavior shaped by personalization logic. Reference flows that map which personalization layers require human review (see the below Figure 3), what data supports suitability gates, and how role-based approvals interact with internal artifacts (model cards, snapshots, audit logs, replay traces) help ensure compliance with the duties articulated in Sections IV and V[155].

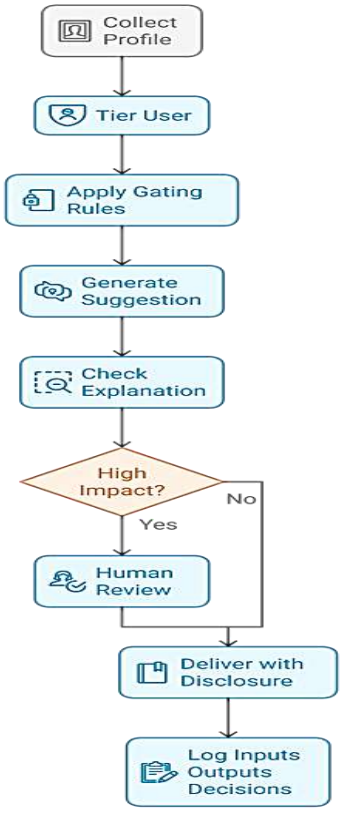

Figure 3: Suitability review flow for personalized recommendations

A consolidated governance checklist enables consistent operational control and audit readiness by translating the legal and technical requirements developed in previous sections into day-to-day supervisory practices. Core elements include model documentation, version control,

monitoring and alerting systems, kill-switch mechanisms, separation of duties, data-lineage verification, consent and purpose review processes, bias and fairness testing, security and trade-secret management, disclosure review and approval, record retention, and customer-suitability procedures (see the below Table 3).

The checklist that follows grounds each item in the sources cited earlier, demonstrating how these controls satisfy disclosure, suitability, privacy, and accountability obligations under the securities law and governance frameworks discussed above.

Table 3: Compliance and governance checklist with grounding

Item	Purpose and link to sources
Model documentation with limits	Clarifies purpose, data sources, lineage, and known limits [3] [4] [6] [9]
Version control and approvals	Records changes and supports rollback [4] [6]
Monitoring and drift alerts	Detects distribution shifts and pipeline faults [4]
Kill switch with scoped access	Stops prompts or models during incidents [4]
Separation of duties	Protects trade secrets and improves accountability [4] [6]
Data lineage and consent flags	Traces provenance and supports minimization and user rights [3] [9]
Disclosure review and parity checks	Reduces selective messaging risk [2] [3] [4] [9]
Bias and fairness testing	Probes proxy risk in features and outcomes [11]
Access control and secret handling	Protects models and features while enabling review [3] [6] [9]
Record retention and replay	Enables audit and incident review [4] [6]
Suitability gating with human checkpoints	Scales prompt with user risk and impact [1] [8] [11]

Three short scenarios stress test the structure:

In this scenario, a ranking model disproportionately weights short-lookback momentum during a market-wide volatility spike. Because short-horizon momentum is highly sensitive to transient price movements, a volatility shock causes the model to elevate a narrow set of tokens to the top of user-facing rankings, even though their recent performance reflects noise rather than stable information. This generates concentrated, unstable surfacing across many user accounts. The resulting misleading impressions raise robustness concerns while potentially failing to meet suitability, oversight, and logging obligations[151].

Under the evaluation criteria developed in Section V, drift and distribution-shift alerts would identify the volatility spike as a model-critical deviation. Replay traces and feature-sensitivity diagnostics would reveal the overweighting of short-horizon momentum and allow supervisors to reconstruct the triggering inputs. Stress testing the model against longer lookback windows and alternative regimes would identify the fragility of the momentum feature. Lastly, a kill-switch mechanism would allow supervisors to suspend prompts or rankings until the behavior is reviewed.

Because the kill-switch affects regulated investor interactions, activation would trigger disclosure and supervisory action discussed in Section IV. Customers would receive clear, non-misleading disclosure notices that certain rankings or prompts have been temporarily disabled, consistent with antifraud obligations under Section 17(a) and the Advisers Act[152], and in compliance with the Guidance's requirement that digital disclosures avoid false impressions about service functionality[153]. It is not necessary to disclose proprietary logic, but users must interpret the suspension of prompts appropriately.

Internally, activation would generate time-stamped supervisory alerts and immutable audit-trail entries documenting the triggering signals, affected model version, and assigned reviewers, in accordance with Rule 206(4)-7's supervisory obligations[154], the PDA Proposal's requirements for conflict-neutralization[155], and recordkeeping obligations[156].

The scenario demonstrates how modeling choices such as overweighting short-run momentum can produce legally significant implications. The system's internal pathways lead to surfacing that risks misleading impressions under the aforementioned regulatory frameworks. The scenario provides a meaningful stress test of the combined risk taxonomy (Section V) and governance architecture outlined in this section, demonstrating the relevance of internal artifacts, lifecycle controls, and supervisory duties and how they interact under realistic market conditions.

Scenario Two concerns selective disclosure, where the system presents upside-down tiles to high-responsiveness cohorts while downplaying equivalent risk language for users with similar financial profiles. This asymmetry raises disclosure-parity, fairness, and privacy concerns because the differential messaging is driven by engagement signals rather than suitability criteria.

Under the legal framework in Section IV, such behavior may violate the "fair and balanced" communication requirements of Rule 206(4)-1[157], and FINRA Rule 2210[158], as well as antifraud duties under the Advisers Act and Section 17(a). This conduct creates misleading impressions through selective emphasis on benefits without providing comparable disclosure of risks. It also risks improper secondary use of behavioral data under Regulation S-P's purpose-limitation requirements[159].

The evaluation criteria in Section V detect this pattern through disclosure-template parity checks, cohort comparisons in audit logs, and replay traces that allow supervisors to confirm whether similarly situated users received conflicting disclosures. Governance structures in Section VI then route the issue to template-approval workflows, risk-sensitive human-in-the-loop review, and role-based controls that block asymmetric templates until reviewed.

User-facing and reviewer-facing alerts would be generated. On the user end, notices would be sent to clients that certain tiles or prompts have been temporarily disabled. On the reviewer end, time-stamped audit-log entries documenting the template version, affected cohorts, and required supervisory approvals would be

implemented, satisfying 206(4)-7's documentation and oversight requirements[160].

This scenario stress-tests the framework's ability to identify disclosure asymmetries, constrain behaviorally targeted messaging, and enforce suitability-aligned and legally compliant investor communication practices.

Scenario Three demonstrates conflicts between SEC-mandated transparency requirements and trade secret & intellectual property protection of proprietary model logic. During a public-facing disclosure, compliance drafts a description to reveal feature constructions and weighting schemes that constitute protected trade secrets. This creates intellectual property risk because competitors could replicate the model, and a secondary privacy risk if feature design indirectly exposes how user data is processed or inferred.

Under the disclosure duties in Section IV, the appropriate response is to provide function-level descriptions and risk explanations[161]. An emphasis on clarifying the model's purpose, assumptions, and limitations while omitting specific feature formulas or parameter values. Governance then requires documenting the review, approval, and access-control decisions associated with the redaction process, ensuring supervisory oversight consistent with Rule 206(4)-7[162] and the controlled-transparency principles described in Section VI[163].

