

DOI: https://doi.org/10.48009/1_iis_2026_132

Real-world asset tokenization: Power asymmetries and AI in BlackRock's treasury bond tokenization

James J. Lee, *Seattle University*, jamesleeseattle@gmail.com

Benjamin Yeo, *Seattle University*, benjkyeo@gmail.com

Abstract

Real-world asset tokenization increasingly connects public blockchain settlement with traditional financial asset management. The March 2024 launch of BUIDL, a tokenized US Treasury fund on Ethereum, marks a convergence of blockchain infrastructure and institutional finance. BUIDL uses smart contracts for on-chain settlement but operates within an ecosystem in which the assets backing the token, along with portfolio decisions surrounding them, are managed through workflows that increasingly incorporate AI-assisted decision-support tools such as BlackRock's Aladdin platform. This creates a two-layer dependency: users gain blockchain transparency for token transfers while remaining exposed to upstream, opaque, potentially AI-assisted human decisions they cannot directly observe. This study analyzes this structure as a problem of power asymmetry in hybrid blockchain systems and introduces the concept of Compassion Debt: the accumulation of deferred accountability, interpretability, and human oversight as AI systems scale. It identifies three dimensions through which AI may amplify power asymmetry in tokenized real-world assets: ecosystem concentration, automation opacity, and deferred oversight. The paper contributes to information systems research by showing that tokenization combined with AI creates governance challenges that neither blockchain nor AI governance frameworks alone can address, and offers implications for firms, regulators, and users.

Keywords: Tokenization, Power Asymmetries, Artificial Intelligence, Blockchain Governance, Compassion Debt, Systemic Risk, Algorithmic Transparency

Introduction

Real-world asset (RWA) tokenization, the representation of traditional financial instruments such as bonds, equities, and funds as digital tokens on public blockchains, has emerged as a major trend at the intersection of traditional finance and decentralized infrastructure. Proponents celebrate it as validation of blockchain technology: round-the-clock global trading, fractional ownership, and programmable settlement, delivered on infrastructure originally designed to enable decentralization (Nakamoto, 2008). However, RWA tokenization reveals a paradox that information systems researchers must examine: the same blockchain infrastructure that enabled decentralization now enables large, centralized institutions to expand their reach while appearing innovative. Moreover, RWA tokenization surfaces a second, subtler dependency. While blockchain settlement operates transparently through smart contracts, the underlying traditional assets are managed off-chain through workflows that increasingly incorporate AI-assisted decision-support systems. Asset managers use such tools to assist human investment professionals with portfolio optimization, risk assessment, and market prediction through human-in-the-loop approaches. As AI-assisted portfolio

management becomes standard across asset management, tokenized real-world assets face potential ecosystem risks that conventional blockchain analysis misses. Recent regulatory and peer-reviewed work on tokenized finance similarly emphasizes that tokenization may improve settlement efficiency, reduce transaction frictions, and expand investor access, while also raising concerns about governance, operational fragility, interconnectedness, investor protection, and financial stability (Bank for International Settlements & Committee on Payments and Market Infrastructures, 2024; Financial Stability Board, 2024; Tanveer et al., 2025; The Board of the International Organization of Securities Commissions, 2025).

In March 2024, the USD Institutional Digital Liquidity Fund (BUIDL), launched by BlackRock on the Ethereum blockchain (*BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network*, 2024), exemplifies this pattern. BUIDL enables 24/7 global trading of tokenized US Treasury bonds; its blockchain layer automates settlement through deterministic smart contracts, while underlying Treasury holdings are managed through portfolio systems that users cannot audit or directly influence. BlackRock publicly describes Aladdin Copilot as an AI-assisted decision-support tool within its broader Aladdin platform (BlackRock, n.d.). This paper does not claim that Aladdin Copilot, or any specific AI tool, is used in the day-to-day management of BUIDL. Rather, BlackRock's public discussion of AI-assisted decision-support is used to illustrate the broader institutional context in which tokenized real-world asset systems may combine transparent blockchain settlement with less visible upstream portfolio, risk-management, and operational workflows. BUIDL is therefore used as a concrete prompt throughout this paper to examine the governance implications that emerge when tokenization combines with AI-assisted asset management.

The BUIDL tokenization provides access to qualified investors subject to onboarding and compliance requirements, albeit concentrating control narrowly. BlackRock serves as the primary manager of the fund, while Bank of New York Mellon acts as custodian (*BlackRock USD Institutional Digital Liquidity Fund Token | BUIDL Price*, n.d.), with no alternate provider available to investors. What distinguishes this system is the ecosystem vulnerability: while BUIDL's smart contracts execute token transfers transparently on-chain, underlying Treasury portfolios in tokenized funds may be managed through systems that can use AI to assist human professionals in portfolio decisions by analyzing market volatility, suggesting risk parameters, and recommending allocation strategies. Users benefit from improved access but may remain dependent on asset managers' human portfolio decisions informed by opaque AI recommendation systems operating upstream, thus creating ecosystem vulnerability even when blockchain settlement is transparent.

This paper examines how AI amplifies power asymmetries in blockchain systems. The paper introduces the concept of Compassion Debt in AI governance, referring to the accumulating gap between the scale at which AI systems act on users and the infrastructure users have to contest, audit, or understand those actions. Compassion debt compounds with scale, remains invisible to the parties best positioned to act on it, and becomes structural once users lose alternatives. This construct is developed in a subsequent section and is used throughout the analysis to organize the findings. The paper contributes to information systems research by building on prior work showing that blockchain-based systems introduce novel governance challenges (Beck et al., 2018), and by demonstrating how AI integration into tokenized systems amplifies these challenges in ways that neither blockchain nor AI governance frameworks alone can fully address. The paper adopts a conceptual case-study approach, analyzing BUIDL's publicly documented architecture to examine the governance implications of AI integration in tokenized real-world asset (RWA) systems more generally. The argument is theoretical and normative rather than empirical: the analysis does not measure AI use at any specific firm. References to BlackRock's upstream AI-assisted processes are treated as plausible illustrations of the broader phenomenon rather than claims about this specific fund, drawing on public BlackRock disclosures, industry documentation of Aladdin and similar platforms, and the IS governance literature.

