

VOLUME 1 (2) - 2026

ISSN: 3021-5811

The Journal of FinTech and Digital Assets.

AN INTERNATIONAL JOURNAL.

JFDA



THE JOURNAL OF FINTECH AND DIGITAL ASSETS.

An International Journal.

The Journal of FinTech and Digital Assets (JFDA) is an international, peer-reviewed scholarly journal dedicated to advancing knowledge at the intersection of financial technology, blockchain innovation, computer science, entrepreneurial finance, management, alternative data and digital asset markets.

Editor-in-Chief

Cuong Nguyen, *Lincoln University*, New Zealand.

Monica Rossolini, *Università degli Studi di Milano-Bicocca*, Italy.

Nuttanan Wichitaksorn, *Auckland University of Technology*, New Zealand.

Kazuo Yamada, *Kyoto University*, Japan.

Associate Editor for this issue:

Van Ha Nguyen, *Foreign Trade University*, Vietnam.

JFDA serves as a global hub for advanced research shaping the digital finance revolution. We welcome manuscripts that push conceptual boundaries, offer empirical insight, and provide solutions to real-world problems.

For further information on the types of submission, please go to [JFDA/Submissions](#)

ISSN 3021-5811

Available online: September 2026

Published for Volume 1(2) - November 2026

© Edition: Lincoln University.

© Text and illustrations: individual contributors, unless otherwise noted.



CONTENTS

Inaugural Editorial

Editors-in-Chief 1

Research Papers

Bridging the Financial Inclusion Gap:
FinTech Innovation and Digital Financial
Services in Rural India and New Zealand.
Syed Matiur Rahman 2–28

Comparison of digital RMB and private
payment platform: Gap Analysis and
Improvement.
Gesheng Chen 29–44

Financial Fraud Detection: Machine
Learning vs. Traditional Rule-Based
Systems - A Comparative Analysis of
Model.
Haixia Du 45–53

Two-Regime Liquidity Recovery After a
Perpetual Futures Liquidation Cascade:
Evidence from Hyper-liquid and the
October-10th-2025 Event.
Lim Boon Chuan 54–68

Report

From Agent Identity to Agent Economy:
Measuring the Operational Readiness of
ERC-8004 AI Agents.
*Rischan Mafrur, and Priagung
Khusumanegara* 69–88

Special Thanks: Yanan Zhao, Tuan Anh Le and Manh Ha
Nguyen for their excellent support of the JFDA
production team.

Foreword:

We stand at a pivotal juncture in the evolution of global finance. The rapid intersection of decentralized infrastructure, artificial intelligence, dynamic machine learning, and central bank innovations is fundamentally altering how value is stored, exchanged, and secured. Yet, as the collection of insightful papers in this issue demonstrates, every advance in efficiency and access brings forth fresh systemic vulnerabilities, governance dilemmas, and behavioral friction points.

This issue brings together a rich tapestry of empirical research that dissects both the promise and the structural friction of our modern financial architecture. Across disparate geographies and technologies, a unified theme emerges: the presence of digital financial infrastructure does not automatically yield economic participation, market resilience, or flawless security.

Consider the dual realities of financial inclusion. Rahman's comparative analysis of digital financial services in rural India and New Zealand highlights the persistent structural and cultural gaps that leave millions unbanked or underserved, asserting that technology must be paired with hybrid, human-centric delivery models to drive real economic access. Similarly, Chen's examination of China's digital RMB reveals how even state-backed technological triumphs can struggle against established consumer habits, proving that utility, trust, and user habituation often outweigh technical availability when competing with legacy private platforms.

Concurrently, the mechanisms safeguarding and powering these digital networks face rigorous scrutiny. Du demonstrates the undeniable superiority of adaptive machine learning algorithms, specifically Random Forests, over rigid, rule-based systems in combating increasingly complex financial fraud. Yet, as transactions become faster and automated, venue-level liquidity remains highly sensitive to systemic shocks. Lim's granular investigation into the October 2025 liquidation cascade on Hyperliquid uncovers a stark "two-regime recovery". While top-level bid-ask spreads may quickly normalize after a crash, underlying market depth contracts persistently, that is, superficial liquidity metrics often mask underlying venue fragile state.

Finally, looking toward the emerging horizon of autonomous systems, Mafrur and Khusumanegara evaluate the ERC-8004 standard for AI agents. Their empirical findings reveal an ecosystem that is currently "registration-heavy but operationally shallow," underscoring the gap between creating agent identity on the blockchain and achieving a fully functional, decentralized agent economy.

Together, these authors offer a sobering yet optimistic roadmap. They remind us that building the future of finance requires equal parts innovative engineering, empirical vigilance, and thoughtful regulatory stewardship. It is our hope that this collection inspires researchers, policymakers, and practitioners alike to look beyond surface-level metrics and design ecosystems that are truly inclusive, resilient, and secure.

With warmest thanks,
Cuong Nguyen
On behalf of the JFDA Editorial Team

Bridging the Financial Inclusion Gap: FinTech Innovation and Digital Financial Services in Rural India and New Zealand

Syed Matiur Rahman*

Author

Syed Matiur Rahman, Auckland Institute of Studies, Auckland, New Zealand.

Email: smr1975nz@gmail.com

* Corresponding author.

Abstract

Financial exclusion remains a critical barrier to economic development, affecting approximately 190 million unbanked adults in rural India and significant populations in developed economies like New Zealand. This research examines how FinTech innovation and digital financial services address financial inclusion challenges across diverse geographical contexts and explores how digital banking can recalibrate sectoral lending inequalities in developed economies. Through comparative analysis of rural India and New Zealand, this study investigates the role of mobile banking, blockchain technology, peer-to-peer lending, digital payment systems, and emerging models such as Business Correspondents (BC), Village Level Entrepreneurs (VLEs), and Banking-as-a-Service (BaaS) in expanding financial access. The research utilises mixed-methods analysis, incorporating quantitative data from the Global Findex Database and qualitative assessments of policy frameworks including India's Jan Dhan Yojana, Aadhaar, and Unified Payments Interface, alongside New Zealand's digital financial inclusion initiatives. Findings reveal that while FinTech adoption demonstrates promise in urban centres, rural areas face persistent challenges, including limited internet connectivity, digital literacy gaps, socio-cultural barriers, and structural inequalities in credit allocation. The study concludes that the future of banking in rural regions hinges on a hybrid model combining FinTech agility with traditional banking stability, supported by proactive regulation, digital infrastructure investment, and culturally responsive design. This research contributes evidence-based recommendations for policymakers, financial institutions, and FinTech entrepreneurs seeking to design inclusive digital financial ecosystems and rebalance sectoral credit toward productive economic sectors.

Keywords: FinTech innovation, Digital financial services, Financial inclusion, Rural banking, Mobile payments, Agricultural finance

JEL Classification: G21, O16, O33

Article Information

Received: 15 February 2026

Revised: 02 April 2026

Accepted: 01 May 2026

Published: June 2026

Citation:

Rahman, S. M. (2026). Bridging the financial inclusion gap: FinTech innovation and digital financial services in rural India and New Zealand. *The Journal of FinTech and Digital Assets*, 1(2), 2–28.

1. Introduction

Financial inclusion represents a cornerstone of sustainable economic development and poverty alleviation globally (Demirgüç-Kunt et al., 2022). The World Bank's Global Findex Database (2021) estimates that approximately 1.4 billion adults worldwide remain unbanked, with 190 million residing in India alone. This exclusion perpetuates cycles of poverty by limiting access to savings mechanisms, credit facilities,

insurance products, and payment systems essential for economic participation (Ozili, 2023). Traditional brick-and-mortar banking infrastructure proves economically unviable in sparsely populated rural areas, creating persistent geographical disparities in financial access (Patwardhan, 2024).

The emergence of Financial Technology (FinTech) has fundamentally transformed the financial services landscape (Arner et al., 2020). Digital technologies, including mobile banking applications, blockchain-based systems, peer-to-peer lending platforms, and digital payment infrastructure, enable cost-effective delivery of financial services to previously underserved populations (Gomber et al., 2018). The proliferation of smartphones, declining data costs, and supportive regulatory frameworks have created favourable conditions for FinTech adoption across economies (Mention, 2019).

India has emerged as a global leader in digital financial innovation, driven by transformative government initiatives, including the Pradhan Mantri Jan Dhan Yojana, the Aadhaar biometric identification system, and the Unified Payments Interface (Reserve Bank of India, 2023). These interventions have facilitated the opening of over 460 million bank accounts. New Zealand presents a complementary case study, with approximately 7%–8% of adults remaining unbanked, a rate that disproportionately affects Māori and Pacific Island communities (Financial Services Council New Zealand, 2022). This comparative analysis illuminates transferable insights for digital financial inclusion across diverse contexts and examines how digital innovations can address sectoral lending imbalances in developed economies like New Zealand, particularly their impact on employment in small state economies.

This paper contends that FinTech's impact on financial inclusion, especially in rural regions, is not inherently stabilising or destabilising but rather depends on appropriate governance, adaptive market structures, and rapid regulatory responses. FinTech has significant potential to enhance economic inclusion and resilience in rural India by tackling obstacles such as poor infrastructure, limited access, and low financial literacy. However, unmanaged risks, including cybersecurity vulnerabilities, digital literacy disparities, and regulatory compliance issues in cash-intensive rural economies, could undermine stability. This research examines FinTech's role in transforming rural banking, its accelerated relevance post-COVID-19, the critical balance between innovation and control, and its potential to rebalance structural credit allocation inequalities and their employment implications.

This study makes several original contributions to the FinTech and financial inclusion literature. It is among the first to conduct a systematic comparative analysis of FinTech-enabled financial inclusion across a developing economy (rural India) and a developed small-state economy (New Zealand) within a unified analytical framework. It extends existing literature by explicitly linking sectoral credit allocation imbalances to employment outcomes, a dimension largely absent from prior FinTech inclusion studies. The dual-context design generates cross-applicable policy insights of direct relevance to regulators, financial institutions, and FinTech entrepreneurs operating in both Global South and developed economy settings.

The remainder of this paper is structured as follows. Section 2 reviews the relevant literature across financial inclusion theory, FinTech modalities, sectoral credit allocation, and the theoretical framework. Section 3 outlines the data sources and methodology. Section 4 presents empirical results and discussion, covering the current state of financial exclusion, FinTech adoption patterns, case studies from both India and New Zealand, regulatory frameworks, sectoral credit inequalities, a merged analysis of FinTech-driven credit recalibration and employment implications, sustainability dimensions, and economic impact. Section 5

discusses cross-contextual implications and policy recommendations. Section 6 concludes with key findings, policy implications, limitations, and directions for future research.

2. Literature Review

2.1 Conceptualising Financial Inclusion

Financial inclusion encompasses the accessibility and usage of formal financial services by all segments of society, particularly marginalised populations (Sarma & Pais, 2011). The concept extends beyond mere account ownership to include the quality, affordability, and appropriateness of financial products (Demirgüç-Kunt & Klapper, 2013). Barriers operate at multiple levels, including supply-side constraints (e.g., lack of infrastructure) and demand-side limitations (e.g., low literacy, distrust) (Park & Mercado, 2018). Rural contexts amplify these challenges through geographical isolation and inadequate infrastructure (Rao & Gupta, 2022).

2.2 FinTech and Financial Services

FinTech represents the convergence of finance and technology, fundamentally disrupting traditional service delivery through digital innovation (Arner et al., 2020). Key modalities include mobile banking, blockchain technologies, peer-to-peer lending, digital payments, and AI-driven credit scoring (Gomber et al., 2018; Chen et al., 2019). Mobile money services have demonstrated efficacy in developing economies, with Kenya's M-Pesa serving as the archetypal success story (Suri & Jack, 2016). The leapfrogging phenomenon characterises the proliferation of mobile financial services in regions lacking traditional banking infrastructure (Lee & Shin, 2018).

In rural India, FinTech innovations such as micro-ATMs, cardless transactions, the Business Correspondent (BC) model, and Village Level Entrepreneurship (VLE) models have significantly increased access to financial services (Uzma & Pratihari, 2019). These solutions often offer lower operating costs and lower fees than conventional banking processes (McKinsey, 2021). AI and data analytics are extensively used to understand rural consumers' needs and behaviours, facilitating the customisation of products and services (Kumar, 2023).

2.3 Sectoral Credit Allocation and Structural Inequalities

Research on sectoral credit allocation in Table 1 reveals persistent structural inequalities in both developing and developed banking systems (Beck et al., 2010). In many economies, including New Zealand, lending disproportionately favours real estate over productive sectors like agriculture and manufacturing (RBNZ, 2023). This misallocation stems from biases in risk assessment, collateral requirements, and information asymmetries, which particularly disadvantage agricultural enterprises (Cole et al., 2011). Digital financial innovations offer potential pathways to recalibrate these imbalances through alternative data, specialised platforms, and embedded finance solutions (BIS, 2022).

Table 1. Sectoral Credit Allocation: India vs. New Zealand Comparison

Sector	India: Credit %	India: GDP %	India: Gap	NZ: Credit %	NZ: GDP %	NZ: Gap
Real Estate	45%	7.3%	+37.7	55%	15.8%	+39.2
Agriculture	13%	17.8%	−4.8	11%	5.2%	+5.8
Manufacturing	12%	17.4%	−5.4	8%	10.1%	−2.1
Services	22%	53.9%	−31.9	18%	65.4%	−47.4
Technology	5%	3.6%	+1.4	4%	3.5%	+0.5
Other	3%	—	—	4%	—	—

Source: Authors' compilation from RBI (2023); RBNZ (2023); World Bank (2022); Stats NZ (2023).

2.4 Credit Allocation and Employment Linkages

The relationship between sectoral credit allocation and employment outcomes represents a critical but understudied dimension in small state economies (Hausmann & Rodrik, 2003). In economies like New Zealand, where agriculture represents a significant employment sector (6.8% of workforce), constrained credit access limits productivity improvements, value-added processing, and job creation potential (Stats NZ, 2023). Research indicates that targeted agricultural credit expansion can generate multiplier employment effects of 2.5–3.0 additional jobs per agricultural job supported (World Bank, 2021). This is particularly relevant for New Zealand's growing unemployment challenges, where youth unemployment (10.2%) and regional disparities exceed OECD averages.

2.5 Theoretical Framework


This research draws upon the Technology Acceptance Model (TAM), which provides insights into factors influencing FinTech adoption (Davis, 1989; Venkatesh & Davis, 2000). Extended TAM frameworks incorporate trust, security concerns, social influence, and facilitating conditions (Venkatesh et al., 2003). Diffusion of Innovation theory explains how new technologies spread through social systems (Rogers, 2003). The conceptual framework illustrates relationships between enabling infrastructure, FinTech innovations, mediating factors, and financial inclusion outcomes in Figure 1.

Additionally, the research considers the tension between innovation and stability. Historical episodes such as the 2008 Global Financial Crisis (GFC) and the 2022 cryptocurrency market collapse illustrate how unregulated innovation can lead to systemic risk (International Monetary Fund, 2009; Reuters, 2022). These lessons underscore the need for governance systems that balance experimentation with systemic risk mitigation (Cámara et al., 2015).

2.6 Summary and Research Gap

The preceding review reveals that while substantial literature exists on FinTech adoption, mobile banking, and financial inclusion in developing economies, particularly India and Sub-Saharan Africa, several important gaps remain. First, comparative studies that systematically analyse FinTech-enabled inclusion across developing and developed small-state economies within a single framework remain scarce. Second, most existing studies focus on access and adoption metrics, with limited attention to the structural credit allocation imbalances that constrain productive sector development, particularly agriculture. Third, the employment implications of FinTech-driven credit recalibration in small state economies, where sectoral employment

concentration creates systemic vulnerability, remain largely unexamined. Fourth, sustainability dimensions of digital finance, including energy economics and climate-resilient agricultural FinTech, represent an emerging frontier insufficiently integrated into the mainstream inclusion literature. This study addresses these gaps by providing a comparative mixed-methods analysis of rural India and New Zealand, connecting FinTech innovation to sectoral credit rebalancing, employment generation, and sustainability outcomes. The findings contribute both theoretical extensions to TAM and Diffusion of Innovation frameworks and practical evidence-based policy recommendations for diverse economic contexts.



Source: Author conceptualization based on literature review

Figure 1. Conceptual Framework: FinTech-Enabled Financial Inclusion

Source: Authors' own conceptualisation based on Davis (1989); Rogers (2003); Venkatesh et al. (2003).

3. Data and Methodology

This study employs a comparative case study methodology integrating quantitative and qualitative approaches to examine FinTech-enabled financial inclusion in rural India and New Zealand. Case study methodology proves appropriate for investigating contemporary phenomena within real-world contexts (Yin, 2018). Primary quantitative data derives from the World Bank Global Findex Database (2021), providing standardised financial inclusion indicators across 123 economies. Supplementary data sources include Reserve Bank of India statistics, National Payments Corporation of India (NPCI) transaction data, Reserve Bank of New Zealand (RBNZ) reports, Statistics New Zealand sectoral lending data, employment statistics, and Payments NZ industry data.

Qualitative data encompasses policy documents, regulatory frameworks, government reports, academic literature, and case studies from both jurisdictions. Case studies include Axis Bank-ITC collaboration, Paytm's rural microfinance initiative, and India Post Payments Bank (IPPB) in India, alongside analyses of New Zealand's digital inclusion initiatives for Māori and Pacific communities and sectoral credit allocation patterns.

Data analysis proceeds through multiple stages, including descriptive statistical analysis, comparative analysis, and thematic analysis (Braun & Clarke, 2006). The framework evaluates four dimensions: Access, Usage, Quality, and Impact (Cámara & Tuesta, 2014), capturing the nuanced reality that account opening is a necessary but insufficient condition for meaningful financial inclusion. Additional analysis examines sectoral credit allocation patterns, FinTech's potential to recalibrate, and projections of employment impacts.

4. Empirical Results and Discussion

4.1 Current State of Financial Exclusion

4.1.1 Rural-Urban Disparities in India


Analysis of Global Findex data in Table 2 reveals substantial progress in account ownership in India, rising from 53% in 2014 to 78% in 2021, with rural account ownership reaching 75%, compared with 83% in urban areas. However, account dormancy rates remain concerning, with approximately 40% of rural accounts experiencing no transactions in the preceding year (World Bank, 2021). Gender disparities persist, as rural women lag rural men in account ownership (72% versus 78%).

Digital payment adoption demonstrates stark urban-rural divides. In India, 57% of rural adults report making or receiving digital payments in 2021, compared to 68% of urban adults (World Bank, 2021). UPI transaction analysis reveals that metropolitan areas account for 65% of transaction volumes despite housing only 32% of India's population (NPCI, 2023). Refer to Figure 2 for more information.

Table 2. Financial Inclusion Indicators: Rural vs. Urban India (2014–2021)

Indicator	2014 Rural	2014 Urban	2017 Rural	2017 Urban	2021 Rural	2021 Urban
Account Ownership (%)	50	68	77	86	75	83
Women Account (%)	43	60	72	80	72	80
Digital Payments (%)	12	30	22	40	57	68
Formal Savings (%)	12	18	17	25	18	28
Formal Borrowing (%)	5	8	6	10	9	14
Dormancy Rate (%)	58	32	48	25	40	20

Source: World Bank Global Findex Database (2014, 2017, 2021).



Source: NPCI (2023); Census of India (2021)

Figure 2. Urban vs. Rural UPI Transaction Distribution in India (2023)

Source: Authors' analysis based on NPCI (2023); Census of India (2021).

4.1.2 New Zealand's Persistent Inclusion Gaps

Despite its advanced economy, New Zealand faces significant gaps in digital financial inclusion, particularly among Māori and Pacific communities. While 99% of adults have bank accounts in Table 3, 7–8% lack functional transaction accounts, with rates rising to 11% for Māori and 14% for Pacific peoples (Payments NZ, 2023). Rural broadband deficiencies affect 13% of households, and digital literacy gaps persist, especially among older adults. Culturally incompatible product design, prioritising individualism over collective values, and historical distrust of financial institutions further inhibit adoption.

In Table 3, New Zealand exhibits similar urban-rural patterns in digital adoption, with 72% overall mobile banking adoption masking substantial variations: only 45% of Pacific peoples and 52% of Māori report regular mobile banking use, compared with 78% of European New Zealanders (Payments NZ, 2023).

4.2 FinTech Adoption Patterns and Drivers

Mobile banking emerges as the primary FinTech modality driving financial inclusion in both contexts. India's mobile banking user base exceeded 425 million in 2023 (Reserve Bank of India, 2023). Analysis reveals that simplified KYC requirements under Aadhaar-based authentication reduced account opening time from weeks to minutes (Muralidharan et al., 2020). Digital payment infrastructure demonstrates remarkable scalability. UPI's real-time settlement processes transactions at approximately 3–5 rupees compared to 15–20

rupees for traditional methods (Sengupta & Vardhan, 2021). Peer-to-peer lending and alternative credit scoring show potential to expand credit access through smartphone data and digital footprints (Kumar & Rao, 2021).

Table 3. Digital Financial Inclusion Indicators: New Zealand by Ethnic Group (2023)

Indicator	European NZ	Māori	Pacific Peoples	Asian NZ	Overall
Bank Account (%)	99	89	86	97	99
Functional Transaction (%)	96	89	86	95	92
Mobile Banking (%)	78	52	45	75	72
Digital Literacy (%)	78	52	45	72	72
Broadband Access (%)	92	82	78	90	87
Trust in FIs (%)	72	48	42	65	65

Source: Authors' compilation from Payments NZ (2023); Financial Services Council NZ (2022); Stats NZ (2023).

In rural India, several FinTech innovations have demonstrated significant impact. Micro-ATMs in Table 4, operated by Business Correspondents (BCs), use Aadhaar verification to provide basic banking services in remote areas. By 2022, RapiPay's network of 500,000 micro-ATMs had reduced transaction costs by 40% and served 10 million rural consumers (Business Standard, 2020). The Business Correspondence (BC) model extends this reach further, with BCs acting as financial intermediaries in branchless areas to facilitate account opening and transaction services. By 2023, 1.2 million BCs were operating at 50% lower cost than traditional branches, serving 300 million rural customers (RBI, 2023).

The Village Level Entrepreneurship (VLE) model is another significant channel, providing access to banking, payments, and government schemes through Common Service Centres (CSCs). By 2023, as shown in Figure 3, 400,000 CSCs had enabled \$10 billion in transactions and generated 1.5 million jobs (Singh, 2023). Banking-as-a-Service (BaaS) platforms further expand the ecosystem by enabling non-banks to offer financial products via APIs; Razorpay's BaaS platform served 500,000 SMEs by 2023, though challenges related to data privacy and regulatory compliance persist (Razorpay, 2023). Additionally, AI and data analytics are being deployed to tailor financial products using alternative data sources. CredAvenue's AI-driven microloans for farmers achieved 90% repayment rates, illustrating the potential for technology-driven credit assessment in agricultural contexts (Times of India, 2023).

4.3 Barriers to Rural FinTech Adoption and Case Studies of Innovation

4.3.1 Barriers to Rural FinTech Adoption

Despite progress, significant barriers constrain rural FinTech adoption. Digital infrastructure deficits represent the most fundamental constraint in rural India, where only 37% of villages have 4G connectivity (Department of Telecommunications, 2023). New Zealand's rural areas experience connectivity challenges, with 13% of rural households lacking adequate broadband (Statistics New Zealand, 2022). Digital literacy gaps constitute critical impediments; only 38% of rural Indians report confidence in using digital financial

services compared to 62% of urban residents (NCAER, 2022). Age-related digital literacy gaps affect both contexts, with adults over 65 demonstrating significantly lower adoption rates.




Figure 3. FinTech Credit Growth: Alternative Lending Platforms (2018–2023)

Source: Authors' compilation from RBI (2023); RBNZ (2023); Kumar & Rao (2021).

Table 4. FinTech Innovation Models for Rural Financial Inclusion in India

Model	Technology	Scale (2023)	Cost Reduction	Users Served	Key Challenge
Micro-ATMs	Aadhaar Biometric	500,000 devices	40%	10M rural	Connectivity
BC Model	Mobile POS + Aadhaar	1.2M agents	50%	300M rural	Agent retention
VLE / CSC	Digital kiosks	400,000 CSCs	35%	\$10B transactions	Sustainability
BaaS	API platforms	500K SMEs	45%	SME finance	Data privacy
AI Credit	ML + alt. data	5M farmers	60%	90% repayment	Bias in algorithms

Source: Authors' compilation from RBI (2023); Singh (2023); Razorpay (2023); Times of India (2023).

Socio-cultural factors influence adoption patterns. Rural Indian communities exhibit a preference for cash transactions, rooted in anonymity and social norms (Sivathanu, 2019). In New Zealand, Māori and Pacific communities may prioritise collective economic values, which may be incompatible with individualistic financial product design, while distrust of mainstream institutions creates adoption barriers (Tapsell & Woods, 2021).

4.3.2 Case Studies of FinTech Innovation in Rural India

Axis Bank–ITC Collaboration in Table 5: In 2022, Axis Bank partnered with ITC's ITCMAARS platform, connecting 4 million farmers. Axis offered tailored loans and savings accounts, achieving 27% year-on-year growth in rural advances, a 12% increase in disbursements, and 16% deposit growth by December 2022 (Axis Bank, 2023). This hybrid model showcases FinTech-bank synergy for inclusion.

Paytm's Rural Microfinance Initiative: Launched in 2016, Paytm's BC model onboarded 15 million rural consumers across 600 districts by 2023 with \$5 billion in transactions, combining micro-ATMs and digital wallets. AI-driven credit scoring enabled microloans with 92% repayment rates. Though early fraud (10%) was reduced by literacy programmes, partnerships with Regional Rural Banks (RRBs) ensured KYC compliance (Paytm, 2023).

India Post Payments Bank (IPPB): Started in 2018, IPPB uses 136,000 post offices to reach 8 million rural consumers by 2023. Partnerships with FinTech like PhonePe facilitated \$1.2 billion in transactions. With 60% lower costs than traditional banks, doorstep banking reduced access barriers (IPPB, 2023).

Table 5. Comparative Analysis of Rural FinTech Case Studies in India

Parameter	Axis Bank–ITC	Paytm Rural Microfinance	India Post Payments Bank
Launch Year	2023	2016	2018
Reach	4M farmers	15M consumers, 600 districts	8M consumers, 136K post offices
Transaction Volume	—	\$5 billion	\$1.2 billion
Key Innovation	FinTech-bank hybrid via ITCMAARS	BC model + AI credit scoring	Leveraging postal network + PhonePe
Growth Metrics	27% rural advance growth	92% repayment rate	60% lower costs than banks
Main Challenge	Scale beyond ITC ecosystem	Early fraud (10%)	Digital literacy of users

Source: Authors' compilation from Axis Bank (2023); Paytm (2023); IPPB (2023).