This scenario highlights how transparency duties and IP constraints can come into direct tension, requiring governance processes capable of satisfying both without compromising supervisory visibility.

In the below Table 4, the heatmap summarizes how each scenario concentrates regulatory concerns across the dimensions developed in Sections IV and V. Scenario 1 places the greatest pressure on robustness, suitability, oversight, and logging due to volatility-driven distortions in ranking and surfacing. Scenario 2 raises elevated concerns in disclosure parity, privacy boundaries, and fairness, reflecting the risks of behaviorally targeted messaging. Scenario 3 centers on intellectual-property and controlled-transparency obligations, with secondary implications for privacy and governance. Together, the scenarios illustrate how different failure modes activate different portions of the proposed risk taxonomy and governance framework, and how internal artifacts and supervisory controls enable effective but contextual regulatory responses.

Table 4: Qualitative heatmap of regulatory concern across scenarios

Dimension	Scenario 1	Scenario 2	Scenario 3
Data provenance and quality	M	M	M
Feature selection and proxy risk	M	H	M
Personalization and suitability logic	H	H	M
Explainability and reviewability	H	H	H
Robustness and overfitting	H	M	L
Leakage and look-ahead bias	M	L	L
Privacy and confidentiality	M	H	M
Trade secret protection	L	L	H
Intellectual property constraints	L	L	H
Human oversight and accountability	H	H	H
Logging and auditability	H	H	H
Governance and change control	H	H	H
Misuse and manipulation pathways	M	H	M

The analysis draws on cross-domain scholarship in privacy, trade secret law, intellectual property, accountability, equality, and automated-systems oversight to structure the governance and evaluation framework[164]. Work from adjacent fields - such as health-care validation, judicial automation, and experimental systems testing - provides analogical support for documentation, challenge procedures, and human oversight without serving as empirical evidence about financial markets or securities practice[165]. All legal interpretations reflect analytical inference grounded in the authorities discussed and are not offered as legal advice[166].

VII. DISCUSSION

The findings of this study indicate a practical convergence between platform engineering and the compliance obligations outlined in Section IV. Predictive ranking and personalized prompts shape the informational environment in which disclosure and suitability operate, reinforcing Section V's argument that system-level behavior – rather than user interface text alone – determines regulatory outcomes. Evidence from credit risk and financial equity indicates that model design and documentation decisions materially influence access and performance, supporting the need for explicit governance of ranking and personalization in crypto platforms[167]. Privacy and surveillance journalism warns that routine data collection and secondary use can impose harm, strengthening the case for documented minimization, purpose limits, and role-based access across the full pipeline[168]. Automation scholarship in judicial administration stresses the pivotal importance of replayable decision paths and clearly defined approval structures, which map cleanly onto gates for validation, staging, deployment, monitoring, and audit in platform life cycles developed in Section VI[169]. Together, these strands support a compliance posture in which documentation, replay, approvals, and disclosure review function as core system artifacts rather than late-stage processes.

Data provenance and feature selection require sustained attention, as Section V demonstrates, because they set the conditions under which all later controls operate. Evidence from credit and equity modeling shows that narrow or biased inputs can skew downstream outputs, while clear documentation of lineage and scope improves reviewability and supervisory understanding[170]. Equality writing further demonstrates that seemingly neutral variables can encode proxy effects that systemically disadvantage certain cohorts[171]. In crypto platforms, such

proxy risk can arise even without explicit demographic data: location signals, device patterns, venue mix, response history, and temporal activity can all serve as indirect pathways through which models learn to rank or prompt users differently. Privacy sources reinforce that data collection should remain necessary and proportional to purpose, and that retention and access should be constrained by policy and role[172]. Accordingly, proxy audits should examine families of features rather than isolated fields and should be repeated after material model updates or cohort shifts, focusing on what the model can infer rather than solely on what the database stores. The SEC's growing attention toward predictive data analytics underscores the regulatory significance of these issues; however, some scholars warn that overly broad rules risk sweeping in benign computational tools and generating compliance uncertainty for emerging digital finance markets[173]. Together, these considerations reinforce Section IV's point that antifraud and suitability duties apply equally to the internal construction of feature sets as to their outward presentation.

Reflecting these concerns, the SEC has expanded its enforcement architecture by creating a Crypto Assets and Cyber Unit dedicated to addressing crypto-asset misconduct and technology-mediated fraud, including the use of AI and machine learning systems[174]. The Commission's launch of "Project Crypto" further signals an effort to articulate clearer regulatory expectations for digital-asset platforms[175].

Explainability and reviewability are central to accountability and to suitability review. As Section V emphasizes, these functions allow supervisors to assess whether a platform's internal decision pathways align with the disclosure and suitability duties described in Section IV. Oversight work in administrative settings indicates that institutional trust depends on the ability to reconstruct a decision using the same inputs, configuration, model version, and the same approval path[176]. In financial recommendation scenarios, the same logic supports replay tools that bind a suggestion to its input bundle and to the user risk profile that governed gating. Health-care surveys and predictive modeling research provide analogies for staged approvals, documentation templates, and validation records that can be adapted to financial-product pipelines without making claims about financial effects[177]. A platform that cannot reproduce a past suggestion on demand will struggle to demonstrate suitability fit for disclosure sufficiency, even if public-facing text is clear. Replay, therefore, serves as the anchor for proportionate accountability, linking the screen seen by the user to the evidence available to the reviewer.

Privacy-by-design interacts with personalization in a way that requires discipline rather than abstention. Surveillance sources emphasize minimization and purpose limits, while platforms often seek increasingly granular, highly detailed profiles to enhance the perceived relevance of personalized rankings and prompts[178]. As the governance structure in Section VI suggests, the appropriate response is to enforce minimization and purpose limits both at the point of collection and within the feature registry, ensuring that model builders query only approved analytical views. Trade-secret journalism reinforces the architecture by requiring that protected features and transformations remain behind role-based access controls, with disclosures to internal reviewers subject to formal approvals and with clear prohibitions on public release[179]. Intellectual property writing further supports function-level transparency – describing purpose, logic categories, and known limits – while avoiding exposure of proprietary formulas or parameter settings[180]. These boundaries do not weaken disclosure provided the platform supplies clear statements of purpose, input-family definitions, gating rules, and known limits to users and reviewers, paired with internal records that contain the confidential technical details necessary for audit.

Incident response and accountability benefit from a small number of well-practiced playbooks. As section IV makes clear, criminal-liability sources make clear that automated systems do not erase responsibility, which in turn necessitates explicit assignment of roles and approval authority before and after deployment[181]. Oversight literature further supports the usage of immutable logs and documented approvals as the backbone of audit and post-incident review procedures[182]. Equality writing suggests that monitoring should incorporate checks for differential effects across cohorts and for parity in disclosure templates when cohorts are similarly situated[183]. Red-team and challenge procedures from Section V also inform incident response by providing pre-deployment evidence of failure modes, leakage pathways, or over-personalization patterns that may later trigger supervisory intervention. A compact response program can therefore address the recurring failure modes identified in the scenarios: Ranking instability during volatility spikes requires drift detection and a kill-switch for prompt families; selective messaging across cohorts requires parity checks in disclosure templates and cohort-based log reviews. Tension between transparency and trade secrets suggests a review that confirms the adequacy of function-level descriptions for users and the presence of protected technical detail in internal records.

