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.2025.13.6.18
Article ID IJIRE-1421, Pages 169-186
www.ijircst.org

Algorithmic Gatekeeping in Crypto Markets: Governing AI-Driven Market Access in Centralized Exchange

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

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

Correspondence should be addressed to Waqar Ahmad; ahmadtarar05@gmail.com

Received: 2 November 2025 Revised: 17 November 2025 Accepted: 30 November 2025

Copyright © 2025 Made Waqar Ahmad et al. This is an open-access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

ABSTRACT- Centralized crypto exchanges increasingly govern market access through AI-dominated, vendor-supplied infrastructure that performs core gatekeeping functions, including transaction surveillance, risk scoring, listing and delisting decisions, and account restrictions. These systems exercise authority comparable to that of regulated financial intermediaries yet operate within a fragmented U.S. regulatory landscape in which jurisdictional uncertainty, enforcement constraints, and reliance on case-by-case litigation have limited the development of binding oversight. This paper argues that the resulting governance gap is primarily a consequence of the absence of institutional frameworks capable of disciplining automated market access as infrastructure. Drawing on recent enforcement actions, vendor concentration practices, and comparative analysis of the European Union's Digital Operational Resilience Act (DORA), the paper demonstrates that effective governance of opaque, automated infrastructure is possible and already implemented in other countries through procedural accountability, documentation, and human oversight, rather than technical explainability. The paper contributes a DORA-inspired governance framework focused on internal controls, auditability, and clearly assigned human responsibility for AI-mediated decision-making within centralized exchanges. The analysis does not seek to resolve underlying classification disputes but instead identifies the legal and institutional conditions under which automated gatekeeping could be brought within a regime of enforceable oversight.

KEYWORDS— Algorithmic Governance, Centralized Crypto Exchanges, Digital Market Access, Financial Infrastructure, Risk-Scoring Systems

I. INTRODUCTION

Centralized Crypto Exchanges (CEX) have a rising dependency on AI-driven and automated back-end infrastructure to perform core governance functions. These systems - often supplied by a small number of third-party vendors - shape which tokens are listed, restricted, deprioritized, or excluded, and which users are permitted to access markets [1],[2]. In practice, decisions about market participation, liquidity, and regulatory exposure are no longer made solely through institutional judgment, but through opaque, automated processes embedded deep within exchange infrastructure. This shift has created the concept of "algorithmic gatekeepers. Existing regulatory frameworks are not designed with these gatekeepers in mind, especially when they govern market access.

As this paper argues, the resulting opacity is further amplified by regulatory fragmentation and the grey area in which crypto exchanges operate. These exchanges are treated as accountable intermediaries in principle [3] but operate with minimal oversight over the automated systems that execute compliance in practice. Risk models may over-classify benign activity, assign disproportionate risk profiles to new or smaller tokens - due to AI models inherently being less efficient with minority classes - or negatively affect market integrity, yet neither developers nor regulators possess meaningful mechanisms to interrogate or reconstruct these determinations. [4]

Where the same surveillance and risk-scoring infrastructure is deployed across multiple exchanges, private vendors effectively exercise system-wide influence over market access without corresponding public-law obligations. [1] Against this backdrop, the paper examines how the EU's preventative regulatory framework - particularly under Digital Operational Resilience Act (DORA) - demonstrates that governance of AI-dominated, third-party infrastructure is both legally feasible and structurally constraining, focusing on an ex-ante approach towards potential algorithmic harm. [5]

The paper is split into a historical overview of the progression of the crypto market, leading to the rise of centralized exchanges and the regulatory landscape surrounding such exchanges. It aims to examine the legal difficulties of regulating crypto markets due to the nature of exchange platforms and examines the crossroads between AI-driven infrastructure and evolving, adaptable regulatory processes better suited to handle the evolving crypto market.

Building on this analysis, the paper aims to address the presence of algorithmic gatekeeping in the current U.S. crypto market, the divided governance stance among the U.S. regulators, the struggles to implement a centralized regulatory system, and the current landscape of CEX regulation, as well as the legal perspectives - wherein CEXs are concerned - shown thus far through the Securities and Exchange Commission's (SEC) enforcement cases. [3] It calls for structural governance

while avoiding hindering innovation or market growth through its core contributions, namely, a proposed legal governance framework grounded in the principles of the EU's Digital Operational Resilience Act that imposes accountability, auditability, and investor-protection-oriented governance on AI-dominated, third-party exchange infrastructure. [5]

II. LITERATURE REVIEW AND IDENTIFIED RESEARCH GAP

A. Existing Literature and The Function of Centralized Exchanges

Early scholarship on the crypto market identifies centralized exchanges as a natural response to the technical constraints of peer-to-peer blockchain systems. Economic and institutional analyses document how early decentralized trading environments struggled with slow settlement, complicated interfaces, and a need for technical literacy to enable regular usage. CEXs were developed to provide continuous trade while offering custodial services that enabled broader participation in crypto markets [6],[2].

Subsequent literature describes centralized exchanges as dominant market intermediaries that control trade execution and asset custody within their own platforms. [3] These exchanges perform roles like traditional financial intermediaries. They coordinate order flow, operate internal order books, and decide which assets may be listed and traded. Empirical studies also show that centralized exchanges concentrate market power. Since they control liquidity and transaction processing, they can influence price formation and, in some cases, enable market manipulation [7],[2].

Current scholarship provides valuable insight into the organizational and economic role of centralized exchanges; it largely treats exchange governance as a structural or market-design issue. While this is a defensible perspective, there is a lack of literature that examines how the infrastructure of centralized exchanges enables market governance. The internal technical systems through which exchanges implement surveillance, listing decisions, and compliance controls are typically assumed rather than examined. There is a lack of literature that frames the increasing presence of automated back-end systems as an active enabler that establishes market governance.

Although scholars acknowledge that exchanges function as gatekeepers over participation, the mechanisms through which such gatekeeping is operationalized - particularly through algorithmic and AI-driven infrastructure - require further analysis. This gap becomes increasingly significant as exchanges rely less on direct human discretion and more on automated systems to perform these core functions.

B. Usage & Technical Constraints of Automation Systems in Financial Markets

A separate body of literature examines the growing role of algorithmic and machine-learning systems in financial markets. Research in this area shows that automated systems frequently outperform human actors in environments characterized by large datasets, rapid decision cycles, and narrow profit margins. Studies on algorithmic trading and financial automation demonstrate that machine-driven strategies can detect patterns, execute trades, and respond to market signals with greater speed and consistency than human operators [8],[9].

Beyond trading, scholars have documented the expansion of machine-learning systems into compliance, surveillance, and risk-assessment functions. These systems are used to classify transactions, detect anomalous behavior, and assign risk scores to assets or participants. Technical research highlights their capacity to process vast volumes of transactional data and identify behavioral patterns that would be difficult for human reviewers to detect at scale [10].

At the same time, this literature identifies persistent limitations in machine-learning systems. It mainly emphasizes issues such as dataset bias, class imbalance, and reduced accuracy for smaller or less-represented categories. When models are trained on uneven or historically skewed data, they tend to perform better for majority classes while misclassifying minority or novel actors. These limitations are not incidental. They are structural features of how machine-learning systems are developed and deployed [4],[10].

Legal and policy-oriented scholarships have begun to examine the challenges raised by automated financial systems regarding accountability. This work focuses on the difficulty of assigning responsibility when decisions are made through complex technical processes rather than direct human judgment. Scholars note that existing regulatory frameworks were designed around identifiable human decision-makers and struggle to accommodate systems whose internal logic is opaque and difficult to reconstruct [11].

Despite these insights, the literature largely treats algorithmic systems as tools that assist human-decision making rather than tools that now perform functions which structure market participation- a role historically assumed by human actors. The use of AI in finance is typically framed in terms of efficiency, accuracy, or risk detection, with less attention paid to how automated systems can shape access, visibility, and outcomes within markets. As a result, the governance implications of delegating market-shaping functions to machine-learning systems remain underdeveloped in existing scholarships.

C. Surveillance, Compliance, and Risk-Scoring Infrastructure

A related strand of scholarship documents how machine-learning tools are deployed to support anti-money laundering, transaction monitoring, and customer-risk assessment across financial institutions. However, this analysis is broad - covering general financial institutions, such as banks, brokerage firms, or moneylending platforms such as Upstart. It does not critically examine the underlying infrastructure present within centralized exchanges [11].

Technical research highlights both the strengths and limitations of such systems. Machine-learning models can process large volumes of transactional data and identify correlations that exceed human capacity. However, scholars emphasize that these systems are shaped by the data on which they are trained and the objectives embedded in their design. As a result, automated surveillance systems often reproduce structural biases,

generate false positives, and disproportionately affect smaller or less-represented actors [4].

Studies of algorithmic harm further show that classification errors are not evenly distributed. At least at present, models trained on imbalanced datasets largely perform better when they work on dominant or well-represented categories. In financial contexts, this can look like inflated risk scores, unwarranted compliance flags, or even market access-shaping judgements such as exclusion from services. These effects are the result of systemic features inherent in machine-learning pipelines [4]. In the context of this paper, the current technical research presents infrastructural risks inherent to machine-learning pipelines; however, the effect on consumer protection has not been researched in depth.

D. Legal Scholarship on Crypto Regulation and Comparative Governance Frameworks

Legal scholarship addressing crypto markets has largely focused on the application of existing securities law frameworks to digital assets, centering on the Howey Test [12] and the classification of tokens as securities, with attention to promoter conduct, the actual economic function of said assets, and purchaser expectations. Scholars tend to debate whether digital assets can be classified as securities within the scope of established U.S. securities law. [13] They criticize the regulation via enforcement approach that has so far been adopted by regulatory bodies such as the SEC, which prioritizes a case-by-case strategy in terms of regulating centralized exchanges as unregistered securities exchanges. [3] These precedents primarily treat regulatory questions as issues of jurisdiction and asset classification rather than as matters of infrastructure governance.

Scholars highlight the absence of a comprehensive statutory framework governing centralized exchanges and note the difficulty of expanding decades-old-securities rhetoric to rapidly evolving technical systems [11]. They rarely address how exchanges are operated internally. As a result, there is a lack of relevant legal analysis of these technical mechanisms within the current U.S. landscape.

In contrast, the increasing body of scholarship analyzes European frameworks for the regulation of automated infrastructure in digital finance. The Digital Operational Resilience Act focuses on how financial entities manage digital systems in practice. In terms of what it defines as proper management, the Resilience Act requires clear documentation, auditability, and most importantly, human responsibility for automated systems. Instead of focusing on classifying assets, DORA aims to regulate how institutions design, use, and oversee digital infrastructure. This includes systems supplied by third-party vendors [5]. A preventative, overall approach such as this makes DORA particularly relevant to this analysis, as it directly addresses the governance risks created when financial entities rely on outsourced, vendor-supplied technological systems.

The chief takeaway from DORA for this analysis is the way it creates traceable accountability through a step-by-step governance system. By implementing clear processes supported by internal and external reviews, it ensures that the final deciding factor is always human governance - not automated gatekeeping. The paper draws inspiration from this governance ideology to construct its own proposed framework.

E. Current Research Gaps and Relevance of this Analysis

In summation, there is indisputable literature that documents the structure of present-day exchange and the risks of AI systems; however, relatively little scholarship examines the relationship between centralized exchanges, third-party AI infrastructure, and market access.