4.3.3 Case Study: FinTech-Driven Financial Inclusion Initiatives in New Zealand

New Zealand's financial inclusion landscape presents a distinct case study in which digital innovation intersects with persistent demographic and geographic disparities. Three initiatives illustrate the trajectory of FinTech-enabled inclusion in this developed-economy context.


BNZ Digital Inclusion Programme: Bank of New Zealand's partnership with the Digital Inclusion Alliance Aotearoa (DIAA) combines subsidised device access, community-based digital literacy training, and simplified digital onboarding to target Māori, Pacific, and low-income communities. The programme has assisted over 40,000 individuals across community hubs, improving mobile banking adoption rates within participating communities (BNZ, 2023; DIAA, 2023). This initiative demonstrates that even in a developed economy, active community-embedded interventions, rather than passive digital platform provision, are required to reach persistently excluded populations.

Kiwibank's Rural Digital Banking Expansion: As a state-owned bank with a mandate for serving underserved communities, Kiwibank has invested in app-based banking services designed to compensate for rural branch closures. Integration with NZ Post infrastructure, paralleling the IPPB model in India, enables cash-in and cash-out services at PostShop locations nationwide, maintaining physical touchpoints while transitioning customers to digital platforms (Kiwibank, 2023). This hybrid phygital model provides a transferable template for developed economies managing the transition from branch networks to digital-first banking.

Māori-Led FinTech Initiatives: Emerging iwi-led FinTech initiatives, facilitated through Te Puni Kōkiri's digital enterprise fund, are developing culturally responsive lending platforms that reflect collective ownership values and integrate whānau-based credit assessment approaches. Early pilots demonstrate higher trust levels and engagement rates among Māori borrowers compared to mainstream banking interfaces, highlighting the centrality of culturally congruent product design (Te Puni Kōkiri, 2023). These initiatives reinforce a cross-contextual finding from both India and New Zealand: community co-creation and cultural alignment are not optional additions to digital financial services, but prerequisites for meaningful inclusion among indigenous and minority communities.

4.4 Policy and Regulatory Frameworks

India's regulatory approach balances innovation promotion with consumer protection. The Reserve Bank of India has established regulatory sandboxes, issued digital lending guidelines, and mandated data localisation (Reserve Bank of India, 2023). UPI's success demonstrates how public digital infrastructure can catalyse private sector innovation. New Zealand's framework emphasises principles-based regulation and technology neutrality. The Digital Identity Services Trust Framework Act (2023) establishes governance for digital identity services (Financial Markets Authority, 2023). Both jurisdictions recognise the importance of balancing innovation with consumer protection and financial stability, illustrated in Figure 4.



Source: RBI (2023); RBNZ (2023); FMA NZ (2023); Author analysis

Figure 4. Regulatory Framework Comparison: India vs. New Zealand

Source: Authors' analysis based on RBI (2023); RBNZ (2023); FMA NZ (2023).

4.5 Sectoral Credit Allocation Inequalities in the New Zealand Banking System

Despite New Zealand's highly developed financial sector, significant inequalities persist in sectoral credit allocation, reflecting structural biases in traditional banking. Analysis of Reserve Bank of New Zealand data in Table 6 reveals a disproportionate concentration of lending in residential property (approximately 55% of total bank lending), while productive sectors like agriculture, manufacturing, and technology receive comparatively limited credit allocation (RBNZ, 2023). This misalignment between credit flows and economic contribution is particularly evident in agriculture, a sector representing approximately 5.2% of GDP but receiving only 11% of business lending (Stats NZ, 2023).

Several structural barriers contribute to these inequalities (see Table 7). First, traditional banks employ standardised risk assessment models that undervalue agricultural and small business viability due to seasonal cycles, commodity price volatility, and perceived information asymmetries. Second, over-reliance on property collateral disadvantages agricultural enterprises whose primary assets, land, livestock, and equipment, are often undervalued or considered illiquid in conventional banking assessments. Third, branch network rationalisation has disproportionately affected rural communities through geographic concentration, reducing relationship banking and local credit assessment capabilities. Fourth, many mainstream banks lack specialised agricultural credit teams, resulting in sectoral expertise gaps that produce conservative lending practices toward this sector. These credit allocation imbalances have significant economic consequences: constrained agricultural innovation, limited value-added processing capacity, reduced export diversification, and heightened vulnerability to property market cycles.

Table 6. New Zealand Sectoral Credit Allocation vs. Economic Contribution

Sector	Credit Share (%)	GDP Share (%)	Employment Share (%)	Credit Gap	Lending (NZ\$ B)
Residential Property	55	15.8	—	+39.2pp	185.0
Agriculture	11	5.2	6.8	+5.8pp	37.0
Manufacturing	8	10.1	10.5	−2.1pp	26.9
Services (Business)	18	65.4	74.2	−47.4pp	60.5
Technology	4	3.5	4.5	+0.5pp	13.5
Other	4	—	4.0	—	13.5
Total	100	100	100	—	336.4

Source: Authors' analysis based on RBNZ (2023); Stats NZ (2023).

Table 7. New Zealand Structural Barriers to Agricultural Credit Access

Barrier	Description	Impact on Agriculture
Risk Assessment Models	Standardised models undervalue seasonal cycles, commodity volatility	Lower credit scores for viable agri-enterprises
Collateral Requirements	Over-reliance on property; agricultural assets undervalued	Farmers unable to access credit despite productive assets
Geographic Concentration	Branch closures in rural areas; reduced relationship banking	Loss of local credit assessment capability
Sectoral Expertise Gaps	Banks lack agricultural credit specialists and sector knowledge	Conservative lending; higher rejection rates

Source: Authors' analysis based on RBNZ (2023); Cole et al. (2011).


4.6 FinTech-Driven Recalibration of Sectoral Lending and Employment Implications in New Zealand

This section focuses exclusively on New Zealand, examining how FinTech innovations can recalibrate the structural credit allocation imbalances identified in Section 4.5 and the consequent employment generation potential for this small state economy.

4.6.1 FinTech Mechanisms for Sectoral Credit Recalibration

FinTech innovations offer transformative potential for recalibrating sectoral lending imbalances through several mechanisms. In the domain of alternative data and AI-driven credit assessment, satellite imagery, soil sensors, and weather tracking enable dynamic assessment of farm productivity and risk. Blockchain-enabled supply chain analytics facilitate tracking of agricultural produce from farm to market, creating verifiable performance data, while machine learning models trained on sector-specific data can identify viable agricultural enterprises overlooked by conventional scoring.

Specialised digital lending platforms represent another important channel. Agricultural FinTech platforms such as Figo (NZ) and AgriDigital (Australia) offer tailored products designed for the unique requirements of farming enterprises. Peer-to-peer agricultural lending platforms connect urban investors with rural enterprises, and equity-based crowdfunding platforms target agri-tech startups seeking capital for agricultural innovation. Digital public infrastructure for agricultural finance includes integrated farmer registries linking digital identities to land records, production history, and subsidy receipts. Digital warehouse receipts based on blockchain technology enable commodity-backed financing, while government-backed digital guarantee systems provide partial credit guarantees via digital platforms. In the area of embedded and supply chain finance, platform-enabled value chain financing allows digital marketplaces to embed financing at the point of transaction. Refer to Figure 5 and Table 8 for more information.



Source: NCAER (2022); World Bank Global Findex (2021); Ngaruiya & Schiffling (2023)

Figure 5. Digital Financial Literacy Levels by Demographic Groups

Source: Authors' analysis based on NCAER (2022); World Bank Global Findex (2021).

Table 8. FinTech Mechanisms for Sectoral Credit Recalibration

Mechanism	Technology	Application	Expected Impact
Alternative Credit Scoring	AI / ML + Satellite + IoT	Farm productivity assessment via non-financial data	30–40% more loan approvals for agri-enterprises
Specialised Platforms	Digital lending platforms	Harvest-cycle aligned repayment (Figured, AgriDigital)	150–200 bps interest margin reduction
Supply Chain Finance	Blockchain + Smart Contracts	Commodity-backed financing; warehouse receipts	Reduced collateral requirements by 40–60%
Peer-to-Peer Lending	P2P platforms	Urban investor to rural enterprise capital flow	NZ\$1–2B additional capital mobilised
Embedded Finance	APIs + BaaS	Financing at point of transaction in value chains	Improved cash flow for 50K+ farmers

Source: Authors' analysis based on BIS (2022); RBNZ (2023).

4.6.2 Employment Implications of Agricultural Credit Recalibration

The recalibration of agricultural lending in New Zealand through FinTech innovations carries significant employment implications for this small state economy. Analysis of sectoral employment multipliers reveals that agriculture exhibits one of the highest employment generation potential per unit of credit extended. The constrained credit allocation to agriculture (11% of business lending) relative to its employment contribution (6.8% of workforce) represents a structural inefficiency with direct unemployment implications.

New Zealand faces rising unemployment, particularly among youth (10.2% unemployment rate for 15–24-year-olds) and in regional areas where agricultural employment is concentrated (Stats NZ, 2023). Māori and Pacific communities experience disproportionately higher unemployment rates (8.3% and 11.7% respectively, versus 4.0% for European New Zealanders). These disparities are exacerbated by limited access to productive credit for agriculture and related value-added industries.

Based on employment multipliers derived from New Zealand Treasury and Ministry for Primary Industries data, every NZ\$1 million in additional agricultural credit can generate 8–12 direct and indirect jobs in primary production, with an additional 15–20 jobs in value-added processing, logistics, and related services. The constrained NZ\$2.1 billion credit gap identified in agriculture (Table 6) thus represents a potential employment deficit of approximately 16,800–25,200 direct jobs and 31,500–42,000 indirect jobs (Table 9).

Different agricultural subsectors exhibit varying employment potentials. Dairy farming generates significant downstream employment in processing (multiplier: 2.8). Horticulture is labour-intensive with strong export potential (multiplier: 4.2). Wine production offers high-value-added opportunities with tourism linkages (multiplier: 3.9). Māori agribusiness, currently underserved, demonstrates strong community employment benefits (multiplier: 4.5). Digital agricultural platforms combined with targeted credit could create 5,000–8,000 entry-level positions in precision agriculture, digital farm management, and agri-tech services over five years. Agricultural credit recalibration would disproportionately benefit regional New Zealand, where 65% of agricultural employment is concentrated in areas including Northland, Gisborne, and Bay of Plenty.

Table 9. Employment Impact Projections: Agricultural Credit Recalibration in New Zealand

Subsector	Employment Multiplier	Direct Jobs (est.)	Indirect Jobs (est.)	Key Beneficiary Groups
Dairy	2.8	4,200–6,300	7,800–10,500	Regional workers; processing staff
Horticulture	4.2	5,000–7,500	10,500–15,000	Youth; seasonal workers; Māori
Wine Production	3.9	2,800–4,200	5,400–7,200	Tourism-linked; regional SMEs
Māori Agribusiness	4.5	3,500–5,200	5,800–7,300	Māori & Pacific communities
Other Agriculture	3.0	1,300–2,000	2,000–3,000	Mixed rural communities
Total	—	16,800–25,200	31,500–42,000	—

Source: Authors' projections based on NZ Treasury; Ministry for Primary Industries; Stats NZ (2023).

4.6.3 Expected Outcomes and Implementation Pathways

Implementation of these digital strategies could potentially increase agriculture's share of business lending from 11% to 18–20% within five years, mobilise an additional NZ\$4–6 billion in agricultural credit, reduce average interest margins by 150–200 basis points, and improve loan approval rates for small-to-medium agricultural enterprises by 30–40%. Table 10 illustrates the expected outcomes of different implementation pathways.

Table 10. Implementation Pathways for FinTech-Driven Agricultural Credit Recalibration in New Zealand

Pathway	Description	Key Actions	Expected Timeline
Regulatory Sandbox for Agricultural FinTech	Development of a specialised regulatory sandbox to test agricultural finance innovations in a controlled environment	Establish sandbox criteria; invite agricultural FinTech startups; evaluate outcomes for regulatory adaptation	Year 1–2
Digital Farm Data Commons	Establishment of a secure, farmer-controlled data repository enabling standardised agricultural data sharing	Design data governance framework; build secure platform; pilot with farming cooperatives	Year 1–3
Sector-Specific Open Banking Standards	Development of APIs enabling agricultural data sharing between financial institutions and FinTech providers	Define API standards; implement data-sharing protocols; ensure privacy compliance	Year 2–4
FinTech-Bank Collaboration Models	Incentive structures for traditional banks to partner with agricultural FinTech specialists	Create co-lending frameworks; develop risk-sharing arrangements; establish joint product development	Year 2–5
Māori-Led Digital Finance Initiatives	Support for iwi partnerships with FinTech developers to create culturally responsive financial products	Fund iwi-FinTech partnerships; co-design platforms reflecting collective values; integrate te ao Māori principles	Year 1–5

Source: Authors' own analysis based on comparative study findings and stakeholder consultation.

4.7 Economic Impact of FinTech-Driven Inclusion: Comparative Evidence from India and Global Contexts

This section examines the broader economic impact of FinTech-driven financial inclusion, drawing primarily on evidence from India, with reference to global comparators including Kenya and Sub-Saharan Africa, to contextualise the scale of digital financial transformation examined in this study.

Financial inclusion significantly drives economic growth (Van et al., 2021) and helps reduce economic inequality (Kling et al., 2022). In India, digital payments rose from 35% to 57% of adults (2014–2021), contributing to poverty reduction (World Bank, 2022). Mobile money lifted approximately 2% of Kenya's population out of poverty between 2008 and 2014 (Islam et al., 2022). FinTech-driven microfinance supported 5 million Indian rural entrepreneurs, generating \$2 billion by 2023 (RBI, 2023).

Digital payments have a catalytic effect; most adults who receive digital payments also make them and are more likely to save, borrow, or store money. Beyond payments, mobile money accounts have become important for saving in regions such as Sub-Saharan Africa, where 15% of adults use them, a share similar to that of formal bank accounts (World Bank, 2022) as shown in Figure 6.




Figure 6. Economic Impact of FinTech-Driven Financial Inclusion in India

Source: Authors' analysis based on World Bank Global Findex Database (2014, 2017, 2021); RBI (2023).

4.8 Regulatory Institutions and Governance: India, New Zealand, and International Frameworks

This section analyses regulatory approaches to FinTech governance, with specific reference to India's Reserve Bank of India Regulatory Sandbox model and international frameworks, before considering implications for New Zealand's regulatory environment.

Regulatory bodies balance FinTech's opportunities and risks. The Bank for International Settlements (BIS) Innovation Hub develops AI-based risk monitoring (BIS, 2023). The RBI's Regulatory Sandbox, launched in 2019, has tested 50 FinTech products, ensuring KYC and cybersecurity compliance (RBI, 2023). In 2021, the sandbox tested a blockchain-based microloan platform that enabled 10,000 farmers to access credit

at 8% interest, compared with 15% from traditional lenders (RBI, 2022). The UK's Financial Conduct Authority (FCA) sandbox supported 200 FinTech entities, reducing entry barriers (FCA, 2023).

The RBI's simplified KYC enabled 100 million rural account openings (RBI, 2023). The EU's AI Act addresses algorithmic biases (European Commission, 2021). However, cybersecurity concerns persist, with 40% of FinTech enhancing security post-COVID (World Bank, 2021). For New Zealand, these international experiences provide actionable templates for developing a sector-specific agricultural FinTech sandbox and a digital identity framework capable of supporting rural financial inclusion at scale.

4.9 Post-COVID-19 Acceleration and Recent Crises

COVID-19 boosted FinTech adoption, with finance app downloads rising 20% globally in 2020 (Fu & Mishra, 2022). In India, digital payments grew to 57% of adults by 2021, with UPI processing approximately \$1.5 trillion in 2022 (RBI, 2023). Rural women accounted for 30% of new account holders (World Bank, 2020).

Recent financial crises highlight ongoing vulnerabilities. The 2023 Silicon Valley Bank collapse underscored digital-first risks (Reuters, 2023). India's 2021 ban on 100 lending apps curbed predatory practices (RBI, 2021). The 2022 FTX cryptocurrency exchange failure, affecting 1 million creditors and \$8 billion in consumer funds, illustrates dangers in insufficiently regulated digital platforms (Reuters, 2022). Table 11 summarises these developments.

Table 11. Post-COVID-19 FinTech Acceleration Indicators

Indicator	Pre-COVID (2019)	Post-COVID (2021–22)	Change
Finance App Downloads (Global)	Baseline	+20% (2020)	Sustained growth
India Digital Payments (% adults)	29%	57%	+28pp
UPI Annual Volume	\$0.4 trillion	\$1.5 trillion (2022)	+275%
Rural Women New Accounts (India)	—	30% of new account holders	Significant gender gain
NZ Mobile Banking Adoption	58%	72%	+14pp
Cybersecurity Investment (FinTech)	Baseline	40% increase post-COVID	Enhanced security focus

Source: Authors' compilation from World Bank (2021, 2022); RBI (2023); Payments NZ (2023); ScienceDirect (2021).

4.10 Sustainability Dimensions of FinTech and Digital Financial Solutions


This section examines the sustainability dimensions of FinTech across two primary geographic contexts - New Zealand and South Asia. Within South Asia, the primary analytical focus is India, given its scale, data availability, and the centrality of Indian FinTech to this study. Where evidence from Bangladesh, Sri Lanka, or Nepal is incorporated, it is explicitly identified to distinguish country-specific findings from broader regional trends.

The integration of sustainability considerations within FinTech and digital financial solutions represents an emerging frontier in both developed and developing contexts, with distinct implications for New Zealand and South Asia. This section examines how digital financial innovations intersect with environmental sustainability, energy economics, and climate resilience, revealing both opportunities and challenges across diverse economic landscapes.

4.10.1 Environmental Sustainability Through Digital Finance

Digital financial services inherently reduce environmental footprints through several mechanisms. In New Zealand, as shown in Figure 8, where the financial sector accounts for approximately 2.3% of national carbon emissions (Ministry for the Environment, 2023), digital banking solutions offer tangible sustainability benefits. Electronic transactions generate only 0.5 grams of CO₂ equivalent compared to 4.5 grams for traditional paper-based transactions (RBNZ, 2022). The transition from physical branches to digital platforms has reduced energy consumption in New Zealand's banking sector by 18% between 2019–2023, while paper usage declined by 62% (Banking Association NZ, 2023).

In South Asia, particularly India, the sustainability implications are more complex. While digital payments reduce paper consumption (UPI transactions saved an estimated 1.2 million trees in 2022), the rapid expansion of digital infrastructure raises concerns about increased energy consumption. India's data centres supporting digital financial services consumed 7.4 TWh in 2022, representing 2.1% of national electricity consumption (IEA, 2023). However, this must be contextualised against the alternative: maintaining physical banking infrastructure in 600,000 villages would require significantly greater resource extraction and energy use.



Source: IEA (2023); EECA (2023); RBNZ (2022); MfE NZ (2023); Author analysis

Figure 8. Renewable Energy Integration in Digital Financial Ecosystems

Source: Authors' analysis based on IEA (2023); EECA (2023); RBNZ (2022); Ministry for the Environment NZ (2023).

4.10.2 Energy Economics of Digital Financial Ecosystems

New Zealand's renewable energy advantage (84% renewable electricity generation) positions its FinTech sector favourably for sustainable operations. Digital banking platforms operating on renewable energy sources demonstrate carbon intensities of 12–18 grams CO₂ per transaction, compared to 35–42 grams in jurisdictions with fossil-fuel-dominated grids (EECA, 2023). In South Asia, India's digital financial infrastructure drives efficiency gains elsewhere in the economy. Digital payments reduce transportation emissions associated with cash handling by approximately 1.2 million metric tons CO₂ annually (NPCI, 2022). Furthermore, pay-as-you-go (PAYG) solar financing in rural India has enabled 8.4 million households to access solar energy, reducing kerosene consumption by 1.8 billion litres annually (International Solar Alliance, 2023). Digital microgrid financing platforms in Bangladesh have facilitated \$320 million in renewable energy investments since 2020 (World Bank, 2023).

4.10.3 Climate Resilience and Agricultural FinTech

The intersection of agricultural FinTech and climate resilience is a critical dimension of sustainability. In New Zealand, as shown in Table 12, where agriculture contributes 48% of greenhouse gas emissions (Ministry for Primary Industries, 2023), digital financial innovations offer pathways to sustainable intensification. Precision agriculture platforms combining IoT sensors, satellite data, and AI analytics enable 15–25% reduction in fertiliser application through targeted dosing, 20–30% water conservation through smart irrigation systems, and 10–18% methane reduction in dairy through optimised feed management. Rabobank's sustainability-linked loans in New Zealand, offering preferential rates for achieving environmental targets, reached NZ\$1.2 billion in 2023 (Rabobank NZ, 2023).

In South Asia, particularly India and Bangladesh, climate-resilient agricultural finance addresses more immediate adaptation challenges. Digital platforms integrating climate risk assessments into credit scoring enable parametric insurance products based on weather data that protect 3.2 million smallholder farmers against climate shocks. Drought-resilient crop financing with integrated advisory services has reduced crop failure rates by 22–35%, while water conservation-linked credit systems have improved water use efficiency by 40% in water-stressed regions.

4.10.4 Sustainable Investment and Green Finance Platforms

FinTech innovations are democratising sustainable investment opportunities. In New Zealand, digital platforms have mobilised NZ\$4.8 billion in green bonds and sustainability-focused investments since 2020, representing 18% of all retail investment flows (Financial Markets Authority, 2023). Key developments include digital green bond platforms enabling retail participation in renewable energy projects, ESG robo-advisors managing NZ\$2.3 billion in assets, and carbon footprint tracking apps integrated with banking services used by 32% of digital banking customers. South Asia presents a different paradigm, where green FinTech in India focuses on solar rooftop financing (850,000 households), electric vehicle financing (210,000 EV purchases), and waste management payments in 42 cities.

4.10.5 Policy and Regulatory Dimensions

New Zealand's approach encompasses climate-related financial disclosure regulations (effective 2023) mandating sustainability reporting for large financial institutions, a green FinTech regulatory sandbox testing

24 sustainability-focused innovations, and a sustainable finance taxonomy guiding NZ\$12 billion in annual green investments. In South Asia, India's Sustainable Finance Task Force (established 2022) is coordinating FinTech and sustainability policies, while Bangladesh's Green Transformation Fund provides USD 500 million for sustainable FinTech initiatives.

Table 12. Sustainability Metrics for Digital Financial Services: New Zealand vs. South Asia

Metric	New Zealand (2023)	India (2023)	Bangladesh (2023)	Comparative Notes
Digital banking energy use (TWh/yr)	0.42	7.4	0.38	NZ efficiency advantages
Renewable energy share in digital infra	84%	28%	12%	Grid composition differences
Carbon intensity per transaction (g CO ₂)	12–18	35–42	45–52	Varies with energy mix

Metric	New Zealand (2023)	India (2023)	Bangladesh (2023)	Comparative Notes
Paper reduction from digital transactions	89%	76%	68%	Maturity differences
E-waste from FinTech devices (tons/yr)	4,200	890,000	62,000	Scale disparities
Trees saved through digital payments	240,000	1.2 million	180,000	Absolute vs relative impact

Source: Authors' compilation from national statistics agencies, central banks, environmental agencies, and industry associations (2023).

4.10.6 Challenges, Trade-offs, and Future Directions

Despite promising developments, significant challenges persist. The digital divide and energy access gap remain critical, as 32% of rural South Asians lack reliable electricity for digital financial services. E-waste from rapid device turnover generates 3.2 million metric tons annually in South Asia. Rebound effects represent a further concern, as digital efficiency gains may increase consumption, partially offsetting environmental benefits. Data centre energy intensity is projected to increase 3–5 times by 2030 in both regions. Future priorities include integrating sustainability metrics into digital credit scoring, developing renewable-powered digital infrastructure in South Asia's off-grid regions, creating circular economy financing platforms with material tracking capabilities, and establishing just transition mechanisms to ensure sustainability benefits reach vulnerable populations.


5. Discussion and Implications

Cross-contextual analysis yields transferable insights. First, digital public infrastructure investments generate network effects that pure market-driven approaches may not achieve. UPI demonstrates how open, interoperable platforms catalyse ecosystem development. Second, simplified identity verification proves critical for expanding access. Third, digital financial inclusion requires comprehensive strategies addressing infrastructure, digital literacy, culturally appropriate design, and consumer protection simultaneously. Fourth, exclusion persists even in developed economies, underscoring that technological sophistication alone cannot

guarantee inclusion. Fifth, structural inequalities in credit allocation require targeted digital interventions to rebalance lending toward productive sectors with significant employment implications for small state economies.

Policy implications span multiple domains as illustrated in Figure 7. Governments must invest in digital infrastructure as foundational public goods. Digital literacy programmes should target vulnerable populations, including women, older adults, and indigenous communities. Regulatory frameworks must balance innovation with robust consumer protection, using tools like regulatory sandboxes to test innovations safely. Financial institutions should prioritise user-centred design, incorporating feedback from target populations. Sector-specific digital finance initiatives should be developed to address credit allocation imbalances, particularly in agriculture, with explicit employment generation targets.

FinTech has emerged as a powerful tool in combating rural poverty by expanding access to financial services, facilitating microfinance, enabling digital payments, promoting agricultural technologies, providing financial literacy, creating employment opportunities, and recalibrating sectoral credit allocation. Ongoing collaboration among FinTech companies, government entities, and financial institutions will be essential to ensure benefits reach the most marginalised regions and productive sectors.



Source: Author analysis based on comparative study findings

Figure 7. Policy Implications for Inclusive Digital Financial Ecosystems

Source: Authors' own analysis based on comparative study findings.

6. Conclusion

This research examined the role of FinTech innovation in advancing financial inclusion in rural India and New Zealand, revealing both transformative potential and persistent implementation challenges. Key findings indicate substantial progress in account ownership and digital payment adoption, yet significant rural-urban and demographic disparities persist. Infrastructure deficits, digital literacy gaps, socio-cultural factors,

and structural inequalities in credit allocation constrain inclusive financial development across different levels of economic development.

India's ambitious digital public infrastructure initiatives demonstrate the feasibility of large-scale transformation, while New Zealand's experience highlights the importance of culturally responsive approaches and the need to address sectoral lending imbalances with their employment implications. FinTech innovations are reshaping rural banking in India, enabling financial inclusion for approximately 300 million rural individuals through technologies such as micro-ATMs, BC models, VLE frameworks, and AI. Case studies including Axis Bank-ITC, Paytm, and IPPB illustrate scalability and transformative potential.

The research further demonstrates how FinTech can rebalance sectoral credit allocation toward productive sectors like agriculture through alternative data, specialised platforms, and digital public infrastructure. In New Zealand, digital innovations offer pathways to increase agriculture's share of business lending from 11% to 18–20%, mobilising additional capital for productive investment while enhancing inclusion for Māori agribusinesses. Critically, this credit recalibration carries significant employment implications for New Zealand's small state economy, potentially addressing structural unemployment through direct job creation in agriculture (16,800–25,200 jobs) and indirect employment in value-added processing and services (31,500–42,000 jobs), with benefits for youth, Māori and Pacific communities, and regional economies.