To keep these registers distinct, the analysis separates publicly documented features of BUIDL from plausible governance risks. Publicly documented features include BUIDL's March 2024 launch on Ethereum, the role of Bank of New York Mellon as custodian and administrator, Securitize's role in the tokenization ecosystem, the fund's disclosed fee structure, and the publicly observable use of token-contract mechanisms for settlement. By contrast, the claim that AI tools materially shape BUIDL's specific upstream portfolio decisions is treated as a plausible governance risk rather than a verified fact. The argument does not depend on establishing that any particular AI tool is used in this fund. It depends on the broader architectural concern that tokenized real-world asset systems can combine transparent blockchain settlement with off-chain institutional decision processes that users cannot directly observe, audit, or contest.

Methodological Approach

This paper uses a conceptual case-study approach to examine the governance implications of AI-assisted decision-support in tokenized real-world asset systems. The analysis is conceptual and normative rather than empirical. It does not seek to measure AI use within BlackRock, verify the internal operation of Aladdin, or establish that specific AI tools are used in the day-to-day management of BUIDL. Instead, BUIDL is used as an illustrative case because its publicly documented structure combines public blockchain settlement with off-chain institutional asset management, custody, compliance, and risk-management processes. This architecture provides a useful setting for examining how power asymmetries may emerge when tokenized assets rely on upstream institutional decision-making that users cannot directly observe, audit, or contest.

The case analysis draws on publicly available documentation about BUIDL, BlackRock's public materials on Aladdin and AI-assisted decision-support, blockchain governance literature, and information systems research on AI governance, automation, opacity, and accountability. The analysis proceeds by identifying the major architectural layers of tokenized real-world asset systems, distinguishing between verified system features and plausible governance risks, and using these distinctions to develop the concept of Compassion Debt. This approach is appropriate for theory development because the objective is to clarify emerging governance mechanisms rather than to test causal hypotheses or evaluate firm-specific internal practices.

Power Asymmetries in System Design

In system design, a critical question is whether the system architecture constrains or enables the concentration of control. This is critical for information systems because architectures, like all technologies, are not neutral (Pacey, 1983). They embed political choices about who can act, ask permission, and bear risk. In blockchain systems, these choices become visible through the distribution of multiple forms of power, including governance power (to determine protocol changes), custody power (to control underlying assets), infrastructure power (to operate critical systems), trading power (to process transactions), and regulatory power (to enforce compliance).

Theoretical Foundations: Power Distribution in Distributed Systems

Distributed systems comprise architectural, logical, and/or political decentralization that refers to how many physical computers the system comprises, how many nodes exist, and how many entities control those nodes (Buterin, 2017). Bitcoin's mining network is largely decentralized but faces concentration risk because hash power is increasingly aggregated in a small number of mining pools (Ekblaw et al., 2016). Ethereum was designed as a decentralized system with distributed participation in consensus (Buterin,

2014). However, empirical analyses of the post-proof-of-stake network show that validator participation is increasingly concentrated among large staking providers such as Lido and Coinbase (Grandjean et al., 2024). The gap between logical and political decentralization creates vulnerability where systems appear distributed but concentrate control.

Traditional finance concentrates power through vertical integration. Banks control custody, execution, settlement, and governance within single institutions. In blockchain systems, decision-making and record validation are performed through decentralized consensus mechanisms rather than centralized control (Beck et al., 2018). Separate entities handle validation, custody, execution, and governance, with protocol rules constraining each. However, hybrid systems like BlackRock's tokenization reintroduce vertical integration while leveraging blockchain's infrastructure. BlackRock controls custody arrangements (with Treasuries held by BNY Mellon as the designated custodian), execution (via smart contracts informed by upstream AI-assisted portfolio decisions), governance (unilateral protocol changes), and compliance (account freezing authority). Only settlement runs on Ethereum, and even that is constrained by BlackRock's ability to pause trading or modify token logic.

Table 1: Power Distribution Across Blockchain Models

Power Dimension	Bitcoin	Ethereum/DAOs	BlackRock BUIDL
Governance	Community consensus (BIPs)	Token-weighted voting	Unilateral (BlackRock)
Custody	Self-custody (user wallets)	Multi-sig/smart contracts	Centralized (BNY Mellon)
Execution	Distributed miners	DEX protocols/validators	Smart contracts (AI upstream)
Infrastructure	Permissionless nodes	Validator sets (semi-open)	Proprietary (closed)
Regulatory Power	Code-based resistance	Partial regulatory capture	Full compliance integration

Five Dimensions of Power in Tokenized Systems

First, governance power determines who can change system rules. Decentralized blockchain systems may still involve hybrid or evolving structures, and token holders may participate in decision processes (Beck et al., 2018). While Bitcoin's governance model assumes consensus among developers, miners, and community, empirical work suggests this is more centralized in practice (Ekblaw et al., 2016). Ethereum protocol development involves open, community-driven processes in which researchers, developers, and other stakeholders evaluate and refine proposed changes. DAOs embed governance in smart contracts with vote-weighted decision-making. BlackRock's BUIDL concentrates governance entirely within the firm: there is no user voting, no community proposal process, and no external oversight mechanism. BlackRock unilaterally decides fee structures, operational parameters, and protocol upgrades.

