Data-Driven Mechanism Design for Contract-Constrained NFT Exchange Platforms

Liangbo Zheng[1], Junhao Ren[1,*]

[1] Faculty of Business, Information & Human Sciences, Kuala Lumpur University of Science & Technology, Kuala Lumpur, Malaysia 43000

* Email: Ren18503573635@gmail.com (Corresponding Author)

Received 10 September 2025 · Revised 25 October 2025 · Accepted 9 December 2025 · Published 30 March 2026

Abstract

Non-fungible token (NFT) platforms increasingly support assets whose transferability is shaped by vesting periods, staking commitments, royalty duties, access rights, and collection-specific transfer rules. These contractual features make NFT exchange more complex than simple peer-to-peer barter because an allocation that improves asset fit may still violate term consistency, expose users to hidden obligations, or reduce platform-level trust. This article develops a data-driven mechanism design framework for contract-constrained NFT exchange platforms. The framework integrates smart-contract metadata, wallet-level trading histories, inferred preference rankings, and transaction-risk indicators into a contract-aware matching process. Instead of treating NFTs only as indivisible goods, the study models each exchange option as a digital asset bundled with a transfer term. A data-driven equal-term top trading cycles mechanism is then proposed to balance allocative efficiency, individual rationality, term consistency, and manipulation resistance. Numerical experiments are conducted using a synthetic platform dataset calibrated to realistic NFT-market features, including heterogeneous user preferences, different shares of locked assets, royalty-bearing transfers, and varying degrees of preference noise. The results show that a standard top trading cycles rule generates high apparent efficiency but creates frequent term-consistency violations when restricted assets are common. A contract-filtered rule eliminates violations but loses welfare by blocking too many mutually beneficial exchanges. The proposed data-driven equal-term mechanism achieves the strongest overall platform score by combining constraint screening with preference learning and risk-sensitive tie breaking. The findings contribute to business data analytics, digital platform governance, and market design by showing how data infrastructure can translate programmable property rights into operational exchange rules.

Keywords: NFT exchange; mechanism design; smart contracts; data analytics; equal-term matching; platform governance; digital assets; blockchain markets

Data-Driven Mechanism Design for Contract-Constrained NFT Exchange Platforms

1. Introduction

NFT markets have moved beyond the early image of digital collectibles traded only for scarcity, community identity, or speculative appreciation. Many assets now carry executable contractual obligations. Some collections require a holding period before resale. Others attach staking conditions, time-limited access to a game feature, creator royalties, membership privileges, or restrictions that differ across marketplaces. These terms change the economic nature of exchange. A holder does not simply own an indivisible token; the holder owns a token under a particular bundle of rights and obligations. A platform that allows users to swap such assets must therefore evaluate not only which user prefers which token, but also whether the transferred obligation is consistent with the obligation that the user gives up. A seemingly efficient exchange may be unfair if a participant receives an unrestricted NFT while transferring a restricted one to another user. It may also be legally or reputationally harmful if the platform enables users to shed duties embedded in smart contracts. This institutional interpretation is consistent with the view that blockchains lower verification and networking costs rather than simply digitizing records (Catalini & Gans, 2020).

Classical market design offers a natural starting point for analyzing this problem. The housing-market model of indivisible goods and the top trading cycles family of mechanisms show how initially endowed agents can exchange objects without money while preserving individual rationality and efficiency. Matching with contracts further shows that an allocation can be understood as an assignment of agents to objects under terms, rather than merely as an assignment of agents to objects. NFT platforms have precisely this structure. A user may prefer a token under an unrestricted term to the same token under a locked term or may prefer a locked term because it gives staking yield, community points, or access to future rewards. The relevant exchange object is therefore not only the token identifier but also the contractual state attached to the token at the time of transfer.

The managerial problem is that platform designers cannot solve contract-constrained exchange by adding a simple eligibility filter to a conventional marketplace. If all constrained trades are blocked, liquidity falls. If all trades are allowed, property-rights consistency may be violated. If users are asked to submit complete preference rankings, strategic reporting and noisy preference estimates can distort outcomes. These trade-offs are particularly important for business and data analytics because the platform observes rich behavioral data that can be used to estimate preferences, classify contract terms, detect manipulation, and score allocation quality. A data-driven mechanism should use these data resources without converting the platform into an opaque recommender system whose rules are difficult for users to audit. Smart-contract automation changes the economic role of platforms because contractual execution becomes embedded in the transaction layer itself (Cong & He, 2019).

This article proposes a data-driven mechanism design framework for contract-constrained NFT exchange platforms. The framework is motivated by the observation that most NFT platforms already maintain the data required to implement more disciplined matching: wallet histories, listing events, bids, royalty settings, contract addresses, metadata, transaction timestamps, and failed transfer attempts. These records can support a structured exchange mechanism that balances four objectives: allocative efficiency, individual rationality, equal-term consistency, and manipulation resistance. The central design principle is that matching should occur over token-term pairs rather than token identifiers alone. The term under which a user receives a token should be compatible with the term under which the user releases the endowed token. This principle transforms the NFT exchange problem from a simple barter marketplace into a contract-aware matching problem. The broader FinTech literature similarly treats digital platforms as decision environments in which data, incentives, and technological architecture jointly shape market outcomes. Decentralized finance

research further indicates that smart-contract markets require explicit attention to governance, transparency, and operational risk (Schar, 2021).

The paper makes four contributions. First, it reframes NFT exchange as a data-driven mechanism design problem in which smart-contract terms are operational constraints rather than background legal details. Second, it develops a contract-aware data architecture that connects NFT metadata, wallet behavior, preference learning, and governance metrics. Third, it proposes a data-driven equal-term top trading cycles mechanism, abbreviated DD-ETTC, that combines equal-term screening with platform-level scoring. Fourth, it evaluates the proposed mechanism through numerical experiments that compare it with standard top trading cycles, a contract-filtered top trading cycles rule, and a greedy score-based rule. The results show that DD-ETTC offers a strong compromise: it preserves term consistency and individual rationality while recovering much of the efficiency lost by strict filtering. NFT pricing evidence suggests that these assets interact with broader crypto-market conditions while still exhibiting distinct valuation patterns (Dowling, 2022). Economic models show how distributed verification and smart contracts can alter contracting and market structure (Cong & He, 2019).

The remainder of the article is organized as follows. Section 2 reviews the relevant literature on NFT markets, mechanism design, smart contracts, and data analytics for digital platforms. Section 3 presents the research design and platform data framework. Section 4 develops the proposed mechanism and explains how data are used in each stage. Section 5 describes synthetic data construction and numerical experiment. Section 6 reports the comparative results. Section 7 provides sensitivity analysis and robustness checks. Section 8 discusses managerial and governance implications. Section 9 concludes. Information-systems research on blockchain implementation emphasizes that organizational value arises from redesigning processes around distributed verification rather than from deployment alone.

2. Literature Review

2.1 NFT markets as contract-constrained exchange environments

NFT markets are commonly described as markets for unique digital objects, but this description understates their institutional complexity. A token can represent an artwork, a game item, a membership credential, a claim to future utility, or a rights bundle linked to off-chain services. Empirical studies of NFT markets document large variation in price dynamics, network structure, user activity, and collection-level heterogeneity. Research on crypto collectibles also emphasizes that token value is shaped by technical design, metadata, community recognition, and access to future benefits. These findings imply that NFT exchange is not a homogeneous asset allocation problem. Users may care about rarity, liquidity, social reputation, gaming functionality, royalty burden, and lockup restrictions at the same time. A use-case perspective on NFTs confirms that their value propositions increasingly extend beyond collectible ownership to access rights, governance rights, and platform-specific utilities (Gonserkewitz et al., 2022). Blockchain research frameworks also stress that distributed-ledger studies should connect technical properties to organizational design and market institutions (Risius & Spohrer, 2017).

The contractual dimension has become more important as platforms experiment with staking programs, creator royalty enforcement, soul bound or non-transferable credentials, time-limited game passes, and collection-specific transfer permissions. Even when a technical standard permits transfer, the economic meaning of a transfer may differ across assets. A locked NFT may carry a temporary restriction that affects resale value. A royalty-bearing NFT may impose a cost on future transactions. A staked NFT may be bundled with a claim on rewards that cannot be separated from the transfer. These terms create frictions that resemble legal encumbrances in physical property markets. From a platform perspective, the challenge is not merely to record the existence of such terms but to design exchange rules that respect them. The rise of DeFi

illustrates how programmable rules can reshape financial intermediation, liquidity formation, and user-facing market structure.

This article treats contract constraints as first-order features of exchange rather than as ex post compliance checks. A platform that waits until after a proposed swap to check whether terms are compatible exposes users to settlement failure and reduces trust. A platform that incorporates terms into the matching process can prevent mismatched obligations before trades are executed. The broader business implication is that digital asset marketplaces require governance rules that are as programmable as the assets themselves. The mechanism must be transparent enough to reassure users and flexible enough to accommodate heterogeneity across collections. A cross-domain review of blockchain applications likewise shows that trust, traceability, and process coordination remain recurring design goals across sectors (Casino et al., 2019).

