



NATIONAL CENTER FOR UNDERSTANDING FUTURE
TRAVEL BEHAVIOR AND DEMAND

Final Project Report

**Blockchain Application on Smart
Transportation Systems**

BY

Mahdiah Allahviranloo

Email: mallahviranloo@ccny.cuny.edu

Nikhita Kannam

Email: nkannam000@citymail.cuny.edu

Grove School of Engineering, Department of Civil Engineering
The City College of New York
160 Convent Avenue, New York, NY, 10031

May 2026

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. N/A	2. Government Accession No. N/A	3. Recipient's Catalog No. N/A	
4. Title and Subtitle Blockchain Application on Smart Transportation Systems		5. Report Date May 30, 2026	
		6. Performing Organization Code N/A	
7. Author(s) Mahdiah Allahviranloo, https://orcid.org/0009-0001-8089-5267 Nikhita Kannam, https://orcid.org/0009-0000-1085-2483		8. Performing Organization Report No. N/A	
9. Performing Organization Name and Address The city college of New York 160 Convent Avenue, New York, NY, 10031		10. Work Unit No. (TRAIS) N/A	
		11. Contract or Grant No. 69A3552344815 and 69A3552348320	
12. Sponsoring Agency Name and Address U.S. Department of Transportation, University Transportation Centers Program, 1200 New Jersey Ave, SE, Washington, DC 20590		13. Type of Report and Period Covered Final Report, 2024-2026	
		14. Sponsoring Agency Code USDOT OST-R	
15. Supplementary Notes N/A or Conducted in cooperation with the U.S. Department of Transportation, Federal Highway Administration.			
16. Abstract Blockchain technology has recently emerged as a transformative innovation, offering decentralized and secure solutions that enhance transparency and efficiency across multiple industries. This research investigates the application of blockchain in urban transportation systems, focusing on the potential to unify payments across multiple transportation providers. A simulation platform was developed leveraging Ethereum smart contracts, Ganache, and Python to model decentralized transactions for MTA and Uber services. The simulation incorporates realistic features including multi-modal trip chaining, time-of-day demand patterns with surge pricing, and a utility-based discrete choice framework. A choice model compares rider behavior across transit modes and payment methods under varying cryptocurrency acceptance caps. Key metrics collected include transaction volumes, fare distributions, denial rates, and market-responsive crypto demand patterns. The findings demonstrate that provider crypto caps function as a payment-rationing mechanism rather than a demand-shaping tool, and that crypto demand is highly elastic with respect to market conditions. These results highlight the viability of blockchain technology for urban transportation payment systems and contribute to the growing research on smart city infrastructure.			
17. Key Words Blockchain, smart contracts, transportation payment systems, cryptocurrency adoption, utility-based choice model, decentralized ledger, smart cities, multi-modal transit		18. Distribution Statement No restrictions.	
19. Security Classif.(of this report) Unclassified	20. Security Classif.(of this page) Unclassified	21. No. of Pages 23	22. Price N/A

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, under Grant No. 69A3552344815 and 69A3552348320 from the U.S. Department of Transportation's University Transportation Centers Program. The U.S. Government assumes no liability for the contents or use thereof.

ACKNOWLEDGMENTS

This research was partially supported by the National Center for Understanding Future Travel Behavior and Demand (TBD), a National University Transportation Center sponsored by the U.S. Department of Transportation (USDOT) under grant numbers 69A3552344815 and 69A3552348320. The authors would like to thank the TBD National Center and USDOT for their support of university-based research in transportation, particularly for the funding provided for this project. The authors also extend their thanks to Laure Vatin for her valuable contributions to the work presented in this report.

TABLE OF CONTENTS

<i>EXECUTIVE SUMMARY</i>	1
<i>INTRODUCTION</i>	2
<i>LITERATURE REVIEW</i>	3
Blockchain in Transportation and Smart Cities	3
Peer-to-Peer Ride-Sharing.....	4
Blockchain in Supply Chains	4
<i>BLOCKCHAIN TECHNOLOGY</i>	5
Bitcoin	5
Ethereum	7
<i>SIMULATION PLATFORM</i>	7
Solidity	8
Ganache.....	8
Truffle.....	9
<i>TRANSPORTATION CHOICE MODEL</i>	9
Model Setup	9
Mode Choice	10
Payment Choice	10
Crypto Budgets (Daily and Monthly).....	11
Crypto Price and EMA50	11
Provider Crypto Caps	12
Simulation Process	12
Scenario Design	12
<i>RESULTS</i>	13
Baseline Outcomes.....	13
Market Context.....	13
<i>DISCUSSION</i>	15
<i>CONCLUSION</i>	16
<i>REFERENCES</i>	17

LIST OF TABLES

Table 1	Comparison of Selected Cryptocurrencies	2
Table 2	Main Variables Used in the Simulation	10
Table 3	Baseline Simulation Outcomes	13

LIST OF FIGURES

Figure 1	Comparison of Traditional Banking and Cryptocurrency Transaction Processes	3
Figure 2	Transactions in Conventional vs. Blockchain-Based Supply Chain Systems	5
Figure 3	Bitcoin Mining Target Calculation	6
Figure 4	SHA-256 Hash Algorithm Process	6
Figure 5	Blockchain Simulation Platform Workflow	8
Figure 6	Solidity Smart Contract Code — Backend Model	8
Figure 7	ETH Price Relative to its 50-day EMA over the Simulation Window	13
Figure 8	Daily Payment Outcomes for MTA and Uber	14

EXECUTIVE SUMMARY

The proliferation of blockchain technology and decentralized digital currencies presents both opportunities and challenges for urban transportation systems. As cryptocurrency adoption grows, transit agencies and mobility providers face fundamental questions about whether and how to integrate cryptocurrency payment options into their fare collection infrastructure. This report presents a rigorous analytical investigation into blockchain-based payment integration for urban transit, conducted under the U.S. Department of Transportation's University Transportation Centers Program.

The research develops a utility-based discrete choice framework to model joint mode-payment decisions of urban travelers. The framework couples a binary logit model for mode selection between public transit (MTA) and ride-hailing (Uber)—parameterized by travel time, monetary cost, and service reliability—with a conditional payment-method model that determines crypto-versus-fiat selection based on individual crypto openness, real-time Ethereum market conditions (spot price relative to a 50-day exponential moving average), and budget feasibility constraints. A distinguishing feature is the incorporation of provider-side acceptance caps—daily limits on cryptocurrency volume that each operator will process—creating a supply-constrained system where excess demand is rationed first-come-first-served. For ride-hailing, the model disaggregates this constraint into a platform-level cap and individual driver-side tolerances, capturing the two-sided market structure where both platform and drivers must consent to crypto acceptance.

The simulation platform is built on Ethereum blockchain infrastructure using Solidity smart contracts deployed via Truffle to a local Ganache network, with Python's Web3.py orchestrating transaction execution. This architecture demonstrates that decentralized ledger technology can serve as the backbone for multi-provider fare settlement, providing transparent, auditable, and tamper-proof transaction records grounded in historical ETH/USD price data.