Second, custody power determines who controls underlying assets. Decentralized custody distributes asset control through multi-signature wallets, threshold cryptography, or proof-of-reserve mechanisms. Users can verify that custodians hold claimed assets. Some protocols eliminate custodians entirely through self-custody or algorithmic escrow. BlackRock's BUIDL uses centralized custody: the underlying Treasury bonds are held by Bank of New York Mellon as custodian (*BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network*, 2024), and users cannot independently verify these holdings. No multi-signature protection exists. No alternative custodian options exist. If BlackRock's custody fails, through insolvency, fraud, or operational error, all BUIDL holders lose access to underlying assets simultaneously. Third, infrastructure power determines who operates critical systems. Decentralized infrastructure distributes operational control across independent node operators. Blockchain systems are designed as decentralized networks of nodes participating in consensus (Buterin, 2014). In this fashion, Ethereum validators run independently, and no single entity can halt the network. In Bitcoin, nodes independently participate in a globally distributed proof-of-work process to produce blocks (Nakamoto, 2008). Hence, no cartel can censor transactions permanently. BlackRock's tokenization centralizes infrastructure across designated partners: the token contract is operated through Securitize as transfer agent, interest distribution

is managed on-chain, on-chain price data is supplied by oracle providers selected within this ecosystem, and BlackRock runs the upstream execution infrastructure. Users depend entirely on BlackRock's infrastructure reliability.

Fourth, trading power determines who processes transactions and at what terms. In decentralized exchanges, execution happens through automated market makers (AMMs) or order books running on chain. Although participants can provide liquidity and execute trades, pricing follows algorithmic rules visible in smart contract code. BlackRock's on-chain token contract and upstream execution processes concentrate trading power, with BlackRock alone deciding order routing, execution timing, and price optimization. Users cannot choose alternate execution venues for primary market transactions. Execution occurs behind a single-institution process where users see final execution prices but not the logic that produced them.

Fifth, regulatory power determines who can enforce compliance and restrict access. Decentralized protocols reduce reliance on trusted intermediaries through peer-to-peer, permissionless architecture (Nakamoto, 2008) such that anyone can participate without approval. Censorship resistance emerges from distribution where no single entity can block transactions. BlackRock's BUIDL embeds regulatory control: AI compliance systems monitor transactions for suspicious patterns, freeze accounts flagged by automated risk models, and enforce Know Your Customer (KYC) requirements through centralized gatekeepers. Thus, regulatory power concentrates in BlackRock's compliance infrastructure.

Power Concentration Through AI Integration

AI systems enable the automation of control mechanisms that operate without human intervention (Yeung, 2018), which can amplify existing power asymmetries. Traditional centralized systems concentrated power but faced scalability constraints: human compliance officers can review limited transaction volumes; human traders can execute limited order flows; human governance requires deliberation time. AI removes these constraints. Compliance algorithms can monitor millions of transactions simultaneously. Trading algorithms can execute thousands of orders per second. Governance automation can implement rule changes instantly. The result is concentration without traditional limiting factors.

Moreover, AI can obscure underlying power structures through opacity (Liu, 2018), potentially masking emerging concentrations of power. Human decision-makers can explain their reasoning. Regulatory compliance officers can justify account freezes. Human traders can describe trading strategies. AI systems using deep learning or reinforcement learning produce choices that even their designers cannot fully explain. This creates accountability gaps where users cannot challenge decisions they cannot understand; regulators cannot audit algorithms they cannot interpret; and competitors cannot verify that AI systems treat participants fairly.

These amplification effects do not operate identically across the five power dimensions identified above. After BUIDL's architecture is specified, three specific mechanisms are identified, ecosystem concentration, automation opacity, and deferred oversight, through which AI simultaneously amplifies governance, custody, infrastructure, trading, and regulatory asymmetries. The five dimensions describe *where* power concentrates in tokenized systems; the three mechanisms describe *how* AI integration intensifies that concentration. These three mechanisms are later used to evaluate alternate governance models.

BlackRock's Treasury Tokenization: Architecture and Control

System Architecture and Operational Flow

BUIDL operates as a hybrid system bridging traditional finance and blockchain infrastructure. The architecture separates into two distinct systems: the blockchain settlement layer and a potentially AI-

assisted portfolio management layer. The blockchain layer includes tokenization (ERC-20 tokens on Ethereum), automated settlement (smart contracts for transfers and redemptions), and interest distribution (on-chain accrual). In tokenized real-world asset funds generally, an AI-assisted decision layer may exist upstream. This layer can include portfolio management systems (where AI tools recommend asset selection and rebalancing to human managers), risk assessment (where AI analyzes volatility and suggests parameter adjustments), and operational decisions (executed through traditional infrastructure based on AI-informed human choices). Users see only the blockchain layer; AI recommendation systems, if employed to inform human decisions, operate upstream and remain opaque to token holders.

The operational flow reveals where AI operates versus where smart contracts operate. Token transfers (buying, selling, transferring) happen on-chain through deterministic smart contracts - no AI involved. However, underlying asset management in tokenized funds happens off-chain and may involve human professionals using AI decision-support tools - systems that could recommend which assets to purchase, suggest allocation strategies, predict market movements, and propose risk adjustments, all operating outside blockchain visibility. Users interact with transparent blockchain settlement but may depend on human portfolio decisions informed by opaque AI recommendation systems affecting the traditional assets backing their tokens - a vulnerability inherent to tokenized real-world assets, not specific to any single fund.

AI Integration Points: Execution, Risk Management, and Custody

Execution. Smart contract execution on the blockchain layer is deterministic, but execution decisions made upstream, including order routing, timing, and venue selection for the underlying Treasury holdings, may involve AI tools that inform human traders. Such systems can recommend which assets to purchase, suggest allocation strategies, predict market movements, and propose risk adjustments. Final decisions remain with human professionals under supervision, but the speed differential between AI analysis and human review can pressure judgment toward accepting algorithmic suggestions without full scrutiny.