Robustness failures, leakage, and look-ahead bias arise when models retrain frequently or consume signals that inadvertently incorporate forward-looking information. Credit-risk literature warns that overfitting to short windows and relying on leaked signals undermines validity and produces unstable outputs[184]. Clinical predictive work often proposes analogies for drift checks and frozen-time splits that test performance on data unavailable during training[185]. As Sections V and VI emphasize, these controls map directly onto the platform life cycle: validation gates are positioned before deployment to block models that fail stability or leakage checks, while audit gates follow monitoring to pause or retire the models that fail drift in production. A simple rule is to permit deployment only when performance advantages hold across frozen splits and representative cohorts – providing a reliable safeguard against spurious gains and helping satisfy the disclosure and suitability duties in Section IV. If such stability cannot be demonstrated, the model should be stopped at the gate, not deferred to later monitoring, because post-deployment remediation cannot undo misleading impressions already created for investors.

Suitability depends on how user tiering and gating rules interact with personalization logic. As Sections IV and V establish, the sequence that moves from profile collection through tiering, gating, explanation, human review, and

logging is what aligns personalized prompts with legal expectations for fit and clarity at the point of decision. Equality sources emphasize the need for disclosure parity to prevent asymmetries that advantage engagement over risk awareness when users are similarly situated[186]. Privacy sources reinforce checks for consistent purpose and access across cohorts, ensuring that personalization does not drift into unauthorized secondary use[187]. The qualitative heatmap derived from the scenarios does not claim to measure probability or magnitude. It functions as a structural tool for directing attention to suitability, robustness, privacy, and accountability in Scenario one and to disclosure parity and privacy in Scenario two, and to trade secret and intellectual property boundaries in Scenario three. Governance teams should treat the heatmap as a standing agenda for release reviews and post-incident assessments, ensuring that each category of risk is systematically examined rather than inferred ad hoc.

Measurement without external datasets is both a constraint and a methodological strength for internal audits because, as Section V outlined, the evaluation criteria proposed rely on artifacts that the platform can generate throughout development and operation. Explainability is measured through input family descriptions and sensitivity notes in the model card. Documentation sufficiency is evaluated through the completeness of model cards, purpose, scope, lineage, limits, retraining cadence, privacy elements documenting minimization and retention, and gating rules[188]. Reviewability is assessed through successful replay using identical inputs, configurations, and model versions. Human oversight is measured via recorded approvals at validation and go-live gates, while auditability is measured by immutable logs with role-based access. Red team and challenge procedures are measured by evidence of frozen time splits, shadow deployments, and recorded outcomes. Incident readiness is demonstrated through an established playbook and the implementation of periodic drills. Health-care and experimental sources provide only process analogies for these checks but do not claim financial outcomes[189]. These measures do not claim to prove or guarantee improved investor outcomes; rather, they supply a legible, reproducible account of the system's behavior – one that reviewers and regulators can follow to evaluate whether disclosure, suitability, and supervisory duties have been met.

Organizational incentives often pull against transparency and documentation. Trade-secret and intellectual property protection concerns may be invoked to resist disclosure, and privacy obligations may be cited to justify limited logging. However, the cited sources throughout this paper point to a different equilibrium. Privacy is strengthened not by the absence of records but by strict purpose limits, minimization, and role-based access controls[190]. Accountability is advanced through replayable paths and recorded approvals, rather than informal operational practice[191]. Intellectual-property interests are preserved through controlled transparency that informs users about function, purpose, and risk while maintaining protected technical details in secure internal records[192]. Equality goals are supported by parity checks and proxy audits rather than assurances or intuition[193]. When applied together, these principles reinforce the governance architecture developed in Section IV: programs that publish a clear set of artifacts for each release improve internal coordination by ensuring that engineers, product teams & managers, compliance personnel, and legal reviewers operate from a shared and reviewable source of truth.

The transferability from health-care and experimental literature remains limited. These sources demonstrate how complex and safety-sensitive domains employ staged approvals, standardized documentation, and prospective challenges to reduce error and support institutional learning[194]. They do not resolve questions of financial impact, legal classification, or market integrity. They do, however, offer mature process patterns – particularly around validation, documentation, human oversight, and decommission planning – that can be adapted to platform settings without overstating their normative reach. Future work should examine whether these process patterns reduce measurable harm in crypto contexts through empirical user studies and market data, and should examine how disclosure and suitability outcomes shift in terms of meeting compliance when platforms adopt the artifacts proposed here.

International Regulators have also confronted similar issues and have developed distinct regulatory approaches. Regulators in the European Union, acting under the Markets in Financial Instruments Directive, already apply fiduciary and suitability obligations to automated investment advisors in portfolio management, affirming – consistent with the cross-jurisdictional comparisons established in Sections IV and V – that these duties attach regardless of whether advice is delivered by a human or an algorithm[195]. The European Securities and Markets Authority has further emphasized the need for greater transparency regarding the role of AI in investment decision-making[196]. In the United Kingdom, regulators apply the "Know your Customer" requirement to algorithmic advisory services, treating automated tools as extensions of adviser obligations[197].

Canada permits robo-advisors and algorithmic models within investment platforms; however, it ultimately lends superior precedence and hence responsibility to human advisors, holding them accountable for any resulting conflicts of disclosure, suitability, or investor protection[198]. Across jurisdictions, the trend is consistent: regulators reaffirm investor protection, disclosure, fairness, transparency, suitability, and principles of equity as principles governing automated advisory systems.

From a legal standpoint, AI and generative models are evolving far faster than regulatory frameworks can evolve to adapt. As Section IV indicates, the SEC and similarly situated regulators exhibit slow doctrinal uptake: the most recent comprehensive robo-adviser guidance dates to 2017, predating generative AI, modern predictive analytics, and deep-learning-based personalization. This temporal gap reflects a broader structural challenge in financial regulation, where obligations grounded in traditional advisory models must be applied to a system whose technical capabilities and behavioral impacts now exceed the assumptions embedded in existing rules[199].

Taken together, the analysis in this section demonstrates that effective disclosure and suitability in predictive crypto-investment platforms cannot be achieved through interface-level, front-end text alone. As Sections IV–VI show, meaningful compliance requires integrated controls

across data provenance, feature selection, model logic, gating decisions, human review, and logging[200]. When implemented together, these components allow a platform to demonstrate how a suggestion was generated, why it passed the relevant gates, how it was presented to the user, and how the system would be paused, corrected, or retired when risks arise. The privacy, trade-secret, intellectual property, equality, criminal-liability, and oversight sources collectively provide the conceptual foundation for designing these controls in a manner that remains reviewable and regulatorily legible[201]. Validation and documentation analogies from health care and related safety-critical fields further supply process structures that may be adapted to platform settings without overstating their normative reach[202].

VIII. LIMITATIONS

This study remains qualitative in scope and uses internal artifacts - such as model cards, input snapshots, logs, and replay traces - as its primary unit of analysis. As established in Sections IV–VII, these artifacts are essential for evaluating disclosure and suitability obligations in predictive systems; however, the framework does not empirically test whether improved disclosures alter user behavior or measurably reduce investor harm. Nor does it quantify how ranking distortions affect market prices, liquidity, or volatility. The analysis does not classify crypto-assets or assess platform status under any specific jurisdictional regime, and its legal framing relies on general principles drawn from antifraud, suitability, and supervisory doctrines rather than jurisdiction-specific determinations. Because the evaluation criteria are designed for document review, code inspection, and logged system behavior, they do not require external datasets and consequently do not provide statistical estimates of risk reduction. These boundaries reflect the paper's conceptual purpose: to articulate how disclosure and suitability duties can be interpreted through internal system mechanics - not to make predictive claims about market effects or regulatory outcomes[209].