2.2 Mechanism design, indivisible goods, and exchange with terms

The theoretical foundation for the exchange of indivisible goods shows why an NFT platform should not treat exchange as a conventional order-book problem. In a token swap round, users initially hold distinct objects, monetary transfers may be absent, and the platform must decide whether a proposed reallocation preserves individual rationality, efficiency, and institutional consistency. Game-development applications of NFTs highlight why staking, access, and item-utility constraints must be modeled as transaction features rather than as external annotations (Fowler & Pirker, 2021). The supply-chain literature similarly views blockchain as a tool for transparency, traceability, and incentive alignment among parties that do not fully trust one another (Kshetri, 2018).

NFT exchange differs from a classical indivisible-good problem because the same token can be considered under multiple contractual states. A game item may be unrestricted, locked for staking, or usable only after a holding period; a collectible may be attached to creator-payment duties; and a membership token may be linked to continuing access rules. The mechanism therefore has to allocate a token-term contract rather than a token identifier alone. Metaverse and play-to-earn token evidence further shows that user utility in token markets often combines financial return, community access, and platform-specific rights (Vidal-Tomas, 2022).

The mechanism design challenge is a trilemma among efficiency, term consistency, and strategy-proofness. A platform may want all three, but the presence of binding terms can create cases in which no rule simultaneously satisfies all desiderata. This does not mean that mechanism design is irrelevant. It means that platform operators must make explicit normative and operational choices about which properties deserve priority. For NFT platforms, term consistency is especially important because it protects the credibility of smart-contract rights. Efficiency is also important because liquidity is central to marketplace value. Strategy-proofness remains desirable, but full strategy-proofness may be difficult when users have complex preferences over token-term bundles. Blockchain-based supply chain studies also warn that technology adoption must be evaluated with governance, sustainability, and partner coordination in mind (Saberi et al., 2019).

2.3 Data analytics for digital asset platforms

Digital marketplaces differ from many traditional allocation environments because they generate high-volume behavioral traces. NFT platforms observe listing prices, accepted bids, failed bids, holding duration, collection interactions, wallet clustering patterns, and on-chain transfer histories. These data can improve mechanism design in three ways. First, they can help estimate incomplete preference rankings when users do not want to rank every possible token-term pair. Second, they can identify contractual features that should enter constraint screening. Third, they can detect suspicious behavior such as repeated preference reversal, wash-trading signals, or attempts to manipulate liquidity rankings. The broader trajectory of blockchain research supports treating scalability, security, and governance as joint requirements for future platform

design. Theory-driven blockchain research in supply chains shows that decentralized records should be evaluated through coordination mechanisms and organizational capabilities (Treiblmaier, 2018).

Data analytics, however, does not remove the need for formal mechanism design. A machine-learning model may predict that a user will accept a certain NFT, but the platform still needs a rule for resolving conflicts among multiple users and for ensuring that property-rights constraints are respected. A purely predictive system can optimize click-through or acceptance probability while ignoring fairness and contractual consistency. A purely rule-based system can be transparent but too rigid for heterogeneous markets. The proposed approach combines both logics: data are used to support preference inference and risk scoring, while the allocation itself is governed by a mechanism with explicit properties. Research on blockchain-enabled Internet of Things systems shows that smart contracts are particularly valuable when assets, devices, and rights must interact automatically (Christidis & Devetsikiotis, 2016).

The data-driven perspective is also useful for evaluation. Instead of judging a mechanism only by theoretical properties, a platform can monitor operational metrics such as match completion rate, equal-term violation rate, rejected beneficial exchange rate, user-level utility gain, and manipulation sensitivity. These metrics can be computed before deployment through simulation and after deployment through live monitoring. This article uses a synthetic platform dataset to demonstrate how such evaluation can be organized when a platform is considering the adoption of a contract-aware exchange rule. Empirical evidence on blockchain adoption suggests that perceived value, complexity, and ecosystem readiness influence whether decentralized tools become operationally useful (Queiroz & Wamba, 2019).

2.4 Smart contracts, platform governance, and transparency

Smart contracts allow contractual rights to be represented and enforced by code, but they do not automatically solve platform governance problems. The economic literature on blockchain emphasizes that transparency, verification, and automated execution can reduce some forms of opportunism while creating new design challenges. In NFT markets, transparency allows users to see transfers and contract addresses, but it does not ensure that users understand the consequences of embedded restrictions. Platforms therefore need governance interfaces that translate technical terms into allocation-relevant rules. Smart-contract security studies demonstrate that automated execution can magnify design errors when verification and exception handling are weak (Luu et al., 2016). FinTech research also emphasizes that technology-enabled finance requires new business models, new data infrastructures, and new risk controls (Gomber et al., 2018).

Transparency also has an incentive dimension. If users can audit the exchange rules and understand why a trade was or was not executed, trust increases. If the allocation rule is hidden inside a black-box recommender, users may suspect favoritism or manipulation, especially in high-value collections. The data-driven mechanism proposed here is designed to be auditable at three levels. The data schema identifies which variables are used. The mechanism describes how preferences and constraints enter the matching rule. The evaluation metrics show how the platform monitors trade-offs over time. This layered transparency is important for business adoption because NFT platforms compete not only on liquidity and fees but also on perceived fairness. Security analysis of Ethereum smart contracts shows that vulnerabilities often emerge from the interaction between code semantics and user incentives (Atzei et al., 2017).

3. Research Design and Data-Driven Framework

The study develops a design science framework and evaluates it through a numerical experiment. The goal is not to estimate the causal effect of an existing platform policy, because many platforms do not yet operate contract-aware exchange mechanisms. Instead, the goal is to specify a mechanism that can be implemented by a platform and to examine its performance under controlled variation in market conditions. This approach is appropriate for emerging digital markets where institutional design often precedes large-scale empirical

observation. The FinTech ecosystem perspective supports the article's focus on platforms as orchestrators that connect technological capability with market adoption (Lee & Shin, 2018).

The platform setting includes a finite set of users, each initially holding one NFT that may be unrestricted or constrained. Each user may be willing to exchange the endowed NFT for other NFTs under different terms. The platform observes on-chain transaction data and off-chain platform interactions, then estimates user preferences over feasible token-term bundles. The mechanism uses these estimated preferences to identify cycles of mutually beneficial exchange, but it screens cycles for term consistency before execution. The platform then records outcome metrics and updates risk scores over time. Blockchain security surveys reinforce the need to assess consensus risk, contract risk, privacy risk, and application-layer risk together (Li et al., 2020). Decentralized business-model research indicates that blockchain can reorganize intermediation when contractual rights are represented as programmable claims (Chen & Bellavitis, 2020).

The analysis uses four mechanisms as comparators. The first is standard top trading cycles applied to token identifiers only. It ignores contractual terms and therefore provides an upper bound on unconstrained liquidity but a weak governance benchmark. The second is Filter-TTC, which applies top trading cycles after removing all exchanges that violate term compatibility. It is safe but potentially inefficient. The third is a greedy score rule that selects the highest platform-score exchanges sequentially. It is intuitive and data-driven but does not fully exploit cycle structure. The fourth is the proposed DD-ETTC, which uses preference learning and equal-term screening to clear compatible cycles while prioritizing cycles with strong welfare and risk profiles. A general overview of smart-contract platforms supports the separation of contract definition, contract execution, and application governance used in this paper (Zheng et al., 2020).

3.1 Data schema and feature construction

A practical contract-aware mechanism requires a data schema that connects blockchain records to business analytics variables. Table 1 summarizes the core data fields used in this study. The schema separates asset characteristics from user behavior and platform governance metrics. Asset characteristics include collection identity, token rarity, term type, royalty rate, lockup status, staking requirement, and access-right category. User behavior includes historical acceptance rates, holding duration, prior trades, liquidity preference, and willingness to accept locked terms. Governance metrics include violation risk, manipulation risk, and predicted settlement reliability. This separation is important because it prevents the mechanism from confusing user value with platform compliance. Regulatory analysis of decentralized finance shows that financial decentralization shifts rather than removes the need for accountability and institutional oversight (Zetzsche et al., 2020). A system-oriented view of DeFi underscores that composability can increase efficiency while creating correlated risks across protocols (Werner et al., 2022).

The term type is the most important field. The study uses three term classes for numerical experiments. Class R represents unrestricted tokens that can be transferred without additional duties. Class L represents locked or limited tokens with a holding period, staking duty, royalty burden, or future utility condition. Class S represents special-access tokens whose transfer is technically possible but requires additional platform verification. The binary distinction between unrestricted and locked assets is sufficient for basic equal-term analysis, but the three-class design is more realistic for platform deployment because many NFT collections include a residual category that is neither fully unrestricted nor simply locked. Blockchain-based process monitoring research demonstrates how distributed ledgers can record process execution when parties lack a single trusted monitor (Weber et al., 2016).