Three principal findings emerge. First, the mode split between public transit and ride-hailing is invariant to crypto cap levels, demonstrating that payment constraints do not induce mode shifting—distinguishing them from traditional supply-side constraints. Second, crypto payment demand is highly elastic with respect to market conditions: riders exhibit "paper-gains spending" behavior when the spot price exceeds its moving average, creating demand volatility exogenous to the transportation system. Third, provider denial rates respond independently to their respective caps, enabling each operator to calibrate crypto acceptance without cross-modal spillover. These findings establish that crypto caps function as a pure payment-rationing mechanism rather than a demand-shaping tool—a distinction with significant implications for transit policy and smart city governance. The report concludes with a research agenda encompassing endogenous mode-payment feedback, stablecoin and Layer-2 integration, adaptive cap mechanisms, welfare analysis, agent-based adoption modeling, consortium blockchain architectures, and equity considerations.

INTRODUCTION

The evolution of cryptocurrency since its inception in 2009 has been remarkable. Initially speculated to be a short-lived scheme, cryptocurrency has since proven its potential for broad application across diverse domains. Cryptocurrency constitutes a form of digital payment independent of traditional banking systems. It operates on a decentralized ledger powered by blockchain technology—a distributed data structure that stores all transactions across a network of computers, where each block contains a list of validated transactions [1]. The distributed architecture provides inherent security and transparency because every network node maintains a complete copy of the blockchain. Most cryptocurrencies employ consensus mechanisms such as Proof of Work (PoW) or Proof of Stake (PoS) to validate transactions and append new blocks to the chain [2].

Table 1 Comparison of Selected Cryptocurrencies

<i>Cryptocurrency</i>	<i>Use Case</i>	<i>Supply Cap</i>	<i>Consensus Mechanism</i>	<i>Avg Transaction Time</i>	<i>Algorithm</i>
<i>Bitcoin (BTC)</i>	Digital Currency	21 Million	Proof of Work (PoW)	10 minutes	SHA-256
<i>Ethereum (ETH)</i>	Smart contracts and dApps	No limit	Proof of Stake (PoS)	12–14 seconds	SHA-3 (Keccak-256)
<i>XRP</i>	Fast, low-cost cross-border payments	No limit	XRP Ledger Consensus Protocol	4 seconds	Ripple Protocol Algorithm

Cryptocurrency differs from traditional banking in several fundamental ways. Traditional banking systems require intermediaries such as banks or payment processors, whereas cryptocurrency enables direct peer-to-peer transactions without an intermediary. Furthermore, cryptocurrency transactions typically complete within minutes, compared to the days often required for traditional bank transfers to settle. These characteristics—disintermediation, speed, and transparency—make blockchain technology particularly compelling for sectors that depend on high-volume, low-value transactions processed across multiple stakeholders, such as urban transportation. Blockchain technology, while predominantly associated with cryptocurrency, is increasingly being adapted across diverse sectors, and transportation systems are no exception. Understanding blockchain's potential applications and benefits in addressing future urban challenges constitutes an emerging field of research that remains under-investigated. What makes blockchain particularly attractive for smart cities is its design scheme and underlying protocols. The decentralized computational architecture enhances the reliability of data transmission across network nodes. Each data transaction between nodes is recorded with a unique identifier and validated by consensus among all agents within the system. Transaction transparency mitigates the risk of propagating inaccurate information throughout the network. Because the system is decentralized, it eliminates single-point-of-failure vulnerabilities—there is no central target for attack—thereby enhancing security and reducing the risk of fraud.

Despite these advantages, a critical gap exists in the literature: while blockchain's technical feasibility for transportation payments has been demonstrated in pilot programs, no rigorous analytical framework has been developed to model how cryptocurrency payment integration interacts with traveler behavior at the system level. Specifically, the question of how provider-side constraints on cryptocurrency acceptance affect rider payment choices, mode selection, and overall system performance remains unanswered. This research addresses that gap by developing a utility-based discrete choice framework that jointly models mode and payment decisions, incorporates market-responsive cryptocurrency demand dynamics, and evaluates the system-level consequences of provider acceptance caps—contributing to the growing body of knowledge on smart city infrastructure and decentralized transportation finance. The remainder of this report is organized as follows. Section 2 reviews the existing literature on blockchain applications in transportation, ride-sharing, and supply chains. Section 3 provides technical background on Bitcoin and

Ethereum. Section 4 describes the simulation platform developed for this project — a local Ethereum blockchain environment built with Solidity smart contracts, Truffle for deployment, and Ganache for emulation — which serves as the transaction-recording backend. Section 5 then presents a transportation choice model that generates realistic travel and payment decisions for a population of riders; these decisions are executed as blockchain transactions on the platform described in Section 4. The interaction between the behavioral model (demand side) and the blockchain infrastructure (supply side) enables the study of how provider-level cryptocurrency acceptance caps affect payment outcomes. Sections 6 and 7 report and discuss the simulation results, and Section 8 outlines future research extensions.

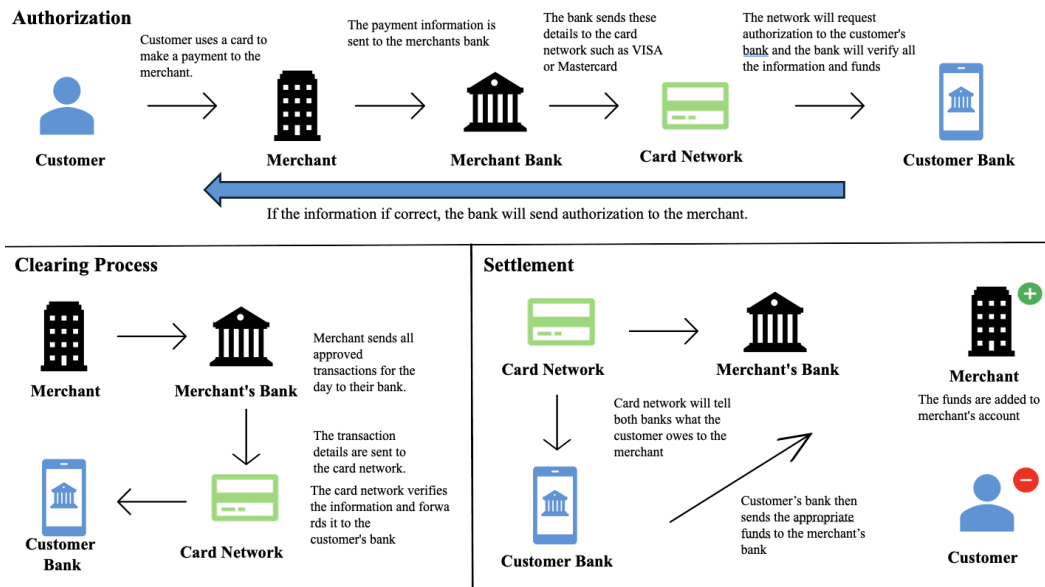


Figure 1 Comparison of Traditional Banking and Cryptocurrency Transaction Processes

LITERATURE REVIEW

Blockchain in Transportation and Smart Cities

Blockchain technology has been increasingly recognized for its potential in supply chain management, transportation, and smart city applications. The technology offers solutions for improving transportation ticketing systems and mobility services across numerous use cases. Rather than tracking physical goods, blockchain can track the tickets that passengers use to ride buses or trains, analogous to bills of lading that accompany shipments of goods with tracking information.