The stability of FinTech-driven systems hinges on robust governance, adaptive market frameworks, and effective regulatory oversight. Historical disruptions and the post-COVID-19 acceleration underscore the need for vigilant regulation to mitigate risks such as cybersecurity threats and digital literacy gaps.

The findings of this study carry several concrete implications for policymakers and financial sector actors. First, governments in both developing and developed economies should prioritise digital public infrastructure, particularly interoperable payment rails and digital identity systems, as foundational investments that catalyse private sector inclusion. Second, regulatory sandboxes should be adapted with sector-specific tracks for agricultural FinTech, enabling controlled testing of alternative credit scoring and embedded finance models that address persistent collateral-based exclusion. Third, financial institutions operating in contexts with significant indigenous or minority communities, such as Māori and Pacific peoples in New Zealand or tribal populations in rural India, should mandate culturally co-designed product development, moving beyond surface-level adaptation toward structural integration of collective economic values. Fourth, policymakers in small state economies should formally incorporate FinTech-enabled credit recalibration into employment and regional development strategies, given the quantified multiplier effects demonstrated in the agricultural sector analysis. Fifth, sustainability-linked financial products and digital green bond platforms require coordinated regulatory support to scale beyond pilot stages and contribute meaningfully to climate resilience in both agricultural and energy contexts.

Future research should examine longitudinal impacts through rigorous evaluations, investigate emerging technologies, expand comparative studies across diverse settings, and analyse the long-term effects of digital interventions on sectoral credit allocation patterns and employment outcomes. Achieving universal financial inclusion requires sustained commitment across governments, financial institutions, FinTech innovators, and communities. By fostering collaborative ecosystems that leverage FinTech's agility alongside traditional banking's stability, societies can build inclusive, resilient financial systems that empower

marginalised populations, rebalance credit toward productive sectors, generate meaningful employment, and drive sustainable, equitable economic growth in both developing and developed contexts.

Data Availability Statement

The data supporting the findings of this study are publicly available. World Bank Global Findex Database data are accessible at <https://www.worldbank.org/en/publication/globalindex>. Reserve Bank of India statistics are available at <https://www.rbi.org.in/>. National Payments Corporation of India data can be accessed at <https://www.npci.org.in/>. Reserve Bank of New Zealand reports and sectoral lending data are available at <https://www.rbnz.govt.nz/>. Statistics New Zealand data at <https://www.stats.govt.nz/>. Payments NZ data are accessible at <https://www.paymentsnz.co.nz/>.

Funding

This research received no external funding.

Acknowledgements

The author gratefully acknowledges Dr. Hanoku Bathula of the University of Auckland for reviewing an earlier version of the manuscript and providing valuable comments and suggestions. Any remaining errors or omissions are the sole responsibility of the author.

Declaration of the Use of Generative AI

Used Grammarly to assist with grammar and sentence-level proofreading.

Conflicts of Interest

The author declares no conflict of interest.

References

- Arner, D. W., Buckley, R. P., Zetsche, D. A., & Veidt, R. (2020). Sustainability, FinTech and financial inclusion. *European Business Organization Law Review*, 21, 7–35. <https://doi.org/10.1007/s40804-020-00183-y>
- Axis Bank. (2023). *Annual report 2022–23: Rural and agricultural banking initiatives*. Axis Bank.
- Bank for International Settlements. (2022). *FinTech and financial inclusion: A review of policy approaches* (BIS Papers No. 117). <https://www.bis.org/publ/bppdf/bispap117.pdf>
- Bank for International Settlements. (2023). *Annual economic report 2023: Innovation Hub activities*. <https://www.bis.org/publ/arpdf/ar2023e.htm>
- Banking Association New Zealand. (2023). *Sustainable banking report 2023: Digital transformation and environmental impact*.
- Beck, T., Demirgüç-Kunt, A., & Levine, R. (2010). Financial institutions and markets across countries and over time. *The World Bank Economic Review*, 24(1), 77–92. <https://doi.org/10.1093/wber/lhp016>

- Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- Business Standard. (2020, October 15). RapiPay expands micro-ATM network to 500,000 devices across rural India. *Business Standard*.
- Chen, M. A., Wu, Q., & Yang, B. (2019). How valuable is FinTech innovation? *The Review of Financial Studies*, 32(5), 2062–2106. <https://doi.org/10.1093/rfs/hhy130>
- Cole, S., Sampson, T., & Zia, B. (2011). Prices or knowledge? What drives demand for financial services in emerging markets? *The Journal of Finance*, 66(6), 1933–1967. <https://doi.org/10.1111/j.1540-6261.2011.01696.x>
- Cámara, N., & Tuesta, D. (2014). *Measuring financial inclusion: A multidimensional index* (BBVA Research Paper, 14(26), 1–40). https://www.bbvaresearch.com/wp-content/uploads/2014/09/WP14-26_Financial-Inclusion.pdf
- Cámara, N., Peña, X., & Tuesta, D. (2015). Factors that matter for financial inclusion: Evidence from Peru. *Aestimatio: The IEB International Journal of Finance*, 10, 10–31. <https://doi.org/10.5605/IEB.10.2>
- Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 13(3), 319–340. <https://doi.org/10.2307/249008>
- Demirgüç-Kunt, A., & Klapper, L. (2013). Measuring financial inclusion: Explaining variation in use of financial services across and within countries. *Brookings Papers on Economic Activity*, 2013(1), 279–340. <https://doi.org/10.1353/eca.2013.0002>
- Demirgüç-Kunt, A., Klapper, L., Singer, D., & Ansar, S. (2022). *The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. World Bank Publications. <https://doi.org/10.1596/978-1-4648-1897-4>
- Department of Telecommunications, Government of India. (2023). *Annual report 2022–23*. Ministry of Communications. <https://dot.gov.in/annual-report>
- Energy Efficiency and Conservation Authority. (2023). *Energy in New Zealand 2023: Digital infrastructure and renewable energy integration*. <https://www.eeca.govt.nz/insights/data-tools/energy-in-new-zealand/>
- European Commission. (2021). *Proposal for a Regulation laying down harmonised rules on artificial intelligence (Artificial Intelligence Act)*. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206>
- Financial Conduct Authority. (2023). *Regulatory sandbox: Lessons learned report*. <https://www.fca.org.uk/firms/innovation/regulatory-sandbox>
- Financial Markets Authority, New Zealand. (2023). *Regulatory sandbox guidance*. <https://www.fma.govt.nz/>
- Financial Services Council New Zealand. (2022). *Financial inclusion in New Zealand: Current state and future directions* (FSC Policy Brief). https://www.fsc.org.nz/site/fsc/files/Policy%20documents/FSC_Financial_Inclusion.pdf
- Fu, J., & Mishra, M. (2022). FinTech in the time of COVID-19: Technological adoption during crises. *Journal of Financial Intermediation*, 46, Article 100945. <https://doi.org/10.1016/j.jfi.2021.100945>
- 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>

- Hausmann, R., & Rodrik, D. (2003). Economic development as self-discovery. *Journal of Development Economics*, 72(2), 603–633. [https://doi.org/10.1016/S0304-3878\(03\)00124-X](https://doi.org/10.1016/S0304-3878(03)00124-X)
- India Post Payments Bank. (2023). *Annual report 2022–23*. India Post Payments Bank. <https://www.ippbonline.com/web/ippb/annual-reports>
- International Energy Agency. (2023). *World energy outlook 2023: Data centres and energy consumption in South Asia*. <https://www.iea.org/reports/world-energy-outlook-2023>
- International Monetary Fund. (2009). *Global financial stability report: Navigating the financial challenges ahead*. <https://www.imf.org/en/Publications/GFSR/Issues/2016/12/31/Navigating-the-Financial-Challenges-Ahead>
- International Solar Alliance. (2023). *Solar energy financing in South Asia: Pay-as-you-go models and rural electrification*. <https://isolaralliance.org/work/reports>
- Islam, A., Muzi, S., & Meza, J. L. (2022). Does mobile money use increase firms' investment? Evidence from enterprise surveys in Kenya, Uganda, and Tanzania. *Small Business Economics*, 59(3), 1065–1090. <https://doi.org/10.1007/s11187-020-00393-1>
- Kling, G., Pesqué-Cela, V., Tian, L., & Luo, D. (2022). A theory of financial inclusion and income inequality. *The European Journal of Finance*, 28(1), 137–157. <https://doi.org/10.1080/1351847X.2020.1792901>
- Kumar, A. (2023). AI and data analytics in rural financial services: Customizing products for underserved populations. *Journal of Financial Technology*, 4(1), 45–68.
- Kumar, A., & Rao, S. (2021). Alternative credit scoring in Indian FinTech: Opportunities and challenges. *Journal of Banking and Financial Technology*, 5(2), 167–183.
- 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>
- McKinsey & Company. (2021). *Financial services for the underserved: Technology-driven approaches in developing economies* (McKinsey Global Institute Report).
- Mention, A. L. (2019). The future of FinTech. *Research-Technology Management*, 62(4), 59–63. <https://doi.org/10.1080/08956308.2019.1613123>
- Ministry for Primary Industries, New Zealand. (2023). *Situation and outlook for primary industries (SOPI) 2023*. <https://www.mpi.govt.nz/resources-and-forms/economic-intelligence/situation-and-outlook-for-primary-industries/>
- Ministry for the Environment, New Zealand. (2023). *New Zealand's greenhouse gas inventory 1990–2021*. <https://environment.govt.nz/publications/new-zealands-greenhouse-gas-inventory-1990-2021/>
- Muralidharan, K., Niehaus, P., & Sukhtankar, S. (2020). *Identity verification standards in welfare programs: Experimental evidence from India* (NBER Working Paper No. 26744). <https://www.nber.org/papers/w26744>
- National Council of Applied Economic Research. (2022). *Digital financial literacy in India: A household survey*. https://www.ncaer.org/publication_detail.php
- National Payments Corporation of India. (2022). *Annual report 2021–22: Digital payments and sustainability metrics*. <https://www.npci.org.in/PDF/npci/annual-report/NPCI-Annual-Report-2021-22.pdf>

- National Payments Corporation of India. (2023). *UPI ecosystem statistics and annual review 2022–23*. <https://www.npci.org.in/what-we-do/upi/product-statistics>
- Ozili, P. K. (2023). FinTech and financial inclusion: A framework. *International Journal of Financial Studies*, 11(1), 17. <https://doi.org/10.3390/ijfs11010017>
- Park, C. Y., & Mercado, R. (2018). *Financial inclusion: New measurement and cross-country impact assessment* (ADB Economics Working Paper Series, No. 539). <https://www.adb.org/publications/financial-inclusion-new-measurement-cross-country-impact-assessment>
- Patwardhan, A. (2024). Financial inclusion through FinTech: Perspectives from rural India. *Journal of Financial Services Marketing*, 29(1), 89–105. <https://doi.org/10.1057/s41264-022-00178-7>
- Payments NZ. (2023). *Annual payments trends report 2023*. Payments NZ. <https://www.paymentsnz.co.nz/resources/industry-reports/>
- Paytm. (2023). *Paytm rural financial services: Impact report 2023*. Paytm.
- Rabobank New Zealand. (2023). *Sustainability-linked lending report 2023: Agricultural finance and environmental outcomes*. Rabobank New Zealand.
- Rao, P., & Gupta, S. (2022). Barriers to financial inclusion in rural India: A systematic review. *International Journal of Rural Management*, 18(2), 234–256. <https://doi.org/10.1177/09730052211040143>
- Razorpay. (2023). *Banking-as-a-service annual report: SME financial inclusion through API-driven platforms*.
- Reserve Bank of India. (2021). *Guidelines on digital lending: Regulatory framework for lending apps*. <https://rbidocs.rbi.org.in/rdocs/notification/PDFs/NOTI1B2B5E6A4A4744C2BBAAD8E7C1B6C49B.PDF>
- Reserve Bank of India. (2022). *Report on regulatory sandbox: Blockchain-based microlending pilot outcomes*. <https://www.rbi.org.in/Scripts/PublicationReportDetails.aspx?UrlPage=&ID=1201>
- Reserve Bank of India. (2023). *Annual report 2022–23*. <https://www.rbi.org.in/Scripts/AnnualReportPublications.aspx?Id=1381>
- Reserve Bank of New Zealand. (2022). *Climate-related financial disclosures and digital banking sustainability metrics*. <https://www.rbnz.govt.nz/financial-stability/climate-change>
- Reserve Bank of New Zealand. (2023). *Financial stability report*. Reserve Bank of New Zealand. <https://www.rbnz.govt.nz/financial-stability/financial-stability-report>
- Reuters. (2022, November 11). FTX cryptocurrency exchange collapse: Timeline and impact assessment. *Reuters*.
- Reuters. (2023, March 13). Silicon Valley Bank collapse triggers global banking concerns. *Reuters*.
- Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.
- Sarma, M., & Pais, J. (2011). Financial inclusion and development. *Journal of International Development*, 23(5), 613–628. <https://doi.org/10.1002/jid.1595>
- Sengupta, R., & Vardhan, H. (2021). UPI transformation: Revolutionising digital payments in India. *Journal of Payments Strategy & Systems*, 15(3), 258–271.

- Singh, R. (2023). Village level entrepreneurship and common service centres: Digital financial inclusion through community-based delivery models. *Indian Journal of Public Administration*, 69(2), 312–328.
- Sivathanu, B. (2019). Adoption of digital payment systems in the era of demonetization in India: An empirical study. *Journal of Science and Technology Policy Management*, 10(1), 143–171. <https://doi.org/10.1108/JSTPM-09-2017-0046>
- Statistics New Zealand. (2022). *Digital divide: Internet access in New Zealand*. <https://www.stats.govt.nz/reports/digital-divide-report/>
- Statistics New Zealand. (2023). *Labour market statistics, sectoral employment data, and regional economic indicators 2023*. <https://www.stats.govt.nz/topics/labour-market>
- Suri, T., & Jack, W. (2016). The long-run poverty and gender impacts of mobile money. *Science*, 354(6317), 1288–1292. <https://doi.org/10.1126/science.aah5309>
- Tapsell, P., & Woods, C. (2021). Social entrepreneurship and innovation in Māori communities. *New Zealand Journal of Social Sciences Online*, 16(1), 45–62. <https://doi.org/10.1080/1177083X.2021.1919268>
- Times of India. (2023, May 22). *CredAvenue's AI-driven microloans achieve 90% repayment among farmers*. *The Times of India*.
- Uzma, S., & Pratihari, S. K. (2019). Financial inclusion through FinTech: A study of micro-ATMs and business correspondents in rural India. *International Journal of Bank Marketing*, 37(4), 1034–1059. <https://doi.org/10.1108/IJBM-02-2018-0036>
- Van, L. T. H., Vo, A. T., Nguyen, N. T., & Vo, D. H. (2021). Financial inclusion and economic growth: An international evidence. *Emerging Markets Finance and Trade*, 57(1), 239–263. <https://doi.org/10.1080/1540496X.2019.1679907>
- Venkatesh, V., & Davis, F. D. (2000). A theoretical extension of the Technology Acceptance Model: Four longitudinal field studies. *Management Science*, 46(2), 186–204. <https://doi.org/10.1287/mnsc.46.2.186.11926>
- Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>
- World Bank. (2020). *COVID-19 and financial inclusion: Impact on digital account adoption in developing economies*. <https://www.worldbank.org/en/topic/financialinclusion/publication/the-global-findex-database-2021>
- World Bank. (2021). *The Global Findex Database 2021*. <https://www.worldbank.org/en/publication/globalfindex>
- World Bank. (2022). *The Global Findex Database 2021: Financial inclusion, digital payments, and resilience in the age of COVID-19*. <https://doi.org/10.1596/978-1-4648-1897-4>
- World Bank. (2023). *Digital finance for climate action: Sustainable FinTech in South Asia*. <https://www.worldbank.org/en/topic/financialinclusion>
- Yin, R. K. (2018). *Case study research and applications: Design and methods* (6th ed.). SAGE Publications. <https://doi.org/10.4135/9781473951556>

Comparison of digital RMB and private payment platform: Gap Analysis and Improvement.

Gesheng Chen*

Authors

Gesheng Chen, Independent Researcher,
Auckland, New Zealand.

Email: suzhoucgs@gmail.com

* Corresponding author.

Abstract

Although the digital RMB in China has made some achievements since its launch in late 2019, there is still a big gap in comparison between the RMB and We Chat as well as AliPay. The study uses the Unified Theory of Acceptance and Use of Technology (UTAUT) and Porter's competition Theory to construct a model to identify how the differences in each key factor affect the willingness to use digital RMB. The research was conducted by collecting 314 questionnaires, testing the reliability and validity of the questionnaires, using SEM modelling to make different attempts on the pathways, and discovering the optimal route, which verified the hypotheses and their relationships among the variables. The results of the model showed that awareness, habit, and safety have a weak positive effect on intention to use digital RMB, while compatibility and trust have a low moderate impact on the intention. In addition, it shows that though facility condition, perceived ease of use, and perceived usefulness have no direct impact on usage intention, they have a direct impact on habit, thus indirectly affecting usage intention. The study also used the K-Means model to group respondents and analysed them with low scores in several dimensions. The study reveals that more than 70% are young and middle-aged people with knowledge, and a certain economic base can understand and have the ability to use digital RMB, but they do not actively use it; and their daily life is still dominated by private payment platforms such as WeChat. This shows another angle that digital RMB lacks obvious advantages and is unattractive.

Keywords: Digital RMB, WeChat, Perceived usefulness, Perceived ease of use
JEL Classification: E42, G21, D91

Article Information

Received: 21 January 2026

Revised: 11 April 2026

Accepted: 09 May 2026

Published: July 2026

Citation:

Gesheng, C. (2026). Comparison of digital RMB and private payment platform: Gap analysis and improvement. *The Journal of FinTech and Digital Assets*, 1(2), 29–44.

1. Introduction

As private payment platforms, WeChat and AliPay account for more than 90% of China's local currency payments market (Liu et al., 2024a). Under these circumstances, the digital RMB was introduced by the government in reaction to the revolutionary change with the goals of updating its financial system, preserving monetary sovereignty, and utilizing the advantages of digital innovation (Wu et al., 2024). However, five years have passed; most users still insist on using WeChat and Ali pay and reluctant to use the digital RMB because the digital RMB has great gaps compared to WeChat and AliPay in terms of overall functionalities. In previous literature, researchers have explored the factors affecting business intelligence (BI) on the usage of the digital RMB, such as those by Liu et al. (2024a) and Liu et al. (2024b).

However, these studies tend to focus on individual factors and lack a comprehensive comparative analysis to show how differences in these factors affect the BI. Meanwhile, these studies didn't analyse the data in depth, thus lacking measures to improve the use of the digital RMB. To fill the research gap in this area, it is necessary to analyse the various factors from another perspective. In the study, using WeChat as a benchmark, a new model is constructed to identify how differences in these factors affect the intention to use the digital RMB based on the unified theory of acceptance and use of technology (UTAUT), combining with Porter's competition theory.

The results of this study are of great practical significance, not only providing valuable insights to the government and proposing targeted measures but also meeting the public's needs and expectations for the digital RMB. Liu et al. (2024b) noted that social groups such as consumers, merchants, and companies are very interested in the potential usage of the digital RMB to improve payment efficiency and protect data security.

Based on the questionnaire data, this study empirically examines the pathway relationships and influence mechanisms between variables by using the SEM modelling, verifies the proposed research hypothesis, and proposes the optimal pathway. Also, the questionnaire data are categorized by using the K-Means model and analyzed in a multi-dimensional manner, which verifies the results of SEM modelling from another angle. The structure of the study is as follows: section 2 will be the literature review, section 3 devotes to data and methodology, and section 4 presents empirical results and discussion; section 5 concludes the study.

2. Literature Review

Prior studies on BI influencing have assumed that the digital RMB is a new technology and then built a model to analyse relationships between variables. However, although the digital RMB adopts new technologies, its performance relative to WeChat has no distinct advantages. So, in this case, the impact on BI may not be as influential as described in the literature, and the actual situation also shows this point. That is to say, a new technology without a competitive advantage is also unable to attract customers. Evidence based on this assumption can be found in the following literature.

First, Liu et al. (2024b) combined the TAM and IDT models to explore the BI of using the digital RMB, but it didn't mention whether the digital RMB has a significant advantage over the existing solutions. An important factor of IDT, i.e., relative advantage, is missing; also, it can be seen from the design of the questionnaire that the whole scale is centred around the digital RMB without comparison to other payment platforms. Second, Liu et al. (2024a) used the UTAUT model to analyse the factors affecting BI and proposed that comparative advantage is an important mediating factor affecting BI, but the article's research assumed that the digital RMB was built on the cutting-edge blockchain technology, which has an inherent comparative advantage. The assumption that it would surely become the preferred method of payment is problematic.

The following research provides different views into the advantages and disadvantages of the digital RMB and how to increase its use. Shen and Hou (2021) note that the digital RMB is guaranteed by creditworthiness, which eliminates the possibility of bank runs. The digital RMB employs a controlled anonymization mechanism that allows the government to track all transaction records, identifying and preventing fraud and ensuring the security of user accounts. So, the digital RMB is similar to cash, without the risk of commercial bank insolvency, and very safe compared to WeChat (Xia, 2021).

The central bank has been strongly supporting the development of the digital RMB, and part of the issuance needs to be in cooperation with commercial banks. However, Huang and Li (2023) do not mention that there will be problems in the cooperation. Shen and Hou (2021) explain that the digital RMB in the field of mobile payment and the commercial banking industry have a competitive relationship, which also affects the incentives of commercial banks to promote the digital RMB.

Parasol (2022) points out that WeChat and AliPay, the two main private payment platforms, have made China almost a cashless society. They are easy to use, based on QR codes to receive and pay, and there is no need for special equipment like POS. So, the digital RMB must be compatible with private payment platforms, which can also increase the use of digital RMB overseas. It is difficult to change customers' habits, but the digital RMB can be used to motivate consumers to use it by waiving transaction fees, etc. Bai et al. (2025) agree with the view that the digital RMB lacks competitiveness compared to WeChat and AliPay because the digital RMB only has a payment function, which can't be compared to private payment platforms' very tempting eco-payment system. Wu et al. (2024) demonstrate that the collaboration between the digital RMB and private payment platforms can lead to mutually beneficial outcomes for the digital RMB. It should focus on expanding its user base through collaboration and take advantage of government policy and policy drivers to build public trust and acceptance.

From the literature above, it can be seen that the digital RMB may not have much impact on BI in terms of PU and PEOU; users are more accustomed to using WeChat. However, the digital RMB may have advantages in terms of safety and trust as it has the endorsement of the government and does not have the risk of clearing. This is because of continuous support from the government in terms of policy and infrastructure; the digital RMB may be comparable to WeChat regarding CMP and FC, which is a positive boost to BI. The promotion of the digital RMB also needs the active cooperation of commercial banks, which will help improve AW. The users will adopt the digital RMB when they know it exists and believe the digital RMB is superior even if they can't feel it now.

Based on the above analysis and reality, the hypotheses are provided as follows: 1) PU, PEOU, and HAB have no comparative advantage and no positive impact on BI; 2) TR and Safety have a comparative advantage and strong positive impact on BI; 3) CMP and FC have a weak positive impact on BI; 4) AW has a weak positive impact on BI; 5) Tax has a negative impact on BI; 6) Try to explore unknown relationships between variables.

To validate the hypotheses, a new model needs to be built to analyse the impact of each factor on BI, while the model is based on UTAUT theory. UTAUT integrates models such as TAM and TPB, which is one of the most powerful models to explain users' acceptance and use of new technologies (Venkatesh et al., 2003). The advantage of UTAUT is its greater integration and extensibility, which not only accommodates different types of variables, but also allows for the examination of pathway relationships between variables through SEM modelling, while the openness of UTAUT allows for the incorporation of different theories to complement.

A new technology adoption may be affected by competition (Porter and Strategy, 1980). For example, whether the digital RMB is more advantageous than WeChat determines the choice of users. By analysing the competition, it will be better understood how the digital RMB can find its position in the Chinese market. Therefore, the new model constructed using UTAUT and Porter's competition theory is theoretically sound and provides solid support for this study at the empirical level.

3. Data and Methodology

3.1 Data

The research uses a questionnaire survey approach, which consists of three parts: the first part of the description, including the purpose of the survey and those who could participate; the second part contains a demographic survey covering gender, age, occupation, income, and educational background; the third section contains 12 questions that allow for the assessment of 10 latent variables, which are 9 independent variables (PU, PEOU, TR, SAFETY, FC, CMP, HAB, AW, TAX.), which are explained in Figure 1, and 1 dependent variable (BI). All questions are on a 5-point Likert scale (1=completely disagree, 5=completely agree) and can directly measure variables.

Data were collected from 13 Mar 2025 to 9 Apr 2025 through the Questionnaire Star platform. 317 questionnaires were returned, eliminating 3 invalid questionnaires, and finally, we obtained 314 valid samples. The survey was promoted through WeChat channels to cover respondents from different backgrounds. The IPs of all participants are checked, and no duplicate IPs were involved. To alleviate the fatigue of the respondents and improve the efficiency of answering questions, all questions were set in a simple and easy-to-understand manner, and the number of questions was reduced so that the questionnaire could be answered in less than 2 minutes. Although this approach may limit the comprehensiveness of the questionnaire, it helps improve completion rates and data quality. To address this limitation, the questionnaire items were optimized by merging similar questions and removing redundant ones. Meanwhile, all questions were adapted from existing literature to ensure that each variable has a core question to measure. The sample covers participants with diverse demographic characteristics in terms of gender, age, education level, income, and occupation, which enhances the representativeness of the data. However, due to the fact that the sampling method is mainly based on convenience, there may still be potential sampling bias, which is considered one of the limitations of this study.

Furthermore, the sample size of this study exceeded the threshold typically recommended by structural equation models, indicating that a sample size of around 200 is usually considered a common benchmark for reliable estimation (Kline, 2023). In addition, the reliability and validity of the questionnaire were evaluated, and the results confirmed that these data are suitable for further analysis.

In the samples shown in Table 1, 48% are male, and 52% are female. The age group of 18-50 years accounts for 73%, 81% have high school or above, and 74% have a monthly income of more than RMB5000. It is clear that most of the participants are relatively young and educated; they are more likely to understand and accept the new technology. And, they have a certain economic base, which indicates that they can use the new technology. It is noticed that 24.2% of government employees attended the survey, and the digital RMB was first implemented in the government sector, so their feedback is valuable to the research.