The reference set used in this study is intentionally scoped toward domains that illuminate privacy governance, equality, accountability, intellectual property, and trade-secret constraints, all of which inform the disclosure and suitability frameworks developed in Sections IV–VII[210]. These sources clarify why internal mechanisms such as data minimization, parity checks, auditability, and controlled transparency are central to determining whether predictive interactions create misleading impressions or produce suitability misalignments[211]. In contrast, the study does not incorporate market-microstructure literature or securities-trading releases, and thus makes no claims regarding execution quality, spreads, depth decay, or order-routing behavior. Health-care and experimental sources are used solely as analogies for process structures - such as staged validation, documentation standards, human oversight, and challenge procedures - and their transferability is limited where incentives, error modes, or time scales diverge from financial-platform environments[212]. The taxonomy's qualitative labels (e.g., low, medium, high) reflect conceptual levels of supervisory concern rather than estimates of probability, magnitude, or expected loss. This orientation reinforces the paper's central focus: assessing how system-level behaviors interface with legal expectations about transparency, suitability, and supervisory accountability, rather than predicting market impacts or quantifying investor harm.

The scenarios in this study are illustrative stress tests rather than empirical accounts; they model plausible failure modes without making claims about frequency or real-world incidence. Because platform architectures differ substantially - some centralizing data engineering and others embedding it within product teams - the proposed controls may require distinct role configurations or governance pathways in practice. The analysis is tailored to consumer-facing platforms that rely on personalization and point-of-decision prompts; institutional environments that emphasize portfolio tools rather than retail prompts may exhibit different risk frontiers and thus interact differently with the disclosure and suitability pathways developed in Sections IV–VII. The framework does not examine cross-venue conflicts, broker-dealer obligations, or the full spectrum of international privacy and consumer-protection regimes, all of which could alter the compliance parameters described earlier[213]. Section V documents that only the provided sources were used; where bibliographic confirmation was unavailable, the item remains as listed, reflecting the conceptual nature of the study[214]. These boundaries underscore that the scenarios are designed to test the internal logic of the proposed taxonomy and governance structure, not to provide jurisdiction-specific predictions or claims of regulatory completeness.

The framework's viability depends on the availability of replayable traces, immutable logs, and role-based approval structures[215]. Without stable identifiers for data versions, model versions, and configuration snapshots, any replay mechanisms - and thus any meaningful supervisory reconstruction — will fail. This study does not attempt to resolve that infrastructural gap; rather, it identifies what a reviewer or regulator would reasonably expect to see within a framework grounded in the disclosure, suitability, and accountability principles analysed in Sections IV–VII[216]. The checklist and gating structures proposed in Section VI can support the development of these artifacts, but they cannot substitute for missing observability within the underlying engineering environment. Moreover, the study does not evaluate the cost of implementing such controls, which may be significant for smaller teams or early-stage platforms. These constraints highlight that, although enforcement risk may vary, the process artifacts emphasized throughout the paper remain essential for demonstrating how internal system behavior aligns with legal expectations and for enabling retrospective review when investor-facing interactions are challenged[217].

IX. CONCLUSION

Predictive ranking and personalization sit at the center of crypto platform engagement. Disclosure and suitability require clarity and fit at the point of decision. Meeting those duties through front-end text alone is unlikely to work. The results support an integrated approach that places governance within data provenance, feature selection, model logic gating, human review, and logging[219]. The figures and tables provide a shared vocabulary for engineering, compliance, and legal teams.

The pipeline figure anchors technical flows. The life cycle figure locates validation, monitoring, and audit gates. The suitability figure links tiering, gating, explanation, and delivery with logging. The taxonomy table identifies where regulatory concern concentrates. The control mapping table shows how documentation, parity checks, lineage, and access control mitigate those risks[220]. The heatmap table helps teams prioritize attention across scenarios.

The paper contributes three practical elements. First, a risk taxonomy that ties provenance, proxy risk, personalization, explainability, robustness, leakage, privacy, trade secrets, intellectual property, human oversight, governance, and misuse pathways to reviewable artifacts[221]. Second, an evaluation design that uses documents, code, and logs, and that can be scored without external datasets. Third, a governance structure with a checklist that links model documentation, version control, monitoring, kill switch, separation of duties, data lineage, consent review, bias and fairness testing, access control, record retention, and suitability review. These elements are grounded in the provided sources and marked as author analysis, where legal interpretation is offered.

The framework does not claim to fix market failures or to predict enforcement outcomes. It provides a way to show how a suggestion reached a screen, why it passed gates, what disclosure accompanied it, and how the process would be paused, corrected, or retired when risks arise. This record helps reviewers understand what happened and why. It also helps builders design with privacy, equality, accountability, intellectual property, and trade secret constraints in mind[222]. The same artifacts can be used for internal training and for post-incident learning. Importantly, the articulation of reviewable records also strengthens alignment with the legal section of regulatory regimes, demonstrating how compliance requirements can be mapped to technical safeguards in practice.

Future work should build empirical evidence on several fronts. Teams can measure how disclosure parity affects behavior in matched cohorts and whether replayable traces shorten incident resolution time. Studies can test whether drift gates reduce the rate of prompt withdrawals after volatility spikes. Work can compare audit readiness across organizations that maintain model cards, lineage records, and parity checks, and organizations that do not. Research can examine how function-level transparency affects user trust while preserving trade secrets. Cross-jurisdiction studies can test the portability of the artifacts under different privacy and consumer law regimes[223]. Shared templates for model cards, disclosure text, parity audits, and replay logs could support comparability across platforms and foster a baseline of responsible practice.

Finally, as predictive algorithms in crypto investment platforms evolve, their governance cannot be separated from legal considerations, particularly in relation to court precedents and emerging case law on disclosure duties, suitability requirements, and investor protection[224]. Embedding these legal insights within technical governance ensures that compliance frameworks remain robust, adaptable, and defensible in both regulatory reviews and judicial scrutiny.

In summary, crypto investment platforms using predictive algorithms are not operating in a legal void; rather, they are touched by various SEC rules and regulations if they deal in securities and or give investment advice. The existing SEC regimes apply holistically. The SEC approach suggests a way in which evolution in the space must be balanced with investor protection. The SEC has started to address this through new rule-making, guidance, and enforcement[225].