Present scholarships offer a holistic overview of the general landscape of centralized exchanges. However, to mitigate risk & harm on the consumer end in the U.S., there must be relevant analysis into what possible shapes regulation and governance may take to ensure consumer protection.

This paper addresses that gap by examining how automation within centralized exchanges reshapes market access and accountability. The paper also proposes a governance framework informed by comparative regulatory practices. By bringing the role of automated infrastructure within centralized exchanges into sharper focus, the paper contributes to existing scholarship by framing these systems as a form of de facto governance tools whose risks and consequences warrant direct regulatory attention.

III. RESEARCH DESIGN AND METHODOLOGY

This paper employs a qualitative, doctrinal, and comparative research design to examine how AI-dominated infrastructure within centralized crypto exchanges reshapes market access and governance.

Rather than conducting empirical testing or technical evaluation of machine-learning models, the analysis focuses on legal structure, institutional practice, and governance implications arising from the deployment of automated systems in exchange operations.

The research draws on three primary categories of sources. First, it relies on relevant legal and regulatory materials, including U.S. securities statutes, enforcement actions, and agency guidance, to establish the current regulatory posture governing centralized exchanges. These materials are used descriptively to identify how existing legal frameworks conceptualize exchange functions, managerial responsibility, and market access, rather than to assert settled jurisdictional conclusions.

Second, the paper incorporates academic literature from economics, computer science, and law to contextualize the technical capabilities and limitations of AI-driven systems used in financial markets. This literature is used to illustrate how machine-learning systems operate in practice, including issues related to opacity, bias, and error propagation. The analysis does not treat these sources as empirical proof of misconduct, but as evidence of the structural features and risks inherent in automated decision-making systems.

Third, the paper adopts a comparative governance approach by examining the European Union's Digital Operational Resilience Act. DORA is used as a reference framework to demonstrate what governance can look like for digital financial infrastructure and as a contrast to the challenges present within U.S. governance. The paper does

not argue for direct adoption of European law within the United States but uses DORA to illustrate the institutional form that binding oversight could take.

Methodologically, the paper applies a functional analysis to centralized exchanges and their internal systems. It evaluates exchange conduct based on the roles automated systems play in surveillance, listing, and exclusion. This approach aligns with existing securities-law analysis, which prioritizes economic reality and functional effect over nominal classification.

The scope of the analysis is intentionally limited but not disengaged from securities law. The paper does not aim to definitively resolve whether digital assets constitute securities or whether centralized crypto exchanges are presently subject to registration requirements. Instead, it engages existing securities-law doctrine, particularly the Howey framework [13] and the SEC's enforcement posture [3] - as an analytical lens for evaluating the economic and legal reality of exchange conduct. This approach allows the paper to examine how centralized exchanges increasingly perform functions analogous to traditional securities platforms without asserting that the surrounding classification questions are legally settled.

A central feature of this research design is its reliance on functional and analogical reasoning. Due to the limitations of modern legal frameworks regarding AI-governance in CEX infrastructures, much of the analysis is rooted in analogy to existing regulatory doctrines and institutional models. These analogies do not intend to make the claim that automated crypto exchanges are legally equivalent to traditional securities exchanges, or that existing legal categories inherently must apply to them. Instead, they are used to a similar economic function, governance effects, and the related market impact. By comparing automated exchange infrastructure to historically regulated financial intermediaries, the paper aims to establish regulatory blind spots created by delegating core governance duties to technological systems rather than collapsing established legal distinctions.

IV. FROM DECENTRALIZED VISION TO CENTRALIZED REALITY: EVOLUTION OF CRYPTO MARKET INFRASTRUCTURE AND ITS REGULATION

A. Peer-to-Peer Platforms & Early Barriers to Trading

The original proposition of crypto-exchange systems, most notably the white papers of Bitcoin and Ethereum, was to facilitate direct peer-to-peer transactions without the need for a trusted third party [14], [15]. While these texts did not address securities regulation directly, their design logic implicitly rejected centralized entities performing functions analogous to those of national securities exchanges, including order matching, listing control, and trading suspension.

However, the implementation of this model imposed substantial practical and technical burdens on users. During the 2009-2013 period, Bitcoin transactions commonly occurred through informal and decentralized channels, such as online forums, peer-to-peer arrangements, and early marketplace platforms. [6] These environments were not conceived as permanent market structures, but as ad hoc arrangements facilitating occasional trades among technically sophisticated users.

What this early period reveals is not a failure of the peer-to-peer ideology but a misalignment between the functional demands of widespread market participation and the limits of decentralized platforms. The model assumed a level of technical literacy, settlement tolerance, and self-custodial capacity that most users could not realistically meet - preventing routine usage. As a result, the early ecosystem lacked infrastructural stability and accessibility characteristic of mature trading venues. These foundational constraints served as the backdrop against which subsequent institutional and technological developments, including the emergence of centralized exchanges, would take shape, ultimately reshaping the trajectory of crypto-market infrastructure in ways that subsequent sections will explore.

B. The Development of Centralized Exchanges

Therefore, user demand for fiat on-ramps and more accessible trading interfaces far exceeded the capacity of early peer-to-peer systems, leading to the rise of Mt. Gox and Bitcoin Market. Unlike their traditional decentralized alternatives, these platforms offered advanced features such as fiat integration, instant account set-up, and user-friendly interfaces, capturing most of the trading volume by 2013 and reflecting a decisive shift toward intermediary-based market infrastructure [6], [2].

During its prime, Mt. Gox became the de facto market infrastructure for the crypto ecosystem [2]. By performing the classic functions of a national securities exchange - coordinating order flow, pairing buyers and sellers through non-discretionary rules, and exerting control over which assets could trade - Mt. Gox and its successors functionally mirrored the conduct that triggers registration obligations under Section 6 of the Securities Exchange Act of 1934, yet none ever registered - a regulatory gray area further highlighted in Section 1.3. [3], [16]. The centralization of these exchange functions was reinforced not only through organizational control but through the technical mechanisms that were embedded within centralized exchanges. Their automated infrastructure rendered these platforms indispensable by providing the operational reliability and usability that decentralized, peer-to-peer mechanisms were structurally unable to deliver.

Mt. Gox operated as an early centralized Bitcoin exchange with automated, exchange-based trade execution, enabling continuous matching of buy and sell orders at scale [17] - a core functional characteristic of conventional securities exchanges. Subsequent exchanges, including Bittrex, Coinbase, and Binance, adopted the same model, institutionalizing automated trade execution as an integral component of centralized exchange architecture [18], [19].

In the literature, centralized exchanges are structurally defined by the internalization of trade execution and asset custody within automated, intermediary-operated platforms [2], [6]. Although implemented through simpler automated systems in the industry's early stages, centralized exchanges quickly embedded non-discretionary gatekeeping functions over market participation. Today, these same functions, particularly asset admission, transaction monitoring, and risk classification, are increasingly executed through and governed by automated systems [1], [4].

The SEC has consistently maintained that the legal consequence of meeting the Howey criteria is unchanged. This regulatory friction would later be formalized in the SEC's framework for "Investment Contract" Analysis of Digital Assets and operationalized through the Bittrex enforcement action discussed below [12], [13], [3].

V. THE FUNCTIONAL REALITY OF CENTRALIZED EXCHANGES

A. Functional Parallels between Centralized Exchanges and Regulated Market Intermediaries

The current landscape of crypto exchanges operates outside of the registration-based compliance framework set forth by Section 6 of the Securities Exchange Act of 1934, which states that entities performing the core functions of a national securities exchange, including, but not limited to, the pairing of buyers and sellers, establishing listing practices, or suspending and delisting securities, must first register with the SEC [16]. Crypto exchanges have generally avoided registration by asserting that the digital assets they list do not constitute securities [3].

This position is addressed directly in the SEC's Framework for "Investment Contract" Analysis of Digital Assets. Firstly, it establishes the fundamental compliance measure that when engaging in the offer, sale, or distribution of digital assets, companies must assess if U.S. federal securities law is applicable and if all engagements are within established compliance [12]. It is important to note that the framework itself does not impose new obligations; rather, it clarifies the SEC's interpretive position.

The Framework further grounds its analysis of whether a digital asset constitutes a "security" in the concept of an investment contract. As defined by the U.S. Supreme Court in SEC v. WJ. Howey Co. [13], an investment contract arises when "there is the investment of money in a common enterprise with a reasonable expectation of profits to be derived from the efforts of others." [12].

The following analysis does not suggest that algorithmic systems are legal actors under securities law. Rather, it employs Howey as a functional analytic lens to examine whether AI-mediated exchange conduct performs the same market-structuring role that the doctrine historically has associated with managerial effort.

The Howey Test is intentionally functional in nature. It comprises four elements. A transaction constitutes an investment contract - and therefore a security - if:

- It involves investment of money.
- This investment is in a common enterprise.
- There is a reasonable expectation of profits.
- These profits are derived from the efforts of others [13].

Courts apply this test by examining the economic substance of the arrangement rather than its formal structure, assessing whether investors rely on external managerial or operational efforts for the realization of profit. The Howey Test provides a relevant legal framework through which U.S. securities law determines whether novel financial arrangements fall within the scope of regulatory oversight based on their practical functions rather than the design of the arrangement [13].

In this paper, Howey is relevant as an analytic reference point for understanding how centralized exchanges, through control over market structure and access, may perform market-shaping functions analogous to the managerial efforts that securities law treats as legally significant. This is not with the aim of presuming a definitive classification that digital assets are securities; merely, it is to shed light on how the function of digital assets is, in practice, like assets, and thus requires similar regulatory oversight.

In the context of digital assets, the decisive issue is frequently the fourth: whether profits are expected to be derived predominantly from the managerial or entrepreneurial efforts of others. When a promoter or Active Participant (as defined by the SEC) retains essential managerial control over network development, governance, liquidity provision, price-support mechanisms, or ongoing operational decisions, purchasers are understood to rely on those efforts for their financial return [12]. The inquiry is objective: courts examine the economic reality of the offering, the distribution plan, and the inducements presented to investors, not purchasers' subjective beliefs [13].

Centralized crypto exchanges satisfy this criterion to the extent that they exercise control over token admission, provide secondary-market liquidity, and (increasingly) delegate these essential mechanisms to proprietary machine-learning systems; the function they possess is similar to the definition of Active Participants on whose efforts investors depend. As the SEC has argued, this dynamic supports the conclusion that many listed digital assets constitute unregistered investment contracts and that exchanges facilitating their trading operate as unregistered national securities exchanges under Sections 5 and 6 of the Securities Exchange Act of 1934 [1], [12], [16].

B. SEC V. Bittrex: Managerial Functions and Automation of Exchange Control

The ambiguity and resulting tension between the functional reality of centralized crypto exchanges and U.S. securities law were addressed in April 2023, documented in the prolific civil charges filed against Bittrex, Inc., and its foreign affiliate, Bittrex Global GmbH.

Bittrex had gained market prominence as a leading exchange following Mt. Gox's collapse [6]. According to the Commission, however, Bittrex operated chronically as "an unregistered national securities exchange, broker-dealer, and clearing agency" [3] - in violation, according to the Commission, of Sections 5 and 6 of the 1934 Exchange Act [16]. What makes the Bittrex action legally significant is the Commission's articulation of why Bittrex's conduct satisfied the elements of an exchange, a broker, and a clearing agency simultaneously.