Cities such as São Paulo and Rio de Janeiro have incorporated Bitcoin as a payment option for public transit to reduce dependency on fiat currencies, which are prone to inflationary fluctuation [5]. In 2019, the Netherlands implemented a ticketing system based on a distributed ledger using a VAI token through the VMC.ai platform. The VAI token acts as a digital currency while VMC.ai provides the blockchain infrastructure that facilitates the decentralized ticketing process.

Currently, most transit systems rely on radio-frequency identification (RFID) or QR-code-based card payments. Blockchain can enhance these systems by ensuring that transactions are secure and tamper-proof [5]. This application demonstrates blockchain's ability to improve transparency and security in transportation systems. Beyond ticketing, blockchain's potential extends to ride-sharing systems and Mobility as a Service (MaaS) platforms.

Peer-to-Peer Ride-Sharing

Peer-to-peer ride-sharing frameworks built on blockchain aim to eliminate intermediaries and provide a more equitable system for both drivers and passengers. These platforms use smart contracts written in Solidity, deployed on the Ethereum blockchain, to enable secure and automated transactions between users [9]. The process involves a passenger requesting a ride, with the request published on the blockchain and then matched with an appropriate driver. The driver proposes a price, and the passenger selects the most suitable option. Smart contracts handle deposits for both parties, verify transaction terms (ensuring both parties fulfill their obligations), and manage deposit releases, refunds, and partial payments during the ride [9]. Upon journey completion, the smart contract validates the trip and processes payment. Supporting tools include Google Maps API and ReactJS for user interface functionality, with MetaMask providing cryptocurrency wallet integration for logins and ride approvals [9]. These systems generate tamper-proof transaction records that build trust among users by leveraging blockchain's immutability.

Blockchain also integrates with Mobility as a Service (MaaS), a platform that consolidates various transportation modes—trains, taxis, and buses—into a unified service. Blockchain enables a secure, decentralized system where transportation providers directly offer routes and tickets to passengers without intermediaries. Passengers select their preferred route, and the selection converts into a smart contract that defines the ride terms. Because of blockchain's decentralized nature, all parties can view and confirm transactions, creating full transparency [6]. Beyond public transportation services, blockchain also addresses challenges in supply chain systems related to traceability and transparency.

Blockchain in Supply Chains

Supply chain systems face significant challenges including lack of transparency, traceability, and accountability. A supply chain involves tracking extensive information—logs, payments, deliveries, and more—making it difficult to trace events and investigate issues when they arise. Customers cannot accurately judge the true value of products and services because supply chains lack transparency. Blockchain mitigates these issues by allowing companies to register information such as transfer of goods, identify parties involved, and record payment details, dates, and locations on an immutable ledger.

Capturing the Details of a Simple Transaction: Conventional vs. Blockchain Systems

The financial ledgers and enterprise resource planning systems now used don't reliably allow the three parties involved in a simple supply-chain transaction to see all the relevant flows of information, inventory, and money. A blockchain system eliminates the blind spots.



Figure 2 Transactions in Conventional vs. Blockchain-Based Supply Chain Systems

These improvements in transparency and traceability demonstrate blockchain's core value proposition. To understand how these capabilities can be harnessed for transportation payment systems, the following section examines the underlying technology—starting with Bitcoin's foundational architecture and progressing to Ethereum's smart contract platform.

BLOCKCHAIN TECHNOLOGY

Blockchain operates as a decentralized ledger that records transactions across a network, ensuring security and transparency. Understanding its foundational concepts is essential for exploring potential applications beyond cryptocurrency, including transportation systems and supply chains. This section examines the key components of blockchain—the structure of a block, the role of cryptographic hashing, and the process of validating and appending blocks to the chain—through an introduction to Bitcoin and Ethereum, the two leading cryptocurrency platforms.

Bitcoin

Cryptocurrency transactions rely on two cryptographic keys: a public key, which serves as the recipient's address, and a private key, used to authorize fund transfers. Initiating a transaction requires three elements: the sender, receiver, and amount. A digital signature verifies the transaction's authenticity. Once initiated, the transaction enters a memory pool until a miner or validator selects it for inclusion in a block, using the

network’s specific validation method to append it to the blockchain.

Understanding Bitcoin requires starting with the structure of a block—a data structure that records a list of transactions. The block header includes the version, timestamp, previous block hash, and a Merkle root. The Merkle root serves as a fingerprint for all transactions within the block, created by recursively hashing transactions in a binary tree structure. This mechanism not only secures the data but also facilitates block mining—the process of adding new blocks to the blockchain while maintaining decentralization. Bitcoin relies on the Proof of Work (PoW) consensus mechanism and the SHA-256 cryptographic algorithm for this process.

Proof of Work is a system where miners compete to solve complex mathematical puzzles to add new blocks to the blockchain. The first miner to find a valid solution receives a Bitcoin reward. Specifically, miners must find a hash value lower than a specified target using the SHA-256 algorithm by iteratively adjusting the nonce in the block header and rehashing until the result meets the target threshold. Mining difficulty is adjusted every 2,016 blocks to maintain an average block time of approximately 10 minutes. If the previous 2,016 blocks were mined faster than this target rate, the difficulty increases (the target is lowered) to make future mining more challenging.



Figure 3 Bitcoin Mining Target Calculation

The SHA-256 algorithm produces a 256-bit hash value from any given input. The process begins by converting the input message into binary representation and padding it to a total length of 512 bits. The padded message then undergoes two sequential processes: the Message Expander (ME) and the Message Compressor (MC).




Figure 4 SHA-256 Hash Algorithm Process

In the ME process, the padded message is divided into 64 data chunks; the first 16 remain unchanged, while

the subsequent 48 are computed using bitwise rotation ($S^n(x)$) and shift ($R^n(x)$) operations. The resulting chunks are then processed through the MC stage, which involves two steps: looping and hashing. Eight initial hash values are iteratively updated through 64 rounds of computation. The final 256-bit hash value is obtained by adding the initial hashes to the computed loop hashes [17]. Each cryptocurrency employs a different hashing algorithm; Ethereum uses the SHA-3 (Keccak-256) algorithm.