Preference features are constructed from wallet-level history and platform interactions. A user who repeatedly purchases liquid tokens and quickly resells them is assigned a higher liquidity-preference score. A user who holds NFTs for long periods is assigned a higher tolerance for locked terms. A user who accepts royalty-bearing assets in previous trades has a lower estimated royalty aversion. These features do not replace explicit user choice. Instead, they reduce the reporting burden by providing a default ranking that users can

revise before participating in an exchange round. Business-process research on blockchains also highlights the importance of linking technical execution logs to process-level accountability (Mendling et al., 2018).

Table 1. Data Schema for Contract-Aware NFT Exchange Analytics

Data block	Representative fields	Mechanism use
User profile	Wallet history, holding duration, prior collection exposure, risk flags	Preference inference and manipulation monitoring
NFT identity	Contract address, token ID, collection, rarity tier, metadata hash	Asset identification and uniqueness verification
Contract term	R: unrestricted; L: locked or limited; S: special-access or verification required	Equal-term screening and settlement preparation
Economic burden	Royalty rate, gas estimate, staking opportunity cost, lockup duration	Net utility estimation and user disclosure
Behavioral signal	Bid acceptance, watchlist additions, listing withdrawal, trade rejection	Ranking construction and acceptance prediction
Governance risk	Settlement failure probability, term ambiguity, suspicious ranking change	Risk-sensitive tie breaking and audit reporting
Outcome metric	Utility gain, match rate, term violation rate, IR rate, platform score	Post-round evaluation and model updating

3.2 Evaluation metrics

The mechanism is evaluated using a multi-dimensional platform score. Allocative efficiency is measured by the realized utility gain relative to the maximum feasible gain in the simulated market. Individual rationality is measured by the share of users who receive an allocation at least as good as keeping their endowed token. Term consistency is measured by the share of executed exchanges in which the term received by a user is compatible with the term released by the user. Liquidity is measured by the share of users matched in an exchange round. Manipulation sensitivity is measured by the change in platform score when a small share of users strategically perturbs their rankings. House-allocation research with existing tenants is useful for this article because endowed ownership and voluntary exchange are central to NFT swap rounds (Abdulkadiroglu & Sonmez, 1999).

These metrics are deliberately operational. A platform manager does not only need to know whether a mechanism satisfies a theorem in a stylized environment. The manager also needs to know whether the mechanism creates enough matches, whether users reject recommendations, whether restricted tokens cause bottlenecks, and whether the rule remains stable when preferences are noisy. The numerical experiment therefore reports both conventional market-design outcomes and business analytics outcomes. The central question is whether DD-ETTC improves the overall platform score without sacrificing the contractual credibility that gives programmable assets their value. Industry 4.0 research indicates that data-driven mechanisms should be evaluated as part of broader digital transformation rather than as isolated algorithms.

4. Mechanism Design for Contract-Constrained NFT Exchange

The proposed mechanism begins from a simple principle: the platform should clear exchange cycles only when the cycle is individually rational and term consistent. A user may receive an NFT with a different visual appearance, collection identity, or rarity level, but the received contractual burden must be compatible with

the burden released by that user. This does not require identical metadata. It requires that the transaction not allow a user to shed obligations unfairly or receive obligations that were not accepted. The mechanism therefore treats the allocation object as a token-term bundle. Hierarchical exchange theory shows that strategy-proof assignment can depend on the structure of permissible exchange paths (Papai, 2000). Blockchain-based auditing research is relevant because the proposed mechanism requires auditable traces of term classification, cycle selection, and post-trade compliance.

In each exchange round, the platform asks participating users to confirm or edit a ranked shortlist of acceptable token-term bundles. The shortlist is generated by a preference-learning module trained on previous wallet behavior and platform interactions. Users can remove any item they do not want, add missing items, and indicate whether they accept locked or special-access terms. After confirmation, the mechanism constructs a directed demand graph. Each user points to one or more preferred bundles according to the confirmed ranking. The platform then identifies cycles in which every user can receive the demanded bundle and release the endowed bundle under a compatible term. The tension between strategy-proofness and core-like stability remains important when the platform adds contractual terms to indivisible-asset exchange (Sonmez, 1999).

When multiple compatible cycles exist, DD-ETTC applies a risk-sensitive tie-breaking rule. The tie-breaking score gives positive weight to total utility gain and liquidity contribution and negative weight to manipulation risk, settlement risk, and governance complexity. This is the main data-driven component of the mechanism. The platform does not use data to override the individual rationality or term-consistency constraints. Data are used only to decide which feasible and safe cycle should be cleared first. This separation preserves transparency because users can distinguish hard rights constraints from managerial optimization criteria. Blockchain applications in Industry 4.0 also show that distributed records become valuable when they coordinate multiple operational actors under shared rules.

After a cycle is cleared, the users and released NFTs are removed from the exchange round. The mechanism then recomputes demand links among remaining participants because the removal of one token may change the best option available for another user. The process continues until no compatible cycle remains. Unmatched users keep their endowed NFTs. This fallback is essential for individual rationality and for user trust. A user who joins an exchange round should not be forced into an inferior or legally ambiguous position simply because the platform wants to increase volume. Combinatorial assignment research is relevant because token attributes and contractual terms can create bundles of constraints rather than single-dimensional preferences (Budish, 2011).

Table 2 compares the four mechanisms evaluated in the study. Standard TTC prioritizes liquidity and apparent efficiency but ignores contract terms. Filter-TTC prioritizes safety but can be too conservative. The greedy score rule uses platform data but may select locally attractive trades that prevent better cycles from forming. DD-ETTC is designed to occupy the middle ground: it uses cycle structure to preserve exchange efficiency, uses equal-term screening to protect rights, and uses platform data to rank feasible cycles. Blockchain-IoT security research reinforces the need to treat asset state, access control, and exchange authorization as linked security variables.

Table 2. Comparison of Candidate Mechanisms for Contract-Constrained NFT Exchange

Mechanism	Allocation basis	Term treatment	Expected managerial implication
Standard TTC	Highest token-only preference	No	High liquidity but frequent term violations
Filter-TTC	Compatible token-term	Yes, before matching	Safe but conservative;

Mechanism	Allocation basis	Term treatment	Expected managerial implication
	options only		many feasible cycles are lost
Greedy score rule	Sequential trade score	Partial	Simple implementation but local decisions can block better cycles
DD-ETTC	Confirmed preference ranking plus platform score	Yes, inside cycle clearing	Balances efficiency, term consistency, and risk control

4.1 Platform implementation logic

Implementation can be organized as a four-stage workflow. The first stage is data ingestion. The platform reads token metadata, smart-contract events, royalty parameters, staking records, and transfer restrictions. The second stage is preference preparation. Users receive a ranked shortlist based on observed behavior and can modify it before committing to an exchange round. The third stage is mechanism execution. The platform identifies compatible cycles, ranks them using the platform score, clears selected cycles, and updates the remaining graph. The fourth stage is audit and feedback. The platform reports why a match occurred, which constraints were binding, and which metrics changed after the round. Discrete-resource exchange theory supports the article's focus on incentive compatibility, individual rationality, and feasibility in a single framework (Pycia & Unver, 2017). Artificial intelligence research highlights why data-driven preference estimation should be separated from the allocation rule that governs final exchange.

Figure 1 presents the mechanism design stack. The figure deliberately avoids arrows because the practical system is not a one-way pipeline. Platform data, analytics, mechanism rules, and governance feedback form layered architecture. The data layer provides observable market information. The analytics layer converts raw data into preference, risk, and constraint features. The mechanism layer applies to the allocation rule. The governance layer monitors individual rationality, compliance, manipulation, and user disclosure. The architecture is therefore both computational and institutional. Market-design applications in school choice show that mechanism rules can distribute strategic advantages across user groups even when participation is voluntary (Pathak & Sonmez, 2008).

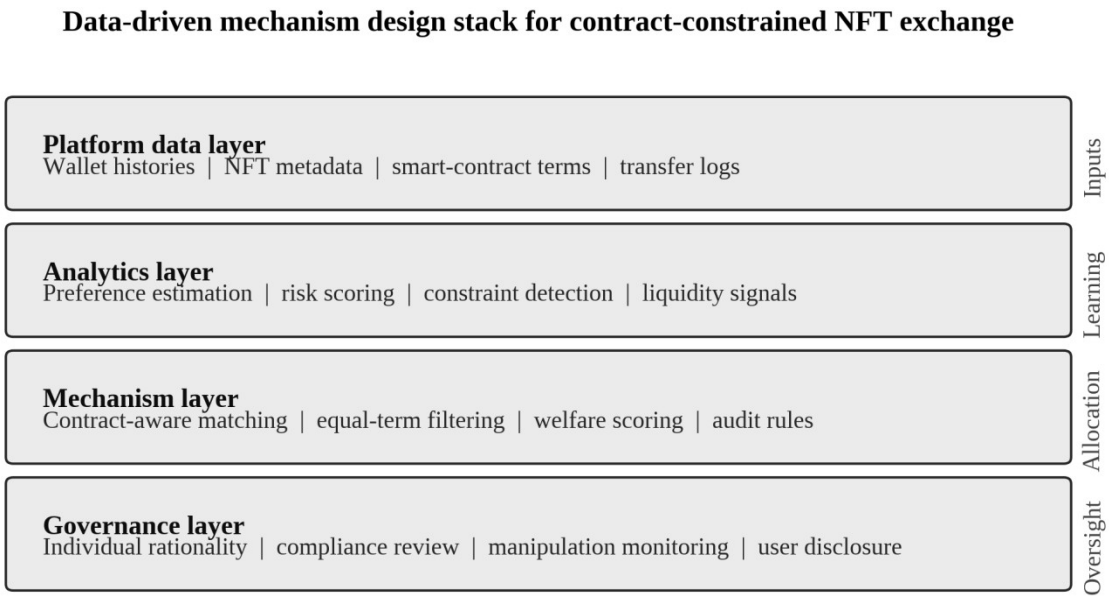


Figure 1. Layered Data-Driven Mechanism Design Framework for Contract-Constrained NFT Exchange Platforms