The complaint alleged three integrated forms of managerial control:

- The consolidation of several buyers and sellers of crypto assets using established, non-discretionary automated matching rules within a single shared order book operated jointly with Bittrex Global - performing the core exchange function of coordinating price and liquidity through established protocols.
- By settling transactions & maintaining custody of customer assets, it performed the core functions of a clearing agency.

- By effecting transactions in crypto assets offered and sold as securities for the accounts of others, it operated as a broker.

The most revealing allegation, however, concerned its discretionary curation of issuer disclosure. According to the SEC, Bittrex's then-CEO William Shihara routinely instructed token issuers to remove statements referencing "price prediction," "expectation of profit," or other investment-forward language that might attract regulatory scrutiny [3]. This conduct is legally critical because it demonstrates managerial shaping of the informational environment, a function relevant to the "effort of others" element of the Howey test. It displays that centralized exchanges perform an active role in shaping the external understanding of tokens, how they are presented, and ultimately, their value - a form of gatekeeping that is representative of their market power.

The doctrinal relationship between the Howey framework and the Bittrex enforcement action is vital. By establishing non-discretionary automated matching rules, maintaining custodial control over customer assets, effecting transactions for the accounts of others, and actively shaping issuer disclosures to minimize the appearance of investment-contract characteristics, Bittrex supplied managerial and operational efforts of the type Howey treats as legally significant [3], [13].

The SEC, therefore, treated Bittrex not as a passive facilitator of peer-to-peer trades but as the central coordinating enterprise whose continued managerial activities satisfies the fourth prong of Howey. The Bittrex complaint thus operates as a concrete factual application of the Howey doctrine, which demonstrates the SEC's position that it is the exchange's functional conduct, rather than the nominal classification of the tokens that renders the traded assets investment contracts [12], [13] and the platform itself an unregistered national securities exchange, broker, and clearing agency [3].

Within the context of this research, Shihara's manipulation of information explains the current salience of this doctrinal framework, as decisions formerly based on explicit human judgment are now increasingly driven through the control and structuring of information. For instance, functions such as monitoring transactional behavior, evaluating token risk, and governing listing criteria are increasingly mediated through automated systems that handle risk-scoring and surveillance within centralized exchanges [2],[5],[20],[21],[22]. In practical execution, these systems perform market-structuring functions analogous to the managerial efforts that Howey treats as legally consequential [13].

Thus, the relevance of the Howey doctrine extends far beyond the classification of digital assets. It exposes a structural continuity between human and algorithmic governance. The managerial efforts that once made exchanges legally significant actors under securities laws have been operationalized into automated, technically complex systems that replicate - and in many cases amplify - their gatekeeping powers. The paper argues that the shift from human discretion to algorithmic implementation does not diminish the doctrinal relevance of these functions; it strengthens it precisely due to the difficulty regulators face when interrogating and evaluating these systems. This continuity forms the legal and conceptual foundation for the technical architecture implemented in crypto exchange platforms today and the governance concerns explored in the following sections.

VI. AI-DRIVEN INFRASTRUCTURE & CONCENTRATED VENDOR AUTHORITY - THE RESULTING LEGAL VACUUM

A. Technical Design of Infrastructure & Inherent Risks

The relationship between automation and crypto exchanges is a very real, observed concept today. It's largely because of the efficiency of these systems that they're so prevalent within most platforms. Raun et. Al demonstrates that when trained on sufficiently large datasets, AI-driven systems such as Maximal Extractable Value(MEV) bots can outperform human strategies in arbitrage MEV auctions [9].

While their respective realms are different (MEV Flash Bot auctions vs CEX managerial accountability), the research shows that Machine Learning (ML) systems are already able to be trained to the point that they possess superior predictive skills over humans in high-stakes financial environments. In securities-law analysis, the economic analogy, not the technical equivalence, is what is relevant when attempting to gauge the scope - both potential and current - of the regulatory concern raised by AI-driven infrastructure.

A relevant example of the capability of AI surveillance in analyzing behavioral patterns and transaction flows to predict wallet ownership is Arkham Intelligence [23]. Although Arkham isn't technically in the exchange environment, its rather proficient capacity to algorithmically infer wallet identities from transactional behavior provides an example of the type of surveillance and attribution tools that centralized exchanges increasingly integrate into their Know Your Transaction (KYT), Anti-Money Laundering (AML), and risk-scoring workflows [23]. The relevance is regulatory. When these systems become a core part of an exchange's infrastructure, they possess authority over market-shaping decisions by essentially handling managerial functions. Arkham's methodologies exemplify how AI-driven behavioral inference can migrate into exchange governance, expanding the scope of information shaping functions that the SEC would ordinarily define as legally significant conduct for supervision, internal controls, and scienter evaluation.

- Performance Advantages of AI-Driven Financial Systems-

Another relevant example is found within Harlev et al., which trained machine-learning models on approximately 200 million transactions across 434 labeled Bitcoin entities and achieved a 77% classification accuracy for unlabeled addresses. [10] Although the authors identify the model's performance issues on smaller or underrepresented entity types as a dataset limitation, this class imbalance bias is an inherent risk present within ML systems. Models trained on imbalanced data inevitably perform better with the majority that's represented within the given dataset [4].

In the context of this paper and of crypto-exchange platforms, when deployed for surveillance, risk scoring, or listing-eligibility tools, these technical difficulties may impact governance. Smaller, newer, or less represented

tokens and entities become disproportionately vulnerable to misclassification, inflated risk scores, or false-positive compliance flags. In this respect, the study's reported weakness is an indicator of how, when market access decisions are mediated through opaque, data-driven classification systems, underrepresented groups face a disadvantage that may affect their overall participation in the market.

- Opacity within Core Governance Functions & the Relationship with Accountability-

One of the key issues that makes AI-driven infrastructure quite difficult to regulate via modern resources lies in the infamous "black box" problem. To facilitate effective oversight would require the ability to understand and evaluate decision-making, yet modern machine-learning systems operate through opaque internal logic that can make interrogation a challenging matter [11]. Rather than following governable, pre-defined rules, these systems rely on adaptive processes based on large datasets [4], even though their programmers may not be capable of fully interpreting their post deployment logic. When deployed in market-facing roles, such opacity enables algorithms to:

- Exercise governance authority over market access by determining which assets meet listing criteria.
- Effectively hinder, frustrate, or complicate regulatory accountability within established regulatory regimes.
- Operate beyond the transparency and explainability conventions assumed in existing regulatory frameworks.

Infrastructural opacity at this level raises fundamental questions about who can be held responsible for an algorithm's decisions, reasoning, and behaviors, and on what the evidentiary basis for the liability is. Further, the opacity of machine-learning-driven decision systems raises a distinct challenge when viewed through the lens of Howey's "efforts of others" requirement [13]. In traditional securities analysis, managerial roles, like promoters or issuers, are easy to identify. When their essential judgment tasks are delegated to opaque algorithmic systems, it becomes hard to impose or even identify liability. The managerial effort still exists, but there is no clear source.

The economic function remains the same: investors rely on an external decision-making process that shapes the asset's economic reality (value & risk profile). What does change is its traceability. In this sense, algorithmic governance complicates established frameworks such as the Howey test because the legally relevant managerial functions are performed by systems whose reasoning cannot be attributed, audited, or meaningfully interrogated.

The empirical evidence above demonstrates that AI has become embedded in the same evaluative and market-shaping functions that historically defined human managerial oversight on centralized exchanges. The result is not the disappearance of managerial control, but a reallocation of it into technical systems that now act as gatekeepers for access, risk, and information flow. The evolution of CEX infrastructure from the crypto ecosystem's original vision of peer-to-peer exchange without intermediaries, to the rise of centralized exchanges and further the emergence of AI systems that now function as invisible intermediaries embedded within the back end, highlights an increasing regulatory challenge. Sections IV and V explore this regulatory concern in detail.

B. AI-driven Risk Scoring: Potential for Market Exclusion due to Technical Limitations

Within an AI-governed CEX infrastructure, ML pipelines rely heavily on historical behavioral data to identify risk classification, perform surveillance, and determine market access. This architecture makes exchange platforms structurally prone to data sparsity bias. When a token lacks sufficient historical features, such as volatility patterns, a given token's liquidity depth, wash-trading signals, abnormal price movement profiles, and order-book stability metrics, the model cannot form stable representations and therefore defaults to conservative, risk-inflated predictions. As demonstrated by Harlev et al, models struggle with minority classes, especially entities with limited transaction histories, which are disproportionately misclassified [10][4].

New tokens without sufficient historical data face a structural issue in AI-driven risk-scoring systems because early classifications become part of the dataset that governs subsequent assessment [4]. When an asset is initially classified with a high-risk score due to incomplete information, that designation reduces liquidity and investor interaction, which in turn reinforces the original classification by generating further data consistent with a high-risk profile. The model thus treats the consequences of its own output as evidence that it made the correct judgement, exhibiting confirmation bias [4]. This can create a temporal feedback loop that progressively degrades the asset's market standing and increases the likelihood of delisting. In this respect, the system functions as a self-validating gatekeeper.

This feedback loop explains the type of adverse decisions made by automated models that the regulators, in SR 11-7-federal rules of guidance on the management of risk models, noted must be carefully tested, controlled, and monitored [24]. SR 11-7 warns that financial institutions must "be attentive to the possible adverse consequences" of models that are "incorrect or misused," and must mitigate those risks through structured model development, independent validation, and governance frameworks.

Although the guidance formally applies to banking organizations, a parallel can be drawn with the underlying risk it addresses - financial harm produced by opaque and unvalidated models [24], such as delisting of assets, is identical to the risks created by proprietary risk-scoring engines that now determine market access on centralized crypto exchanges. As SEC v. Bittrex makes clear, the Commission's stance is that the economic function of these exchanges is the deciding factor on whether they are to be treated as unregistered securities exchanges [3]. Yet unlike regulated financial institutions, they operate outside any supervisory regime that could demand model validation or remediation of harmful feedback loops. The result is a model-driven exclusion that governs market access but without the procedural safeguards that apply to regulated entities.

Exchanges and blockchain analytics firms deploy internal monitoring or adjustment mechanisms [11],[6],[7]. These internal mechanisms rely on the organization of "large amounts of transaction data and use proprietary algorithms

to identify transaction patterns. However, centralized platforms often maintain internal ledgers that are not externally visible. Scholars have proposed institutional measures such as forensic nodes to enable supervisory access or the handling of internal ledgers by regulators. [25]

These mechanisms are voluntary and lack the mandatory governance, audit trails, and supervisory review that regulate model use in traditional financial institutions. Thus, when risk-scoring engines penalize tokens due to sparse data, structural bias, model inaccuracy, or misclassifications, there is no regulatory framework through which accountability can be imposed or errors can be remediated.