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

FOOTNOTES

1. U.S. Sec. & Exch. Comm'n (SEC), Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990, Advisers Act Release No. IA-6353 (July 26, 2023) (PDA Proposal); Shlomit Wagman, Algorithmic Consumer Manipulation, 95 N.Y.U. L. Rev. 563 (2020); U.S. Sec. & Exch. Comm'n (SEC), In re Wealthfront Advisers, LLC, Advisers Act Release No. IA-5086 (Dec. 21, 2018).
2. SEC, Guidance Update: Robo-Advisers, Div. of Inv. Mgmt., No. 2017-02 (Feb. 2017); Marcos López de Prado, Advances in Financial Machine Learning (Wiley 2018); U.S. Sec. & Exch. Comm'n (SEC), In re Hedgeable, Inc., Advisers Act Release No. IA-5087 (Dec. 21, 2018).
3. PDA Proposal, supra note 1, at 20–24; Rory Van Loo, The Emerging Automated Administrative State, 58 Harv. J. on Legis. 55 (2021); Wagman, supra note 1, at 586–90.
4. Ahmed Raza, Credit, Code, and Consequence: How AI Is Reshaping Risk Assessment and Financial Equity, 2 EuroVantage J. Artif. Intell. 79 (2025); López de Prado, supra note 2, at 53–61; Ahmed Raza et al., Artificial Intelligence and Criminal Liability: Rethinking Criminal Liability in the Era of Automated Decision Making, 2 Int'l J. Contemp. Issues Soc. Sci. 1529 (2023).
5. PDA Proposal, supra note 1, at 12–17; Directive 2014/65/EU of the European Parliament and of the Council (MiFID II), art. 25; Ahmed Raza, AI and Privacy: Navigating a World of Constant Surveillance, 1 EuroVantage J. Artif. Intell. 74 (2024).
6. Regulation S-P: Privacy of Consumer Financial Information and Safeguarding Customer Information, Exchange Act Release No. 34-100155 (May 16, 2024); Wagman, supra note 1; Van Loo, supra note 3, at 70–73.
7. Ahmed Raza et al., Automation in Judicial Administration: Evaluating the Role of Artificial Intelligence, Int'l J. Contemp. Issues Soc. Sci. 1544 (2023); SEC, Guidance Update: Robo-Advisers, supra note 2, at 3–4.
8. PDA Proposal, supra note 1, at 20–24; U.S. Sec. & Exch. Comm'n (SEC), In re Wealthfront Advisers, LLC, Advisers Act Release No. IA-5086 (Dec. 21, 2018).
9. U.S. Sec. & Exch. Comm'n (SEC), In re Hedgeable, Inc., Advisers Act Release No. IA-5087 (Dec. 21, 2018); PDA Proposal, supra note 1, at 27–32; SEC, Guidance Update: Robo-Advisers, supra note 2, at 3–5.
10. SEC, Guidance Update: Robo-Advisers, supra note 2, at 2–4.
11. Directive 2014/65/EU (MiFID II), arts. 24–25; U.K. Financial Conduct Authority, Guidance on Algorithmic Trading Compliance (2022).
12. Van Loo, supra note 3, at 78–83.
13. PDA Proposal, supra note 1, at 10–12.
14. Ahmed Raza et al., Automation in Judicial Administration: Evaluating the Role of Artificial Intelligence, Int'l J. Contemp. Issues Soc. Sci. 1544, 1547–49 (2023).

15. Michidmaa Arikhad et al., AI-Powered Solutions for Precision Healthcare: Focusing on Heart and Brain Disorders, 15 Rev. Inteligencia Artificial Med. 1278 (2024).
16. SEC, Guidance Update: Robo-Advisers, supra note 2, at 4–6.
17. PDA Proposal, supra note 1, at 28–30.
18. PDA Proposal, supra note 1, at 12–18; Wagman, supra note 1; U.S. Sec. & Exch. Comm'n (SEC), In re Wealthfront Advisers, LLC, Advisers Act Release No. IA-5086 (Dec. 21, 2018).
19. SEC, Guidance Update: Robo-Advisers, supra note 2, at 3–6; PDA Proposal, supra note 1, at 28–35; Van Loo, supra note 3, at 70–73.
20. Regulation S-P, supra note 6, at 34–39; Directive 2014/65/EU (MiFID II), art. 25; Raza, Credit, Code, and Consequence, supra note 4, at 82–85.
21. Marcos López de Prado, Advances in Financial Machine Learning 17–25 (Wiley 2018).
22. SEC, Guidance Update: Robo-Advisers, supra note 2, at 2–4.
23. López de Prado, supra note 21, at 43–68.
24. Van Loo, supra note 3, at 66–70.
25. López de Prado, supra note 21, at 95–118.
26. SEC, Guidance Update: Robo-Advisers, supra note 22, at 5–7.
27. Regulation S-P, supra note 6, at 45–49.
28. SEC, Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, Rel. No. IA-6389 (July 2023), at 13–17.
29. Id. at 22–28.
30. SEC, Guidance Update: Robo-Advisers, supra note 22, at 7–10.
31. PDA Proposal, supra note 1, at 30–35.
32. U.S. Sec. & Exch. Comm'n (SEC), Guidance Update: Robo-Advisers, Div. of Inv. Mgmt., No. 2017-02 (Feb. 2017).
33. U.S. Sec. & Exch. Comm'n (SEC), Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990, Advisers Act Release No. IA-6461 (July 26, 2023).
34. U.S. Sec. & Exch. Comm'n (SEC), Regulation S-P: Privacy of Consumer Financial Information, Final Rule, Exchange Act Release No. 34-100155 (May 16, 2024).
35. U.S. Sec. & Exch. Comm'n (SEC), In re Wealthfront Advisers, LLC, Advisers Act Release No. IA-5086 (Dec. 21, 2018).
36. U.S. Sec. & Exch. Comm'n (SEC), In re Hedgeable, Inc., Advisers Act Release No. IA-5087 (Dec. 21, 2018).
37. European Securities and Markets Authority, MiFID II: Article 25 — Assessment of Suitability and Appropriateness and Reporting to Clients (2014).
38. U.K. Financial Conduct Authority, Guidance on the Use of Digital Engagement Practices by Investment Platforms (2022).
39. Ahmed Raza, The Application of Artificial Intelligence in Credit Risk Evaluation: Obstacles and Opportunities in the Path to Financial Justice, 3 Ctr. for Mgmt. Sci. Rsch. 240 (2025).
40. Raza, Credit, Code, and Consequence, supra note 4.
41. A. Raza et al., Balancing Privacy and Technological Advancement in AI: A Comprehensive Analysis of the U.S. Perspective, 3 Int'l J. Contemp. Issues Soc. Sci. 3732 (2024).
42. Raza, AI and Privacy, supra note 5.
43. A. Raza et al., Automation in Judicial Administration, supra note 14.
44. M. A. Chohan et al., Artificial Intelligence and Intellectual Property Rights: From Content Creation to Ownership, 2 Int'l J. Soc. Scis. Bull. 2514 (2024).
45. Ahmed Raza, Trade Secrets as a Substitute for AI Protection: A Critical Investigation into Different Dimensions of Trade Secrets, 2 Int'l J. Soc. Scis. Bull. 2506 (2024). Ahmed Raza & N. Bashir, Artificial Intelligence as a Creator and Inventor: Legal Challenges and Protections in Copyright, Patent, and Trademark Law, 2 Int'l J. Contemp. Issues Soc. Sci. 1598 (2023).
46. Ahmed Raza et al., Artificial Intelligence and Criminal Liability, supra note 4.
47. A. Raza, Equality Before Law and Equal Protection of Law: Contextualizing Its Evolution in Pakistan, Pak. L.J. (2023).
48. M. Arikhad et al., AI-Driven Innovations in Cardiac and Neurological Healthcare: Redefining Diagnosis and Treatment, 19 Rev. Esp. Doc. Cient. 124 (2024).
49. M. Arikhad et al., The Role of Artificial Intelligence in Advancing Heart and Brain Disease Management, 19 Rev. Esp. Doc. Cient. 137 (2024).
50. H. Rafi et al., Transforming Cardiovascular and Neurological Care with AI, 15 Rev. Inteligencia Artificial Med. 1264 (2024).
51. M. Arikhad et al., AI-Powered Solutions for Precision Healthcare, supra note 49.
52. H. Khan et al., Revolutionizing Heart and Brain Healthcare with Artificial Intelligence, 15 Rev. Inteligencia Artificial Med. 1289 (2024).
53. S. Rasool et al., AI-Based Predictive Tools for Managing and Preventing Cardiovascular Diseases, 13 Int'l J. Innovative Res. Computer Sci. & Tech. 76 (2025).
54. A. M. Khaja et al., Predictive Modeling for Chemotherapy Response Using Machine Learning, 13 Int'l J. Innovative Res. Computer Sci. & Tech. 62 (2025).
55. H. Rafi et al., Comparative Effectiveness of Agmatine and Choline Treatment…, 15 Curr. Clin. Pharmacol. 251 (2020).
56. H. Rafi & M. Farhan, Dapoxetine: An Innovative Approach in Therapeutic Management in an Animal Model of Depression, 2 Pak. J. Pharm. Sci. 15 (2015).
57. I. Bhatti et al., Use of ICT Technologies for the Assistance of Disabled Migrants in the USA, 18 Rev. Esp. Doc. Cient. 66 (2024).
58. SEC, SEC Investor Protection Mission, U.S. Sec. & Exch. Comm'n.
59. SEC v. Capital Gains Research Bureau, Inc., 375 U.S. 180 (1963).
60. Investment Advisers Act Rule 206(4)-1.
61. FINRA Rule 2111.
62. Investment Advisers Act Rule 206(4)-7.
63. U.S. Sec. & Exch. Comm'n (SEC), Sample Letter to Companies Regarding Crypto Asset Disclosures (Dec. 2022).
64. U.S. Sec. & Exch. Comm'n (SEC), Guidance Update: Robo-Advisers (2017).Securities Act of 1933 § 17(a), 15 U.S.C. § 77q(a).
65. See SEC v. Capital Gains Research Bureau, Inc., 375 U.S. 180, 191–92 (1963).
66. PDA Proposal, supra note 1, at 12–18.
67. PDA Proposal, supra note 1, at 30–35.
68. PDA Proposal, supra note 1, at 5.
69. PDA Proposal, supra note 1, at 14–16.
70. FINRA Rule 2210(d)(1)(A).
71. PDA Proposal, supra note 68, at 22–24.
72. SEC, Proposed Rule: Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990, at 184–87 (July 26, 2023).
73. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 4 (Feb. 2017).
74. Id. at 5–6.
75. Id. at 2–3.
76. Id. at 7–9.
77. Id. at 7.
78. Id. at 8.
79. Id. at 9–10.
80. Id. at 2–3.
81. Id. at 4–5.
82. Id. at 4.
83. Id. at 5–6.