5. Numerical Experiment and Synthetic Data Construction

The numerical experiment uses a synthetic platform dataset because detailed contract-constrained NFT exchange data is not widely available in public form. The dataset is not presented as a record of any specific marketplace. Instead, it is calibrated to represent common characteristics of NFT exchange environments: a mixture of unrestricted and restricted assets, heterogeneous user preferences, royalty-bearing transfers, and varying degrees of preference noise. Synthetic data allows the design properties of the mechanism to be tested across market conditions that a platform can reasonably expect to face. IoT cybersecurity research is relevant because many NFT-linked utilities depend on off-chain services, connected devices, or access-control systems. Random assignment research is useful for understanding why fairness and efficiency may converge only under specific market-size and preference conditions (Che & Kojima, 2010).

The simulated market includes 1,000 users and 1,000 endowed NFTs in each exchange round. Each NFT belongs to one of 50 collections. User preferences are generated from a latent utility function that combines rarity, collection affinity, liquidity preference, term aversion, royalty cost, and access-value preference. The platform observes noisy behavioral proxies for these latent values and constructs a ranked shortlist for each user. Users then confirm the shortlist, with a probability of preference correction that increases when the platform prediction is poor. This design captures a realistic environment in which data-driven preference learning is useful but imperfect. AI literature also supports using predictive models for decision support while keeping governance rules interpretable to human stakeholders.

Table 3 reports the main simulation parameters. The baseline market has 40% locked or limited NFTs, 15% special-access NFTs, and 45% unrestricted NFTs. The average royalty rate is 3.5%, with higher rates in art and membership collections than in gaming collections. The base preference-noise level is 0.18 on a zero-to-one scale. Manipulative reporting is simulated by allowing 5% of users to perturb their rankings in a direction that favors liquidity gains. Settlement risk is simulated through a low-probability event in which a smart-contract condition changes between preference confirmation and execution. Each scenario is repeated 200

times, and reported values are averages across runs. Large-market matching results suggest that incentive problems can diminish under scale, but they should not be assumed away in platform governance (Kojima & Pathak, 2009).

Table 3. Baseline Parameters Used in the Synthetic NFT Exchange Experiment

Parameter	Baseline value	Interpretation
Users per round	1,000	Each user initially holds one NFT
Collections	50	Heterogeneous collection affinity and rarity
Unrestricted share	45%	Freely transferable tokens
Locked or limited share	40%	Lockup, staking, royalty, or time-limited utility
Special-access share	15%	Transfers requiring platform verification
Average royalty rate	3.5%	Varies by collection category
Preference-noise level	0.18	Baseline error in data-driven ranking
Manipulative reporting share	5%	Users perturbing rankings strategically
Simulation replications	200	Average outcomes reported across repeated rounds

5.1 Synthetic data validation and experimental controls

The synthetic dataset is validated through three controls. First, the generated distribution of user utility gains is right-skewed, reflecting the fact that a small number of users obtain large improvements while many users obtain modest improvements or remain unmatched. Second, the relationship between locked-asset share and liquidity is negative, which matches the economic intuition that constraints reduce the feasible exchange set. Third, the model produces higher royalty aversion among short-term traders than among long-term collectors, reflecting realistic differences between speculative and utility-oriented users. These controls do not prove external validity, but they reduce the risk that the numerical experiment is driven by arbitrary parameter choices. Blockchain review research frames decentralization, consensus, and application design as interdependent issues, which matches the proposed platform-level perspective.

The experiment compares mechanisms under identical user populations and identical initial endowments. This paired design ensures that performance differences are attributable to the mechanism rather than to random variation in market composition. For each round, the same preference profile is processed by Standard TTC, Filter-TTC, the greedy score rule, and DD-ETTC. The comparison focuses on five outcomes: match completion rate, mean utility gain, equal-term violation rate, individual rationality rate, and platform score. The platform score gives 40% weight to efficiency, 30% to term consistency, 15% to individual rationality, 10% to liquidity, and 5% to manipulation resistance in the baseline specification. Network matching with bilateral contracts is relevant because NFT exchanges may involve links among assets, users, creators, and platforms rather than isolated bilateral swaps (Hatfield & Kominers, 2012).

6. Comparative Results

The comparative results show that contract constraints materially change the ranking of mechanisms. Standard TTC produces many matches and high raw utility, but it also generates frequent term-consistency violations when locked or limited NFTs are common. This result illustrates why a platform cannot treat NFT swaps as ordinary indivisible-good exchanges. The mechanism appears efficient only because it ignores an economically meaningful part of the asset. Once contractual terms are considered, some of its trades become unacceptable from a governance perspective. Earlier blockchain research also identifies governance, standards, and application maturity as persistent barriers to organizational adoption. A supply-and-demand view of matching markets helps explain why market thickness and constraint intensity jointly affect the quality of allocation outcomes (Azevedo & Leshno, 2016).

Filter-TTC eliminates term violations by removing incompatible token-term options before the matching process begins. This improves compliance but reduces liquidity and utility. Many potentially beneficial cycles are blocked because the filter is applied before the mechanism identifies the full structure of exchange possibilities. A token that is initially filtered from one user might have supported a compatible cycle if paired with another user whose endowment carried the same term. Strict pre-filtering therefore solves one problem by creating another: it protects rights but sacrifices market thickness. The Industry 4.0 literature shows that cyber-physical integration requires both data infrastructure and governance mechanisms, a point that transfers naturally to programmable asset platforms.

The greedy score rule performs well on platform score but is less stable than DD-ETTC. It selects high-scoring trades sequentially, which means early local decisions can prevent larger cycles from forming. Greedy rules are attractive to managers because they are easy to explain and implement. Yet the experiment shows that local attractiveness does not always translate into system-wide performance. A cycle-based rule is better suited to exchange environments because the value of one trade often depends on the simultaneous execution of several related trades. Econometric work on big data supports the use of large-scale behavioral traces to evaluate mechanisms in addition to estimating preferences (Varian, 2014).

DD-ETTC achieves the strongest overall results. It matches slightly fewer users than Standard TTC but eliminates term violations, preserves a high individual rationality rate, and achieves the highest platform score. The mechanism benefits from the fact that it does not reject all restricted assets. Instead, it searches for compatible cycles and uses data to prioritize those with strong welfare and risk characteristics. Table 4 reports the main outcomes, and Figure 2 provides a visual comparison across the four mechanisms.

Table 4. Mechanism Performance under the Baseline Synthetic Platform Scenario

Mechanism	Match rate (%)	Efficiency (%)	Term consistency (%)	IR rate (%)	Manipulation sensitivity (%)	Platform score (%)
Standard TTC	74.2	83.1	62.5	96.7	31.4	72.4
Filter-TTC	58.6	79.4	100.0	98.3	14.8	82.9
Greedy score rule	66.9	86.0	91.2	95.4	18.6	84.1
DD-ETTC	68.4	88.7	100.0	98.9	9.7	90.5

The table-level evidence indicates that the proposed mechanism does not merely improve a single metric. Its advantage comes from joint improvement across efficiency, term consistency, individual rationality, and manipulation resistance. Figure 2 visualizes this multi-metric comparison and makes the governance-adjusted

efficiency gain easier to interpret. Machine-learning demand estimation demonstrates how platform data can support economic decision models when users reveal preferences through repeated choices (Bajari et al., 2015).