Once a miner produces a SHA-256 hash below the target, they may add their block to the blockchain and receive a mining reward. Because mining requires substantial computational energy and time, rewards serve as the primary incentive for miners to maintain the network. For Bitcoin, the reward derives from a coinbase transaction—the first transaction in each block—through which the miner claims the block reward. The current set reward is 3.125 BTC, supplemented by the cumulative transaction fees from all transactions in the mined block. When users initiate transfers, they may attach a transaction fee; higher fees increase the probability that miners will prioritize that transaction. Bitcoin maintains a hard supply cap of 21 million coins.

Ethereum

Ethereum transitioned to a Proof of Stake (PoS) consensus mechanism—termed “Ethereum 2.0”—in September 2022, replacing the energy-intensive Proof of Work model. Under PoS, validators are selected based on the amount of cryptocurrency (ETH) they stake as collateral, demonstrating commitment to the network. Becoming a validator requires depositing 32 ETH into a designated staking contract. Validators verify transactions, with one randomly selected every 12 seconds to propose a new block. Other validators review and attest to the proposed block; upon acceptance, it becomes part of the blockchain, and the proposing validator earns rewards in additional ETH [19].

Transaction execution in Ethereum PoS proceeds as follows: a user initiates a transaction, specifies the gas fee (a tip to the validator), and signs it with their private key. Ethereum’s execution client verifies that the sender holds sufficient ether and validates the transaction using the SHA-3 (Keccak-256) algorithm [12]. Once validated, the transaction enters the mempool and awaits selection by a validator. The chosen validator bundles mempool transactions into a block, computes the new blockchain state, and broadcasts the proposed block to the network. Other validators review the proposal to ensure compliance with protocol rules. If approved, the block is appended to the blockchain and the proposer receives an ether reward.

Ethereum employs a checkpoint system to finalize transactions. Validators vote on specific blocks (checkpoints) at the start of each epoch, which comprises 32 slots. Once two-thirds of validators agree on a checkpoint’s validity, the block is finalized and becomes immutable—reversal would require destroying a significant portion of staked ether [19].

Ethereum’s framework enables the development of decentralized applications (dApps), extending blockchain’s utility beyond financial transactions into transportation systems and supply chains. Incorporating blockchain in these domains addresses challenges in data management, transaction security, and system efficiency. For example, a smart city could deploy blockchain to track and manage transportation payments or monitor supply chain logistics. Data generated by users and transportation systems must be stored securely and managed efficiently to support informed decision-making and seamless operations.

SIMULATION PLATFORM

This project employs blockchain technology to enhance transaction management within transportation systems. By simulating real-world scenarios, the project demonstrates how blockchain can securely store and manage transaction data, improving transparency and efficiency for both agencies and users. The simulation incorporates a suite of tools and frameworks designed for blockchain development, enabling the

design, execution, and real-time monitoring of smart contracts. Specifically, the project utilizes Ganache, Truffle, Solidity, and Python (Web3.py) to write and interact with smart contracts and simulate blockchain transactions within transportation payment systems.

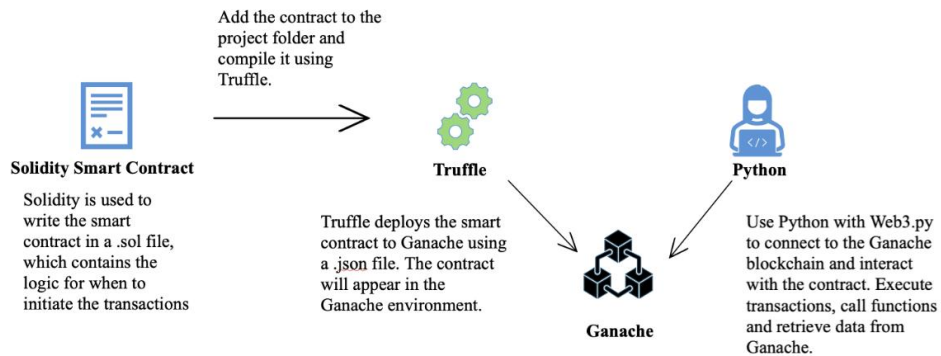


Figure 5 Blockchain Simulation Platform Workflow

Solidity

Solidity is a programming language specifically designed for the Ethereum Virtual Machine (EVM), used to write smart contracts that manage payments between parties—such as transit agencies and riders. In this project, smart contracts were developed in Solidity to define transaction rules for MTA and Uber payment simulations. The contracts specify elements including service providers, users, fare structures, and transaction logic. The source code is written in .sol files, developed in Visual Studio Code or Remix (an Ethereum-specific IDE).

ACCOUNT	BALANCE	TX COUNT	HEX
0xaF71B8c6E5eE126C70CE2E3F85d01094ce2027E7	100.00 ETH	0	0
0x54980de5E7d9003F6850F8F997AA17e3B46A2a0	100.00 ETH	0	1
0x58A524937687D4e948448efA33Ae09129903650	100.00 ETH	0	2
0xe3c66FCE767a8e96f6E1A3888B42fC973fF864A	100.00 ETH	0	3
0x0646a9c377b1036752fC18824223cd8108A095e	100.00 ETH	0	4
0x091c698f090244eb313931192E57C009E2B0cd2	100.00 ETH	0	5

Figure 6 Solidity Smart Contract Code Backend model

Ganache

Ganache is a local blockchain emulator that provides an environment for testing smart contracts and monitoring blockchain activity. Its interface comprises multiple sections: the Accounts tab lists Ethereum accounts automatically generated by Ganache, each initialized with a unique address and a balance of 100 ETH. The TX Count column tracks transactions per account. Additional tabs include Blocks (displaying every block added to the chain), Transactions (listing each transaction with its hash), and Contracts (showing deployed contract status and metadata).

In the simulation, the first two accounts served as transit service providers (MTA and Uber), while the remaining accounts represented passengers. This configuration models a realistic multi-provider transportation payment ecosystem where riders can transact using cryptocurrency.

Truffle

Truffle is a development framework for Ethereum that simplifies compiling and deploying smart contracts. After authoring a contract in Solidity, Truffle compiles it into bytecode and deploys it to the local blockchain environment. The deployment workflow involves initializing a Truffle project, placing the Solidity contract into the contracts directory, and configuring Ganache to connect to the project via the generated configuration file. Compilation uses ‘truffle compile’ and deployment uses ‘truffle migrate --network deployment’. These steps connect smart contracts to the Ganache network, rendering them visible and operational.

Python’s Web3.py library facilitates interaction with deployed contracts on Ganache. The library connects to Ganache via its URL endpoint, enabling data reads from the blockchain, transaction initiation, and contract state modifications. Transactions initiated through Python—such as payments for subway rides and ride-hailing services—are recorded on the blockchain. The contract’s Application Binary Interface (ABI), loaded from a JSON file, defines available functions and events. Using the deployed contract’s address, an instance is created for invoking functions and monitoring events.

Transaction data and associated state changes (account balance updates and new block additions) were monitored through Ganache and exported to JSON and CSV files for analysis. These files include transaction hashes, sender and receiver addresses, timestamps, gas usage, and block information.