84. Id. at 6.
85. Id. at 6–7.
86. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 2 (Feb. 2017).
87. Id. at 3.
88. Id. at 3–4.
89. Id. at 1–2.
90. Id. at 8.
91. Id. at 8–9.
92. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 8 (Feb. 2017).
93. Id. at 4–5.
94. Id. at 8–9.
95. Id. at 9.
96. Id. at 9–10.
97. Id. at 10–11.
98. Id. at 9–11.
99. Id. at 11–12.
100. U.S. Sec. & Exch. Comm'n (SEC), In re Wealthfront Advisers LLC, Advisers Act Release No. IA-5085, at 1 (Dec. 21, 2018).
101. U.S. Sec. & Exch. Comm'n (SEC), In re Hedgeable, Inc., Advisers Act Release No. IA-4960, at 1 (Apr. 17, 2018).
102. In re Wealthfront Advisers LLC, supra note 103, at 2–4.
103. Id. at 4–6.
104. Id. at 6–7.
105. Id. at 7–8.
106. Aaron v. SEC, 446 U.S. 680, 696–97 (1980).
107. In re Hedgeable, Inc., supra note 104, at 2–3.
108. Id. at 3–5.
109. Id. at 3.
110. Id. at 3–4.
111. Id. at 4–5.
112. Id. at 5–6.
113. In re Wealthfront Advisers LLC, supra note 103, at 4–7; In re Hedgeable, Inc., supra note 104, at 3–6.
114. See In re Wealthfront Advisers LLC, supra note 116; In re Hedgeable, Inc., supra note 116.
115. See In re Wealthfront Advisers LLC, supra note 116, at 7–8.
116. Arrieta et al., Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges, 58 Info. Fusion 82 (2020); Barocas & Selbst, Big Data's Disparate Impact, 104 Calif. L. Rev. 671 (2016).
117. Mitchell et al., Model Cards for Model Reporting, in Proc. Conf. Fairness, Accountability & Transparency 220 (2019); Raji et al., Audit Logging for Accountability in Automated Decision Systems, FAT* (2020).
118. Hardt et al., Equality of Opportunity in Supervised Learning, in Advances in Neural Info. Processing Sys. 3315 (2016); Goodman & Flaxman, European Union Regulations on Algorithmic Decision-Making and a "Right to Explanation," 38 AI & Soc. 35 (2017).
119. Khandani, Kim & Lo, Consumer Credit-Risk Models via Machine Learning Algorithms, 72 J. Banking & Fin. 301 (2010); Kleinberg et al., Prediction Policy Problems, 81 Am. Econ. Rev. Papers & Proc. 98 (2015).
120. See Aaron v. SEC, 446 U.S. 680 (1980); Advisers Act § 206(4).
121. Directive 2014/65/EU (MiFID II), art. 25 (2014).
122. Raza, AI and Privacy, supra note 5.
123. Goodman & Flaxman, supra note 121.
124. Chohan et al., Artificial Intelligence and Intellectual Property Rights, supra note 44; Raza, Trade Secrets, supra note 45; Raza & Bashir, Artificial Intelligence as a Creator and Inventor, supra note 46.
125. SEC, Proposed Rule: Conflicts of Interest Associated With the Use of Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990, at 120–25 (2023).
126. Raza et al., Automation in Judicial Administration, supra note 14.
127. See supra–Section V.
128. SEC, Proposed Rule: Conflicts of Interest Associated With the Use of Predictive Data Analytics, supra note 128, at 23–28.
129. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 4 (Feb. 2017).
130. Van Loo, supra note 3; Arikhad et al., supra note 49; Raza et al., supra note 14; López de Prado, supra note 21; SEC, Guidance Update: Robo-Advisers, supra note 2.
131. Investment Advisers Act Rule 206(4)-7; SEC, Exchange Act Release No. 34-97990, supra note 131, at 52–59.
132. Wagman, supra note 1; Regulation S-P, supra note 6; Raza, AI and Privacy, supra note 5.
133. In re Wealthfront Advisers LLC, supra note 103, at 4–7; In re Hedgeable, Inc., supra note 104, at 3–6.
134. Raza et al., Automation in Judicial Administration, supra note 14.
135. SEC, Exchange Act Release No. 34-97990, supra note 131, at 60–65.
136. Id.; Mitchell et al., supra note 120; Raji et al., supra note 120.
137. Van Loo, supra note 3; Aaron v. SEC, supra note 109.
138. See supra–Section V.
139. SEC, Proposed Rule: Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990, at 60–65 (2023).
140. Raza, AI and Privacy, supra note 5; Raza et al., Automation in Judicial Administration, supra note 14; Raza & Bashir, Artificial Intelligence as a Creator and Inventor, supra note 46.
141. SEC, Proposed Rule: Predictive Data Analytics by Broker-Dealers and Investment Advisers, Exchange Act Release No. 34-97990, at 40–47 (2023).
142. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 4–6 (Feb. 2017).
143. SEC, Exchange Act Release No. 34-97990, supra note 142, at 90–98.
144. Regulation SCI, 17 C.F.R. pts. 240, 242, 249.
145. NIST, Framework for Improving Critical Infrastructure Cybersecurity (2018).
146. Regulation (EU) 2022/2554 (Digital Operational Resilience Act).
147. NIST, supra note 148; Regulation (EU) 2022/2554, supra note 149.
148. See Hardt et al., supra note 121; Goodman & Flaxman, supra note 121.
149. See Raza, AI and Privacy, supra note 5; Wagman, supra note 1.
150. See Chohan et al., supra note 44; Raza, Trade Secrets, supra note 45; Raza & Bashir, supra note 46.
151. See Raza et al., Automation in Judicial Administration, supra note 14.
152. See supra-Sections IV–V.
153. Advisers Act Rule 206(4)-7; SEC, Proposed Rule: Predictive Data Analytics, Exchange Act Release No. 34-97990, at 40–47 (2023).
154. Securities Act § 17(a); Advisers Act § 206.
155. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 3–6 (Feb. 2017).
156. Advisers Act Rule 206(4)-7.
157. SEC, Exchange Act Release No. 34-97990, supra note 156, at 90–98.
158. Id. at 52–59.
159. Advisers Act Rule 206(4)-1.
160. FINRA Rule 2210.
161. Regulation S-P, 17 C.F.R. pt. 248.
162. Advisers Act Rule 206(4)-7.
163. SEC, Guidance Update: Robo-Advisers, No. 2017-02, at 4 (Feb. 2017).
164. Ahmed Raza, Credit, Code, and Consequence: How AI Is Reshaping Risk Assessment and Financial Equity, 2