3.2 Methodology

After collecting data using a questionnaire, the study empirically analyses the main factors affecting users' adoption of digital RMB versus WeChat, and the methodology includes the following steps:

3.2.1 Tests of the reliability of questionnaires

Internal consistency of assessment scales using Cronbach's Alpha ensures that all questions on the same scale correlate well. DeVellis and Thorpe (2021), in the guidelines for scale development, show that reliability in the range of 0.7 to 0.95 is usually acceptable, indicating that the questionnaire is reliable.

Table 1. Demographic Information

Title	Options	Frequency / Percentage (%)
Gender	Male	152 / 48.41%
	Female	162 / 51.59%
Age	Under 18	0 / —
	Years of age 18–25	71 / 22.61%
	Years of age 26–30	68 / 21.66%
	Years of age 31–40	46 / 14.65%
	Years of age 41–50	44 / 14.01%
	Years of age 51–60	43 / 13.69%
	Years of over age 60	42 / 13.38%
Education	Middle school and below	59 / 18.79%
	High school	47 / 14.97%
	College	73 / 23.25%
	Bachelor	84 / 26.75%
	Postgraduate and above	51 / 16.24%
Income	Less than 5000 yuan	81 / 25.8%
	5000–10,000 yuan	59 / 18.79%
	10,001–20,000 yuan	59 / 20.06%
	20,001–50,000 yuan	59 / 18.15%
	Over 50,000	59 / 17.2%
Occupation	Student	0 / 0.00%
	Employee (Private Sector)	95 / 30.25%
	Employee (Public Sector)	76 / 24.2%
	Self-employed (business owner)	68 / 21.66%
	Freelancer	62 / 19.75%
	Retired	9 / 2.87%
	Other	4 / 1.27%

KMO measure and Bartlett's test are used in evaluating the suitability of data for factor analysis. KMO checks whether the overall correlation of variables is strong enough. The higher the value, the higher the quality of data. While Bartlett's test checks whether the correlation is significant and whether the variables are suitable for factor analysis. DeVellis and Thorpe (2021) note that a KMO value greater than 0.6 is accepted, and Bartlett's test should be significant to proceed, which means a p-value less than 0.05.

The relationships between variables, designed questions, and hypotheses are shown in Figure 1.

UTAUT Variables	Questionnaire (Likert 1-5)	Hypotheses
Perceived Usefulness (PU)	Do you think DIGITAL RMB can improve payment efficiency and reduce transaction costs compared to WeChat Pay?	No comparative advantage and no positive impact on BI
	Do you think DIGITAL RMB is more suitable for various payment scenarios (e.g., cross-border payments) compared to WeChat Pay?	
Perceived Ease of Use (PEOU)	Do you think DIGITAL RMB is easier to understand, and use compared to WeChat Pay?	No comparative advantage and no positive impact on BI
	Do you think DIGITAL RMB is more convenient to use across different devices (e.g., mobile phones, PCs) compared to WeChat Pay?	
Trust (TR)	Compared to WeChat Pay, do you trust the central bank-issued DIGITAL RMB more?	digital RMB has a positive impact to BI
Perceived Security (SAFETY)	Compared to WeChat Pay, do you think DIGITAL RMB accounts are more secure?	digital RMB has a positive impact to BI
Facilitating Conditions (FC)	Do you think DIGITAL RMB can better integrate with existing payment methods (e.g., bank cards, Alipay) compared to WeChat Pay?	No comparative advantage and no positive impact on BI
Compatibility (CMP)	If DIGITAL RMB becomes compatible with WeChat Pay in the future, would you be more willing to use DIGITAL RMB?	No comparative advantage and no positive impact on BI
Habit (HAB)	Compared to WeChat Pay, have you already become accustomed to using DIGITAL RMB for	No comparative advantage and no positive impact on BI
Awareness (AW)	Compared to WeChat Pay, do you feel you have a sufficient understanding of DIGITAL RMB's technical principles and application scenarios?	Positive impact to BI
Tax Concern (TAX)	Compared to WeChat Pay, do you think DIGITAL RMB will reduce your personal tax burden?	Negative impact to BI
Behavioral Intention (BI)	Would you prefer DIGITAL RMB over WeChat Pay if only one payment method exists in the future?	

Figure 1. Variables, Questions, and Hypotheses

3.2.2 Tests of the validity of questionnaires

The Exploratory Factor Analysis (EFA) method is applied to identify the potential factors that affect the willingness to use the digital RMB. EFA helps us check the validity. If the factor loading is generally low, it means that the questions do not measure the latent variables well, and the questionnaire needs to be adjusted. The method is often used in the fields of psychology, sociology, and market research (Field, 2024). According to the recommendations of Hair et al. (2010) and Field (2024), when the factor loading is greater than 0.5, it means the variables significantly reflect the factor; loading between 0.3 and 0.5 indicates that the variable is somewhat representative; lastly, loading less than 0.3 means weak and is recommended to be removed. CR and AVE are used to verify the internal consistency of the factors and aggregated validity, assessing whether the factor is consistently measured by its observed variables. Hair et al. (2010) suggest that CR >0.7 indicates that the factor has good internal consistency, and AVE >0.5 indicates that an acceptable level of convergent validity is achieved.

3.2.3 SEM modelling fit test

The test is to confirm that the SEM modelling truly and reliably reflects the relationships between the variables, in which the overall model is reasonable, and the pathways between the variables support the data itself. Fit indices include degrees of freedom of the Chi-square (χ^2 df), comparative fit index (CFI), Tucker-Lewis index (TLI), and root mean square error of approximation (RMSEA). Hu and Bentler (1999) note that a good model fit is indicated by CFI and TLI being greater than 0.9, RMSEA being less than 0.08, and SRMR being less than 0.

3.2.4 SEM modeling

SEM modeling is used to investigate the relationships between the variables. Because UTAUT variables are latent variables such as PU, PEOU, TR, etc. They can't be measured directly but need to be measured by a scale. While SEM modeling can handle the situation, SEM can also deal with multiple causal pathways at the same time. In contrast, traditional regression can only handle a single causality (Kline, 2023). In the pathway analysis, the strength and direction of the relationships between variables are determined by coefficients (estimates) and the level of significance (p-value). Generally, a p-value < 0.05 indicates that there is a statistically significant relationship between variables, and on how to determine the effect size, Cohen (2013) notes that $[0.1 < \text{estimate} < 0.3]$ is a small effect; $0.3 < \text{estimate} < 0.5$ is a medium effect, and a significant average degree of influence; estimate > 0.5 is very strong and dominant pathways effect.

3.2.5 K-Means Modeling

K-Means can automatically group users based on their scores on each factor, and it can reveal whether there are different groups of users with different perspectives. By grouping users, it can predict which groups are prone to accepting and adopting the digital RMB and which groups need to be persuaded further. Once low-scoring groups have been identified, cross-analysing them with demographic variables (gender, age, occupation, income). It will answer which demographic groups are more resistant to digital RMB, which groups are more concerned about safety and tax, and therefore, the segmentation allows policymakers to develop targeted strategies rather than a one-size-fits-all approach.

4. Empirical Results and Discussion

4.1 The test results of the reliability of the questionnaire

The test results show that Cronbach's Alpha is 0.858, greater than 0.7, indicating that these questions are consistent and reliable. KMO is 0.88, greater than 0.6, indicating that the correlation of variables is strong, being suitable for factor analysis. Bartlett's P value is less than 0.05, indicating that the overall correlation is significant and the correlation is not due to chance.

4.2 The test results of the validity of the questionnaire

Five factors were extracted according to the factor analysis model, and the results are shown in Table 2.

- Factor 1, PEOU and HAB load highly (0.829 and 0.854, respectively); CR/AVE: 0.709.
- Factor 2, CMP loads at 0.71; SAFEFY and FC load at (0.465 and 0.43, respectively); CR/AVE: 0.5.
- Factor 3, AW shows the highest loading (0.99); CR/AVE: 0.981.
- Factor 4, Trust loads highly at 0.879; CR/AVE: 0.772.
- Factor 5, PU loads at 0.732; CR/AVE: 0.536.
- Tax loading does not meet the standard and can be considered for removal.

Table 2. Factor Loading

Factor Loadings:	F1	F2	F3	F4	F5
PU	0.084909	0.131616	0.071017	-0.089235	0.732096
PEOU	0.829464	-0.299725	-0.020896	0.153614	0.211088
Trust	0.049988	0.201821	0.026366	0.878613	-0.092531
SAFETY	-0.021877	0.465125	-0.127829	0.098714	0.223318
FC	0.188473	0.431577	-0.001770	-0.106376	0.113595
CMP	-0.070609	0.707426	0.013216	0.045764	-0.027526
HAB	0.853984	0.168281	-0.048826	-0.081563	-0.188539
AW	-0.017580	-0.009814	0.990500	0.020358	0.056229
TAX	0.365062	0.198051	0.120783	0.011299	0.077356
BI	-0.005759	0.598025	0.061235	0.079368	-0.055864

From the above results, it can be seen that the loadings of PEOU, HAB, AW, and Trust are all greater than 0.5, respectively, indicating these variables can represent factor 1, factor 3, and factor 4 well. The CV and AVE are all greater than 0.7, which indicates that the internal consistency of the factor is good and the aggregation validity has reached an acceptable level. For factor 5, the loading of PU is greater than 0.5, although the CR is 0.536, which is slightly lower than 0.7; the convergent validity is at an acceptable level. For factor 2, although the loadings of SAFETY and FC are slightly lower than 0.5 and AVE just reaches 0.5, the variables of SAFETY and FC can still be taken into account because this study belongs to the field of social psychology. Hair et al. (2010) note that although lower factor loadings may indicate that the variable of the observation is of poorer quality, it does not necessarily imply that the variable is related to the underlying variable, but it does not necessarily mean that the relationship between the variable and the underlying variable is completely invalid. In many cases, lower factor loading still affects the target variable through other measures of the underlying factor or multipath relationships.

Table 3. Explained Variance

Factor	Eigenvalue	Variance Proportion	Cumulative Variance
1	1.60160751	0.16016075	0.16016075
2	1.47621451	0.14762145	0.3077822
3	1.02450635	0.10245063	0.41023284
4	0.84017039	0.08401704	0.49424988
5	0.70043084	0.07004308	0.56429296

As can be seen from Table 3, according to Kaiser's Criterion, if the Eigenvalue is less than 1, the factor should not be retained, but without factor 4 and factor 5, the cumulative explanation rate is only 49.24%, while with them, the cumulative explanation rate can reach 56%. Hair et al. (2010) state that for social science research, a cumulative explanation rate of 50% -60% can be considered as a good factor, so factors 4 and 5 should be retained.

To sum up, the above results show that the factors of this scale have good internal consistency, and their convergent validity is all at an acceptable level, indicating that the scale is well structured and accurately measured.

4.3 SEM Modelling

Starting from the most basic path, verifying that each variable has a direct effect on the BI. The fit test results show Table 4. The χ^2/DoF is close to 0, P-value > 0.05, indicating that the gap between the model and the data is not significant. RMSEA is close to 0, indicating a favourable range of <0.08. Furthermore, the CFI to TLI indexes exceed 0.9, which is a strong indication of exceptional fit. As a result, the SEM model is deemed suitable for analysing the factors of the digital RMB adoption.

Table 4. SEM fit test

	DoF		DoF Baseline	chi2		p-value		chi2 Baseline	CFI
Value	36		44	0.000054		1		919.590165	1.041115
	GFI	AGFI	NFI	TLI	RMSEA	AIC	BIC	LogLik	
Value	1.0	1.0	1.0	1.050252	0	18.0	51.744537	1.715155E-07	

As can be seen from the path diagram after the SEM model run, see below Figure 2, the P-value of AW, CMP, and Trust is less than 0.05, indicating that there is a positive influence on BI. However, the influence of AW and CMP is weaker, while Trust is relatively stronger. SAFETY is equal to 0.05, indicating that it is approaching significance. The P-values of FC, HAB, PU, and PEOU are greater than 0.05, indicating that there is no statistically significant relationship with BI.




Figure 2. Pathway 1

From Pathway 1, PU, PEOU, and HAB have no impact on BI, which is in line with the hypothesis; AW and CMP have a positive impact on BI, but the estimates are 0.144 and 0.176, respectively. This means the impact is weak, which is in line with the hypothesis. Although Trust has a positive impact on the BI with an estimate of 0.223, still less than 0.3, the impact is not as large as hypothesized. The relationship of SAFETY on BI is close to significance, and the estimate is 0.11. Again, it does not meet the hypothesis. Similarly, FC hasn't lived up to its intended hypothesis. Although the paths constructed in this exploration performed well in terms of the overall fit indices, some of the paths do not reach statistical significance, and the coefficients are low. This suggests that the relationships between these variables are weak and have limited practical influence. Therefore, it is necessary to further optimize the model.

Baron and Kenny (1986) propose a classical theory in social science research that even if the direct effect between certain variables is not significant, a relationship may be observed by indirectly influencing the target variable through a mediating variable. Following this line of thought, although PU, PEOU, HAB, and

FC have no significant relationship with BI, there may be a mediating variable that acts as a bridge. Limayem et al. (2007) point out that PEOU and PU are key factors influencing habit formation, especially in the early stages of technology use. They also emphasize that FC reduces the barriers to repeated use of the technology, thus making it easier to form a habit, which, once formed, further influences behavioural intentions. Therefore, based on these theories, a new pathway can be constructed whereby HAB can be used as a mediating variable, and PU, PEOU, and FC influence BI through HAB.

To validate the new pathway, it is necessary to redo the fit test. The fit test results are shown in Table 5. The χ^2/DoF is close to 0.78, the P-value is 0.815(>0.05), RMSEA is close to 0, and the CFI and TLI indices all exceed 0.9, which means it is an exceptional fit, and the SEM model is deemed suitable for analysis.

Table 5. SEM fit test2

	DoF		DoF Baseline	chi2		p-value		Chi2 Baseline	CFI
Value	35		43	27.4355 33		0.81531		919.590171	1.00862 9
	GFI	AGFI	NFI	TLI	RMSEA	AIC	BIC		
Value	0.970165	0.9963 346	0.97016 5	1.01060 2	0	19.825251	57.3191 81		
	Loglik								
Value	0.087374								

After running the SEM modeling, a new pathway 2 diagram can be seen in Figure 3. Compared to pathway 1, the common point is that AW, CMP, SAFETY, and Trust all have a weak positive effect on BI, with SAFETY having a more statistically significant relationship in pathway 2; while the difference is as follows:

- 1) The p-value of HAB is all less than 0.05 in pathway 2, indicating a statistically significant relationship with BI.
- 2) In pathway 1, HAB does not affect BI, but the combined effect of FC, PU, and PEOU makes HAB have a positive effect on BI; the estimates of FC and PU are less than 0.3, the impact is weak, while PEOU's estimate is 0.44 and greater than 0.3, indicating PEOU has a moderate effect on HAB.

Limayem et al. (2007) provide a new way of thinking, and it can be seen that the variables that do not have a direct effect on BI in pathway 1 do not mean that it is invalid. They may have an important impact through indirect paths; from the theoretical point of view, the UTAUT model itself takes into account multiple paths and variable interactions, and in particular with a particular emphasis on the mediating effects. The change in SAFETY is interesting in the process of model reshaping, which may be due to the fact that certain variables are highly correlated with each other. This, if placed in the model at the same time, would create problems such as multicollinearity, leading to a loss of significance in the original path. After reshaping, the other variables are deleted or adjusted, and the model frees up the true role of the variables. Hair Jr et al. (2021) note that some variables mask the role of the other variables, and that deletion or mediation can make the model clearer.




Figure 2. Pathway 2

4.4 K-Means Modelling

After determining the optimal K-value of 5 by the elbow method, the model automatically divides the data into 5 groups; refer to Table 6. The scores of groups 0, 1, and 2 are all lower than 4, so the data of these 3 groups can be put together for further analysis.

These three groups totalled 222 persons, or 71% of all respondents. Out of 222 persons, 51% are female, and 49% are male, which is relatively balanced, indicating that the low acceptance is not limited to a certain gender. And, 74% are between 18 and 50 years old, 80% have a high school education or above, and 76% have a monthly income of more than RMB 5000, which indicates that the group is an educated young group, they are capable of understanding and using digital RMB, and have the ability to pay and consume but they are reluctant to actively use the digital RMB.

Table 6. Clustering

Group	PU	PEOU	Trust	SAFETY	FC
0	3.681373	3.593137	3.470588	3.578431	3.186275
1	3.006667	2.960000	2.906667	2.786667	2.880000
2	1.822222	2.044444	2.133333	2.066667	2.022222
3	4.272152	3.829114	4.012658	3.873418	4.189873
4	2.346154	1.653846	4.692308	4.538462	2.923077
Group	CMP	HAB	AW	TAX	BI
0	3.117647	3.470588	3.637255	3.686275	3.274510
1	2.893333	2.933333	2.813333	2.680000	2.906667
2	1.933333	1.800000	2.000000	2.000000	2.088889
3	4.126582	4.025316	3.848101	3.987342	4.202532
4	4.769231	1.615385	2.307692	2.076923	4.076923

They do not feel that the digital RMB is more convenient, and there is no difference in the way they use it; this group of people has higher expectations of the “new technology”. Interestingly, 23% of government employees also scored low. Their salaries in the pilot area are paid through the digital RMB, so theoretically, they should be the “pioneers”. However, their consumption and money transfers are still dominated by WeChat, so it is not feasible to rely on the government to push the digital RMB without market -driven and

user-perceived benefits. It can be seen that K-Means validates pathway 2 of SEM modelling from another angle in Table 7.

Table 7. Characteristics for the low-score cluster

Feature	Category	Percentage
Gender	Male	49%
	Female	51%
Age Group	18–50 years	74%
	Above 50 years	26%
Education Level	High school or above	80%
	Below high school	20%
Income Level	Monthly income >5000	76%
	Monthly income <5000	24%
Occupation	Company employee	30%
	Gov employee	23%
	Business owner	24%
	Freelancer	19%
	Retirement	3%
	Others	1%

4.5 Discussion on the results

AW has a weak positive impact on BI. Alalwan et al. (2017) state that AW does not necessarily bring a strong perception of relative advantage immediately but rather serves as a sort of preparation to increase the presence of the innovation in the mind of the user. When the market changes or the product is optimized, the user can accept it. Therefore, it's not useless for the government to publicize the digital RMB; users more or less think it is a new trend that can't be ignored and will not feel strange and resistant. This forms a mild positive influence, but whether it translates into action depends on other variables.

For CMP, the digital RMB is currently compatible with WeChat, the interface is friendly and easy to use, can be used for offline code scanning. This indeed lowers the threshold of technology adoption, but it does not form a significant comparative advantage, thus does not materially improve BI. Porter and Strategy (1980) point out that to gain the benefit, the key is to find out the competitor's weaknesses and utilize them to build one's own differentiation. Corresponding to the digital RMB, it is good to take advantage of the fact that WeChat's inability to interoperate with Alipay while the digital RMB can be used as an intermediary and so on.

SAFETY and Trust have a positive impact on BI, but the effect is weaker rather than stronger, which is a bit surprising. Shen and Hou (2021) and Parasol (2022) both state that the government is strengthening the regulation of private payment platforms to deal with any potential risks, and that this proactive approach enhances its trust. On the other hand, it causes the advantages of Trust and SAFETY of digital RMB not to be as obvious as they could be. This explains why there is a deviation from expected assumptions. In addition, Culnan and Bies (2003) note that users are not unconcerned about company-controlled data, but are more sensitive to the consequences of government-controlled data. This confirms the reality that many users trust WeChat more because they care about whether the future use of state-controlled currency is more likely to be regulated and tracked, and that feeling undermines trust in state-controlled currency.

From the path analysis, PU, PEOU, and FC are difficult to form a direct driving force on the BI, but the three variables through HAB have a weak influence on the BI. It can be explained that the government in the past five years has enhanced PU through red packets and social services. Also, it simplifies the interface to enhance PEOU by being compatible with WeChat and Alipay. It also expands the application scenarios and provides technical support to enhance FC, so that users unknowingly begin to get used to using the digital RMB. At present, the ecological construction of digital RMB is still at an early stage; if PU, PEOU, and FC are further optimized, the path strength of HAB is expected to be lifted. Bai et al. (2025) also mention that by the end of May 2024, the total amount of digital RMB transactions has reached 6.6 trillion RMB, and the government has continued to carry out digital RMB on a wider scale, which shows that digital RMB has made great progress and people are gradually accepting and starting to use it in their daily life.

4.6 Improvement of digital RMB

The study provides important practical implications for promoting the digital RMB usage. The next step is to explore how to improve the willingness to adopt the digital RMB along seven dimensions.

Commercial banks have a wide range of service resources and are able to promote it at a low cost. However, due to competitive relationships between the digital RMB and banks' mobile payments, which lead to their low motivation, Policymakers should strengthen cooperation with banks and provide incentives to them to increase AW. WeChat currently has high withdrawal fees and is not interoperable with other private platforms. The digital RMB can fully utilize its advantage of being compatible with them and expand its user base through free or low fees, to increase HAB and CMP.

Digital RMB can provide financial management and points exchange to increase PU. This may have an impact on the bank, but it can cooperate with the bank. The bank provides the underlying services while the digital RMB is only a channel; both sides can realize a win-win situation.

The government could prioritize the digital RMB usage in social insurance and tax subsidies, public healthcare reimbursement, and the use of the digital RMB in the government's daily office system or travel reimbursement to increase PU, FC, and HAB.

WeChat customer service mainly relies on robots, such as handling account blocking and emergencies, which are inefficient. The digital RMB can utilize the strong platform of banks and the government to provide high-quality customer service to enhance FC. Increase application scenarios, encouraging more merchants to access the digital RMB and provide points incentives to increase FC. The registration process can also be simplified in terms of new user registrations with the help of WeChat and Alipay authentication to improve CMP and PEOU.

Beyond its practical significance, this study also offers insights for policymakers. Governments should leverage their strengths by collaborating with financial institutions and promoting the adoption of the digital RMB in the public sector. Through targeted outreach and educational initiatives, along with enhanced regulatory support, public trust can be strengthened. These measures will effectively facilitate the widespread application of the digital RMB.

5. Conclusion

The digital RMB has been promoted for almost 5 years; although progress has been made, its popularity is still low. The factors affecting BI haven't been studied in depth. To fill the research gap, this study constructs a theoretical model combining UTAUT and Porter's competition theory with a new perspective to investigate the effects of the differences in the main factors on the BI of the digital RMB, adopting SEM modelling to explore the optimal path. The model was tested through a survey of 314 digital RMB users. The results show that the proposed theoretical model can effectively explain the related BI. Awareness, HAB, and SAFETY have a weak positive effect on the BI, while CMP and Trust have a low moderate impact on the BI. However, FC, PU, and PEOU do not directly affect the BI but directly affect HAB, then affect the BI through the mediator of HAB. Also, from the results of K-Means modelling, it is clear that without market incentives and the drive of user-perceived benefits, government policies alone can't achieve good results. Based on the above findings, the research puts forward some insightful suggestions to build, improve, and refine the digital RMB, to accelerate the application of digital RMB.

There are some limitations in the study. Firstly, the factors that influence the BI may be more complex, for example, the effect of the tax factor. Although this factor is removed in this study, it does not mean that the factor is invalid. Secondly, there may be other plausible paths that can better explain the BI. Nonetheless, these limitations provide opportunities for future research to deepen the understanding of the topic.

Data Availability Statement

The data supporting the findings of this study are not publicly available due to privacy and ethical considerations related to the survey respondents but are available from the corresponding author upon reasonable request.

Funding

This study received no external funding.

Acknowledgements

I would like to thank the supervisor for his valuable guidance and constructive suggestions throughout the completion of this study.

Declaration of the Use of Generative AI

I used ChatGPT to support brainstorming and preliminary literature searches. The manuscript was independently written by me, and all academic judgments, analyses, and conclusions are solely mine.

Conflicts of Interest

I declare no conflicts of interest.

References

- Alalwan, A. A., Dwivedi, Y. K., & Rana, N. P. (2017). Factors influencing adoption of mobile banking by Jordanian bank customers: Extending UTAUT2 with trust. *International Journal of Information Management*, 37(3), 99–110. <https://doi.org/10.1016/j.ijinfomgt.2017.01.002>
- Bai, H., Cong, L. W., Luo, M., & Xie, P. (2025). Adoption of central bank digital currencies: Initial evidence from China. *Journal of Corporate Finance*, Article 102735. <https://doi.org/10.1016/j.jcorpfin.2024.102735>
- Baron, R. M., & Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6), 1173–1182. <https://doi.org/10.1037/0022-3514.51.6.1173>
- Cohen, J. (2013). *Statistical power analysis for the behavioral sciences*. Routledge.
- Culnan, M. J., & Bies, R. J. (2003). Consumer privacy: Balancing economic and justice considerations. *Journal of Social Issues*, 59(2), 323–342. <https://doi.org/10.1111/1540-4560.00067>
- DeVellis, R. F., & Thorpe, C. T. (2021). *Scale development: Theory and applications* (6th ed.). Sage Publications.
- Field, A. (2024). *Discovering statistics using IBM SPSS Statistics* (6th ed.). Sage Publications.
- Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2010). *Multivariate data analysis: A global perspective* (7th ed.). Pearson.
- Hair Jr., J. F., Hult, G. T. M., Ringle, C. M., Sarstedt, M., Danks, N. P., & Ray, S. (2021). *Partial least squares structural equation modeling (PLS-SEM) using R: A workbook*. Springer Nature.
- Hu, L.-t., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling: A Multidisciplinary Journal*, 6(1), 1–55. <https://doi.org/10.1080/10705519909540118>
- Huang, R. H., & Li, S. X. (2023). China's pursuit of central bank digital currency: Reasons, prospects and implications. *Banking & Finance Law Review*, 39(3), 409–443.
- Kline, R. B. (2023). *Principles and practice of structural equation modeling* (5th ed.). The Guilford Press.
- Limayem, M., Hirt, S. G., & Cheung, C. M. (2007). How habit limits the predictive power of intention: The case of information systems continuance. *MIS Quarterly*, 31(4), 705–737. <https://doi.org/10.2307/25148817>
- Liu, M., Li, R. Y. M., & Deeprasert, J. (2024a). Factors that affect individuals in using digital currency electronic payment in China: SEM and fsQCA approaches. *International Review of Economics & Finance*, 95, Article 103418. <https://doi.org/10.1016/j.iref.2024.103418>
- Liu, X., Wang, Q., Wu, G., & Zhang, C. (2024b). Determinants of individuals' intentions to use central bank digital currency: Evidence from China. *Technology Analysis & Strategic Management*, 36(9), 2213–2227. <https://doi.org/10.1080/09537325.2022.2125948>
- Parasol, M. (2022). China's digital yuan: Reining in Alipay and WeChat Pay. *Banking & Finance Law Review*, 37(2), 265–303.

Porter, M. E. (1980). *Competitive strategy: Techniques for analyzing industries and competitors*. Free Press.