Together, these four components — Solidity, Truffle, Ganache, and Python — form an integrated simulation backend capable of recording, validating, and querying transportation payment transactions on a local blockchain. However, a blockchain ledger alone does not generate realistic travel behavior. To study how cryptocurrency acceptance policies affect ridership and payment outcomes, a demand-side model is required — one that produces the daily mode choices and payment decisions that populate the blockchain with transactions. The following section introduces a utility-based transportation choice model that serves this role. Each simulated rider's trip — including mode selection, payment method, fare amount, and acceptance or denial — is ultimately executed as a smart contract transaction on the Ganache network, producing an immutable, auditable record of every payment event. The choice model thus acts as the demand engine, while the blockchain platform acts as the transaction ledger and enforcement layer.

TRANSPORTATION CHOICE MODEL

The blockchain simulation platform described in the previous section provides the infrastructure for recording and enforcing transportation payment transactions. This section develops the behavioral demand model that drives those transactions. Specifically, a discrete choice model is implemented for a population of urban riders comparing two modes of transportation: MTA subway (a single system-wide operator) and Uber (a multi-driver ride-hailing platform). Each day, riders select both a transportation mode and a payment method (fiat currency or cryptocurrency). Each rider maximizes personal utility while adhering to cryptocurrency spending budgets. Additionally, each transportation provider independently sets a cap on the total daily dollar volume of cryptocurrency it will accept — a constraint enforced at the smart contract level on the blockchain. The model is designed to elucidate how the rider population shifts between transit modes and payment methods as daily cryptocurrency caps vary, with every accepted and denied transaction recorded on the Ganache blockchain for auditability and analysis.

Model Setup

Let i denote a rider, d denote an Uber driver, and t denote a day. Each rider chooses between MTA and Uber,

then selects whether to pay with fiat or cryptocurrency. Fiat payments are subject to a daily cash budget set substantially higher than the crypto budget, so it rarely binds; however, neither payment method is unlimited. A crypto payment is accepted only if the rider retains sufficient crypto budget for that day and the provider still has remaining crypto-acceptance capacity.

Table 1 Main Variables Used in the Simulation

VARIABLE	MEANING
U_{imt}	Utility of mode m for rider i on day t
TT_{imt}	Travel time for rider i on mode m on day t (minutes)
C_{imt}	Cost of using mode m for rider i on day t (dollars)
R_m	Reliability score of mode m
κ_i	Crypto openness of rider i
P_{imt}	Probability that rider i chooses mode m on day t
P_t	Cryptocurrency price on day t
$EMA50_t$	50-day exponential moving average of the crypto price
B_{it}^d	Daily crypto budget of rider i on day t
B_{it}^m	Remaining monthly crypto budget of rider i on day t
B_{it}^s	Daily cash budget of rider i on day t
$K_{MTA,t}^C$	Total crypto cap accepted by MTA on day t
$K_{uber,t}^C$	Platform-level crypto cap accepted by Uber on day t

Mode Choice

The first component of the model is the rider’s transportation mode choice. Each rider selects between MTA and Uber based on mode-specific utility, which depends on travel time, cost, and reliability:

$$U_{imt} = \alpha^0 + \alpha^1 \cdot TT_{imt} + \alpha_2 \cdot C_{imt} + \alpha_3 \cdot R_m$$

The coefficient signs follow standard travel behavior: travel time and cost exert negative effects on utility (riders prefer shorter and cheaper trips), while reliability exerts a positive effect (riders prefer dependable service).

The probability of choosing MTA is computed using a binary logit model:

$$P_{\{i,MTA,t\}} = \exp(U_{\{i,MTA,t\}}) / [\exp(U_{\{i,MTA,t\}}) + \exp(U_{\{i,Uber,t\}})]$$

Payment Choice

After selecting a mode, the rider decides whether to pay with fiat or cryptocurrency. Let C denote crypto payment. The marginal probability that a rider pays in crypto on day t is:

$$P_{\{i,t\}}(C) = P_{\{i,MTA,t\}} \cdot P_{\{i,t\}}(C | MTA) + P_{\{i,uber,t\}} \cdot P_{\{i,t\}}(C | Uber)$$

The overall probability of paying in crypto depends on both the mode choice and the conditional payment choice. A rider may prefer crypto for one mode but not the other.

Since actual payment-choice data are unavailable, the crypto payment probability is generated within the simulation. Each rider is assigned a crypto openness value, κ_i , representing their general willingness to use cryptocurrency—higher values indicate greater openness.

The model also incorporates daily crypto market conditions, measured by comparing the current ETH price (P_t) to its 50-day exponential moving average ($EMA50_t$). The ratio ($P_t/EMA50_t - 1$) indicates whether ETH is trading above or below its recent trend. These components combine into a latent payment score:

$$s_{\{i,t\}}^C = \alpha_4 \cdot \kappa_i + \gamma \cdot (P_t/EMA50_t - 1)$$

where α_4 controls the influence of personal crypto openness and γ controls the sensitivity to the price-EMA gap. The score is converted into a crypto preference probability using the logistic function:

$$P_{\{i,t\}}(C | m) = 1 / [1 + \exp(-s_{\{i,t\}}^C)]$$

This probability represents the rider's preference for crypto payment based on openness and market conditions. Whether crypto is actually used depends additionally on affordability: if the rider's crypto budget cannot cover the fare, the trip defaults to fiat regardless of preference, and the budget decreases with each crypto payment.

Crypto Budgets (Daily and Monthly)

Each rider maintains two separate cryptocurrency budgets: a daily budget for routine transit expenses and a monthly pool for larger expenditures. The daily crypto budget $B_{\{i,t\}}^d$ refreshes each morning and is primarily consumed by smaller transactions such as MTA fares. This budget adjusts dynamically based on market performance:

$$B_{\{i,t\}}^d = B_i^{\{d,0\}} \cdot \left(\frac{P_t}{EMA50_t} \right)^\eta$$

Here, $B_i^{\{d,0\}}$ is the rider's baseline daily budget, drawn from a lognormal distribution centered at \$12/day. The parameter $\eta = 2$ represents price elasticity, making the budget sensitive to market movements. The multiplier is bounded between 0.25 and 2.50 to prevent collapse during market crashes or unrealistic growth during price surges.

The behavioral logic follows the rider's perception of value. When P_t exceeds $EMA50_t$, riders view their crypto holdings as overvalued and increase spending willingness (multiplier > 1). Conversely, when the price falls below the average, riders prefer to hold their assets, reducing the spending budget (multiplier < 1).

The monthly crypto budget $B_{\{i,t\}}^m$ finances larger, occasional purchases such as Uber rides. Each rider begins with a monthly base $B_i^{\{m,0\}}$ drawn from a lognormal distribution centered at \$80. Unlike the daily budget, this pool does not refresh daily; it depletes with each crypto Uber fare and resets only at the start of the next month.