Risk management. AI tools may inform position sizing, stress testing, and parameter adjustment at the asset-management layer by analyzing portfolio volatility, identifying correlated exposures, and suggesting hedging strategies. Compliance and risk-management systems may also incorporate automated monitoring for regulatory requirements: pattern recognition for suspicious transactions, automated KYC checks, and dynamic adjustment of transaction thresholds. Both classes of system can affect token transfers or account access without producing visible explanations to those affected.

Custody. Custody operations are largely deterministic, where assets are held, ledgers are reconciled, statements are issued. However, AI-assisted compliance systems can affect custody outcomes by flagging accounts or transactions for review, potentially restricting access to assets even when underlying custody arrangements remain stable. Custody is therefore implicated indirectly: AI does not provide custody to assets. Instead, AI-assisted decisions may determine who can access them and when.

Across all three areas, the specific algorithms, thresholds, and decision processes typically remain proprietary to the asset manager and its technology partners. This creates information asymmetry: users see transparent on-chain settlement but lack visibility into potentially AI-informed decisions affecting the assets backing their tokens, raising accountability challenges when automated systems affect user access without clear explanation or appeal mechanisms.

Control Distribution Analysis

Analyzing control distribution across BUIDL's architecture reveals stark asymmetries. BlackRock controls: (1) tokenization logic, that is, the smart contract code that defines BUIDL token behavior, which uses an upgradeable proxy pattern allowing the contract owner to redirect all token logic to a new implementation unilaterally (the BUIDL token contract source code is publicly viewable at Etherscan, address

0x7712c34205737192402172409a8f7cce8aa2aec); (2) custody of all underlying Treasury bonds, held by Bank of New York Mellon as custodian (*BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network*, 2024) in centralized custody with no user verification or alternative custodian for the underlying assets; (3) fee structure: BlackRock sets management fees unilaterally, up to 0.50% annually (*BlackRock USD Institutional Digital Liquidity Fund Token | BUIDL Price*, n.d.) and adjustable without user consent; (4) protocol upgrades where the contract owner can modify token contracts, change interest distribution logic, or alter operational parameters through admin keys without governance votes; (5) operational decisions including pausing trading, modifying interest calculations, changing authorized participants, or restricting access. In addition, BlackRock's broader asset-management environment includes AI-assisted decision-support tools that may inform execution, risk-management, and portfolio workflows (BlackRock, n.d.). To the extent that such tools shape upstream decisions affecting tokenized assets, they remain proprietary systems that users cannot audit, configure, or opt out of while using BUIDL.

Users control: (1) buying or selling tokens on secondary markets that are subject to the asset manager's operational decisions about whether trading is enabled; (2) transferring tokens to other addresses that are subject to compliance requirements, constituting the entirety of user control. They lack all governance rights, custody choices, portfolio oversight, compliance appeals, and direct asset access. This distribution is a key site at which compassion debt accrues: every operational decision made without a user-facing accountability mechanism widens the gap between the system's reach and the infrastructure available for users to respond.

Ethereum validators control: (1) settlement finality for token transfers. Validators include BUIDL transactions in blocks and achieve consensus on state transitions; (2) theoretical censorship resistance. Validators can refuse to process BUIDL transactions, though this is economically irrational and technically difficult to coordinate. However, even this limited validator control is constrained by compliance requirements. If regulators demand transaction censorship, asset managers can implement it through the token contract without validator cooperation. The contract can maintain a blacklist, reject transfers involving flagged addresses, or halt all trading. Here, validator censorship resistance becomes irrelevant when the token contract enforces compliance at the application layer.

This distribution creates structural dependence. Users enjoy improved trading access but may remain dependent on upstream institutional processes, including potentially AI-assisted systems, that they cannot directly audit or challenge. If employed, the dependence may compound through AI-assisted decision systems in which users depend on institutional integrity and potentially opaque algorithms they cannot inspect or challenge. If portfolio management systems produce suboptimal results, users have no recourse. If risk management systems halt trading at critical moments, they cannot override the decision. If compliance systems freeze accounts incorrectly, they cannot appeal to human judgment. The system optimizes for efficiency and access but substantially constrains user autonomy.

Evidence from BUIDL's Operating History

BUIDL's operating history provides concrete examples of how the governance dynamics described above can manifest in practice. By mid-2025, BUIDL had become one of the dominant tokenized Treasury products, with substantial market share and deployment across multiple blockchains. These chain deployments, new share classes, and ecosystem integrations were determined by BlackRock and Securitize rather than by token-holder governance, illustrating the concentrated governance power described above.

BUIDL's redemption activity also illustrates how concentration can create correlated-behavior risks. In 2025, BUIDL experienced a substantial net outflow event that coincided with outflows from competing tokenized Treasury products. Although such redemptions do not by themselves demonstrate governance

failure, they show how large institutional flows in tokenized Treasury products can move together and create risks that users observe only after portfolio, liquidity, or redemption decisions have already been made upstream.

BUIDL's growing use as collateral and as infrastructure for other venues further extends these governance dependencies. When a tokenized fund becomes accepted as collateral or embedded in downstream protocols, users of those connected systems may inherit dependencies on upstream custody, compliance, redemption, and portfolio-management decisions that they cannot directly inspect. This makes the governance problem ecosystem-level rather than confined to one fund.

Ecosystem AI Vulnerability in Tokenized Assets

AI-assisted systems in tokenized real-world assets amplify power asymmetries through three dimensions: ecosystem concentration, automation opacity, and deferred oversight. Each dimension operates independently but they combine to create governance challenges that exceed what purely human systems can produce.