Shen, W., & Hou, L. (2021). China's central bank digital currency and its impacts on monetary policy and payment competition: Game changer or regulatory toolkit? *Computer Law & Security Review*, 41, Article 105577. <https://doi.org/10.1016/j.clsr.2021.105577>

Venkatesh, V., Morris, M. G., Davis, G. B., & Davis, F. D. (2003). User acceptance of information technology: Toward a unified view. *MIS Quarterly*, 27(3), 425–478. <https://doi.org/10.2307/30036540>

Wu, W., Chen, X., Zvarych, R., & Huang, W. (2024). The Stackelberg duel between central bank digital currencies and private payment titans in China. *Technological Forecasting and Social Change*, 200, Article 123169. <https://doi.org/10.1016/j.techfore.2023.123169>

Xia, M. (2021). In search of the perfect coin: China's approach towards cryptocurrency and its own central bank digital currency. *Banking & Finance Law Review*, 36(3), 419–456.

Financial Fraud Detection: Machine Learning vs. Traditional Rule-Based Systems - A Comparative Analysis of Model.

Haixia Du*

Authors

Haixia Du, Independent Researcher,
Auckland, New Zealand.

Email: cassiedu0827@gmail.com

* Corresponding author.

Abstract

With the rapid development and widespread use of electronic payments and digital financial transactions, financial fraud has become increasingly common. It has already caused significant threats and losses to customers and financial institutions. How to detect financial fraudulent activities in a timely and accurate manner has become crucial, as it can not only reduce financial losses but also improves customers' trust in the financial system. Traditional rule-based systems rely on manually defined thresholds and patterns, which can identify anomalous behaviors to a certain extent, but it is difficult to identify complex fraudulent behaviors in the face of continuously changing fraud patterns. In contrast, machine learning provides a dynamic, data-driven approach that is more adaptive and accurate, and it can learn hidden patterns from large datasets. This study uses a simulated credit card transaction dataset generated by the Sparkov tool, trained and tested under the same dataset conditions. Standard pre-processing steps were applied to compare a traditional rule-based system with four machine learning models, including logistic regression, decision trees, random forests and support vector machines (SVM), and to use metrics such as accuracy, precision, recall, F1-score and AUC-ROC to evaluate the model performance. The experimental results show that among all the methods, the random forest model in machine learning performs the best in terms of overall evaluation, showing strong fraud detection ability and low false alarm rate. The decision tree model also performs relatively well with high recall. Although the SVM model has high recall, it falls behind the F1-score and precision. Logistic regression has relatively high recall but very low precision, meaning it produces many false positives. On the other hand, the traditional rule-based system is relatively highly accurate, but its precision and recall are poor, which means it has a high false alarm rate and the actual detection effect is limited. Overall, the results show that machine learning models, especially Random Forests, perform well compared to the traditional rule-based system under the same data conditions. The findings indicate that machine learning models outperform traditional rule-based systems under the same conditions and are more suitable for financial fraud detection in practice.

Keywords: Financial fraud detection, Credit card transactions, Machine learning, Rule-based systems, Random forest

JEL Classification: C45, G21, G28

Article Information

Received: 26 January 2026

Revised: 28 June 2026

Accepted: 10 July 2026

Published: September 2026

Citation:

Du, H. (2026). Financial fraud detection: Machine learning vs traditional rule-based systems - A comparative analysis of model. *The Journal of FinTech and Digital Assets*, 1(2), 45–53.

1. Introduction

Financial fraud issues have become increasingly serious and appear more frequently in daily life. With the growth of different types of fraud, such as credit card theft, forged identities, email scams, and fake transactions, financial fraud not only causes large financial losses to customers and financial institutions but

also creates a crisis of trust in the financial system. Traditional financial systems commonly rely on rule-based fraud detection methods, such as screening transaction behaviors using predefined logical rules. However, the limitations of this approach have become more evident, as it is difficult to quickly and accurately identify fraud in the presence of complex and continuously evolving fraudulent behaviors (West and Bhattacharya, 2016).

Machine learning models provide an alternative to traditional rule-based systems. Unlike rule-based approaches, machine learning models can automatically learn fraud patterns from large volumes of historical transaction data and demonstrate stronger adaptability and predictive capability (Ngai et al., 2011). Many studies have shown that machine learning models achieve higher precision and recall in fraud detection. However, after reviewing existing literature, it can be observed that there is still a lack of experimental research that systematically compares traditional rule-based systems with different machine learning models under the same dataset and evaluation conditions.

This study aims to address this gap by using a publicly available simulated credit card transaction dataset. Under the same dataset and evaluation framework, a traditional rule-based system is compared with several machine learning models, including logistic regression, decision tree, random forest, and support vector machine. The models are evaluated using accuracy, precision, recall, F1-score, and AUC, to identify more effective and adaptive approaches for financial fraud detection.

The remainder of this paper is structured as follows. Section 2 reviews the existing literature on financial fraud detection. Section 3 describes the dataset. Section 4 outlines the methodology. Section 5 presents empirical results. Section 6 concludes the study.

2. Literature Review

Financial fraud has long been one of the key topics studied in the field of financial learning. Early fraud detection systems mainly relied on traditional rule-based methods, which require predefined thresholds to identify suspicious transactions. However, with constantly evolving fraud techniques, traditional rule-based detection methods often lack flexibility and practical applicability.

West and Bhattacharya (2016) conducted a comprehensive review of traditional rule-based systems and pointed out that these detection systems exhibit low responsiveness and limited adaptability. They also emphasized the advantages and potential of machine learning approaches in dealing with dynamic fraud patterns. This perspective provides important theoretical support for the present study. Ngai et al. (2011) proposed an operational classification framework to systematically organize the application of data mining techniques in financial fraud detection. They classified methods such as decision trees, neural networks, and support vector machines, which provided a theoretical basis for algorithm selection and background for this research study.

Bhattacharyya et al. (2011) analyzed the performance of multiple supervised learning models using real-world credit card transaction data. Their results showed that non-linear models typically outperform linear models in fraud detection tasks, which guided this study to prioritize non-linear models such as decision trees and random forests. Ryman-Tubb et al. (2018), by investigating the practical application of artificial intelligence in payment card fraud detection, identified several key challenges in current industry practices, including model interpretability, real-time responsiveness, and system integration. Their findings indicate that model accuracy alone is insufficient, and feasibility in real systems must also be considered for practical deployment.

Al-Hashedi and Magalingam (2021) provided an overview of data mining methods applied to financial fraud detection between 2009 and 2019, summarizing a wide range of mainstream fraud detection approaches. Riskiyadi (2024) and Ramzan and Lokanan (2024) demonstrated the broad applicability of machine learning methods by applying them to financial statement fraud and accounting fraud detection, respectively. However, these studies largely lacked experimental designs that directly compare machine learning approaches with traditional methods. This limitation highlights a current research gap and aligns with the objective of the present study.

A frequently discussed issue identified in the literature is data imbalance. As fraudulent transactions usually account for only a small proportion of total transaction volumes, predictive models tend to favor non-fraudulent cases. Carcillo et al. (2021) proposed a hybrid approach combining unsupervised and supervised learning to address data imbalance and improve model robustness on highly imbalanced datasets. This literature also motivated the use of the Synthetic Minority Oversampling Technique (SMOTE) to oversample the dataset to enhance the model's ability to detect fraudulent behavior.

Afriyie et al. (2023) validated the application of supervised machine learning algorithms in real-time credit card fraud detection and provided useful references for selecting appropriate model evaluation metrics. Stojanović et al. (2021) focused on the strategic application of machine learning for fraud detection in the FinTech domain and emphasized the importance of model performance in real-world applications. This further reinforces the motivation to conduct multi-model comparative experiments, particularly to identify a balance between deplorability and performance.

Although existing studies provide strong theoretical foundations and guidance for methodological selection, most research evaluates only a single machine learning model or conducts limited comparisons. Systematic comparative studies between traditional rule-based systems and machine learning methods under the same dataset and evaluation conditions remain limited. Therefore, this study aims to address this gap by comparing multiple approaches to better understand their respective strengths and weaknesses and to identify more effective solutions for financial fraud detection.

3. Data and Methodology

Before presenting the dataset and experimental procedures, the basic principles of the fraud detection methods are briefly introduced. This study compares a traditional rule-based system with four supervised machine learning algorithms, including logistic regression, decision tree, random forest, and support vector machine (SVM). These methods were selected because they represent different approaches to fraud detection and allow a comparison of their performance under the same experimental conditions (Ngai et al., 2011; Bhattacharyya et al., 2011).

Logistic regression is a supervised learning algorithm widely used for binary classification problems. It estimates the probability that a transaction is fraudulent by using the logistic (sigmoid) function. The probability is calculated as:

$$P(y = 1|x) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 x_1 + \dots + \beta_n x_n)}} \quad (1)$$

where $P(y=1|x)$ represents the probability that a transaction is fraudulent, $x_1, \dots, x_n$ are the predictor variables, and $\beta_0, \beta_1, \dots, \beta_n$ are the model coefficients. If the predicted probability is greater than a predefined threshold, the transaction is classified as fraudulent. According to Ngai et al. (2011), logistic regression has a simple structure and is easy to interpret. It is therefore commonly used as a baseline model for comparison with other machine learning algorithms.

A decision tree is a nonlinear supervised learning algorithm that classifies data by splitting the dataset based on selected feature values. During model construction, the Gini index is commonly used to evaluate the quality of each split. The Gini index is calculated as:

$$Gini = 1 - \sum_{i=1}^C p_i^2 \quad (2)$$

where p_i represents the proportion of observations belonging to class i . The process continues until a predefined stopping criterion is reached. A decision tree is easy to understand because its classification process can be represented as a series of decision rules. However, when applied to complex datasets, decision trees may become overly complex and are more likely to overfit the data (Bhattacharyya et al., 2011).

Random forest is an ensemble learning algorithm based on decision trees. It constructs multiple decision trees using bootstrap sampling and randomly selected feature subsets. The final classification result is determined by majority voting. Compared with a single decision tree, random forest provides more stable predictions, reduces the risk of overfitting, and achieves better classification performance. Therefore, random forest is widely used in financial fraud detection because of its strong classification performance (Bhattacharyya et al., 2011).

Support vector machine (SVM) is a supervised learning algorithm that classifies data by finding the optimal separating hyperplane. For more complex classification problems, kernel functions can be used to map the data into a higher-dimensional feature space and construct a nonlinear decision boundary. SVM has strong classification capability. However, SVM has a high computational cost when processing large datasets, which may limit its application in some situations (Ngai et al., 2011).

Unlike machine learning models, traditional rule-based systems rely on manually defined IF–THEN rules and predefined thresholds to identify suspicious transactions. These rules are usually developed based on expert knowledge and historical fraud cases. When a transaction meets the predefined conditions, it is flagged for further investigation. Although rule-based systems are simple to implement and easy to understand, they cannot automatically learn from historical data and therefore have limited ability to adapt to new fraud patterns. West and Bhattacharyya (2016) pointed out that the performance of rule-based systems gradually declines as fraud methods continue to evolve.

The following section describes the dataset and experimental procedures used in this study. The dataset used in this study is obtained from a fraud detection benchmark provided by Amazon Science. It is a simulated credit card transaction dataset generated using the Sparkov tool, designed to reflect real-world transaction behaviour with a clear and well-defined structure. The dataset captures key behavioural characteristics of credit card transactions while avoiding the disclosure of sensitive user information, making it suitable for training and testing financial fraud detection models.

The dataset is pre-divided into two subsets: *fraudTrain.csv*, which is used for model training, and *fraudTest.csv*, which is used for model testing and performance evaluation. Each transaction contains multiple feature variables, such as transaction amount, merchant category, and user gender, as well as a binary target variable, *is_fraud*, indicating whether a transaction is fraudulent (1) or non-fraudulent (0). This predefined data split facilitates model development and ensures consistent performance comparison across different algorithms. Both subsets preserve the original distribution of fraudulent and non-fraudulent transactions.

Several factors motivated the selection of this simulated dataset. First, it closely mimics real-world credit card transaction behaviour. Second, it avoids issues related to customer privacy and data confidentiality. In addition, the use of a publicly available dataset enhances the reproducibility of this research, enabling comparisons with existing studies and supporting future academic work.

This study adopts a quantitative research methodology based on secondary data and incorporates binary classification supervised learning models to systematically evaluate the performance of various fraud detection approaches. A descriptive comparative research design is employed to compare the performance of traditional rule-based detection methods with multiple machine learning algorithms under the same dataset and experimental conditions, with the aim of identifying more reliable and adaptable solutions for real-world financial fraud detection.

The sample used in this study is derived from the simulated transaction dataset provided by Amazon Science, which contains approximately 200,000 transaction records covering both normal and fraudulent behaviour. Each transaction record is independent, making the dataset suitable for supervised learning modelling.

Prior to model training, standard data preprocessing procedures were applied, including handling missing and duplicate values, encoding categorical variables into numerical representations, and standardising continuous variables such as transaction amounts. As fraudulent transactions represent only a small proportion of the dataset, a severe class imbalance problem exists. To address this issue, the Synthetic Minority Oversampling Technique (SMOTE) was applied to the training set to synthetically generate minority class samples, aiming to achieve a 1:1 ratio between fraudulent and non-fraudulent transactions. As a result, the training dataset was expanded to approximately 2.56 million samples, reducing bias toward the majority class and improving the model's ability to identify fraudulent behaviour.

This study constructs and compares a traditional rule-based detection system and several machine learning models. The selected machine learning algorithms include logistic regression, decision tree, random forest, and support vector machine. All models are trained on the same training dataset and tested on the same testing dataset to ensure comparability of evaluation results.

Model performance is evaluated using consistent metrics, including accuracy, precision, recall, F1-score, and the area under the ROC curve (AUC-ROC). Performance assessment considers overall metrics such as accuracy and AUC, while placing particular emphasis on fraud-sensitive metrics such as recall and F1-score.

As this study uses a publicly available simulated dataset, no field surveys or participant recruitment are required. This reduces data collection time and allows the research to focus on model construction and performance evaluation. The dataset is structurally complete, contains rich features, and does not involve real user information or sensitive data, ensuring compliance with legal and ethical requirements.

The main challenge encountered during the experiment relates to the training efficiency of the support vector machine (SVM) model. In particular, SVM with an RBF kernel requires substantial computational resources. Training the model on the full dataset resulted in memory overflow and program crashes in the Google Colab environment. To address this issue, the sample size was reduced, and different cross-validation strategies were tested. The final approach involved using a subset of 30,000 samples combined with 10-fold cross-validation. Consequently, all models were trained on the same 30,000-sample subset to maintain comparability while addressing computational constraints. The results obtained from this subset were stable and reproducible (random_state = 42), indicating that it was both representative of fraud patterns and computationally feasible.

Overall, the dataset and methodology provide a solid foundation for evaluating fraud detection systems and ensure consistent experimental conditions for comparing different fraud detection methods.

4. Empirical Results and Discussion

To evaluate the effectiveness of traditional rule-based systems and machine learning-based fraud detection methods, this study conducts a systematic comparison of five models: one traditional rule-based system and four supervised learning algorithms, namely logistic regression, decision tree, random forest, and support vector machine. All models are evaluated under consistent experimental conditions using the same test dataset and performance metrics, including accuracy, precision, recall, F1-score, and AUC-ROC, as demonstrated in Figures 1 and 2.

The random forest model achieves the best overall performance across all evaluation metrics. It records the highest recall, F1-score, and AUC, indicating strong capability in identifying fraudulent transactions while maintaining a relatively low false alarm rate. This result suggests that the random forest model provides the most balanced and practical fraud detection performance among the evaluated approaches.

The decision tree model also performs well, achieving high recall and good accuracy. However, its F1-score is lower than that of the random forest model, and its precision is slightly weaker, indicating a less balanced performance and a higher likelihood of false positives. Despite this limitation, its relatively simple structure makes it a potentially interpretable solution in practical applications.

The support vector machine (SVM) model performs well in terms of recall and AUC, demonstrating strong fraud detection capability. However, its precision and F1-score are relatively low, suggesting that although it identifies many fraudulent transactions, it also produces a higher false alarm rate, which may limit its effectiveness in real-world deployment. In addition, due to the high computational requirements of SVM for large datasets, the model initially crashed or timed out during training. By limiting the sample size to 30,000 and applying 10-fold cross-validation, stable and reproducible results were obtained. Although the performance is acceptable under these conditions, the cost-performance ratio of SVM remains inferior to that of the random forest model.

The logistic regression model performs the worst among the four machine learning methods. Although its recall is relatively high, its precision is only 2.29%, resulting in an F1-score of 0.044. This indicates that the model classifies nearly all transactions as fraudulent, leading to an extremely high false positive rate and very low practical usability.

The traditional rule-based system achieves high overall accuracy but exhibits very low recall and the lowest F1-score among all evaluated models. This highlights its limitations in fraud detection and reflects its lack of adaptability to complex and diverse fraud patterns, as fraudulent behaviours often deviate from static predefined rules.

Overall, the results demonstrate that machine learning models significantly outperform traditional rule-based systems in financial fraud detection. Among all evaluated methods, the random forest model shows the best balance between precision and recall, providing strong empirical support for its application in practical fraud detection systems. These findings also offer useful insights for future research exploring more advanced approaches, such as deep learning or hybrid detection models.

	Model	Accuracy	Precision	Recall	F1	AUC
1	LogisticRegression	0.876567	0.022890	0.743124	0.044412	0.908218
2	DecisionTreeClassifier	0.975903	0.131375	0.934266	0.230358	0.955165
3	RandomForestClassifier	0.984935	0.196688	0.941259	0.325383	0.994853
4	SupportVectorClassifier	0.946604	0.063269	0.929604	0.118475	0.980979
5	Rule-Based System	0.946604	0.063269	0.929604	0.118475	0.938137

Figure 1. Model comparison (created by the author, 20/06/2025)




Figure 2. Comparison of model performance metrics (created by the author, 20/06/2025)

5. Conclusion

This study focuses on financial fraud detection and conducts a comprehensive comparative analysis of a traditional rule-based system and four supervised machine learning models, namely logistic regression, decision tree, random forest, and support vector machine, under consistent data conditions. A publicly available and well-structured dataset is employed, and all models are trained and tested under identical experimental settings. Multiple evaluation metrics, including accuracy, precision, recall, F1-score, and AUC, are used to ensure fairness and reliability in performance comparison.

The results clearly demonstrate that machine learning models significantly outperform the traditional rule-based system across key performance metrics, particularly in terms of fraud detection capability and false positive control. Among the evaluated models, the random forest model achieves the best overall performance. It accurately identifies most fraudulent transactions while effectively reducing false positives, making it highly suitable for real-world applications. The decision tree model, with its relatively simple structure, also performs well in terms of recall, which makes it a feasible option in scenarios where interpretability is a primary

requirement. The support vector machine shows stable performance and strong recognition ability; however, it requires substantial computational resources and is therefore less suitable for large-scale deployment in environments with limited capacity. Although acceptable results are achieved through sample down-sampling and cross-validation, deployment cost remains an important constraint.

Compared with machine learning models, the overall detection capability of the traditional rule-based system is relatively weak. Although it achieves high accuracy, it exhibits low recall and F1-score, indicating a high rate of missed fraud cases. This highlights the difficulty of relying on static rules to address the diversity and variability of modern financial fraud patterns. In comparison, traditional rule-based systems are more appropriate as preliminary screening mechanisms rather than primary fraud detection tools.

Overall, integrating machine learning models into financial fraud detection systems can bring significant benefits in terms of improved accuracy, flexibility, and adaptability. To enhance overall detection performance, financial institutions are encouraged to move beyond static rule-based systems and actively explore data-driven and intelligence-based solutions. Future research may extend to areas such as the introduction of deep learning models, real-time streaming data processing, or the development of hybrid approaches that combine rule-based logic with machine learning techniques. When selecting detection models, it is also essential to consider the balance between performance, resource constraints, and real-world deployment scenarios, in order to avoid one-size-fits-all strategies. Through continuous optimisation of algorithms and system architectures, financial fraud detection systems can become more resilient and better equipped to address increasingly complex and dynamic fraud challenges.

Funding

This research received no external funding.

Acknowledgements

The author would like to express sincere gratitude to Dr. David Gabauer and Dr. Cuong Nguyen for their guidance, continuous support, and valuable feedback throughout this research. The author also thanks the anonymous reviewers for their constructive comments and insightful suggestions, which have greatly improved the quality of this manuscript.

Declaration of the Use of Generative AI

Generative AI tools were used solely to improve the language, grammar, and readability of this manuscript. The research design, data analysis, interpretation of the results, and all academic decisions and final decisions related to this manuscript were undertaken solely by the author. The author takes full responsibility for the accuracy, originality, and integrity of this manuscript.

Conflict of Interest

The author declares no conflict of interest.

References

- Afriyie, J. K., Tawiah, K., Pels, W. A., Addai-Henne, S., Dwamena, H. A., Owiredue, E. O., Ayeh, S. A., and Eshun, J. (2023). A supervised machine learning algorithm for detecting and predicting fraud in credit card transactions. *Decision Analytics Journal*, 6:100163.
- Al-Hashedi, K. G. and Magalingam, P. (2021). Financial fraud detection applying data mining techniques: A comprehensive review from 2009 to 2019. *Computer Science Review*, 40:100402.
- Bhattacharyya, S., Jha, S., Tharakunnel, K., and Westland, J. C. (2011). Data mining for credit card fraud: A comparative study. *Decision Support Systems*, 50(3):602–613.
- Carcillo, F., Le Borgne, Y.-A., Caelen, O., Kessaci, Y., Oblé, F., and Bontempi, G. (2021). Combining unsupervised and supervised learning in credit card fraud detection. *Information Sciences*, 557:317–331.
- Ngai, E. W., Hu, Y., Wong, Y. H., Chen, Y., and Sun, X. (2011). The application of data mining techniques in financial fraud detection: A classification framework and an academic review of literature. *Decision Support Systems*, 50(3):559–569.
- Ramzan, S. and Lokanan, M. (2024). The application of machine learning to study fraud in the accounting literature. *Journal of Accounting Literature*.
- Riskiyadi, M. (2024). Detecting future financial statement fraud using a machine learning model in indonesia: a comparative study. *Asian Review of Accounting*, 32(3):394–422.
- Ryman-Tubb, N. F., Krause, P., and Garn, W. (2018). How artificial intelligence and machine learning research impacts payment card fraud detection: A survey and industry benchmark. *Engineering Applications of Artificial Intelligence*, 76:130–157.
- Stojanović, B., Božić, J., Hofer-Schmitz, K., Nahrgang, K., Weber, A., Badii, A., Sundaram, M., Jordan, E., and Runevic, J. (2021). Follow the trail: Machine learning for fraud detection in fintech applications. *Sensors*, 21(5):1594.
- West, J. and Bhattacharya, M. (2016). Intelligent financial fraud detection: a comprehensive review. *Computers & Security*, 57:47–66.

Two-Regime Liquidity Recovery After a Perpetual Futures Liquidation Cascade: Evidence from Hyperliquid and the October-10th-2025 Event.

Lim Boon Chuan*

Authors

Lim Boon Chuan, Independent

Researcher, Singapore.

Email: boonchuan@singapore.to

* Corresponding author.

Abstract

We document the inside-market response of Hyperliquid - the largest decentralized perpetual futures venue - to the October 10, 2025 system-wide liquidation cascade. Using the venue's public L2 order-book archive, we construct a minute-level panel covering BTC, ETH, SOL, XRP, and DOGE perpetual contracts across a 92-day window. There are three main findings. First, the cascade was a tightly compressed intraday event. Median quoted spread pooled across assets reached 9.4 basis points at hour 21 UTC, and the 95th percentile reached 87.5 bps, with intraday recovery well under way by end of day. Second, across the 23-day post-event window prior to a separate market-wide selloff, quoted spreads on all five contracts returned to within 1.2 times pre-event levels, and three of five to within 1.1 times. Third, inside-quote depth contracted persistently over the same window, with four of five contracts ending between 66 and 79 percent of pre-event levels. We interpret these results as two-regime recovery, in which pricing recovers while inventory commitment persistently contracts. Spread-only metrics systematically understate post-event impairment of venue liquidity.

Keywords: Decentralized exchange, Market microstructure, Perpetual futures, Liquidation cascades, Market resilience

JEL Classification: G14, G15, G23

Article Information

Received: 24 April 2026

Revised: 16 June 2026

Accepted: 20 June 2026

Published: September 2026

Citation:

Lim, B. C. (2026). Two-regime liquidity recovery after a perpetual futures liquidation cascade: Evidence from Hyperliquid and the October-10th-2025 event. *The Journal of FinTech and Digital Assets*, 1(2), 54–68.

1. Introduction

Cryptocurrency derivatives have grown from a marginal segment into the largest and most actively traded part of the digital-asset market. Trading in these instruments now routinely exceeds spot trading in notional volume, and the segment spans futures, options, and perpetual contracts listed across both centralized and decentralized venues. Among these instrument types, the perpetual future - a futures-style contract with no expiry, held in line with the underlying through a periodic funding payment - has become the dominant vehicle for leveraged exposure to cryptocurrencies. Against this backdrop, we focus on the perpetual-futures segment. Perpetual futures are the dominant derivative class in cryptocurrency markets, with global notional volumes that have exceeded sixty trillion US dollars annually. Until recently, all meaningful liquidity in this segment resided on centralized exchanges. Hyperliquid, launched in 2023 and scaled materially through 2024 and 2025, is the first decentralized venue to attain comparable order-book liquidity. Its market share of

decentralized perpetual volume reached approximately 60-65 percent by mid-2025, and its single-venue daily volume regularly exceeds five billion US dollars. The protocol's distinguishing feature is a consensus-layer ordering rule. Within each block, the validator software is required to process cancel orders and post-only limit orders before processing good-til-cancel and immediate-or-cancel orders. The rationale is grounded in microstructure theory. When prices move, market makers want to cancel stale quotes, and if a faster taker is processed first the maker is picked off. Cancel prioritization eliminates this specific source of adverse selection on resting orders by giving the maker a structural ordering advantage.

On October 10, 2025, the cryptocurrency derivatives market experienced a system-wide liquidation cascade. Across centralized and decentralized perpetual venues, leveraged positions were liquidated at unusual scale; aggregate liquidation volume exceeded nineteen billion US dollars within a 24-hour window. The cascade has been the subject of subsequent industry commentary and a formal analysis of the protocol's automated deleveraging mechanism (Chitra, 2025). What has not been documented is the inside-market response on Hyperliquid in detail - how quickly spreads widened, how quickly they recovered, and whether the recovery was symmetric between quoted prices and posted sizes.

This paper provides that documentation. Using Hyperliquid's public L2 order-book archive, we construct a minute-level panel of inside-market metrics for the five most-traded perpetual contracts (BTC, ETH, SOL, XRP, and DOGE) across a 92-day window centered on the October 10 event. The dataset comprises 531,305-minute observations. Because the entire analysis is venue-internal, all comparisons are against the venue's own pre-event baseline; no cross-venue spread inference is required.