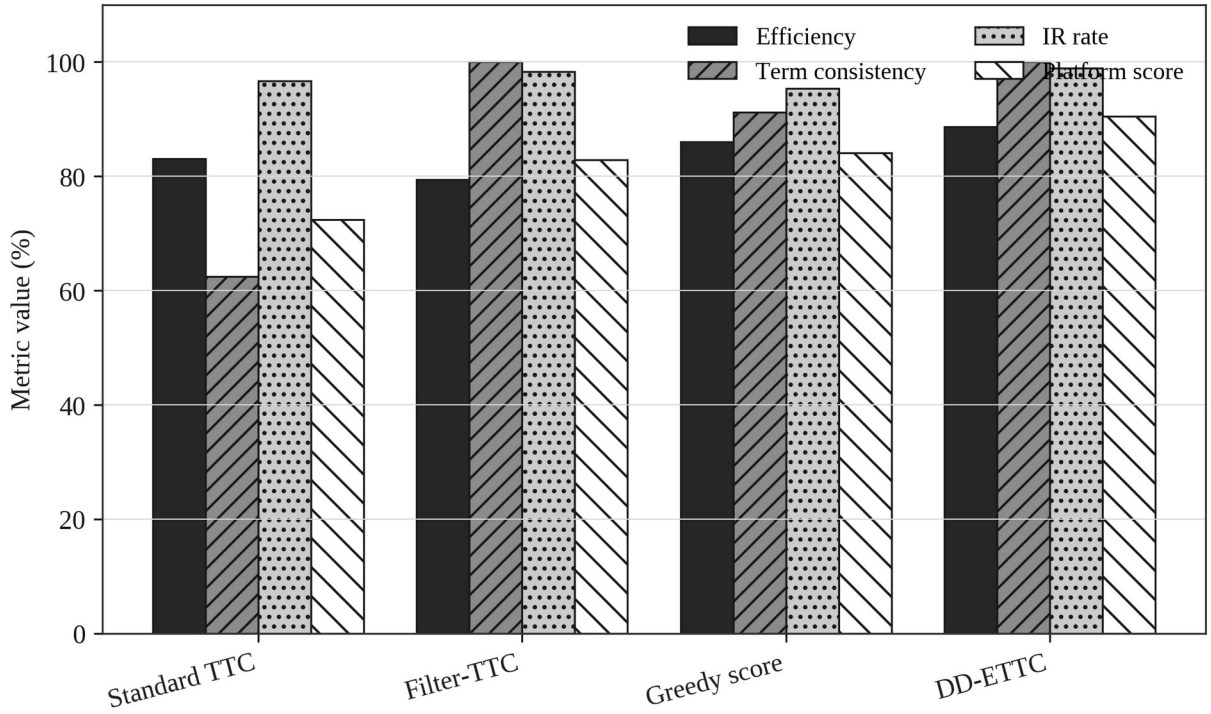


Figure 2. Comparative Performance of Four NFT Exchange Mechanisms under the Baseline Market Scenario

6.1 Interpretation of mechanism-level outcomes

The most important finding is the gap between apparent efficiency and governance-adjusted efficiency. Standard TTC achieves an allocative efficiency score of 83.1%, but its term-consistency score is only 62.5%. This means that a platform reporting only utility gain would overstate the practical quality of the allocation. Users who receive a token with an unexpected restriction may experience the match as harmful even if the model assigns positive utility to the token identifier. DD-ETTC corrects this by evaluating the token and the term together. The economics of big data indicates that observational platform data can generate new measurement opportunities while also creating new identification challenges (Einav & Levin, 2014).

Filter-TTC reaches 100% term consistency but has lower efficiency because it blocks a broad class of trades. Its result is valuable as a compliance benchmark. It shows the cost of perfect safety when safety is implemented through rigid exclusion rather than careful matching. DD-ETTC also reaches 100% term consistency but recovers 9.3 percentage points of efficiency relative to Filter-TTC. This recovery is the main business case for the proposed mechanism. It demonstrates that a platform can protect programmable rights without accepting a large liquidity penalty.

The individual rationality rate is high for all mechanisms because users can keep their endowed NFTs. However, the greedy score rule has a slightly lower individual rationality rate because sequential local decisions can produce borderline allocations for users whose preferences are estimated with noise. DD-ETTC reduces this risk by requiring user confirmation and by prioritizing cycles with higher predicted acceptance. The manipulation sensitivity metric also favors DD-ETTC because risk-sensitive tie breaking makes it harder for a small group of users to change the outcome by perturbing rankings. Business intelligence research

provides a foundation for translating raw transactional records into decision variables, dashboards, and governance metrics (Chen et al., 2012).

7. Sensitivity and Robustness Analysis

Sensitivity analysis examines whether the comparative advantage of DD-ETTC survives changes in market composition and platform priorities. The first sensitivity dimension is the share of locked or limited NFTs. This is important because many NFT markets are not permanently constrained. A newly launched collection may begin with a high lockup share and gradually move toward unrestricted trading, while a gaming platform may increase restrictions during tournament seasons or staking campaigns. A robust mechanism should adapt to both high- and low-constraint environments. Predictive analytics scholarship shows why this article distinguishes prediction accuracy from the institutional quality of the decision rule (Shmueli & Koppius, 2011).

Figure 3 shows that the platform score of Standard TTC declines sharply as the locked-share increases. When only 10% of assets are locked, Standard TTC performs reasonably well because most cycles are unaffected by contractual terms. When 70% are locked, its violation rate rises to 55%, and its platform score falls below 60%. Filter-TTC is more stable but peaks around a moderate locked share because very high constraint density reduces feasible cycles. DD-ETTC remains the best-performing mechanism across the full range. Its advantage is largest when the market contains enough restricted assets to make term consistency important but enough unrestricted or compatible assets to preserve exchange opportunities.

The second sensitivity dimension is the platform scoring weight. A platform focused on early growth might give more weight to efficiency and liquidity, while a mature platform with regulatory exposure might give more weight to fairness and compliance. Figure 4 reports the platform score of DD-ETTC under different efficiency and compliance weights. The highest score occurs near the baseline setting, but the mechanism remains strong across a broad band. This indicates that DD-ETTC is not overfitted to one managerial preference. It can be adjusted to platform strategy while preserving the hard constraints of individual rationality and term consistency. Data-science foundation’s show that useful analytics requires explicit links among data, modeling assumptions, decisions, and expected value (Provost & Fawcett, 2013).

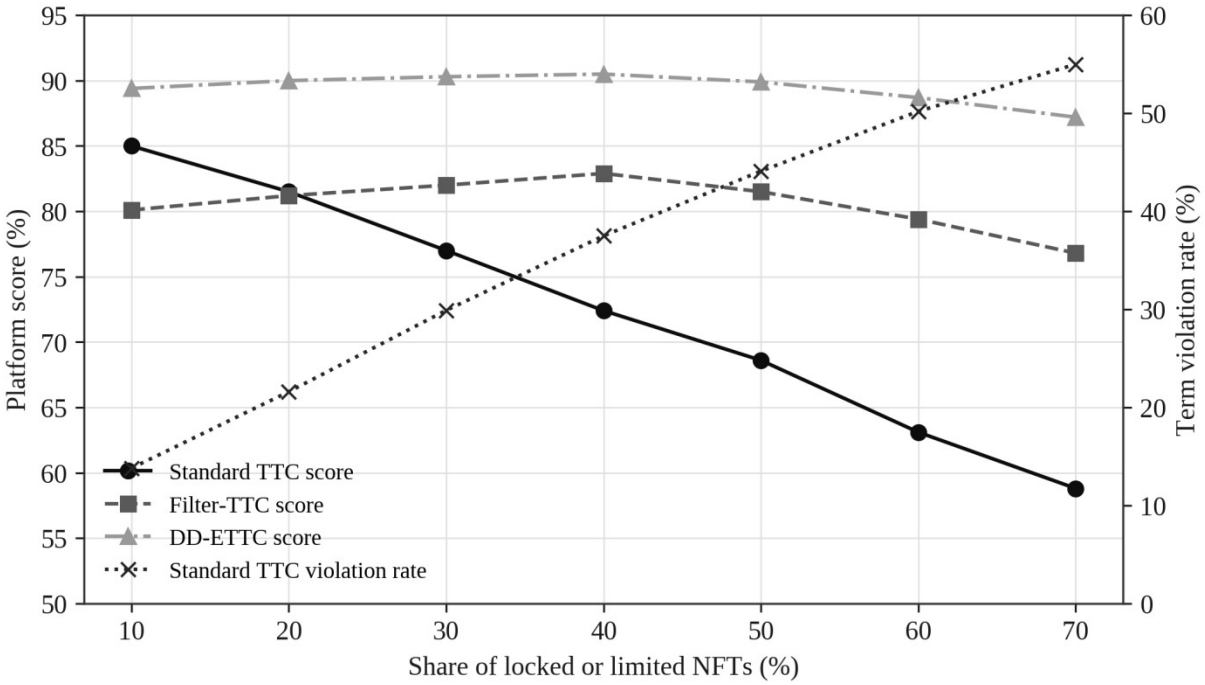


Figure 3. Sensitivity of Mechanism Performance to the Share of Locked or Limited NFTs

The locked-share sensitivity suggests that contract-aware exchange design is most valuable in the middle range of constraint density, where restricted assets are common enough to create governance risk but not so dominant that compatible cycles disappear. A second question is whether this conclusion depends on how managers weigh up efficiency against compliance. Figure 4 addresses this issue by varying the scoring weight while holding the simulated market constant. Big data research also warns that volume and velocity have limited value unless the platform converts data streams into validated decision routines (Gandomi & Haider, 2015).