C. Misclassification risks, Model biases, and the Resulting Need for Human Governance

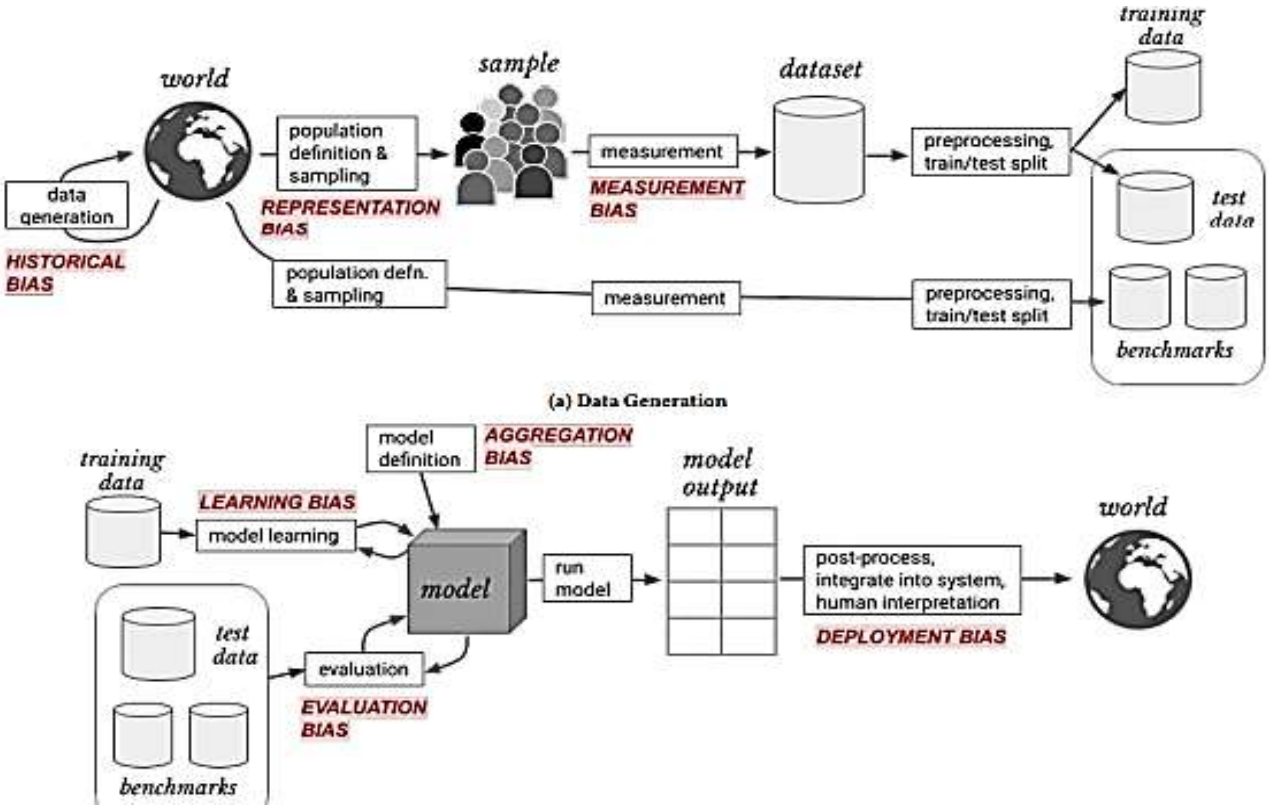

Figure 1: Sources of bias across the algorithmic decision pipeline {reproduced from H. S. Suresh and J. V. Guttag, "A framework for understanding sources of harm throughout the machine learning life cycle" [4]}

As Figure 1 depicts, machine-learning systems are vulnerable to multiple biases throughout each development stage. Each step of a model's lifetime is fraught with technical danger. From data gathering and portrayal to training, collection, assessment, and deployment. These include biases in the data's history, how it is measured, how it is represented, how it is aggregated, how it is evaluated, and how the model learns.

These biases tend to surface during rather critical stages. The potential for shaping market participation arises when these models are given the authority to exercise governance functions based on their conclusions. If, hypothetically, the model is trained on data that amplifies minority-class distortion among existing or potential assets, or when automated systems produce over-compliance because the system has been trained not to raise regulatory concern or scrutiny, each stage of development inherently starts to carry legal significance [4].

These classifications effectively function as regulatory judgments. They shape who can join a platform, who is removed, on what basis they're removed, which transactions are flagged, and which assets qualify for trading. Systems that carry this manner of authoritative weight should be subjected to appropriate oversight and governance, but due to their opaque nature, they're largely protected from external review.

By structuring factors that govern the eligibility of an asset on a platform, or its classification as high-risk or low-risk, exchanges perform market-shaping roles that securities law has historically treated as legally consequential, supporting the SEC's position that many listed digital assets may, in fact, constitute investment contracts [16].

The deeper issue is in the central nature of blockchain investigation tools. A handful of vendors are responsible for supplying tools to the entire industry - meaning these issues are a global industry-wide infrastructural risk. A lack of public data on actual exchange-level outcomes amplifies these concerns. The section below explores this in further detail [26].

Finally, while compliance tools supplied by vendors like Chainalysis, Scorechain, and Elliptic provide valuable assistance in meeting regulatory expectations [27], [28], and are no doubt necessary due to the sheer complexity of centralized exchanges, they also raise governance concerns. False positives, excess alerts, and errors in flagging risk all reinforce the need for human judgment and governance structures capable of reviewing and

correcting automated determinations, rather than deferring to them as authoritative.

D. Market-Wide Vendor Dependency for Key Infrastructure

Vendors like Chainalysis dominate the compliance and surveillance infrastructure of the cryptocurrency economy and act as industry-leading platforms for "blockchain intelligence, compliance, and risk management." [29] Additionally, while it may have begun as a blockchain analytics tool, it is now more accurate to describe it as a dominant U.S. market force within the world of financial crime prevention across digital assets.

Chainalysis has become a one-stop compliance and surveillance infrastructure provider for U.S. centralized exchanges. Its services and tools cover a wide area of crypto risk management [27].

Similarly, several other suppliers also offer key systems and functions that are either a main part of the exchange infrastructure or supplement it [29]. For example, Reactor maps fund flows across 25+ blockchains, 17 million assets, and more than 100 bridge protocols [30].

Major U.S. exchanges such as Coinbase and Binance rely on these systems simultaneously for wallet attribution, ongoing monitoring, and automated compliance adherence, effectively outsourcing core regulatory functions to a single private vendor.

Binance has remained a long-term user of the systematics and analytics tools provided by vendors like Chainalysis, Elliptic, and TRM. The company has incorporated architecture from the named vendors into its surveillance and compliance infrastructure [27].

Similarly, TRM Labs is also a major player within the crypto ecosystem [29], particularly on issues that extend beyond a single jurisdiction. Its tools are commonly used to track regulatory developments across countries, assess sanctions exposure, and analyze cross-border financial crime risks. This capacity to situate on-chain activity within a broader geopolitical and regulatory context has made TRM a significant reference point for exchanges operating in multiple legal environments. In its recent review of crypto policy developments, it is revealed that TRM Labs processed and reviewed 30 jurisdictions representing 70% of crypto global exposure in 2025 alone, analyzing upcoming crypto enforcement and regulatory actions [31].

In a 2025 report, Binance cited Chainalysis and TRM Labs as relying on similar components, such as "clustering wallet addresses, classifying transactions, and detecting suspicious behavior." [31].

Since 2019, Binance has employed Elliptic, but it is not the only one. Elliptic now claims to, following the addition of eighty-nine new crypto assets to its analytics platform, monitoring 97% of all crypto assets by global trading volume, giving it the broadest screening footprint in the industry [32]. Elliptic illustrates a parallel form of market dominance in which market access is shaped not by statutory rules but by the coverage and classificatory power of a single private vendor.

Centralized exchanges now depend on a small group of analytics vendors - Chainalysis, TRM Labs, Elliptic - to supply the surveillance, attribution, and risk-scoring systems through which regulatory obligations are satisfied [26], [29], [30], [31],[32]. However, there are currently no clear U.S. rules or legal precedents on how to regulate the conduct of vendors in the crypto ecosystem, or how much governance responsibility exchanges like Binance and Coinbase can delegate to them. This absence of direct regulation does not imply that such reliance has no place. Rather, it exposes a structural misalignment between the industry's compliance architecture and the regulatory & supervisory principles that govern all other parts of the U.S. financial system. The Office of the Comptroller of the Currency's (OCC) Interagency Guidance on Third-Party Risk Management clearly establishes that institutions remain fully responsible for harms or supervisory failures arising from outsourced tools, and the Federal Reserve's Model Risk Framework warns against unvalidated or misused external models [33].

- Regulatory Expectations from Existing Legal Frameworks & Precedents

Even if exchanges are not banks, regulators increasingly treat them as entities with similar compliance responsibilities. Additionally, the SEC's perspective on whether digital assets constitute securities and whether CEXs can be defined as National Securities Exchanges has been extensively outlined in Section IV.

The central perspective several regulators have taken is that while centralized exchanges are not banks or any traditional financial institution, they do require similar compliance responsibilities [3],[20],[21],[24]. The SEC's outlook - as discussed in Section IV - is further demonstrated by its enforcement actions against major centralized exchanges [20], [21], the Bank Secrecy Act's guidelines for Virtual Asset Service Providers [34] (which from here on out will be referred to as VASPs), and the Department of Justice's findings in United States v. Binance [35].

These principles reveal a regulatory vacuum: exchanges have built a shadow governance layer that relies on opaque third-party systems in ways that would never be acceptable in traditional financial supervision, yet no agency oversees the vendors themselves. It is this unregulated concentration of authority - not any specific statutory violation - that makes major U.S. CEX's and their reliance on third-party automation infrastructure a governance risk for the crypto ecosystem.

While sections IV and V demonstrate the plausibility of applying existing securities law concepts to AI-driven exchange conduct, the following sections examine the institutional and jurisdictional constraints that have prevented those concepts from evolving into enforceable regulatory frameworks.

VII. EXISTING REGULATORY FRAMEWORKS & THE GREY ZONE IN WHICH CENTRALIZED EXCHANGES OPERATE

A. Modern-Day Regulatory Constraints

Currently, several regulatory guidelines and frameworks govern algorithmic fairness in systems that implement automated processes. However, at present, a significant number of said frameworks and guidelines are proposed legislation, not yet enforced within federal and state laws. The Algorithmic Accountability Act (AAA), a proposed guideline, was never enacted by Congress. However, the

underlying logic validates the regulatory vacuum this paper highlights. Guidelines such as the AAA, as proposed legislation, reflect the United States fragmented, agency-by-agency approach to AI regulation and the absence of institutional consensus about how such oversight would operate.

The Algorithmic Accountability Act encompasses both entities that "use automated decision processes to make critical decisions themselves" and "entities that build the technology for use by others to make critical decisions. Most major CEXs fall into this definition and within the scope of the AAA [36].

The issue does not lie in identifying what responsible governance for automated infrastructure would look like. The 2022 AI Bill of Rights outlines principles for algorithmic discrimination, and it exists as one of many proposed frameworks highlighting the need for transparent practices [37]. However, without proper enforcement, they are just theoretical frameworks that lack the appropriate regulatory power.

The current conundrum is derived from the lack of a comprehensive federal AI regulatory framework and the absence of a clear, central enforcement agency for violations of established regulatory guidelines. CEXs do not fall into any one organization's clear jurisdiction. Simultaneously, the rapid growth of the industry cannot be understated. Centralized exchanges and crypto-asset markets have become integral to the present financial and economic landscape of the U.S. While they cannot remain outside a unified regulatory perimeter, a supervisory regime must be constructed in a manner that balances technological innovation and ensures investor-protection and market-integrity principles that anchor U.S. financial regulation.