There are three main findings obtained from the analysis. First, the cascade was a tightly compressed intraday event. Hour-by-hour analysis of October 10 UTC shows the inside spread remaining at normal levels (0.37 to 0.44 basis points pooled across all five assets) through hour 14, beginning to widen at hour 15, and peaking at hour 21 (median 9.38 bps, p95 87.48 bps, maximum 241.8 bps). By hour 23, median spread had declined to 1.08 bps. The bulk of the venue's stress event was contained within a three-hour window.

Second, spread recovery across the 23-day window between the cascade and a separate market-wide selloff beginning November 3, 2025 was close to complete. By November 2, 2025, daily median spreads returned to between 1.06- and 1.19-times pre-event baseline, with three of five contracts within 1.10 times. A subsequent market-wide selloff pushed spread trajectories back into the 1.3 to 1.5 range; we treat this second event explicitly in Section 4 and demonstrate that the clean-window recovery finding is robust to its exclusion.

Third, inside-quote depth contracted persistently over the same clean window, even as spreads recovered. Four of five contracts showed further depth contraction from the first week to the third week post-event, with end-of-window depth ratios between 0.66 and 0.79 of pre-event levels. The depth contraction preceded the November selloff and is therefore attributable to the cascade itself.

We term this pattern two-regime recovery, in which pricing recovers quickly and nearly completely while inventory commitment contracts persistently. The pattern is consistent with the inventory-management framework (Ho & Stoll, 1981; Stoll, 1978), in which a tail event raises perceived inventory risk and the rational maker response is size reduction rather than price markup. For execution-quality measurement, the implication is that spread-only metrics systematically understate post-event impairment of venue liquidity.

These findings contribute to two literatures. First, to the empirical microstructure of decentralized derivatives venues, which until recently has been data constrained. The public archive that Hyperliquid publishes makes inside-market analysis at this granularity feasible from outside the venue, at minimal cost. Second, to the literature on market resilience after stress events, where almost all empirical work has focused on equity venues (Anand & Venkataraman, 2016; Bessembinder et al., 2015; Brogaard et al., 2018). We provide what we believe is the first detailed inside-market resilience study on a major perpetual futures venue, centralized or decentralized.

The remainder of the paper proceeds as follows. Section 2 reviews related literature. Section 3 describes the institutional setting, data, and methodology. Section 4 presents empirical results and discussion, including the cascade event study, recovery dynamics, depth contraction findings, and robustness analyses. Section 5 concludes.

2. Literature Review

2.1. Cryptocurrency Market Microstructure

The microstructure of cryptocurrency markets has been an active research area since the emergence of derivatives trading on these assets. Foley et al. (2019) document order-flow dynamics on early Bitcoin spot exchanges. Brauneis et al. (2021) provide a comprehensive treatment of liquidity measurement in cryptocurrency markets, comparing transaction-cost proxies against high-frequency benchmarks across major exchanges. He et al. (2024) study perpetual-spot arbitrage and the associated execution costs. Ruan and Streltsov (2024) analyze how the introduction and removal of perpetual contracts affect spot-market liquidity and price efficiency. None of these papers analyze inside-market behavior on a decentralized perpetual venue, which is the focus of our work.

On-chain venues introduce distinct microstructure features absent from centralized exchanges, namely deterministic block-level execution, validator co-location, and protocol-defined order priorities. Chitra (2025) provide the first formal model of automated deleveraging mechanisms for perpetual exchanges and analyze them empirically using the Hyperliquid October 10 cascade. Their work establishes a trilemma between solvency, revenue, and trader fairness in the design of forced-liquidation mechanisms. The present paper is complementary, analyzing the inside-market response to the same event from the perspective of voluntary market-maker behavior rather than protocol-enforced liquidations.

2.2. Market Resilience After Stress Events

A substantial literature studies how equity markets respond to and recover from stress events. Brogaard et al. (2018) document liquidity dynamics around extreme price moves on US equity venues, finding that high-frequency traders supply liquidity during ordinary stress but withdraw during extreme stress. Anand and Venkataraman (2016) study recovery dynamics after market-wide circuit-breaker triggers. Bessembinder et al. (2015) examine the persistence of spread widening following corporate-event-driven volatility. The shared finding across these papers is that recovery is asymmetric. Spreads widen quickly in stress and narrow more slowly, with the slowness driven by maker risk-aversion rather than by underlying volatility persistence.

A common limitation of this literature is that it focuses almost exclusively on the spread dimension of execution quality, with depth treated as a secondary consideration or omitted entirely. The two-regime finding

documented here suggests this focus misses the more persistent dimension of post-event impairment in at least one class of venue. To our knowledge no paper in this literature has examined a perpetual futures venue, and no paper has examined a decentralized venue. The Hyperliquid October 10 cascade is well-suited to extending this literature because it occurred on a venue with continuous (24/7) trading and no circuit breakers, providing a clean observation of maker-driven recovery dynamics absent venue-imposed interventions.

2.3. Inventory Management and Maker Responses to Risk

The theoretical foundation for the two-regime interpretation is the inventory-management framework of Stoll (1978) and Ho and Stoll (1981), in which market makers set quotes as a function of both expected adverse selection and the cost of holding inventory exposed to price risk. In this framework, a tail event that raises the perceived variance of inventory can rationally produce a separable response, in which competitive quotes are maintained to capture order flow while the size at those quotes is reduced to limit inventory accumulation. The two-regime pattern documented here is consistent with this model. The literature has previously documented the model's predictions in other settings (Hasbrouck & Sofianos, 1993; Madhavan & Smidt, 1991), but we are not aware of a direct decomposition of post-event recovery into pricing and inventory dimensions separately for a major perpetual futures venue.

2.4. Cancel Prioritization and Maker Protection

The literature on order-priority rules is long but has been relatively quiet on cancel prioritization specifically. Hendershott and Moulton (2011) document that the introduction of NYSE's Hybrid Market — which reduced the time-priority advantage of slow-moving liquidity providers — measurably worsened execution quality for retail-sized orders, suggesting that protections for resting orders matter for inside-market quality. Foucault et al. (2013) provide a theoretical treatment of cancel-related fees and their effect on quote-update behavior. Hyperliquid's design implements a particularly aggressive form of cancel prioritization, under which all cancels are processed at the consensus layer before any matching, regardless of submission time. To our knowledge no centralized exchange operates this rule in its full form.

3. Data and Methodology

3.1. Institutional Setting: Hyperliquid

Hyperliquid is a layer-one blockchain whose primary application is a perpetual-futures decentralized exchange. The protocol uses HyperBFT, a modified HotStuff consensus algorithm with sub-second block times. All order-book operations - placement, modification, cancellation, and matching - execute on-chain within consensus blocks. Validators are tightly co-located, which keeps consensus latency low enough to support order-book trading workflows familiar from centralized exchanges. Trading is fee-funded rather than gas-funded, eliminating one source of execution variance present on other on-chain venues.

The protocol's matching rules differ from a standard continuous limit order book in one substantive way. Within each block, the validator software processes cancel orders and post-only limit orders before processing good-till-cancel and immediate-or-cancel orders. When the underlying spot price moves, a market maker with resting quotes faces an immediate adverse-selection threat, since any taker arriving in the same block is more likely to be informed about the price move than the average trader. The maker's natural response is to cancel and re-quote. Without cancel prioritization, a sufficiently fast taker can submit a marketable order

in the same block as the maker's cancel, and have it processed first - the maker is picked off at a stale price. Cancel prioritization eliminates this specific failure mode.

The Hyperliquid Foundation publishes a daily archive of L2 order-book snapshots to a public S3 bucket (Hyperliquid Foundation, 2024). As measured in our sample, snapshots are produced at a median inter-snapshot gap of 0.543 seconds, yielding approximately 110 snapshots per asset per minute. Each snapshot contains the top 20 levels on each side of the book. Trade data are available separately via the Hyperliquid API but are not included in the public S3 archive; the analysis in this paper uses only book-snapshot data and does not require trade-level inputs.

3.2. The October 10 Cascade and the November 3-4 Selloff

On October 10, 2025, the cryptocurrency perpetual futures market experienced a system-wide liquidation cascade. Across centralized and decentralized venues, leveraged positions were forcibly closed at scale; aggregate liquidation volume across all major venues exceeded nineteen billion US dollars within a 24-hour window. Hyperliquid's automated deleveraging mechanism was triggered for several contracts, and the venue's inside market widened sharply for a multi-hour window beginning approximately 21:00 UTC.

Twenty-five days after the cascade, on November 3-4, 2025, the market experienced a smaller but still substantial selloff. Bitcoin declined from approximately \$107,000 on November 1 to \$99,000-\$104,000 on November 4, with daily moves in excess of five percent and renewed liquidations across major perpetual venues. The November event is an order of magnitude smaller than October 10 in absolute liquidation volume but is large enough to produce visible widening in our post-event recovery trajectory. We treat the cascade as the primary event and the November selloff as a confound to be explicitly controlled for in the empirical analysis.

3.3. Data and Sample Construction

Our sample covers September 15 through December 15, 2025, a 92-day window centered on the cascade event. We restrict the asset universe to the five perpetual contracts with continuous existence on Hyperliquid throughout the sample and with sufficient daily volume to support meaningful inside-market inference, namely BTC, ETH, SOL, XRP, and DOGE.

We obtain Hyperliquid L2 order-book snapshots from the public archive bucket. The bucket is requester-pays, meaning the data are free at source but the requester bears the egress cost; total data-transfer cost for our sample was approximately ten US dollars. Snapshots are stored as line-delimited JSON within hourly LZ4-compressed files. We extracted, parsed, and re-stored each asset-day as a Parquet file, retaining the original timestamps and the full twenty-level book on each side.

From each L2 snapshot we compute the quoted spread at the inside (in basis points relative to the midpoint), the inside-quote depth in US dollars (the dollar size resting at the best bid plus the dollar size resting at the best ask), and the simulated round-trip execution cost for marketable orders at \$50K, \$100K, and \$500K notional sizes. Realized volatility is computed as the annualized standard deviation of one-second log midpoint returns within each minute, expressed in basis points. All metrics are aggregated to one-minute resolution by taking the first value within each minute bin.

After construction, the venue-internal panel contains 531,305-minute observations across 78 dates and 5 assets. Coverage is essentially complete in the pre-event baseline window (October 3-9, 2025: 10,080 minutes per asset) and in the primary post-event window (October 11 through November 2, 2025: approximately 25,000 minutes per asset accounting for a six-day archive gap October 17-22). The Hyperliquid public archive also published no L2 snapshots for November 18-25, 2025; observations for those date ranges are missing from the extended recovery analysis.

3.4. Methodology

The analysis proceeds in five steps. First, we characterize the venue's normal-state inside market using the pre-event baseline window. Second, we compute hour-by-hour inside-market statistics on October 10 UTC to characterize the intraday cascade evolution. Third, we compute daily median spread and depth ratios relative to the pre-event baseline for each asset across the 51-day post-event window, smoothed with seven-day rolling medians. Fourth, we re-run the recovery analysis restricted to the 23-day clean window (October 11 to November 2, 2025) that excludes the November 3-4 selloff, to verify that the spread and depth findings are attributable to the cascade rather than to the subsequent event. Fifth, we compute time-to-threshold recovery measures for spreads at 1.5x, 1.2x, and 1.1x of pre-event baseline.

We define the metrics used in the paper as follows. For a given L2 snapshot, let a and b denote the best (highest) bid price and best (lowest) ask price, with associated resting sizes q_b and q_a in base units. The midpoint is $m = (a + b)/2$. The quoted spread in basis points is $S = 10,000 \times (a - b)/m$. The inside-quote depth in US dollars is $D = (q_b \times b) + (q_a \times a)$, the dollar size resting at the best bid plus that resting at the best ask. Snapshot-level values are aggregated to one-minute resolution by taking the first valid observation in each minute bin, yielding minute series S_t and D_t for each asset.

For each asset, the pre-event baseline median spread is the median of S_t over all minute bins in the baseline window (October 3-9, 2025), denoted $\bar{S}_0$; the baseline median depth $\bar{D}_0$ is defined analogously. The daily median spread on day d is the median of S_t over the minute bins in that day, $M_S(d) = \text{median}\{S_t: t \in d\}$, and the daily median depth $M_D(d) = \text{median}\{D_t: t \in d\}$. The reported spread ratio and depth ratio on day d are $R_S(d) = M_S(d)/\bar{S}_0$ and $R_D(d) = M_D(d)/\bar{D}_0$, so that a value of 1 corresponds to the pre-event baseline. Reported series are smoothed with a centred seven-day rolling median of the daily ratios. Percentiles (for example the 95th percentile of spread, p_{95}) are computed as the empirical percentile of the minute-level S_t within the relevant hour or day bin, and the maximum is the largest minute-level S_t in that bin. Realized volatility within a minute is the annualized standard deviation of one-second log midpoint returns in that minute, expressed in basis points. All analysis code and derived data products are made publicly available; see the Data Availability Statement.

4. Empirical Results and Discussion

4.1. Pre-Event Inside-Market Characterization

The cross-sectional pattern is consistent with stylized facts about market quality across asset capitalizations. BTC quotes a tight inside spread (median 0.088 bps) backed by approximately \$1.8 million of inside depth; DOGE quotes a wider spread (0.562 bps) backed by approximately \$28,000 of inside depth. The five assets span roughly an order of magnitude in pre-event spread and approximately three orders of magnitude in pre-event depth. On a 24/7 venue with no opening or closing auctions, median quoted spreads

pooled across assets vary by only seven percent across UTC hours of the day and six percent across days of the week, contrasting sharply with conventional equity venues (Wood et al., 1985). The flatness is itself an institutional feature of continuous global trading without cohort handovers. Table 1 reports the pre-event baselines across the five assets.

Table 1. Pre-event inside-market baselines (October 3-9, 2025).

Asset	Median spread (bps)	Mean spread (bps)	Median depth (USD)	Median realized vol (bps)
BTC	0.088	0.097	1,825,661	2.11
ETH	0.238	0.248	1,400,559	3.65
SOL	0.467	0.479	444,653	5.57
XRP	0.369	0.385	77,696	4.49
DOGE	0.562	0.602	28,135	6.46

Note: Each cell is computed across 10,080-minute observations per asset (7 days $\times$ 1,440 minutes). Inside depth is the dollar value resting at the best bid plus the best ask, summed across both sides. Realized volatility is the annualized standard deviation of one-second log midpoint returns within each minute, in basis points.

4.2 Intraday Evolution on October 10

Figure 1 presents the minute-level evolution of inside spread on October 10, 2025 UTC for each of the five assets. Spreads on all five assets remained near baseline through approximately hour 13. A first elevation is visible around hour 15 (corresponding to a brief BTC drawdown), partial recovery through hour 17, and a second elevation around hour 19. The principal cascade peak occurred in hour 21 UTC, with all five assets simultaneously experiencing spread widening of one to two orders of magnitude.




Figure 1. Minute-level inside spread on October 10, 2025 UTC, by asset. Five-minute rolling median, log scale. Vertical dashed line indicates the principal cascade peak at hour 21 UTC.

Table 2 reports the hour-by-hour evolution pooled across assets.

Table 2. Hour-by-hour inside-market evolution, October 10, 2025 UTC.

Hour UTC	Median spread (bps)	p95 spread (bps)	Max spread (bps)	Median depth (USD)	Median rv (bps)
00	0.372	0.566	0.639	415,716	3.43
04	0.378	0.667	1.014	604,670	4.59
08	0.381	0.658	0.904	520,120	4.37
12	0.368	0.577	0.728	493,769	3.44
13	0.399	0.802	1.071	548,614	6.64
14	0.441	0.802	4.135	597,182	7.27
15	0.724	1.983	4.968	463,643	16.60
16	0.497	1.009	1.568	623,944	9.89
17	0.485	0.936	1.188	487,329	8.99
18	0.439	0.960	1.178	515,262	8.04
19	0.588	1.291	2.852	566,483	12.78
20	0.534	3.030	13.692	514,913	8.66
21	9.376	87.476	241.819	391,473	150.62
22	1.709	18.804	44.351	375,291	42.46
23	1.080	5.576	8.713	217,889	31.14

Note: Statistics are pooled across all five assets within each UTC hour bin. Hours 1, 2, 3, 5, 6, 7, 9, 10, 11 are similar to hour 0 and omitted for brevity.

The first 14 hours of the day were essentially indistinguishable from a normal trading day. Beginning hour 15, both spread and realized volatility elevated noticeably, and through hours 16 to 20 the inside market was moderately stressed. The peak occurred in hour 21, when median spread reached 9.38 bps (more than 25 times normal levels), the 95th percentile reached 87.5 bps, the maximum minute observation reached 241.8 bps, and median realized volatility reached 150.6 bps (more than 30 times normal). By hour 22 the median spread had already declined to 1.71 bps, and by hour 23 to 1.08 bps.

4.3 Spread Recovery in the Clean Post-Event Window

Because a separate market-wide selloff began on November 3, 2025, our primary recovery analysis focuses on the 23-day clean window between the cascade and the second event, that is, October 11 through November 2, 2025. Table 3 reports the key recovery metrics in this window.

Table 3. Recovery metrics in the clean post-event window (October 11 to November 2, 2025).

Asset	Max spread ratio	End-of-window spread ratio	Days to $\leq 1.2x$	Days to $\leq 1.1x$
BTC	1.31x	1.07x	6	15
ETH	1.41x	1.13x	6	15
SOL	1.43x	1.19x	14	Not reached
XRP	1.82x	1.14x	15	Not reached
DOGE	1.77x	1.06x	13	14

Note: Spread ratio is the daily median quoted spread relative to the pre-event baseline (October 3-9, 2025). “Max” refers to the maximum daily median observed during the 23-day window. “End-of-window” is the spread ratio on November 2, 2025. Threshold-recovery columns report the number of days after October 10

required to first reach the threshold; “Not reached” indicates the threshold was not crossed within the 23-day window.

The recovery in quoted spread is rapid and close to complete. Within 6 days, BTC and ETH were within 1.2 times pre-event spread; within 13-15 days, BTC, ETH, and DOGE were within 1.1 times. By the end of the 23-day window, BTC (1.07x) and DOGE (1.06x) were essentially fully recovered, ETH (1.13x) was close, and the two laggards (SOL at 1.19x and XRP at 1.14x) were within approximately 20 percent. This pattern contrasts with the more common stylized finding that stress-event spread widening persists for extended periods; on Hyperliquid, the spread dimension of execution quality was substantially restored within a month of a major cascade.

Figure 2 presents the seven-day rolling spread ratio for each asset across the full 51-day post-event window.




Figure 2. Daily spread ratio relative to pre-event baseline, seven-day rolling median, by asset. Days 1-23 correspond to the clean post-event window. Days 24 onward are confounded by the November 3-4 market-wide selloff.

4.4 Depth Contraction - The Persistent Dimension

The recovery story in depth is qualitatively different from the spread story. Table 4 reports depth dynamics in the same clean post-event window.

Four of five assets (all except XRP) showed further depth contraction from the first week to the third week of the clean post-event window. BTC depth declined by 16.2 percentage points, DOGE by 14.7 percentage points, ETH by 9.8 percentage points, SOL by 2.7 percentage points. Only XRP showed any depth recovery, and the improvement was marginal (+1.5 pp) from an already-low starting level (0.76x). Four of five assets ended the clean window with depth materially below pre-event levels, and three were below 0.80x on the days 17-23 median measure. The depth contraction preceded the November 3-4 event and is therefore attributable to the cascade itself, not to the subsequent disturbance.

Table 4. Depth contraction in the clean post-event window (October 11 to November 2, 2025).

Asset	Days 1-7 median depth ratio	Days 17-23 median depth ratio	Change (percentage points)
BTC	0.93x	0.77x	-16.2
ETH	0.88x	0.79x	-9.8
SOL	0.93x	0.90x	-2.7
XRP	0.76x	0.78x	+1.5
DOGE	0.97x	0.83x	-14.7

Note: Depth ratio is the daily median inside-quote depth relative to the pre-event baseline. “Days 1-7 median” aggregates October 11-16 (with October 17 missing due to archive gap). “Days 17-23 median” aggregates October 27 to November 2.

Figure 3 extends the depth analysis through the full 51-day post-event window. The pattern differs markedly from the spread figure, in that depth trajectories are monotonically declining through the entire window with no visible second-event signature. Two contracts (ETH and DOGE) cross the 0.5x line by the end of the observation period — inside depth cut by half relative to pre-event seven weeks later.




Figure 3. Daily depth ratio relative to pre-event baseline, seven-day rolling median, by asset. Unlike the spread series in Figure 2, depth shows monotonic deterioration over the entire 51-day window, with no visible second-event inflection.

4.5 Cross-Sectional Pattern

Figure 4 presents the cross-sectional pattern at 30 days post-cascade. The left panel plots spread widening against pre-event spread; the right panel plots depth contraction against the same.




Figure 4. Cross-sectional pattern of persistent post-event effects at 30 days. Left: spread widening (% vs pre-event). Right: depth contraction (% vs pre-event).

The depth cross-section is the cleaner finding and is interpretable as follows. BTC and ETH (the two largest contracts by pre-event depth) exhibit the largest depth contractions (38 and 42 percent). The pattern is consistent with an interpretation in which tail-risk perception increases proportionally to notional exposure, in that makers committed the most inventory to the largest contracts pre-event and therefore had the most inventory to reduce post-event. The smallest contract (DOGE) exhibits the smallest percentage depth contraction, but this should not be interpreted as DOGE being least affected — DOGE had the least depth to begin with.

4.6 Discussion: Two-Regime Recovery

The central empirical finding of this paper is the two-regime recovery pattern, in which pricing recovers quickly and nearly completely within three weeks of a cascade while inventory commitment contracts persistently over the same window and continues to deteriorate over the full observation period. The pattern is consistent with the inventory-management framework of Stoll (1978) and Ho and Stoll (1981), in which market makers set both quoted prices and committed sizes as separable decisions responding to different risk considerations. A tail event raises the perceived variance of inventory holdings, and the rational maker response is to reduce size exposure rather than demand a wider spread - particularly in a competitive market where quoted prices are disciplined by the possibility of being undercut.

The pattern has implications for how venue resilience should be measured. A metric that examines only the spread dimension - as is conventional in most empirical microstructure work - will systematically understate the post-event impairment of venue liquidity when makers respond primarily through size reduction. A complete resilience assessment requires monitoring both dimensions, with particular attention to depth which the evidence here suggests is the more persistent and potentially more consequential dimension for large-order execution quality.

Hyperliquid's distinguishing market-design feature is the consensus-layer cancel prioritization rule, motivated by maker protection. The findings are consistent with this protection operating during normal-state trading, in that pre-event spreads are tight relative to the venue's age and intraday patterns are flat (Section

4.1). The rapid spread recovery documented in Section 4.3 is also consistent with an environment where makers feel confident re-entering quoted prices quickly post-stress. But the persistent depth contraction (Section 4.4) suggests that cancel prioritization, while beneficial for normal-state market quality and for fast pricing recovery, does not insulate the venue from the inventory-commitment effects of a tail event. Venue resilience and venue-design quality are distinct properties.

There are some limitations. First, the Hyperliquid public archive publishes only the top twenty levels of the book on each side. Wider-book dynamics (depth at, say, 25 basis points outside the inside) are not directly observable from the public data; the depth metric we report captures top-of-book commitment specifically. Second, the archive published no L2 snapshots for October 17-22 and November 18-25, 2025. The first gap falls within the clean post-event window and reduces our coverage of the 7-12 day post-cascade range. Third, we cannot directly identify individual market makers; our analysis is silent on whether the persistent inventory contraction reflects a few large makers exiting or many small makers each reducing exposure marginally. Fourth, the November 3-4 selloff confounds post-event analysis beyond day 23; our primary recovery findings are restricted to the clean window. Fifth, the analysis is venue-internal by design and does not address whether Hyperliquid's recovery dynamics are typical or atypical relative to centralized perpetual venues experiencing the same cascade.

5. Conclusion

This paper documents the inside-market response of Hyperliquid - the largest decentralized perpetual futures venue - to the October 10, 2025 system-wide liquidation cascade. Using the venue's public L2 archive, we construct a minute-level panel spanning five major perpetual contracts across a 92-day window centered on the event. The analysis yields three main findings.

First, the cascade was a tightly compressed intraday event. Median quoted spread pooled across assets reached 9.4 basis points at hour 21 UTC and the 95th percentile reached 87.5 basis points, with intraday recovery well under way by end of day. Second, in the 23-day clean post-event window prior to a separate market-wide selloff, quoted spreads recovered substantially - three of five contracts within 1.1 times pre-event levels and all within 1.2 times. Third, inside-quote depth contracted persistently over the same clean window, with four of five contracts showing further depth reduction from days 1-7 to days 17-23 post-event and none fully recovering over the full 51-day observation window.

The contrast between the two findings - spread recovery near complete while depth contraction persists - motivates what we term two-regime liquidity recovery. Market makers maintain competitive quoted prices but reduce the inventory they commit at those quotes following a tail event. Three implications follow. For microstructure theory, the finding is consistent with the classic inventory-management framework's prediction that pricing and sizing are separable maker decisions responding to different risk considerations. For empirical market-resilience research, the finding implies that spread-only metrics systematically understate the impairment of venue liquidity following stress events and that depth metrics should be a standard component of resilience analysis.

For decentralized derivatives venue design, the finding establishes that a venue specifically engineered to protect quoting market makers can still experience multi-week impairment of inventory commitment after a tail event; venue resilience and venue-design quality are distinct properties. For investors and traders, the

finding carries a practical implication. In the weeks after a cascade, a tight quoted spread can give a misleading impression that liquidity has normalized, while the depth available at the inside remains materially reduced. Investors who size orders against the quoted spread alone risk underestimating the price impact and execution cost of larger trades during this recovery phase. Those executing size in the post-event window may benefit from monitoring inside depth directly, working orders more patiently, or scaling trade sizes to the reduced depth rather than to the apparently recovered spread.

Beyond the specific findings, this paper also takes a methodological position. The public order-book archives that on-chain venues publish offer opportunities for transparent and replicable empirical microstructure research that are difficult to match in centralized venues that typically restrict tick-level access. We hope the results encourage future replication and extension on other on-chain venues with similar archive practices.

Data Availability Statement

All code, derived data products, and figures used in this study are publicly available at the project repository: <https://github.com/boonchuan/hyperliquid>. The raw Hyperliquid L2 order-book snapshots analyzed are publicly available from Hyperliquid's S3 archive at `s3://hyperliquid-archive` (requester-pays bucket; total egress cost for the full sample was approximately ten US dollars). The repository contains the full analysis pipeline - data-pull scripts, metrics construction, cascade event study, and figure generation - that reproduces every table, figure, and numerical claim in this paper from the public raw data.