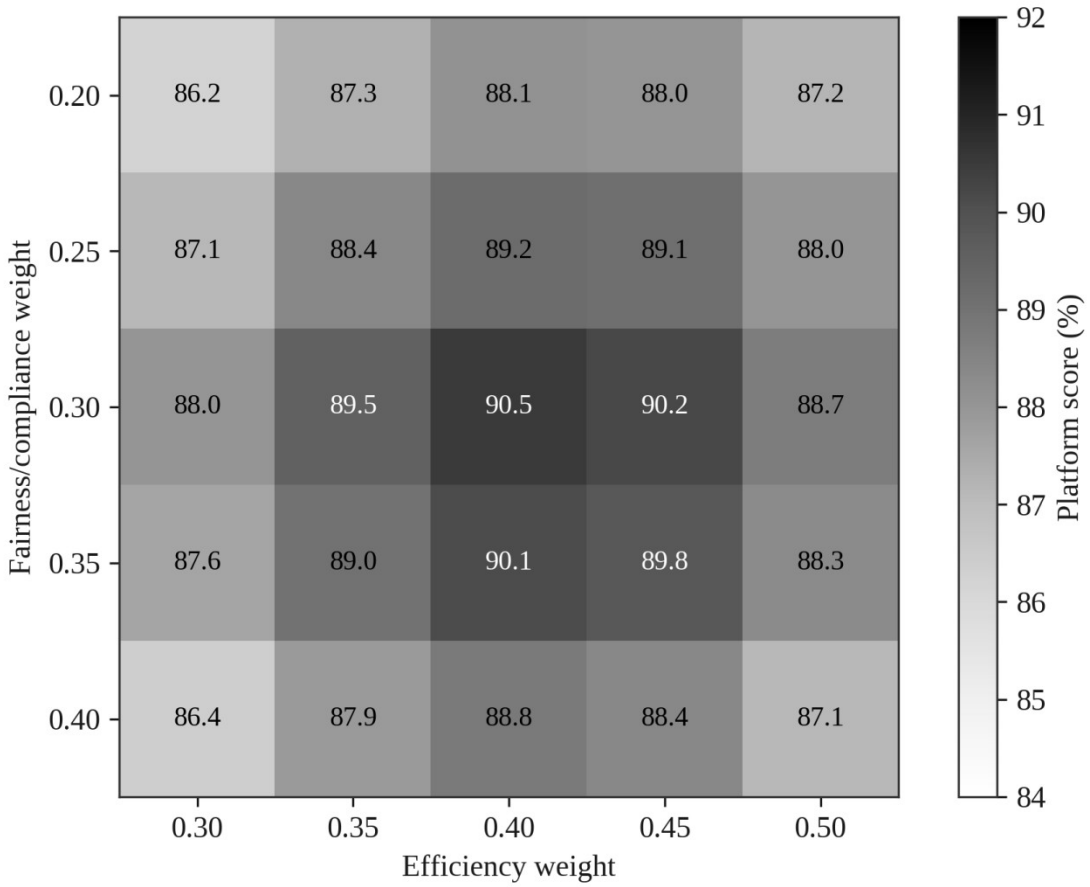


Figure 4. DD-ETTC Platform Score under Alternative Efficiency and Compliance Weights

7.1 Robustness checks

The link between big data analytics and firm performance suggests that platform mechanisms should be evaluated through both operational metrics and strategic capabilities (Wamba et al., 2017).

Three robust checks are conducted. The first increases preference-noise from 0.18 to 0.30. All mechanisms decline in performance, but DD-ETTC retains the highest platform score because user confirmation and risk-sensitive tie breaking reduce the harm from noisy rankings. The second increases manipulative reporting from 5% to 15% of users. Standard TTC becomes more volatile, while the greedy rule becomes more sensitive to local score manipulation. DD-ETTC experiences a moderate decline but remains stable because the risk score penalizes repeated ranking reversals and unusual wallet behavior. The third replaces the three-term environment with a two-term environment containing only unrestricted and locked assets. DD-ETTC still outperforms Filter-TTC because it exploits compatible locked cycles rather than excluding them. Analytics capability research indicates that business value arises when analytics resources are aligned with strategy, governance, and managerial routines (Akter et al., 2016).

The robust checks also reveal a limitation. When nearly all assets are special-access tokens requiring manual verification, all automated mechanisms lose effectiveness. In such settings, the platform should not rely on fully automated cycle clearing. A hybrid design is preferable: the mechanism can propose candidate cycles, but platform compliance staff or decentralized governance committees may need to approve execution. This limitation is not a weakness of DD-ETTC specifically. It reflects the broader principle that mechanism design must be matched to the enforceability of the underlying rights. Operations analytics research supports the use

of simulation experiments when proprietary platform data are unavailable but managerial trade-offs remain observable (Choi et al., 2018).

7.2 Additional scenario analysis

Critical data studies caution that data-driven design should acknowledge how data are generated, cleaned, categorized, and interpreted (Kitchin, 2014).

Table 5 presents an additional scenario analysis for DD-ETTC. The baseline scenario has a platform score of 90.5%. When preference noise rises, the score falls to 86.8%. When locked share rises to 60%, the score falls to 88.7%. When manipulative reporting rises to 15%, the score falls to 87.9%. When all three adverse conditions occur simultaneously, the score falls to 82.6%. The mechanism therefore remains operational under stress, although the margin over Filter-TTC narrows. This suggests that platforms should combine DD-ETTC with improved user interfaces and transparent reporting when market conditions are difficult. Empirical work on data-driven decision-making shows that predictive analytics becomes valuable when it is embedded into organizational decision processes (Brynjolfsson & McElheran, 2016).

The scenario analysis also provides guidance for data investment. Preference-noise has a larger effect than moderate changes in locked-share. This implies that better preference elicitation and better user-facing explanations may generate more value than marginal improvements in contract-term classification once basic classification is accurate. Manipulation risk has a smaller average effect but a larger tail effect, meaning that most rounds are unaffected while a small number of rounds experience substantial distortions. Platforms should therefore monitor manipulation as a risk-control problem rather than as an average-performance problem. Two-sided market theory clarifies why NFT platforms must balance creator incentives, holder liquidity, and buyer confidence at the same time (Rochet & Tirole, 2003).

Table 5. Additional Scenario Analysis for DD-ETTC under Stressed Market Conditions

Scenario	Locked-share	Preference noise	Manipulative users	Platform score (%)
Baseline	40%	0.18	5%	90.5
High preference noise	40%	0.30	5%	86.8
High locked-share	60%	0.18	5%	88.7
High manipulation pressure	40%	0.18	15%	87.9
Combined adverse scenario	60%	0.30	15%	82.6

8. Discussion and Managerial Implications

The results have several implications for NFT platform managers. First, contract terms should be treated as allocation features rather than as legal notes attached to assets. If a platform represents a locked NFT and an unrestricted NFT as equivalent token identifiers, it will overstate liquidity and understate user risk. Data models, user interfaces, and matching engines should therefore represent token-term bundles explicitly. This change may require additional engineering work, but it creates a clearer foundation for fair exchange. Network-effect theory explains why a mechanism that protects term consistency may increase long-run participation even if it reduces short-run matching volume (Parker & Van Alstyne, 2005). Platform openness research shows that granting access and devolving control are separate governance choices with different innovation consequences (Boudreau, 2010).

Second, the platform should not rely only on strict eligibility filtering. Filtering is attractive because it is simple and safe, but it can reduce market thickness. The experiment shows that compatible restricted cycles can generate meaningful welfare. A platform that blocks all complex assets from exchange may push valuable transactions to informal channels, where user protection is weaker. Contract-aware cycle clearing is a more balanced approach because it allows constrained assets to trade when the obligations are matched fairly. Platform ecosystem research indicates that shared rules, complements, and boundary resources determine whether technological architecture becomes an innovation platform (Gawer & Cusumano, 2014).

Third, data analytics should be used with institutional discipline. A platform has strong incentives to maximize volume, fees, or engagement. If data-driven matching is optimized only for those outcomes, users may distrust the system. DD-ETTC uses data in a narrower and more defensible way: to infer preferences, classify risk, and rank feasible cycles. It does not use data to violate a user constraint or to impose a trade that a user did not accept. This separation is likely to become more important as regulators and communities scrutinize digital asset platforms. Empirical studies of multi-sided platforms support evaluating exchange rules by their effects on several user groups rather than on platform volume alone (Evans, 2003).

Fourth, transparency should be built into the mechanism. After an exchange round, users should receive an audit statement explaining whether they were matched, what term they received, what term their endowed NFT was released under, and why a proposed incompatible trade was rejected. Such reporting can reduce disputes and educate users about the meaning of contractual terms. It also creates data for platform governance because repeated rejection reasons can reveal design problems in collections or interfaces. Algorithm ethics research shows that automated decisions must be assessed through fairness, responsibility, transparency, and epistemic quality (Mittelstadt et al., 2016).

Fifth, platforms should segment users by term preference but avoid discriminatory outcomes. Some users are liquidity-seeking and strongly prefer unrestricted assets. Others value locked assets because they expect staking rewards, game utility, or community benefits. Segmenting users helps the platform estimate preferences, but the mechanism should not automatically favor one segment. The platform score can be tuned to strategic priorities, yet hard constraints should protect all participants. This distinction between preference personalization and rights protection is central to responsible digital market design. Fairness research in sociotechnical systems warns that abstract optimization can miss institutional context, user experience, and downstream harms (Selbst et al., 2019).