Although the current U.S. landscape prioritizes a commitment to balancing industry growth with investor protection, existing approaches to digital-asset enforcement have not fully addressed the governance and accountability gaps surrounding centralized exchanges.

B. The Importance of Centralized Frameworks Due to Regulatory Fragmentation

Centralized exchanges are not entirely unregulated entities either. Most major U.S.-facing CEXs are registered federally as Money Service Businesses with the Financial Crimes Enforcement Network (FinCen) and hold state-level money-transmitter licenses, placing them within the scope of the federal anti-money laundering and counter-terrorist financing obligations under the Bank Secrecy Act [36].

The BSA requires CEXs to implement AML programs, customer identification procedures, transaction monitoring, and suspicious-activity reporting, and it imposes recordkeeping duties designed to detect illicit finance. While these obligations are operationally significant, they regulate financial crime controls, not the underlying exchange functions or the automated infrastructures that execute them. The BSA does not mandate model validation, algorithmic transparency, vendor oversight, or systems-governance requirements comparable to those applied to banks' risk-management models. As a result, even though CEXs are formally compliant within the BSA's framework, the statute does not reach the opaque, vendor-supplied back-end architecture that structures market access and risk classification, rendering that infrastructure outside of federal oversight [34].

The current U.S. administration has emphasized easing the regulatory pressure on the digital-asset sector with the aim of shifting away from aggressive enforcement and moving towards a growth-oriented approach that reflects the current integral importance of CEXs and cryptocurrency. This is best exemplified through the issuance of Executive Order 14178, a presidential order that aims to "support the responsible growth and use of digital assets," reflecting the growing importance of the crypto market [38].

It is important to stress that both the SEC's proactive enforcement approach and the growth-focused approach emphasized in Executive Order 14178 are critical foundations upon which future governance solutions must be developed. It is precisely due to the growth of the cryptocurrency market that regulations must prioritize industry growth and innovation; however, the opposite implication is also relevant: as the market increases rapidly and exponentially, the need for comprehensive, investor-protection-oriented regulatory frameworks rises, especially as technically unaware users continue to enter the market.

The DOJ memorandum has positive implications for industry growth and innovation. It reduces the risk of sudden or expansive criminal enforcement due to evolving legal interpretations. Historically, this has created deterrent effects that can slow product development and venture investment. The memo declines to use the Securities Act, Exchange Act, or Commodity Exchange Act as a framework for criminal prosecution in ambiguous digital assets cases [38]. It gives firms greater confidence in innovating while channeling enforcement toward more traditional fraud-based theories, allowing legitimate businesses to focus on technological development rather than navigating uncertain jurisdictional boundaries.

However, a subsequent concern raised by Executive Order 14178 is that this policy has the potential to strengthen the regulatory gap in which CEXs operate - the same regulatory gap addressed by this paper. It expressly instructs prosecutors not to bring the very registration-based offenses that anchor the SEC's theory of liability in its cases against Bittrex, Binance, and Coinbase [38]. The memo directs that prosecutors "should not charge regulatory violations in cases involving digital assets, including unlicensed money transmission, unregistered securities offering violations, unregistered broker-dealer violations, and other violations of registration requirements unless the DOJ can prove, with evidence, that such violations were willful" [38].

The DOJ's adoption of a willfulness requirement - despite such proof not being required by statute - effectively raises the threshold in terms of evidence for criminal enforcement to a level that is difficult to meet in AI-mediated systems. In the absence of DOJ-backed criminal enforcement, the SEC's registration theories function as interpretive frameworks rather than binding constraints, shaping doctrine but lacking the deterrent force necessary to discipline market behavior.

Without defined litigation, there are limited mechanisms to challenge or scrutinize the automated decision systems that modern centralized exchanges depend on. These systems are typically supplied by third-party vendors and include wallet-clustering algorithms, risk-scoring tools,

KYT filters, sanctions-screening, and behavioral surveillance models.

Because these systems are deep within the actual infrastructure of major exchanges, their operation is largely protected or hidden from external review. Of course, companies do their own compliance checks and due processes; however, this doesn't eliminate the need for proper, neutral external assessment. As a result, they function in an even stronger legal vacuum than before. Decisions that would usually be made by human management are effectively determined by these third-party automated systems and major exchanges.

This dynamic becomes incredibly significant when the fact that major platforms have historically tried to operate outside existing regulatory frameworks. (FN - Cite Bittrex/Coinbase/Binance). Those frameworks impose fairness, disclosure, and suitability obligations on traditional financial institutions. These obligations, however, aren't consistently applied to automated systems within crypto markets, even if they're applied to the exchanges themselves. However, to separate the result (market shaping decisions) from the process that enables it is redundant.

Historically, market access has depended on institutional decisions about legality, risk, and eligibility. Those decisions have been shaped by statutes, administrative rules, and enforcement actions. The DOJ memo constrains the use of registration-based theories in criminal enforcement and removes a key source of prosecutorial leverage, without displacing the underlying civil doctrines or statutory framework. The current divisive legal environment, along with the lack of a central body responsible for enforcement, regulation, and remediation, strengthens the paper's argument, which rests on CEXs operating within a regulatory vacuum. Coupled with their opaque, rapidly evolving automated systems, this only further complicates the creation of potential regulation frameworks or policies for these platforms.

- Further Analysis of Previous SEC Enforcement Actions

SEC v. Binance and SEC v. Coinbase were the SEC's most aggressive attempts to classify a centralized exchange as an unregistered securities enterprise. The Commission claimed that both exchanges operate as broker-dealers, exchanges, and clearing agencies, specifically citing opaque internal controls as one of several primary issues [21], [22]. In line with this research's analysis of CEXs as functionally analogous to traditional financial institutions, the Securities and Exchange Commission's view of Binance as a centrally coordinated enterprise highlights CEX's role as an algorithmic gatekeeper over transactional processes and market access [21].

In the matter of Coinbase, the SEC identified 13 crypto assets as meeting the definition of investment contracts, while also acknowledging that Coinbase acted as a centralized intermediary facilitating trading in assets considered securities [22]. This is, perhaps, the clearest recognition of the authoritative power major CEXs hold regarding market access that can be inferred from existing enforcement examples.

The grey zone arises from the fact that while CEXs have been viewed as centralized intermediaries, the AI-driven systems that govern their infrastructure are not ever analyzed or regulated. The lack of discussion relevant to third-party vendor-based architecture and automation is a regulatory blind spot. It is not merely enough to identify CEXs as carrying out the functions of a traditional exchange if the opaque, automated infrastructure that facilitates those functions is not also brought under regulatory examination.

Although the SEC's theory of CEX liability survives in principle, the absence of sustained enforcement limits its practical reach. When enforcement pathways narrow, courts have fewer opportunities to develop the regulatory authority and framework needed to assess questions such as market exclusion or third-party risk. A legal theory unsupported by enforcement cannot mature into regulatory precedent. As a result, the circumstances under which a court might scrutinize the internal functioning of exchange infrastructure become substantially reduced, leaving key aspects of their operations unlikely to receive judicial examination.

The question then arises, who is to be held accountable when an automated system unfairly enforces market exclusion? When the system's logic cannot be reconstructed, when even small "coding errors" can result in years' worth of transactions going unmonitored [39], and when existing regulations for third-party risk management cannot apply to CEXs due to their undefined legal identity? These questions form the basis of the governance solutions proposed in Section V.

C. Comparison with Global Responses

In contrast to the splintered U.S. approach towards crypto regulation, the E.U. has so far provided the most comprehensive attempt to provide infrastructural-level regulation regarding third-party systems. Historically, the U.S. has had a regulation-by-enforcement approach, while the E.U. 's Digital Operational Resilience Act creates a preventative regulatory perimeter that governs both crypto intermediaries and automated systems through which they operate.

When examining AI infrastructure provided by third parties and the subsequent risks, the DORA is a relevant framework upon which future governance solutions may potentially be based. Firstly, DORA applies to centralized crypto exchanges given that they are registered and authorized as Crypto-Asset Service Providers under MiCA [40]. This makes their internal third-party infrastructure, which the E.U refers to as Information & Communication Technology systems (ICTs), eligible for the compliance and governance requirements outlined in the Digital Operational Resilience Act [5].

DORA directly addresses the infrastructural risks identified in this paper by subjecting third-party digital systems to ex ante governance obligations, an approach largely absent from the U.S. regulatory framework. Rather than constraining innovation or industry growth, DORA establishes structural safeguards - focused on dependency, accountability, and resilience - to ensure that the outsourced technological infrastructure operates in a manner consistent with fairness and regulatory compliance.

Secondly, it is important to understand the principle of "digital operational resilience." The EU understands digital operational resilience as "being vital for ensuring financial stability and market integrity in the digital age".

Fair market access, non-discriminatory system infrastructure (whether third-party or otherwise), and investor protection can all be defined as the preservation of market integrity. To satisfy this standard, the E.U., under DORA, stresses the importance of internal governance and control frameworks for "effective management of ICT risk" [5]. It delegates this responsibility to the management body of the (given) financial entity.

Internal governance, per DORA, may appear as internal reporting channels at a corporate level, which would enable management to be informed of the impact of implementing third-party ICT services in critical functions, with a relevant risk analysis summary to properly evaluate aid impact, complete with "response, recovery, and corrective measures" as stated in Article V. [5] The effect of such solutions is to create traceable accountability and named responsibility, something currently lacking within third-party supplied infrastructure present in major U.S. CEXs.

The EU's DORA shows that it is possible to regulate & govern complex, automated financial infrastructure directly, without categorizing financial entities such as CEXs into traditional banking models or hampering industry growth and innovation through reactive, aggressive enforcement measures. Rather, on page 30, DORA requires financial entities to maintain a "sound, comprehensive, and well-documented ICT risk management framework" [5] encompassing not only internal systems but the full range of information assets, protocols, and tools upon which market operations depend. Crucially, this framework is subject to regular, yearly review, independent internal audits, supervisory access, and resilience testing. In contrast to the U.S. approach, where automated surveillance, risk-scoring, and the resulting impact on market access operate largely outside regulatory oversight, the DORA treats infrastructure itself as a site of governance rather than a neutral technical factor.

Most relevant to this paper's focus on vendor dominance, DORA explicitly anticipates and regulates multi-vendor dependency. As per Article 28, it permits the outsourcing of critical ICT functions while making clear that regulated entities "remain fully responsible" [5] for compliance, even where verification tasks are delegated to external providers.

This allocation of responsibility directly addresses the accountability gap identified in the U.S. crypto markets, where exchanges rely on third-party analytics, risk scoring, and surveillance tools yet face no formal obligation to document, audit, or explain how these systems shape surveillance and market access. By requiring documentation of vendor dependencies, independent control functions, auditability, and supervisory disclosure, DORA demonstrates a preventative regulatory model capable of disciplining automated back-end infrastructure before harm occurs. These are precisely the operational safeguards whose absence in the U.S. crypto context explains how automated systems can reshape market access without contemporaneous accountability or supervisory visibility.

In practice, DORA operationalizes these principles through internal governance of AI-driven exchange infrastructure that requires changes to vendor-supplied systems - such as updates to the infrastructure, adjustments to risk-scoring thresholds, or modifications to transaction-surveillance logic - trigger formal internal processes. Where such changes result in impacts on market access, such as account freezes, transaction blocks, delisting, or the exclusion of users from trading, senior management must be informed [5]. Further, the downstream impact on market access must be assessed, and the incident must be documented and remediated. Similarly, where multiple compliance and surveillance functions depend on a single vendor's infrastructure, concentration risk must be identified, evaluated, and recorded, particularly where a vendor's error or model failure could systemically restrict participation across the platform [5]. The DORA effectively acts as an ex-ante framework in stark contrast to the U.S.'s reactive approach, focusing on preventing infrastructural risk and harm rather than regulating and enforcing remediation after said harm has occurred.