EuroVantage J. Artif. Intell. 79 (2025); Marcos López de Prado, Advances in Financial Machine Learning (2018).

165. Ahmed Raza, AI and Privacy: Navigating a World of Constant Surveillance, 1 EuroVantage J. Artif. Intell. 74 (2024); Shlomit Wagman, Algorithmic Consumer Manipulation, 95 N.Y.U. L. Rev. 563 (2020).
166. Ahmed Raza et al., Automation in Judicial Administration: Evaluating the Role of Artificial Intelligence, Int'l J. Contemp. Issues Soc. Sci. 1544 (2023).
167. Raza, Credit, Code, and Consequence, supra note 167; López de Prado, supra note 167.
168. Hardt et al., supra note 121.
169. Raza, AI and Privacy, supra note 168; Wagman, supra note 168.
170. Michidmaa Arikhad et al., AI-Powered Solutions for Precision Healthcare, 15 Rev. Inteligencia Artificial Med. 1278 (2024).
171. SEC, Crypto Assets and Cyber Unit, https://www.sec.gov/divisions/enforcement/crypto-assets-and-cyber-unit.
172. SEC, Project Crypto, https://www.sec.gov/litigation/special-studies/project-crypto.
173. Raza et al., Automation in Judicial Administration, supra note 169.
174. Arikhad et al., supra note 173; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
175. Wagman, supra note 168; Raza, AI and Privacy, supra note 168.
176. Regulation S-P, Exchange Act Release No. 34-100155 (2024).
177. Raza & Bashir, supra note 46; SEC, Guidance Update: Robo-Advisers, supra note 166.
178. Aaron v. SEC, 446 U.S. 680 (1980).
179. Raza et al., Automation in Judicial Administration, supra note 169.
180. Goodman & Flaxman, supra note 121.
181. López de Prado, supra note 167; Raza, Credit, Code, and Consequence, supra note 167.
182. Regulation S-P, supra note 179; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
183. Goodman & Flaxman, supra note 121.
184. Wagman, supra note 168; Raza, AI and Privacy, supra note 168.
185. Wagman, supra note 187; Raza et al., Automation in Judicial Administration, supra note 182; Regulation S-P, supra note 179; Raza, AI and Privacy, supra note 168.
186. Rory Van Loo, The Emerging Automated Administrative State, 58 Harv. J. on Legis. 55 (2021);
187. Conflicts of Interest Associated With the Use of Predictive Data Analytics, Exchange Act Release No. 34-97990, Advisers Act Release No. IA-6353 (2023);
188. SEC, Guidance Update: Robo-Advisers, supra note 166; Regulation S-P, supra note 179; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
189. Wagman, supra note 168; Raza, AI and Privacy, supra note 168.
190. Raza et al., Automation in Judicial Administration, supra note 169; Aaron v. SEC, supra note 181.
191. IM Guidance Update: Robo-Advisers, SEC Div. Inv. Mgmt., No. 2017-02 (Feb. 2017); Regulation S-P, supra note 179; Raza & Bashir, supra note 46.
192. Goodman & Flaxman, supra note 121.
193. Van Loo, supra note 189; Conflicts of Interest Associated With the Use of Predictive Data Analytics, supra note 189; Arikhad et al., supra note 173; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
194. Directive 2014/65/EU (MiFID II); ESMA & EBA, Guidelines on Suitability.
195. European Securities & Markets Authority, Artificial Intelligence in EU Securities Markets: Supervisory Considerations.
196. U.K. Financial Conduct Authority, Know Your Customer (KYC) and Suitability Requirements.
197. Canadian Securities Administrators, CSA Staff Notice on Robo-Advisers.
198. SEC, Guidance Update: Robo-Advisers, No. 2017-02 (Feb. 2017).
Wagman, supra note 168; Raza et al., Automation in Judicial Administration, supra note 169; Regulation S-P, supra note 179; Raza & Bashir, supra note 46; Aaron v. SEC, supra note 181; Goodman & Flaxman, supra note 121.
199. Van Loo, supra note 189; Conflicts of Interest Associated With the Use of Predictive Data Analytics, supra note 189; Regulation S-P, supra note 179; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
200. See supra-Sections IV–VII.
201. Wagman, supra note 168; Raza, AI and Privacy, supra note 168; Goodman & Flaxman, supra note 121; Regulation S-P, supra note 179; Raza & Bashir, supra note 46; Aaron v. SEC, supra note 181.
202. See id.; see also Raza et al., Automation in Judicial Administration, supra note 169.
203. Van Loo, supra note 189; Conflicts of Interest Associated With the Use of Predictive Data Analytics, supra note 189; Arikhad et al., supra note 173; Regulation S-P, supra note 179; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
204. Wagman, supra note 168; Raza, AI and Privacy, supra note 168.
205. See supra–Section V.
206. Wagman, supra note 168; Raza et al., Automation in Judicial Administration, supra note 169; Raza, AI and Privacy, supra note 168; Aaron v. SEC, supra note 181; Goodman & Flaxman, supra note 121.
207. Regulation S-P, supra note 179; Raza & Bashir, supra note 46; Wagman, supra note 168; Raza, AI and Privacy, supra note 168; Goodman & Flaxman, supra note 121; Raza et al., Automation in Judicial Administration, supra note 169.
208. Wagman, supra note 168; Raza et al., Automation in Judicial Administration, supra note 169; Raza, AI and Privacy, supra note 168.
209. López de Prado, supra note 167; Wagman, supra note 168; Raza et al., Automation in Judicial Administration, supra note 169; Raza, AI and Privacy, supra note 168; Goodman & Flaxman, supra note 121.
210. Wagman, supra note 168; Raza, AI and Privacy, supra note 168; Goodman & Flaxman, supra note 121; Raza et al., Automation in Judicial Administration, supra note 169; Regulation S-P, supra note 179; Raza & Bashir, supra note 46.
211. Van Loo, supra note 189; Wagman, supra note 168; Raza et al., Automation in Judicial Administration, supra note 169; Regulation S-P, supra note 179; Raza & Bashir, supra note 46; Goodman & Flaxman, supra note 121; Aaron v. SEC, supra note 181.
212. Wagman, supra note 168; Raza, AI and Privacy, supra note 168; Goodman & Flaxman, supra note 121; Aaron v. SEC, supra note 181; Regulation S-P, supra note 179; Raza & Bashir, supra note 46.
213. Van Loo, supra note 189; Conflicts of Interest Associated With the Use of Predictive Data Analytics, supra note 189; Regulation S-P, supra note 179; López de Prado, supra note 167; SEC, Guidance Update: Robo-Advisers, supra note 166.
214. Raza et al., Automation in Judicial Administration, supra note 169; Aaron v. SEC, supra note 181; Goodman & Flaxman, supra note 121.
215. SEC, Guidance Update: Robo-Advisers, No. 2017-02 (Feb. 2017); Conflicts of Interest Associated With the Use of Predictive Data Analytics, supra note 189; SEC, Crypto Assets and Cyber Unit, supra note 174; SEC, Project Crypto, supra note 175.