Funding

This research received no external funding. The author self-funded all data-access costs (approximately ten US dollars in AWS egress fees) and computational resources used in this study.

Acknowledgements

The author thanks the Hyperliquid Foundation for maintaining the public L2 order-book archive that made this study possible.

Declaration of the Use of Generative AI

During the preparation of this work, the author used Claude (Anthropic) for drafting prose, editing, generating Python code for data analysis, and structuring the manuscript. All empirical analysis, data pulls, methodology choices, and interpretations are the author's own. Every number, table, and figure in the paper was generated from the author's own data pipeline running on the author's machine; AI-generated text was verified against the underlying data and revised by the author throughout. The AI tool was not used to generate citations or references; all references were verified by the author against primary sources. The AI system is not an author and bears no responsibility for the content. After using this tool, the author reviewed, revised, and verified all content to ensure accuracy and accepts full responsibility for all aspects of the publication.

Conflicts of Interest

The author declares no conflicts of interest. The author has no financial relationship with, position at, or advisory role for Hyperliquid, the Hyperliquid Foundation, Binance, or any of the venue operators referenced in this paper. The author is an independent researcher and received no funding from any party with an interest in the findings.

References

- Anand, A., & Venkataraman, K. (2016). Market conditions, fragility, and the economics of market making. *Journal of Financial Economics*, 121(2), 327–349. <https://doi.org/10.1016/j.jfineco.2016.03.006>
- Bessembinder, H., Hao, J., & Zheng, K. (2015). Market making contracts, firm value, and the IPO decision. *Journal of Finance*, 70(5), 1953–2002. <https://doi.org/10.1111/jofi.12285>
- Brauneis, A., Mestel, R., Riordan, R., & Theissen, E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124, 106041. <https://doi.org/10.1016/j.jbankfin.2020.106041>
- Brogaard, J., Carrion, A., Moyaert, T., Riordan, R., Shkilko, A., & Sokolov, K. (2018). High frequency trading and extreme price movements. *Journal of Financial Economics*, 128(2), 253–265. <https://doi.org/10.1016/j.jfineco.2018.02.002>
- Chitra, T. (2025). Autodeleveraging: Impossibilities and optimization [Working paper].
- Foley, S., Karlsen, J. R., & Putnins, T. J. (2019). Sex, drugs, and Bitcoin: How much illegal activity is financed through cryptocurrencies? *Review of Financial Studies*, 32(5), 1798–1853. <https://doi.org/10.1093/rfs/hhz015>
- Foucault, T., Kadan, O., & Kandel, E. (2013). Liquidity cycles and make/take fees in electronic markets. *Journal of Finance*, 68(1), 299–341. <https://doi.org/10.1111/j.1540-6261.2012.01801.x>
- Hasbrouck, J., & Sofianos, G. (1993). The trades of market makers: An empirical analysis of NYSE specialists. *Journal of Finance*, 48(5), 1565–1593. <https://doi.org/10.1111/j.1540-6261.1993.tb05121.x>
- He, S., Manela, A., Ross, O., & von Wachter, V. (2024). Fundamentals of perpetual futures [Working paper]. Washington University in St. Louis.
- Hendershott, T., & Moulton, P. C. (2011). Automation, speed, and stock market quality: The NYSE's Hybrid. *Journal of Financial Markets*, 14(4), 568–604. <https://doi.org/10.1016/j.finmar.2011.02.003>
- Ho, T., & Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1), 47–73. [https://doi.org/10.1016/0304-405X\(81\)90020-9](https://doi.org/10.1016/0304-405X(81)90020-9)
- Hyperliquid Foundation. (2024). Latency and transaction ordering on Hyperliquid. <https://hyperliquid.gitbook.io>
- Madhavan, A., & Smidt, S. (1991). A Bayesian model of intraday specialist pricing. *Journal of Financial Economics*, 30(1), 99–134. [https://doi.org/10.1016/0304-405X\(91\)90023-D](https://doi.org/10.1016/0304-405X(91)90023-D)

Ruan, Q., & Streltsov, A. (2024). The emerging market of cryptocurrencies and perpetual contracts [Working paper]. Cornell University.

Stoll, H. R. (1978). The supply of dealer services in securities markets. *Journal of Finance*, 33(4), 1133–1151. <https://doi.org/10.1111/j.1540-6261.1978.tb02053.x>

Wood, R. A., McInish, T. H., & Ord, J. K. (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance*, 40(3), 723–739. <https://doi.org/10.1111/j.1540-6261.1985.tb04996.x>

From Agent Identity to Agent Economy: Measuring the Operational Readiness of ERC-8004 AI Agents.

Rischan Mafrur * and Priagung Khusumanegara

Authors

Rischan Mafrur,
Lecturer at School of Computer, Data and
Mathematical Sciences (CDMS), Western
Sydney University, Sydney, Australia.
Email: R.Mafrur@westernsydney.edu.au

Priagung Khusumanegara,
Telkomsel, Indonesia.
Email:
priagung_khusumanegara@telkomsel.co.id

* Corresponding author.

Abstract

This paper examines whether blockchain-registered AI agents demonstrate operational readiness beyond identity registration. Using a dataset of ERC-8004 agents on Ethereum, we construct an agent-level feature table covering identity status, metadata, service declarations, reputation feedback, transfers, and cross-chain registration. We develop an operational readiness framework based on observable evidence layers and complement it with network analysis of owner-agent, feedback-client, wallet-transfer, and combined evidence relationships. The results show that early ERC-8004 adoption is registration-heavy but operationally shallow. While the identity layer is visible at scale, metadata availability, service exposure, reputation formation, and cross-chain evidence remain limited. Ownership and feedback activity are also highly concentrated, suggesting that early participation is shaped by a small number of high-activity wallets and clients. The network analysis further shows that richer operational evidence clusters around a small subset of agents rather than being broadly distributed across the ecosystem. The findings suggest that ERC-8004 provides an important identity layer for decentralized AI agents, but the transition from agent identity to agent economy remains incomplete.

Keywords: ERC-8004, AI agents, Decentralized identity, Agent economy
JEL Classification: C55, L86, O33

Article Information

Received: 14 July 2026.
Revised: 09 August 2026.
Accepted: 15 August 2026.
Published: September 2026.

Citation:

Mafrur, R. & Khusumanegara, P. (2027). From Agent Identity to Agent Economy: Measuring the Operational Readiness of ERC-8004 AI Agents. *The Journal of FinTech and Digital Assets*, 1(2), 69–88.

1. Introduction

Autonomous AI agents are increasingly discussed as future economic actors that can search for information, coordinate with other agents, execute transactions, and provide services with limited human intervention (Kirana & Havidz, 2020; Tomašev et al., 2025; Xu, 2026). In decentralized environments, this vision depends not only on advances in artificial intelligence, but also on infrastructure that allows agents to be identified, discovered, paid, evaluated, and trusted. Blockchain platforms such as Ethereum provide a natural setting for this infrastructure because they support machine-readable identity, programmable accounts, transparent transaction records, and public trust mechanisms (Al Jasem et al., 2025; Cao, 2022; Xu, 2026). Within this setting, ERC-8004 (De Rossi et al., 2025) has emerged as a proposed Ethereum standard for trustless AI agent infrastructure. It assigns each agent a unique on-chain identity, represented through an ERC-721 style token, and links this identity to metadata, and reputation. In principle, such infrastructure could support an agent economy in which autonomous agents offer services, transact with one another, build reputation, and participate in decentralized coordination (Liu, 2026a).

Yet the existence of a technical standard does not by itself demonstrate that an agent economy is operating in practice. A registered agent may simply be an on-chain identity with little or no metadata, no declared service endpoint, no reputation history, and no evidence of interaction. This distinction is important because prior work on AI and blockchain warns that tokenized AI projects can create an illusion of decentralization when claims of autonomy, intelligence, or trust are not supported by verifiable usage, transparent execution, or effective governance (Mafrur, 2025; Romandini et al., 2025; Yu et al., 2026). In this context, the key empirical question is not whether ERC-8004 agents exist, but whether they show observable signs of operational readiness. The following questions remain unknown:

1. To what extent do ERC-8004 agents progress beyond on-chain identity registration toward observable operational readiness across metadata, service, reputation, and cross-chain evidence layers?
2. How distributed or concentrated are ownership, reputation feedback, service infrastructure, and interaction patterns within the ERC-8004 ecosystem?
3. Which observable agent characteristics are associated with the formation of reputation feedback in the ERC-8004 ecosystem?

This paper addresses these questions by examining the transition from agent identity to agent economy. Using a dataset of 10,000 ERC-8004 agents on Ethereum, we analyse covering identity status, metadata, service declarations, reputation feedback, transfer activity, and cross-chain registration. We then develop an operational readiness framework based on observable evidence layers. In this framework, an agent progresses from basic identity registration toward richer forms of readiness when it provides metadata, declares services, accumulates feedback, appears in cross-chain records, or participates in observable transfer activity. This approach allows us to distinguish between formal registration and more substantive evidence of ecosystem participation.

Our empirical analysis shows that early ERC-8004 adoption is registration-heavy but operationally shallow. Although the dataset contains 10,000 registered agents, only 67 agents expose service records, only 628 agents receive reputation feedback, and only 19 agents combine metadata, services, feedback, and cross-chain registration. Ownership and reputation activity are also highly concentrated: the top 10 owner wallets hold 51.40% of agents, while the largest feedback client contributes 65.82% of all feedback records. These findings suggest that ERC-8004 currently functions more strongly as an identity-registration layer than as a mature agent economy with broad service discovery, distributed reputation formation, and multi-party interaction.

To complement the readiness indicators, we also conduct a network analysis of the ERC-8004 ecosystem. We construct owner-agent, feedback-client-agent, wallet-transfer, and combined evidence networks to examine how agents are connected to wallets, feedback providers, service domains, and cross-chain records. The network results reinforce the main descriptive findings. Richer operational evidence is clustered around a small subset of agents, while most registered agents remain disconnected from the higher layers of service, feedback, and cross-chain activity. This network perspective is important because an agent economy should not only contain many registered identities, but should also exhibit distributed relationships among agents, users, service endpoints, and trust mechanisms.

This paper makes three contributions. First, it provides one of the first empirical assessments of ERC-8004 as a blockchain-based agent identity and trust standard. Second, it proposes a layered operational readiness framework for decentralized AI agents, linking identity, metadata, services, feedback, reputation,

and cross-chain registration. Third, it introduces a network-based perspective on agent ecosystem maturity by visualizing how agents connect to owners, feedback clients, service domains, and other infrastructure signals. Together, these contributions provide a practical way to evaluate whether blockchain-registered AI agents are progressing beyond identity creation toward an operational agent economy.

The research proceeds as follows. Section 2 provides the related literature review, followed by concept framework in Section 3, data and methodology in Section 4. Section 5 discusses the results, and Section 6 concludes the report.

2. Literature Review

2.1 Autonomous AI Agents and Blockchain-Based Agent Infrastructure

Recent advances in large language models and multi-agent systems have renewed interest in autonomous software agents that can interpret goals, reason over context, use tools, and act with limited human intervention (S. Fan & Min, 2026; T. Fan et al., 2025). In conventional digital environments, these capabilities are often evaluated through task completion, planning accuracy, or tool-use performance. In Web3 environments, however, the requirements are more demanding. Agents must interpret on-chain state, interact with smart contracts, manage wallet-related operations, and avoid irreversible transaction errors (S. Fan & Min, 2026; Karim et al., 2025; Romandini et al., 2025). As a result, blockchain-based agents cannot be evaluated only by whether they generate plausible instructions or recommendations. They must also be assessed by whether they can operate safely, reliably, and transparently in transaction-based environments.

The existing literature illustrates both the promise and the limits of this transition. Language-model agents can translate natural language instructions into structured on-chain operations, but reliable execution depends on intent extraction, instruction decomposition, state grounding, calibration, and error handling (S. Fan & Min, 2026). Evidence from autonomous trading and financial-market agents also shows that general language-model capability does not automatically translate into reliable domain-specific decision-making (T. Fan et al., 2025). These findings suggest that autonomy is not a binary property. An agent may be technically capable of generating actions, but still lack the surrounding infrastructure needed for safe and trustworthy participation in blockchain markets.

Security and privacy research reinforces this point. Blockchain AI agents can be understood along a spectrum from conversational systems to instruction-following systems, to goal-directed systems that execute actions or manage objectives over time (Chaffer et al., 2024; Ranjan et al., 2025; Romandini et al., 2025). As agents move along this spectrum, risks increase. Goal-directed agents may introduce risks related to erroneous behavior, prompt injection, context manipulation, privacy leakage, wallet exposure, market instability, and loss of human oversight (Romandini et al., 2025). These risks make trust infrastructure central to any agent economy. Before agents can transact or coordinate at scale, other actors must be able to identify them, discover their services, assess their past behavior, and verify whether their outputs or claims are reliable.

Identity is therefore a foundational requirement for autonomous economic interaction. Humans rely on legal personhood, institutional records, and social identity to open accounts, enter contracts, and build reputation. Autonomous agents do not naturally possess these mechanisms. Instead, machine actors require cryptographic identity and verifiable credentials that can be interpreted by other machines (Tomašev et al., 2025; Xu, 2026). Ethereum's token standards provide part of this foundation. The non-fungible token model

allows unique digital entities to be represented as distinct on-chain objects, making it suitable for persistent agent identities (Dafflon et al., 2023; Entriken et al., 2018). ERC-8004 extends this logic by proposing a registry-based architecture for trustless AI agent infrastructure (De Rossi et al., 2025; Liu, 2026a).

ERC-8004 provides a framework through which agents can be represented, discovered, evaluated, and validated using Ethereum-based records. The standard is organized around three linked components: identity, feedback, and reputation. The Identity Registry provides each agent with a unique on-chain identity, represented through an ERC-721 style token. The Reputation Registry allows counterparties to submit feedback about agents. The significance of ERC-8004 lies in this connection between identity and trust. However, the existence of a standard does not mean that registered agents are operationally ready. A functioning agent economy would require agents to have accessible metadata, declared services, meaningful feedback and interaction patterns that extend beyond isolated minting. This distinction between formal registration and operational use motivates the empirical focus of this paper.

2.2 From Agent Identity to Agent Economy: Feedback, Reputation, and Operational Readiness

The idea of an agent economy has received increasing attention in recent work on decentralized AI and autonomous commerce. In this view, virtual agent economies are environments in which AI agents transact, negotiate and coordinate at speeds and scales beyond direct human oversight (Tomašev et al., 2025). Such systems require careful design around market permeability, reputation, resource allocation, credit assignment, and governance (Tomašev et al., 2025). Related work also argues that agents require a layered blockchain foundation that combines physical infrastructure, identity and agency, economic settlement, and collective governance (Xu, 2026).

Empirical studies of blockchain-based AI agents suggest that the ecosystem remains early, fragmented, and unevenly developed. Existing evidence shows that autonomous AI agent projects in decentralized finance already span trading, analytics, infrastructure, community engagement, and entertainment, but many of these projects remain heterogeneous and experimental rather than operationally mature (Ante, 2026). Studies of DeFi investment agents similarly show that tokenized agent markets can attract substantial valuations while providing limited evidence of verifiable autonomous execution, sustained performance, or stakeholder alignment (Yu et al., 2026). Related work on AI-based crypto projects further cautions that decentralization claims may become an illusion when core computation, governance remains off-chain and weakly verifiable (Mafrur, 2025). Taken together, this literature suggests that agent economies should not be inferred from branding, tokenization, or identity creation alone.

Operational readiness therefore provides a useful lens for distinguishing formal registration from actual ecosystem maturity. In blockchain settings, readiness implies active participation, observable use, and meaningful interaction rather than the mere existence of an on-chain record (Yu et al., 2026). In the AI agent context, this means that readiness should be understood as a progression from identity infrastructure toward economic functionality. An agent may first appear as a registered identity, but it becomes more operationally meaningful only when it provides accessible metadata, declares services, receives reputation feedback, participates in cross-chain or transfer activity, and generates evidence that other actors can inspect or evaluate. This paper builds on this logic by defining observable readiness layers and examining whether ERC-8004 agents move beyond identity registration toward richer forms of ecosystem participation.

A mature agent economy requires observable participation across multiple layers. Agents must be discoverable, able to expose services, accumulate reputation, and interact with other agents or users in ways that can be audited. Reputation is particularly important because it connects identity with historical performance. For autonomous agents, these mechanisms must be replaced or supplemented by machine-readable records that can be verified and reused across systems (Tomašev et al., 2025; Xu, 2026). On-chain reputation offers one possible approach because feedback can be recorded publicly, associated with a persistent agent identity, and queried by other participants.

At the same time, reputation systems must be interpreted carefully. A feedback record is useful only if it reflects meaningful interaction, credible evaluation, and sufficient independence between the evaluator and the agent. Existing work argues that agent markets require robust reputation mechanisms to support trust, specialization, and collaboration, and that historical performance can become a form of reputation capital or trust signal for autonomous agents (Tomašev et al., 2025; Xu, 2026). Related security research also emphasizes the need for accountability and auditability because blockchain actions can be irreversible and financially consequential (Romandini et al., 2025). These studies suggest that reputation should not be measured only by whether feedback exists. It should also be assessed by who provides feedback, how concentrated feedback activity is, what dimensions are being evaluated, and whether reputation is distributed across a wider ecosystem.

The literature also highlights broader risks that may affect the pace and shape of agent infrastructure adoption. Autonomous and decentralized agents may introduce security vulnerabilities, including prompt injection, key leakage, malicious tool use, and unintended transaction execution (Romandini et al., 2025). They may also create privacy risks when agents process sensitive user instructions, wallet data, or off-chain context (Romandini et al., 2025; Tomašev et al., 2025). In financial settings, autonomous agents could amplify market volatility if many agents react to similar signals, execute similar strategies, or interact through automated markets (Ante, 2026). These risks suggest that sparse adoption of agent infrastructure may reflect not only technical immaturity, but also caution around trust, governance, and safety.

The existing literature therefore leaves an important empirical gap. Prior work has developed conceptual models of agent economies, evaluated individual autonomous agents, and discussed the risks of blockchain-based AI agents. However, there is still limited evidence on whether the infrastructure proposed for agent economies is being used in practice. In particular, it remains unclear whether ERC-8004 agents are merely being registered as identities, or whether they are developing the metadata, services, reputation records, and network relationships needed for operational readiness. This paper addresses that gap by examining observable ERC-8004 activity on Ethereum mainnet. Instead of asking only whether agents exist, we assess whether their on-chain and metadata records show evidence of progression from agent identity toward an operational agent economy.

3. Conceptual Framework

This paper conceptualizes ERC-8004 agents through a layered operational readiness framework. The framework starts from a simple premise: an agent economy cannot be inferred from identity registration alone. A blockchain-registered agent may exist as an on-chain object, but this does not necessarily mean that it is discoverable, service-capable, trusted or economically active. Operational readiness therefore refers to the extent to which an agent progresses from basic identity registration toward observable forms of ecosystem participation. Building on prior work on agent economies, decentralized identity, reputation, and blockchain-

based agent infrastructure (Romandini et al., 2025; Tomašev et al., 2025; Xu, 2026), we define readiness as the movement from agent identity toward usable and verifiable agent activity.

The framework contains several connected layers. At the foundation is identity readiness, where an agent is registered on-chain and assigned a persistent identity through an ERC-721 style token. This layer makes the agent visible, but it is only a minimal condition for participation. The next layer is metadata readiness, where the agent provides structured information such as a name, description, status, image, capability claims, or interaction-related attributes. Metadata is important because it links a cryptographic identity to a usable description of what the agent is and what it claims to do. The third layer is service readiness, where the agent declares endpoints, interfaces, or service capabilities that allow other actors to interact with it. This may include web endpoints, documentation links, communication protocols, or agent-native service formats such as MCP, A2A, OASF, and related endpoint structures. The fourth layer is reputation readiness, where feedback records connect the agent's identity and service claims to evidence of past interaction. Reputation is therefore a key bridge between registration and trust formation.

To operationalize this framework, identity registration is treated as the baseline condition because all agents in the sample are registered agents by construction. We then measure whether each agent shows additional observable evidence beyond this baseline. Specifically, we construct five Boolean indicators: *has meta data*, which captures whether an agent has observable metadata evidence; *has service*, which captures whether an agent declares at least one service record; *has feedback*, which captures whether an agent has received at least one feedback record; *has cross chain*, which captures whether an agent appears in cross-chain registration records; and *has transfer*, which captures whether an agent has observed transfer activity. These indicators correspond to observable forms of ecosystem participation, but they should be interpreted carefully. In particular, transfer activity indicates ownership movement or custody changes, not direct evidence of autonomous behavior.

We combine these indicators into an observable evidence score. The score is calculated as the sum of the five Boolean indicators, excluding identity registration itself. An agent with a score of zero is registered but has no additional observable evidence layer. An agent with a score of one has one additional evidence layer, such as metadata or feedback. Higher scores indicate that the agent is connected to multiple forms of observable ecosystem participation. This scoring approach provides a simple and transparent way to distinguish registration-only agents from agents with richer operational signals.

4. Data and Methodology

We analyse a compiled dataset of ERC-8004 agents introduced by Liu (2026a, 2026b). The sample consists of agent IDs 0–9999, covering the first 10,000 agents registered on Ethereum mainnet. The on-chain data were collected from block 24,339,925 on January 29, 2026, to block 24,839,925 on April 9, 2026. The dataset is provided as CSV tables: *agents_core*, *agent_metadata*, *agent_services*, *agent_reputation_summary*, *agent_feedback_records*, *mint_economics*, *transfer_history*, and *crosschain_registrations* (Liu, 2026a, 2026b). These tables allow us to reconstruct agent-level records covering identity, metadata, services, reputation, minting, transfer, and cross-chain evidence.

We perform data cleaning and integration as follows. All Ethereum addresses are lowercased to ensure consistent matching across tables. Boolean fields, such as *active_flag*, are converted into Boolean values. For

each agent, we merge information from the raw tables into a single feature record. We extract agent-level counts and indicators, including the number of service records, the number of unique service types, the number of feedback records, the number of unique feedback clients, and feedback values in both raw and decimal formats. We also extract ownership and transfer-related variables, including owner wallet, transfer count, and recipient diversity. Ownership concentration metrics, including the Herfindahl-Hirschman Index (HHI) and Gini coefficient, are computed from the distribution of agent counts across owner wallets.

Building on the integrated feature table, we construct agent-level readiness indicators. Identity registration is treated as the baseline condition because all agents in the sample are registered agents by construction. We then define additional observable evidence indicators: *has_metadata*, which captures whether an agent has observable metadata evidence; *has_service*, which captures whether an agent declares at least one service record; *has_feedback*, which captures whether an agent has received at least one feedback record; *has_crosschain*, which captures whether an agent appears in cross-chain registration records; and *has_transfer*, which captures whether an agent has observed transfer activity. These indicators allow us to distinguish agents that remain at the level of basic identity registration from those that show additional signs of ecosystem participation.

We combine these indicators into an observable evidence score. The score is calculated as the sum of the five Boolean indicators, excluding identity registration itself. An agent with a score of zero is registered but has no additional observable evidence layer. Higher values indicate that the agent is connected to multiple forms of observable activity, such as metadata, services, feedback, cross-chain registration, or transfer activity. This score provides a transparent way to measure whether ERC-8004 agents are progressing from identity registration toward operational readiness.

In addition to the descriptive readiness analysis, we estimate an exploratory logistic regression to examine which observable agent characteristics are associated with reputation formation. The dependent variable is *has_feedback*, which equals one if an agent has received at least one feedback record and zero otherwise. This outcome is used because feedback is one of the clearest observable signals that an agent has entered the reputation layer of ERC-8004. The model is specified as follows:

$$\Pr(\text{has_feedback}_i = 1) = \Lambda(\alpha + \beta X_i + \gamma Z_i), (1)$$

where $\Lambda(\cdot)$ denotes the logistic function, X_i represents readiness-related variables for agent i , and Z_i represents controls related to minting and ownership. The readiness variables include metadata indicators, service indicators, cross-chain registration, transfer activity, active metadata status, and lifecycle metadata linkage. The minting and ownership controls include log-transformed mint cost, gas used, mint batch size, and owner-group categories.

The regression is included as a supplementary association analysis. The purpose is to identify whether agents with richer observable infrastructure are more likely to receive feedback. This is important because the descriptive analysis shows how many agents reach each readiness layer, while the regression examines which observed features are associated with entry into the reputation layer. We use log-transformed variables for skewed count and cost measures, including service records, unique service types, transfer counts, minting cost, gas used, and mint batch size. A balanced logistic regression is used to account for the imbalance between agents with and without feedback. The coefficients are interpreted as associations conditional on the other

variables in the model and should not be interpreted causally, since feedback may be shaped by off-chain review processes, dominant clients, platform-specific mechanisms, or early-stage ecosystem sparsity.

Finally, we complement the readiness indicators and regression analysis with network analysis. The network analysis examines whether ERC-8004 activity is broadly distributed across many participants or concentrated around a small number of wallets, clients, service domains, and cross-chain records. This is important because a mature agent economy should not only contain many registered identities, but should also exhibit distributed relationships across identity, service, reputation, and interaction layers.




Figure 1. Operational readiness funnel for ERC-8004 agents. The figure shows how the sample narrows from identity registration to metadata, services, feedback, and cross-chain registration. The final category, full evidence, refers to agents that simultaneously have observable metadata, at least one service record, at least one feedback record, and at least one cross-chain registration record.

5. Results

5.1 Descriptive analysis

The descriptive results reveal a substantial gap between the conceptual promise of ERC-8004 and its current empirical implementation. Although identity registration is clearly observable, with 10,000 agents created and discoverable on-chain, the additional dimensions required for operational readiness remain weakly developed. This finding is consistent with (Mafrur, 2025), which cautions that blockchain-based AI projects may appear decentralised at the architectural or narrative level while lacking substantive ecosystem development in practice. In the case of ERC-8004, the evidence suggests that registration alone does not translate into meaningful participation, service provision, or trust formation. Most agents remain concentrated at the identity layer, rather than progressing toward the higher layers required for a functioning agent economy.

Figure 1 presents the readiness funnel, which shows how the population of agents narrows as additional layers of evidence are required. The funnel structure provides a systematic way to evaluate whether ERC-8004 agents are moving beyond basic on-chain identity and toward operational participation. The results

indicate that only a small subset of agents show evidence of richer metadata, service readiness, or feedback activity. This pattern suggests that many ERC-8004 registrations currently resemble inactive identities rather than economically active agents. The high level of ownership concentration reinforces this interpretation. Instead of being distributed across a broad base of end users, ERC-8004 agents appear to be held disproportionately by a small number of addresses, which may correspond to developers, early deployers, or speculative holders as shown in Figure 2a.