VIII. PROPOSED GOVERNANCE FRAMEWORK

Rather than proposing novel governance structures, this paper draws on existing frameworks from across the globe, such as DORA, to extend governance practices & strategies onto the issue of algorithmic gatekeeping within crypto exchange platforms.

Namely, the DORA's proposals outlined in Articles 5 and 6 of having a management body, ICT-relevant staff training, the establishment of a separate control function for overseeing ICT risk, annual risk management framework evaluation alongside incident responses, all previously mentioned governance tactics being subject to internal audits, and the documentation of all ICT-supported business functions [5].

A. Constructing Theoretical Governance: Strategic Takeaways from DORA

What DORA suggests is rigorous implementation of human oversight to properly implement risk management. While it chiefly advocatesactive human oversight, it does not imply that the systems it deems ICTs should not be used [5]. DORA grounds human oversight in formal governance obligations. As stated in Article 5, by requiring financial entities to assign responsibility for ICT risk to an independent control function, segregated from operational and audit roles [5], DORA establishes that automated systems must remain subject to identifiable human authority. This is exactly the kind of named accountability that is currently missing from U.S. regulatory frameworks. Human oversight should be perceived as an institutional and governance requirement. Management bodies must bear the ultimate responsibility for risks created by digital infrastructure, including outsourced and third-party systems.

This model directly addresses the shadow governance layer. Due to the opacity and complexity of modern AI infrastructure within CEXs, it becomes increasingly difficult to delegate accountability to one singular party. By assigning formal oversight and ultimate responsibility for ICT risk to the management body, this framework prevents AI-driven systems from operating as de facto regulators within the exchange. Human oversight

functions serve to diffuse concentrated algorithmic authority by requiring the following actions:

- At the organizational level, the exchange would be required to maintain a documented risk management framework that explicitly identifies which automated systems influence critical functions, such as surveillance, risk scoring, listing, and delisting. This documentation would describe and document the areas of business impacted or managed by the vendor-supplied infrastructure.
- Internal recordkeeping. Internal records must require that exchanges reconstruct what system acted, under what conditions, and with what risk assumptions. These records are reviewed periodically through internal audits conducted by personnel with expertise in the systems supplied by third parties, and the findings must feed into corrective action plans. Regulators are also entitled to these internal records. This is to ensure that automated systems cannot possess or exercise concentrated authority without leaving an institutional record.

This model provides a clear template for governing AI-driven infrastructure that now performs gatekeeping functions traditionally associated with regulated intermediaries.

B. Designated Governance Roles within Centralized Exchanges

Under the DORA-inspired oversight framework, delisting decisions triggered by automated risk-scoring engines would be subject to escalation and review by an independent control function, rather than being executed solely based on vendor-supplied outputs. Where a delisting is later found to have resulted from improper, incomplete, or flawed data, biased models, or unvalidated threshold changes, the framework would require documented remediation, management accountability, and internal review of the system logic that produced the decision.

This control unit would consist of:

- A Designated Algorithmic Risk Officer - The risk officer holds ultimate responsibility for risks arising from automated systems that affect or reshape market access, including risk scoring, surveillance, listing, delisting, and account restrictions. This role is not operational and does not review individual decisions or system updates. Instead, it sets risk tolerance for algorithmic gatekeeping, determines which categories of decisions may be automated, and formally approves conditions under which automated systems are permitted to exercise authority, acting as the final layer of management that automated infrastructure must pass through before deployment. This role is built to ensure that responsibility for algorithmic outcomes cannot be displaced onto technical teams or third-party vendors.
- Systems Implementation Unit - This unit is responsible for defining and enforcing the scope of authority granted to AI-driven systems prior to deployment. Its function is to determine what actions a system may take autonomously, what actions require human confirmation, and what thresholds or conditions trigger escalation. This unit reviews system integrations, vendor updates, and architectural changes solely to ensure that automated systems operate within pre-approved boundaries and conditions. It does not evaluate outcomes or investigate harm; its value lies in preventing excessive or unreviewable authority from being embedded into system design.
- Change & Incident Review Channel - To facilitate the previously proposed internal recordkeeping measure, this channel would document any harmful outcome - false positives, unjustified delisting, systemic exclusions - must be logged and routed here. This channel's responsibility is not to redesign systems, reassess risk tolerance, or make final authoritative decisions, but to create a traceable internal record that triggers remediation, management awareness, and potential supervisory disclosure. This prevents algorithmic governance from operating in opaque, ungoverned manners - without institutional evidence or accountability.
- Internal Audit Team - This is a separate control function independent of the above three actions. Whereas the above three would work in tandem to provide the necessary data, records, and evaluations in order to satisfy a practical risk management framework, an internal audit team would be tasked with ensuring governance rules were followed, compliance was adhered to, and consumer-protection was prioritized by assessing the preventative, corrective, and regulatory measures within the AI infrastructure and centralized exchange. Internal audits ensure that algorithmic governance mechanisms function as intended. Ideally, a bi-annual report and review process to assess whether automated systems continue to operate within approved authority boundaries, enabling and implementing ex-ante measures before algorithmic harm can affect customer experience or market integrity.

A natural concern is that the logic of complex AI systems cannot be fully reconstructed, raising the familiar "black box" problem [11]. The framework proposed in this paper - drawing directly from DORA - does not seek to eliminate this opacity or mandate technical explainability. Instead, it operates governance around it. By requiring procedural documentation, decision authority, data categories, thresholds, updates, and incident outcomes, the framework generates a reconstructed institutional record of how automated decisions affecting market access are produced, without demanding insight into proprietary or highly complex model internals. This enables oversight based on procedural accountability rather than algorithmic transparency.

- Governance Committee Oversight and Periodic Review-

In addition to the internal oversight and audit functions outlined above, the proposed framework would be strengthened by the establishment of a standing governance committee responsible for periodic, high-level review of automated market-access infrastructure. This committee would operate independently from day-to-day compliance, engineering, and audit functions, and would report directly to senior management or the board. Its purpose would not be to revisit individual enforcement actions or technical model performance, but to evaluate whether the exchange's governance architecture remains

adequate considering how automated systems operate in practice.

First, the governance committee would be responsible for implementing and conducting periodic governance adequacy reviews. This would be on a bi-annual or annual basis. Additionally, reviews may be necessary after any major compliance violations, infrastructure failures, or crucial system updates. These reviews would assess whether existing oversight roles continue to function as intended, or whether authority has effectively become concentrated back into vendor systems or automated processes. The focus of this review would be structural rather than technical: whether escalation thresholds, approval mechanisms, and accountability assignments remain sufficient given observed system behavior and organizational incentives.

Second, the committee would perform a market-access impact review based on aggregated outcomes rather than individual decisions. This review would examine patterns in listings, delisting, account restrictions, and liquidity access over the review period to identify whether automated systems have produced unfair exclusion, over-compliance effects, or systemic bias against particular classes of assets or users. By focusing on cumulative effects rather than case-specific errors, this review would allow institutional awareness of the degree to which automated governance affects market participation.

Third, the governance committee would be responsible for a concentration and dependency assessment addressing the exchange's reliance on third-party vendors, shared data sources, and standardized risk-scoring models. This assessment would evaluate whether dependency on a limited set of vendors or tools has increased systemic vulnerability, synchronized decision-making across platforms, or reduced the exchange's human control over market-access determinations. Where such risks are identified, the committee would be tasked with recommending governance adjustments, such as diversification strategies, enhanced oversight requirements, or revised escalation protocols.

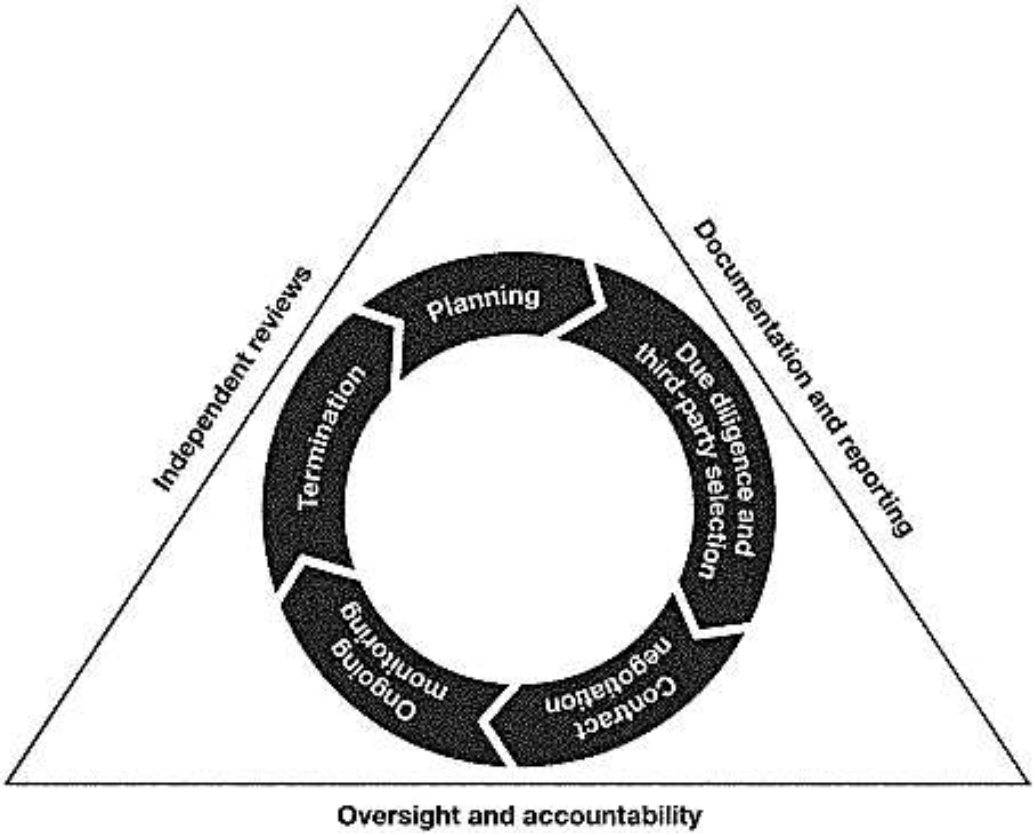

Figure 2: Stages of risk management life cycle for third-party relationships (Reproduced from the Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency, "Final Interagency Guidance on Third-Party Relationships: Risk Management," Federal Register, vol. 88, no. 111, p. 183, Jun. 9, 2023 [41]}.

Figure 2 depicts a possible governance framework. This is implemented through a continuous lifecycle that is anchored in oversight and accountability, emphasizing cyclical governance functions. These functions align with the proposed governance framework as they are similarly rooted in ongoing surveillance of automated infrastructure. Importantly, human oversight, such as the one proposed by this paper, creates foundational constraints across all stages.