Ecosystem Concentration: Upstream Control and Data Network Effects

AI-assisted decision-making in asset management may concentrate upstream control through mechanisms unavailable to traditional human-only operations. This concentration operates through two reinforcing dynamics - economies of scale and data network effects. AI systems scale without the resource constraints that limit human organizations. Traditional compliance operations scale linearly - doubling the user base requires roughly doubling compliance staff. Traditional trading operations face similar constraints where more order flow requires more traders, more supervisors, more infrastructure. These scaling constraints naturally limit power concentration: organizations that grow too large become inefficient, thus creating opportunities for competitors.

AI-assisted systems remove these constraints. A compliance algorithm can monitor one million transactions as easily as one thousand, since the computational cost scales sublinearly with volume. Portfolio optimization systems can analyze unlimited asset combinations. The marginal cost per analysis approaches zero as complexity increases. This creates economies of scale that favor concentration. The largest institutions with the most sophisticated AI gain cost advantages that smaller competitors cannot match. Moreover, AI systems benefit from data network effects that compound concentration. The more transactions processed through the platform, the more data collected, the better the AI recommendation systems become, the more attractive the platform is to users, which generates more transactions and more data.

In tokenized Treasury funds, AI systems analyzing billions of dollars in historical trades can generate superior portfolio recommendations compared to smaller competitors with limited transaction histories. This creates a feedback loop. Market leaders with AI advantages compound those advantages through data accumulation. New entrants cannot compete because they lack the historical data needed to train effective AI systems. Users cannot switch to alternatives because competitors' AI systems may perform worse due to limited data. The concentration extends upstream from blockchain visibility. When tokenized funds face large redemptions, AI decision-support systems in asset management could analyze market signals in milliseconds and generate recommendations, potentially suggesting portfolio rebalancing, risk adjustments, or allocation changes that human professionals then review and execute.

The speed differential between AI analysis and human review creates pressure. While BUIDL's smart contracts execute token transfers transparently, AI recommendation systems, if used, can generate consequential portfolio suggestions affecting underlying assets faster than human review processes operate, creating pressure on human decision-makers to accept AI recommendations without full scrutiny. Human oversight cannot match these timescales. By the time human investment professionals fully review market conditions and AI recommendations, market opportunities may have passed. By the time a compliance officer investigates a suspicious transaction, AI systems have processed thousands more.

This speed differential creates accountability gaps. Traditional oversight mechanisms assume humans can review decisions before implementation. Traders submit orders to compliance review; custody decisions undergo legal analysis; risk management interventions receive executive approval. AI systems collapse these review cycles. Execution happens before oversight occurs. The result is post-hoc accountability at best: humans review what AI did rather than approving what AI should do. However, post-hoc review is ineffective for time-sensitive decisions. A suboptimal execution cannot be reversed. A frozen account cannot be unfrozen retroactively without harming the user. A halted trading period creates losses that review cannot undo.

In the same redemption example, AI systems can specifically recommend which underlying assets to sell, in what sequence, and through which channels. These are recommendations that human managers may follow under time pressure, affecting redemption pricing for token holders. These AI-generated recommendations and subsequent human decisions happen upstream, before blockchain settlement executes, thus creating an information asymmetry. Token holders see final redemption amounts on the blockchain but cannot audit AI recommendation systems and human decision processes operating upstream that influence those amounts. This is a fundamental vulnerability in tokenized real-world assets. No transparency requirement extends to Aladdin's decision logic.

Automation Opacity: The Black Box Problem in Portfolio Management

AI decision-support tools for risk assessment and portfolio optimization used by asset managers are typically proprietary systems not disclosed to token holders of tokenized funds, creating opacity at the portfolio management layer even when blockchain settlement is transparent. If employed, AI recommendation systems that inform human decisions on asset allocation, collateral requirements, and stress responses will operate as black boxes to token holders. Traditional rule-based systems use explicit logic: if condition X, then action Y. Humans can audit the rules, verify they are followed, and understand why specific decisions occurred. Machine learning systems learn implicit patterns from data. A deep neural network might use millions of parameters across hundreds of layers to produce a trading decision. The decision is mathematically deterministic. Given the same input, the network produces identical outputs. The computation, however, is humanly incomprehensible. No one, including the network's designers, can explain in simple terms why the algorithm chose a specific action.

This opacity creates accountability failures across the ecosystem. Token holders cannot verify that AI-assisted portfolio decisions act in their interests. Systems might optimize for institutional risk metrics rather than user yield. Regulators cannot audit proprietary AI recommendation systems for systemic risk. Such processes might concentrate risk or discriminate against user categories without detection, as bias emerges from learned patterns rather than explicit rules that auditors can review.

Opacity also enables plausible deniability. Asset managers can claim algorithms recommended actions based on market conditions, while maintaining deniability because even their engineers cannot fully explain specific AI recommendations. This creates an accountability vacuum: users cannot challenge decisions they

cannot understand; institutions cannot be held responsible for algorithmic behavior they cannot predict; regulators cannot enforce rules against systems operating through learned patterns.

The black box problem compounds over time. As AI systems learn from new data, their behavior drifts. An execution algorithm trained on 2023 market conditions may behave differently in 2024 after observing new patterns. Risk management AI might tighten restrictions after experiencing market stress. Compliance systems might flag new transaction patterns as suspicious based on recent fraud cases. These behavioral changes happen continuously, automatically, without documentation or explanation. Users interact with constantly evolving AI systems whose logic is always opaque and continuously shifting. This opacity is the algorithmic-layer manifestation of compassion debt: the accountability infrastructure required to interpret these systems is missing, and that gap compounds as the systems evolve.

Deferred Oversight: Human Decisions Operating Outside Blockchain Transparency

Human investment decisions informed by opaque AI tools operate outside blockchain transparency, creating a two-layer system where only one layer is auditable. In traditional asset management, oversight happens through direct human review. Traders submit orders to compliance review; portfolio decisions undergo committee approval; risk management interventions receive executive authorization. These review mechanisms assume transparency such that humans can examine the reasoning, challenge the logic, and approve or reject decisions before implementation.