Affordability depends on these budgets. MTA fares are processed against the daily pool (requiring $B_{\{i,t\}}^d \geq C_{\{i,MTA,t\}}$). Uber fares, being larger, may draw from either the daily allowance or the remaining monthly pool ($B_{\{i,t\}}^d \geq C_{\{i,Uber,t\}}$ or $B_{\{i,t\}}^m \geq C_{\{i,Uber,t\}}$). If neither pool covers the fare, the affordability check fails and payment defaults to fiat.

Each rider also maintains a daily cash budget $B_{\{i,t\}}^{\$}$, drawn from a lognormal distribution centered at \$300/day. This is substantially higher than the crypto budget and rarely binds, but ensures that cash is also treated as a finite resource.

Crypto Price and EMA50

The model uses historical ETH/USD prices. For each simulation day, the 50-day exponential moving average is calculated from preceding price data and serves as the rider's reference price. The EMA assigns greater weight to recent prices than a simple moving average, consistent with standard financial trading

indicators. A 50-day warm-up window precedes the simulation horizon to ensure $EMA50_t$ is well-formed at simulation start.

Provider Crypto Caps

Each transportation provider sets a daily limit on cryptocurrency acceptance. For MTA, this limit is straightforward: MTA operates as a single system, so every crypto-paid trip counts toward one daily cap. Once the cap is reached, no further crypto payments are accepted for MTA that day.

Uber presents greater complexity because each fare is split between the platform (40%) and the driver (60%). A crypto payment for Uber requires that both the platform cap and the driver-side capacity remain available. This structural difference means an Uber crypto transaction can be denied even when the rider prefers crypto, if either the platform cap or driver capacity is exhausted.

The driver-side constraint is constructed from individual tolerances. Each driver d has a daily crypto tolerance $\tau_d \geq 0$ in dollars, measured against the driver's fare share. Forty percent of drivers refuse crypto entirely ($\tau_d = 0$); the remaining sixty percent draw tolerances uniformly between \$10 and \$100. Tolerances replenish daily.

When a cap binds, crypto requests are rationed on a first-come-first-served basis: riders are processed in random arrival order each day, and once remaining capacity is exhausted, subsequent crypto requests are denied and the rider falls back to fiat.

Simulation Process

The simulation runs over multiple days following a sequential procedure. Each day, the model first reads the ETH price and updates the 50-day EMA. This price information updates each rider's daily crypto budget—spending willingness varies day-to-day based on market conditions.

Next, the model generates daily travel conditions: each rider is assigned travel times and costs for both modes. MTA fare remains fixed, while Uber cost varies due to day-to-day fluctuations and surge pricing. These values determine mode utilities and the resulting choice probabilities.

After computing the mode choice, the model determines payment preference based on the rider's crypto openness, remaining budget, and fare affordability. Riders are then processed in random order to simulate stochastic arrival throughout the day. Fiat trips are completed normally. Crypto trips for MTA are checked against the MTA cap; for Uber, both the platform cap and driver-side capacity must have room. Denied crypto requests revert to fiat.

At the end of each simulated day, the model records: the number of MTA and Uber trips, crypto and fiat payment counts, denied crypto requests, and cap utilization for each provider. Each accepted transaction is simultaneously submitted to the deployed smart contract on Ganache, producing on-chain records that mirror the simulation's internal state. This dual-layer architecture — behavioral model on top, blockchain ledger below — ensures that all payment outcomes are independently verifiable through the blockchain's transaction history.

Scenario Design

The simulation compares multiple crypto acceptance scenarios by varying provider cap levels. For MTA, the daily crypto cap is set to low, medium, and high levels to observe system response under different acceptance thresholds. For Uber, scenarios are more granular because both the platform cap and driver-side capacity can be varied independently. The model can also adjust driver willingness distributions—for instance, scenarios where most drivers refuse crypto versus scenarios with widespread driver acceptance at higher tolerance levels.

For each scenario, the model reports: mode shares, crypto and fiat payment volumes, denied crypto requests, total crypto received by each provider, denial rates per mode, and the gap between crypto demand and crypto fulfillment. This gap quantifies how much potential crypto payment was blocked by provider caps.

RESULTS

This section reports simulation outcomes at the baseline scenario and across the range of crypto cap settings. The baseline uses a daily MTA crypto cap of \$2,000 and a daily Uber platform crypto cap of \$800, with 2,000 riders simulated over 30 days for a total of 60,000 trips. All results are direct model outputs rather than input assumptions.

Baseline Outcomes

The model produces a mode split of approximately 92% MTA and 8% Uber. This split emerges from the utility comparison rather than being imposed: MTA is substantially cheaper than Uber, and although Uber offers shorter travel times, the cost differential dominates for most riders. Within each mode, riders partition into three groups: those paying in crypto, those choosing fiat voluntarily, and those forced into fiat because the crypto cap was already exhausted (denied).

Table 2 Baseline Simulation Outcomes (MTA cap \$2,000, Uber cap \$800)

<i>Outcome</i>	<i>Value</i>
<i>Total trips</i>	60,000
<i>MTA mode share</i>	91.7%
<i>Uber mode share</i>	8.3%
<i>MTA crypto trips</i>	19,740
<i>MTA fiat trips (by choice)</i>	22,330
<i>MTA denied (forced to fiat)</i>	12,951
<i>Uber crypto trips</i>	1,636
<i>Uber fiat trips (by choice)</i>	2,436
<i>Uber denied (forced to fiat)</i>	907
<i>MTA denial rate</i>	39.6%
<i>Uber denial rate</i>	35.7%

The denial rate represents the share of riders who preferred crypto payment but were blocked because the provider's cap was already exhausted. At baseline, approximately 40% of crypto-preferring MTA riders and 36% of crypto-preferring Uber riders are denied and must pay in fiat. These denied trips constitute the central quantity of interest: they measure unmet demand for crypto payment created solely by provider caps.

Market Context

The ETH price over the simulation window is displayed alongside its 50-day exponential moving average. During roughly the first half of the window, the price sits below the moving average; it then crosses above and remains there. This crossing is consequential because each rider's daily crypto budget scales with the price-to-EMA ratio. When the price exceeds the average, riders perceive their crypto as overvalued and become more willing to spend, causing crypto demand to rise sharply after the crossing.