REFERENCES

[1] R. Ahmed Raza, "Credit, code, and consequence: How AI is reshaping risk assessment and financial equity," *EuroVantage Journal of Artificial Intelligence*, vol. 2, p. 79, 2025. Available from: https://evjai.com/index.php/evjai/article/view/30

[2] M. A. Chohan *et al*., "Artificial intelligence and intellectual property rights: From content creation to ownership," *International Journal of Social Sciences Bulletin*, vol. 2, p. 2514, 2024. Available from: https://tinyurl.com/yau78cdd

[3] R. Ahmed Raza, "AI and privacy: Navigating a world of constant surveillance," *EuroVantage Journal of Artificial Intelligence*, vol. 1, p. 74, 2024. Available from: https://evjai.com/index.php/evjai/article/view/28

[4] R. Ahmed Raza *et al*., "Automation in judicial administration: Evaluating the role of artificial intelligence," *International Journal of Contemporary Issues in Social Sciences*, p. 1544, 2023. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5209960

[5] H. Rafi and M. Farhan, "Dapoxetine: An innovative approach in therapeutic management in an animal model of depression," *Pakistan Journal of Pharmaceutical Sciences*, vol. 2, p. 15, 2015. Available from: https://tinyurl.com/bdeswcrk

[6] R. Ahmed Raza, "Trade secrets as a substitute for AI protection: A critical investigation into different dimensions of trade secrets," *International Journal of Social Sciences Bulletin*, vol. 2, p. 2506, 2024. Available from: https://tinyurl.com/mvkwswwm

[7] R. Ahmed Raza *et al*., "Artificial intelligence as a creator and inventor: Legal challenges and protections in copyright, patent, and trademark law," *International Journal of Contemporary Issues in Social Sciences*, vol. 2, p. 1598, 2023. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5376048

[8] R. Ahmed Raza, "The application of artificial intelligence in credit risk evaluation: Obstacles and opportunities in path to financial justice," *Center for Management Science Research*, vol. 3, p. 240, 2025. Available from: https://tinyurl.com/5t55347c

[9] Raza, B. Munir, G. Ali, M. A. Othi, and R. A. Hussain, "Balancing privacy and technological advancement in AI: A comprehensive analysis of the U.S. perspective," *International Journal of Contemporary Issues in Social Sciences*, vol. 3, p. 3732, 2024. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5198259

[10] R. Ahmed Raza, "Navigating the intersection of artificial intelligence and law in healthcare: Complications and corrections," *Policy Review Journal*, vol. 2, p. 2364, 2024. Available from: https://tinyurl.com/43dezt5m

[11] Raza, "Equality before law and equal protection of law: Contextualizing its evolution in Pakistan," *Pakistan Law Journal*, 2023. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5200799

[12] M. Arikhad, M. Waqar, A. H. Khan, and A. Sultana, "AI-driven innovations in cardiac and neurological healthcare: Redefining diagnosis and treatment," *Revista Española de Documentación Científica*, vol. 19, p. 124, 2024. Available from: https://tinyurl.com/2s3kf37a

[13] M. Arikhad, M. Waqar, A. H. Khan, and A. Sultana, "The role of artificial intelligence in advancing heart and brain disease management," *Revista Española de Documentación Científica*, vol. 19, p. 137, 2024.

[14] M. Arikhad *et al*., "AI-powered solutions for precision healthcare: Focusing on heart and brain disorders," *Revista de Inteligencia Artificial Médica*, vol. 15, p. 1278, 2024. Available from: https://tinyurl.com/488j89x7

[15] M. Arikhad *et al*., "Transforming cardiovascular and neurological care with AI: A paradigm shift in medicine," *Revista de Inteligencia Artificial Médica*, vol. 15, p. 1264, 2024.

[16] Y. Hayat *et al*., "Artificial intelligence in early detection of neurological disorders," *International Journal of Innovative Research in Computer Science & Technology*, vol. 13, p. 67, 2025. Available from: https://tinyurl.com/3eun3xus

[17] S. Rasool, A. M. Khaja, Y. Hayat, and A. H. Khan, "AI-based predictive tools for managing and preventing cardiovascular diseases," *International Journal of Innovative Research in Computer Science & Technology*, vol. 13, p. 76, 2025. Available from: https://tinyurl.com/wjrb3nsv

[18] M. Khaja, M. Arikhad, Y. Hayat, and S. Rasool, "Predictive modeling for chemotherapy response using machine learning," *International Journal of Innovative Research in Computer Science & Technology*, vol. 13, p. 62, 2025. Available from: https://tinyurl.com/3fwve96m

[19] H. Rafi *et al*., "Comparative effectiveness of agmatine and choline treatment in rats with cognitive impairment induced by aluminum chloride and forced swim stress," *Current Clinical Pharmacology*, vol. 15, p. 251, 2020. Available from: https://doi.org/10.2174/1574884714666191016152143

[20] M. López de Prado, *Advances in Financial Machine Learning*, 1st ed., Hoboken, NJ, USA: Wiley, 2018, ch. 1.Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3104847

[21] R. Van Loo, "The emerging automated administrative state," *Harvard Journal on Legislation*, vol. 58, p. 55, 2021. Available from: https://www.jstor.org/stable/27024713

[22] S. Wagman, "Algorithmic consumer manipulation," *New York University Law Review*, vol. 95, p. 563, 2020. Available from: https://tinyurl.com/4wba78hz