Even with the speed differential described above, traditional oversight mechanisms shift from ex-ante approval to ex-post review. Humans end up reviewing what occurred rather than approving what should occur. This matters structurally because it creates a two-layer system. The blockchain layer provides complete transparency for token transfers: every transaction recorded, timestamped, and publicly auditable. But the portfolio management layer operates upstream in traditional finance systems, potentially using proprietary AI recommendation tools. Token holders see blockchain outcomes (redemption amounts, interest distributions) but cannot audit the AI-assisted decision processes that produced those outcomes.

Oversight is deferred to post-hoc review, while the most consequential decisions, including which assets to buy, when to rebalance, how to manage risk, happen outside blockchain transparency entirely. The combination of ecosystem concentration, automation opacity, and deferred oversight creates power concentration mechanisms unavailable to traditional systems. Each dimension amplifies the others: concentration provides data advantages that improve AI performance; opacity prevents accountability mechanisms from constraining concentrated systems; deferred oversight enables both concentration and opacity to operate outside blockchain transparency. The result is concentration that appears natural and efficient but reflects architectural choices favoring incumbents. The deferred oversight pattern is what makes compassion debt structural rather than transient: the longer ex-post review remains the only available mechanism, the more the gap between automated action and human accountability becomes load bearing for the system's operation.

Systemic Risks from Combined Dimensions

The amplification of asymmetries through AI creates systemic risks that transcend individual user harms. The first is correlated failures. If upstream AI-assisted systems used by a dominant asset manager were to malfunction, through software bugs, adversarial attacks, or unexpected market conditions, users of the associated tokenized products could experience simultaneous failures. There is neither diversity of trading strategies, redundancy in risk management, nor alternate custody systems. A single AI failure point can affect billions of dollars in assets across all token holders simultaneously.

The second is algorithmic collusion. As more institutions deploy similar AI systems for tokenization, their models may converge on identical strategies. Execution models trained on the same market data might all

attempt to sell Treasury bonds simultaneously during stress. Risk management systems might all halt trading at the same volatility threshold. Custody systems might all freeze accounts based on similar pattern-matching rules. This convergence creates fragility: what appears to be independent institutional decisions is actually correlated algorithmic behavior, producing system-wide instability that no single institution intended.

The third is regulatory arbitrage through AI. Institutions can use algorithmic complexity to evade regulatory intent while maintaining technical compliance. Compliance systems might systematically disadvantage certain user categories through patterns that violate anti-discrimination principles but don't violate explicit rules. Portfolio management systems might engage in timing strategies that technically comply with requirements but disadvantage certain users through implementation choices. Regulators cannot challenge these behaviors because the algorithms' opacity prevents proving intent or demonstrating systematic patterns.

Comparative Analysis: AI Governance Models in Tokenization

To understand the implications of centralized approaches to AI-integrated tokenization, building on foundational work on blockchain governance (Beck et al., 2018), this section compares alternate governance models that can shape this emerging space. These alternatives exist in theory and, in some cases, in practice through experimental protocols and decentralized finance (DeFi) projects. The comparison reveals a spectrum of approaches - with different tradeoffs between efficiency, transparency, and user authority. Each model is evaluated against the three AI amplification mechanisms identified earlier, specifically, how each manages ecosystem concentration, automation opacity, and deferred oversight, before considering efficiency tradeoffs and the resulting compassion debt.

The centralized AI-assisted model, exemplified by traditional asset managers entering tokenization, concentrates algorithmic decision-support within the institution. AI systems provide recommendations to human professionals, but the systems remain proprietary, interpretability is limited to internal audits, and users have no oversight rights. This model may maximize execution speed and operational efficiency but creates information asymmetry. Users must trust the institution's AI-assisted decision processes without verification.

The federated AI model distributes algorithmic decision-making across multiple institutions or node operators. Examples include decentralized oracle networks (Chainlink) and distributed execution protocols. This model improves resilience and reduces single points of failure but introduces coordination complexity and potential for manipulation if node operators collude.

The open algorithm model makes AI logic publicly auditable. Protocols publish their algorithms, training data, and model parameters, allowing independent verification. Users can inspect how decisions are made, propose improvements, and fork the system if they disagree with governance choices. This model maximizes transparency but may expose proprietary strategies or create competitive disadvantages.

The user-controlled AI model gives users direct oversight over algorithmic determinations affecting them. Users can set preferences, opt out of certain AI behaviors, or require human review. This model maximizes user autonomy but may reduce system efficiency and increase complexity.

Table 2 compares these governance models across key dimensions relevant to power asymmetries and compassion debt. The table shows clear tradeoffs: models that maximize efficiency and centralized

control tend to minimize transparency and user authority, while models that maximize transparency and user control tend to sacrifice efficiency.

The comparative analysis reveals that centralized approaches prioritize efficiency and control over transparency and user authority. This choice may be rational from an institutional perspective. Centralized AI-assisted portfolio management can maximize profit margins and operational speed. However, from a societal perspective, it creates potential systemic risks. As more financial infrastructure migrates to AI-integrated tokenization, concentration in a few institutions with opaque AI-assisted decision processes can create fragility, potential for abuse, and limited recourse for users.

Table 2: Comparative Framework - AI Governance Models in Tokenization

Governance Model	Ecosystem Concentration	Automation Opacity	Deferred Oversight	Efficiency	Compassion Debt
Centralized AI-Assisted (Traditional Asset Managers)	Single institution, AI supports human decisions	High (proprietary, internal only)	Deferred/Post-hoc only	Highest	Highest (complete opacity)
Federated AI	Distributed nodes	Moderate (aggregate outputs visible)	Partial (node-level)	Moderate	Moderate (distributed opacity)
Open Algorithm	Publicly auditable	Low (code and data public)	Full ex-ante (verifiable)	Lower (coordination costs)	Low (accountability visible)
User-Controlled AI	Individual users	Minimal (user-defined)	User-controlled (opt-in/opt-out)	Lowest (fragmentation)	Minimal (users accountable)

There are alternate models that can balance efficiency with accountability. Information systems researchers should examine whether hybrid models that combine elements from different approaches can achieve better outcomes than purely centralized approaches, and what organizational and technical barriers prevent their adoption in traditional finance entering blockchain spaces.