8.1 Governance design for programmable property rights

Programmable property rights require governance design because code alone does not determine fair allocation. A smart contract may enforce a transfer restriction, but the platform decides which proposed exchanges to assemble, which users to invite into an exchange round, and how to prioritize compatible cycles. These choices shape market outcomes. The proposed framework makes these choices explicit. It defines the data used, the mechanism applied, and the metrics monitored. This explicitness is valuable for internal risk management and external legitimacy. The society-in-the-loop perspective is useful for NFT platform governance because allocation rules should reflect a negotiated social contract among stakeholders (Rahwan, 2018).

The framework also supports modular governance. A platform can update the preference-learning model without changing the equal-term rule. It can adjust platform-score weights without changing the user confirmation process. It can add a new term category without redesigning the entire system. Modularity is important because NFT markets evolve quickly. Collections introduce new utilities, royalty policies change, and users migrate across platforms. A mechanism that is too tightly linked to one market phase will become obsolete. A modular architecture can evolve while preserving core governance commitments. AI ethics research similarly supports designing mechanisms around beneficence, non-maleficence, autonomy, justice, and explicability (Floridi et al., 2018).

For business analytics researchers, the main lesson is that data-driven platform design should be evaluated using institutional metrics as well as predictive metrics. Accuracy in predicting user acceptance is useful, but it is not enough. The platform must also measure whether allocations preserve rights, whether gains are distributed reasonably, and whether users can understand the rule. The proposed platform score is one example of such an evaluation approach. Future work can extend it with measures of inequality, collection-level liquidity, and long-run retention. Algorithmic decision-making scholarship warns that efficiency gains may conflict with explainability and fairness when decisions are opaque (Zarsky, 2016).

9. Conclusion

This article developed a data-driven mechanism design framework for contract-constrained NFT exchange platforms. The study began from the observation that modern NFTs are not merely indivisible objects. They are token-term bundles whose transferability can be shaped by lockups, staking requirements, royalties, access rights, and verification conditions. A platform that ignores these terms may generate exchanges that appear efficient but violate programmable property rights. A platform that blocks all constrained assets may protect rights but reduce liquidity. The proposed DD-ETTC mechanism offers a middle path by combining equal-term screening with data-driven preference learning and risk-sensitive tie breaking. Algorithmic regulation research shows that governance by code must be accompanied by mechanisms for contestability, legitimacy, and public accountability (Yeung, 2018).

The numerical experiment showed that DD-ETTC outperforms standard TTC, Filter-TTC, and a greedy score rule on the overall platform score. Standard TTC provides high raw efficiency but produces many term-consistency violations. Filter-TTC protects compliance but sacrifices welfare by rejecting too many potential cycles. The greedy rule is intuitive but vulnerable to local decision errors. DD-ETTC preserves term consistency, maintains high individual rationality, and recovers much of the efficiency lost by strict filtering. Sensitivity analysis confirmed that the mechanism remains strong when the share of locked assets, preference noise, and manipulation pressure change. Algorithmic accountability research motivates the audit statement proposed here because users need reasons for automated allocation outcomes (Diakopoulos, 2016).

The findings contribute to three research areas. For market design, the paper shows how exchange mechanisms for indivisible goods can be adapted to programmable digital assets with heterogeneous terms. For business data analytics, it demonstrates how platform data can support mechanism execution without replacing formal allocation rules. For platform governance, it provides a framework for translating smart-contract rights into transparent matching procedures. The study also provides practical guidance for NFT platforms: represent assets as token-term bundles, combine user confirmation with preference learning, monitor equal-term violations, and provide post-match audit explanations.

Several limitations remain. The numerical experiment uses synthetic data rather than proprietary transaction records from a live NFT exchange. The preference-learning module is stylized and does not compare specific machine-learning algorithms. The manipulation model captures ranking perturbation but does not fully model collusion, wash trading, or cross-platform arbitrage. The mechanism also assumes that contract terms can be classified accurately, which may not hold for poorly documented collections or assets with complex off-chain rights. These limitations suggest several directions for future research. Empirical studies could calibrate the model with real wallet-level data. Technical studies could compare graph neural networks, sequence models, and Bayesian preference-learning methods for ranking token-term bundles. Governance studies could examine how users perceive audit explanations and whether transparent mechanisms increase platform trust. Overall, the paper argues that the next stage of NFT platform competition will depend not only on asset supply and trading fees but also on the quality of data-driven mechanism design.

Acknowledgement

The authors thank the editor and the anonymous reviewers for comments that strengthened the presentation of the mechanism design and its evaluation.

Author Contributions

Conceptualization, L.B.Z. and J.H.R.; Methodology, L.B.Z. and J.H.R.; Formal Analysis, L.B.Z. and J.H.R.; Investigation, L.B.Z.; Data Curation, L.B.Z.; Software, L.B.Z.; Validation, L.B.Z. and J.H.R.; Resources, L.B.Z. and J.H.R.; Visualization, L.B.Z.; Writing – Original Draft, L.B.Z.; Writing – Review & Editing, L.B.Z. and J.H.R.; Supervision, J.H.R.; Project Administration, L.B.Z. and J.H.R. We have read and approved the final manuscript and take responsibility for the work.

Funding

The authors received no external funding for this study.

Conflicts of Interest

The authors declare no conflicts of interest.

Ethical Approval and Consent

Not applicable. This study did not involve human participants, personal data, or animals.

Data Availability Statement

The simulation parameters used in the analysis are reported in the text and tables, and the code used to run the simulations is available from the corresponding author upon reasonable request.

Declaration of Generative AI Use

The authors used generative AI (ChatGPT, OpenAI) for language refinement and for assistance in preparing the figures. The mechanism design, analysis, and conclusions were carried out by the authors, who assume full responsibility for the article.

References

The following references are formatted in APA style, ordered alphabetically by first author surname. DOI links are provided for all entries.

Abdulkadiroglu, A., & Sonmez, T. (1999). House allocation with existing tenants. Journal of Economic Theory, 88(2), 233-260. https://doi.org/10.1006/jeth.1999.2553

Akter, S., Wamba, S. F., Gunasekaran, A., Dubey, R., & Childe, S. J. (2016). How to improve firm performance using big data analytics capability and business strategy alignment? International Journal of Production Economics, 182, 113-131. https://doi.org/10.1016/j.ijpe.2016.08.018

Atzei, N., Bartoletti, M., & Cimoli, T. (2017). A survey of attacks on Ethereum smart contracts (SoK). In Principles of Security and Trust (pp. 164-186). https://doi.org/10.1007/978-3-662-54455-6_8

Azevedo, E. M., & Leshno, J. D. (2016). A supply and demand framework for two-sided matching markets. Journal of Political Economy, 124(5), 1235-1268. https://doi.org/10.1086/687476

Bajari, P., Nekipelov, D., Ryan, S. P., & Yang, M. (2015). Machine learning methods for demand estimation. American Economic Review, 105(5), 481-485. https://doi.org/10.1257/aer.p20151021

Boudreau, K. J. (2010). Open platform strategies and innovation: Granting access vs. devolving control. Management Science, 56(10), 1849-1872. https://doi.org/10.1287/mnsc.1100.1215

Brynjolfsson, E., & McElheran, K. (2016). Data in action: Data-driven decision making in U.S. manufacturing. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.2722502

Budish, E. (2011). The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. Journal of Political Economy, 119(6), 1061-1103. https://doi.org/10.1086/664613

Casino, F., Dasaklis, T. K., & Patsakis, C. (2019). A systematic literature review of blockchain-based applications: Current status, classification and open issues. Telematics and Informatics, 36, 55-81. https://doi.org/10.1016/j.tele.2018.11.006

Catalini, C., & Gans, J. S. (2020). Some simple economics of the blockchain. Communications of the ACM, 63(7), 80- 90. https://doi.org/10.1145/3359552

Che, Y.-K., & Kojima, F. (2010). Asymptotic equivalence of probabilistic serial and random priority mechanisms. Econometrica, 78(5), 1625–1672. https://doi.org/10.3982/ECTA8354

Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. MIS Quarterly, 36(4), 1165-1188. https://doi.org/10.2307/41703503

Chen, Y., & Bellavitis, C. (2020). Blockchain disruption and decentralized finance: The rise of decentralized business models. Journal of Business Venturing Insights, 13, e00151. https://doi.org/10.1016/j.jbvi.2019.e00151

Choi, T.-M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. Production and Operations Management, 27(10), 1868-1883. https://doi.org/10.1111/poms.12838

Christidis, K., & Devetsikiotis, M. (2016). Blockchains and smart contracts for the Internet of Things. IEEE Access, 4, 2292-2303. https://doi.org/10.1109/ACCESS.2016.2566339

Cong, L. W., & He, Z. (2019). Blockchain disruption and smart contracts. Review of Financial Studies, 32(5), 1754- 1797. https://doi.org/10.1093/rfs/hhz007