Feedback activity provides additional evidence of limited ecosystem maturity. In a healthy agent economy, reputation and feedback mechanisms should support trust formation through broad, repeated, and diverse participation. In the current ERC-8004 ecosystem, however, feedback is sparse and highly concentrated, with a substantial share originating from a single address as shown in Figure 2b. Figure 3 provides supplementary evidence on which observable agent characteristics are associated with entering the feedback layer. The strongest positive associations are observed for *lifecycle_metadata_linked* and *has_crosschain*, suggesting that agents with broader observable infrastructure are more likely to receive feedback. Smaller positive associations are also observed for service intensity variables, including *loglp_service_records* and *loglp_unique_service_types*. However, the binary indicators *has_metadata* and *has_service* have negative coefficients. This should not be interpreted as evidence that metadata or services reduce feedback. Rather, it likely reflects sparsity, overlap among readiness indicators, and the uneven structure of the early ERC-8004 dataset. Overall, the regression supports the descriptive interpretation that feedback formation is associated with richer observable infrastructure, but the relationship remains selective and should be interpreted as association rather than causality.

Overall, the descriptive evidence indicates that ERC-8004 is currently functioning more as an on-chain identity registry than as a mature infrastructure layer for an autonomous agent economy. At this stage, the ERC-8004 ecosystem appears to have established the lower layer of identity creation, but the higher layers of service activity, reputation formation, and economic coordination remain underdeveloped.

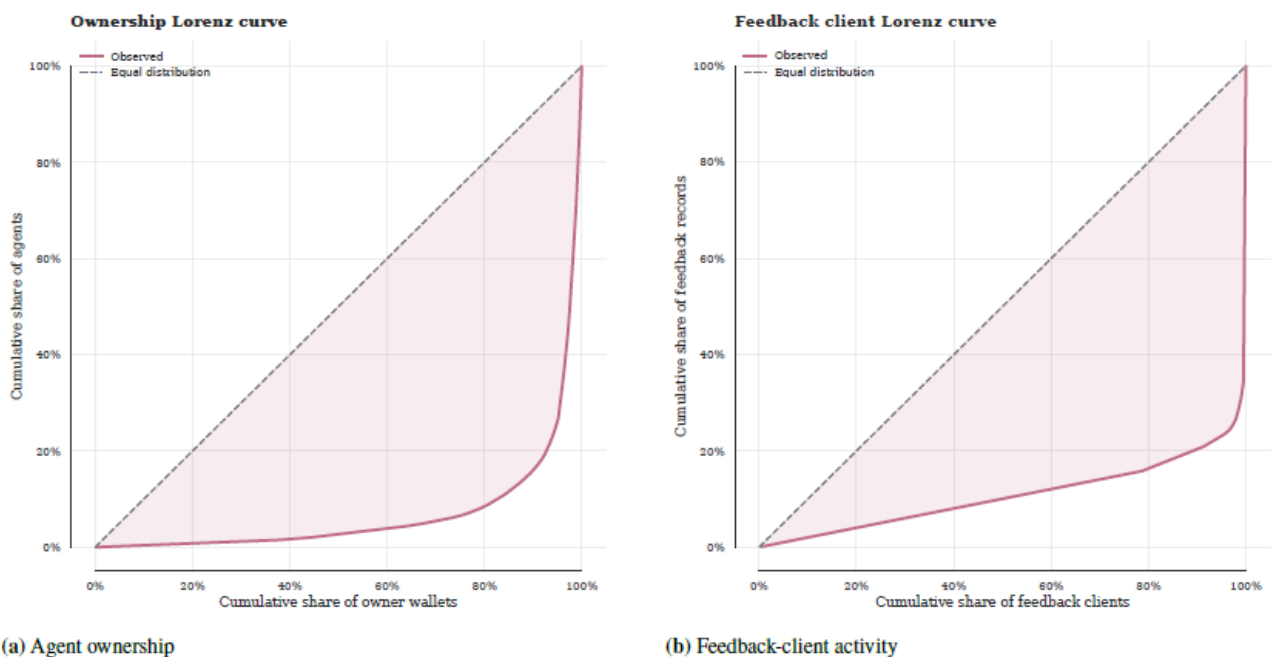


Figure 2. Lorenz curves for ERC-8004 ownership and reputation concentration. Panel (a) shows the concentration of agent ownership across owner wallets, while Panel (b) shows the concentration of reputation

feedback across feedback-client addresses. Together, the figures indicate that both identity ownership and reputation activity are unevenly distributed in the early ERC-8004 ecosystem.

5.2 Network Structure of ERC-8004 Agent Activity

To complement the tabular readiness indicators, we construct four network views of the ERC-8004 agent ecosystem. These plots are designed to show whether agent activity is broadly distributed or concentrated around a small number of wallets, feedback clients, service domains, and transfer paths. The network analysis is exploratory and should be interpreted as structural evidence of ecosystem organization rather than proof of autonomous agent behavior.

5.2.1 Owner-to-Agent Network

Figure 4 visualizes the relationship between owner wallets and ERC-8004 agent identities. Orange nodes represent owner wallets, while blue nodes represent agent IDs. The plot highlights a clear hub-and-spoke structure, where a small number of wallets are connected to many agents. This means that agent registration is not evenly distributed across the ecosystem. Instead, many agents appear to be held or deployed by a limited group of high-volume owner wallets. This pattern is important because a mature agent economy would ideally show broader participation across many independent owners, rather than relying heavily on a small number of concentrated deployers.

Table 1 provides quantitative support for the network structure shown in Figure 4. The dataset contains 394 unique owner wallets for 10,000 registered agents, which already indicates that many agents are grouped under relatively few owners. The mean number of agents per owner is 25.38, but the median is only 3. This large gap between the mean and median shows that the ownership distribution is highly skewed: most wallets hold only a small number of agents, while a small number of wallets hold very large numbers of agents. The maximum value confirms this concentration, with the largest owner wallet holding 779 agents. The concentration metrics further reinforce this interpretation. The ownership Gini coefficient is 0.863, indicating a highly unequal distribution of agent ownership across wallets. The HHI value of 0.034 also suggests non-trivial concentration, although ownership is not monopolized by a single wallet. Together, the network plot and the concentration table show that ERC-8004 identity formation is still dominated by a small set of high-volume owners or deployers. This does not necessarily imply harmful centralization, since early-stage ecosystems often involve developer wallets, testing wallets, or infrastructure providers. However, it does indicate that ERC-8004 adoption is not yet broadly distributed across a large base of independent participants.

Overall, the owner-to-agent network shows that ERC-8004 identity registration is active but unevenly distributed. The presence of 10,000 registered agents may suggest rapid adoption at first glance, but the ownership structure reveals a more concentrated pattern. This finding supports the broader argument of the paper: early ERC-8004 activity is visible on-chain, but the ecosystem has not yet reached a mature stage of broad and distributed participation.

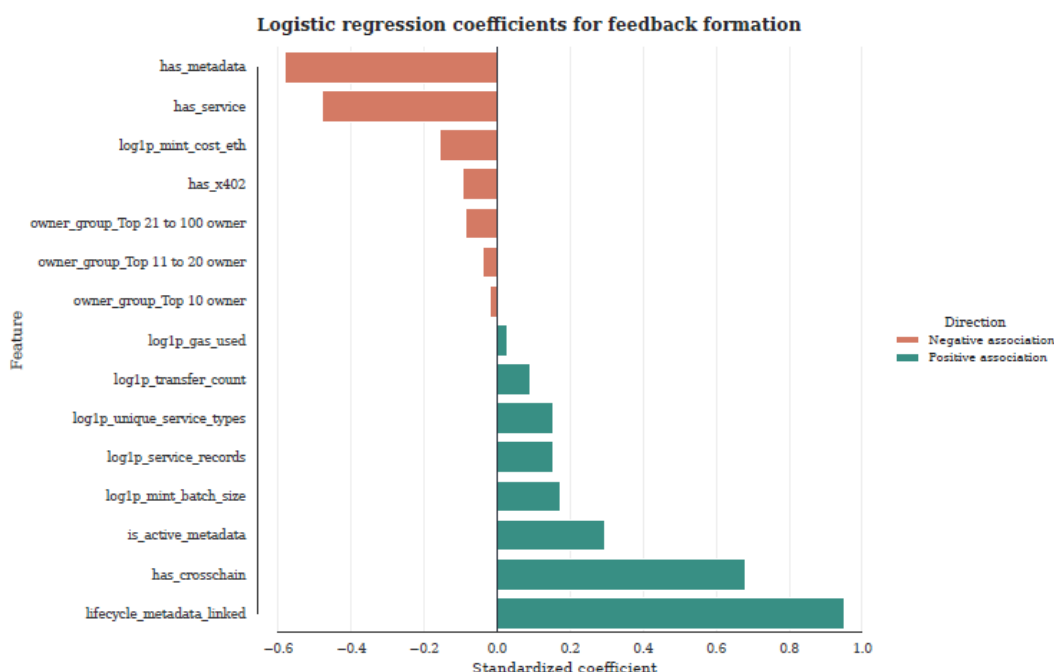


Figure 3. Standardized logistic regression coefficients associated with feedback formation. The model is exploratory and should be interpreted as evidence of association rather than causality.




Figure 4. Owner-to-agent network for ERC-8004 agents. Orange nodes represent owner wallets and blue nodes represent agent identities. The visualization shows that many agent identities are connected to a small number of high-volume owner wallets, indicating concentrated ownership in the early ERC-8004 ecosystem.

Table 1. Ownership concentration among ERC-8004 agent owner wallets.

Metric	Value
Unique owner wallets	394
Ownership HHI	0.034283
Ownership Gini	0.863005
Mean agents per owner	25.380711
Median agents per owner	3
Maximum agents owned by one wallet	779

5.2.2 Feedback Client-to-Agent Network

Figure 5 examines the structure of reputation activity in the ERC-8004 ecosystem. Panel (a) visualizes the relationship between feedback clients and agents that received reputation feedback. Green nodes represent feedback client addresses, while blue nodes represent agent identities. The network shows a highly uneven structure: a small number of feedback clients are connected to many agents, while most agents have limited or no visible feedback relationships. This indicates that reputation formation is not yet broadly distributed across many independent reviewers.

Figure 5, Panel (b) complements the network visualization by showing the most active feedback clients. The concentration is substantial. Feedback records are observed for only 628 agents, equivalent to 6.28% of the 10,000 registered agents. Moreover, the majority of feedback activity originates from a single client address, which contributes 645 records, or 65.8% of all feedback entries. The top five feedback clients together account for 92.4% of feedback records. This pattern suggests that the reputation layer is currently driven by a very small number of high-activity clients rather than by broad ecosystem participation.

Table 2. Summary of ERC-8004 reputation feedback records.

Metric	Value
Total feedback records	980
Agents with feedback records	628
Unique feedback clients	197
Largest client feedback records	645
Largest client share (%)	65.82
Feedback client HHI	0.435800
Feedback client Gini	0.782513
Agents with more than one feedback client	18

This concentration is also reflected in the feedback-provider concentration metrics. The Herfindahl-Hirschman Index for feedback clients is 0.436, while the Gini coefficient is 0.783. These values indicate a highly unequal distribution of feedback activity across client addresses. In practical terms, this means that although the ERC-8004 Reputation Registry is being used, its early use remains narrow and centralized around a few dominant feedback providers. From an operational readiness perspective, this is important because reputation systems become more meaningful when feedback comes from diverse, repeated, and independent

counterparties. A reputation layer dominated by a small number of clients may reflect early testing, developer-driven interaction, or platform-specific feedback processes rather than mature market-wide trust formation.

Overall, the feedback network reinforces the broader finding that ERC-8004 has not yet developed a mature and distributed reputation layer. The existence of feedback records indicates that the Reputation Registry is operational at a basic level, but the limited agent coverage and extreme client concentration suggest that reputation formation remains early-stage. This finding is consistent with the paper's broader argument that ERC-8004 currently functions more strongly as an identity and registration infrastructure than as a fully developed agent economy with distributed trust formation.

5.2.3 Wallet Transfer Network

Figure 6 visualizes wallet-to-wallet transfer relationships observed in the transfer history table. Purple nodes represent wallets and directed edges represent observed transfers between wallets. The graph provides a structural view of agent ownership movement, but it should not be interpreted as evidence of agent autonomy. Transfers may be initiated by human operators, scripts, marketplaces, or infrastructure processes. The main insight is that transfer activity is observable but limited in interpretability. It provides evidence of movement in ownership or custody, but not direct evidence that agents are performing independent economic actions.




Figure 5. Feedback concentration in the ERC-8004 reputation layer. Panel (a) shows the network between feedback clients and agents, where green nodes represent feedback client addresses and blue nodes represent agents. Panel (b) reports the most active feedback clients. Together, the figures show that reputation activity is concentrated around a small number of client addresses, suggesting that feedback formation remains limited and uneven in the early ERC-8004 ecosystem.




Figure 6. Wallet transfer network for ERC-8004 agent identities. Purple nodes represent wallets and directed edges represent observed transfer relationships. The figure captures ownership movement across wallets, but transfer activity should not be interpreted as proof of autonomous agent behavior.

5.2.4 Combined Agent Evidence Network

Figure 7 presents the combined agent evidence network, which integrates several observable readiness layers into a single structural view. In this network, agent nodes are connected to owner wallets, feedback clients, service endpoint domains, and cross-chain namespaces whenever these relationships are observed in the dataset. This visualization is useful because it moves beyond counting agents in separate readiness categories and instead shows how different evidence layers cluster around specific agents.

The network reveals that richer operational evidence is concentrated among a small subset of agents. Only a limited number of agents are connected simultaneously to multiple evidence sources, such as ownership, feedback, service domains, and cross-chain registration. Most registered agents do not appear prominently in this richer network because they lack one or more of the observable layers required for multi-dimensional participation. This pattern is consistent with the readiness funnel and supports the main empirical finding that ERC-8004 adoption is registration-heavy but operationally shallow. In other words, while the registry contains many agent identities, only a small group of agents shows evidence of broader ecosystem engagement.




Figure 7. Combined agent evidence network for ERC-8004 agents. The network links agents to owner wallets, feedback clients, service endpoint domains, and cross-chain namespaces. The figure shows that richer operational evidence is concentrated among a small subset of agents.

Figure 8 summarizes both the distribution of declared service types and the endpoint domains used in service metadata. Figure 8, Panel (a) shows that web endpoints dominate the service layer, accounting for 56.3% of service records. Other service categories appear much less frequently: A2A accounts for 8.0%, EMAIL for 7.1%, OASF for 7.1%, MCP for 2.7%, X402 for 1.8%, and HTTP-JSON-RPC for only 0.9%. Since MCP refers to the Model Context Protocol and A2A refers to agent-to-agent interaction, their low representation suggests that most service declarations do not yet reflect a mature inter-agent communication layer.

Figure 8, Panel (b) shows the endpoint-domain distribution. The most common domain is *marketplace.olas.network*, followed by smaller clusters of domains such as *github.com*, *x.com*, *useotto.xyz*, *acp.virtuals.io*, and other project-specific endpoints. This distribution suggests that service metadata is not yet widely diversified across many independent infrastructure providers. Instead, a small number of domains account for a large share of declared endpoints. The service layer therefore appears to be present but narrow, with many service records pointing to web or platform-specific endpoints rather than to a broad ecosystem of operational agent interfaces.




Figure 8. Declared service infrastructure among ERC-8004 agents with service records. Panel (a) reports the distribution of service types, showing that web endpoints dominate the observed service layer. Panel (b) reports the most common endpoint domains, showing that service declarations are concentrated around a small number of domains. Together, the figures indicate that the service layer is present but remains limited in diversity and depth.

Taken all things together, the combined evidence network and service-layer figures provide a structural view of ERC-8004's current maturity. The combined network shows that only a small subset of agents connects across multiple observable readiness layers. The service-type and endpoint-domain distributions show that even among agents with service records, declared services are concentrated in a small number of formats and domains. These findings reinforce the interpretation that ERC-8004 currently functions more strongly as an identity-registration layer than as a mature agent economy with broad service discoverability, distributed trust formation, and multi-party interaction.

Overall, the network analysis complements the readiness funnel by showing not only how many agents reach each layer, but also how those layers are connected. The owner-to-agent network indicates concentrated identity ownership, the feedback client-to-agent network indicates concentrated reputation activity, and the combined evidence network shows that multi-layer readiness is limited to a small portion of the ecosystem. This does not mean that ERC-8004 lacks potential. Rather, it shows that the current empirical footprint of ERC-8004 remains early-stage: identity registration is visible at scale, but the service, reputation, and interaction layers needed for a mature agent economy remain sparse and unevenly distributed.

5.3 Limitations and Interpretation Boundaries

The findings of this study should be interpreted within several important limitations. First, the analysis is based on an early snapshot of ERC-8004 deployment on Ethereum mainnet, covering agent registrations from January to April 2026. Since ERC-8004 remains an emerging standard, the observed registration-heavy and operationally shallow structure should be interpreted as a baseline assessment of early adoption rather than a definitive characterization of the long-term ecosystem. The distribution of metadata, services, reputation, cross-chain activity, and other forms of operational evidence may change as the specification matures, tooling improves, and additional agents are registered.

Second, operational readiness in this study is measured using observable on-chain records and metadata that could be retrieved and parsed from the available dataset. These records allow us to examine registration, metadata availability, service declarations, feedback, transfers, cross-chain registration, and network structure, but they do not provide direct observation of all off-chain agent activity. Consequently, the absence of observable service or reputation evidence should not necessarily be interpreted as evidence that an agent has no off-chain functionality. Although the dataset contains 10,000 registered agents, only a small subset has fully parsed metadata and declared services. Inferences regarding agent capabilities, service readiness, and endpoint diversity are therefore necessarily limited. Similarly, feedback records are available for only a minority of agents and are highly concentrated among a small number of client addresses, limiting the extent to which current feedback activity can be interpreted as evidence of broad and distributed reputation formation.

Finally, the regression and network analyses are descriptive and exploratory. The logistic regression identifies agent characteristics associated with receiving feedback, but it does not establish causal relationships. For example, metadata linkage may be associated with a greater likelihood of receiving feedback because it improves agent visibility, but it is also possible that more actively developed or promoted agents are simultaneously more likely to provide metadata and receive feedback. Similarly, the network analysis identifies concentration and clustering across ownership, feedback, services, and cross-chain evidence, but it cannot determine the off-chain mechanisms that generate these structures. Transfer activity should also be interpreted cautiously because ownership movement does not provide direct evidence of autonomous economic behaviour. These limitations do not alter the central descriptive finding that early ERC-8004 activity is concentrated primarily at the identity-registration layer, but they define the boundaries within which the observed patterns should be interpreted. The findings therefore provide an empirical baseline for monitoring whether ERC-8004 progresses toward broader service provision, distributed reputation formation, and richer operational agent infrastructure as the ecosystem develops.

6. Conclusion

This paper examined whether ERC-8004 agents on Ethereum demonstrate operational readiness beyond basic identity registration. Using a dataset of the first 10,000 ERC-8004 agents, we constructed an agent-level feature table covering identity status, metadata, service declarations, reputation feedback, transfers, cross-chain registration, mint economics, and network relationships. We then proposed a layered readiness framework to assess whether registered agents move from on-chain identity toward discoverability, service exposure, reputation formation, and broader ecosystem participation.

The findings show that early ERC-8004 adoption is registration-heavy but operationally shallow. While the identity layer has reached visible scale, higher readiness layers remain limited. Only a small share of agents has parsed metadata, service records, feedback, or cross-chain registrations, and only 19 from 10,000 agents combine metadata, services, feedback, and cross-chain registration. Ownership and reputation activity are also highly concentrated, with a small number of wallets and feedback clients accounting for a large share of observed activity. The network analysis reinforces this conclusion by showing that richer operational evidence clusters around a small subset of agents rather than being broadly distributed across the ecosystem.

These results suggest that ERC-8004 currently functions more strongly as an identity-registration layer than as a mature agent economy. This does not diminish the importance of the standard. On the contrary, ERC-8004 provides a useful foundation for representing agents, linking them to metadata, and building reputation.

However, the empirical evidence indicates that the transition from agent identity to agent economy remains incomplete. A functioning agent economy will require more than registered identities. It will require reliable metadata, active service endpoints, diverse reputation formation, and distributed relationships among agents, users, wallets, and service providers.

The paper contributes to the literature on decentralized AI agents and blockchain-based agent economies by offering an empirical framework for measuring operational readiness. Rather than assuming that agent registration implies agent functionality, the framework distinguishes between identity creation and observable evidence of use. This distinction is important for researchers, developers, and standard designers because it highlights where current infrastructure is already visible and where substantial gaps remain. As ERC-8004 and related standards evolve, ongoing empirical measurement will be essential for tracking whether blockchain-registered agents become active participants in decentralized economies or remain primarily static identity tokens.

Appendix

List of Abbreviations

Abbreviation	Definition
A2A	Agent2Agent Protocol
ACP	Agent Commerce Protocol
AI	Artificial Intelligence
CSV	Comma-Separated Values
DeFi	Decentralized Finance
ENS	Ethereum Name Service
ERC-721	Ethereum Request for Comments 721 (Non-Fungible Token Standard)
ERC-8004	Ethereum Request for Comments 8004 (Trustless Agents)
HHI	Herfindahl-Hirschman Index
HTTP	Hypertext Transfer Protocol
JSON-RPC	JavaScript Object Notation Remote Procedure Call
MCP	Model Context Protocol
OASF	Open Agentic Schema Framework
x402	HTTP 402-based payment protocol

Data Availability Statement

The data used in this study are publicly available from Harvard Dataverse. This study uses the third-party dataset by Liu, Yulin (2026), *Replication Data for: "A dataset of the first 10000 blockchain registered AI agents on Ethereum"*, Harvard Dataverse, V1, DOI: [10.7910/DVN/HJZW8Q](https://doi.org/10.7910/DVN/HJZW8Q). The dataset is available at <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/HJZW8Q>. The present paper provides an independent empirical analysis and interpretation of this dataset; it does not introduce a new primary dataset.

Acknowledgements

The authors gratefully acknowledge Yulin Liu for making the dataset Replication Data for: “*A dataset of the first 10000 blockchain registered AI agents on Ethereum*” publicly available through Harvard Dataverse. The availability of this dataset made the independent empirical analysis conducted in this study possible.

Declaration of the Use of Generative AI

During the preparation and revision of this manuscript, the authors used ChatGPT (OpenAI, GPT-5.6 Sol) solely for English-language polishing and grammar correction. The authors reviewed and revised all AI-assisted edits and take full responsibility for the accuracy, integrity, and final content of the manuscript.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- Al Jasem, M. S., De Clark, T., & Shrestha, A. K. (2025). Toward decentralized intelligence: A systematic literature review of blockchain-enabled AI systems. *Information*, 16(9), Article 765. <https://doi.org/10.3390/info16090765>
- Ante, L. (2026). Autonomous AI agents in decentralized finance: Market dynamics, application areas, and theoretical implications. *Technological Forecasting and Social Change*, 228, Article 124669. <https://doi.org/10.1016/j.techfore.2026.124669>
- Cao, L. (2022). Decentralized AI: Edge intelligence and smart blockchain, metaverse, Web3, and DeSci. *IEEE Intelligent Systems*, 37(3), 6–19. <https://doi.org/10.1109/MIS.2022.3181504>
- Chaffer, T. J., Goins, I., Von, C., Okusanya, B., Cotlage, D., & Goldston, J. (2024). *Decentralized governance of autonomous AI agents*. arXiv. <https://arxiv.org/abs/2412.17114>
- Dafflon, J., Breidenbach, B., & Kistner, M. (2023). *ERC-6551: Non-fungible token bound accounts* (Ethereum Improvement Proposal ERC-6551). <https://eips.ethereum.org/EIPS/eip-6551>
- De Rossi, M., Crapis, D., Ellis, J., & Reppel, E. (2025). *ERC-8004: Trustless agents* (Ethereum Improvement Proposal EIP-8004). <https://eips.ethereum.org/EIPS/eip-8004>
- Entriiken, W., Shirley, D., Evans, J., & Sachs, N. (2018). *ERC-721: Non-fungible token standard* (Ethereum Improvement Proposal EIP-721). <https://eips.ethereum.org/EIPS/eip-721>
- Fan, S., & Min, T. (2026). Web3Agent: Automating on-chain operations via natural language interfaces. *ACM Transactions on the Web*, 20(1), Article 9. <https://doi.org/10.1145/3777446>
- Fan, T., Yang, Y., Jiang, Y., Zhang, Y., Chen, Y., & Huang, C. (2025). *AI-Trader: Benchmarking autonomous agents in real-time financial markets*. arXiv. <https://arxiv.org/abs/2512.10971>
- Karim, M. M., Van, D. H., Khan, S., Qu, Q., & Kholodov, Y. (2025). AI agents meet blockchain: A survey on secure and scalable collaboration for multi-agents. *Future Internet*, 17(2), Article 57. <https://doi.org/10.3390/fi17020057>

Kirana, M. Y., & Havidz, S. A. H. (2020). Financial literacy and mobile payment usage as financial inclusion determinants. In *2020 International Conference on Information Management and Technology (ICIMTech)* (pp. 905–910). IEEE.

Liu, Y. (2026a). *A dataset of early blockchain-registered AI agents on Ethereum*. arXiv. <https://arxiv.org/abs/2604.22652>

Liu, Y. (2026b). *Replication data for: A dataset of the first 10000 blockchain registered AI agents on Ethereum* [Data set]. Harvard Dataverse. <https://doi.org/10.7910/DVN/HJZW8Q>

Mafrur, R. (2025). AI-based crypto tokens: The illusion of decentralized AI? *IET Blockchain*, 5, Article e70015. <https://doi.org/10.1049/blc2.70015>

Ranjan, R., Gupta, S., & Singh, S. N. (2025). *Loka protocol: A decentralized framework for trustworthy and ethical AI agent ecosystems*. arXiv. <https://arxiv.org/abs/2504.10915>

Romandini, N., Mazzocca, C., Otsuki, K., & Montanari, R. (2025). SoK: Security and privacy of AI agents for blockchain. In *2025 7th International Conference on Blockchain Computing and Applications (BCCA)* (pp. 708–720). IEEE. <https://doi.org/10.1109/BCCA66705.2025.11229689>

Tomašev, N., Franklin, M., Leibo, J. Z., Jacobs, J., Cunningham, W. A., Gabriel, I., & Osindero, S. (2025). *Virtual agent economies*. arXiv. <https://arxiv.org/abs/2509.10147>

Xu, M. (2026). *The agent economy: A blockchain-based foundation for autonomous AI agents*. arXiv. <https://arxiv.org/abs/2602.14219>

Yu, J., Zhao, A., & Sui, D. (2026). *Paper agents, paper gains: An empirical analysis of DeFi investment agents*. arXiv. <https://arxiv.org/abs/2605.29174>