The function of this committee would enable governance at the institutional level, rather than within isolated technical or compliance processes. By embedding periodic, independent review into the organizational structure of the exchange, the governance committee provides a mechanism for observing how automated authority concentrates over time and for recalibrating the governance framework accordingly. This audit supplements the role of the internal audit team and

operational controls by addressing systemic change, rather than procedural compliance alone.

- Decision Traceability and Data Provenance-

In addition to incident logging and escalation mechanisms, governance should require decision traceability: a layer of evidence that preserves the data categories, sources, system versions, and contextual parameters used when automated systems produce market-access outcomes. The goal is not to force model explainability, but to allow reconstruction of the decision inputs that produced a harmful outcome so that faulty, biased, or stale datasets can be identified and remediated.

This approach has precedents in regulatory practice. The EU's Digital Operational Resilience Act (DORA) requires financial entities to maintain governance, documentation, and oversight of ICT systems - including lifecycle management, operational records, and third-party oversight - thereby enabling retrospective review of ICT-supported decisions [5]. In the U.S., supervisory guidance on model risk management similarly emphasizes robust model development, validation, and governance to guard against adverse consequences from incorrect or misused models - a principle that supports maintaining records of model inputs and operational context [24].

National supervisors have acted on these governance gaps. The UK Financial Conduct Authority's reviews of algorithmic trading and related market-structure activity have repeatedly highlighted deficiencies in documentation, governance, and the capacity of compliance teams to oversee automated strategies - leading to formal interventions and attestation requirements in some cases [42].

Similarly, Australia's ASIC has warned that the rapid adoption of AI by licensees has created governance gaps and signaled plans for tighter supervisory focus on how firms govern AI-driven decision systems [43]. This reflects a global acknowledgement of the need for adaptable governance in automation.

Operationalizing decision traceability in exchanges would therefore require:

- Mandatory recording of the categories of data sources used for material decisions (including vendor classifications)
- Immutable versioning and timestamping of system updates and threshold changes, and
- Indexed incident links connecting harmful outcomes to the exact inputs and system state at the time of decision.

These records would enable internal audit units, the governance committee, and supervisors to distinguish harms arising from defective data or vendor inputs from those arising from governance failures or human error - while stopping short of mandating proprietary model disclosure.

C. *Market Access, Vendor Concentration, and Power Allocation*

Crucially, this documentation is paired with a human-governed control structure that reallocates authority away from the automated back end. The Algorithmic Risk Officer retains final responsibility for whether categories of automated decisions are permitted at all. The Systems Implementation Unit implements ex-ante actions that ensure harmful outcomes are logged, escalated, and remediated rather than silently absorbed. Internal audit then independently verifies that these governance obligations are observed in practice.

Taken together, the framework directly addresses the concentration of authority in third-party vendors supplying surveillance and risk-scoring infrastructure. Vendors such as Chainalysis, Elliptic, and TRM Labs increasingly supply standardized tools across multiple exchanges, creating the risk that a single model error or classification bias can propagate across the market [29], [31], [32].

Under the proposed framework, vendor dependency is wholly and operationally the responsibility of the exchanges that employ them. Exchanges remain accountable for how those systems are deployed, what authority they exercise, and how their outputs affect market access. Concentration risk - where multiple critical functions depend on a single vendor - must be identified, documented, and reviewed, particularly where failure could result in systemic exclusion or impact market integrity, access, and fairness.

By re-centering authority in human governance structures, the framework mitigates automated delisting, over-compliance distortions, and unreviewable exclusionary practices without prohibiting vendor participation or technological innovation.

D. *Implementation of Proposed Governance Structure within the Existing Legal Structures*

The governance framework proposed in this section does not exist in a vacuum. Its feasibility and limits must be understood against the institutional posture of U.S. financial regulation as it currently operates. The preceding sections have shown how AI-dominated, vendor-supplied infrastructure has come to possess the authority it does over market access within CEXs, and how the absence of binding oversight has allowed automated systems to create needless exclusion without accountability.

The paper acknowledges that the proposed governance framework is heavily inspired by the DORA legislation; however, it does not attempt to import a European framework as it is into the U.S. landscape. Rather, the existence of the DORA legislation provides critical evidence for the argument that although AI infrastructure and centralized exchanges are inherently technical and more complex innovations are advancing rapidly, this complexity does not constitute an insurmountable barrier towards a measured regulation by law enforcement agencies.

DORA does not require technical explainability of machine-learning systems. While it does clearly claim that technically proficient staff must be employed to assess these systems [5], it does not mandate access to proprietary model internals. Instead, it governs authority allocation, documentation, accountability, and resilience. In doing so, it demonstrates that opacity is not a concrete barrier to governance, but a condition around which effective governance can be constructed. While the paper does not claim that the proposed governance will be effective in regulating AI-dominated infrastructure entirely, it contributes a theoretical base rooted in accountability, transparency, preventative measures, and human-based governance subject to internal audits and

continuous refinement to adapt to the evolving nature of both AI and centralized exchanges.

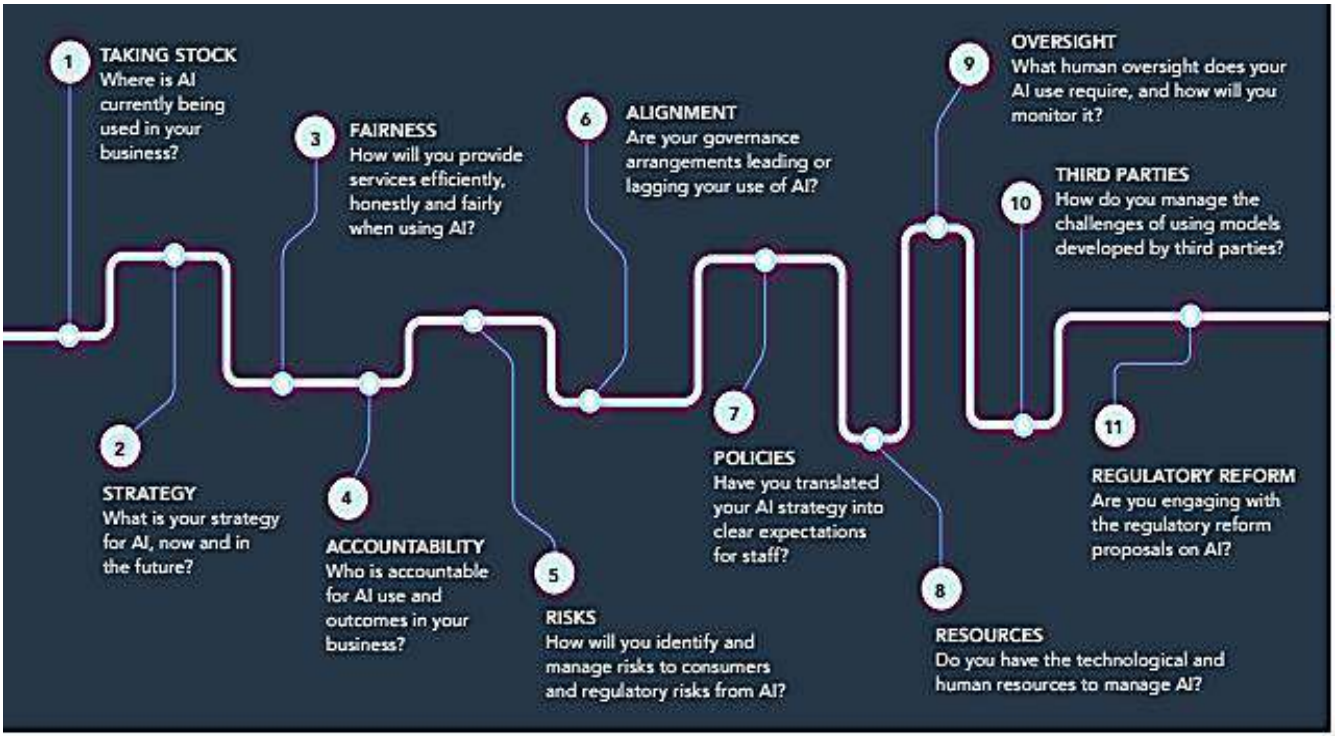

Figure 3: Governance considerations for organizational use of artificial intelligence (Reproduced from: Australian Securities and Investments Commission, *Report 798: ASIC review of automated advice providers*, Oct. 29, 2024 [43]).

The above Figure 3 depicts fundamental considerations that should be centered in an effective governance strategy or framework. It provides a comprehensive overview of what factors must be considered when implementing AI, how that translates into regulatory reform, and how to evaluate and assess third-party risk management.

- An Overview of the SEC's Enforcement Posture-

Historically, the approach the SEC has adopted has been that of reactive, case-by-case enforcement rather than ex ante rulemaking. This strategy was primarily rooted in proving specific platforms function as unregistered securities exchanges, broker-dealers, or clearing agencies [3], [20], [21], often grounded in the functional application of the Howey test [13]. Hypothetically, were CEXs to be classified as securities exchanges from the next day, and all subsequent compliance and regulatory scrutiny were to be enforced, this may still not properly address the governance of exchange infrastructure itself.

There are several reasons for this. First, the SEC's litigation does not address ongoing structural conditions. It addresses violations of the Exchange Act of 1934 through registration-based enforcement theories [16] that, under the DOJ memorandum, are no longer viable as a basis for criminal prosecution without proof of willfulness [38]. The automated infrastructure outlined in the previous sections operates continuously, shaping market access through cumulative, often invisible decisions. Case-specific enforcement cannot meaningfully interrogate these systems as systems.

Second, the SEC's enforcement posture, reliant on lawsuits, has become fragile. Due to recent setbacks and procedural dismissals, there is heightened skepticism towards accepting expansive interpretations of decades-old securities statutes in the absence of clear congressional authorization. This has heightened the legal risk associated with any effort to formalize a broad regulatory framework for crypto-market infrastructure.

In this environment, formal rulemaking poses a particular challenge for the SEC. Unlike litigation, rulemaking would require the agency to publicly articulate and defend a comprehensive account of why AI-driven exchange infrastructure - such as automated listing systems, surveillance engines, and vendor-supplied risk models - falls within its regulatory authority. Once a binding rule is created, that theory would apply across the industry and be subject to immediate judicial review. Such a rule would likely attract aggressive legal challenges. The Commission, therefore, has strong incentives to avoid committing to infrastructure-level governance through rulemaking, even if such a lack of commitment may reinforce the market-shaping authority that centralized exchanges currently have through the delegation of governance functions to automated infrastructure.

- Road to Future Governance: Extending Contemporary Regulatory Ideology to Automated Exchange Infrastructure-

The framework proposed in this paper would require formal rulemaking under the Administrative Procedure Act to carry legal consequences. Governance of automated market access cannot be achieved through non-binding guidance or informal expectations without reproducing the very vacuum this paper identifies.

The legal authority that would most plausibly support such rulemaking, once jurisdictional coverage is established, would be grounded in the Securities Exchange Act's civil recordkeeping and internal-controls provision, commonly codified at 15 U.S.C § 78q(a) of the Securities Exchange Act of 1934 [16]. This authorizes the Commission to

prescribe rules requiring covered entities to make and keep records and reports necessary or appropriate for investor protection and market integrity.