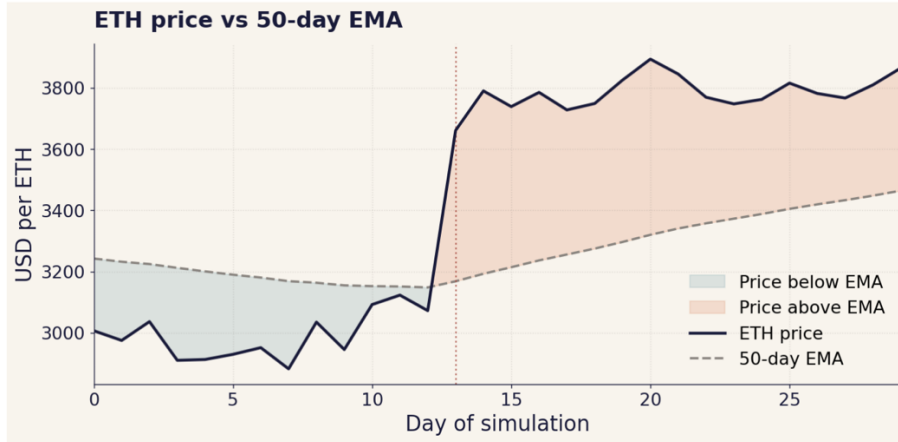


Figure 7 ETH price relative to its 50-day EMA over the simulation window (May–June 2024)

The daily payment outcomes reveal how many riders paid in crypto, how many were forced to switch to fiat, and how many chose fiat voluntarily. Before the price crosses the moving average, crypto demand remains sufficiently low that the cap rarely binds and almost no riders are denied. After the crossing, crypto demand surges, the cap becomes binding nearly every day, and a substantial block of riders (the rationed band) is pushed into fiat. This constitutes the core dynamic of the model: the cap has minimal effect when crypto demand is low but becomes a hard constraint precisely when demand is high.

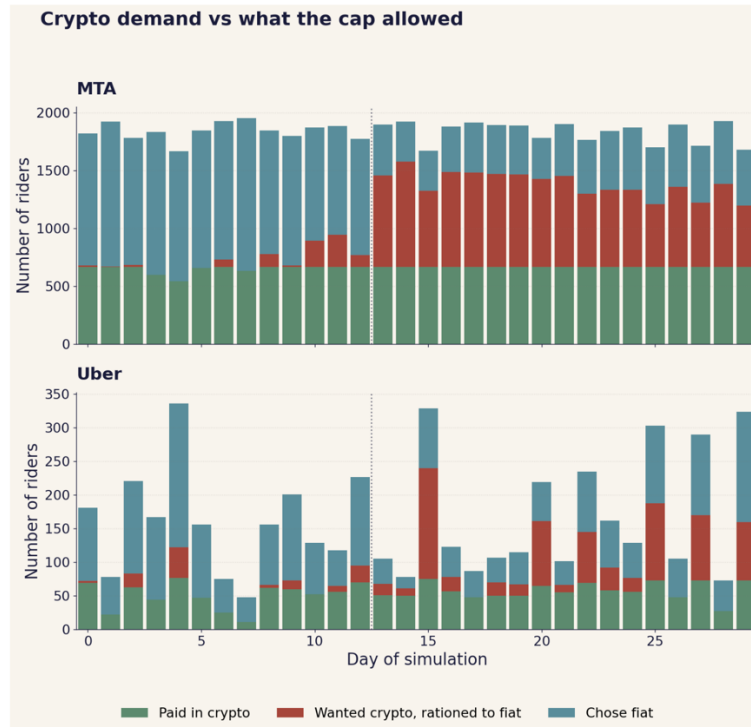


Figure 8 Daily payment outcomes for MTA (top) and Uber (bottom)

Sweeping the MTA cap from \$250 to \$6,000 drives the MTA denial rate from approximately 85% down to 0%, and sweeping the Uber cap from \$150 to \$4,000 produces an analogous reduction for Uber. Across all combinations tested, the MTA denial rate responds only to the MTA cap and the Uber denial rate responds only to the Uber cap; the two are independent. The overall mode split remains at roughly 92%/8% in every

scenario because, in this specification, the cap affects whether a rider can pay in crypto but does not enter the mode utility function—it does not shift riders between MTA and Uber. The welfare consequence of the cap therefore operates entirely through payment rationing rather than mode shifting.

DISCUSSION

The results of this study carry meaningful implications for how we think about blockchain-based payment systems in urban transit. Perhaps the most important finding is that provider-side cryptocurrency acceptance caps behave as a pure rationing mechanism rather than a demand-shaping tool. In classical transportation economics, supply-side constraints — road capacity limits, parking scarcity — reliably push travelers toward alternative modes. Payment-method constraints, however, operate on a fundamentally different axis. They restrict how a rider pays without changing the attractiveness of the ride itself. A rider denied crypto on the subway does not suddenly find Uber more appealing; she simply pays cash and boards the same train. This means transit agencies cannot use crypto caps as an indirect lever for mode management — their effect is confined entirely to redistributing payment-method welfare among riders who already chose that mode.

The behavioral dimension adds a layer of complexity that traditional fare systems never encounter. Crypto demand turns out to be highly elastic with respect to market conditions: when the Ethereum spot price trades above its 50-day moving average, riders perceive unrealized gains in their holdings and become markedly more willing to spend them on transit fares. This "paper-gains spending" pattern, well documented in behavioral finance, introduces payment demand volatility that has nothing to do with ridership itself. A transit agency could see crypto payment requests double in a single week not because more people are riding, but because the crypto market rallied. Fixed daily caps become binding precisely during these surges, creating substantial pockets of unmet demand that agencies must anticipate even though they cannot predict them from ridership trends alone.

On the policy side, the independence of MTA and Uber denial rates offers a simplifying insight: each provider can calibrate its own crypto acceptance threshold without worrying about spillover to the other. Raising or lowering MTA's cap has no measurable effect on Uber's denial rate or vice versa. This decouples the regulatory design problem — a transit authority does not need to coordinate its cap-setting with private ride-hailing operators. However, this independence also means that no single provider can unilaterally unlock the full potential of decentralized payments across the urban mobility ecosystem. System-wide improvement requires simultaneous action by all participants, which raises governance questions about who sets targets and how providers are incentivized to expand crypto access together.

It is important to acknowledge that the financial analysis presented here represents an early-stage modeling effort. The mode-choice parameters, the crypto-openness distribution, and the budget assumptions are grounded in reasonable approximations drawn from transportation literature and cryptocurrency market behavior, but they have not yet been validated against real-world fare and ridership data from either MTA or Uber. The provider caps used in the simulations — 2,000 per day for MTA, 800 per day for Uber — are hypothetical scenarios meant to illuminate system dynamics rather than prescriptive recommendations calibrated to actual institutional risk tolerances. The results should therefore be read as a proof of concept that establishes the structural relationships and identifies the key behavioral mechanisms rather than as definitive operational guidance.

Several directions naturally follow from these findings. The current model treats mode choice and payment choice as independent: a rider who keeps getting denied crypto on MTA does not eventually switch to Uber or abandon crypto altogether. In reality, people learn from repeated frustration. A natural next step is to introduce a feedback mechanism where past denial experiences erode the perceived utility of that mode-payment combination over time, allowing the model to capture mode shifting driven by payment constraints — something the current specification explicitly rules out. Whether this feedback produces meaningful

reallocation or merely small perturbations is an empirical question worth resolving, because it determines whether caps have any indirect mode-shifting power once learning is accounted for.