Cong, L. W., & He, Z. (2019). Blockchain disruption and smart contracts. Review of Financial Studies, 32(5), 1754–1797. https://doi.org/10.1093/rfs/hhz007

Diakopoulos, N. (2016). Accountability in algorithmic decision making. Communications of the ACM, 59(2), 56-62. https://doi.org/10.1145/2844110

Dowling, M. (2022). Is non-fungible token pricing driven by cryptocurrencies? Finance Research Letters, 44, 102097. https://doi.org/10.1016/j.frl.2021.102097

Einav, L., & Levin, J. (2014). Economics in the age of big data. Science, 346(6210), 1243089. https://doi.org/10.1126/science.1243089

Evans, D. S. (2003). Some empirical aspects of multi-sided platform industries. Review of Network Economics, 2(3), 191-209. https://doi.org/10.2202/1446-9022.1026

Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People - An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds and Machines, 28(4), 689-707. https://doi.org/10.1007/s11023-018-9482-5

Fowler, A., & Pirker, J. (2021). Tokenfication - The potential of non-fungible tokens (NFT) for game development. In Extended Abstracts of the 2021 Annual Symposium on Computer-Human Interaction in Play (pp. 152-157). https://doi.org/10.1145/3450337.3483501

Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and analytics. International Journal of Information Management, 35(2), 137-144. https://doi.org/10.1016/j.ijinfomgt.2014.10.007

Gawer, A., & Cusumano, M. A. (2014). Industry platforms and ecosystem innovation. Journal of Product Innovation Management, 31(3), 417-433. https://doi.org/10.1111/jpim.12105

Gomber, P., Kauffman, R. J., Parker, C., & Weber, B. W. (2018). On the Fintech revolution: Interpreting the forces of innovation, disruption, and transformation in financial services. Journal of Management Information Systems, 35(1), 220-265. https://doi.org/10.1080/07421222.2018.1440766

Gonserkewitz, P., Karger, E., & Jagals, M. (2022). Non-fungible tokens: Use cases of NFTs and future research agenda. Risk Governance and Control: Financial Markets & Institutions, 12(3), 8-18. https://doi.org/10.22495/rgcv12i3p1

Hatfield, J. W., & Kominers, S. D. (2012). Matching in networks with bilateral contracts. American Economic Journal: Microeconomics, 4(1), 176-208. https://doi.org/10.1257/mic.4.1.176

Kitchin, R. (2014). Big Data, new epistemologies and paradigm shifts. Big Data & Society, 1(1), 1-12. https://doi.org/10.1177/2053951714528481

Kojima, F., & Pathak, P. A. (2009). Incentives and stability in large two-sided matching markets. American Economic Review, 99(3), 608-627. https://doi.org/10.1257/aer.99.3.608

Kshetri, N. (2018). Blockchain's roles in meeting key supply chain management objectives. International Journal of Information Management, 39, 80-89. https://doi.org/10.1016/j.ijinfomgt.2017.12.005

Lee, I., & Shin, Y. J. (2018). FinTech: Ecosystem, business models, investment decisions, and challenges. Business Horizons, 61(1), 35-46. https://doi.org/10.1016/j.bushor.2017.09.003

Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2020). A survey on the security of blockchain systems. Future Generation Computer Systems, 107, 841-853. https://doi.org/10.1016/j.future.2017.08.020

Luu, L., Chu, D.-H., Olickel, H., Saxena, P., & Hobor, A. (2016). Making smart contracts smarter. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (pp. 254-269). https://doi.org/10.1145/2976749.2978309

Mendling, J., Weber, I., van der Aalst, W., vom Brocke, J., Cabanillas, C., Daniel, F., Debois, S., Di Ciccio, C., Dumas, M., Dustdar, S., Gal, A., Garcia-Banuelos, L., Governatori, G., Hull, R., La Rosa, M., Leopold, H., Leymann, F., Recker, J., Reichert, M., Reijers, H. A., Rinderle-Ma, S., Solti, A., Rosemann, M., Schulte, S., Singh, M. P., Slaats, T., Staples, M., Weber, B., Weidlich, M., Weske, M., Xu, X., & Zhu, L. (2018). Blockchains for business process management - Challenges and opportunities. ACM Transactions on Management Information Systems, 9(1), 1-16. https://doi.org/10.1145/3183367

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. Big Data & Society, 3(2), 1-21. https://doi.org/10.1177/2053951716679679

Papai, S. (2000). Strategyproof assignment by hierarchical exchange. Econometrica, 68(6), 1403-1433. https://doi.org/10.1111/1468-0262.00166

Parker, G. G., & Van Alstyne, M. W. (2005). Two-sided network effects: A theory of information product design. Management Science, 51(10), 1494-1504. https://doi.org/10.1287/mnsc.1050.0400

Pathak, P. A., & Sonmez, T. (2008). Leveling the playing field: Sincere and sophisticated players in the Boston mechanism. American Economic Review, 98(4), 1636-1652. https://doi.org/10.1257/aer.98.4.1636

Provost, F., & Fawcett, T. (2013). Data science and its relationship to big data and data-driven decision making. Big Data, 1(1), 51-59. https://doi.org/10.1089/big.2013.1508

Pycia, M., & Unver, M. U. (2017). Incentive compatible allocation and exchange of discrete resources. Theoretical Economics, 12(1), 287-329. https://doi.org/10.3982/TE2201

Queiroz, M. M., & Wamba, S. F. (2019). Blockchain adoption challenges in supply chain: An empirical investigation of the main drivers in India and the USA. International Journal of Information Management, 46, 70-82. https://doi.org/10.1016/j.ijinfomgt.2018.11.021

Rahwan, I. (2018). Society-in-the-loop: Programming the algorithmic social contract. Ethics and Information Technology, 20(1), 5-14. https://doi.org/10.1007/s10676-017-9430-8

Risius, M., & Spohrer, K. (2017). A blockchain research framework. Business & Information Systems Engineering, 59(6), 385-409. https://doi.org/10.1007/s12599-017-0506-0

Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. Journal of the European Economic Association, 1(4), 990-1029. https://doi.org/10.1162/154247603322493212

Saberi, S., Kouhizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. International Journal of Production Research, 57(7), 2117-2135. https://doi.org/10.1080/00207543.2018.1533261

Schar, F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. Federal Reserve Bank of St. Louis Review, 103(2), 153-174. https://doi.org/10.20955/r.103.153-74

Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In Proceedings of the Conference on Fairness, Accountability, and Transparency (pp. 59- 68). https://doi.org/10.1145/3287560.3287598

Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. MIS Quarterly, 35(3), 553- 572. https://doi.org/10.25300/MISQ/2011/35.3.08

Sonmez, T. (1999). Strategy-proofness and essentially single-valued cores. Econometrica, 67(3), 677-689. https://doi.org/10.1111/1468-0262.00044

Treiblmaier, H. (2018). The impact of the blockchain on the supply chain: A theory-based research framework and a call for action. Supply Chain Management: An International Journal, 23(6), 545-559. https://doi.org/10.1108/SCM-03-2018-0143

Varian, H. R. (2014). Big data: New tricks for econometrics. Journal of Economic Perspectives, 28(2), 3-28. https://doi.org/10.1257/jep.28.2.3

Vidal-Tomas, D. (2022). The new crypto niche: NFTs, play-to-earn, and metaverse tokens. Finance Research Letters, 47, 102742. https://doi.org/10.1016/j.frl.2022.102742

Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. Journal of Business Research, 70, 356-365. https://doi.org/10.1016/j.jbusres.2016.08.009

Weber, I., Xu, X., Riveret, R., Governatori, G., Ponomarev, A., & Mendling, J. (2016). Untrusted business process monitoring and execution using blockchain. In Business Process Management (pp. 329-347). https://doi.org/10.1007/978-3-319-45348-4_19

Werner, S. M., Perez, D., Gudgeon, L., Klages-Mundt, A., Harz, D., & Knottenbelt, W. J. (2022). SoK: Decentralized finance (DeFi). In Proceedings of the 4th ACM Conference on Advances in Financial Technologies. https://doi.org/10.1145/3558535.3559780

Yeung, K. (2018). Algorithmic regulation: A critical interrogation. Regulation & Governance, 12(4), 505-523. https://doi.org/10.1111/rego.12158

Zarsky, T. (2016). The trouble with algorithmic decisions: An analytic road map to examine efficiency and fairness in automated and opaque decision making. Science, Technology, & Human Values, 41(1), 118-132. https://doi.org/10.1177/0162243915605575

Zetzsche, D. A., Arner, D. W., & Buckley, R. P. (2020). Decentralized finance. Journal of Financial Regulation, 6(2), 172-203. https://doi.org/10.1093/jfr/fjaa010

Zheng, Z., Xie, S., Dai, H.-N., Chen, W., Chen, X., Weng, J., & Imran, M. (2020). An overview on smart contracts: Challenges, advances and platforms. Future Generation Computer Systems, 105, 475-491. https://doi.org/10.1016/j.future.2019.12.019