Compassion Debt in AI Governance

Compassion Debt in AI governance is defined as the accumulating gap between the scale at which automated systems act on users and the infrastructure available for those users to contest, audit, or understand those actions. Compassion debt accrues whenever an AI system makes or meaningfully shapes a consequential decision affecting a user without a matching mechanism for explanation, appeal, or human override. Every automated decision made without explanation, every account frozen without a path to appeal, every algorithm deployed without interpretability adds to the debt.

The construct is modeled on the software engineering notion of technical debt (Cunningham, 1992), which is the observation that decisions to defer structural work produce short-run velocity gains but compound into long-run costs. Compassion debt operates through an analogous mechanism: institutions gain operational efficiency by deploying AI systems without interpretability, appeal, or override infrastructure, and the cost of building that infrastructure grows as the systems scale, as retrofitting becomes harder, and as users lose alternatives. Unlike technical debt, however, compassion debt is paid primarily by users rather than by the institution that incurred it. The externality is the core problem.

BUIDL, viewed as one node in a broader tokenized-RWA ecosystem, accrues compassion debt at three levels. At the algorithmic level, the AI tools that plausibly inform upstream portfolio decisions produce recommendations whose internal logic is not accessible to human reviewers, let alone to token holders. At the operational level, automated compliance systems can freeze accounts or reject transfers with limited human oversight, leaving users to experience blocks with no clear appeal path. At the governance level, BlackRock retains substantial authority to modify AI systems, change parameters, or deploy new algorithms without user consent or notification.

Three properties make compassion debt particularly difficult to reverse. First, it compounds; as the user base grows, the cost of retrofitting accountability mechanisms (interpretability, appeal processes, override capability) grows faster than linearly. Adding interpretability to established models requires retraining, validation, and potential performance trade-offs. Creating appeal processes at scale requires staffing and infrastructure the institution has no competitive incentive to build. Second, it is invisible to the parties best positioned to act on it. Users often do not know what oversight exists until they need it, and regulators often cannot see the debt accumulating because the systems are proprietary. Third, it becomes structural. Once users have no alternatives, the institutional cost of leaving the debt unpaid drops to zero.

Compassion debt is not a claim about intent. An institution can deploy AI-assisted systems in good faith, with no interest in exploiting users, and still accrue compassion debt as an emergent property of scale. The construct directs analytical attention to the gap itself, between the capabilities of automated systems and the infrastructure users have to respond to them, rather than to the motives of those who deployed the systems. For information systems research, the construct provides a way to measure and compare governance maturity across AI-integrated platforms, and to ask which architectural choices make the debt easier or harder to pay down.

Potential Objections

A reasonable defender of BUIDL's architecture may argue that tokenization meaningfully expands institutional access to Treasury markets, that BlackRock's AI-assisted decision tools may produce superior risk-adjusted returns through superior data integration, that BNY Mellon's role as custodian offers regulatory protections that decentralized custody currently lacks, and that access restrictions and compliance requirements reflect broader securities law constraints rather than discretionary architectural choices. These benefits are real, and the argument is not that tokenization is undesirable. The claim is narrower: even granting these benefits, the asymmetric governance structure created by combining centralized AI-assisted asset management with public-blockchain settlement deserves analytical attention. Compassion debt is a property of the architecture, not a verdict on the institution that built it.

Organizational and Policy Implications

For Organizations Considering AI-Integrated Tokenization

Firms tokenizing traditional assets should explicitly decide whether to integrate AI decision-making and, if so, what accountability mechanisms to embed. BlackRock's choice prioritizes efficiency and profit optimization. An alternative would be to require interpretability, such that algorithms must explain their decisions; users can audit execution; appeals processes exist for contested decisions. Organizations must be transparent about this choice and account for its power implications.

Concretely, this means publishing governance disclosures that identify what AI-assisted tools are deployed, what decisions they inform, what data sources they rely on, and what override mechanisms exist. Existing frameworks offer useful starting points. The NIST AI Risk Management Framework provides a framework

for documenting AI system purpose, context, risk, validity, reliability, and monitoring (National Institute of Standards and Technology, 2023). Algorithmic impact assessments can complement this approach by documenting how automated or AI-assisted systems may affect users, counterparties, and downstream markets. Regulatory guidance on tokenized financial assets, including IOSCO's recommendations, also points toward applying market-integrity and investor-protection principles to tokenized finance (The Board of the International Organization of Securities Commissions, 2025).

These strategies map directly onto the three mechanisms identified in this paper. Interpretability requirements, documentation, and algorithmic impact assessments address automation opacity. Audit rights, escalation procedures, and human override mechanisms address deferred oversight. Competitive custody arrangements, multi-provider infrastructure, and stronger disclosure around service-provider dependencies address ecosystem concentration. These measures do not eliminate the efficiency benefits of tokenization, but they can reduce compassion debt by making accountability infrastructure part of the system design rather than an after-the-fact correction.

For Regulators and Policymakers

Regulators should require impact assessments on algorithmic governance in tokenized systems. As AI-driven custody and execution become systemic, regulators must ensure: (1) interpretability requirements (algorithms must explain decisions); (2) audit rights (users/regulators can inspect algorithms); (3) override capabilities (humans can intervene); (4) appeal processes for algorithmic determinations. Without these safeguards, AI-driven centralization becomes invisible concentration of power.