The volatility problem — where crypto demand spikes whenever ETH trades above its moving average — could potentially be mitigated by offering stablecoin alternatives like USDC alongside volatile assets. If riders can pay with a price-stable cryptocurrency, the EMA-driven demand surges that cause caps to bind might soften considerably. But it is also possible that stablecoins simply cannibalize fiat payments rather than alleviating pressure on the cap, since riders with high crypto openness may already be spending through volatile assets. Testing this empirically will clarify whether stablecoins are part of the solution or merely rearrange existing payment flows without reducing denial rates.

Finally, the equity dimension demands serious attention. Not everyone has access to cryptocurrency, and the populations most dependent on public transit — lower-income riders, older adults, those without smartphones or bank accounts — are the least likely to hold crypto. If blockchain payment systems eventually offer benefits such as loyalty rewards, dynamic discounting, or priority processing, there is a real risk of creating a two-tier transit experience where crypto-enabled riders accumulate advantages that cash-dependent riders cannot access. Future iterations of this work will model heterogeneous populations stratified by income and technology access to examine whether and how blockchain payment integration might widen or narrow existing transportation equity gaps. Addressing these questions is essential before any transit agency moves from proof-of-concept to deployment.

CONCLUSION

This study presents a utility-based discrete choice framework for analyzing cryptocurrency payment integration in urban transportation systems. By coupling a logit mode-choice model with a market-responsive crypto payment preference model, the research provides the first systematic analysis of how provider-side acceptance constraints interact with rider-side demand dynamics in a multi-modal transit network. The framework jointly captures mode selection driven by travel time, cost, and reliability alongside payment-method selection governed by individual crypto openness, real-time Ethereum market conditions, and layered budget constraints—offering a holistic representation of the decision architecture facing urban travelers in a blockchain-enabled payment ecosystem.

Three principal findings emerge from the analysis. First, rider demand for cryptocurrency payments is highly elastic with respect to market conditions: when the Ethereum spot price exceeds its 50-day exponential moving average, riders perceive unrealized gains and significantly increase their propensity to spend crypto on transit fares—a behavioral pattern consistent with the disposition effect documented in behavioral finance. This "paper-gains spending" mechanism introduces demand volatility into the transit payment system that is entirely exogenous to ridership patterns, representing a novel challenge for revenue planning that traditional fare systems do not face. Second, provider acceptance caps function as a pure rationing mechanism: they determine which riders can pay in crypto based on arrival order but do not influence the underlying mode choice between public transit and ride-hailing. The mode split remains stable regardless of cap levels, confirming that payment constraints operate on a fundamentally separate dimension from transportation utility—a theoretical distinction that differentiates payment caps from classical supply-side constraints known to induce mode shifting. Third, provider denial rates are mutually independent, responding only to their respective caps, which simplifies the policy design problem but necessitates coordinated action across all operators for system-wide improvements in crypto accessibility.

The practical implications of these findings are significant for transit agencies and urban planners contemplating blockchain integration. Agencies must anticipate that crypto payment demand will exhibit market-driven volatility uncorrelated with traditional ridership patterns, and that fixed daily caps will bind precisely during periods of high demand—creating substantial unmet demand among crypto-preferring

riders who must revert to fiat. The rationing interpretation implies that crypto caps cannot serve as indirect instruments for demand management; instead, their welfare effects operate entirely through payment-method allocation. Adaptive cap mechanisms that respond dynamically to market conditions, stablecoin integration that eliminates price-EMA-driven demand spikes, and Layer-2 scaling solutions that reduce transaction costs for low-fare rides represent promising pathways for aligning provider infrastructure with the inherently volatile demand for cryptocurrency payments.

Ultimately, while blockchain technology offers a transparent, auditable, and decentralized framework for multi-provider transit payment systems—as demonstrated by the Ethereum-based simulation platform developed in this research—its practical success depends on careful mechanism design that balances provider risk exposure against rider welfare, addresses equity concerns for populations without crypto access, and accommodates the unique demand dynamics introduced by cryptocurrency market fluctuations. This study provides the analytical foundation for these design decisions and establishes a research agenda at the intersection of transportation demand modeling, decentralized finance, and smart city infrastructure. Future work will pursue endogenous mode-payment feedback loops, formal welfare quantification of payment rationing, agent-based adoption dynamics with peer effects, and consortium blockchain architectures for pooled multi-provider risk management.

REFERENCES

- [1] S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” 2008. <https://bitcoin.org/bitcoin.pdf>.
- [2] A. Hayes, “Blockchain Facts: What It Is, How It Works, and How It Can Be Used,” Investopedia, Sep. 16, 2024.
- [3] R. Yousuf, Z. Jeelani, D. A. Khan, O. Bhat, and T. A. Teli, “Consensus Algorithms in Blockchain-Based Cryptocurrencies,” ResearchGate, 2021.
- [4] WisdomTree, “Bitcoin vs. Traditional Currencies,” WisdomTree Investments.
- [5] International Association of Public Transport (UITP), “Distributed Ledger Technology in Public Transport: Use Cases for Blockchain,” Brussels, Belgium, Mar. 2022.
- [6] T. H. Nguyen, J. Partala, and S. Pirttikangas, “Blockchain-based Mobility-as-a-Service,” ResearchGate, May 2019.
- [7] P. Helo and A. H. M. Shamsuzzoha, “Real-time supply chain—A blockchain architecture for project deliveries,” ScienceDirect, 2019.
- [8] F. Longo, L. Nicoletti, A. Padovano, G. d’Atri, and M. Forte, “Blockchain-enabled supply chain: An experimental study,” ScienceDirect, 2019.
- [9] T. S. Ruprah, A. Patil, Y. Nanavare, B. Khade, and A. Solapure, “Peer to Peer Ridesharing using Blockchain Technology,” ResearchGate, 2024.
- [10] B. Dickson, “Blockchain has the potential to revolutionize the supply chain,” TechCrunch, Nov. 24, 2016.
- [11] S. Chauhan, R. Singh, and R. Bansal, “Role of Blockchain Technology in Building Transparent Supply Chain Management,” Jun. 2022.
- [12] NIST and M. Dworkin, “SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions,” Federal Inf. Process. Stds. (NIST FIPS), 2015.
- [13] H. Min, “Blockchain technology for enhancing supply chain resilience,” ScienceDirect, Dec. 2018.
- [14] “How does blockchain work?” Stanford Online. <https://online.stanford.edu/how-does-blockchain-work>.
- [15] The Investopedia Team, “What Is a Block in the Crypto Blockchain, and How Does It Work?” Investopedia, May 24, 2024.
- [16] G. Walker, “Learn Me A Bitcoin,” Learn Me A Bitcoin, Dec. 2024. <https://learnmeabitcoin.com/>.
- [17] T. H. Tran, P. H. Luan, and Y. Nakashima, “A High-Performance Multimem SHA-256 Accelerator for Society 5.0,” IEEE Access, Mar. 2021.
- [18] V. Buterin, “Ethereum: A Next-Generation Smart Contract and Decentralized Application Platform,” 2014.
- [19] P. Wackerow, “Intro to Ethereum,” Ethereum.org, Jun. 26, 2024.