The proposed framework is not premised on a definitive resolution of whether centralized crypto exchanges constitute registered securities exchanges under the existing law. For the purpose of this analysis, § 78q(a) of the 1934 Exchange Act [16] is invoked not to assert present jurisdiction over CEXs. Rather, it demonstrates the form that binding oversight of the AI-dominated market would necessarily be implemented should such platforms be brought under a regulatory perimeter, whether through statutory clarification, judicial interpretation, or future rulemaking. The current inability to impose internal-control obligations on exchange infrastructure without first resolving jurisdictional coverage underscores the central governance failure this paper identifies.

In summation, the regulatory acts from which the governance framework draws inspiration prove that opacity and technological complexity are not sufficient barriers to regulation. Rather, they show that effective oversight can be constructed around opaque systems through institutional design, accountability, and procedural governance. The framework clarifies what binding oversight of automated market access would require once jurisdictional authority is established, without presuming that such authority is presently settled.

By specifying the institutional form that oversight of AI-dominated market access would require, this framework shifts the regulatory question from whether automated gatekeeping can be governed to what legal conditions must exist for such governance to become possible.

IX. CONCLUSION

Centralized crypto exchanges increasingly rely on automated systems to determine who exactly gets to participate in the market and on what terms. These systems, in practice, exercise the kind of governance authority that has always been within the realm of human oversight. However, they do this without the appropriate regulatory scrutiny or accountability mechanisms.

This paper argued that existing legal debates - largely focused on questions of asset classification, registration issues, and enforcement authority - do not fully encompass the governance implications of AI-dominated exchange infrastructure. By examining the functional role of automated systems through the lens of securities law doctrine, enforcement practice, global regulatory responses, and discourse, the analysis shows that centralized exchanges now perform market-structuring functions akin to those associated with financial intermediaries. The move from human decision-making to automated infrastructure simply situates these functions within technical systems that are opaque and difficult to assess via existing regulatory tools and resources.

In response, the paper proposes its core contribution in the form of a governance framework inspired by the Digital Operational Resilience Act (DORA). Rather than requiring technical transparency or the disclosure of the internal logic of these systems, it instead centers on auditability and adaptability. With clearly defined human responsibility at the core, this approach aims to demonstrate that automated infrastructure, despite its opaqueness, can be meaningfully governed without mandating explainability or redesign.

The secondary contribution of this paper is a reframing of the challenge posed by centralized exchanges. The problem is not the absence of regulatory tools or whether automated infrastructure can be governed. Rather, the problem lies in how to extend current tools and direct them towards these automated systems. By clarifying what effective governance would need, this paper shifts attention away from doctrinal uncertainty and toward the structural conditions necessary for accountability within AI-dominated back-end systems. As both the crypto market and automated systems increasingly evolve, so too does their capacity for impacting and shaping market access, participation, and the overall consumer experience. Therefore, addressing this governance gap becomes essential to maintaining market integrity and adapting consumer protection in the face of these evolving technological realities.

CONFLICTS OF INTEREST

The authors declare that they have no conflicts of interest.

REFERENCES

[1] Chainalysis, “Chainalysis,” Chainalysis, 2025. Available from: https://www.chainalysis.com/

[2] S. Hägele, “Centralized exchanges vs. decentralized exchanges in cryptocurrency markets: A systematic literature review,” *Electronic Markets*, vol. 34, art. no. 33, 2024. Available from: https://link.springer.com/article/10.1007/s12525-024-00714-2

[3] U.S. Securities and Exchange Commission, “SEC v. Bittrex, Inc. and Bittrex Global GmbH, Complaint,” Comp. No. 3:23-cv-04495, United States District Court for the Northern District of California, Aug. 28, 2023.

[4] H. S. Suresh and J. V. Guttag, “A framework for understanding sources of harm throughout the machine learning life cycle,” *arXiv preprint*, Jan. 2019. Available from: https://doi.org/10.1145/3465416.3483305

[5] European Union, “Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector,” *Official Journal of the European Union*, Dec. 2022. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5525144

[6] R. Böhme, N. Christin, B. Edelman, and T. Moore, “Bitcoin: Economics, technology, and governance,” *Journal of Economic Perspectives*, vol. 29, no. 2, pp. 213–238, 2015. Available from: https://www.aeaweb.org/articles?id=10.1257/jep.29.2.213

[7] N. Gandal, J. T. Hamrick, T. Moore, and T. Oberman, “Price manipulation in the bitcoin ecosystem,” *Journal of Monetary Economics*, vol. 95, pp. 86–96, 2018. Available from: https://doi.org/10.1016/j.jmoneco.2017.12.004

[8] A. Kirilenko and A. W. Lo, “Moore’s Law versus Murphy’s Law: Algorithmic trading and its discontents,” *Journal of Economic Perspectives*, vol. 27, no. 2, pp. 51–72, 2013. Available from: https://www.aeaweb.org/articles?id=10.1257/jep.27.2.51

[9] Raun, B. Estermann, L. Zhou, K. Qin, R. Wattenhofer, A. Gervais, and Y. Wang, “Leveraging machine learning for bidding strategies in miner extractable value (MEV) auctions,” *IACR Cryptology ePrint Archive*, Report 2023/1281, 2023. Available from: https://eprint.iacr.org/2023/1281

[10] M. Harlev, H. S. Yin, K. C. Langenheldt, R. R. Mukkamala, and R. K. Vatrapu, "Breaking bad: De-anonymizing entity types on the bitcoin blockchain using supervised machine learning," in *Proc. 51st Hawaii Int. Conf. on System Sciences (HICSS)*, Jan. 2018, pp. 1–10. Available from: https://tinyurl.com/2fhxbt6a

[11] G.-G. S. Fletcher and M. M. Le, "The future of AI accountability in the financial markets," *Journal of Entertainment and Technology Law*, vol. 24, pp. 14–36, 2021. Available from: https://tinyurl.com/4hwvehat

[12] U.S. Securities and Exchange Commission, "Framework for 'Investment Contract' Analysis of Digital Assets," Apr. 3, 2019. Available from: https://tinyurl.com/aarehn88

[13] SEC v. W. J. Howey Co., 328 U.S. 293 (1946). Available from: https://supreme.justia.com/cases/federal/us/328/293/

[14] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," White paper, 2008. Available from: https://tinyurl.com/5d58c2b8

[15] V. Buterin, "Ethereum: A next-generation smart contract and decentralized application platform," White paper, 2014. Available from: https://tinyurl.com/ye2936cd

[16] U.S. Congress, "Securities Exchange Act of 1934," 15 U.S.C. § 78a et seq., 1934. Available from: https://www.jstor.org/stable/791859

[17] P. Saggese, "Arbitrage in the bitcoin ecosystem: An investigation of the Mt. Gox exchange platform," Ph.D. dissertation, IMT School for Advanced Studies Lucca, Italy, 2021. Available from: http://e-theses.imtlucca.it/329/

[18] S. Foley, W. Krekel, V. Mollica, and J. Svec, "Not so fast: Identifying and remediating slow and imprecise cryptocurrency exchange data," *Finance Research Letters*, vol. 51, art. no. 103401, Jan. 2023. Available from: https://doi.org/10.1016/j.frl.2022.103401

[19] Coinbase, "Exchange matching engine," Coinbase Documentation, 2025. Available from: https://docs.cdp.coinbase.com/exchange/concepts/matching-engine

[20] U.S. Securities and Exchange Commission, "SEC v. Binance Holdings Ltd., BAM Trading Services Inc., and Changpeng Zhao, Complaint," Comp. No. 1:23-cv-01599, United States District Court for the District of Columbia, June 5, 2023.

[21] U.S. Securities and Exchange Commission, "SEC v. Coinbase, Inc., Complaint," Comp. No. 1:23-cv-04738, United States District Court for the Southern District of New York, June 6, 2023.

[22] W. Liang *et al.*, "AI and blockchain for AML: A policy and technology convergence to combat crypto-enabled financial crimes," pp. 3–20, June 4, 2025. Available from: https://tinyurl.com/469revpz

[23] R. Mafrur, "AI-based crypto tokens: The illusion of decentralized AI?" *IET Blockchain*, vol. 5, no. 1, 2025. Available from: https://doi.org/10.1049/blc2.70015

[24] Board of Governors of the Federal Reserve System, "Supervisory guidance on model risk management (SR 11-7)," Apr. 4, 2011. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3332484

[25] H. Yue, "Regulating crypto money laundering: An assessment of current regulatory responses and potentials for technology-based solutions," *Stanford Journal of Blockchain Law & Policy*, June 30, 2025. Available from: https://stanford-jblp.pubpub.org/pub/crypto-laundering/release/1

[26] McKinsey & Company, "The stable door opens: How tokenized cash enables next-gen payments," May 20, 2025. Available from: https://tinyurl.com/3p658ebk

[27] Chainalysis, "2025 Crypto Crime Report," 2025. Available from: https://go.chainalysis.com/2025-Crypto-Crime-Report.html

[28] Elliptic, "Typologies Report 2024: Preventing financial crime in crypto assets," 2024. Available from: https://www.elliptic.co/resources/elliptic-typologies-report-2024

[29] Wright, *Executive Brief: Chainalysis*, Forrester Research, Aug. 15, 2025. Available from: https://www.fourester.com/fourester/executive-brief-chainalysis

[30] Chainalysis, "Reactor: Crypto & blockchain investigations," 2025. Available from: https://www.chainalysis.com/product/reactor/

[31] TRM Labs, *Global Crypto Policy Review & Outlook 2025/26*, Dec. 3, 2025. Available from: https://tinyurl.com/mty7kytt

[32] Elliptic, "Elliptic expands coverage to over 97% of all cryptoassets," May 20, 2020. Available from: https://www.elliptic.co/media-center/elliptic-expands-coverage

[33] Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency, "Interagency guidance on third-party relationships: Risk management," June 7, 2023. Available from: https://www.federalreserve.gov/supervisionreg/srletters/sr2304a1.pdf

[34] Bank Secrecy Act, 31 U.S.C. §§ 5311–5336, 1970. Available from: https://tinyurl.com/9wva67a9

[35] U.S. Department of Justice, "Binance and CEO plead guilty to federal charges in $4B resolution," Nov. 21, 2023. Available from: https://tinyurl.com/478zrxd8

[36] Algorithmic Accountability Act of 2022, S. 3572, 117th Cong., 2022. Available from: https://www.congress.gov/bill/117th-congress/senate-bill/3572

[37] White House Office of Science and Technology Policy, *Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People*, Oct. 2022.

[38] J. Trump, "Executive Order 14178—Strengthening American leadership in digital financial technology," Jan. 23, 2025. Available from: https://tinyurl.com/bdfz2669

[39] S. Campbell and R. Taylor, "Crypto exchange Coinbase fined €21.5mn by Ireland's central bank," *Financial Times*, Nov. 6, 2025. Available from: https://www.ft.com/content/21f15153-5f72-4860-97be-74b6c1ad0a60

[40] Regulation (EU) 2023/1114 of the European Parliament and of the Council of 31 May 2023 on Markets in Crypto-assets (MiCA), *Official Journal of the European Union*, OJ L 150, June 9, 2023.

[41] Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, and Office of the Comptroller of the Currency, "Final interagency guidance on third-party relationships: Risk management," *Federal Register*, vol. 88, no. 111, June 9, 2023. Available from: https://tinyurl.com/2kme9wj4

[42] Financial Conduct Authority, "Algorithmic trading compliance in wholesale markets: Multi-firm review," 2024.

[43] Australian Securities and Investments Commission, *Report 798: ASIC review of automated advice providers*, Oct. 29, 2024. Available from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5040343