Existing frameworks already provide partial coverage. SEC oversight of investment companies governs the underlying funds, federal banking guidance on model risk management addresses some algorithmic concerns, and securities laws govern investor disclosure. However, gaps remain: most existing frameworks were designed for human decision-making at organizational scale, not algorithmic decision-making at population scale. Jurisdictional fragmentation across crypto-asset rules in different markets further complicates consistent oversight, particularly for cross-chain tokenized funds whose operations span regulatory regimes.

For Users and Investors

Users should understand that improved access does not imply power. Using BUIDL provides access to Treasury trading, but upstream portfolio decisions, potentially informed by AI recommendation tools, happen without user visibility or recourse. Users should demand transparency: What systems execute their trades? How are custody decisions made? What oversight exists? Until these questions are answered, users accept structural dependence on BlackRock's upstream institutional processes, including potentially AI-assisted systems that may not be directly visible to token holders.

Practical due-diligence questions investors can ask include: What AI systems are used in this fund? What audit rights exist for investors? How are disputes resolved when algorithmic decisions appear to disadvantage users? Third-party audits and transparency reports can shift institutional incentives by making opacity a competitive disadvantage rather than a default. The leverage to extract such information differs sharply between institutional and retail investors. Collective action through institutional channels (consortia, fiduciary intermediaries, asset-allocator due diligence) is largely available only to the former, which is itself a structural feature of compassion debt in this context.

Conclusion

The BUIDL case reveals how blockchain can be leveraged to expand institutional reach while avoiding governance constraints that decentralized systems impose. The system improves user access dramatically but concentrates operational control predominantly in BlackRock. More critically, AI-assisted decision-making amplifies this concentration while obscuring it through automation. Users enjoy improved trading access but may remain dependent on upstream institutional processes, including potentially AI-assisted systems, that they cannot directly audit or challenge.

This paper introduces the concept of Compassion Debt in AI governance. This is the accumulating gap between the reach of AI systems and the infrastructure users have to respond to them. Information systems researchers should examine tokenization beyond a monolithic phenomenon to cover a spectrum of governance choices. Some systems distribute both access and control, while others distribute access while concentrating control (Chen et al., 2021).

Future research should examine whether alternative tokenization models, such as decentralized custody, community governance, competitive custody options, mandatory interpretability, can achieve accessibility without concentration. As traditional finance migrates to blockchain and AI becomes embedded in settlement infrastructure, these questions become urgent. The convergence of blockchain, institutional control, and artificial intelligence defines the next phase of financial infrastructure. Understanding its power consequences is essential for building equitable systems.

References

- Bank for International Settlements, & Committee on Payments and Market Infrastructures. (2024). Tokenisation in the context of money and other assets: Concepts and implications for central banks [Report to the G20]. Bank for International Settlements. <https://www.bis.org/cpmi/publ/d225.htm>
- Beck, R., Müller-Bloch, C., & King, J. L. (2018). Governance in the blockchain economy: A framework and research agenda. *Journal of the Association for Information Systems*, 19(10). <https://aisel.aisnet.org/jais/vol19/iss10/1>
- BlackRock. (n.d.). Introducing Generative AI by Aladdin® | Aladdin Copilot. BlackRock. Retrieved April 10, 2026, from <https://www.blackrock.com/aladdin/solutions/aladdin-copilot>
- BlackRock Launches Its First Tokenized Fund, BUIDL, on the Ethereum Network. (2024, March 20). <https://www.businesswire.com/news/home/20240320771318/en/BlackRock-Launches-Its-First-Tokenized-Fund-BUIDL-on-the-Ethereum-Network>
- BlackRock USD Institutional Digital Liquidity Fund Token | BUIDL Price. (n.d.). Retrieved April 29, 2026, from <https://stomarket.com/sto/blackrock-usd-institutional-digital-liquidity-fund-buidl>
- Buterin, V. (2014). A next-generation smart contract and decentralized application platform. White Paper, 3(37), 2–1.
- Buterin, V. (2017, February 6). The meaning of decentralization. Medium. <https://medium.com/@VitalikButerin/the-meaning-of-decentralization-a0c92b76a274>

- Chen, Y., Richter, J. I., & Patel, P. C. (2021). Decentralized governance of digital platforms. *Journal of Management*, 47(5), 1305–1337.
- Cunningham, W. (1992). The WyCash portfolio management system. *ACM Sigplan Oops Messenger*, 4(2), 29–30.
- Ekblaw, A., Barabas, C., Harvey-Buschel, J., & Lippman, A. (2016). Bitcoin and the myth of decentralization: Socio-technical proposals for restoring network integrity. 18–23.
- Financial Stability Board. (2024). The financial stability implications of tokenisation. Financial Stability Board. <https://www.fsb.org/2024/10/the-financial-stability-implications-of-tokenisation/>
- Grandjean, D., Heimbach, L., & Wattenhofer, R. (2024). Ethereum proof-of-stake consensus layer: Participation and decentralization. 253–280.
- Liu, H.-Y. (2018). The power structure of artificial intelligence. *Law, Innovation and Technology*, 10(2), 197–229.
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system.
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (pp. 100–101). National Institute of Standards and Technology. <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- Pacey, A. (1983). *The Culture of Technology*. MIT Press.
- Tanveer, U., Ishaq, S., & Hoang, T. G. (2025). Tokenized assets in a decentralized economy: Balancing efficiency, value, and risks. *International Journal of Production Economics*, 282, 109554. <https://doi.org/10.1016/j.ijpe.2025.109554>
- The Board of the International Organization of Securities Commissions. (2025). *Tokenization of financial assets*. The Board of the International Organization of Securities Commissions. <https://www.iosco.org/library/pubdocs/pdf/IOSCOPD809.pdf>
- Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*, 12(4), 505–523.